

## TOWARDS THE FUTURE INTERNET

The image on the cover is a partial map of the Internet based on the OPTE project started by Barrett Lyon ([www.blyon.com](http://www.blyon.com)) who kindly let us use it for the front cover of this book. In this graph the lines connect nodes representing IP addresses of some indicative Class C networks color-coded according to their corresponding allocation. For more information see <http://www.opte.org/maps/>.

# Towards the Future Internet

A European Research Perspective

Edited by

Georgios Tselentis

John Domingue

Alex Galis

Anastasius Gavras

David Hausheer

Srdjan Krco

Volkmar Lotz

and

Theodore Zahariadis

**IOS**  
*Press*

Amsterdam • Berlin • Tokyo • Washington, DC

© 2009 The authors and IOS Press.

All rights reserved. No part of this book may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, without prior written permission from the publisher.

ISBN 978-1-60750-007-0

Library of Congress Control Number: 2009925664

*Publisher*

IOS Press BV  
Nieuwe Hemweg 6B  
1013 BG Amsterdam  
Netherlands  
fax: +31 20 687 0019  
e-mail: [order@iospress.nl](mailto:order@iospress.nl)

*Distributor in the UK and Ireland*

Gazelle Books Services Ltd.  
White Cross Mills  
Hightown  
Lancaster LA1 4XS  
United Kingdom  
fax: +44 1524 63232  
e-mail: [sales@gazellebooks.co.uk](mailto:sales@gazellebooks.co.uk)

*Distributor in the USA and Canada*

IOS Press, Inc.  
4502 Rachael Manor Drive  
Fairfax, VA 22032  
USA  
fax: +1 703 323 3668  
e-mail: [iosbooks@iospress.com](mailto:iosbooks@iospress.com)

LEGAL NOTICE

The publisher is not responsible for the use which might be made of the following information.

PRINTED IN THE NETHERLANDS

## Preface

The Internet is a remarkable catalyst for creativity, collaboration and innovation providing us today with amazing possibilities that just two decades ago it would have been impossible to imagine; and yet we are not amazed! It is only 20 years ago that Tim Berners-Lee invented the Web and two years later, CERN publicized the new World Wide Web project. If one could take a trip back in time with a time machine and say to people that today even a child can access for free a satellite image of any place on earth, interact with other people from anywhere and query trillions of data all over the globe with a simple click on his/her computer they would have said that this is science fiction!

Our challenge today is to prepare a similar trip into the future: what will be the Internet in ten-twenty years from now and what more amazing things will it offer to people? But before trying to see how the future will look like, we need to consider some important challenges that the Internet faces today.

If we consider Internet like one big machine, we should note that it has been working all these years without witnessing a major overall failure, showing a remarkable resilience for a human-made technology. However, Internet provides its services on the basis of “best effort” (i.e. there is no guarantee of delivering those services) and “over provisioning” (i.e. to be sure that we get a certain quality of services we need to keep available all time an important amount of resources). Internet was never designed to serve massive scale applications with guaranteed quality of service and security. Emerging technologies like streaming high quality video and running 3D applications face severe constraints to run seamlessly anytime, everywhere, with good quality of services. Thus, if we want to continue the growth, improve the quality and provide the affordable basic access, new business models have to be put in place to make Internet sustainable.

European scientists proved that they are at the forefront of internet research already since the invention of the web. But the challenges are huge and complex and cannot be dealt in isolation. The European Future Internet Assembly is the vehicle to a fruitful scientific dialogue bringing together the different scientific disciplines that contribute to the Future Internet development with scientists from more than 90 research projects funded until today with about 300 million euros under the 7th Framework Programme. Another 400 million euros will be made available in the near future. These amounts coupled with private investments bring the total investment to more than a billion euros. This is an important investment showing Europe’s commitment to address the challenges of the future Internet.

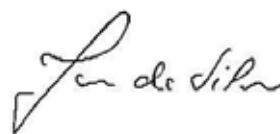
This book is a peer-reviewed collection of scientific papers addressing some of the challenges ahead that will shape the Internet of the Future. The selected papers are

representative of the research carried out by EU-funded projects in the field. European scientists are working hard to make the journey to the Future Internet as exciting and as fruitful as was the trip that brought us the amazing achievements of today. We invite you to read their visions and join them in their effort so Europe can fully benefit from the exciting opportunities in front of us.

Mário Campolargo  
Director F – Emerging Technologies and  
Infrastructures



João Da Silva  
Director D – Converged Networks and  
Services



## **Editorial Board**

Coordination: Georgios TSELENTIS, European Commission

John DOMINGUE – Knowledge Media Institute, The Open University & STI International

Alex GALIS – University College London

Anastasius GAVRAS – Eurescom

David HAUSHEER – University of Zurich

Srdjan KRKO – Ericsson

Volkmar LOTZ – SAP Research

Theodore ZAHARIADIS – Synelixis/TEI of Chalkida

## Reviewers

Nancy ALONISTIOTI	University of Piraeus
Federico ALVAREZ	Universidad Politecnica de Madrid
Pascal BISSON	Thales Group
Mike BONIFACE	University of Southampton – IT Innovation Centre
Jan CAMENISH	IBM Research Zurich Research Laboratory
Guillermo CISNEROS	ETSIT – UPM
Petros DARAS	Informatics and Telematics Institute
Stefano DE PANFILIS	Ingegneria Informatica SpA
Panagiotis DEMESTICHAS	University of Piraeus
Jordi DOMINGO-PASCUAL	Universitat Politecnica de Catalunya
John DOMINGUE	Knowledge Media Institute the Open University
Schahram DUSTDAR	Vienna University of Technology
Dieter FENSEL	University of Innsbruck
Mike FISHER	British Telecom
Vincenzo FOGLIATI	Telespazio
Jan FURMAN	Cesnet
Alex GALIS	University College London
Anastacius GAVRAS	Eurescom
Steve HAILES	University College London
Stephan HALLER	SAP
David HAUSHEER	University of Zurich
Juanjo HIERRO	Telefonica Investigacion y Desarrollo
Erik HUIZER	ICS-Department of Information & Computing Sciences
Valerie ISSARNY	INRIA
Ebroul IZQUIERDO	Queen Mary University of London
Adam KAPOVITZ	Eurescom
David KENNEDY	Eurescom
Peter KIRSTEIN	University College London
Srdjan KRKO	Ericsson
Javier LOPEZ-MUNOZ	University of Malaga
Thomas MAGEDANZ	Fraunhofer Fokus
Petri MAHONEN	RWTH Aachen University
Daniele MIORANDI	Create-Net
Pekka NIKANDER	Helsinki Institute for Information Technology
Andrew OLIPHANT	British Telecom
Dimitri PAPADIMITRIOU	Alcatel-Lucent Bell
Bernhard PLATTNER	ETH Zurich
Klaus POHL	University of Duisburg-Essen
Aiko PRAS	University of Twente
Christian PREHOFER	Nokia Research
Mirko PRESSER	University of Surrey
John STRASSNER	Waterford Institute of Technology
Panagiotis TRAKADAS	ADAE
Phuoc TRAN-GHIA	University of Wuerzburg
Paolo TRAVERSO	IRST
Theodore ZAHARIADIS	Synelixis Ltd



## Introduction

### 1. CURRENT INTERNET

The current Internet designed 40 years ago is today the most important information, service and networking infrastructure providing the mechanisms for the digital society at large to function as an integrated entity. This infrastructure is evolving rapidly with the transition from “sharing” in Web 1.0 (Web) to “contributing” in Web 2.0 (user generated content) to “co-creating” in Web 3.0 (collaborative production, semantic Web).

The current Internet has been founded on a basic architectural premise, that is: a simple network service can be used as a universal means to interconnect intelligent end systems. The current Internet is centred on the network layer being capable of dynamically selecting a path from the originating source of a packet to its ultimate destination, with no guarantees of packet delivery or traffic characteristics. The end-to-end argument has served to maintain the desire for this simplicity. The continuation of simplicity in the network has pushed complexity into the endpoints, which has allowed the Internet to reach an impressive scale in terms of inter-connected devices. However, while the scale has not yet reached its limits, the growth of functionality and the growth of size have both slowed down. It is now a common belief that the current Internet would reach soon both its architectural capability limits and its capacity limits (i.e. in addressing, in reachability, for new demands on Quality of Service, Service and Application provisioning, etc).

Although the current Internet, as a ubiquitous and universal means for communication and computation, has been extraordinarily successful, there are still many unsolved problems and challenges several of which have basic aspects. Many of these aspects could not have been foreseen 40 years ago when the first parts of the Internet were built, but these need to be addressed now. The very success of the Internet is now creating obstacles to future innovation of both the networking technology that lies at the Internet’s core and the services that use it. In addition, the ossification of the Internet makes the introduction and deployment of new network technologies and services both difficult and costly.

We are faced with an Internet that is good at delivering packets, but shows a level of inflexibility at network layer and a lack of built-in facilities to support any non-basic functionality. The aspects, which we consider to be missing, are:

- Inherent network management functionality, specifically self-management functionality.
- Facilities for the addition of new functionality, including capability for activating a new service on-demand, network functionality, or protocol (i.e. addressing the ossification bottleneck).

- Facilities for the large scale provisioning, management and deployment of services; support for a high-level integration between services and networks.
- Facilities for orchestration of security, reliability, robustness, mobility, context, service support, and management for both the communication resources and the services' resources.
- Mobility of networks, services, and devices.
- Facilities to support Quality of Service (QoS) and Service Level Agreements (SLAs).
- Trust Management and Security; Privacy and data-protection mechanisms of distributed data.
- An adequate addressing scheme, where identity and location are not embedded in the same address.
- Facilities to interact with the physical world and seamlessly use the physical context information to enhance and improve existing services and to create the new ones.
- Socio-economic aspects including the need for appropriate incentives, viable business models, legal and regulative issues, and the need for security and privacy.
- Energy awareness.

Those challenges and capabilities envisaged for the Future Internet are addressed by several research areas as depicted in Figure 1 and are presented in the following sections.




Figure 1 – Future Internet Research Areas

This book contains 32 selected papers presenting a variety of European research results aimed at progressing the current Internet. It offers, above all, a vision of the future rather than an account of deployed solutions. It presents representative research results in seven interrelated area of research for Future Internet: 1.1 Socio-economics; 1.2 Trust and Identity; 1.3 Experimental Research; 1.4 Management and Service-aware networking Architectures; 1.5. Service Offers; 1.6. Content Networks; 1.7 Real World Internet.

## **1.1 Future Internet Socio-economics Research Challenges**

Far reaching technological innovations like those envisioned by the Future Internet Assembly can only be successfully deployed, if their non-technical issues and business potential are taken into account. Therefore, the consideration of socio-economic aspects in the Future Internet is of key importance. Socio-economics is a multi-disciplinary field which cuts across all domains of the Future Internet, including networks, services, and content. The research challenges faced in this context are manifold. Suitable pricing and business models need to be designed which will provide appropriate incentives for network operators, service providers, and end-users. Moreover, legal and regulative issues such as network neutrality, privacy, and digital rights have to be addressed. Finally, important challenges include the increasing degree of mobility in social life and the need for security.

To this end, this book includes a number of contributions in the area of Future Internet Socio-economics. Paper **#1** provides an overview on the socio-economic challenges and perspectives, which have been identified and discussed by the Future Internet Socio-Economics (FISE) working group. Moreover, Paper **#30** presents and discusses a number of socio-economic scenarios from which it derives design recommendations for Real World Internet applications and beyond. Paper **#2** analyzes further challenges of the Internet evolution in regard to technology development and information society attitude by evaluating the dynamics of change against the outstanding aspects of the current situation. Finally, two specific architectural approaches are presented. Paper **#3** proposes an economic traffic management architecture to provide appropriate incentives for users, ISPs, and overlay providers in the Future Internet, while Paper **#8** discusses the Trilogy architecture which aims to jointly integrate both the technical and socio-economic aspects into a single solution.

## **1.2 Future Internet Trust and Identity Research Challenges**

The demand for security, trust establishment and privacy protection for the Internet has increased the same way as its usage scenarios have changed. While its protocols have been originally designed for specific-purpose communication between benign and honest peers, the Internet developed and continues to develop towards a platform used by everyone for a multitude of private and business purposes, with its openness and the criticality and value of the transactions conducted over it creating an incentive for

malicious entities to tamper with it. While lots of security technologies for the Internet are available, the Future Internet imposes additional challenges on security and trust, mainly due to its pervasiveness, its scale, and the lifelong involvement of its users. The challenges include, for instance, the design of identity management systems capable of dealing with billions of entities and their different roles in the Future Internet, the trustworthiness and control of distributed applications based on services offered through open service delivery platforms, and the secure and trusted interaction with real-world objects and entities through sensor and actuator network infrastructures.

The three papers selected for the Security, Trust and Identity section of this book exemplify these challenges, and the proposed solutions provide significant contributions towards a secure and trusted Future Internet. Paper #4 investigates into the trustworthiness of distributed applications that involve potentially malicious third-party services or client software. The architecture is based on the notion of “security by contract”, where a contract describes security relevant behaviour of an application that can be evaluated against the security requirements of the platform the application is running on. The architecture supports the integration of several policy enforcement techniques to ensure that the constraints imposed by the contract are met.

Wireless sensor networks allow connecting software applications and physical objects. While this opens a full range of exciting new usage scenarios within the Future Internet, it has to be considered that because of their openness, exposure and limited resources sensor networks are susceptible to attacks. Paper #5 proposes a protocol stack that integrates trust and reputation mechanisms to enable trusted routing for sensor networks and that supports secure service discovery. The integration of an intrusion detection system and self-configuration capabilities cope with the fact that attackers still might be able to surpass the first line of defence.

The opportunities the Future Internet opens for individuals and businesses will only be taken if users keep control over their personal and private data. Identity management and privacy protection is not only a technical issue, but has to take legal constraints into account. Paper #6 shows how European legislation and users’ privacy needs impacts the design of Identity Management Systems, and discusses the concept of Virtual Identity as an enabler of solutions achieving transparency and control.

### **1.3 Future Internet Research and Experimentation Challenges**

The Internet itself has been the largest-scale laboratory for emerging applications and services. However, it is evident that it cannot be considered as a testbed for the basic protocols, architectures and services. The shape of a Future Internet Research and Experimentation (FIRE) facility is emerging that will provide an open playground for the kind of research that cannot be conducted on a critical infrastructure like today’s Internet. This facility is being gradually built, based on a federation of existing relevant testbeds, and will grow according to the research needs specified by the related

“customer” research projects. The technical infrastructure supporting the federation of existing testbeds is described in Paper #7.

At the same time there is an increasing demand from academia and industry to bridge the gap between long-term research and large-scale experimentation, which can be done through experimentally driven research consisting of iterative cycles of research, design and experimentation of new networking and service architectures and paradigms for the Future Internet addressing all levels. Experimentally-driven research is suggested to address broad system-level research which views the Future Internet as a complex system and which proves and exploits the full value of the facility doing truly multidisciplinary experimental research, testing new internet architectures and paradigms and allowing for socio-economic impact assessment of future changes to the Internet.

The FIRE facility is planned to support research on the Future Internet and its services by including testbeds for different stages of the research and development cycle - from proof-of-concept type testbeds to pre-commercial testbeds and to support testing the impact of changes to the Internet not only in technical but also in socio-economic terms. Furthermore it should cover all levels from fast network connectivity to service architectures at different levels taking a holistic system view such as to support security architectures as described in Paper #4, or home area networks as described in the Paper #15. The facility should become a sustainable research infrastructure for the Future Internet serving both industry and academia in their Future Internet related research and to overcome limited availability of testbeds both in time and geographic reach.

#### **1.4 Future Internet Management and Service-aware Networking Architectures (MANA) Research Challenges**

MANA research covers the management, the networking, the service-aware networking including service platform technologies and systems, which form the infrastructure for Future Internets.

The current trend for networks is that they are becoming service-aware. Service awareness itself has many aspects, including the delivery of content and service logic, fulfilment of business and other service characteristics such as Quality of Service (QoS) and Service Level Agreements (SLA) and the optimisation of the network resources during the service delivery. Conversely, services themselves are becoming network-aware. Networking-awareness means that services are executed and managed within network execution environments and that both services and network resources can be managed uniformly in an integrated way. The design of both Networks and Services is moving forward to include higher levels of automation, autonomicity, including self-management.

The Future Internets would be built as service-aware and self-aware federated networks, which provide built-in and orchestrated aspects such as: context, reliability, robustness, mobility, security, service support, and self-management of the communication resources and the services. Such aspects suggest a transition from a service-agnostic Internet to a new service-aware and self-aware Internet, in which self-awareness is serving the purpose of communication and computation by means of enhanced in-network and in-service decisions

In order to achieve the objective of being service-aware and network-aware and to overcome the ossification of the current Internet, papers selected for this book envisage various novel solutions for the Future Internet including:

- Clean slate approaches to communication networks as presented in the Papers **#8, #9, #10**
- Self-managed networks as presented in the Papers **# 11, # 12, # 13, #14**
- Service Clouds approaches as presented in the Papers **#21, #22**
- QoS and SLA for harmonized management as presented in the Papers **#25, #27, #24, #32, #17**
- Context Networking approaches as presented in the Papers **#23, #31**
- Optical Networking approaches as presented in the Papers **#15, #16**

## **1.5 Future Internet Service Offer Research Challenges**

Although several perspectives can be taken on a new generation network there is near universal agreement that a service-centric view will be central. Within some quarters it is assumed that above a certain level of abstraction everything can be viewed as a service leading to the concept of the “Internet of Services”. Several papers examine this concept of from several angles. In Paper **#17** the authors argue that within the Future Internet we should consider ‘real services’ rather than software services and the outlines a conceptual architecture based on this premise. Taking a business perspective over services is also a central theme of Paper **#18** which describes a multi-level approach to SLA management covering business, software and infrastructure layers. In Paper **#17** authors additionally argue that the “Internet of Services” will be similar to the “Internet of Web pages” and we can learn from the success of the latter. This theme of using the principles of the Web is continued and expanded in Paper **#19**. Here the authors postulate a “Service Web” which combines the principles underlying SOA, semantics, the Web and Web 2.0 to create a Web where billions of services are managed and used in a seamless fashion. Web 2.0 continues as a theme in Paper **#20** where a platform for user-centric service creation and management is described.

Within infrastructures virtualisation, abstracting away from the details of the underlying hardware and software, is a key concept in providing services in a usable

scalable fashion. Paper **#21** outlines the concept of virtual execution environments as a way of partitioning virtualized computational resources and describes how these could support the Future Internet effort. Paper **#22** also outlines a virtualisation approach, which increases mobility and performance and compares this to Grid and Cloud based infrastructures.

A key component of any architecture is data storage and Paper **#14** outlines a scalable self-managing data store built on a structure overlay network utilising a distributed transaction protocol.

Context is an issue that occurs again and again in discussions on the Future Internet, normally encompassing location, user preferences and constraints and devices involved in an interaction. Paper **#23** presents a generic view to federating context related services and makes a case for a new European context-brokering industry.

How might the above pieces fit together? NEXOF-RA project, aims to develop a reference architecture for NESSI the European Technology Platform for software and services. Early on in the development of NEXOF-RA terminology arose as an important issue. The processes, results and lessons learned in developing the NEXOF glossary are outlined in Paper **#32**.

As we can see from this book, multi-disciplinarity is an important aspect for the whole of the Future Internet. Paper **#24** outlines multi-disciplinary research challenges for the “Internet of Services” and describes a research framework to address these.

## **1.6 Future Content Networks Research Challenges**

As the Internet becomes more capable of transporting high bandwidth communications’ signals, with low delay and enhanced Perceived Quality of Service (PQoS), demand for ever richer real-time communication modalities will arise. 2D and multi-viewpoint video and stereo audio will give way to 3D and stereoscopic video and audio-wave field synthesis. Based on the vision of the papers that follow, current networks and infrastructures for content and communication will be gradually replaced by the Future Content Internet, which would accommodate seamless end-to-end multi-media communications across a complex combination of network constituents, such as personal area networks, body area networks, home networks, fixed access networks, mobile access networks, metro networks and core networks. Several approaches are opened for consideration:

- Seamless end-to-end multi-media communications across complex combinations of networks as presented in the Paper **#9**.
- Peer-to-peer, hybrid or fully distributed networking paradigms for content handling as presented in the Paper **#25**

- Content codification and adaptation to facilitate content capabilities as presented in the Papers **#26**, **#27**.

A summary of the major challenges towards the Future 3D Media Internet can be found in the Paper **#28**.

## 1.7 Real World Internet Research Challenges

Integration of the physical world into the Future Internet is another important aspect addressed in the book. Sensors, actuators, Radio-frequency identification (RFID) enabled items, and generally heterogeneous network enabled machines and everyday items will be integrated into the fabric of the Future Internet, merging the digital and the physical worlds and enabling a range of new and enhanced Internet services and applications. Current sensor and actuator network deployments are still relatively sparse and built as independent, vertically closed solutions. To achieve real integration of the physical and the digital world it is necessary to interconnect these individual installations, to enable their seamless interaction with other Internet users and to provide unified access to the information and services they provide.

Design of scalable architecture that incorporates mechanisms to enable easy convergence and interoperability of heterogeneous sensor and actuator networks with the Internet; provision of access to context information and actuation services in a unified manner, while ensuring adequate security properties; efficient communication protocols able to efficiently cope with the new traffic and traffic patterns as well as mechanisms and protocols that deal with the consequences of mobility of edge networks are some of the main RWI challenges. In Paper **#29**, the properties and respective challenges of RWI are discussed, the design goals for a RWI architecture described and components for an initial RWI architecture and the interactions between those outlined and identified.

The socio-economic aspects of the RWI are addressed in Paper **#30**. Several scenarios are described in an effort to anticipate the evolution of the RWI with the goal to support design of a system that will be accepted by all stakeholders involved and that creates value for the users. Taking a multidisciplinary view, the potential benefits of the RWI services to individual users and the society as a whole and a business perspective of the dynamic deployment of RWI services are addressed.

The security and trust aspects of wireless sensor networks are addressed in Paper **#5**. The main challenges addressed by the paper are design of a modular, scalable, secure and trusted networking protocol stack able to offer self-configuration and secure roaming of data and services across insecure infrastructure of heterogeneous Wireless Sensor Networks (WSNs), with particular focus on trusted route selection, secure service discovery, and intrusion detection.



## 2. FUTURE INTERNET ASSEMBLY (FIA)

Future Internet research and development threads have recently been gaining momentum all over the world and as such the international race to create a new generation Internet is in full swing. On 31 March 2008, the Future Internet Assembly (FIA) was kicked off at the Bled conference (e.g. [www.future-internet.eu/events/future-internet-conference-bled-2008.html](http://www.future-internet.eu/events/future-internet-conference-bled-2008.html)) organised by the European Commission and the Slovenian European Union (EU) Presidency as the means to enable fundamental and systemic innovation in Europe in networking and services for the realization of the Future Internet within the timescale of 2020. FIA includes most of the EU FP7 research projects associated with Future Internet. The Assembly is structured to permit open interaction and cross-fertilization across technical domains and also has the following strategic goals:

- A joint strategic research agenda for the Future Internet encompassing common actions and requirements
- Fostering common research deliverables and results creating value for the EU research projects concerned
- Developing a consolidated calendar of events aiming at avoiding fragmentation of research efforts in Europe

## ACKNOWLEDGMENT

This book reports on advances in Internet and it is a joint effort of the people who are active in European Union funded research projects. The book editorial group wish to thank to the papers' authors, who worked hard and timely to produced and edit 32 selected papers from a group of 64 papers submitted to the call for papers. We also thank to the book reviewers who helped with the selection of and with the quality improvement of the papers by providing valuable recommendations.

Finally, we would like to thank to Joao da Silva and Mario Campolargo for their drive and wisdom in establishing FIA as a key player in the research market place for Future Internet, to Georgios Tselentis, Paulo de Sousa, Yves Paindaveine, Anne-Marie Sassen, Remy Bayou, Manuel Mateo, Isidro Laso for their support and encouragement for the work on the book. They actively supported the progress of the book and therefore favourably and constructively affected its content. Special thanks also go to Sandra Giovanelli who patiently assured the interface with authors and the collection of contributions during all the months this book was prepared.

28<sup>th</sup> February 2009  
Future Internet Assembly caretakers and book editorial group

Alex Galis - University College London, United Kingdom  
John Domingue – Knowledge Media Institute, The Open University and STI  
International, United Kingdom  
Theodore Zahariadis – Synelixis/TEI of Chalkida, Greece  
Anastasius Gavras – Eurescom, Germany  
David Hausheer – University of Zurich, Switzerland  
Volkmar Lotz – SAP Research, France  
Srdjan Krco – Ericsson, Ireland

## Contents

Preface	v
<i>Mário Campolargo and João Da Silva</i>	
Editorial Board	vii
Reviewers	viii
Introduction	ix
<i>Alex Galis, John Domingue, Theodore Zahariadis, Anastasius Gavras, David Hausheer, Volkmar Lotz and Srdjan Krco</i>	
1. Future Internet Socio-Economics – Challenges and Perspectives	1
<i>David Hausheer, Pekka Nikander, Vincenzo Fogliati, Klaus Wünnel, Maria Ángeles Callejo, Santiago Ristol Jorba, Spiros Spirou, Latif Ladid, Wolfgang Kleinwächter, Burkhard Stiller, Malte Behrmann, Mike Boniface, Costas Courcoubetis and Man-Sze Li</i>	
2. Challenges of Internet Evolution: Attitude and Technology	12
<i>Jon Mikel Rubina</i>	
3. An Economic Traffic Management Approach to Enable the TripleWin for Users, ISPs, and Overlay Providers	24
<i>Tobias Hoßfeld, David Hausheer, Fabio Hecht, Frank Lehrieder, Simon Oechsner, Ioanna Papafili, Peter Racz, Sergios Soursos, Dirk Staehle, George D. Stamoulis, Phuoc Tran-Gia and Burkhard Stiller</i>	
4. A Security Architecture for Web 2.0 Applications	35
<i>Lieven Desmet, Wouter Joosen, Fabio Massacci, Katsiaryna Naliuka, Pieter Philippaerts, Frank Piessens, Ida Siahaan and Dries Vanoverberghe</i>	
5. Securing Wireless Sensor Networks Towards a Trusted “Internet of Things”	47
<i>Theodore Zahariadis, Panagiotis Trakadas, Helen Leligou, Kostas Papadopoylos, Evangelos Ladis, Christos Tselikis, Charalampos Vangelatos, Lionel Besson, Jukka Manner, Michalis Loupis, Federico Alvarez and Yannis Papaefstathiou</i>	
6. Privacy-Enabled Identity Management in the Future Internet	57
<i>Christoph Sorge, Joao Girao and Amardeo Sarma</i>	
7. Control of Resources in Pan-European Testbed Federation	67
<i>Anastasius Gavras, Halid Hrasnica, Sebastian Wahle, David Lozano, Denis Mischler and Spyros Denazis</i>	
8. The Trilogy Architecture for the Future Internet	79
<i>Louise Burness, Philip Eardley and Robert Hancock</i>	
9. A Future Internet Embracing the Wireless World	91
<i>Henrik Abramowicz, Norbert Niebert, Stephan Baucke, Martin Johnsson, Börje Ohlman, Mario Kind, Klaus Wuenstel, Hagen Woesner and Jürgen Quittek</i>	

10. The Publish/Subscribe Internet Routing Paradigm (PSIRP): Designing the Future Internet Architecture	102
<i>Sasu Tarkoma, Mark Ain and Kari Visala</i>	
11. Management Architecture and Systems for Future Internet Networks	112
<i>A. Galis, S. Denazis, A. Bassi, P. Giacomini, A. Berl, A. Fischer, H. de Meer, J. Srassner, S. Davy, D. Macedo, G. Pujolle, J.R. Loyola, J. Serrat, L. Lefevre and A. Cheniour</i>	
12. Towards a Future Internet: Node Collaboration for Autonomic Communication	123
<i>Tanja Zseby, Thomas Hirsch, Michael Kleis and Radu Popescu-Zeletin</i>	
13. Creating a Viable Evolution Path Towards Self-Managing Future Internet via a Standardizable Reference Model for Autonomic Network Engineering	136
<i>Ranganai Chaparadza, Symeon Papavassiliou, Timotheos Kastrinogiannis, Martin Vigoureux, Emmanuel Dotaro, Alan Davy, Kevin Quinn, Michał Wódczak, Andras Toth, Athanassios Liakopoulos and Mick Wilson</i>	
14. A Scalable, Transactional Data Store for Future Internet Services	148
<i>Alexander Reinefeld, Florian Schintke, Thorsten Schütt and Seif Haridi</i>	
15. Future Internet in Home Area Networks: Towards Optical Solutions?	160
<i>Roberto Gaudino, Daniel Cardenas, Martial Bellec, Benoit Charbonnier, Noella Evanno, Philippe Guignard, Sylvain Meyer, Anna Pizzinat, Ingo Möllers and Dieter Jäger</i>	
16. DICONET: Future Generation Transparent Networking with Dynamic Impairment Awareness	173
<i>Ioannis Tomkos, Yvan Pointurier, Siamak Azodolmolky, Michael Eiselt, Thierry Zami, Radoslaw Piesiewicz, Chava Vijaya Saradhi, Matthias Gunkel, Uri Mahlab, Ming Chen, Yabin Ye, Mario Pickavet, Maurice Gagnaire, Emmanouel Varvarigos, Josep Solé Pareta, Reza Nejabati, Yixuan Qin and Dimitra Simeonidou</i>	
17. From Software Services to a Future Internet of Services	183
<i>Marco Pistore, Paolo Traverso, Massimo Paolucci and Matthias Wagner</i>	
18. Multi-Level SLAs for Harmonized Management in the Future Internet	193
<i>Wolfgang Theilmann and Luciano Baresi</i>	
19. The Service Web: A Web of Billions of Services	203
<i>John Domingue, Dieter Fensel, John Davies, Rafael González-Cabero and Carlos Pedrinaci</i>	
20. User-Centric Future Internet and Telecommunication Services	217
<i>Carlos Baladrón, Javier Aguiar, Belén Carro, Laurent-Walter Goix, Alberto León Martín, Paolo Falcarin and Jürgen Sienel</i>	
21. Design for Future Internet Service Infrastructures	227
<i>B. Rochwerger, A. Galis, D. Breitgand, E. Levy, J.A. Cáceres, I.M. Llorente, Y. Wolfsthal, M. Wusthoff, S. Clayman, C. Chapman, W. Emmerich, E. Elmroth and R.S. Montero</i>	

22. Above the Clouds: From Grids to Service-Oriented Operating Systems <i>Lutz Schubert, Alexander Kipp and Stefan Wesner</i>	238
23. Future Internet: Towards Context Information Brokering <i>M. Oskar van Deventer, Paul Tilanus, Mike Schenk, Eelco Cramer and Joost Adriaanse</i>	250
24. S-Cube: Addressing Multidisciplinary Research Challenges for the Internet of Services <i>Elisabetta Di Nitto, Dimka Karastoyanova, Andreas Metzger, Michael Parkin, Marco Pistore, Klaus Pohl, Fabrizio Silvestri and Willem-Jan Van den Heuvel</i>	263
25. Survey on P2P Overlay Streaming Clients <i>Alexandro Sentinelli, Luca Celetto, Damien Lefol, Claudio Palazzi, Giovanni Pau, Theodore Zahariadis and Ahola Jari</i>	273
26. Content Adaptation Issues in the Future Internet <i>Theodore Zahariadis, Catherine Lamy-Bergot, Thomas Schierl, Karsten Grüneberg, Luca Celetto and Christian Timmerer</i>	283
27. QoE and *-Awareness in the Future Internet <i>Fidel Liberal, Jose-Oscar Fajardo and Harilaos Koumaras</i>	293
28. A Future Perspective on the 3D Media Internet <i>Petros Daras and Federico Alvarez</i>	303
29. Towards an Architecture for a Real World Internet <i>Alexander Gluhak, Martin Bauer, Frederic Montagut, Vlad Stirbu, Mattias Johansson, Jesus Bernat Vercher and Mirko Presser</i>	313
30. Roadmap for Real World Internet Applications – Socioeconomic Scenarios and Design Recommendations <i>Fabrice Forest, Olivier Lavoisy, Markus Eurich, Jilles Van Gurp and Duncan Wilson</i>	325
31. Context-Aware Systems and Implications for Future Internet <i>Nigel Baker, Madiha Zafar, Boris Moltchanov and Michael Knappmeyer</i>	335
32. Agreeing Upon SOA Terminology – Lessons Learned <i>Vanessa Stricker, André Heuer, Johannes Maria Zaha, Klaus Pohl and Stefano de Panfilis</i>	345
Subject Index	355
Author Index	357

# Future Internet Socio-Economics – Challenges and Perspectives

David HAUSHEER <sup>a,1</sup>, Pekka NIKANDER <sup>b</sup>, Vincenzo FOGLIATI <sup>c</sup>,  
Klaus WÜNSTEL <sup>d</sup>, María Ángeles CALLEJO <sup>e</sup>, Santiago Ristol JORBA <sup>f</sup>,  
Spiros SPIROU <sup>g</sup>, Latif LADID <sup>h</sup>, Wolfgang KLEINWÄCHTER <sup>i</sup>,  
Burkhard STILLER <sup>a</sup>, Malte BEHRMANN <sup>j</sup>, Mike BONIFACE <sup>k</sup>,  
Costas COURCOUBETIS <sup>l</sup> and Man-Sze LI <sup>m</sup>

<sup>a</sup> *Department of Informatics IFI, University of Zurich, Switzerland*

<sup>b</sup> *Helsinki Institute for Information Technology and Ericsson Research, Finland*

<sup>c</sup> *Telespazio, Italy*

<sup>d</sup> *Alcatel-Lucent Bell Labs, Germany*

<sup>e</sup> *Telefónica Investigación y Desarrollo, Madrid, Spain*

<sup>f</sup> *Atos Origin, Barcelona, Spain*

<sup>g</sup> *Intracom Telecom, Greece*

<sup>h</sup> *University of Luxembourg, Luxembourg*

<sup>i</sup> *Department for Media and Information Sciences, University of Aarhus, Denmark*

<sup>j</sup> *German National Association of Game Developers GAME, Germany*

<sup>k</sup> *University of Southampton, IT Innovation Center, UK*

<sup>l</sup> *Department of Informatics, Athens University of Economics and Business, Greece*

<sup>m</sup> *IC Focus, UK*

**Abstract.** Socio-economics aims to understand the interplay between the society, economy, markets, institutions, self-interest, and moral commitments. It is a multi-disciplinary field using methods from economics, psychology, sociology, history, and even anthropology. Socio-economics of networks have been studied for over 30 years, but mostly in the context of social networks instead of the underlying communication networks. The aim of this paper is to present and discuss challenges and perspectives related to “socio-economic” issues in the Future Internet. It is hoped that this will lead to new insights on how to structure the architecture and services in the Internet of the future.

**Keywords.** Socio-economics, Future Internet, networks, services, users, providers, business models, pricing, markets, QoS, trust, user identity, privacy, content, user behaviour, P2P networks, standardization, regulations, value chains, customization, Internet governance

## 1. Introduction

The Future Internet Assembly (FIA [4]) is a European initiative that has recently been established with the goal to shape the Internet of the future. This initiative, which is backed

---

<sup>1</sup>Corresponding Author: David Hausheer, University of Zurich, IFI, Binzmühlestrasse 14, CH-8050 Zurich, Switzerland; E-mail: hausheer@ifi.uzh.ch

by a number of European research projects under the EU Seventh Framework Programme (FP7), follows similar activities in the US (GENI [6], FIND [2]), Japan (AKARI [1]), and Korea (FIF [3]). Over the past decades, the Internet has grown and evolved to unprecedented size. However, its architecture is still based on the original design principles for an academic network in a “friendly” environment. Since then, the Internet has changed enormously both in size and in the way it is being used. In addition to the academic usage, the Internet is now used as a business platform and has become a central part of social life. The types of applications running “over” the Internet exhibit more and more variety and put new requirements to the network providers for how to run and manage their networks.

With an increasing number of users, providers, and services, the Internet of the future is facing problems, including but not limited to issues like scalability and address space limitation. To address these problems, possible solutions like those envisioned by the projects of the FIA and the Future Internet initiatives world-wide range from evolutionary to revolutionary approaches. However, far reaching technological innovations can only be successfully deployed, if their non-technical issues and business potential are taken into account. Any new technology, no matter how excellent, can only succeed in the market if it satisfies, in a sustainable way, the needs of current or potential future users.

Therefore, the aim of this paper is to study the *socio-economic* challenges and perspectives related to the Future Internet. The overall socio-economic context is an important one, as it can significantly boost or hamper the success of an innovation – issues include the “degree of mobility” in the life-style, the balance of “privacy vs. sharing”, the need for security, the importance ascribed to health, and the distribution of wealth. The impact of new technologies on various segments of society such as the young or the elderly has to be appraised with the aim of maximizing benefits for the users (cf. [10], [11]).

The remainder of this paper is organized as follows. Section 2 outlines general socio-economic aspects which are of interest with respect to the Future Internet. Furthermore, Section 3 provides different positions and research orientations highlighting the socio-economic challenges that are faced by the Internet of the future, while Section 4 identifies possible perspectives that can be gained and describes the possible integration paths towards the Future Internet. Finally, Section 5 concludes this paper.

## 2. General Aspects

Future Internet socio-economics studies the relationship between any sort of economic activity and the social life of users. Important socio-economic aspects include markets of Internet Service Providers (ISPs) and Telecommunication Providers, ISP peering agreements and/or transit contracts, as well as customer usage behaviors and selections of content. A study of all these aspects needs to address emerging and disruptive technologies, which effect the user/customer to provider relation, and has to include investigations of (European) regulations for the e-services market and security regulations, as well as the physical environment of e-services in terms of availability – world-wide vs. highly focused (cities) – and dependability for commercial services. This approach will enable to determine (if possible) the economic growth, providers’ revenue maximization, and customers’ benefits.

Traditionally the field of telecommunications is strictly regulated, as opposed to the Internet. When Internet and telecommunications merge towards the Networks of the Future, the required minimum set of regulatory constraints must be defined (*e.g.*, issues like privacy, dependability, or neutrality). Therefore, a view on policy and governance of the network itself is especially prominent today with regard to address assignment and network neutrality. To create the basis for new socio-economic opportunities, efficient solutions on critical issues like security, mobility, quality-of-service (QoS), and economical use of energy are demanded. These requirements ask for innovations on the network layer with focus on topics such as network virtualization, embedded network management capability, functionally rich communication paths, and networks of information.

There are several trends that can be observed over the last 30 years of Internet use. Stakeholders responsible for governance have changed from governments and academics to businesses trying to extract commercial value, where connectivity (ISPs) run as a commercial activity rather than by governments and academics. Moreover, trust between users has reduced dramatically with many users not knowing how to protect themselves and new service models requiring a proper performance level to give an acceptable user experience. Original developers were concerned with large scale system failures but now face quite different attacks, therefore, increasing the necessity of social and legal considerations of participating in “connecting” endpoints.

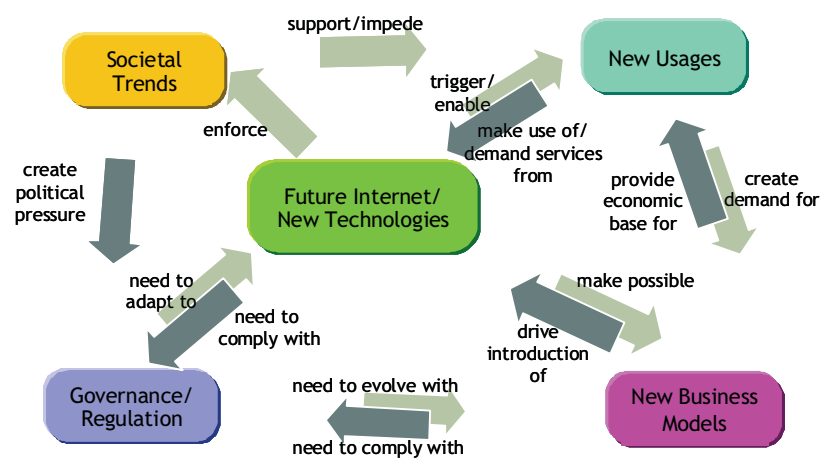


Figure 1. Relationship between technical and non-technical driving forces of the Future Internet

Drawing the appropriate conclusions from the mutual dependence of these technical issues with the non-technical driving forces will be the key for deployment of future Internet innovations. However, as all these technical and non-technical drivers are various, manifold, heterogeneous, it is difficult to draw weighed requirements from each of the drivers in such a complex system of meshed relationships (see Fig. 1). Similarly, it is very difficult to appraise the individual contribution of each specific driver to the overall success or failure of an innovation.



### 3. Challenges

Based on the above considerations, socio-economics challenges can be identified in all domains of the Future Internet including the areas of networks, services, and content. In the following, Section 3.1 discusses the economic challenge faced by all three areas, addressing issues such as who pays, along with consideration for costs, pricing, and benefits. Furthermore, Section 3.2 presents the social challenge with respect to a universal and trustworthy Internet service. Finally, Section 3.3 outlines important socio-economic aspects related to content services provided over the Internet.

#### 3.1. *A Healthy Internet Ecosystem?*

There are many indications that today's Internet technology – combined with the pricing schemes that it supports – does not align well with economics and results in a number of inefficiencies. For instance it is criticized that it does not promote enough competition at the interconnection level, does not generate incentives for network expansion, and the revenue streams generated favor only certain stakeholders in the value chain, which may create incentives for network operators to interfere with user applications without the consent of the users.

The Internet is founded on the simple premise of resource sharing. Shared communication links are more efficient than dedicated connections that lie idle much of the time. Hence the rules applied for sharing are extremely vital for the healthy operation of the Internet ecosystem and directly affect the value of the network to its users. This fact presents a great number of challenges to the Internet research community which can only be addressed by merging the disciplines of computer science and economics. The key question is: what is wrong with today's Internet sharing technologies? Are these consistent with economics? More specifically, since TCP is the dominant sharing technology, is TCP sensible from an economic point of view? Is deep packet inspection (DPI) technology good or bad for the Internet community? Which network sharing technologies justify the end-to-end (E2E) paradigm from an economics perspective? What is required to make peer-to-peer (P2P) a blessing instead of a curse? Are there bad applications or just inefficient combinations of sharing technologies and pricing schemes?

There is no simple answer to these questions as they are very closely related. It relates to the right definition of the nervous system of the network, i.e. the information (the economic signals) that should be generated by the network in order to allow for building the appropriate mechanisms on top of that. And all this should be incentive compatible, i.e. the network should have the incentive to offer this information, and the users should have the incentive to choose and deploy the rest of the mechanisms that would base their operation and decisions on this information. Then the right pricing structures will prevail – for interconnection and for charging application traffic. At the resulting equilibrium, which is again an economics concept, the complete ecosystem will benefit. It is hence clear that designing the right nervous system maximizes the chances to lead to a good equilibrium. But is this existent in today's Internet? Are all these issues about ineffective pricing systems, networks that police selfishly user applications, and weak incentives for network expansion natural to occur or because a bad network nervous system is in place right now?

The well discussed case of P2P traffic generated the amazing spiral of tussles between network operators and users. Network providers offered several pricing schemes

but flat-rate prevailed. P2P users took advantage of that and by increasing the number of their parallel TCP connections absorbed all the bandwidth of the network making performance dismal for the traditional interactive web-browsing users. Then network providers added capacity, but this got into the black hole since almost all of it was again used by the P2P traffic. Therefore, the P2P users were considered as a curse for the network. But were they doing something wrong? Not really, since they were just operating according to the rules of the game. Finally, the operators decided to control this gas-like ever expanding P2P traffic using volume caps. This improved the situation but created other bad side-effects, since it was throttling down the interactive fraction of the traffic as well. Hence, a more refined control on the application traffic was needed and so the operators started to deploy DPI boxes. These on one hand improved the traffic situation since they only throttled P2P traffic. But on the other hand, these boxes became a tool to deploy other discrimination policies like killing VoIP traffic and degrading performance for traffic from competing networks and applications. Therefore, the users reacted by encrypting their traffic to stop DPI boxes overwriting their TCP connections and networks made DPI boxes smarter to overcome this problem. And this will continue...

It is clear that this network policy is not fixing the problem, it is rather trying to hide it. But what is the real problem? TCP seems to be definitely part of it, since it does not take into account the overall connection activity over time. And combined with a poor pricing scheme it provides the incentives for P2P users to open as many connections as possible. Ideally, a good flow control scheme for P2P traffic would detect the times that demand for bandwidth by interactive users is high and push most of its transmissions at times that this demand will be lower. Such a strategy would make everybody better off. But unfortunately TCP is not equipped to behave like that.

So what is needed for all this to work better? Some recent research indicates that the basic technology that must be deployed by network operators is congestion signaling in a way that this information is visible by the networks themselves. Today congestion signals exist in the form of lost packets, but these are known only at the edges of the network and are invisible to the network itself. Network signaling technologies have been designed that allow this information to flow in an incentive compatible way throughout the Internet. Then pricing structures based on congestion volume will provide for the incentives that the right congestion control algorithms will be deployed by the end systems. Using this new nervous system each application will be charged for the cost it imposes on the rest of the applications that share the Internet. Pushing the information about true cost back to the economic agents that are responsible may be the necessary and sufficient condition according to economics. Then each agent (including the network operators) will be driven by the right incentives to make the appropriate actions and choices, in a completely distributed fashion. In this new context BitTorrent will be a blessing instead of a curse because it actually helps balancing load better and discover unused capacity for the network operator.

### *3.2. The Social Challenge of a Universal and Trustworthy Internet*

Besides the economic dimension, the Internet faces an important social challenge. The current Internet penetration has reached 20% worldwide and should reach 30% by 2015 and 50% by 2020. Broadband access to telecommunication network capacity and services must be guaranteed “anywhere-anytime” to universally exploit the Internet –

present and future – which is becoming a fundamental service that communities use and rely upon. As such, the Future Internet shall be able – among others – to support daily life in developed countries as well as within developing countries. Telecommunication infrastructures must be conceived to guarantee access to the Future Internet also where currently it is poor.

However, the IP address space is depleting fast with only 15% left and expected to be exhausted by 2010. This may not only be the end of the end-to-end (E2E) model, but also the end of the Internet itself. To fix this problem of the current Internet is a big and large-scale task and challenge. With virtually unlimited address space, the new Internet Protocol IPv6 has been designed to cater for the many deployment scenarios, starting with an extension of the packet technology and, therefore, supporting IPv4 with transition models to keep IPv4 working even for ever, and then to cater for new uses and new models that require a combination of features that were not tightly designed or scalable in IPv4 like IP mobility, E2E connectivity, E2E services, and ad hoc services; to the extreme scenario where IP becomes a commodity service enabling lowest cost deployment of large-scale sensor networks, RFID tags, IP in the car, to any imaginable scenario where networking adds value to commodity.

Moreover, with the Internet getting more and more universal, the issue of trust in users and services will become a key challenge in the Future Internet. The Internet nowadays has 1 billion users that access to more than 30 billion pages, most of them basically static, and only 30% of them built by companies. However the number of public web services is only around 30.000 [9] since most of the services are in-house and, therefore, restricted to close environments. It is clear that in this context trust is easy to manage. However, as mobile, wireless, optical, and broadband communication infrastructures become even bigger and more interdependent, the number of Web services is expected to grow exponentially in the years to come. In particular, more companies will publish their offers as services accessible through the Web inspired by the success of examples like Amazon, eBay, and Yahoo. Moreover, the Web 2.0 has popularized concepts such as mash-ups and RSS feeds and thereby illustrated comparatively simple means for business networking and business flexibility. Additionally, efforts to turn the Web into a general platform for accessing and interconnecting arbitrary devices and services are maturing and may turn billions of diverse devices such as household appliances into accessible Web services. Finally, humans may act as “mechanical Turks” acting like a Web service to acquire the input data required to perform a certain requested task. These trends lead to a future Internet of billions of services in a network of equals – large enterprises, small and medium enterprises (SMEs), and citizens – in which these services will be indistinctively produced or consumed by “prosumers”.

In this new context trust will become a major issue and Web 2.0 technologies are already starting to support trust and reputation within and between computers and humans. The use of technologies like FOAF (Friend Of A Friend) in the field of services will allow software agents, and humans, to gain information on the reliability and reputation of a service. For example, if a known and used service (*e.g.*, a flight booking service) has been combined successfully with a geographical location service, it is important that the user is able to gain any information on the newly introduced service.

### 3.3. Socio-Economic Aspects of Content Production and System Customization

In particular content is becoming more and more important for the Future (Media) Internet, including technologies for content manipulation and transmission, as well as content creation. Content itself sometimes acts as a technology driver. For example, today content is created on multiple devices, largely self-organized by communities and centered on aggregators. Business models, human machine interfaces, and the cultural implications affect the technological success. Content is also increasingly responsible for the definition of technology standards, leading to synergies between different delivery platforms and different media forms. Know-how on cameras, recorders, production equipment, and displays are the “vanished sciences” for Europe.

Increased customisation is an opportunity to empower consumers through greater choice and for innovative providers to create new business models and associated revenue streams. However, customisation also exposes existing business models to additional risks. Consumers will exploit technologies in ways that cannot be envisaged and in some cases new usage will break existing business models either legally or illegally. In the music industry it can be seen that the cost of manufacture and distribution of digital music reduced to zero and that IT networks actually can destroy value, with faster networks destroying value more efficiently. In addition, the Internet’s ability to increase the level of indirection from real-world consequences (*e.g.*, criminal prosecution) through anonymity and limited regulation means that normal emotions (*e.g.*, fear) used to temper unacceptable behaviour and risk taking, are experienced to a lesser degree than similar face-to-face interactions (*e.g.*, stealing a CD from a shop compared to downloading an illegal MP3 file).

Of course today customisation is limited, and in fact the most successful business models focus on controlling all of the pillars of the supply chain (like iTunes) rather than offering flexibility. Even with limited flexibility significant threats exist to some industries, especially businesses dependant upon value creation from digital information assets. A new foundation for ICT services may be needed for economic sustainability of the Future Internet. Many challenges exist that require real innovation. For example, how can business models evolve to support underlying notions of value and value creation in the new marketplaces, and the role of ICT as potentially a utility (*e.g.*, interoperability services) and value enabler? How can stakeholders assess and mitigate economic threats introduced by greater customisation in service offerings? How can the service economy provide the right amount of choice and customisation to avoid monopolies and to support economic growth? These are all challenging questions that need to be answered for the Future Internet to become a successful technology.

## 4. Perspectives

Based on the above challenges, this section outlines new insights and perspectives and tries to describe possible integration paths towards the Future Internet. Section 4.1 outlines the need for multi-stake-holder societal control through Internet governance. Furthermore, Section 4.2 presents implications of new generations based on the expected evolution of user behaviour. Finally, Section 4.3 discusses the question of one or multiple Internets as potential paths to a global Future Internet.

#### 4.1. *The Need for Internet Governance*

The governance aspect will become an important part of the Future Internet, in particular if it comes to issues with a public policy component addressing aspects like privacy, security, freedom of expression, intellectual property rights, and data protection (cf. [8]). The term “Internet Governance” emerged in the early 1990s to describe elements of the management of critical Internet resources, in particular the domain name system (DNS), but also IP addresses, root servers, and Internet Protocols. It was based on a conceptual understanding of self regulation and private sector leadership, of “Governance without Government”. One concrete result of this discussion was the establishment of the “Internet Corporation for Assigned Names and Numbers” (ICANN) in 1998.

In 2002, when the UN World Summit on the Information Society (WSIS) started, Internet governance became the issue of very controversial diplomatic negotiations. The government of the Peoples Republic of China argued that private sector leadership was good for one million Internet users but governmental leadership would be needed for one billion Internet users. The controversy “private sector leadership vs. governmental leadership” led to the establishment of the UN Working Group on Internet Governance (WGIG). WGIG defined “Internet Governance” as the development and application by governments, the private sector and civil society, in their respective roles, of shared principles, norms, rules, decision-making procedures, and programs that shape the evolution and use of the Internet.

WGIG rejected the idea of one “single leading Internet organization” but proposed a concept of “Multistakeholderism” and a “Multilayer Multiplayer Mechanism” (M3) where the various stakeholders are involved according to their specific roles and responsibilities. Stakeholders should enhance their cooperation via better communication, coordination and collaboration (C3). As a result, the Internet Governance Forum (IGF) was established as a UN led innovative multistakeholder discussion platform without a decision making capacity. Multistakeholderism should be taken as a guiding principle for NGN governance, including the governance of the ONS and the “Internet of Things”.

#### 4.2. *Expected Evolution of User Behaviour*

Additionally, in order to guarantee the success of the Future Internet design, the end users’ behavior needs to be characterized. Therefore, it is important to analyze the current usage trends as well as to infer the end users’ expectations. Taking a look at today’s Internet usage, it is easy to realize that multiple applications/services (*e.g.*, social networks, P2P, streaming, gaming, voice, videoconference) are being used from multiple end users’ devices (*e.g.*, PC, game consoles, mobile phones, digital assistants). Of course all these services are not used by the same population segments; the Generation X mainly uses Web and e-Mail applications from a PC, while the Generation Y is characterized to be all day connected to access any service. Therefore, this generation mainly values the reliability and ubiquity of the connectivity as a service itself. Moreover, it seems that a new generation of people born in the 21st century can be distinguished, being its main attribute their capability to adapt to continuous changes and their wide knowledge of the technology. This new generation will be ready to use new advanced network services, able to meet the new application performance requirements that could be expected in the future.

From a technical perspective, the Internet traffic analysis is a powerful tool to really know the evolution of the users' preferences as well as their impact on the global amount of traffic that must be carried on future networks. Taking into account the IP traffic study reported in [7], the amount of traffic that needs more than best effort capabilities is growing up.

#### 4.3. *One or Many Internets?*

Finally, a critical issue in the Future Internet research is the current proliferation of separate efforts due to the various initiatives world-wide. This may on one hand be good for innovation, as it can produce more ideas. However, if initiatives remain separate throughout the development of the Future Internet, many technologically incompatible Internets could emerge. In contrast to the current global Internet, these separate Internets could cause market fragmentation and even social seclusion. To avoid such adverse possibilities, design and implementation of the global Future Internet should proceed with a growing degree of cooperation between initiatives.

The origin of current initiatives for the Future Internet is largely correlated with the Internet connection density. It seems that people who use the Internet more are those who want to change it. Most of the Future Internet initiatives like the FIA, FIND, AKARI, and FIF are accompanied by test-bed projects so that researchers can conduct large-scale experiments without affecting the current Internet. There have been some investigative interactions between these initiatives, but so far nothing concrete has emerged for future collaboration. Is this simply bureaucracy or a deeper pathogenicity? After two decades of demonstration, the importance of the Internet is now widely acknowledged. This knowledge has long transcended the esoteric circle of researchers and technologists and now spans almost all groups and levels of society. Value usually begets the desire for exclusive control of the value-generating artifact. To summarize the situation, the stakes are high and the stakeholders are many. So, governments, organizations, corporations, and communities could conceivably promote Future Internet research agendas that would allow them independent control over the produced system. The mere separation of Future Internet initiatives, if left unchecked, could become a schism leading to many incompatible Future Internets.

A market of competing Future Internets might seem beneficial, but technological fights, like the recent battle for the future DVD format, often result in long and expensive stalemates. During such a Future Internet standoff, end-users will be baffled as to which Internet they should join, eventually raising its value and decreasing the value of the other Internets. Equipment manufacturers will either have to choose sides risking a wrong decision, or support all competing technologies and risk creating unstable products. Network operators and service providers will similarly have to gamble on which Internet to support. Even researchers will not know on which area to publish. Competing Future Internets will have difficulties in reaching critical mass and so reaping the benefits of positive network externalities. In the meantime, effort will be duplicated, the market will be fragmented, and cost will increase for all interested parties.

What is then a possible path toward a global Future Internet? A potential plan would initially keep the Future Internet separated to stimulate competition on the generation of concepts. Peer review would also quickly eliminate fundamentally flawed ideas. Then, as surviving ideas move to design, initiatives should collaborate on large-scale, realistic



experiments by joining their test-beds. The interconnection of separate test-beds will proceed gradually to construct a global Future Internet test-bed while blurring the borders of separate initiatives. Higher-level stakeholders, like businesses and governments, will increasingly be able to experiment on the global test-bed, without affecting the current Internet. If these experiments demonstrate real socio-economic value, the global test-bed will naturally become the seed of the global Future Internet.

## **5. Conclusions**

As a general conclusion, it can be stated that the end users' behavior is hard to predict due to the wide variety of services and applications being used; however, due to this wide Internet service portfolio, the end users really perceive the value of the connectivity as a service that should be reliable, ubiquitous, secure, neutral, flexible and able to adapt to new traffic performance demands. These attributes will be required to manage the new traffic demands, according to the end users' requests, and also able to provide the reliability and capabilities required to emerging services such as e-health or sensor networks.

Taking into account this brief analysis, it is hard to think on an Internet just based on the vertical integration of the services that could lead to the balkanization of the Internet, since it would not be possible to address the new emerging services and their requirements in such a close environment; in fact, it is proposed to build up an innovation friendly framework able to support the new traffic demands with higher capabilities in both wired and wireless networks, with guaranteed network performance and with open interfaces to allow the end users to show their preferences.

Socio-economic forces could push Future Internet research towards an arena of competing Future Internets. This scenario would hinder further industrial growth and innovation, limit business opportunities, and confuse end-users. To avoid such problems, separate Future Internet initiatives should be encouraged to compete during the concept generation phase and to collaborate during the design and construction phases. Europe has considerable experience in aligning separate efforts, so it should swiftly promote an increasingly cooperative strategy for the development of the Future Internet.

## **Acknowledgments**

This work has been performed partially in the framework of the EU IST projects SmoothIT, PSIRP, ISI, 4WARD, SOA4ALL, EFIPSANS, EURO-NF, 4NEM, IRMOS, and COIN. Valuable input to this work has been received from the FISE working group [5]. Additionally, the authors would like to acknowledge discussions with all of their colleagues and project partners.

## **References**

- [1] AKARI: Architecture Design Project for New Generation Network. <http://akari-project.nict.go.jp/eng/index2.htm>, Last accessed Nov 2008.
- [2] FIND: Future Internet Design. <http://www.nets-find.net/>, Last accessed Nov 2008.

- [3] FIF: Future Internet Forum. <http://fif.kr/>, Last accessed Nov 2008.
- [4] Future Internet Portal. <http://www.future-internet.eu/>, Last accessed July 2008.
- [5] Future Internet Socio-Economics (FISE) Working Group. <http://www.smoothit.org/wiki/pmwiki.php/FISE>, Last accessed July 2008.
- [6] GENI: Global Environment for Network Innovations. <http://www.geni.net/>, Last accessed Nov 2008.
- [7] Global IP Traffic Forecast and Methodology: 2006-2011, Cisco white paper, 2008.
- [8] D. Lazer: Regulatory Capitalism as a Networked Order: The International System as an Informational Network. *Annals of the American Academy of Political and Social Science*, Vol. 598 (1), pp. 52-66, 2005.
- [9] The seekda Web Services Search Engine. <http://seekda.com/>, Last accessed Nov 2008.
- [10] B. Wellman: Computer Networks As Social Networks. *Science*, Vol. 293, pp. 2031-2034, 2001.
- [11] B. Wellman. Connecting Community: On- and Offline. *Context*, Vol. 3 (4), pp. 22-28, 2004.



## Challenges of Internet Evolution: Attitude and Technology

Jon Mikel Rubina

European Software Institute, Zamudio, Spain  
jonmikel.rubina@esi.es

**Abstract.** This paper analyzes two challenges of Internet evolution by evaluating the dynamics of change against the outstanding aspects of the current situation. To understand these challenges, a model that exposes the factors that drive the dynamics of Internet evolution is discussed first. The conclusion drawn is that attitude and technology are the two vectors that lead the evolution of Internet. Secondly, the contemporary context with regard to technology development and information society attitudes is analyzed. Finally, it is concluded that the necessity of a paradigm shift in the attitude of organizations and the threat of unsustainable technology development are two of the main challenges towards the future Internet.

**Keywords:** future internet evolution attitude technology socioeconomic

### 1 Introduction

There is a great turmoil about what the Internet is going to become in the future. Within a clear dichotomy formed between “clean-slate” and evolutionary approaches, this study examines the socioeconomic challenges that must be confronted in an evolutionary context of the Internet as described in [7]. However, it is not the purpose of this paper to vindicate the evolutionary perspective but to clarify its social and economic drivers provided that evolution is a fact.

This text holds the thesis that the evolution towards the future Internet faces some challenges or threats that may decelerate development and lead it to an impasse. To understand the challenges that threaten the future Internet, this paper outlines an evolution model of Internet from a socioeconomic perspective. This model extends the views exposed in [5], [17] and it helps understand the importance of attitudes and technology.

Afterwards, comparing this model against the current socioeconomic situation with regard to attitudes and Internet technology leads to the conclusion that there are two major risks that may impede the actual evolution of Internet. On the one hand, the poor capitalization on Internet technologies is presented as a major drawback of current Internet trends. And a paradigm shift on organizations’ behavior is suggested as a solution to overcome this situation. On the other hand, the unsustainable way in which Internet technologies evolve in certain cases is stressed. Thus, it is proposed to

add emphasis on a horizontal basic requirement for technology development: efficiency.

Although this paper renders these two challenges essential, it would be too partial to deem them the only main issue towards the realization of the future Internet. Thus, it is acknowledged that there are many other aspects which are out of the scope of this analysis that will have to be solved to achieve the future Internet.

## **2 Internet Socioeconomic Evolution Model**


According to economic theory, in a competitive environment with limited resources, rational species' behaviors change so that individuals can maximize a utility function dependant on these resources and the satisfaction of their needs. The behaviors of rational individuals are constrained by their attitudes. And attitude is understood in this text as the disposition, tendency or orientation of mind with regard to a thing, particularly with regard to technology. These attitudes determine the range of behaviors that are valid according to mental disposition. To let a wider range of behaviors be assessed by a rational individual, a change in attitude towards the elements of reality involved in the satisfaction of their needs is needed.

It is in the very basic nature of the *Homo sapiens* to use and develop tools that make use of a certain technology which assists in performing an operation to satisfy human beings' needs. For a given set of attitudes, technology evolves to satisfy the needs by means of the behaviors allowed by that set of attitudes; i.e. attitudes demand new technologies to enable the behaviors that they allow so that those behaviors that better perform the satisfaction of needs can be executed. There is a limit at which better behaviors cannot be enabled by any new technology, because they cannot be conceived according to existing attitudes. At this limit, the only way to achieve better performance in a competitive environment is a change in attitudes. When the attitudes involved in the satisfaction of needs evolve, new behaviors are considered and consequently new technologies that enable these new behaviors can be developed. Alternatively, if new technologies are not developed there is a limit at which, regardless the number of new attitudes created they cannot materialize the new behaviors allowed, as these are not enabled by existing technologies. Thus, the evolution cycle gets disrupted since those new behaviors which cannot be materialized can also not demand new technologies.

The foundations of this simple co-evolution model are that attitudes demand new technologies to improve the satisfaction of needs and, reciprocally, the new technologies enable the improvements allowed by new attitudes. This reciprocity in the evolution of technology and attitudes is represented in Fig. 1 as two interleaved helixes one representing technology evolution and the other one representing attitude evolution. Sometimes the evolution towards new attitudes is supported by existing technologies that enable new conceptions of new behaviors to satisfy needs. In other occasions, new technologies are supported by existing attitudes that demand advances which enable the achievement of those behaviors which, although accepted, are not technologically feasible.

Internet is a tool involving technologies that is used by users to satisfy their needs according to their attitudes. But this tool has another two characteristics that determine the need to complete this model with two more requirements. These

characteristics are: (1) Internet locates at multiple locations both geographically and within society and (2) Internet has no single ownership.



**Fig. 1.** Co-evolution of technology and attitudes

Changes are said to be adopted by Internet when they are involved in the behavior of a wide community of Internet users when they use Internet. It is not the point of this paper to determine what the “critical mass” of users required is; but it is clear there will be a minimum number. If that minimum critical mass is not reached in a certain period of time, changes will stagnate or disappear; and cannot be said to be adopted by Internet.

According to the two Internet characteristics, two requirements [17] exist for changes on attitudes and technologies to be adopted by Internet. (1) Universality: new technology can be used and new attitudes can be adopted by any user anywhere; and (2) Independency: New technology can be used and new attitudes can be adopted by some users even if others do not use or adopt them, i.e. there is no need to orchestrate change.

In addition to these two requirements, the Internet evolution model is completed by two fundamental characteristics: creativity and economic feasibility. The first one must be present whenever a change occurs intentionally. Unless it is a hazardous change, creative minds conceive changes out of their knowledge and experience. The second one is necessary so that evolution is possible, according to economy, in a context of competition for scarce resources, i.e. so that the underlying agents of evolution have the economic capabilities to invest on new technologies and adopt new attitudes obtaining a benefit out of them. Therefore, creativity, economic feasibility, universality and independency are necessary requirements for the co-evolution explained before to be realized. They ensure that changes will occur, that they will be financed, that they will not be reduced to a specific community and that they have the potential to be virally spread through the Internet.

### 3 About Attitudes

In the model presented, attitudes are said to be one of the vectors of the Internet evolution. The following sections underline a duality on the evolution of attitudes in information society. The analysis of this aspect reveals a challenging burden for Internet evolution. In according to this, a solution to overcome this situation is described.

### 3.1 Attitude Duality in Information Society: Individuals and Organizations

User attitudes are rewarded some benefits on using the Internet and they demand new technologies. These users constitute the so-called information society. According to their interest, two types of users can be identified in information society. These types are (1) individuals that pursue to satisfy their particular needs and (2) those that want to satisfy the objectives of a given organization they belong to. In terms of goals, there are two interests using the Internet: interests of individuals and interests of organizations.

Both individuals and organizations use the Internet to satisfy their needs. But in recent times, and particularly during the evolution of the web towards the web 2.0 [15], individuals prove to obtain far more benefits than organizations do. Blogs, wikis, social networks, file sharing, podcasts, online applications and other Internet progresses help individuals to innovate in satisfying their needs more efficiently; accordingly, the success of these new technologies and behaviors coalesces into their increasing number of users [19]. Furthermore, it is widely admitted that these new technologies have co-evolved together with attitudes [15]. Among others, the main components of this shift in the attitude of individuals are: participation, collaboration, confidence and sharing. These attitudes, enabled by existing technologies (HTML or JavaScript), have subsequently demanded the development of new ones (folksonomies, AJAX or RSS)

In spite of the success of new technologies among individual users, this success does not seem to spread likewise through enterprises. A recent report [1] has warned that, although companies are introduced to web 2.0, not all of them manage to find its benefits. Sometimes, they even abandon these incipient tools, because they do not always turn out to be representing a greater value for their businesses.

This study recognizes that certain organizations are undeniably profitable thanks to new Internet technologies, but this benefit tends to focus on the *tertiary* technological sector, rather than pervasively spread through every sector of economy including the *primary* and *secondary* sectors. The problem is that societal evolution of new economic sectors fundamentally relies on higher productivity rates in the primary and secondary sectors.

### 3.2 Duality Analysis from the Internet Evolution Model Perspective

This text argues that the reason for these two paces of the reward obtained from new Internet technologies is the absence of a paradigm shift in the attitudes of organizations. This statement is based on the aforementioned Internet evolution model: organizational attitudes towards Internet are not changing both to demand new Internet technologies tailored to their needs and to capitalize on the benefits of existing developed technologies. A paradigm shift in attitudes would demand new technologies to enable new behaviors that would help organizations improve their competitiveness. Scarcity of new attitudes not only does not let new technologies enable new behaviors but it does not even let new behaviors emerge and capitalize on existing technologies.

This problem presents two aspects: dragging technology to meet organizational attitudes and taking advantage of existing technologies. Firstly, Internet technologies are being demanded mainly by new attitudes of individual users, not by organizations;

consequently, these technologies favor improvements of the behaviors of individual users instead of those of organizations. Secondly, attitudes in organizations do not evolve fast enough to meet the opportunities offered by new technologies and to take advantage of the current technological status quo.

Organizations play a fundamental role within information society, supporting the economic power of other organizations and of individuals. If most of the organizations continue to fail realizing the benefits of new Internet technologies, as in [1], and if evolution continues to focus on individual users and on isolated types of technological organizations within the tertiary sector, according to the model proposed, it is claimed that Internet technology evolution and its benefits may reach an impasse. Two of the requirements for Internet evolution are threatened: universality and economic feasibility. Firstly, lack of universality leads to unbalanced distribution of the benefits of the Internet among sectors what, in turn, deteriorates economic feasibility even more. Secondly, the economic feasibility of the Internet evolution currently is mainly limited to the economic power of organizations within the technological tertiary sector and that of individual users. These are economically dependant on the performance of other organizations in the primary and secondary sectors. As explained before, the performance and competitiveness of these organizations are not strongly favoured by current Internet trends, thus, the economic capabilities required to ensure the economic feasibility of investments in new Internet technologies and attitudes are limited.

On the other hand, the problem does not consist only of the fact that organizations are not intensively compensated by new Internet technologies, but of the fact that these technologies can be counterproductive. These new Internet technologies reduce the value of industries by restructuring traditional forces of competition: increasing competitive rivalry, reducing entry barriers, shifting bargaining power to clients, etc [16]. At certain occasions, these conditions turn price into the only choice on determining what the competitive advantage is. As a result, competition on price is making the headroom between cost and price lower, what in turn reduces the value created for the industry.

In summary, Internet evolution may stagnate due to waning universality and a technological evolution that is not economically sustainable for the whole information society: a bubble in which some organizations are supporting new Internet technologies that do not seem to pay off.

### **3.3 Proposed Solution: Paradigm Shift in Organizational Attitudes**

It has been seen that new Internet technologies not only may not reward [1] but may even weaken [16] the organizations that actually support that evolution. The model exposed holds that this problem is due to an unbalanced co-evolution of Internet in information society: attitudes of individual users are broadly changing but that is not the case of organizations in general.

The obvious primary solution consists of promoting innovative paradigm shifts on the attitudes of organizations towards Internet. These changes would enable new behaviors to let them perform better their objectives. This would allow organizations to shape the evolution of Internet technologies to fit their needs.

Nonetheless, although the web 2.0 is a representative tool to exemplify the problem, new attitudes not only engine the adoption and the evolution towards the

web 2.0; but they also drive the requirements for other technologies of the future Internet (e.g. business processes instantiation, collaborative frameworks, complex service architectures, etc.)

Together with this, this paper proposes an interim secondary solution that consists of promoting paradigm shifts on the attitudes to enable new behaviors capable of better exploiting the benefits of the existing Internet technologies that were previously motivated by individual users. To attain these benefits, changes that have been experimented by individual users are suggested: collaboration, confidence, participation and sharing.

However, this paradigm shift does not consist of applying, on an “as is” basis, patterns existing in communities of individuals to enterprises. Competition in business is much fiercer. Collaboration, confidence, participation and sharing must be reconciled with competitiveness. New behaviours have to assess how to create value out of confidence, sharing, participation and collaboration. The modifications that these new attitudes may represent for the forces of competition must be evaluated and new operational advantages must be pursued on the adoption of these new attitudes. Consequently, the shift will be implemented by means of behaviours complementary to those of individual users involving brokering roles and arbitration policies among other mechanisms.

Theoretically, a lack of collaboration and confidence make organizations incur in transaction costs [4]. Reducing transaction costs means increasing value for companies and providing operational advantages. This reduction can be achieved by means of new attitudes of collaboration and confidence that synergize with new Internet technologies. It is important to note that these new attitudes are intended to be applied to existing and new relationships intra and inter organizations, thereupon the paradigm shift proposed goes beyond (but along the same line of) the hype of other proposals, as Enterprise 2.0 [13], which are more focused on social networking applied to organizations’ intranets.

A recent success story involving a paradigm shift in organizational attitudes is the [www.gridconnection.co.uk](http://www.gridconnection.co.uk) web site [11]. This website lets any energy promoter obtain feasibility assessments for new renewable energy plants at different geographical locations in terms of the most efficient alternative to connect these renewable energy sources to the electricity grid. This application has managed to create value for the energy promoter, for the application provider, for the electricity providers and for the society at large. The owners of the existing electricity networks have been called to participate in this project, they have had to collaborate and they have shared the data of their networks. Now, an increasing number of energy promoters relies on the data and the assessments provided by this website to undertake new renewable energy projects.

Certainly, if a paradigm shift pervades organizations, performance of every sector of economy can be rewarded by new Internet technologies and can pull Internet technology evolution. Other businesses may appear out of new attitudes and new value can be created for economy with these new attitudes. Unexplored, unexploited and unmerged information might be the clue to follow so that companies can capitalize on new information technologies and gain competitive advantage. For instance, some use cases that claim the necessity of this shift are: freight transportation companies still doing sometimes their return journeys without load, consumers unable to track the value chain of products they buy, wealth flows on financial market remaining hidden to financial bodies, commodity markets



plummeting upon excessive production and rising upon drought, lack of exhaustive metrics on business processes enabling better managerial decisions, etc.

## 4 About Technologies

The Internet socioeconomic model developed unveils that, if attitudes regularly demand new technologies, the other vector of Internet co-evolution will evolve whenever universality, creativity, independency and economic feasibility are present.

Next, a simple pattern of technology evolution is developed to understand the vulnerabilities of Internet technology. This pattern analyses technology evolution from an economic perspective. Afterwards, threats on this pattern are exposed and a key guideline to avoid them on Internet evolution is posed. It should be understood that the length of this paper limits this dissertation to a theoretical level and use cases are accordingly relegated to brief mentions.

### 4.1 Technology Efficiency Pattern

Let us define Internet technology as the branch of knowledge that develops technical means (i.e. tools) using available resources to provide a new functionality for the Internet. Therefore, technical means transform resources into these new functionalities.

Environments in which Internet technologies are used contain a limited pool of resources  $\Gamma \in \mathfrak{R}$  available. As a consequence, as scarcity is present on Internet technology, economic principles apply and limit the evolution of technology. These resources are memory, bandwidth, microprocessor execution time, etc. Obviously, these resources are limited because of several factors, basically economical and technological. For instance, bandwidth for a given physical medium is technologically limited. It can be increased by aggregating in parallel several units of the physical medium. But the number of units that can be aggregated is limited by the price of the physical medium, thus, it is economically limited.

Technologies within these environments may depend on the functionality provided by other technologies to perform theirs (e.g. a software application technology relies on the technology of the operating system; or an application protocol relies on the services of a transport protocol) or may not (e.g. Firefox and OpenOffice sharing execution time; or HTTP and FTP protocols sharing bandwidth)

The relative consumption  $\eta \in \mathfrak{R}$  of a technology is defined as the number of resources  $\rho \in \mathfrak{R}$  used on applying that technology in a given environment with a pool of resources  $\Gamma$ , i.e.  $\eta = \rho/\Gamma$ .

Internet technologies are used within complex environments at which available resources are shared among several technologies (e.g. a computer with microprocessor execution time and physical memory available for all the software technologies hosted; or a transmission medium offering bandwidth to the transport layer and to the application layer).

According to economic principles, **the scarcity condition** holds: the aggregated relative consumption of an environment  $\eta_e$  with respect to a certain type of resource  $\rho$

is lower or equal to one.  $\eta_e = \sum \eta = \sum \rho / \Gamma \leq 1$ . Consequently, the number of technologies that can be integrated in an environment is limited; and so is the number of functionalities that can be provided by the technologies developed within that environment. Therefore, technology for a certain environment will evolve until it exhausts all the resources available.

As the resources available to Internet technologies are usually measured in terms of information units (e.g. bits, bytes, symbols), the efficiency of Internet technologies is measured as the amount of information units that a technology actually uses from a given amount of information units provided to implement a functionality. Information units are calculated as the entropy of the variables or set of symbols  $x_1, \dots, x_n$  used by a technology for which a probability mass function is defined  $p(x) \forall x \in \{x_1, \dots, x_n\}$ .

Furthermore, the **efficiency of an Internet technology**  $\eta_i$  is defined as the relationship of the entropy of the variables used by the technology  $H_i(p(x_1), \dots, p(x_n))$  with respect to the maximum entropy available for the given resources  $H_{i,max}(1/n, \dots, 1/n)$ , i.e.  $\eta_i = H_i / H_{i,max} \leq 1$ . It is obvious, that for a constant level of efficiency, to attain greater functionalities an Internet technology needs more entropy, thus more available entropy and more resources.

For instance, a storage technology providing error checking functionality may have 10 bits available to store information. A technology  $Ta$  that implements error checking by means of redundancy will be capable of storing 5 bits; its efficiency is  $\eta_a = 5/10$ . Another technology  $Tb$  that implements error checking by means of XOR'ing 9 bits into the last bit will be capable of storing 9 bits; its efficiency is  $\eta_b = 9/10$ . On comparing these two technologies the necessity of better or lower error checking functionalities has to be compared against the amount of resources available.

The **aggregated efficiency of an environment** used by several technologies is defined as  $\eta_e = \sum H_i(p(x_1), \dots, p(x_{pi})) / H_{e,max}(1/\sum \rho_p, \dots, 1/\sum \rho_i) \leq 1$ . It represents the capability of a set of technologies to use the resources that have been required from its environment.

For a given distribution of resources among technologies, it can be demonstrated that  $\eta_e \propto \eta_i$ , i.e. the aggregated efficiency of a set of Internet technologies in an environment is higher when the efficiency of each technology used is higher. Equally, it can be demonstrated that the aggregated efficiency is higher when the correlation between symbols among technologies is lower, i.e. technologies do not replicate functionalities.

Consequently, to maximize the usage of the limited pool of resources  $\Gamma$ , (1) the efficiency of each technology with respect to the resources it uses must be maximized and (2) the correlation of functionalities among technologies must be minimized.

To attain the maximum benefit of a pool of resources, a set of information technologies has to be implemented. This set is not usually implemented instantly but progressively. As it has been explained previously, this progress is implemented through a technology evolution process.

To achieve a certain set of functionalities a given aggregated entropy is required  $\sum H_i(p(x_1), \dots, p(x_{pi}))$ . If each technology within the set is not implemented with maximum efficiency  $\eta_i$ , the resulting aggregated efficiency  $\eta_e$  will not be maximum. Therefore, the amount of resources  $\sum \rho$  required to provide the available entropy that is needed to implement the functionalities will be higher. But the amount of resources



available is limited due to the scarcity condition  $\sum \rho \leq I$ . This means that for a given number of resources available, the number of functionalities that can be obtained is limited by the level of aggregated efficiency. The aggregated efficiency, in turn, depends on the correlation among technologies and the efficiency of each technology.

In summary, the evolution of Internet towards new functionalities is limited by the aggregated efficiency of Internet technologies on using existing limited resources. If this efficiency is not observed, Internet technology evolution may reach physical feasibility limits before providing benefits the information society can take advantage of.

#### 4.2 Harnessing Sustainable Internet Technology Evolution

The threat summarized in the technology efficiency pattern is relative. It is relative to the number of resources available. If the number of resources available is very big, almost infinite, regardless the aggregated efficiency the number of functionalities reachable can be said to be unlimited.

But there are some cases at which resources are clearly limited, for example: the bandwidth and the queue size at a router. In these cases, the appearance of technologies that provide a functionality successfully but inefficiently impede the development of future technologies as they are cannibalizing shared resources. This leads technology evolution to an impasse at which the only exits are either looking for more resources or reengineering inefficient technologies.

The problems of sustainability on Internet growth have been widely discussed [14] but are not always taken into consideration while developing new technologies. And that is the reason that may materialize the caveats reported. Technology efficiency is not always a primary requirement on developing new technologies, what has been traditionally attributed to a greater focus on functionality rather than on the technical implementation: *“The details of technical design are not regarded as relevant in themselves, so long as the capabilities continue to evolve to meet the demands of growing markets for enhanced electronic communication capabilities.”* [3]

Actual realizations of these risks have been recently witnessed with regard to the evolution of web services technologies. These technologies resemble the pattern just explained and have been claimed to suffer from low levels of aggregated efficiency [10] It has been on applying web services to embedded systems, with limited bandwidth, storage and processing capabilities, when the threat has proven to be real. The answer to this problem has required to reengineer this technology for devices WS4D [20] According to the technology evolution pattern presented, this analysis states that this expensive reengineering process that may hold back web services evolution, could have been avoided if aggregated efficiency had been a primary requirement of its design. Equally, as web services result to be inefficient when it comes communication resources; programming trends in favor of “reflection” [12] and reflective applications, apart from providing fabulous functionalities, do jeopardize efficiency when it comes to processing resources.

In summary, these risks can be overcome by taking aggregated efficiency of a technology as a fundamental requirement: (1) watching out for low levels of entropy and (2) coordinating the correlation of existing technologies. The latter requirement can be easily tracked following the maxim *“something is not perfect when nothing else has to be added but when nothing has to be removed”* [2]. The former one can be

observed by ensuring that things not only do work but their use of resources is efficient.

### **4.3 The Good Access Point Paradox**

Balancing between functionality and efficiency is one of the reasons because of which efficiency may be relegated. This is necessary on designing and adopting new technologies. But the criteria used to justify the necessity of using more resources to implement functionalities and, consequently, reduce the efficiency must be consistent with themselves in the long term. Being consistent with themselves in the long term means that applying those criteria will not lead to the necessity to apply those criteria again in the future and consequently consume more resources. Applying the same criteria indefinitely would paradoxically lead to resource starvation and, at the very end, it would not solve the problem.

A popular ([8], [18] and more) example of these criteria is the so called “firewall friendliness”, here coined “the good access point”. When web services technology dawned, it was common to have firewalls protecting networks from the open Internet. It is common to have them now too, but we will see that they have had to become more complex. A great variety of pieces of information is transmitted through Internet and it has to be handled by different pieces of logic. This leads to multiplexing the information being transmitted, e.g. type of transport protocol, TCP port or type of application protocol. This multiplexing information is known in terms of the OSI reference architecture as the Service Access Point (SAP). Some of these types of information are deemed to be risky and firewalls use the multiplexing information as a criterion to prevent these risky data from being transmitted to and from the networks they protect. Habitually, security policies establish that the rules of these firewalls should reject every type of information but explicitly allowed SAP’s.

This context propitiated that, on developing overlaying communication architectures [3] that required transmitting data through a great number of networks and firewalls, several problems occurred because of the security policies of intermediary firewalls that tended to reject traffic at every access point but those “good access points” like HTTP, HTTPS, POP or SMTP. These problems were exhibited as criteria that favored the use of web services technology to pipe any type of application. Web services work generally with HTTP, a “good access point”. It is a firewall friendly technology, so it will not present the problems that customized approaches with weird TCP ports present. With web services the HTTP protocol is not only an Internet browsing protocol anymore, it enables piping anything: good or bad, as TCP or UDP do.

Applying this criterion is not consistent and does not solve the problem it is meant to address. According to these criteria every application wanting to go through networks should preferably be piped with web services technology. As a consequence a growing number of different applications have joined this technology: very few applications reject the opportunity to avoid firewalls. But, as the number of applications transferring information through web services and the HTTP port increases, so do increase the variety of pieces of information transmitted through them. Many times, the only criterion for choosing web services as a piping technology instead of custom TCP sockets was just that: firewall friendliness. Furthermore, security concerns start to appear as even malicious applications start to use web

services. Immediately, new firewall technologies are born to tackle this new security threats that the simple browsing HTTP protocol could not even have imagined years ago: the XML firewalls ([6], [9]). The unavoidable consequence is that some overlaying architectures based on web services begin to have the same problems motivated by intermediary firewalls that former architectures had. They begin to have those problems that made them choose web services as a piping technology. Should we now choose a “good XML access point” and pipe all the traffic through it? It would solve the problem, would not it? Or would it just delay it and consume multiplexing entropy inefficiently? Perhaps, the right approach to this would be to adequately do the hard job of modifying security policies to accept new overlying applications. At the same time, the real advantages of web services other than “firewall friendliness” should be deemed before choosing the technology of our application.

## 5 Conclusion

This analysis harmonizes the ideas of several Internet evolution models into a single one. These ideas have been shaped as a co-evolution of two evolving vectors: technology and attitude. Both are conditioned by universality, independency, economic feasibility and creativity. A principal conclusion of this model is that challenges of any of the vectors affect the evolution of the other one.

For each vector an outstanding concern has been argued. (1) Attitudes towards Internet are characterized by an unbalanced evolution between individuals and organizations. (2) Internet technology proves not to be always committed to efficiency principles. Since both concerns affect the economic feasibility of the future Internet, two guidelines are suggested to overcome them: a paradigm shift in organizational attitudes towards Internet and making efficiency a primary cross-cutting requirement of Internet technology development.

Finally, it should be remarked that it would be unbalanced and partial to hold that these two are the only challenges to take into account. On the contrary, they should be comprehensively addressed together with other challenges of the future Internet which may result from other analyses.

## Acknowledgements

Thanks to my colleagues Xabier Larrucea and Stefan Schuster for their fruitful suggestions and comments.

## References

1. Bughin, J., Manyika, J., Miller, A.: Building the Web 2.0 Enterprise: McKinsey Global Survey Results. Retrieved 22 October 2008 from: <http://www.mckinseyquarterly.com>

- /Information\_Technology/Management/Building\_the\_Web\_20\_Enterprise\_McKinsey\_Global\_Survey\_2174. The McKinsey Quarterly (July) (2008)
2. Callon, R.: The twelve networking truths. RFC 1925 (1996)
  3. Clark, D., Lehr, B., Bauer, S., Faratin, P., Sami, R., Wroclawski, J.: Overlay Networks and the Future of the Internet. In: *Communications And Strategies*, vol. 63, num. 3, pp. 109--129. ITS Communications And Strategies (2006)
  4. Coase, R.H.: The nature of the firm. In: *Economica*, vol. 4, num. 16, pp 386--405. Blackwell Synergy (1937)
  5. Consoli, D.: Co-Evolution Of Capabilities And Preferences In The Adoption Of New Technologies. In: *Technology Analysis and Strategic Management*, vol. 20, num. 4, pp. 409--425. Routledge (2008)
  6. Delessy-Gassant, N., Fernandez, E.B., Rajput, S., Larrondo-Petrie, M.M.: Patterns for application firewalls. In: *Proceedings of the 11th Conference on Pattern Languages of Programs* (2004)
  7. Dovrolis, C.: What Would Darwin Think About Clean-Slate Architectures? In: *Proceedings of the 2008 conference on Applications, technologies, architectures, and protocols for computer communications*, vol. 38, num. 1, pp. 29--34. ACM (2008)
  8. Fensel, D., Bussler, C.: The Web Service Modelling Framework WSMF. In: *Electronic Commerce Research and Applications*, vol. 1, num. 2, pp. 113--137. Elsevier (2002)
  9. Fernandez, E.B.: Two patterns for web services security. In: *Proceedings of 2004 International Symposium on Web Services and Applications*. Las Vegas, USA (2004)
  10. Gray, N.A.B.: Comparison of Web Services, Java-RMI, and CORBA service implementations. In: *Fifth Australasian Workshop on Software and System Architectures*, in conjunction with ASWEC 2004. Melbourne, Australia (2004)
  11. Gridconnection web site. Retrieved 27 October 2008 from <http://www.gridconnection.co.uk>
  12. Maes, P.: Concepts and Experiments in Computational Reflection. In: *ACM SIGPLAN Notices*, vol. 22., num. 12, pp. 147--155. ACM, New York, USA (1987)
  13. McAfee, A.P.: Enterprise 2.0: The Dawn of Emergent Collaboration. In: *IEEE Engineering Management review*, vol. 34, num. 3, pp. 38-47. IEEE Press (2006)
  14. Oliveira, R., Zhang, B., Zhang, L.: Observing the evolution of Internet AS Topology. In: *Proceedings of the 2007 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 313--324. ACM (2007)
  15. O'Reilly, T.: What is Web 2.0: Design Patterns and Business Models for the Next Generation of Software. Retrieved 26 October 2008 from <http://www.oreillynet.com/pub/a/oreilly/tim/news/2005/09/30/what-is-web-20.html>. O'Reilly (2005).
  16. Porter, M.E.: Strategy and the Internet. In: *Harvard Business Review*, vol. 79, num. 3, pp. 63--78 (2001)
  17. Ratnasamy, S., Shenker, S., McCanne, S.: Towards an Evolvable Internet Architecture. In: *Proceedings of the 2005 conference on Applications, technologies, architectures, and protocols for computer communications*, pp. 313--324. ACM, New York, USA (2005)
  18. Roy, J., Ramanujan, A.: Understanding web services. In: *IT Professional*, vol. 3, num. 6, pp. 69--73. IEEE Computer Society (2001)
  19. TOP 25 Web 2.0 Sites. Retrieved 27 October 2008 from <http://www.ebizmba.com/articles/user-generated-content>. eBizMBA (October) (2008)
  20. Zeeb, E., Bobek, A., Bohn, H., Pruter, S., Pohl, A., Krumm, H., Luck, I., Golatowski, F., Timmermann, D.: WS4D: SOA-Toolkits Making Embedded Systems Ready For Web Services. In: *3rd International Conference on Open Source Systems*, Limerick, Ireland (2007)

# An Economic Traffic Management Approach to Enable the TripleWin for Users, ISPs, and Overlay Providers

Tobias Hoßfeld<sup>1</sup>, David Hausheer<sup>2</sup>, Fabio Hecht<sup>2</sup>, Frank Lehrieder<sup>1</sup>, Simon Oechsner<sup>1</sup>, Ioanna Papafili<sup>3</sup>, Peter Racz<sup>2</sup>, Sergios Sourdos<sup>3</sup>, Dirk Staehle<sup>1</sup>, George D. Stamoulis<sup>3</sup>, Phuoc Tran-Gia<sup>1</sup>, Burkhard Stiller<sup>2</sup>

<sup>1</sup> University of Würzburg, Institute of Computer Science, Germany

<sup>2</sup> University of Zurich, Department of Informatics IFI, Switzerland

<sup>3</sup> Athens University of Economics and Business, Department of Informatics, Greece

**Abstract.** Socio-economic aspects play an increasingly important role in the Future Internet. To enable a TripleWin situation for the involved players, i.e. the end users, the ISPs and telecommunication operators, and the service providers, a new, incentive-based concept is proposed referred to as Economic Traffic Management (ETM). It aims at reducing costs within the network while improving the Quality-of-Experience (QoE) for end users. In particular, peer-to-peer (P2P) overlay applications generate a large amount of costs due to inter-domain traffic. ETM solution approaches have to take into account (a) the traffic patterns stemming from the overlay application, (b) the charging models for transit traffic, and (c) the applicability and efficiency of the proposed solution. The complex interaction between these three components and its consequences is demonstrated on selected examples. As a result it is shown that different ETM approaches have to be combined for an overall solution. To this end, the paper derives functional and non-functional requirements for designing ETM and provides a suitable architecture enabling the implementation of a TripleWin solution.

**Keywords.** Economic Traffic Management, inter-domain traffic, locality-awareness, P2P VoD, SmoothIT architecture

## 1. Introduction

One of the key applications in today's and probably also the Future Internet is P2P content distribution. This ranges from straightforward file-sharing to the more recent P2P-based video streaming, which is projected to rise even more in popularity. The large numbers of users these systems attract as well as the huge amount of data they efficiently distribute shows their scalability and is one reason for their success.

However, these same features create new difficulties for network and Internet Service Providers (ISPs). Overlay connections used by the P2P networks are up to now generally network agnostic and therefore wasteful with resources. Especially the liberate use of inter-domain connections, i.e., the transit links between ISPs, causes a high cost. Also, the high probability that a connection spans the networks of several providers complicates the provision of end-to-end traffic management needed for specific services. This is compounded by the fact that most P2P overlays do not use

one single connection to provide such a service, e.g., one video stream, but many at the same time. The number and quality of these flows changes dynamically with the overlay's topology and population. Supporting application diversity by service differentiation poses a challenge to the provider, since the user's experienced service quality (Quality of Experience QoE) depends on the individual flows' QoS. Currently, the IETF ALTO (Application-Layer Traffic Optimization) Working Group has been established which also identifies the above mentioned problems and focuses on improving P2P performance and lowering ISP costs.

To resolve this, several approaches exist. One is to simply stifle P2P traffic as a provider, with the aim to improve the performance of other services and to lower its cost. However, this strategy runs the danger of having a detrimental effect on customer satisfaction, and is therefore not an optimal choice.

A different solution to the problem described above is Economic Traffic Management (ETM), see [Fern08,O+08]. It operates under the assumption that all parties involved (ISP, user and, if applicable, overlay provider) will participate voluntarily in a management scheme they all profit from. As a consequence, this scheme has to provide incentives to these players to cooperate, so that in the end, a 'TripleWin', i.e., a win-win-win situation is created. In the following, we demonstrate the complex interaction between ETM solution approaches, the charging models used by ISPs, and overlay traffic characteristics and evaluate qualitatively their applicability. We discuss (1) the ETM mechanisms locality promotion and ISP-owned peers, (2) the 95<sup>th</sup> percentile rule as charging model, and (3) the applications BitTorrent-based file sharing and VoD. Based on this discussion, functional and non-functional requirements for the design of an overall ETM solution achieving the TripleWin situation is derived.

Finally, the SmoothIT overall architecture takes into account these requirements and allows the implementation of various ETM approaches. This is put into context of related work and existing other projects in the same area of interest.

## 2. ETM Solution Approaches

In the following, we discuss two different ETM solution approaches, locality promotion and ISP-owned peers (IoP). Locality promotion aims at reducing inter-domain traffic by fostering the exchange of data among the peers within one domain. Therefore, topology information is exchanged between underlay and overlay resolving the information asymmetry. The IoP approach aims at increasing QoE and reliability of a content distribution network by providing additional capacity supporting the dissemination of popular contents.

### 2.1. Locality Promotion / Locality Awareness

Locality promotion aims at keeping traffic in the same Autonomous System (AS), instead of having long connections spanning several networks. The goal is to lower costs that are created by two effects of these long-reaching connections.

The first is that connections that are long in terms of hops consume more transmission, routing and switching resources than shorter ones. The second is that inter-carrier traffic has to be paid, depending on the agreement between the two ISPs in question. Therefore, the promotion of overlay traffic locality is a key issue under an operator's perspective since it may reduce both network investments and transit costs.



To illustrate this, consider a P2P overlay network where a peer knows of several sources for a (partial) file download. Some of these are in the same AS as the local peer, some are in an AS the local peer's ISP has a peering agreement with, and some can be reached only via transit links to other AS. Normal overlay selection mechanisms do however not consider this information (also because they do not possess it), therefore the data transfer connections are largely random from the ISP's point of view. Depending on the number and location of the potential sources, probability is high that some of these connections use expensive transit links.

Here, one method of promoting locality is to provide an information interface for the overlay, such as the SmoothIT Information Service (SIS) described in Section 7. It can be queried by a peer to provide an informed ranking of the potential sources based on locality, as well as providing a better QoS to users of this service, in order to improve their experienced service quality.

## 2.2. ISP-owned Peers

This approach does, in comparison to the locality promotion, not necessitate cooperation with the overlay to work. It works by adding more resources to the overlay network in the form of ISP-controlled peers, so that the usage of these resources, such as storage and upload capacity, may be utilized to the ISPs advantage. In a variation of this scheme, also a common peer that shows certain characteristics may be granted some of these resources by the ISP (e.g., Highly Active Peers (HAP)). In any case, this peer offers the same functionality as any other peer. However, there may still be some minor modifications, such as increasing the number of upload slots of that peer. Still, this peer is not perceived as a special entity in the system by the other peers.

By participating in the normal data exchange of the overlay, the ISP-owned Peer (IoP) is able to download and cache content and provide it to other peers as a seed. Since it should provide a higher bandwidth due to its allocated resources, it should be a popular overlay neighbor, depending on the mechanism the overlay uses. The tit-for-tat peer selection algorithm, e.g., prefers good uploaders. Apart from the fact that this addition of resources to the overlay improves its performance, it also allows for influencing traffic. The IoP serves as a traffic attractor, and additionally has the possibility to use locality-aware mechanisms such as biased peer selection.

If the overlay wants to cooperate with an ISP using this strategy, the tracker in a BitTorrent-like architecture is in a perfect position to announce or at least propose the IoP as a regular overlay neighbor to the normal peers. Also, helpful information like the global popularity of content useful to cache, or the size and distribution of swarms could be provided to the ISP. Otherwise, the ISP would have to monitor such information by implementing additional modules for the IoP.

## 3. Charging Models

The Internet is characterized as an informal hierarchy of operators [MPD+] and structured in three tiers. The first level includes the *Tier 1 IAPs* (Internet Access Providers), which interconnect with each other and form the "backbone" of the Internet. In the second level, we have the *Tier 2 ISPs*, which usually have some network of their own, limited to a geographic region, and they partly rely on Tier 1 IAPs to obtain world-wide access, by purchasing some level of transit. As, no traffic symmetry is

expected between a Tier 1 IAP and a Tier 2 ISP, the Tier 2 provider pays to Tier 1 provider an amount specified by the details of the traffic rules in place. Depending on the status, the Tier 2 ISP could be charged for the volume of inbound traffic or for the difference between inbound and outbound traffic. Such a charging rule is referred to as *transit agreement*. At the same time, Tier 2 ISPs can have interconnection agreements with other Tier 2 ISPs, so-called *peering agreements*, without any charging due to symmetric traffic exchange. *Tier 3 ISPs* are purely re-sellers of Internet access services and rely solely to interconnection agreements with Tier 2 ISPs to gain Internet access.

Therefore, a Tier 2 ISP play an interesting role regarding TripleWin. One of a Tier 2 ISP's objectives is to reduce his interconnection costs which has to be achieved by a viable ETM solution. Throughout the paper, when we refer to an ISP we consider the case of a Tier 2 ISP.

### 3.1. The 95<sup>th</sup> Percentile Rule

One of the most prominent charging models for transit agreements is the 95<sup>th</sup> percentile rule. Let us consider a typical Tier 2 – Tier 1 transit agreement. The difference  $d_x$  between inbound and outbound data throughput of the Tier 2 ISP is measured for every  $\Delta t=5$  minutes time slice  $x$ . Then, the Tier 2 ISP is charged on the monthly 95<sup>th</sup> percentile of the differences per time slice, formally  $C_{month} = P95\{d_x\} \cdot Price/Mbps$ .

One of the main tasks for ETM related to the 95<sup>th</sup> percentile pricing scheme is to determine which parameters affect the generated costs, so that we can devise ETM mechanisms that decrease those costs for the ISP. One parameter that can be influenced and has direct effects on the costs is the amount of traffic flowing in both directions. One way to do is applying locality awareness, as already mentioned. However, depending on the application (file sharing or video-on-demand) locality may decrease the traffic either in both directions (e.g. due to tit-for-tat for file-sharing) or only in one direction (e.g. using give-to-get for video-on-demand). In addition, the different overlay applications will show different traffic patterns.

### 3.2. Estimation of Parameter Sensitivity

Figure 1 shows the actual effects of different traffic patterns, as well as the effect of the duration  $\Delta t$  of the observation window. On the x-axis,  $\Delta t$  (in seconds) is given while the 95<sup>th</sup> percentile value is depicted on the y-axis.




Figure 1. Impact of different distributions and duration of capturing window




Figure 2. Measurement of population sizes of BitTorrent swarms



We simulate different types of traffic pattern where packets arrive according to a Poisson process, while the packet sizes follow a binomial and a lognormal distribution respectively. The simulations were repeated 1000 times and the average 95th percentile out of the repetitions as well as the 95% confidence intervals are given. As a result, we see that the traffic pattern but also the length of the observation window play an important role in the calculation of the 95<sup>th</sup> percentile value.

#### 4. Overlay Traffic Characteristics

The traffic generated by an overlay depends on several factors, such as the application class it belongs to and the size of the overlay. In the following, we will shortly illustrate these issues.

##### 4.1. Differences between Overlay Application Classes

We compare the main mechanisms influencing the traffic behavior of two overlay classes, namely BitTorrent-like file-sharing and Video-on-Demand (VoD). They differ mainly in the peer and chunk selection processes used. For file-sharing, a Least-Shared First (LSF) strategy is used for chunk selection, while the peer selection is based on the tit-for-tat principle. In contrast, a VoD system has to also take into account the playout deadlines of chunks, and, due to the different playout positions of different peers, has to rely on a strategy like give-to-get for peer selection. Additionally, the peer selection here is much more important for the QoE of the end users, also because a video's quality is much more sensitive to QoS parameters, such as the minimum bandwidth achieved for a stream.

Another consequence of the different peer selection algorithms is that the different ETM approaches do not achieve the same effect, i.e., for a tit-for-tat strategy, traffic locality will lower both incoming and outgoing traffic. If the charging scheme works on the difference between the volumes of these two traffic streams, the effect on the ISP's cost may be minimal. Since give-to-get is not symmetric, the effect of the same approach in a VoD system can be expected to be much larger.

##### 4.2. Measurement of BitTorrent Swarm Sizes

In a BitTorrent-like system, each separate piece of content (e.g., movie, archive or document) has its own overlay, the so-called swarm, where it is distributed. To estimate the effectiveness of ETM measures in one such swarm, its size has to be known. Figure 2 shows results from a study of typical swarm sizes of BitTorrent swarms. The measurements were conducted in August 2009 and 63,867 BitTorrent swarms offering video contents were investigated. It shows that a Pareto-principle governs the total peer distribution: a large share of the peers can be found in a small number of swarms. This has several consequences for the ETM mechanisms employed in such a system. Depending on the swarm size an IoP or HAP is participating in, its optimal position and allotted resources vary. Also, more popular content can be cached by participating in larger swarms, while on the other hand the increase of resources by adding one or a small number of IoPs to such a swarm may be negligible.

For locality promotion, the impact of these results is different. For this mechanism, it is important that swarms are large in order to provide local alternatives for remote

neighbours. If a swarm is too small, only few peers are actually in the same network, prohibiting an overlay restructuring due to a lack of choices.

### 5. Exemplary ETM Solution for BitTorrent using Locality Promotion and IoPs

This section will discuss the mutual interdependency of ETM, charging scheme, and overlay characteristic at the example of BitTorrent file sharing. We have seen that there are many small BitTorrent and few very large BitTorrent swarms. As a consequence, BitTorrent swarms of all sizes contribute significantly to the overall BitTorrent traffic, and for an ETM to work efficiently it has to tackle swarms of – if possible – all sizes. Locality promotion potentially reduces the inter-domain traffic for large swarms. However, there is a critical number of peers per swarm required within an ISP's network in order to successfully use locality promotion for achieving a substantial reduction in inter-domain traffic without simultaneously decreasing the overlay's performance and thus violating the win-win-win maxim of ETM.

The IoP provides a solution for smaller swarm sizes. However, the number of swarms it can support simultaneously is limited mainly by the available storage. The achieved inter-domain traffic reduction per required storage capacity i.e. per CAPEX shrinks dramatically with the swarm size. A simulation study using ns-2 has been performed based on [E07]. In order to investigate the effects of locality promotion and IoP, simulations were performed on a 50 peer swarm subdivided 35 to 15 in two networks A and B. Locality promotion in B achieves a symmetric inter-domain traffic reduction of about 15%. Additionally inserting an IoP in network B leads to an asymmetric change reducing the ingress traffic of B by 45% while simultaneously increasing the ingress traffic of A by 55%.




Figure 3. Simulation results of difference between inbound and outbound traffic



Figure 4. Qualitative illustration of inter-play of Locality Promotion and IoP

Now, the charging model comes into play when evaluating the usefulness of the ETMs. Our exemplary charging scheme is sensitive to asymmetric changes only while symmetric changes have no direct impact on the costs. Figure 3 shows the difference of inbound and outbound traffic for network A. Locality promotion does not affect this difference while the IoP clearly shifts the traffic difference in the favor of A deploying the ETMs. When interpreting A as a Tier 3-ISP and B as a Tier 2-ISP the use of an IoP is beneficial for the Tier 3-ISP. Whether there is an actual monetary benefit, on the one hand depends on the OPEX and CAPEX for the ETM and on the other hand on the exact charging model and the achieved traffic reduction which again depends on the swarm sizes.

For an efficient operation of ETMs we propose a combined and purposeful usage of ETM mechanisms depending on the present overlay application and its characteristics. E.g., the IoP is good for small to medium swarms. Its efficiency can be improved by a more directed use which is enabled by a SIS like architecture that resolves the information asymmetry. Figure 4 gives a qualitative impression on the regions of operation for the different ETM mechanisms. A qualitative statement on the efficiency of IoP and locality interplay is still subject to further studies.

## 6. Requirements for the ETM System Design

The requirements for the ETM system design are derived from the overall goals, the scope of solution approaches, and the overlay traffic characteristics as discussed in the selected examples above. Table 1 outlines main functional and non-functional requirements identified for the SmoothIT architecture.

Table 1. Functional and non-functional requirements for the ETM system design

<i>ID</i>	<i>Functional Requirement</i>	<i>ID</i>	<i>Non-functional Requirement</i>
R.1.	Improving P2P application performance while reducing the network traffic	R.10.	Easy deployment
R.2.	Incentive-compatibility for all player involved	R.11.	Extensibility for new overlay applications
R.3.	Support of different overlay applications	R.12.	Scalability in terms of large end-user population.
R.4.	Interface supporting various optimization schemes	R.13.	Efficiency of the SIS operation
R.5.	QoS support	R.14.	Robustness of the SIS against malicious behavior and against dynamic behavior
R.6.	Different operations: user anonymity mode (free services), user aware mode (premium services)	R.15.	Security: Secure communication between SIS entities and between SIS and overlay application
R.7.	Inter-domain support	R.16.	Standard compliance: The SIS shall use and based on standard protocols where applicable.
R.8.	OAM (Operation and Management) support	R.17.	Transparency: The SIS shall not apply Deep Packet Inspection (DPI).
R.9.	Mobile network support		

A key functional requirement is to manage overlay traffic and to improve P2P application performance in a way that results in a TripleWin situation for all involved entities. Additionally, the architecture shall support different overlay applications in an integrated manner, while providing various optimization schemes. The architecture shall also allow inter-domain interactions between ISPs. The integration of QoS and OAM (Operation and Management) support are also highly relevant in an operational environment. Finally, the architecture shall provide special optimization schemes for mobile networks.

In terms of non-functional requirements, to be able to deploy the SmoothIT architecture in a real world environment, an easy deployment is essential. The

architecture and protocols used shall be extensible in order to be able to integrate new overlay applications, optimization schemes, and metrics. Furthermore, scalability, efficiency, and robustness are essential for a large-scale operational environment. Additionally, security mechanisms shall be integrated into the architecture and data privacy and regulations shall be considered. Finally, the architecture shall use existing standard protocols, where applicable.

## 7. SmoothIT Architecture Design

The SmoothIT architecture provides a flexible service that can be used with a variety of ETM schemes. The service – termed SmoothIT Information Service (SIS) – is an instantiation of the overall architecture developed. It includes interfaces for the bidirectional communication between ISPs and overlay applications and among different ISPs. The SIS is able to provide information about policy, locality, congestion, charging, and QoS, in order to help overlay applications decide how to efficiently build and maintain the overlay network. The SIS has to be provided in a way that prevents the exploitation of the SIS information by malicious users/applications and enables the ISP to hide sensitive network status information.

Figure 5 shows the functional components of the SmoothIT architecture. The SIS server provides the SIS service to the SIS Client and it is the core component of the architecture. It is responsible for providing support to the overlay formation process, aiming at optimizing both overlay and ISP interests. The SIS server is deployed by each ISP and provides the SIS service within the network of the ISP. Additionally, a SIS server can request preference information from other SIS servers deployed in other network domains. This server-to-server interaction is used to provide refined preference information for inter-domain connections.

The Admin component enables the ISP to configure SIS internal parameters over the administrative interface, such as the QoS classes of services, the maximum number of flows that can be prioritized, or how locality must be enforced. This is configured using a graphical management interface, which the network administrator can use to review the current configuration parameters and modify them as necessary.

The Metering component is responsible for performing network measurements. It can be composed of one or more modules that can perform different measurements. The results are combined by the SIS server, according to settings configured by the Admin component. The BGP Information Module, for example, retrieves locality information from the BGP protocol. It takes into account the local preference, the AS hop count, and the MED (Multi Exit Discriminator) BGP attributes.

The Security component provides access control to the SIS service, i.e. authentication and authorization for the SIS Client, the Admin, and other SIS servers. Additionally, data confidentiality and integrity are ensured by secure communication channels between components.

The Configuration database (Config DB) stores ISP policies and is responsible for any information that an ISP can configure for the SIS architecture. The database, which may be based on an existing repository of the ISP, contains information about different types of business relations the ISP has with other ISPs and performance related metrics corresponding to each network.

The QoS component checks the availability of network resources and guarantees resources requested by overlay applications, as well as enforces QoS policies in the

network. The QoS component is important to enable the provisioning of premium services to end-users and overlay providers.




Figure 5. SmoothIT architecture

## 8. Related Work and Projects

Based on the review of related work and projects, we provide a concise comparison to the ETM and TripleWin approach described above and highlight the key differences.

The biased neighbor selection [BCC+06] relies on the popular file sharing tool BitTorrent [BT]. It adjusts the original protocol [C08], where a central tracker server provides new peers with a random set of other peers, so that new peers preferentially receive addresses of peers located in the same network domain. To this end, the tracker needs information about the network domain of all peers. The evaluation shows that biased neighbor selection reduces download times of files and inter-domain traffic.

Aggarwal et al. [ASF07] propose to use an oracle service to achieve the aforementioned goals. Peers send a list of potential neighbors to the Oracle Service which orders it according to the preferences of the network provider. Different metrics for the ordering are possible, e.g., the number of autonomous system (AS, [HB96]) hops on the path from the peer to its potential neighbor. Consequently, it is more application-independent than biased neighbor selection and allows a more fine-grained differentiation of possible connections. Simulations have shown that the oracle service significantly reduces the number of inter-domain connections and the amount of inter-domain traffic. A more detailed evaluation is given in [AAF08].

The network topology information desk service (NTIDS) [B03] is very similar to the oracle service. The network provider offers an information desk service which can be queried by peers. The service sorts a list of IP addresses of potential neighbors for a specific peer according to the network providers' preferences. For this purpose, it uses the routing information to determine the proximity of the given addresses.

The architecture of the P4P project [XKS+07, XYK+08], "Proactive Network Provider Participation for P2P", also provides an entity which maintains information about the network topology and location of peers. This entity is called iTracker and controlled by the network provider and be queried by a tracker or the peers depending

if the system is tracker-based or tracker-less. The iTracker assigns all peers to locality clusters and keeps cluster-to-cluster preferences reflecting the preferences of the network provider, e.g., cost or the traffic policies. Simulations show that using an iTracker reduces inter-domain traffic and improves application performance.

This paper's ETM approach shows similarities with these approaches discussed, like an interaction between overlay applications and the underlying network in the form of a generic information service, the provisioning of locality information, and the reduction of inter-domain traffic. However, there are several key differences: (1) The ETM approach considers not only locality information, but additional approaches, such as applying differentiated pricing or providing QoS-enabled services. (2) Related projects consider mainly pure file-sharing. In contrast, TripleWin aims at developing a solution applicable for different overlay application types, like file sharing and video streaming, and, therefore, will consider different application characteristics. (3) Related projects rely on a cooperation of all involved players. But this may not be possible, if not all players can benefit from the solution proposed. Therefore, TripleWin focuses on incentives to achieve collaboration among players. This also involves investigating inter-domain interactions in closer detail. Finally, ETM within the SmoothIT project follows an overall picture, where the relationship between charging and incentive models, overlay application traffic characteristics, and ETM solutions is essential.

## 9. Conclusions

In this work we have introduced the concept of economic traffic management (ETM). In particular, we have shown that a successful application of ETM depends not only on the solution approaches, but also on the charging model and the overlay application. The complex interaction between these three components and its consequences is demonstrated on selected examples. As ETM solution approaches we consider locality promotion to reduce inter-domain traffic and ISP-owned peers to enhance the available resources within the network. The 95<sup>th</sup> percentile rule charging scheme is presented as a prominent example for transit agreements. The impact of different overlay applications are elaborated for BitTorrent-based file sharing and video-on-demand. We conclude that the spectrum of applicability of a single ETM mechanism is rather limited. Only the intelligent and coordinated usage of various ETM mechanisms can successfully exploit the optimization potential currently present in uncoordinated underlay-agnostic overlays. The paper derives functional and non-functional requirements for designing ETM and provides a suitable architecture, as well as a protocol proposal enabling the implementation of such a TripleWin solution. Finally, differences to existing projects and related work are outlined.

Our current research studies focus on investigating the mutual interdependency of ETM, charging schemes, and overlay applications, not only qualitatively, but also quantitatively. In detail, we have to investigate the overall costs when keeping traffic locally; on one hand we save inter-domain traffic and costs, on the other hand we may increase the number of "internal" hops and therefore costs or we might introduce additional components, like IoPs or SIS which consume CAPEX and OPEX costs. Furthermore, we have to keep in mind the diversity of transit agreements and their impact on the ISP's costs for overlay traffic. Depending on the results, we may propose pricing rules for Tier 1, Tier 2 or Tier 3 providers, specifically designed for achieving a TripleWin situation when applying ETM. Methods from game theory seems to be



appropriate for identifying and validating them. The proposed architecture will be implemented and the usefulness as well as the feasibility of the different ETM solutions will be demonstrated in experimental facilities.

### Acknowledgements

This work has been performed partially in the framework of the EU ICT Project SmoothIT (FP7-2007-ICT-216259). Furthermore, the authors would like to thank all SmoothIT project members for providing insights and their discussions.

### References

- [AAF08] V. Aggarwal, O. Akonjang, A. Feldmann. *Improving User and ISP Experience through ISP-aided P2P Locality*. 11th IEEE GI'08, 2008.
- [AFS07] V. Aggarwal, A. Feldmann, C. Scheideler. *Can ISPs and P2P systems co-operate for improved performance?* ACM SIGCOMM Computer Communications Review (CCR), 37(3):29-40, 2007.
- [BCC+06] R. Bindal, P. Cao, W. Chan, J. Medval, G. Suwala, T. Bates, A. Zhang. *Improving Traffic Locality in BitTorrent via Biased Neighbor Selection*. IEEE ICDCS'06, 2006.
- [BT] BitTorrent. <http://www.bittorrent.com>. Accessed in October 2008
- [B03] S. Brecht: *Investigation of application layer routing*. Master's thesis, University of Federal Armed Forces Munich, Germany, Supervisors: P. Racz, B. Stiller, 2003.
- [C08] B. Cohen, *The BitTorrent protocol specification*, 2008.  
[http://www.bittorrent.org/beps/bep\\_0003.html](http://www.bittorrent.org/beps/bep_0003.html), accessed October 2008.
- [E07] K. Eger, Simulation of BitTorrent Peer-to-Peer (P2P) Networks in ns-2. <http://www.tu-harburg.de/et6/research/bittorrentsim/>, 2007.
- [Fern08] J.P. Fernandez-Palacios Gimenez et. al, A New Approach for Managing Traffic of Overlay Applications of the SmoothIT Project. AIMS 2008
- [HB96] J. Hawkinson and T. Bates: *Guidelines for creation, selection, and registration of an Autonomous System (AS)*. IETF RFC 1930, March 1996.
- [MPD+] B. Mitchell, P. Paterson, M. Dodd, P. Reynolds, P. Waters, R. Nich, *Economic study on IP interworking*. White paper, 2007.
- [O+08] S. Oechsner et. al, *A framework of economic traffic management employing self-organization overlay mechanisms*. IWSOS'08, 2008.
- [XKS+07] H. Xie, A. Krishnamurthy, A. Silberschatz, Y. R. Yang: *P4P: Explicit Communications for Cooperative Control Between P2P and Network Providers*. [http://www.dcia.info/documents/P4P\\_Overview.pdf](http://www.dcia.info/documents/P4P_Overview.pdf), 2007.
- [XYK+08] H. Xie, Y. R. Yang, A. Krishnamurthy, Y. Liu, and A. Silberschatz. *P4P: Provider Portal for Applications*. ACM SIGCOMM, 2008.

## A Security Architecture for Web 2.0 Applications

Lieven Desmet<sup>1</sup>, Wouter Joosen<sup>1</sup>, Fabio Massacci<sup>2</sup>, Katsiaryna Naliuka<sup>2</sup>, Pieter Philippaerts<sup>1</sup>, Frank Piessens<sup>1</sup>, Ida Siahaan<sup>2</sup>, Dries Vanoverberghe<sup>1</sup>

<sup>1</sup> DistriNet Research Group, Department of Computer Science  
Katholieke Universiteit Leuven, Celestijnlaan 200A, B-3001 Leuven, Belgium

<sup>2</sup> Department of Information and Communication Technology  
Universit di Trento, Via Sommarive 14, I-38050 Povo (Trento), Italy

**Abstract.** The problem of supporting the secure execution of potentially malicious third-party applications has received a considerable amount of attention in the past decade. In this paper we describe a security architecture for Web 2.0 applications that supports the flexible integration of a variety of advanced technologies for such secure execution of applications, including run-time monitoring, static verification and proof-carrying code. The architecture also supports the execution of legacy applications that have not been developed to take advantage of our architecture, though it can provide better performance and additional services for applications that are architecture-aware. A prototype of the proposed architecture has been built that offers substantial security benefits compared to standard (state-of-practice) security architectures, even for legacy applications.

### 1 Introduction

The business model of traditional web-services architectures is based on a very simple assumption: the good guys develop their .NET or Java application, expose it on the web (normally to make money as an application service provider), and then spend the rest of their life letting other good guys use it while stopping bad guys from misusing it.

The business trend of outsourcing the non-core part of business processes [8] or the construction of virtual organizations [10] have slightly complicated this initially simple picture. With virtual organizations for fulfilling the same high level goal

1. different service components are dynamically chosen (possibly using different data) and
2. different partners are chosen (possibly with different service level agreements or trust levels).

This requires different security models, policies, infrastructures and trust establishment mechanisms (see e.g. [13, 12]). A large part of the WS security standards (WS-Federation, WS-Trust, WS-Security) are geared to solve some of these problems.

Still, the assumption is the same: *the application developer and the platform owner are on the same side of trust border*. Traditional books on secure coding [11] or the .NET security handbook [14] are permeated with this assumption. However, in the landscape of Web 2.0, users are downloading a multitude of communicating applications ranging from P2P clients to desktop search engines, each plowing through the



user's platform, and communicating with the rest of the world. Most of these applications have been developed by people that the user never knew existed, before installing the application.

The (once) enthusiastic customers of UK Channel 4 on demand services 4oD might tell how use of the third-party services affects the customer's experience [18]. Downloading a client that allows you to watch movies at a very cheap rate seems like an excellent deal. Only in the fine print of the legal terms of use (well hidden from the download now, nowhere in the FAQs and after a long scrolling down of legalese) you find something you would like to know:

If you download Content to your computer, during the Licence Period, we may upload this from your computer (using part of your upstream bandwidth) for the purpose of transferring Content to other users of the Service. Please contact your Internet Service Provider ("ISP") if you have any queries on this.

As one of the unfortunate users of the system noticed, there is no need of contacting your ISP. He will contact you pretty soon.

To deal with the untrusted code, either .NET [14] or Java [15] exploit the mechanism of permissions. Permissions are assigned to enable execution of potentially dangerous or costly functionality, such as starting various types of connections. The drawback of permissions is that after assigning a permission the user has very limited control over how the permission is used. The consequence is that either applications are sandboxed (and thus can do almost nothing), or the user decides that they are trusted and once the permission is received then they can do almost everything.

The mechanism of signed assemblies from trusted third parties might solve the problem. Unfortunately a signature means that the signatory vouches for the software, but there is no clear definition of what guarantees it offers. The 4oD example or the incidents in the mobile phone domain [21] show that this security model is inappropriate.

In this paper, we describe the architecture of the runtime environment on the Web 2.0 platform that we have already developed for .NET (both the desktop and the compact edition). The architecture integrates in a very flexible way several state-of-the-art policy enforcement technologies, such as proof-carrying code and inlined reference monitors. In addition, the security architecture offers additional support for application contracts and the security-by-contract paradigm. Thanks to the combination of different enforcement techniques and the support for application contracts, our security architecture is able to provide policy enforcement for legacy applications, as well as architecture-aware applications. However, the latter set of applications have a smaller runtime performance penalty, which is an important characteristic for resource-restricted environments such as mobile Web 2.0 platforms. In addition, a first prototype implementation of the proposed security architecture is available for Windows based desktops, and for Windows mobile platforms with the .NET Compact framework (so it is also suitable for mobile devices).

The remainder of the paper is structured as follows. Section 2 provides some background information on the security-by-contract paradigm, existing policy enforcement techniques, and policy languages. Next, our flexible security architecture for Web 2.0 platforms is presented in Section 3, and Section 4 describes our prototype implementation. In Section 5, the advantages and disadvantages of the presented architecture are discussed. Finally, the presented work is related to existing research, and we offer a conclusion.

## 2 Background

The architecture described in this paper is largely based on the research results of the European project S3MS [23]. In this section, we describe the key notion of *security-by-contract* underlying the S3MS project, and we briefly discuss the policy enforcement techniques and policy languages considered in that project.

A key ingredient in the S3MS approach is the notion of “security-by-contract”. Web 2.0 applications can possibly come with a *security contract* that specifies their security-relevant behavior [4]. Technically, a contract is a security automaton in the sense of Schneider and Erlingsson [6], and it specifies an upper bound on the security-relevant behavior of the application: the sequences of security-relevant events that an application can generate are all in the language accepted by the security automaton. Web 2.0 platforms are equipped with a *security policy*, a security automaton that specifies the behavior that is considered acceptable by the platform owner. The key task of the S3MS environment is to ensure that all applications will comply with the platform security policy. To achieve this, the system can make use of the contract associated with the application (if it has one), and of a variety of policy enforcement technologies.

### 2.1 Policy Enforcement Techniques

The research community has developed a variety of countermeasures to address the threat of untrusted code. These countermeasures are typically based on runtime monitoring [6], static analysis [19], or a combination of both [28]. We briefly review here the technologies supported in the S3MS system. It must be noted, however, that the system so new technologies can be implemented as needed.

**Cryptographic signatures.** The simplest way to solve the lack of trust is to use *cryptographic signatures*. The application is signed, and is distributed along with this signature. After receiving this application, the signature can be used to verify the source and integrity of the application. Traditionally, when a third party signs an application, it means that this third party certifies the application is well-behaved. Adding the notion of a contract, as is done in the S3MS approach, allows us to add more meaning to claims on well-behavior: the signature means that the application respects the supplied contract.

**Inline reference monitoring.** With *inline reference monitoring* [6], a program is rewritten so that security checks are inserted inside an untrusted application. When the application is executed, these checks monitor the behavior of the application and prevent it from violating the policy. It is an easy way to secure an application when it has not been developed with a security policy in mind or when all other techniques have failed. The biggest challenge for inline reference monitors is to make sure that the application can not circumvent the inlined security checks.

**Proof-carrying code.** An alternative way to enforce a security policy is to statically verify that an application does not violate this policy. As producing the proof is normally too complicated to be done on the user side, the *proof carrying code* concept [19], allows the verification to be done off-line by the developer, or by an expert in the field. Then the application is distributed together with the proof. Before allowing the execution of the application, a proof-checker verifies that the proof is correct for the

application. Because proof-checking is usually much easier than making the proof, this step becomes feasible on Web 2.0 platforms.

**Contract-policy matching.** Finally, when application contracts (called application models in [25]) are available, *contract-policy matching* [6, 17, 25] is an approach to decide whether or not the contract is acceptable. At the deployment of the application the contract acts as an intermediary between the application and the policy of the platform. First, a matching step checks whether all security-relevant behavior allowed by the contract is also allowed by the policy. If this is the case, any of the other enforcement techniques can be used to make sure that the application complies to the contract (as the contract of the application is known in advance this step can be made off-line).

## 2.2 Policy Languages

In this paper, we make a clear distinction between application contracts and platform policies. Both are security automata, but the first ones are associated with a particular application, while the latter ones are associated with a platform.

A security automaton [24] is a Büchi automaton – the extension of the notion of finite state automaton to infinite inputs. A security automaton specifies the set of acceptable sequences of *security relevant events* as the language accepted by the automaton.

In the S3MS system, a policy identifies a subset of the methods of the platform API as security relevant methods. Typical examples are the methods to open network connections or to read files. Security relevant events in the S3MS system are the invocations of these methods by the untrusted application, as well as the returns from such invocations. Hence, a security automaton specifies the acceptable sequences of method invocations and returns on security relevant methods from the platform API.

Security automata have to be specified by means of a policy language. The S3MS system is designed to support multiple policy languages, including policy languages that support multiple runs of the same application. The actual prototype implementation supports already two languages: automata-based specification language ConSpec [2] and logic-based language 2D-LTL [16].

## 3 System Architecture

In this section, our service-oriented security architecture for Web 2.0 platforms is presented. First, we enumerate the most important architectural requirements for the security architecture. Next, subsection 3.1 gives an overview of our architecture, and highlights three important architectural scenario's. The following three subsections discuss some architectural decisions in more detail.

Before presenting and discussing our flexible service-oriented security architecture, the most important architectural requirements are briefly discussed.

**Secure execution of third-party applications** The architecture should give high assurance that applications that have been processed by the security system can never break the platform security policy. This is the key functional requirement of our architecture.

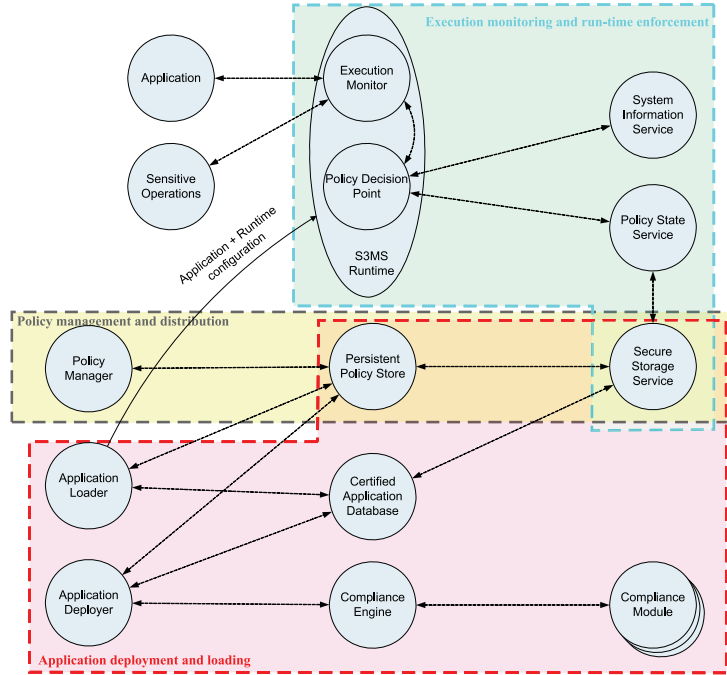


Fig. 1. Detailed architecture overview

**Flexible integration of enforcement techniques** The security architecture should integrate seamlessly the set of enforcement techniques discussed in Sec. 2. In addition, the security architecture should provide a flexible framework for adding, configuring or removing additional enforcement techniques.

**Optimized for resource-restricted platforms** The security architecture needs to be optimized for use on resource-restricted Web 2.0 platforms such as personal digital assistants or SmartPhones. These platform typically have limited memory and processing power, and restricted battery capacity. The architecture should secure the execution of applications with a minimal performance penalty during the application execution, without compromising security during network disconnectivity.

**Compatible with legacy applications** To be compatible with existing applications, it is important that the security architecture supports the secure execution of legacy applications that are unaware of the architecture. Of course, the fact that an application is architecture-unaware may impact performance.

In the following section, an overview of our security architecture for Web 2.0 platforms is presented. As will be explained further, each of the enumerated architectural requirements has impacted the overall architecture.

### 3.1 Overview

The security architecture is built upon the notion of “security-by-contract”. Web 2.0 platforms can select a security policy from a list of available policies, specifying an upper bound on the security-relevant behavior of applications. In addition, applications can be distributed with a security contract, specifying their security-relevant behavior.

The three main scenarios are: policy management and distribution, application deployment and loading, and execution monitoring and runtime enforcement.

**Policy management and distribution** This scenario is responsible for the management of different platform policies, and their distribution to Web 2.0 platforms.

**Application deployment and loading** This scenario is responsible for verifying the compliance of a particular application with the platform policy before this application is executed.

**Execution monitoring and runtime enforcement** This scenario is responsible for enforcing the adherence of a running application to the policy of the platform in the case where the previous scenario has decided that this is necessary.

The three scenarios operate on two different platforms: on the platform of the policy provider and on the Web 2.0 platform.

**Policy provider.** Within the S3MS security architecture, the policies are managed by the *Policy Provider* and a specific policy can be pushed to a particular Web 2.0 platform. The policy provider could for instance be a company that supplies its employees with Web 2.0-capable devices, but wishes to enforce a uniform policy on all these devices. It could also be an advanced end-user that wants to manage the policy using a PC that can connect to his Web 2.0 platform.


**Web 2.0 platform.** The Web 2.0 platform stores the policy and is responsible for deploying, loading and running applications. If necessary, it also applies execution monitoring and run-time enforcement to ensure compliance to the device policy.

Figure 1 shows an architectural overview of the entire platform, and of the software entities that are involved in these three scenarios.

### 3.2 Policy Representations

Our architectural requirements ask for flexibility in policy enforcement techniques used, as well as for resource conscious implementations. However, the different enforcement techniques impose different constraints on optimized policy representations, and hence it is hard, or even impossible, to find one optimal representation that is suitable for each technique.

For instance, in a runtime monitor, the decision whether an event is allowed or aborted relies only on the current state of the policy. Therefore, an efficient representation for runtime enforcement only contains the current state, and methods for each event that check against the state, and update it. On the other hand, contract/policy matching checks whether or not the behavior allowed by the contract is a subset of the behavior allowed by the policy. For this task, a full graph representation may be required.



**Fig. 2.** Compilation of a policy package

To deal with this problem, our architecture introduces the notion of *policy packages*. A policy package is a set of optimized representations (each geared towards a different enforcement technique) of the same policy.

The policy writer is responsible for distributing the policy packages to the policy broker service. To do so, he uses a policy compiler to transform a given policy specification (e.g. in ConSpec or 2D-LTL) to a policy package (figure 2). This policy package is then sent to the policy broker. In a similar vein, representation of application contracts can be optimized towards different application-policy matching algorithms, and hence contracts are supplied in our architecture in the form of a similar contract package.

### 3.3 The Deployment Framework: Support for a Variety of Policy Enforcement Techniques.

In order to achieve a powerful security architecture that can incorporate a variety of state-of-the-art policy enforcement techniques, our architecture includes a very configurable and extensible compliance engine. At deployment time of an application, this compliance engine attempts to ensure the compliance of the application by means of all supported enforcement technologies. Each enforcement technology is represented as a *compliance module*.

Each compliance module selects the policy representation it will use and sends it, together with the application and optional metadata, to a certification technology service.

Each compliance verification technology is encapsulated in a *ComplianceModule*. To verify the compliance of an application with a policy, the *Process(PolicyPackage policy, Application app)* method is executed on such a compliance module. The module selects the policy representation it will use and sends it, together with the application and optional metadata, to a certification technology service. The result of the method indicates whether or not the compliance verification is successful. As a side effect of executing the process method, the application can be altered (e.g. instrumented with an inline reference monitor). The compliance engine instantiates the different compliance modules and applies them sequentially until the compliance of the application with the policy is ensured.

The order in which the compliance modules are applied, and their particular configuration is made policy-specific. This policy-specific configuration is part of the policy

package. In this way, policies can optimize the certification process by favoring or excluding compliance modules (e.g. because they are optimal for the given policy, or because they are inappropriate).

## 4 Prototype Implementation

We have implemented a first prototype of this security architecture in the full .NET Framework, and in the .NET Compact Framework on Windows Mobile 5. We've chosen to do an implementation on a mobile device, because handheld devices are becoming ever more popular and have already made an entrance to the Web 2.0 world (sometimes called *the Mobile Web 2.0*). In addition, using mobile devices forces us to consider resource-restricted platforms.

Our prototype includes policy compilers for both ConSpec, a policy specification language based on security automata, as well as 2D-LTL, a bi-dimensional temporal logic language. The compiler outputs a policy representation package that includes three policy representations, one suitable for inlining, one suitable for signature verification, and one suitable for matching. The corresponding compliance modules that use these representations are also implemented. Compliance modules need not be aware of the source policy language. For instance, our runtime execution monitor uses the same Policy Decision Point interface irrespectively of the policy specification language used.

Our current inliner implementation uses caller side inlining. Because caller side inlining needs to find the target of a method call statically, it is harder to ensure complete mediation. Under certain assumptions about the platform libraries, and with certain restrictions on the applications being monitored, our inlining algorithms has been proven correct [27]. The main assumption about the platform library is that it will not create delegates (function pointers) to security relevant methods and return these to the untrusted application. This seems to be a realistic assumption for the .NET libraries. The main restriction we impose on applications is that we forbid the use of reflection, a common restriction in code access security systems.

A final noteworthy implementation aspect of our mobile prototype is the way we ensure that only compliant applications can be executed on the mobile device, i.e. only after the application successfully passes the deployment scenario. Instead of maintaining and enforcing a Certified Application Database, we decided to rely on the underlying security model of Windows Mobile 5.0 in our prototype. The *Locked or Third-Party-Signed* configuration in Windows Mobile allows a mobile device to be locked so that only applications signed with a trusted certificate can run [1]. By adding a policy-specific certificate to the trusted key store, and by signing applications with that certificate after successfully passing the deployment scenario, we ensure that non-compliant applications will never be executed on the protected mobile device.

## 5 Discussion

In this section, we offer a brief preliminary evaluation of the presented security architecture based on the architectural requirements set forth in Section 3.



**Secure execution of third-party applications** Our architecture assumes that the individual compliance modules and certification technology services are secure: a buggy or unreliable implementation can validate an application that does not comply with the platform policy. This is a weakness, but the cost of building in redundancy (e.g. requiring two independent compliance modules to validate an application) is too high. Apart from this weakness, our architecture supports high assurance of security through a simple and well-defined compliance validation process, and through the precise definitions of the guarantees offered by security checks. An example is the treatment of cryptographic signatures. Our security architecture relies on cryptographic signatures in several places. But a key difference with the use of cryptographic signatures in the current .NET and Java security architectures is the fact that the semantics of a signature in our system are always clearly and unambiguously defined. A signature on an application with a contract means that the trusted third party attests to the fact that the application complies with the contract, and this is a formally defined statement.

**Flexible integration of enforcement techniques** The architecture incorporates different techniques such as cryptographic signatures, contract/policy matching and inlining of a reference monitor. Thanks to the pluggable compliance modules and the concept of policy packages the framework can easily be extended with additional enforcement technologies. In addition, the policy configuration allows for policy-specific configuration of the different compliance modules, including configuration of the order in which they are applied to applications.

**Optimized for resource-restricted platforms** Proving that an application will never violate a give system policy is typically a relatively hard problem. This might become problematic because of the resource-restricted nature of Web 2.0 platforms. Thanks to the service-oriented architecture, most of this work can be outsourced from the platform to the contract certifier service. These high-performance certifier servers can process multiple requests simultaneous and can sometimes, depending on the policy enforcement technique that's being used, use caching mechanisms to speed up the certification of an application by using results obtained from the certification of a previous application.

**Compatible with legacy applications** Because of the use of a general-applicable fallback compliance module (e.g. inlining of a reference monitor), the architecture can also ensure security for architecture-unaware legacy applications.

Based on this preliminary evaluation, the presented architecture looks promising. However, a more in-depth architectural evaluation and validation is necessary for a more grounded conclusion. We see three important tracks for further evaluation of the presented architecture.

First, *an extensive architectural evaluation* of the proposed architecture is necessary. For instance, an architectural trade-off analysis (such as ATAM [22]) with the different stakeholders involved (such as end users, telecom operators, Web 2.0 application developers and vendors, platform vendors and security experts), can evaluate and refine several architectural trade-offs. Second, it is necessary to perform an *end-to-end threat analysis* of the proposed architecture. Based on these results, a risk assessment will identify the most important security risks and will provide additional input for the



architectural trade-off analysis. Third, a *further integration of existing enforcement techniques* in the prototype architecture is needed to validate the flexibility of the framework design.

## 6 Related Work

There is a huge body of related work that deals with specific policy enforcement technologies for untrusted applications. This research area is too broad to discuss here. Some of the key technologies were briefly discussed in Section 2. A more complete survey of relevant technologies can be found in one of the deliverables of the S3MS project [26].

Even more closely related are those research projects that have designed and implemented working systems building on one or more of the technologies discussed above. Naccio [7] and PoET/PSlang [5] were pioneering implementations of runtime monitors. Polymer [3] is also based mainly on runtime monitoring, but the policy that is enforced can depend on the signatures that are present on the code. Model-carrying code (MCC) [25] is an enforcement technique that is very related to the contract matching based enforcement used in the S3MS project. In MCC, an application comes with a *model* of its security relevant behavior, and hence models are basically the same as contracts. The MCC paper describes a system design where models are extracted from the application by the code producer. The code consumer uses the model to select a matching policy, and enforces the model at runtime. Mobile [9] is an extension to the .NET Common Intermediate Language that supports certified inline reference monitoring. Certifying compilers [20] use similar techniques like proof carrying code, but they include type system information instead of proofs.

We should also mention here the existing framework for enforcing information flow control in the Web 2.0 setting DIFSA-J [29]. This approach allows rewriting the third-party programs at the bytecode level to insert the inline reference monitor that controls access to the sensitive information in the databases and prevents leakage to undesirable parties.

## 7 Conclusion

We proposed a flexible security architecture for Web 2.0 platforms built upon the notion of “security-by-contract”. In a very extensible way, the architecture integrates a variety of state-of-the-art technologies for secure execution of Web 2.0 applications, and supports different policy specification languages. In addition, the proposed architecture also supports the secure execution of legacy applications, although a better runtime performance is achieved for security-by-contract-aware applications.

## 8 Acknowledgments

This research is partially funded by the Interuniversity Attraction Poles Programme Belgian State, Belgian Science Policy, and by the European Union under the FP6 Programme.

## References

1. Windows Mobile 5.0 application security. Available at <http://msdn2.microsoft.com/en-us/library/ms839681.aspx>, 2005.
2. I. Aktug and K. Naliuka. Conspec – a formal language for policy specification. In *Proc. of the 1st Int. Workshop on Run Time Enforcement for Mobile and Distributed Systems (REM2007)*, 2007.
3. Lujo Bauer, Jay Ligatti, and David Walker. Composing security policies with Polymer. In *Proc. of the ACM SIGPLAN 2005 Conf. on Prog. Lang. Design and Implementation*, pages 305–314, 2005.
4. N. Dragoni, F. Massacci, K. Naliuka, and I. Siahaan. Security-by-contract: Toward a semantics for digital signatures on mobil code. 2007.
5. Ú Erlingsson. *The Inlined Reference Monitor Approach to Security Policy Enforcement*. PhD thesis, Cornell University, 2004.
6. U. Erlingsson and F. B. Schneider. IRM Enforcement of Java Stack Inspection. In *Proc. of Symp. on Sec. and Privacy*, 2000.
7. David Evans and Andrew Twyman. Flexible policy-directed code safety. In *Proc. of Symp. on Sec. and Privacy*, pages 32–45, 1999.
8. Greg Goth. The ins and outs of it outsourcing. *IT Professional*, 1(1):11–14, 1999.
9. K.W. Hamlen, G. Morrisett, and F.B. Schneider. Certified in-lined reference monitoring on .net. In *Proc. of the 2006 workshop on Prog. Lang. and Analysis for Security*, pages 7–16, 2006.
10. Charles Handy. Trust and the virtual organization. *Harvard Business Review*, 73(3):40–50, 1995.
11. M. Howard and D. Leblanc. *Writing Secure Code*. Microsoft Press, 2003.
12. Y. Karabulut, F. Kerschbaum, F. Massacci, P. Robinson, and A. Yautsiukhin. Security and trust in it business outsourcing: a manifesto. In *Proc. of STM-06*, volume 179, pages 47–58. ENTCS, 2007.
13. G. Karjoth, B. Pfizmann, M. Schunter, and M. Waidner. Service-oriented assurance comprehensive security by explicit assurances. In *Proc. of QoP-05*. Springer, 2004.
14. Brian LaMacchia and Sebastian Lange. *.NET Framework security*. Addison Wesley, 2002.
15. L.Gong and G.Ellison. *Inside Java(TM) 2 Platform Security: Architecture, API Design, and Implementation*. Pearson Education, 2003.
16. F. Massacci and K. Naliuka. Multi-session security monitoring for mobile code. Technical Report DIT-06-067, UNITN, 2006.
17. F. Massacci and I. Siahaan. Matching midlet’s security claims with a platform security policy using automata modulo theory. In *Proc. of The 12th Nordic Workshop on Secure IT Systems (NordSec’07)*, 2007.
18. Ian Morris. Channel 4’s 4od: Tv on demand, at a price. *CNET Networks Crave Webzine*, 2007.
19. G.C. Necula. Proof-carrying code. In *Proc. of the 23rd ACM SIGPLAN-SIGACT Symp. on Princ. of Prog. Lang.*, pages 106–119, 1997.
20. George C. Necula and Peter Lee. The design and implementation of a certifying compiler. In *Proc. of the ACM SIGPLAN 1998 Conf. on Prog. Lang. Design and Implementation*, number 5, pages 333–344, 1998.
21. Bill Ray. Symbian signing is no protection from spyware. Available at [http://www.theregister.co.uk/2007/05/23/symbian\\_signed\\_spyware/](http://www.theregister.co.uk/2007/05/23/symbian_signed_spyware/), 2007.
22. R.Kazman, M.Klein, and P. Clements. Atam: Method for architecture evaluation. Technical Report CMU/SEI-2000-TR-004, CMU/SEI, 2000.

23. S3MS. Security of software and services for mobile systems. <http://www.s3ms.org/>, 2007.
24. F.B. Schneider. Enforceable security policies. *J. of the ACM*, 3(1):30–50, 2000.
25. R. Sekar, V.N. Venkatakrisnan, S. Basu, S. Bhatkar, and D.C. DuVarney. Model-carrying code: a practical approach for safe execution of untrusted applications. In *Proc. of the 19th ACM Symp. on Operating Sys. Princ.*, pages 15–28, 2003.
26. D. Vanoverberghe, F. Piessens, T. Quillinan, F. Martinelli, and P. Mori. Run-time compliance state of the art. Public Deliverable of EU Research project D4.1.0/D4.2.0, S3MS- Security of Software and Services for Mobile Systems, November 2006.
27. Dries Vanoverberghe and Frank Piessens. A caller-side inline reference monitor for an object-oriented intermediate language. In *Proc. of the 10th IFIP Int. Conf. on Form. Methods for Open Object-based Distributed Systems (FMOODS'08)*, volume 5051, pages 240–258, 2008.
28. David Walker. A type system for expressive security policies. In *Symp. on Principles of Programming Languages*, pages 254–267, 2000.
29. Sachiko Yoshihama, Takeo Yoshizawa, Yuji Watanabe, Michiharu Kudoh, and Kazuko Oyanagi. Dynamic information flow control architecture for web applications. In *Computer Security - ESORICS*, pages 267–282, 2007.

# Securing wireless sensor networks towards a trusted “Internet of Things”

THEODORE ZAHARIADIS<sup>1\*</sup>, PANAGIOTIS TRAKADAS<sup>1</sup>, HELEN LELIGOU<sup>1</sup>,  
KOSTAS PAPADOPOYLOS<sup>1</sup>, EVANGELOS LADIS<sup>2</sup>, CHRISTOS TSELIKIS<sup>2</sup>,  
CHARALAMPOS VANGELATOS<sup>2</sup>, LIONEL BESSON<sup>3</sup>, JUKKA MANNER<sup>4</sup>,  
MICHALIS LOUPIS<sup>5</sup>, FEDERICO ALVAREZ<sup>6</sup>, YANNIS PAPAEFSTATHIOU<sup>7</sup>

*1 Technological Educational Institute of Chalkida*

*2 Hellenic Aerospace Industry S.A.*

*3 Thales Communications France*

*4 University of Helsinki*

*5 Northern Venture Ltd*

*6 University Politechnico de Madrid*

*7 Telecommunication Systems Institute*

**Abstract:** Wireless Sensor Networks (WSN) are quickly gaining popularity due to the fact that they are potentially low-cost solutions, which can be used in a variety of application areas. However, they are also highly susceptible to attacks, due to both the open and distributed nature of the network and the limited resources of the nodes. In this paper, we propose a modular, scalable, secure and trusted networking protocol stack, able to offer self-configuration and secure roaming of data and services over multiple administrative domains and across insecure infrastructures of heterogeneous WSNs. The focus is on trusted route selection, secure service discovery, and intrusion detection, while critical parts of the security functionality may be implemented in low-cost reconfigurable hardware modules, as a defense measurement against side channel attacks.

**Keywords**—security, wireless sensor networks, trust model, service discovery

## 1. Introduction

The past few years, Wireless Sensor Networks (WSN) (also called Web of Sensors) have demonstrated great potential, as integral components of solutions, applicable in a variety of domains, such as environmental observation, surveillance, military monitoring, smart (home) environments and ambient assisting living. The evolving “Internet of Thing” raises even more the already high hopes of WSN and has attracted the interest of the research community worldwide.

Before their wide deployment however, WSNs have to solve some significant problems, including energy management and security [1] [2]. In fact, the initial driving impetus for the development of sensor networks has been military applications, where

---

\* Contact: Th. Zahariadis, TEI of Chalkida, Psachna, GR34400, Greece, zahariad@teihal.gr

security requirements are at their highest. Strong security requirements for such applications are often combined with a hostile and physically unprotected environment. For commercial applications of WSNs, the issue of privacy protection is as important as secure and reliable functioning of a network.

The work presented in this article is carried out in the framework of the AWISSENET (Ad-hoc PAN and Wireless Sensor SEcure NETwork) project, which is partially funded by the European Commission Information and Communication Technologies Program. The project targets the design, implementation and validation of a scalable, secure and context-aware networking protocol stack, able to offer self-configuration and secure roaming of data and services over multiple administrative domains and across insecure infrastructures of heterogeneous ad-hoc & wireless tiny sensory networks. The rest of the paper is organized as follows: we first discuss the intricacies of the problem. In section 3 we present our approach to enhance security in routing, service discovery and intrusion detection, while in section 4 we present the test bed. Finally, in section 5, conclusions are drawn.

## 2. The intricacies of securing wireless sensor networks

Security in WSN denotes protection of information and resources from attacks and misbehaviors, while maintaining an acceptable level of operation even in the case of adverse conditions. The security requirements list is too long, but unfortunately the same applies for the security attack list [2]. Several countermeasures have been proposed addressing specific types of attack; however, only a few proposals try to address multiple attacks, the main reason being the limited resources of sensor nodes.

WSNs share similarities and differences with ad-hoc wireless networks. The main similarity is the multi-hop communication nature, while the main differences are the usually much larger number of nodes and the node constraints in computational, transmission, energy and memory/storage resources. Moreover, WSN are often deployed in open, potentially harsh environments, where they are left unattended for a long period of time after their deployment, allowing physical attacks, such as node capture and tampering [3]. So, the possible attacks range from the physical layer up to the application layer, where aggregation and in-network processing often require trust relationships, between sensor nodes that are not typically assumed in ad-hoc networks. The impedimenta in securing WSN can be classified in two main categories: those introduced by the restricted node architecture and those stemming from the wireless media and the specific sensor network characteristics, as shown in Table 1.

**Table 1:** Restrictions stemming from the node and the network characteristics

Node restrictions	Network - channel restrictions
Low Data rates and small packet size impede the exchange of extra information needed to implement security schemes.	Unreliable communication due to the unreliable of low capacity wireless link with transmission collisions.
Limited Processing Power: 8-bit or 16-bit processor architecture, clock up to 8MHz.	The lack of central infrastructure and the unattended operation of WSN obstructs the implementation of well-established security techniques (e.g. PKI).

Restricted Energy Resources:  Battery powered sensors (in many cases not even rechargeable)	The WSN topology is characterized by a high scale in terms of the number of the participating nodes and in terms of the network density. Also, topology changes occur due to random battery outages and link failures.
---	--

Therefore, in WSN the increased vulnerabilities mandate the design and implementation of a secure network protocol stack taking into account the severe limitations which are inherent in the restricted WSN environment.

### 3. WSN Security Approach

To efficiently address the security problem in WSNs, we propose a modular, secure sensor node “toolbox”, by addressing three key research topics: a) discovery, evaluation and selection of trusted routes, b) secure service discovery, and c) intrusion detection, intruder identification and network recovery. Special emphasis is placed on reducing the footprint, the power consumption and the operating system requirements of the toolbox, to render it adaptable to a large variety of mobile/nomadic devices and tiny sensor nodes, and highly secure against side-attacks.

#### 3.1. Discovery, evaluation and selection of trusted routes

Most sensor nodes have limited communication capabilities and rather short transmission range. Thus, in most cases, they communicate using multi-hop forwarding schemes: they have to forward the packets from node to node, until they reach their final destination. Moreover, WSN have a time-varying networking topology: either because the sensors are randomly deployed or they are moving, or because they are battery powered and each sensor’s lifetime may vary based on the networking activity, dramatically changing the WSN network topology anytime. A wide variety of routing algorithms has been proposed in the literature, efficiently dealing with both the multi-hop forwarding communication and the dynamic topological changes of the WSN.

From a security point of view however, these characteristics turn WSN into an extremely vulnerable environment. A significant number of security attacks target the routing procedure, where malicious nodes either deny to forward their neighbors’ traffic or on purpose advertise fake routes to attract traffic (to forward it to a colluding adversary or just drop it) or declare a fake identity or even modify the traveling messages, both carrying user data and routing protocol information [5]. Hence, the communication security depends heavily on the proper choice of the path used to reach the destination; thus it is important for a node to know the reliability of a route. To achieve trusted routing, it is necessary to design and implement a trust management system to compute the trustworthiness of the participating nodes and detect a node that is misbehaving, either faulty or maliciously. This information can then be exploited to specify a protocol for secure path selection.

We propose a secure routing mechanism combining a geographical routing protocol with a decentralized trust management scheme which can incorporate traditional security measures (e.g. encryption) to safeguard data integrity, confidentiality and node authentication in order to mitigate routing attacks identified in WSN deployments. The adoption of geographical routing prevents a number of routing attacks dealing with advertisement of attractive paths, since in geographical routing the

nodes only announce (broadcast) their position to their one hop neighbour. This intricacy is also the key for scalability and efficient mobility support. As regards the proposed trust scheme, each node is responsible for computing its own trust value for each neighboring sensor node in the network, either collecting events from direct relations with this node (first-hand information), or by collecting indirect trust values from its one hop neighbours (second-hand information). In this concept, every node can build a trust relation with its neighbors, based on the collection of actions (events) performed by other nodes in the neighborhood.

The types of events that we propose each node should monitor are:

**Packet forwarding:** This metric is based on overhearing that a packet has been forwarded (event type E1) and the “packet precision- integrity” (event E3) is checked.

**Network layer ACK.** Each node will monitor whether its message has reached a higher layer node in the proposed architecture by counting the network layer acks received (E2). This is a powerful tool especially when combined with cryptography.

**Authentication – Confidentiality – Integrity.** A node can collect trust information about neighbouring nodes during interactions regarding the proper use of the security measures applied. This behaviour as well as the result of the authentication process is coded as packet Precision-Integrity (event type E3), Authentication (event type E4) and Cryptography-Confidentiality (event type E5).

**Reputation Scheme.** Another way of evaluating the behaviour of a neighbour is by observing its behaviour: If it replies to reputation request messages, it is rated high due to its willingness to participate in the procedure (event type E6). If node C has proposed node B but interaction between A and B is unsuccessful, then A will decrease the direct trust value of node C, since its reputation value about node B has been proven false (event type E7 value).

**Log History.** This type keeps the success or failure of the last  $n$  events (where  $n=16$  or  $n=32$ ). This metric aims in protection against on-off attacks, where malicious nodes try to confuse their neighbours by partially forwarding their messages. It also allows for the detection of the beginning of a malicious behaviour, since when the trust value drop, it can easily be checked whether this is due to past or recent behaviour.

**Other Events.** There is a large set of network events, ranging from hardware-related situations to application layer behaviours that can be used as inputs for the trust management system. For example, for geographic routing protocols, some metrics like the distance of each node to the sink node (E11) may be used.

**Table 2:** Direct Trust Table structure

Trust metric	Maintained Information	
Forwarding (E1)	# of Success	# of Failures
Network-ACK (E2)	# of Success	# of Failures
Packet precision- Integrity (E3)	# of Success	# of Failures
Authentication (E4)	# of Success	# of Failures
Cryptography-Confidentiality (E5)	# of Success	# of Failures
Reputation RES (E6)	# of Response	# of request
Reputation Validation (E7)	Value	
Remaining Energy (E8)	Value	
Network ACK History Log (E9)	1 0 1 1 0 1 0 0 1 1 0 1 0 1 1 1	
Number of Interactions (E10)	Value	
Distance to the sink node (E11)	Value	



The structure of the Trust Table that stores the trust values is shown in Table 2. Each node with  $k$  neighbouring nodes will store  $k$  Trust Tables. Thus, the table size should be as small as possible, while keeping the most important information. In order to take the final forwarding decision, the trust values will be combined with factors like the distance to base station, number of hops to base station and node confidence. This is outside the scope of the current paper. In a heterogeneous sensor environment a subset of the above described events can be monitored and used to evaluate a node's trustworthiness based on different sensor node types and capabilities.

### 3.1.1. Direct Trust Evaluation

For each one of the first 6 events of Table 1, node's A Trust regarding node B, i.e.  $T_i^{A,B}$ , can be calculated:

$$T_i^{A,B} = \frac{a_i S_i^{A,B} - b_i F_i^{A,B}}{a_i S_i^{A,B} + b_i F_i^{A,B}} \quad (1)$$

Where:  $S_i^{A,B}$  and  $F_i^{A,B}$  are the successful and failed type  $i$  events that A has measured for B,  $a_i$  and  $b_i$  represent the weight/significance of a success vs. failure and their values will be evaluated using computer simulations.

For the History Log (E9), we propose a simple pattern matching technique which will help towards either calculating the trust value or categorizing the neighbouring nodes activity. The number of interactions (shown as E10) is a measure of confidence. A high confidence value means that the target has passed a large number of tests that the issuer has set, or that the issuer has interacted with the target for a long time, and the node is sure that the value of the neighbouring node is more certain. The algorithm of trust evaluation is more sensitive in the beginning of the interactions period (since confidence value is small, one fault should have a large impact in trust value), while as confidence value increases, the impact (either on positive or negative events) is smoother. Thus, we define a confidence factor, like in the next equation:

$$C^{A,B} = 1 - \frac{1}{n_{oi} + a_{10}} \quad (2)$$

Where  $n_{oi}$  indicates the number of interactions with B and  $a_{10}$  is a factor whose value will be evaluated during simulation testing. This confidence factor can be proved useful, especially during the beginning of network operation.

In case of geographic routing algorithms, a proper metric is the distance of the neighboring node to the sink (E11). The closer a node to the sink, the greater the value added to the final direct trust of the node.

Finally, node's A Direct Trust value for its neighboring node B, i.e.  $DT^{A,B}$  with  $k$  event types (in our case  $k=10$ ) can be calculated according to the following equation:

$$DT^{A,B} = C^{A,B} \left( \sum_{i=1}^k W_i * T_i^{A,B} \right) \quad (3)$$

Where:  $W_i$  is the weighting factor for each one of the  $k$  event types,  $T_i^{A,B}$  is node's A trust value of event  $i$  regarding node B. The use of the weighting factors is a very important feature of the adopted trust model. By using these weighting factors, during the simulation and validation process, we'll be able to categorize the severity of each one of the events that will have a different impact on the direct trust value.



### 3.1.2. Indirect trust/Reputation evaluation

There are several cases where a node (e.g. node A) needs the trust opinion of its neighbouring nodes (e.g. node C, D, E) regarding a specific node (node B). Examples of such cases may be the discovery of a new node appeared during a HELLO message or when direct trust value of node B is neutral (its value is neither large nor small). In the proposed trust model, a node A may find the indirect trust/reputation value of a node B i.e. the  $IT^{A,B}$  by combining the direct trust values (reputation values) of its neighboring nodes, as shown in the following equation:

$$IT^{A,B} = \sum_{j=1}^n W(DT^{A,N_j}) DT^{N_j,B} \quad (4)$$

Where,  $n$  is the number of neighbouring nodes to A,  $N_j$  are the neighbouring nodes to A,  $DT^{N_j,B}$  is node's  $N_j$  reputation value of node B,  $W(DT^{A,N_j})$  is a weighting factor reflecting node's A direct trust value of node  $N_j$

As in the previous section, we use different weighting factors for each node regarding the events described above. For example, if node's C direct trust value (evaluated by node A) is large and also node C is frequently sending responses to node's A requests, then its weighting factor is large. The reputation value  $DT^{N_j,B}$  that the neighbouring nodes propagate to the interested node are kept to the Reputation – Indirect Trust Table, thus the interested node can check the correctness of their answers on next route discovery phase and modify the direct trust values of the neighbours  $W(DT^{A,N_j})$  accordingly (e.g. increase the direct trust value of a node who gave a reputation that was proved correct). This is the reason of the direct trust value selection, instead of the sum of direct and indirect trust values. The metrics that allow node A to evaluate node's B trustworthiness in this case are the node's direct trust value, which includes its responsiveness in the reputation scheme implementation as well as the provided reputation value.

### 3.1.3. Total Trust evaluation

The total trust evaluation node A of node B, i.e.  $TT^{A,B}$  is performed by applying the following equation:

$$TT^{A,B} = W(DT^{A,B}) DT^{A,B} + W(IT^{A,B}) IT^{A,B} \quad (5)$$

Where  $DT^{A,B}$  and  $IT^{A,B}$  are A's direct and indirect trust values of B,  $W(DT^{A,B})$  and  $W(IT^{A,B})$  are a weighting factors reflecting A's direct and indirect trust value of B. Since A can be sure only about the first-hand information, the weighting factor of the Direct Trust Value will be larger than the weighting factor of the Indirect Trust value.

### 3.2. Secure Service Discovery

Automated service discovery is an important functionality in a dynamic environment such as WSNs, e.g., sensor nodes need to find where the sinks are. Yet, the variable connectivity of the nodes coupled with a hostile environment may cause non-uniform or even false service information propagation. In such environment, the service discovery scheme may be based on a push-style information dissemination method, or on pulling information out of the network when needed. In the former case,

client nodes passively monitor the network and get to know about possible services, e.g., this is how IPv6 Router Advertisements work. Client nodes can also actively query the network for services; in IPv6, a node can send a Router Solicitation and force routers to tell about them. We can also use hybrid controlled dissemination where information about a certain service is not stored everywhere in the network, but at one, or a handful, of nodes; this can reduce the signaling load in the network but brings new authorization concerns, e.g., how can we trust the information coming from some intermediate node. The discovery, whether passive or active, can be further enhanced through coupling with route discovery; we can piggy-back service announcement and query messages in routing protocol messaging and thus, at the same time, gain knowledge of routes and available services.

Important and very difficult challenges in service discovery are the authentication of parties involved in the service discovery, and the confidentiality of the service signaling. Even if the content of a service is public information, we still want to confirm the source of the information and make sure the content is valid. For example, a news service or announcements at an airport are public information, but users probably want to be sure that the source is truly the news company or the airport authority, and not someone fooling around on purpose, or even trying to send unsolicited spam. In other services, only authorized requesters are allowed to receive an answer to a query and the requester needs some guarantees that the service is valid. The attacks against a service discovery system are listed in Table 3. Their detection is difficult. In the fake services case, we may not know where the information was altered, while in the Advertisement and query flooding case can be coupled with Cybil attack which makes it harder to identify. Finally, Listening & Profiling is a passive attack and in general we can not detect it.

**Table 3:** Attacks against service discovery

Attack	The attacker actions	Consequences
Fake services	An attacker responds to service queries even if it doesn't have the service or provides misleading information	Battery drain, unintentional Denial of Service
Advertisement and query flooding	An attacker sends massive number of advertisements or queries	The network spends resources in forwarding the messages  The recipients spend CPU cycles and energy in receiving and processing the messages
Listening & Profiling	The attacker observes and profiles both the service provider and the client. By passively listening to the communication	Sensitive information is received by an unauthorized entity

The simple solution here is to have the right certificates and encryption keys at the receiver to verify or decrypt information. If the communication is fully encrypted, the attacker must first fight the encryption before anything else can be done; thus, only some sort of flooding attacks can be performed but if we assign per-host rate limits in routing we can reduce the effect. The more challenging situation is with unencrypted messaging, when we need some mechanism to get the right certificates.

The deployment scenarios in WSNs are very challenging, since we have different sensors and need to make services available between them. Typical fixed, or even ad-

hoc, network protocols can not be employed. Thus, we investigate and work on a hybrid system, which includes alternative signaling mechanisms, proxies and information caching, coupled with a trust mechanism between entities. Since routing in this environment has also similar challenges to the above, we investigate ways to couple routing and service discovery together and use a unified trust management scheme to counter the various attacks.

### 3.3. Intrusion detection, intruder identification and recovery

Key management, secure routing and secure services can be considered as a first line of defense, aimed at preventing malicious nodes to break into the network or to retrieve confidential information. However, there is a non-negligible possibility that an intruder, finally becomes successful. Thus, a second line of defense is needed, able to detect third party’s attacks and raise alarms, even if the attacks haven’t been experienced before. The proposed WSN Intrusion Detection System (WIDS) takes care of this role. WIDS differ in many ways from the one used in legacy networks. In order to achieve an efficient, secure and lightweight WIDS, the proposed system uses innovative architectures and algorithms that we outline hereafter.

**Network Architecture.** Usual IDS are typically *stand-alone IDS*, where each node runs an independent intrusion detector. Such systems are very limited in AWSNs, since local audit data is not enough to have a good comprehension of what is happening on the network. Cooperation between the different nodes is compulsory in order to achieve efficient detection, because local evidences are inconclusive. Since the network infrastructures that AWSNs can be configured to are either flat or multilayered, the same approach can be used for intrusion detection systems. *Hierarchical IDS* are systems where specific nodes are in charge of monitoring their neighbours, with various level of cooperation between cluster heads [5]. *Distributed IDS* meet the decentralized nature of ad hoc wireless sensor networks, where each node is responsible for collecting local audit data, and this knowledge is shared globally in order to carry out a global intrusion detection system [6], [7]. We propose a mixed approach, where the WIDS inside a cluster will be fully distributed, and cluster heads are responsible for exchanges and decisions at the upper level.

**Collecting audit data.** Data is collected by local agents analyzing local sources of information, which can be hardware or network based. Due to the wireless, ad-hoc nature of the network, nodes don’t only analyse packets sent to them, but can also overhear traffic passing from a neighbouring node and act as a watchdog, detecting nodes forwarding selectively packets, or modifying them [8].

**Intrusion Detection.** Detection techniques can be either *Misuse detection* (where audit data is compared with known attack patterns), *Anomaly detection* (detect when the network behaviour differs from ‘normal’ behaviour, established via automated training) and *Specification-based detection* (similar to the former one, but the correct behaviour is manually defined) [7]. The proposed WIDS will have a modular software architecture where one can plug different detection mechanism (available as plug-ins), depending on the hardware of the nodes and what one intends to monitor.

**Detection & Recovery.** Once a local IDS agent has raised an alarm internally, the next question is who is going to make the final decision that a node is effectively an intruder. *Independent decision-making systems* are usually used in cluster-based architectures because they leave the decision that a node is effectively an intruder to certain nodes. The alternative solution is *Cooperative Intrusion Detection Systems*.


When an attack seems to have been detected, the node appeals to neighbouring nodes in order to output a global decision. The proposed WIDS uses a mixed approach, being cooperative inside a cluster, but relying on cluster heads at the upper level. Once a node has been assessed as malicious, the alert is forwarded to the modules in charge of isolating it, like the routing module and secure service discovery.

#### 3.4. Highly secure nodes

The nodes in a sensor network are by nature distributed and thus, in the vast majority of the cases, they are very vulnerable to side attacks. These attacks are based on measuring characteristics of the processing activity on a node such as power consumption, electromagnetic emission, timing. By analyzing those measurements the attacker may recover all or part of the secret information stored in the node (e.g. the key used in the majority of the security algorithms). AWISSENET designs a new node architecture which utilizes the newly introduced extremely low-cost and low power Field Programmable Gate Arrays; this node is practically invulnerable to such side attacks. The implementation will be mainly based on two methods: Dual Rail encoding and Masking [10]. These implementation techniques, together with architectural design methodologies such as spreading processing tasks in random time periods, render the measurements performed by the side attacker useless for him/her.

#### 4. The AWISSENET test bed

The efficiency of the proposed toolbox will be first assessed through exhaustive simulations and then the system will be validated in a trial involving 100 sensor nodes [10]. The trial consists of different “sensor nodes islands” involving different technologies which will be linked together using an IP network. As shown in Figure 1 we are setting up 4 different linked scenarios: one homogeneous and one heterogeneous mote islands, one RFID island and one multimedia island.



**Figure 1:** The AWISSENET Trial architecture

The aim of the trial is twofold: demonstrate the correct and efficient working of the technologies against attacks and prove the applicability to Personal Area Network and

sensor application scenarios. We aim at validating our developments in 4 different environments: industry, home, roads and disaster recovery. For this purposes each island will be equipped with adequate sensor types in each node which can be used to validate the environment we are testing. For example, the multimedia island in the home environment using smart cameras or microphones attached to the nodes can be used for creating a trustworthy and secure surveillance system which can demonstrate the applicability of the proposed solution. For the security validation, we are contemplating also the testing of cross domains and cross island communications which will give the final conclusions of the reliability and trustworthiness of the solutions described in the paper.

## 5. Conclusions

Ad-hoc personal area networks (PAN) and wireless sensor networks impose new challenges on the design of security tools which are more imperative than ever due to their unattended operation in open environments. We propose to implement and pack a set of security-aware protocols from the network to the application layer in a flexible security toolbox which can then be used in a variety of wireless devices [11]. The goal is to efficiently defend against a great number of attacks including side channel attacks focusing on those dealing with service discovery, routing, and intrusion detection. Our concept will be validated through a large and heterogeneous test-bed.

**Acknowledgment:** The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 211998 (AWISSENET project).

## References

- [1] K. Papadopoulos, S. Voliotis, A. Ktena, P. Trakadas, Th. Zahariadis, "Security Aspects in Wireless Sensor Networks," Int. Conference on Telecommunications and Multimedia (TEMU 2008), Ierapetra, Crete, Greece, 16-18 July 2008
- [2] Ivan Stojmenovic, "Handbook of Sensor Networks: Algorithms and Architectures", John Wiley & Sons, 2005, Ch.1.
- [3] V. C. Giruka, M. Singhal, J. Royalty, S. Varanasi, "Security in wireless sensor networks", *Wirel. Communications Mob. Comput.* 2008; 8:1–24.
- [4] Bounpadith Kannhavong, Hidehisa Nakayama, Yoshiaki Nemoto, And Nei Kato, Abbas Jamalipour, "A survey of routing attacks in mobile ad hoc networks", *IEEE Wireless Communications*, October 2007, pp. 85-91.
- [5] O. Kachirski, R. Guha, D. Schwartz, S. Stoecklin, and E. Yilmaz, "Case-based agents for packet-level intrusion detection in ad hoc networks," in *Proc. of the 17th Int. Symposium on Computer and Information Sciences*. CRC Press, October 2002, pp. 315–320.
- [6] Ioanna Stamouli, "Real-time Intrusion Detection for Ad hoc Networks", Master of Science dissertation, University of Dublin, 2003.
- [7] K. Ioannis, T. Dimitriou, F. C. Freiling, "Towards Intrusion Detection in Wireless Sensor Networks", 13<sup>th</sup> European Wireless Conference, Paris, April 1997
- [8] R. Roman, J. Zhou, J. Lopez, "Applying Intrusion Detection Systems to Wireless Sensor Networks", *Consumer Communications and Networking Conference*, 2006, pp. 640-644.
- [9] K. Tiri and I. Verbauwhede, "A Logic Level Design Methodology for a Secure DPA Resistant ASIC or FPGA Implementation," in *Proc. of Design Automation and Test in Europe Conference (DATE 2004)*, pp. 246-251, 2004
- [10] P. Trakadas, T. Zahariadis, H.C. Leligou, S. Voliotis, K. Papadopoulos, "AWISSENET: Setting up a Secure Wireless Sensor Network," 50th International Symposium ELMAR-2008, focused on Mobile Multimedia, Zadar, Croatia, 10-13 September 2008, pp. 519-523
- [11] K. Papadopoulos, S. Voliotis, H.C. Leligou, D. Bargiotas, P. Trakadas, Th. Zahariadis, "A Lightweight Trust Model for Wireless Sensor Networks," *Numerical Analysis and Applied Mathematics (ICNAAM 2008)*, Kos, Greece, 16-20 September 2008, pp.420-423

## Privacy-enabled identity management in the Future Internet

Christoph Sorge, Joao Girao, and Amardeo Sarma

NEC Laboratories Europe, Kurfürsten-Anlage 36, 69115 Heidelberg, Germany  
{christoph.sorge,joao.girao,sarma}@nw.neclab.eu

**Abstract.** Identity management has the potential to play a major role in the Future Internet as an enabling technology that integrates services with transport infrastructures. The use of partial identities can improve users' control over their personal data and enhance privacy, but revealed data can be used by service providers more than strictly needed, such as for user profiling. Without adequate precautions, service and identity providers could pick up profile and service usage data during the authentication process. Such privacy concerns have gained attention of late, and legal constraints are expected to follow to further limit data transfer from identity providers to service providers. The paper discusses legal constraints from a European perspective, taking into account the possibility of network operators acting as identity providers, and approaches to enhancing privacy. We show how European legislation and user needs for privacy affect the design of Identity Management Systems and outline consequences as well as opportunities and directions for future research.

*Keywords:* Identity Management, Future Internet, Privacy, Law

### 1 Introduction

Identity management has been defined as “managing various partial identities (usually denoted by pseudonyms) of an individual, i.e., administration of identity attributes including the development and choice of the partial identity and pseudonym to be (re-)used in a specific context or role.” [1] This means that identity management solutions help a user managing which attributes (like age or name) to reveal, e.g. to service providers he or she interacts with.

Privacy was an underlying motivation to develop the Virtual Identity (VID) concept in the EU IST project DAIDALOS<sup>1</sup>, which has been adopted in the EU IST project MAGNET Beyond<sup>2</sup> and currently being developed further in the EU ICT SWIFT<sup>3</sup> project of the 7th Framework Programme. A major focus in these

<sup>1</sup> Designing Advanced network Interfaces for the Delivery and Administration of Location independent, Optimised personal Services; see <http://www.ist-daidalos.org> and [2] for an overview of the project

<sup>2</sup> My personal Adaptive Global NET; see <http://www.ist-magnet.org>

<sup>3</sup> Secure Widespread Identities for Federated Telecommunications; see <http://www.ist-swift.org>

projects was to extend privacy support to the network layers. This paper builds on the work done here and addresses the legal dimension, in addition to privacy, to make identity management a part of the Future Internet, taking into account the often conflicting interests of users, providers and legislation and the trade-offs involved. A future identity management solution should support convergence and provide network operators and service providers with the potential for new business and revenue while complying with privacy and legal requirements.

The rest of the paper is structured as follows: Section 2 introduces the role of identities in the Future Internet. Section 3 discusses privacy from the viewpoints of the user, legislation and service providers followed by Section 4 on legal consequences. Based on initial work in DAIDALOS and SWIFT, Section 5 looks at some of the coming challenges. The paper concludes with Section 6.

## **2 Role of Identities in the Future Internet**

In the original design of the Internet, the management of identities did not play an important role. Few computers with many users were connected, usually via a single interface. The number of computers, many with multiple interfaces, has exploded, as has the number of users, who now own or use multiple computers or devices. Handling all identities involved in the use of an Internet-based service is quite complex—especially if privacy considerations are taken into account. An IP or MAC address may uniquely identify the user, which may be undesirable.

### **2.1 Identities of communication endpoints**

Nodes participating in the current Internet have an IP address that serves both as identifier and locator. Researchers have pointed out possible weaknesses of this double function [3]. The Host Identity Protocol Architecture [4] introduces an additional namespace for the identification of endpoints; among other advantages, this means that mobile endpoints can keep their identities even when attaching to a different network and getting a new IP address. In the HIP architecture, transport layer connections are bound to host identities instead of IP addresses. The entity providing the mapping from a host identifier to an IP address (e.g. a HIP rendezvous server [5]) might be considered as a very simple identity provider. But this does not address the issue of the actual entity that is the endpoint of a communication. If it is not the device itself, but e.g. a person or an organization, there may in general be an  $n:m$  mapping of entities to hosts, and the host becomes the first intermediate hop. In the following, the term user is often used in the more general sense of an entity that may not be human.

### **2.2 Cross-layer privacy with identity management**

The DAIDALOS/SWIFT Virtual Identity (VID) concept goes one step further. All personal data required for the use of a specific service is associated to a VID [6]. The VID itself is a collection of data associated with a specific user role and

purpose. This contained data is not necessarily located in one place—in fact it usually is not, but the overall VID concept puts the user in charge of which parts of the data are revealed in any interaction. It is important to note that, from a privacy perspective, it is insufficient to only associate service-level data to the VID. Data required by the transport infrastructure must also be revealed. However, this data should not be usable to link or identify virtual identities, for example by knowing the IP address.

DAIDALOS has developed a simplified “bootstrap” procedure to associate addresses at various levels. This allows several virtual identities of the same or different users to access services from a single device via the network. Different MAC and IP addresses do not allow an attacker to relate the virtual identities to the same user or device. Data retention requirements may need additional methods that resolve these addresses to a specific user if required, for example, by a court order, while protecting this data from unauthorized access.

SWIFT further leverages identity management as an enabling technology to integrate services with transport infrastructures and beyond, being applicable even to physical services, such as entering a building, buying articles or borrowing a book from a library. Managed identities can play a crucial role for Future Internet applications—improving usability, simplifying the composition of services by multiple providers, and enabling service personalization.

### 3 Privacy and conflicting interests

In recent years, several instances where data of users have been “lost” or compromised has further raised privacy awareness. At the same time, privacy must be seen in the context of conflicts both of interest and of priority of the concerned.

Regarding priority, the user herself may prefer usability to privacy or actually wish to reveal data. Here, the user must be in a position to both understand and control the intended privacy level. At the same time, law enforcement authorities may want to limit privacy in some cases, e.g. to prevent crime. Finally, service providers themselves may need some data, e.g. to ensure payment, but often want to collect data beyond that and for other purposes. Not willing to disclose some attributes, such as age or an account number may lead to the service failing.

Identity management (IdM) solutions can be installed on the user’s side, where an application decides which identity and credentials to use for authentication towards different services. Service providers (SPs) can also support identity management. Identity providers (IdP) are additional players, and service providers may choose to trust these IdPs instead of having identity management facilities themselves. They then accept the user’s authentication towards the IdP and rely on attributes provided by the IdP. Such attributes may be necessary to provide a service, such as the capability to pay or an age category. Network operators are in a particularly good position to take on the role of an IdP. They have identified their customers—either due to a legal obligation, or because service provisioning and billing make it necessary to know the customer’s name and address. Banks or credit card companies may also take up such a trusted role, though they may not command the same amount of trust to deliver services.



However, the IdP's knowledge also raises privacy concerns, and especially network operators are often subject to particularly strict privacy regulations. They may also be forced to retain data to later uncover criminal activities. Therefore, we must look at the regulatory framework in addition to the privacy needs of the customers. Future identity management solutions will need to find a balanced solution that caters for the legal requirements, enhances business and supports the users' privacy needs. Some trade-offs will be needed, such as when users do not wish to reveal data that is not necessary for service delivery.

### 3.1 Privacy legislation basics

Data privacy (or data protection) legislation differs greatly from country to country, but there is a common principle: The aim of this legislation is to protect personal data, i.e. data that can be attributed to a person (the "subject"). This data is sometimes also referred to as "personally identifiable information" (PII). *Anonymous* data are not protected. This is not necessarily true for *pseudonymous* data, as a pseudonym might be uncovered.

National legislations have often been inspired by the OECD privacy guidelines [7]. Among these guidelines is the purpose specification principle: Data may only be collected for a specific purpose, and may only be used for other purposes if authorized by the law or the subject (the person the data are about).

In the European Union, this principle has been adopted by the 1995 data privacy directive [8, article 6, section 1b]. The directive allows processing of personal data only in a number of cases. In particular, data processing is allowed with the subject's consent or if it is necessary to fulfill an obligation arising from a contract with the subject (see article 7, section b, of the directive). As a further exception, the directive allows data processing if it "is necessary for the purposes of the legitimate interests pursued by the controller or by the third party or parties to whom the data are disclosed, except where such interests are overridden by the interests for fundamental rights and freedoms of the data subject [...]" (article 7, section f). This clause seems to be a very far-reaching exception at first glance. However, it calls for a balancing of the protected interests of controller and data subject and does not allow arbitrary data processing. A legitimate interest that is not overridden by the subject's interests may exist in case of abuse tracking, but typically not during a service provider's normal operation. In addition, article 14, section a, of the directive gives the data subject a right to object the processing of personal data in this case.

The U.S., in contrast, have not adopted the OECD privacy guidelines. A number of sector-specific laws regulate the processing of personal data, particularly by government agencies.

### 3.2 International considerations

For identity management systems deployed internationally, it is important to determine the applicable law and the resulting restrictions concerning data processing and data transfer. The question of applicable law is dealt with by national


regulations. As a principle, service provider and user can choose which law to apply. Otherwise, the law of the country will be applied to which the contract has the closest relationship; this would typically be the service provider's country. Yet, consumer protection laws—including privacy legislation—are applied by some countries independent of this choice of law. As a consequence, service providers have to consider both the privacy legislation of their own country and the legislation of their customers' countries.

The European directive allows the free flow of data within the European Union. The transfer of personal data out of the European Union is subject to restrictions. With a few exceptions, such a data transfer is only allowed

- with the user's explicit consent,
- if the country to which the data is to be transferred has an appropriate privacy protection standard compared to the one of the EU,
- or if it is ensured that personal data and user's privacy are sufficiently protected as in the case of US companies that comply with the Safe Harbor Agreement.

### 3.3 Network operators and service providers

The European directive 2002/58/EC [9], dealing specifically with telecommunications privacy, mainly protects the secrecy of communications and of traffic data. Acting as identity provider does not require the revelation of the content of communications. Not even metadata about users' communication behaviour needs to be revealed. The SP will get to know some of this information—in particular, the user's IP address—as it is necessary to communicate with the user, but this fact is independent of identity provisioning. As a consequence, if a network operator wants to act as an IdP, restrictions specific to its role as a network operator do not apply. As an exception, using location data as identity attributes to be made available to the service provider requires the user's consent (article 9 of the directive 2002/58/EC). Other than that, only directive 95/46/EC is applicable, its consequences being discussed in section 4. Figure 1



**Fig. 1.** A basic Identity Management scenario

depicts a basic identity management scenario: A user authenticates towards his identity provider with his identity U. When the user wants to use a service, the IdP confirms to the service provider SP that the user has been authenticated,

and transmits an identity  $U'$ . As  $U'$  does not consist of all of the user's attributes, it is also referred to as a partial (or virtual) identity.

- The IdP is assumed to have a superset of all user attributes required by the service providers. This may include personal data, such as name and address, phone number, age, ...
- The SP has limited identity information. However, it does have detailed data about service usage. Depending on the kind of data and national legislations, this data may be considered personal. The European data privacy directive considers any “information relating to an identified or identifiable natural person” to be personal data. IP addresses are widely believed to be covered by this definition (see, e.g., [10]), but this is disputed.

Without a separate identity provider, both the service usage and identity information would be stored in the same location, i.e. at the service provider.

## 4 Legal consequences

Legal restrictions may be imposed on the collection, the use, or the transfer of personal data. The European data privacy directive, however, does not make a distinction between the three. From the communications perspective, we are most interested in the transfer of personal data between the three players depicted in Figure 1. As discussed in Section 3.1, the data privacy directive adopts the purpose specification principle: Personal data has to be “collected for specified, explicit and legitimate purposes and not further processed in a way incompatible with those purposes” (article 6 section 1b of the directive).

### 4.1 Restrictions on data processing

Despite identity management efforts, personal data are likely to be transferred *from the user to the SP*: The SP typically knows a partial identity consisting of attributes transmitted by the IdP, and the user's IP address (being considered as personal data). Therefore, the SP may collect data only for a specific purpose—typically service provisioning itself. At first glance, this result seems unsatisfactory from the SP's point of view: Using the identity management system does not allow the SP to collect more data than with classic authentication schemes. Still, there are two major advantages of using identity providers. Firstly, the SP need not take care of authentication itself. Secondly, the IdP can provide (possibly short-lived) partial identities and refrain from retaining a mapping to the user's “real-world” identity. This way, the SP may retain data longer than necessary for service provisioning and even bind them to the partial identity, as long as a mapping to an actual person is impossible. In particular, this includes deleting the IP address as soon as possible.

Transfer of personal data *from the user to the IdP*, too, is worth considering. If the user enters into a contract with the IdP about identity provisioning, the IdP may collect all data necessary to fulfill its contract. If, on the other hand,

an existing network operator (or a service provider) also wants to act as an IdP, it has to ask for the user's consent.

Finally, there is the relationship *between IdP and SP*. Once again, we differentiate two cases: If the user has concluded an identity provisioning contract with the IdP, the transfer of personal data from the IdP to the SP is typically covered by this contract. However, there may be no such contract. If a network operator chooses to act as an IdP, it could simply decide to give identity information about its users to all SPs it considers trustworthy. This would be a comfortable means of authentication for the user. It may even serve the fulfillment of a contract—but this is true only for the contract between user and SP. However, this kind of data transfer requires the user's explicit consent or the existence of a contract about identity provisioning.

There may also be reasons for the SP to transfer personal data to the identity provider. This may be helpful if the IdP performs additional tasks: For example, the IdP could store reputation information about the user, or store usage information in order to facilitate personalization across several services or for billing purposes. Once again, it depends on contractual relationships whether this is allowed. Consider the example of an identity provider that also stores user profiles for service personalization. These user profiles would be transferred to service providers. Though this may be convenient, it is typically not needed to fulfil the SP's contract with the user, and the IdP should explicitly ask for the user's consent.

## 4.2 European Data protection working party

The European Commission receives advice from a commission of experts, the so-called Article 29 Data Protection Working Party. In 2003, this working party adopted a document on "on-line authentication services" [11], being a subset of identity management systems. The report was mainly concerned with Microsoft's Passport service. As a result of the working party's statements, Microsoft decided to perform a number of changes to this service. However, the results of the report are applicable to a wider range of systems<sup>4</sup>:

- A centralized system storing personal data may both lead to security risks and facilitate abuse of that data.
- The users should be in control of which data each SP receives from the IdP and vice versa.
- The use of a single unique identifier could enable SPs to build user profiles by exchanging information, using the identifier as a key. Users should also be able to access their own unique identifiers. The working party favors the Liberty Alliance Approach, which does not require a single unique identifier for a user.
- The contractual framework between service providers and identity providers plays an important role; contracts should make each party's obligations concerning the processing of personal data explicit.

---

<sup>4</sup> Our list contains a selection of the most relevant results.

These concerns go, in part, beyond what is codified in the directive 95/46/EC, and therefore, they are purely recommendations. For example, centralized data storage is not illegal despite the consequent risk of abuse. Identity management solutions should take the working party's opinion into account.

### 4.3 Additional obligations

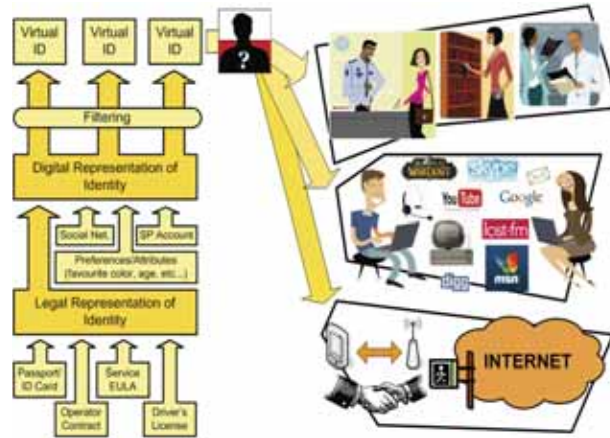
In addition to the presented restrictions on data processing, more obligations must be met: The data subject must be given certain information, specified in articles 10 and 11 of the European data privacy directive, when data are collected. The subject has a right to obtain information about all personal data processed about him and about other parties who receive this data. National legislations may even go beyond these requirements. Germany's Telemediengesetz (telemedia act), for example, defines rules for the *design* of online services processing personal data, not just for processing itself.

Regulations also apply to the form of a user's consent to data processing. Article 2 of directive 95/46/EC defines this consent to be "any freely given specific and informed indication of his wishes by which the data subject signifies his agreement to personal data relating to him being processed". Note in particular that consent must be "informed", meaning that the user must know what his consent is about. National legislation can be more specific. In Germany, for example, the federal data protection act (Bundesdatenschutzgesetz) requires the consent to be highlighted if it is given in writing together with other declarations—it may not be hidden in the fine print. In case of electronic services, the German telemedia act (Telemediengesetz) allows giving the consent in an electronic form, but requires ensuring that it is based on an intentional and unambiguous act. In addition, the consent must be protocolled and made available to the user for future reference.

## 5 Directions for Identities in the Digital World

To fulfil restrictions on handling personal information, rights of the user and legal requirements, the future Identity Management system must cover a wide-range of scenarios. It must also be flexible to adapt itself to the pan-European diversity in culture and legal requirements. IdM thus becomes an inherent part to all our movements online. A vertical digital identity solution should allow us to clearly scope the information we would like to provide to a certain service and keep this information coherent throughout the lifetime of the digital identity. This lifecycle may be long, spanning years of employment, or extremely short such as the minutes it takes to perform a transaction online.

Figure 2 demonstrates how one may build a digital identity and how that identity may be used. On the left, we can observe several pillars on which a digital identity is built: attributes and relations. While some are strong and even legally binding, others may be ephemeral and even self asserted. Once one gathers and categorizes this different data, one can build a Virtual Identity which is no more



**Fig. 2.** Building your online Virtual Identity

than a clearly defined subset. To ensure privacy, such as in the real world, these subsets should be based on the categories which most distinguish one's activities online and the user is not always the best person to define that.

Once a digital identity is established, its use and inherent privacy benefits can be collected when accessing a network, or using a device, when connecting to a service, or communicating with others, and even in the real world, where a closer binding between the real person and the digital person is many times enforced even without the knowledge or control of the user (e.g. health systems).

We see the future of digital identities on three paths: Vertical Use, Data Definition and Usability. *Vertical Use* refers to networks, services and the world outside the cloud perceiving the user's digital ID in a consistent and privacy aware manner. *Data Definition* refers to identity data, its distribution and exchange, taking into account both cultural and technology aspects. *Usability* must be built-in to the system so that the user does not explicitly handle the use of their digital IDs. A much higher level of transparency is expected both in the way a service handles the user data and the mechanisms which enforce the user's privacy protection.

## 6 Conclusion

Identity management can be considered a great chance to improve users' privacy, as well as having a major potential to enhance and enrich Future Internet solutions, putting the user into the centre of considerations. Privacy, fulfilling both user wishes and legal requirements, must be taken into account from the very start and in particular in the design of identity management systems. We discussed requirements with an emphasis on European law, showing another dimension that must be considered. The European Union is among the regions with the strictest privacy regulations. However, since these regulations affect re-

strictions on the transfer of personal data out of the EU, service providers outside the EU cannot neglect privacy issues when doing business with Europeans.

A solution is needed that improves on the way users interact with the Digital World via the Future Internet, while achieving transparency concerning the processing of users' personal data, and to process this data only for a specific purpose the user agrees with. The Virtual Identity concept developed in DAIDALOS and SWIFT can be a major enabler both for privacy and for convergence of networks, services and even physical interactions. Besides catering for privacy needs in all its dimensions, we expect a converged vertical use supported by a cross-layer approach, a well structured data solution for Identity Management and a strong focus on usability with identities as the end-points of future communications to be key elements of the Future Internet.

## References

1. Pfitzmann, A., Hansen, M.: Anonymity, unlinkability, undetectability, unobservability, pseudonymity, and identity management—a consolidated proposal for terminology. Version 0.31, available from [http://dud.inf.tu-dresden.de/literatur/Anon\\_Terminology\\_v0.31.pdf](http://dud.inf.tu-dresden.de/literatur/Anon_Terminology_v0.31.pdf) (2008)
2. Aguiar, R.L., Sarma, A., Bijwaard, D., Marchetti, L., Pacyna, P., Pascotto, R.: Pervasiveness in a competitive multi-operator environment: the DAIDALOS project. *IEEE Communications Magazine* **45**(10) (October 2007) 22–26
3. Saltzer, J.: On the Naming and Binding of Network Destinations. RFC 1498 (Informational) (August 1993) Previously published in *Local Computer Networks*, North-Holland Publishing Company, Amsterdam, 1982, pp. 311–317.
4. Moskowitz, R., Nikander, P.: Host Identity Protocol (HIP) Architecture. RFC 4423 (Informational) (May 2006)
5. Laganier, J., Eggert, L.: Host Identity Protocol (HIP) Rendezvous Extension. RFC 5204 (Experimental) (April 2008)
6. Clarke, J., Butler, S., Hauser, C., Neubauer, M., Robertson, P., Orazem, I., Blazic, A.J., Williams, H., Yang, Y.: Security and privacy in a pervasive world. In: *EURESCOM Summit 2005, Ubiquitous Services and Applications—Exploiting the Potential*, Conference Proceedings, 27–29 April 2005, Heidelberg, Germany. (2005) 315–322
7. Organisation for Economic Co-Operation and Development: OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data (September 1980) Available from [http://www.oecd.org/document/18/0,2340,en\\_2649\\_34255\\_1815186\\_119820\\_1\\_1\\_1,00.html](http://www.oecd.org/document/18/0,2340,en_2649_34255_1815186_119820_1_1_1,00.html).
8. European Community: Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data (1995)
9. European Community: Directive 2002/58/EC of the European Parliament and of the Council of 12 July 2002 concerning the processing of personal data and the protection of privacy in the electronic communications sector (Directive on privacy and electronic communications)
10. Article 29 Data Protection Working Party: Privacy on the Internet: An integrated EU Approach to On-line Data Protection, 5063/00/EN, WP 37 (2000)
11. Article 29 Data Protection Working Party: Working Document on on-line authentication services, 10054/03/EN, WP 68 (2003)

# Control of Resources in Pan-European Testbed Federation

Anastasios GAVRAS<sup>a</sup>, Halid HRASNICA<sup>a</sup>, Sebastian WAHLE<sup>b</sup>, David LOZANO<sup>c</sup>,  
Denis MISCHLER<sup>d</sup>, Spyros DENAZIS<sup>e</sup>

<sup>a</sup>*Eurescom GmbH, Germany*, <sup>b</sup>*Fraunhofer FOKUS, Germany*, <sup>c</sup>*Telefonica I+D, Spain*,  
<sup>d</sup>*Thomson, France*, <sup>e</sup>*University of Patras, Greece*

**Abstract.** The Pan-European laboratory – Panlab – is based on federation of distributed testbeds that are interconnected, providing access to required platforms, networks and services for broad interoperability testing and enabling the trial and evaluation of service concepts, technologies, system solutions and business models. In this context a testbed federation is the interconnection of two or more independent testbeds for the temporary creation of a richer environment for testing and experimentation, and for the increased multilateral benefit of the users of the individual independent testbeds. The technical infrastructure that supports the federation is based on a web service interface through which available testing resources can be queried, provisioned and controlled. Descriptions of the available resources are stored in a repository, and a processing engine is able to identify, locate and provision the requested testing infrastructure, based on the testing users' requirements, in order to dynamically create the required testing environment. The concept is implemented using a gateway approach at the border of each federated testbed. Each testbed is an independent administrative domain and implements a reference point specification in its gateway.

**Keywords.** Testbed, Federation, Federated Testbeds, Reference Point, Resource Control, Panlab, Teagle

## Introduction

The Pan-European laboratory – Panlab – is based on federation of distributed testbeds that are interconnected, providing access to required platforms, networks and services for broad interoperability testing and enabling the trial and evaluation of service concepts, technologies, system solutions and business models. In this context a testbed federation is the interconnection of two or more independent testbeds for the temporary creation of a richer environment for testing and experimentation, and for the increased multilateral benefit of the users of the individual independent testbeds.

Rapid development of Information and Communications Technologies (ICT) in the last decades has been ensured by significant efforts performed by the corresponding research community world-wide. Both theoretical research, e.g. based on mathematical analysis and simulations, and research based on experiments contributed significantly to the recent technological developments. Meanwhile, the complexity of ICT systems, for example networks, devices, applied methods and algorithms, has increased in order to ensure their proper operation. Therefore, to be able to develop and assess new concepts and achievements in complex environments, researchers and engineers are increasingly looking for opportunities to implement their concepts in testing systems



and in this way obtain quickly the solutions and optimizations that can be implemented in production systems. Thus, with the recent developments in the ICT area, the necessity for experimental research carried out in the form of large scale experiments and testing is significantly growing among the research and engineering communities. On the other hand, in the ICT area, as the main driver of common global developments, there is a need to ensure world-wide experiments and testing ensuring that developed solutions can be applied anywhere as well as empowering wider groups of researchers to have access to various experimental and testing facilities at a global scale.

In order to meet the requirements – enlarged experimental opportunities and remote testing – the Panlab concept [1] has been created in Europe to form a mechanism that enables early-phase testing and interoperability trials as widely and deeply through the layers and players of telecommunications, as possible. In order to boost European testing, we must have the means to dynamically provision testbeds according to customer requests. This can be achieved by means of a new functionality capable of composing, managing and refining testbed resources. This constitutes the primary objective of the pan-European laboratory for networks and services, which implements the Panlab concept [2].

The entire mechanism, the rules and procedures of how to achieve the effective testing collaboration, have been developed in the Panlab project at a high abstraction level. The considered mechanisms include legal and operational requirements on the Panlab concept as well as requirements on technical infrastructure to be established in order to realize the Panlab concept, as summarized in [3].

In this paper, we describe how the technical infrastructure controls the distributed testing resources in the Pan-European laboratory federation in order to dynamically create appropriate testing environments. It is organized as follows: First, Sec. 1 presents the Panlab concept, including an introduction of the Panlab roles and the integration of testbeds, as well as Teagle as the main collaboration tool. Sec. 2 presents the seven operational stages of the Panlab concept. The enabler functionalities of the technical infrastructure of the federation are described in Sec. 3, while Sec. 4 describes the reference points for partitioning the identified responsibilities of the different administrative domains. Sec. 5 presents the main components responsible for configuration, control and interconnection of components across multiple administrative domains. Finally in Sec. 6, the conclusions of the paper are presented.

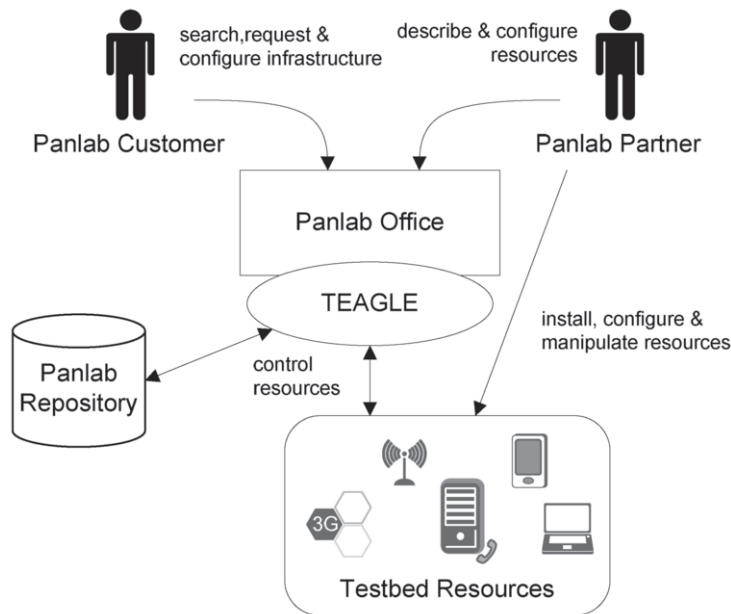
## 1. General Panlab Concept

### 1.1. Components and Roles within the Panlab Concept

The Panlab infrastructure has to ensure interconnection of different distributed testbeds to provide services to its customers, in order to do various kind of testing. Coordination of the testing activities, for example infrastructure configuration, ensuring necessary interconnection of customers and testbeds, and the overall control and maintenance of the environment, is ensured by the so-called Panlab Office. Thus, we can outline the following main roles of the Panlab concept, as is presented in figure 1:

- Panlab Partner – an entity that participates in Panlab activities by providing infrastructural elements and services necessary to provide testing services. Panlab partners are connected to the Panlab Office for offering functionality to the customers.

- Panlab customer – an entity that uses a service provided by the Panlab office, typically to carry out R&D activities, implement and evaluate new technologies, products, or services, benefiting from Panlab testing offerings.
- Panlab Office – an entity that realizes a brokering service for the test facilities, coordinating and supporting the Panlab organization. It is responsible for coordinating the provision of the testing infrastructures and services, partly by using tools and, when possible, web interfaces, but also directly coordinating the Panlab Partner test-sites and the communication path between them and customers.



**Figure 1.** Panlab components and roles


In addition to the roles defined above, the Panlab architecture relies on several other architectural components that will be introduced in the course of the following sections. Among those are the Panlab search and composition engine “Teagle” and a Panlab repository that stores testbed descriptions and testing results.

In a first phase of Panlab operation, various operational steps, ensuring the creation and realization of testing projects, are executed manually by personnel of the Panlab Office and involved partners and customers. Thus, the testbed metadata held in the Panlab repository is entered manually as well as testing configurations, etc. In order to achieve a near to full automation of all Panlab related processes, the so-called Teagle tool will be developed, in order to offer, (among other functionalities) an online form where the testbed representatives can enter the relevant data describing the testbed [4] and its resources, Panlab customers may then search the Panlab repository to find suitable resources needed for doing their tests.

### 1.2. Integration of Testbeds

In a general use case applied to the Panlab architecture, two or more Panlab components have to be interconnected in order to ensure realization of a particular

testing project (figure 2). The established connection between the different testbeds must serve a specific objective, i.e. it must serve the interactions between (i) applications, (ii) signalling control and service support, or (iii) user data transport. Connections can be requested to serve one or more of these objectives.



**Figure 2.** Gateway concept

One of the tasks of the gateways is to match property requirements to the connectivity service with the properties of the available connectivity. In many cases, Virtual Overlay Network (VON) technologies can be used to connect resources and sites with a common set of connectivity properties. Especially Virtual Private Network (VPN) or Virtual LAN (VLAN) technologies are well established means to create a logically dedicated network for a specific purpose. From the federation point of view, a logical connectivity support function must be implemented, which is able to control the gateways, located at the edge of the individual testbeds, to establish the requested connections to the peer site or sites. Thus, interconnection of Panlab components is ensured by establishing connections among gateways of respective individual components representing separated administrative domains, where all other interconnection functions remain under control of these components.

### 1.3. Teagle Tool providing Search & Match Functionality

As mentioned before, the Panlab architecture relies on a tool called “Teagle” which is a web instance that provides the means for a customer to express the testing needs and get feedback on where, how and when testing can take place. Teagle enables finding a suitable site for one’s testing needs within a database of partner testbeds (the Panlab repository). The objective of Teagle is to manage the complete set-up of a desired infrastructure. This includes necessary resource reservations and interconnection of elements to serve a specific testing need. Thus, by using Teagle, the customers can select desired technologies and features for a test configuration to be set up. Functionalities of Teagle tools are described in [3] and further specific details can be found in the documents available at [1].

## 2. Operational Stages of the Panlab Concept

This section aims at exemplifying key operational aspects as they emerge from the overall Panlab concept and use cases. To this end, we have identified seven operational stages, the purpose of which is to provide the context that specific operations are defined and executed in:

1. **Customer Interaction:** This stage comprises the interactions between a customer requesting the provisioning of a testbed and the corresponding tests to be carried out when the testbed has been provisioned. This interaction takes place through the Teagle and may formally end with a Service Level Agreement (SLA). This SLA can then be used by Teagle to analyse customer requests and then to find the proper testbeds in the Panlab repository that match the customer request. Then this SLA becomes a binding contract.
2. **Testbed Discovery:** This stage may well be seen as a precursor of the previous stage or as a distinct part of customer interaction. In the former case a customer before interacting with Teagle may search on his own the testbed repository in order to find for himself what is available in Panlab. This means that proper Graphical User Interfaces (GUI) guide the customer through the Panlab offerings and/or provide examples of the available technologies. In the latter case, Teagle simply searches through the Testbed repository in order to find, with or without customer's collaboration, suitable technologies.
3. **Testbed Provisioning:** This stage starts when the customer and the Panlab office have both agreed on the SLA, and Teagle can now initiate the provisioning of the testbed environment before it is delivered for use to the customer. This is an entirely Panlab office responsibility. Provisioning is carried out through a number of interfaces implemented by the Panlab control architectural elements. One major aspect of testbed provisioning is the wide variety of configuration operations that may need to be performed on testbed components. Therefore, configuration critically depends on control interfaces available in the testbed devices and they raise problems of interoperability for controlling these devices. Resolving them requires that the testbed components implement open or standard interfaces for their configuration. In case of proprietary control interfaces, there must be functionality in place that performs mappings of configuration operations on to the proprietary control interfaces. This description gives rise to one of the main architectural components of the Panlab architecture, namely, the PTM. Finally, all these operations are performed in a secure environment.
4. **Usage and Management of Testbed by the Customer:** When entering this stage, the testbed has been configured, deployed and handed over to the customer and his users or test suites for its actual use. This means that a number of management interfaces may be exported to the customer so that he can further tune the testbed internal operations according to user needs or test requirements. At this stage the testbed operations pertaining to specified tests and users' requirements, become the responsibility of the customer whereas the overall welfare (security, fault tolerance, SLA conformance etc.) remains at Panlab's office responsibility. Any additional operation that falls outside the scope of the contracted SLA must be re-negotiated and re-provisioned.
5. **Monitoring and Collection of Test Data:** Monitoring services and collection of test data represent an important part of the overall Panlab services as it

provides the means to process and analyse the behaviour of the product for which the customer has requested the testbed in the first place. Monitoring in Panlab may be carried out either by the customer by deploying monitoring mechanisms customised for his proprietary tests, or on behalf of the customer when he needs common monitoring mechanisms e.g. packet traces, sampling etc. In the former case, we assume that monitoring functionality is part of the components contributed to the testbed by the customer whereas, in the latter case, the monitoring mechanisms form an integral part of the testbed offering and as such they also undergo deployment and/or configuration during the provisioning stage. The same mechanisms may also be used for other activities, e.g. Quality Assurance or SLA conformance. Finally, there are proper interfaces, protocols and resources, for the collection and transport of monitoring data to repositories either in the customer or Panlab premises, so that they may become available for further processing and analysis.

6. **Processing and Accessing Test Data:** After completing the tests, data should be collected and stored in repositories for further processing by the customer or on behalf of the customer according to his needs. Access to these data may be controlled by certain policies. To this end, a customer may decide to make the collected data publicly available or keep them for his own purposes. In due course, we expect that the collected data will become a valuable asset of the Panlab office and as such it is envisioned that these data should become available to other customers even if they do not require the deployment of a testbed. This involves the definition of common formats to read the data as well as tools for carrying out analysis e.g. statistical, anomaly detection etc. Accordingly, the Panlab office through Teagle may consider this stage as an additional service to testbed provisioning and a distinct service on its own.
7. **Quality Assurance:** This stage comprises a series of Panlab functionalities running at the background of any testbed operations and aiming at guaranteeing the welfare operation of the testbed infrastructures and conformance to contracts by both sides, namely, customer and Panlab. Such functionality ranges from security to monitoring as well as proofs of conformance to contract terms.

### 3. Federation Enablers

For federating different testbeds residing in autonomous administrative domains and provide composite infrastructures using resources across the boundaries of the domains, three major blocks of functionality need to be provided:

- Resource Description,
- Service Exposure, and
- Service Composition.

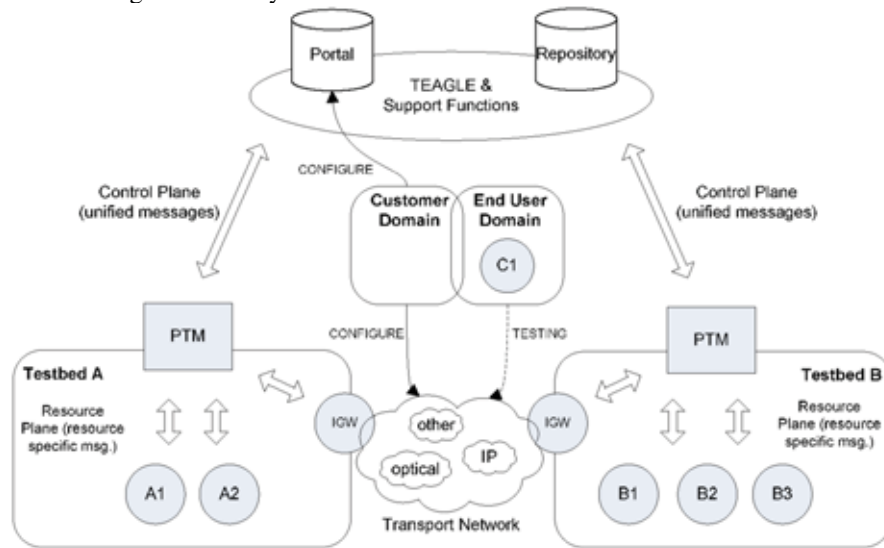
These functionalities are considered the enablers for the federation, since they represent essential capabilities of the technical infrastructure of the federation.

### 3.1. Resource Description

Whatever is offered by a testbed towards the federation must be described according to a common information model in order to allow Teagle to understand the functionality of resources, offer them towards external users and provide compositions of distributed functionalities. This means that Panlab Partner resource offerings become services to be consumed by the customer. The services need to be described uniformly to allow Teagle as a broker residing between the customer and the resource, to match requests to concrete functionalities. The TMF SID [8] provides a good starting point for a common information model that might need to be extended to meet the Panlab requirements and to deal with the high heterogeneity of Panlab Partner testbed resources.

### 3.2. Service Exposure

The basic service exposure mechanism foreseen for the Panlab architecture is shown in figure 3. A1 and A2 are functionalities that are offered by testbed A towards the federation. As is described above, the PTM provides the mapping from federation level commands to resource specific commands and the IGW is responsible for providing and controlling connectivity to other administrative domains.



**Figure 3.** Service exposure

The PTM exposes the functionalities A1 and A2 as services to the federation. As an example consider A1 to be an application server and A2 a web service that allows to send IMS instant messages. The PTM would then expose a service that allows the set up and configuration of A1. Also the PTM would expose the web service which simply results in exposing the endpoint of A2.

The testbeds shall encapsulate functionality in atomic building blocks that can then be used in a combinational manner across the entire federation. This means that functionalities offered by the individual testbeds are represented as atomic components that act as building blocks during the composition of a testbed. In this context, a testbed instance acts as a container of the selected building blocks. Accordingly, different

combinations of building blocks may result in different testbed instances dictated by customer requests. This is in line with the general concept of Service Oriented Architectures (SOA) and shall ultimately lead to a service market. However, in contrast to open, “internet-style” service markets, the principle demonstrated here relies on centralized support functions and legal entities that ease the establishment of business relationships between multiple stakeholders. For example the negotiation of Non-Disclosure Agreements (NDA) and contacts is simplified through agreed templates and a trusted relationship between the Panlab Partner and the Panlab Office.

### 3.3. Service Composition

Teagle shall offer a service composition functionality that allows the orchestration of (atomic) building blocks offered by different testbeds. In order to do so, as stated above, it requires a solid description on what is available. A repository shall hold descriptions of the available resources in the federation and instructions on how to invoke them. Teagle displays the content to the customer, provides user accounts and offers a search tool for browsing federation resources. Once the customer has looked up and identified interesting functionality to be provided by the federation, a composition tool shall provide a workflow that defines how to provision the desired virtual environment as a composition of building blocks from different testbeds. In this regard several aspects are important:

- Pre- and post-conditions of building blocks,
- Timing across the entire workflow (which operation goes first, second, etc.),
- Dynamic probing of availability.

The field of service orchestration is still subject to extensive research boosted by the success of service oriented architectures. Panlab will make use of current state of the art technologies in this area and seeks to contribute with scientific results to this important field of research. However, first Panlab implementation stages are foreseen to rely on many manual processes while fully automated service composition remains the grand vision to be achieved in later concept implementation stages.

## 4. Reference points

The infrastructure supports network agnostic service provisioning by separating the testing service logic from the underlying networking infrastructure. It identifies and defines the necessary interfaces for resource negotiation at the technology and administrative domain borders. It defines the necessary messages at the federation level that are mapped via a Panlab Testbed Manager (PTM) to intra-testbed messages to automate provisioning (figure 4). The messages defined at the federation level are based on the New Generation Operations Systems and Software (NGOSS) [5] framework. Reference Points (RP) mark an interface specification between Panlab entities.

Each participating testbed implements at its border the necessary functionality in the PTM, to be able to receive and interpret control messages from the federation. The PTM is responsible for the clear separation between the mechanisms for services provisioning at the federation level, from the mechanisms needed to map these services onto the network infrastructure.



Furthermore, each participating testbed implements an Interworking Gateway (IWG) for the technologies it supports at the testing usage level (figure 5). The IWG covers only the protocols and technologies that are meaningful in the context of usage of its own components and resources.




Figure 4. Gateway concept implementation – PTM




Figure 5. Gateway concept implementation – IGW

The concept of reference points is essential for partitioning the responsibilities that each administrative domain has to fulfil. Each RP has to satisfy a number of well identified requirements that are related to the technical or business relationship manifested by each RP. The figure 6 below identifies the complete set of reference points that are related to control and usage of a resource inside a testbed as seen from outside a testbed (administrative domain). From this perspective the relevant actors are the federation domain, represented by Teagle, the End-User domain, the customer domain, as well as the neighbour testbed domain, represented by the shaded IGW in figure 6. For completeness the internal reference points from the PTM to testbed components are also illustrated. The main reason for this is that the PTM component maybe a generic component provided by the federation.

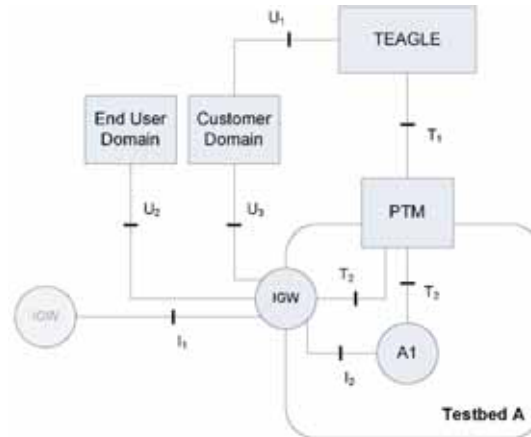


Figure 6. Reference points

**T1 - Teagle to PTM** – This RP refers to the common control plane that is established across the entire federation and is based on SOAP Web Services. Via this reference point, the configuration of resources (A1) is passed to the PTM.

**T2 - PTM to Testbed Components** – The PTM is responsible for managing all exposed components inside its testbed. It does so by applying a device driver concept. A device driver for a component is usually implemented by the testbed provider according to the functionality the component exposes to Teagle. Device drivers are



translating the Teagle operations and executing the operation on the associated testbed component (A1). This is an internal RP.

**U1 - Customer to Teagle** – This RP provides customer access to Teagle services, such as searching for specific functionality, setting up a desired testbed, configuring the testbed for a specific test case, retrieving test results etc. This RP will be implemented as a web interface and is not used during the test execution phase.

**U2 - End-User Domain to IGW** – This RP provides End-User access to a configured testbed and the services that are subject to test. This RP is in the usage plane and is not used for configuration, control or management.

**U3 - Customer Domain to IGW** – This RP provides customer access to a configured testbed during the execution phase of a test. This is necessary when adjustments are needed to a provisioned testbed after its initial configuration.

**I1 - IGW to IGW** – This RP interconnects two neighbour testbeds at the usage plane. Over this RP usage testing data are transported. The implementation of this RP depends on the requested testbed.

**I2 - IGW to Testbed Components** – This RP interconnects the testing service component with the IGW. The implementation of this RP depends on the requested testbed. This is an internal RP.

Although the concept of RPs is defined for the administrative domain borders, i.e. between the testbeds, or between the testbeds and the federation, the internal RPs T2 and I2 are also identified, because the PTM and IGW can be generic components that may be provided by the federation infrastructure. This can be the case if for example T2 has to implement a standard SNMP interface and the IGW is a standard VPN endpoint.

## 5. Resource Configuration and Control

### 5.1. Panlab Testbed Manager

The most important component for the implementation of the infrastructure is the Panlab Testbed manager (PTM) that has been presented in the previous section. It is implementing the interactions at the control layer between the set-up and configuration requests by Teagle and the components in the testbed it manages. The PTM is therefore capable to interact with these functions in a web service client/server way and to adapt commands received as web services into control commands applicable to the testbed component. Control commands can be of 2 types:

- Generic commands such as “create entity” which is the command to instantiate and make resource reservation of the component to which the command is addressed.
- Specific commands such as “include profile data into data base” which is a command to a database to declare new users.

The PTM is also interacting with the IGW in order to prepare the connectivity with other components. It is part of the adaptation process of the initial web services command to issue the right control commands to the gateway to allow the interconnection between several components. This can include simple address set up but also pin-holing through a firewall, or parameter setting for other address translation operations.

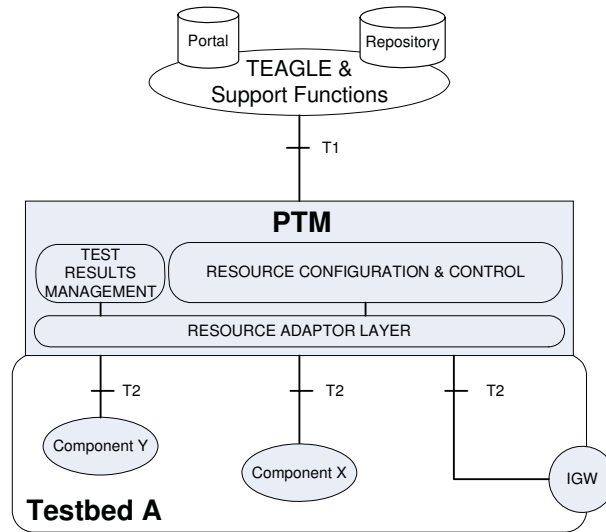


Figure 7. Panlab Testbed Manager

Testbed components may implement heterogeneous technologies at different levels. Although there are no requirements on specific technologies for testbed components, any testbed component must be individually controllable by the PTM. This means that, apart from the functionalities it provides at the usage plane, any testbed component needs to:

- Be individually identifiable, addressable and reachable from the PTM
- Provide booking, provisioning and configuring means to control its operation at the usage plane level.
- Optionally provide monitoring capabilities both to inform the PTM about its current state and to deliver data on test results towards the PTM.

### 5.2. Interconnection gateway

The main requirement for interconnecting testbeds is to ensure that only the nodes configured as part of a testbed are allowed to communicate. Furthermore it is necessary to hide the internal network topology reducing complexity and allowing the partners to dynamically provision multiple overlay networks if so requested. This is the responsibility of the Interconnection Gateway (IGW).

The main purpose of the IGW is to interconnect testbeds and components inside the testbed with the peers that are part of the configuration. Technically seen this component is a border gateway function and policy enforcement point for each testbed. For the IP layer the IGW acts as a dynamically configurable entity implementing:

- L2 connection/isolation of testbed devices
- multi-VPN endpoint and link to central VPN concentrator
- firewall and filter
- application proxy and/or network address translation

## 6. Conclusions

The work presented here is a summary of the conceptual design of the technical infrastructure to prove that federation is a model for the establishment of a long-term sustainable large scale and diverse testing infrastructure for telecommunications technologies, services and applications in Europe. Beyond the demonstration of the technical feasibility of the service related mechanisms described in this paper, the future work includes research towards the fully automated provisioning of composite testbeds across the whole infrastructure.

To support the long-term sustainability of the federation, future work will develop and elaborate on the mechanisms to combine and accommodate potential clean slate approaches. In particular the work is focused on the architectural requirements to facilitate the separation of the “provisioning platform” from the underlying infrastructure as a means to accommodate approaches based on different architectural mindsets.

## Acknowledgements

Parts of this work have been developed within the Panlab Specific Support Action and have received funding by the European Commission’s Sixth Framework Programme. The authors of this paper would like to thank the Panlab Consortium partners for the good collaboration. Also parts of the work presented here are early achievements of the Project PII which receives funding from the European Commission’s Seventh Framework Programme. The mechanisms described are being implemented in the course of the PII project. The authors thank the PII partners for the fruitful discussions and cooperation. The work of PII is aligned with the objectives of the Future Internet Research and Experimentation (FIRE) initiative [10].

## References

- [1] [www.panlab.net](http://www.panlab.net) – website of Panlab and PII European projects, supported by the European Commission in framework programmes FP6 (2001-2006) and FP7 (2007-2013).
- [2] A Gavras, H. Brüggemann, D. Witaszek, K. Sunell, J. Jimenez, “Pan European Laboratory for Next Generation Networks and Services”, TridentCom 2006 2nd International IEEE/Create-Net Conference on Testbeds and Research Infrastructures for the Development of Networks and Communities, Barcelona, Spain, 1-3 March 2006
- [3] S. Wahle, A. Gavras, F. Gouveia, H. Hrasnica, T. Magedanz, “Network Domain Federation – Infrastructure for Federated Testbeds”, NEM Summit 2008, Saint-Malo, France, 13-15 October 2008
- [4] The Panlab Repository website, <http://www.panlab.net/testbed-repository.html>
- [5] TeleManagement Forum NGOSS Release 6.0, <http://www.tmforum.org/page31331.aspx>
- [6] FP6 Panlab Deliverable 2.3, User and operations manual
- [7] Fraunhofer FOKUS Open SOA Telco Playground website, [www.opensoaplayground.org](http://www.opensoaplayground.org)
- [8] The TeleManagement Forum Shared Information Data, <http://www.tmforum.org/browse.aspx?catID=2008>
- [9] TeleManagement Forum, <http://www.tmf.org>
- [10] A. Gavras, A. Karila, S. Fdida, M. May, M. Potts, Future internet research and experimentation: the FIRE initiative, in ACM SIGCOMM Computer Communication Review, ISSN:0146-4833, Volume 37, Issue 3 (July 2007), pages 89-92

# The Trilogy Architecture for the Future Internet

Louise BURNES<sup>a</sup>, Philip EARDLEY<sup>a</sup>, Robert HANCOCK<sup>b</sup>

<sup>a</sup>*BT Innovate, Martlesham Heath, Ipswich, UK*

<sup>b</sup>*Roke Manor Research Limited, Romsey, UK*

**Abstract.** Socio-economic aspects are not intrinsic to the current Internet architecture and so they are handled extrinsically. This has led to increasing distortions and stresses; two examples are inter-domain scaling problems (a symptom of the way multihoming and traffic engineering are handled) and deep packet inspection (a symptom of the lack of resource accountability). The Trilogy architecture jointly integrates both the technical and socio-economic aspects into a single solution: it is thus *designed for tussle*. A Future Internet that follows the Trilogy vision should automatically be able to adapt to the changes in society's demands on the Internet as they occur without requiring permanent redesign.

**Keywords.** future Internet, architecture, resource accountability, resource pooling, multipath TCP, Trilogy project

## 1. Introduction

The current Internet is a remarkable phenomenon, not only technically, but from an economic and social perspective as well. It has grown to become the network of choice for a huge variety of distributed applications, starting from email and file transfer, through a whole range of business to consumer and business to business e-commerce systems, all the way to online news and entertainment - for many people, displacing traditional print and broadcast media. And it has achieved all this almost entirely by accident: the happy decision by the original Internet designers to implement a minimal, best-efforts data transfer capability has allowed the Internet to incorporate technological advances and adapt to new application requirements with extraordinary speed and economy.

Where the Internet has been less successful - and this has its roots in the restricted priorities of the designers [Clark88] rather than fundamental flaws in the design - is in accommodating the conflicting economic interests of its participants, both between different operators of the network infrastructure and between the users and providers of network services. This problem was identified in [Clark05] as the "tussle in cyberspace", and arises increasingly often as the demands on the Internet increase.

In the routing area, inter-domain architectures using BGP are tractable while the domain level topology away from the edge is small in scale and organisationally largely hierarchical. However, as the scale increases and topologies become more strongly meshed (for example, because of both direct peering between edge providers

and end-site multihoming), the lack of economic control over provider-provider interactions becomes more apparent. This lack shows itself technically as growth in router tables sizes and churn rates, and the inability of individual operators to manage this load without harming end to end reachability.

The Internet copes well with the demands of best efforts applications, where the network capacity can be engineered to match reasonable expectations on long-term average data transfer requirements. In recent years, these assumptions have been less and less valid. Dependency on the Internet for more demanding applications (voice, physical infrastructure control, critical business operations) has mushroomed, at the same time as peer to peer traffic has demonstrated its ability to absorb all available bandwidth and more. While integrated services approaches can solve these problems in special scenarios, they cannot scale to general purpose Internet deployments because of the technical, and even more importantly administrative, complexity of marshalling the resources of all the application user and network providers involved.

The Trilogy project in the EU 7th Framework programme has been established to address these issues for the Future Internet. The project is based on the view that the major architectural concept behind the Internet - that of a simple, ubiquitous, transparent data delivery infrastructure, which can accommodate innovation at the link and application levels - remains sound today and for the future. The focus is therefore on the network and transport layers of the standard protocol stack, but taking an economic and social perspective as well as pure engineering research. In fact, one goal of the project is to develop an architecture that embeds both technical functions and the socio-economic forces that influence them; this paper represents the first stage in the development of this architecture. Our vision is that the Future Internet will thereby automatically adapt to the changes in society's demands on the Internet as they occur without requiring permanent redesign.

In this paper we make three contributions:

- We hypothesise a baseline Trilogy architecture, which is comparable in scope to the current Internet network and transport layers, but with a subtly different internal structure. In particular, we divide the functions involving the networking infrastructure into two planes, for reachability and forwarding, and distinguish them from the transport services to which the network is totally transparent. (Section 2)
- We propose an accountability framework that is based on congestion volume. Accountability enables a rational basis for sharing resources amongst users on an Internet that is a playground for competing users, applications and businesses. (Section 3)
- We propose the concept of resource pooling, which enables separate network resources to behave like a single large pooled resource, and we propose a technique to achieve this: an end-to-end multipath transport protocol. Resource pooling enables better resilience and efficiency. (Section 4)

This paper includes some selected key pieces of our architecture (more extensive details are elsewhere). We stress that this is our initial architecture; it will change as we validate it through our own activities and through feedback from other researchers and network designers – hence we highly welcome comments on the paper.

## 2. Baseline Trilogy Architecture

A fundamental assumption of the architecture is that it is based on a minimal packet delivery service. The ideal case is that packets are entirely self-describing, meaning that other concepts such as connections, flows or sessions are higher level constructs, invisible to the delivery service, and the network delivers each packet independently of every other.

We decompose the packet delivery functionality into two parts:

- the reachability plane: responsible for hop-by-hop outgoing link selection and hence enabling network-wide reachability
- the forwarding plane: responsible for deciding how the transmission resource on each link is apportioned between packets.

Distinct from the packet delivery service, we identify the functions that are implemented in a pure end-to-end fashion:

- transport services: functions, such as reliability, flow control or message framing, that are totally invisible to the packet delivery service.




Figure 1: Trilogy baseline architecture

The reachability and forwarding planes are separate; together they achieve the packet delivery service. The key identifier space is the destination locator, which is handled in detail only by the reachability plane. The forwarding plane treats locators as opaque tags which only have to be tested for equality, in order to test for path consistency. The transport services use endpoint identifiers that label the communicating parties, but these are totally independent of the locators used in the reachability plane.

In the rest of this section we briefly describe these three elements. The accountability framework and end-to-end multipath (Sections 3 & 4) are refinements of this baseline architecture. [Del3] contains further details, including other refinements of the baseline architecture and potential extensions

### *2.1. The Reachability Plane*

The reachability plane is responsible for hop-by-hop outgoing link selection. We initially take the information needed to route each packet to be the destination address or locator. Every allocated locator is reachable by default, and functions such as (D)DoS protection must be implemented at the receiver end system. We assume that the reachability service is implemented by a set of autonomous, interconnected administrations or domains. The internal structure of each domain may be non-trivial, but is externally invisible: all that is exposed is information about which locator spaces are reachable and under what circumstances.

Open questions include:

- Are locators from a single global namespace? This approach has the greatest engineering simplicity, and indeed is the only approach totally compatible with the packets being fully self-describing. However, it is also the root cause of many of the Internet stresses, in that it immediately places all network participants (at least, packet sources and destinations if not forwarding infrastructure) into a single “tussle space”. The Loc/ID work in the IRTF’S Routing Research Group [RRG] is essentially exploring trade-offs of possible relaxations.
- What additional identifier spaces are used to manage the global topology, and how? We seek to minimise the use of locators in describing topology, to avoid extending that tussle space into inter-network operations. Inter-domain topology should ideally be described in terms of different identifiers, such as the Autonomous System (AS) number of BGP. Note that in interior routing protocols, router identifiers are typically IP addresses, but there is no actual need to couple these identifiers to the locators in the traffic being routed.
- What reachability information or control is shared, beyond the locator spaces themselves? It is notable that much of BGP operation is concerned with traffic engineering (for load balancing), which is an area where the project is exploring solutions in a different part of the architecture (see Section 4).
- What other packet information influences the path? Note that the path itself is not visible in the packet format, because we believe that the end nodes should be isolated from the network infrastructure. As well as the destination locator, we allow that other identifiers visible at the packet level may influence routing explicitly (path selector bits and service class identifiers). The reason is to enhance path consistency between packets, which is an important property for endpoint-managed resource control algorithms.

### *2.2. The Forwarding Plane*

The forwarding plane decides how transmission resource on each link are apportioned between packets. Information about resource allocation along the path can either be implicit (delay or loss) or explicit (‘marking’ in the packet header); end systems either measure the implicit path properties or read the explicit information, and modify their sending rates in response. The allowed set of responses is left very open at this stage, and specific issues of accountability are discussed further in Section 3.

The fundamental packet delivery service is best-efforts: the network delivers packets to their destination, without guarantees on ordering or throughput. However, there is an implicit requirement that an end system generating a sequence of packets must be able to depend on path consistency, at least over a scale of a round trip time or more.

Looking at the information managed by the forwarding plane, we identify two major open issues:

- What is the level of path property information provided by the network? The optimum situation, consistent with the principle of the reachability plane (that each packet is self-describing), is that each packet is marked with a complete description of forwarding resources available on the path, but this could impose a significant per-packet overhead (in size and forwarding cost). More constrained encodings impose stronger requirements for path consistency and stability for resource control loops to be effective.
- How does the forwarding plane distinguish traffic types? We do, at a minimum, assume that any domain can define certain service classes with different forwarding performance, and that end systems can select between these by embedding information in their packets. However, it is not clear if these service classes can or should be globally defined, or whether they are agreed only at interconnection points – between end systems and networks, and between networks themselves.

### *2.3. The Transport Services*

Transport services are the means by which the packet delivery functions are actually exercised. Re-engineering of the standard transport services is therefore a main method to exploit the functions of the Trilogy architecture. They are implemented in a pure end-to-end fashion and define the communications service offered to applications (email, messaging, file sharing, multimedia, ...). Included are functions such as reliability, flow control or message framing, which are totally invisible to the packet delivery service. The identifier spaces involved are also totally separate from the reachability plane's locators; indeed a single node might use several different identifier families, and they may be implicit rather than explicit.

## **3. Accountability Framework**

### *3.1. From statistical multiplexing to resource accountability*

Since the earliest days of telecommunications operators have used statistical multiplexing to maximise the number of customers that can share a backhaul link. This reflects the fact that it is not economically viable to provide each user with an end-to-end dedicated link or circuit at the speed they require. Statistical multiplexing assumes that at any one time only a handful of users will be actively using any given link. This allows telecoms companies to share resources between users deep in the core and thus



take advantage of significant cost savings. The approach has worked well for phone networks. Of course there's a small chance that at a particular moment there isn't enough capacity and so the new phone call is blocked; hence phone operators have to balance the number of customers, their grade of service and the amount of capacity.

Packet switched networks are the logical extension of this drive for efficiency in the network. Statistical multiplexing is done packet-by-packet rather than call-by-call. TCP's job is to perform this multiplexing – the sharing of capacity amongst the users. Historically it has worked well, however it has proved inadequate with the rise of new sorts of application, peer-to-peer being perhaps the most prominent, but also other like voice and business critical VPNs. P2P, for example, undermines the assumption on which statistical multiplexing is built. Nowadays there are some users who download (and upload) content 24 hours a day, 7 days a week. With the growth in P2P TV this can only get worse.

A closer analysis [Briscoe07] reveals that there are actually several interconnected issues:

1. A wider range of users: there are now some users who consume orders of magnitude more bandwidth than others – far more extreme than when everyone just used 'interactive' applications like web browsing. Today less than 1% of customers can generate several 10s% of the traffic.

2. Different types of utility: The TCP algorithm assumes a particular utility function (how much a flow values a byte of information). TCP's utility has the following characteristics:

- Convexity: twice the bit rate has less than twice the utility
- Equality: all TCP flows have the same utility function
- Immediacy: what matters is the bit rate 'now' (within a round trip time); it's irrelevant what the bit rate was or what it will be.

However, P2P's utility function is very different from TCP's: what matters is how long it takes for the whole film to finish downloading.

3. Non-cooperativeness: Statistical muxing on the Internet also assumes that applications use the TCP algorithm to reduce their rate when there is congestion. But using TCP is voluntary – there are no laws about it! – although luckily (most) application writers build in the use of TCP. However, P2P applications open many TCP flows, typically tens, which of course squeeze the bandwidth available for applications like web browsers that open only one or a few flows. See left hand side of Fig 2.

After the above discussion, the reader may be wondering why they still get a reasonable web browsing experience. Why don't P2P applications get all the bandwidth by opening an ever greater number of flows? Why don't some applications use congestion control that's more aggressive than TCP? (After all, congestion control is a game of chicken; whoever blinks last wins the most bandwidth.) Why don't we have a Tragedy of the Commons?

The reason is that ISPs have introduced DPI. Deep Packet Inspection boxes control the balance of resources between users and between applications. The basic idea is that a device "inspects" each packet and then takes some appropriate action, for example limiting the fraction of bandwidth available for P2P. This leaves more bandwidth available for 'interactive' applications – see centre of Fig 2. ISPs also count how much

traffic each user transfers and cap it according to their fair usage policy. Incidentally, from an operator perspective DPI also allows investment in a capacity upgrade knowing that it won't be "wasted" on P2P.

However, DPI has drawbacks, not least that P2P applications try to disguise themselves as interactive ones, so the DPI has to be cleverer; P2P applications then disguise themselves further (eg using encryption) and then we're in an arms race. So DPI is really a sticking plaster. What is needed is a proper architectural solution for deciding how to allocate resources.




Figure 2: Resource sharing.

Left: TCP sharing (base case). Middle: Limit volume (DPI). Right: Limit congestion volume (this paper)

### 3.2. Accountability framework for resource allocation

We believe that the architectural answer is "accountability for congestion volume". What does this mean? Congestion is what happens when too much traffic meets too little capacity. It is the adverse impact that your traffic has on others (and vice-versa), ie its externality. Congestion volume is simply the congestion integrated over time and data, so that all the congestion you cause counts (however many flows you create, whatever route, whatever the activity factor etc). For creating an accountability framework [Argyaki07], it is important to know the "rest-of-path congestion" [Laskowski06], as explained below; at present, congestion is only known about on an end-to-end basis (by TCP) or at a single point in the network.

We are developing a framework [Briscoe08] containing the elements required to achieve accountability based on congestion volume:

1. congestion information on each packet. This has already been standardised as ECN and is implemented on all routers.
2. a new mechanism such as re-feedback that gives the ability of any point in the network to calculate the congestion on the rest of the path. 'Rest of path congestion' is the total congestion suffered between a particular point in the network and the destination.
3. a policer at the first ingress, to catch those trying to cause more congestion than they're allowed under their contract. [Jacquet08]
4. a dropper at the last egress, to catch those trying to cheat by under-declaring their rest-of-path congestion

5. border gateways, which count the congestion volume in each direction between two adjacent ISPs; this is a bulk measurement (not per flow). There would be inter-ISP contracts, similar to those today for bandwidth.

6. weighted congestion control. The end host runs an algorithm that reacts to congestion but weighted according to the priority of a particular flow. This achieves true end-to-end QoS.

Some immediate implications of this approach:

1. We envisage that a certain amount of congestion volume would form part of the broadband contract (fair usage policy). The end user will almost certainly not be charged for every single byte of congestion volume that their traffic causes.

2. Software on the user's computer will automatically prioritise their traffic, so that their most important applications tend to use up their congestion volume allowance. If there is no congestion on the path, then you can go as fast as you like, since you don't affect anyone else. If there is congestion, then the user should choose their priorities, since only the user really understands the importance of a particular data flow (DPI makes an educated guess, but it may be wrong). See right hand side of Fig 2.

3. It also enables other types of utility to be taken into account. For instance, a mobile terminal might want to save battery power by transmitting fast in short bursts.

4. It enables everyone to get a better experience, compared to DPI and volume capping. The lower right picture shows that both the interactive applications run faster and the P2P downloads finish earlier.

5. Visibility of rest-of-path congestion enables networks to traffic engineer based on how much congestion there is in other networks. This is analogous to an in-car navigation system learning about hold-ups on the rest of your planned route, and so being able to recommend a diversion.

6. The above points emphasise the wide-ranging impact of the accountability framework; it is a mistake to see it as about "just saving operators money".

#### **4. Resource Pooling and Multi-path Transport**

The concept of resource pooling is that separate network resources behave like a single large, pooled resource. Indeed the concept underlies the general principle that resilience should be achieved through redundancy and diversity which led to the use of packet switching. In today's environment, we would like to expand resource pooling across multiple links, because the Internet is much more interconnected than in the past. We are investigating an end-to-end multipath-capable transport protocol as a means to achieve resource pooling [Wischik08]. The concept is simple: enable a single logical connection to use multiple paths simultaneously. The main motivation behind this is to improve the resilience of the network through the use of redundancy.

In the recent discussions, for example [SHIM6], it has been assumed that the sender and/or receiver have multiple addresses associated with different network access technologies or providers. After the initial handshake, the sender and receiver exchange IP addresses and sub-flows can be created using different combinations of source and

destination address. Packets that are lost from one sub-flow may be re-transmitted over any of the other sub-flows.

Reasons for the recent renewed interest [Handley08] may include the fact that recent theoretical work has studied how the congestion response should be managed in such a multi-flow environment [Kelly05], [Massoulie07]. This shows that the congestion response of each sub-flow should be coupled together. One implication is that whilst rate (strictly the TCP window) increase is as normal, decrease is more aggressive. This has the result of allocating rates efficiently between the sub-flows. The correct response also ensures network stability. This is important when a path fails or appears; then data should not be moved suddenly from the failed path (or onto the new path), but instead the window should gradually increase over several round trips, whilst communication can meanwhile continue along existing paths.

The primary benefit of such a scheme is improved resilience in the case of failure. However the benefits are broader because such a mechanism also makes the network better able to handle localized surges in traffic and maximizes the utilization of the network [Handley08]. It automatically achieves a load balanced network. These features arise because multi-path transport effectively pools the network's resources. The upper part of Fig 3 shows how gains are made if resources – bandwidth – can be shared in a single pipe. This is realized today with packet switching. The lower part of Fig 3 shows how this concept is extended to a multi-path environment.




Figure 3: Resource pooling. over single or multiple bit pipes

However, it is important to consider how end-to-end multipath resource pooling will interact (and possibly conflict) with other resource pooling mechanisms that are already in use today. For example, networks may multi-home for resilience and then traffic engineer over these multiple links to utilize the resources properly. The mechanisms used today (essentially BGP routing) unfortunately lead to more scalability and churn problems for the routing system, which decreases the reliability of the whole network, driving more users towards multi-homing; a vicious circle. Peer-to-peer networks also attempt resource pooling in order to maximize their performance –

pooling upstream capacity and also pooling data availability by spreading data over multiple unreliable servers. Content delivery networks also attempt resource pooling, spreading load between multiple servers. Unfortunately the resource pooling of these different entities may be in conflict, for example an ISP lowest cost path choice may well be different from the peer-to-peer user's high bandwidth path choice. Also today we see a much greater range of applications all competing for the same resources – simple data transfers, long-lived bulk flows and voice all co-exist. Users are competing ever more aggressively for a share (fair or otherwise!) of the resources. The costs of this type of conflict can be arbitrarily high [Acemoglu07][Roughgarden02].

Architecturally, it seems that control of multi-path at the transport layer is optimal. When a network operator decides to re-route traffic to achieve better load balancing, the sudden change in traffic patterns could lead to congestion elsewhere in the network (and may in turn lead to traffic engineering downstream attempting to force the traffic back towards the original route). This suggests that control of multipath below the transport layer is too low to ensure network stability and safety. But multi-path management could be offered generically to many applications; having a standardized well understood mechanism may go some way to ensuring that conflicts can be managed. Hence having the functionality in the transport layer seems best.

However, we consider that the network providers need a way to influence the path choice taken – there needs to be some kind of feedback from the network to the users to ensure that the economic stability of the Internet is not lost. One initial idea for how this might be achieved is by the network provider adding ECN congestion marks [Wischik08]; lots of marks will encourage traffic to move away from the non-preferred, expensive link. On the other hand, it seems likely that a multi-path capability that resides with end-users will foster better competition between providers.

There are still outstanding architectural questions:

- This type of resource pooling is most effective for data and bulk transport; it is much less suited to jitter-sensitive applications.
- The end points need to be able to use multiple paths – what is the best mechanism? Provider aggregatable addresses have address management issues which are non-trivial if the access network is multi-homed rather than the end host.
- How many paths do the hosts need to access? The tentative answer is that just a few paths are needed [Mitzenmacher01] – provided they are the right paths [PGB08]! How much benefit could be gained if these were known to be disjoint?

## 5. Conclusions and next steps

### 5.1. Conclusions

The original Internet was an academic research network that has proved immensely successful – so successful that managing its rapid growth has been much more important than finishing the research! But increasingly the pressures on the Internet's growth can be attributed to the missing pieces. The missing pieces are well known, and it is hardly surprising that they were not considered critical to the early Internet. The early Internet was not composed of participants with different economic interests; the degree of interconnectivity was low as physical resources were so expensive and scarce that the concepts of having a large number of paths and multiple points of connection were almost unthinkable. The world is a different place today.

We believe that the basic Internet architecture is a very good starting point. Some concepts, like connectionless packet switching, are as good today as ever. Some concepts, resource pooling for instance, need extending and repositioning within the architecture to cover the wider range of resources that are available today. And some concepts, such as resource accountability, need to be added in from scratch. Overall we strive to provide a technical solution that should be able to respond to the changes in society, and that allows different economic models representing different business regimes and indeed different societies to co-exist within the single global network.

We have presented an architecture based on design for tussle. It is radical in that the architecture is in many ways a small, yet defined, step from the current architecture. We have identified the need for separation of reachability (path discovery) and forwarding (path resource management), and also that end hosts and routers should be involved in both processes in a coordinated manner – rather than in the confrontational manner of today. We believe that the addition of accountability should improve user-to-user interactions, network-to-user interactions and network-to-network interactions by enabling all interested parties to actually understand the global network behaviour and understand how their actions influence this behaviour.

### 5.2. Next steps

As is clear from our open questions, there is still much that needs to be done. We hope that exposing the architecture at this early stage will help foster debate. Meaningful validation of architectural work is always a challenge and we believe that we have reached a point where further progress will now depend on evaluation of concrete case studies or new proposals. We are developing candidate technical proposals in the areas of reachability and resource control, and hope to study their interactions practically. We will evaluate these both for architectural compatibility on the one hand, and simplicity and performance in realistic environments on the other. This combination of engineering evaluation, mathematical analysis, and simulation will be used to refine both the solutions and the architecture within which they fit.

Because the architecture is close to the current system, we hope that migration is plausible. The 'multipath TCP' protocol is very similar to existing transport protocols and so does not require any changes to the network. It depends on end hosts deploying

software pair-wise, and since the end systems directly benefit this migration is very plausible. Migration towards the accountability framework is subject to on-going study.

One of our key motivations is try and incorporate flexibility, so that the architecture can adapt for local business and operational needs. Our aim is to ensure that the future network can be more tolerant of the demands of society - by adding in design for tussle, specifically resource accountability, we hope this architecture will be valid for another 30 years.

## References

- [Acemoglu07] Acemoglu, R. Johari, and A. Ozdaglar, "Partially optimal routing. IEEE Journal of selected areas in communications", 2007
- [Argyaki07] K Argyraki, P Maniatis, O Irzak, S Ashish & S Shenker, "Loss and Delay Accountability for the Internet," In Proc. IEEE ICNP'07 (Oct 2007)
- [Briscoe07] B Briscoe, T Moncaster & L Burness, "Problem Statement: We Don't Have To Do Fairness Ourselves" IETF I-D draft-briscoe-tsvwg-relax-fairness-00.txt (work in progress, Nov 2007)
- [Briscoe08] Bob Briscoe, A Jacquet, T Moncaster & A Smith, "Re-ECN: Adding Accountability for Causing Congestion to TCP/IP" IETF I-D draft-briscoe-tsvwg-re-ecn-tcp-05.txt (work in progress, Jan 2008)
- [Clark88] D. Clark, "The design philosophy of the Darpa Internet protocols," In Proc. ACM SIGCOMM, Vancouver, BC, Canada, Sept. 1988.
- [Clark05] D. Clark, J. Wroclawski, K. Sollins, R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", IEEE/ACM Transactions on Networking, 13(3), p. 462-475, June 2005.
- [Del3] Trilogy Project Initial overall architecture Report, available from <http://trilogy-project.org/publications/deliverables.html> , August 2008
- [FORCES] <http://www.ietf.org/html.charters/forces-charter.html>
- [Handley08] "Multipath TCP and the resource pooling principle, Mark Handley, Damon Wischik and Marcelo Bagnulo Braun, IETF Dublin TSVAREA presentation, 2008
- [Jacquet08] "Policing Freedom to Use the Internet Resource Pool", Jacquet, Briscoe and Moncaster; Re-Arch CoNEXT workshop, Dec 2008
- [Kelly05] F. Kelly , T. Voice, "Stability of end-to-end algorithms for joint routing and rate control", ACM SIGCOMM Computer Communication Review, v.35 n.2, April 2005
- [Laskowski06] P Laskowski & J Chuang, "Network Monitors and Contracting Systems: Competition and Innovation," In Proc. SIGCOMM'06, ACM CCR 36(4)183--194 (2006)
- [Massoulié07] P. B. Key, L. Massoulié, D. F. Towsley, "Path Selection and Multipath Congestion Control", INFOCOM 2007, 143-151
- [Mitzenmacher01] M. Mitzenmacher, "The Power of Two Choices in Randomized Load Balancing", IEEE Transactions on Parallel and Distributed Systems, Volume 12 , Issue 10 (October 2001)
- [PBG08] P. Brighten Godfrey, "Balls and bins with structure: balanced allocations on hypergraphs", Symposium on Discrete Algorithms, San Francisco, 511-517, 2008
- [Roughgarden02] T. Roughgarden and E. Tardos, "How bad is selfish routing?", Journal of the ACM, 2002
- [RRG] Routing research group
- [SHIM6] <http://tools.ietf.org/wg/shim6/>
- [Thaler00] D. Thaler, C. Hopps, "Multipath Issues in Unicast and Multicast Next-Hop Selection", RFC 2991, November 2000
- [Wischik08] "The Resource Pooling Principle", Damon Wischik, Mark Handley and Marcelo Bagnulo Braun. ACM/SIGCOMM CCR, Oct 2008

## Acknowledgements

The research results presented herein have received support from Trilogy (<http://www.trilogy-project.eu>), a research project (ICT-216372) partially funded by the European Community under its Seventh Framework Programme. The views expressed here are those of the author(s) only. The European Commission is not liable for any use that may be made of the information in this document.



## **A Future Internet Embracing the Wireless World**

Henrik ABRAMOWICZ<sup>1</sup>, Norbert NIEBERT<sup>1</sup>, Stephan BAUCKE<sup>1</sup>, Martin JOHNSON<sup>1</sup>,  
Börje OHLMAN<sup>1</sup>, Mario KIND<sup>2</sup>, Klaus WUENSTEL<sup>3</sup>, Hagen WOESNER<sup>4</sup>, Jürgen  
QUITTEK<sup>5</sup>

<sup>1</sup>Ericsson Research, Corporate Unit, {henrik.abramowicz@ericsson.com}

<sup>2</sup>T-Systems, <sup>3</sup>Alcatel-Lucent, <sup>4</sup>Telecommunication Networks Group (TKN),

<sup>5</sup>NEC Laboratories Europe

**Abstract:** In this paper we describe several approaches to address the challenges of the network of the future especially from a mobile and wireless perspective. Our main hypothesis is that the Future Internet must be designed for the environment of applications and transport media of the 21st century, vastly different from the initial Internet's life space. One major requirement is the inherent support for mobile and wireless usage. A Future Internet should allow for the fast creation of diverse network designs and paradigms and must also support their co-existence at run-time. We observe that a pure evolutionary path from the current Internet design is unlikely to be the fastest way, if at all possible, to address, in a satisfactory manner, major issues like the handling of mobile users, information access and delivery, wide area sensor network applications, high management complexity and malicious traffic that hamper network performance already today. We detail the scenarios and business use cases that lead the development in the FP7 4WARD project towards a framework for the Future Internet.

**Keywords:** Future Internet, Network Architecture, Network Virtualisation, Self-Management, Information-centric Networking.

### **1 Introduction**

Driven by the encouragement for new approaches from some of the "fathers of the Internet" (e.g. [4], [7]) and early experimental testbeds (see e.g. [14]), the discussion on the "Network of the Future" is gaining in intensity due to increasing concerns about the inability of the current Internet to address a number of important issues affecting present and future services and to the impetus provided by "clean slate design" research initiatives launched in the US, Europe and Asia. Many problems, arising from mobile and wireless perspectives, with the current network architecture have been recognized for a long time but have not received a satisfactory solution (see e.g. [1]). The issues like security, manageability, dependability, mobility, etc. result both from initial design flaws as well as the wide set of applications over the Internet that could not be envisioned from the beginning. In this paper, we present the approach taken within the 4WARD project ([www.4ward-project.eu](http://www.4ward-project.eu)) to address these problems by researching different aspects of the Future Internet design.



In section 2 we first discuss societal and business forces that must guide our technical choices. Section 3 introduces our ideas on the key technical components for a Future Internet consisting of an architectural framework, network of information, flexible transport paths, network virtualisation and self-management. We end the paper with a short conclusions section.

## 2 Motivation and scenarios for the Future Internet

The Internet was initially developed for a limited number of trusted nodes interconnected by copper based transmission technology implemented supporting applications like file transfer and message exchange. The initial architecture developed for this purpose was essentially simple but open for new applications. Its evolution has led to a tremendous success – the Internet as we know it today. It is however far from clear that it is still the optimally evolvable solution, able to meet the challenges of fibre optics and radio transmission technology, real-time multimedia and file-sharing applications and exposure to an untrustworthy world. Furthermore the Internet, starting as a simple set of protocols and rules, has over the decades reached a state of high complexity with regard to interoperability, routing, configuration and management.

Within the research community the need for change is largely acknowledged although there is not yet agreement on how this change should take place. Some propose a *clean slate* approach, which aims at investigating new architectural concepts with new requirements in mind and which initially doesn't need to consider legacy, while others are advocating an evolutionary approach, introducing new solutions incrementally. It seems likely that both approaches will migrate current Internet technologies towards a Future Internet.

### 2.1 Scenarios for the Future Internet

The identification of key drivers is one of the most difficult prerequisites in the development of the Future Internet. By analysing the key driving forces and challenges in the Future Internet business environment, the 4WARD scenarios were built. These scenarios cover aspects of technical as well as non-technical areas.

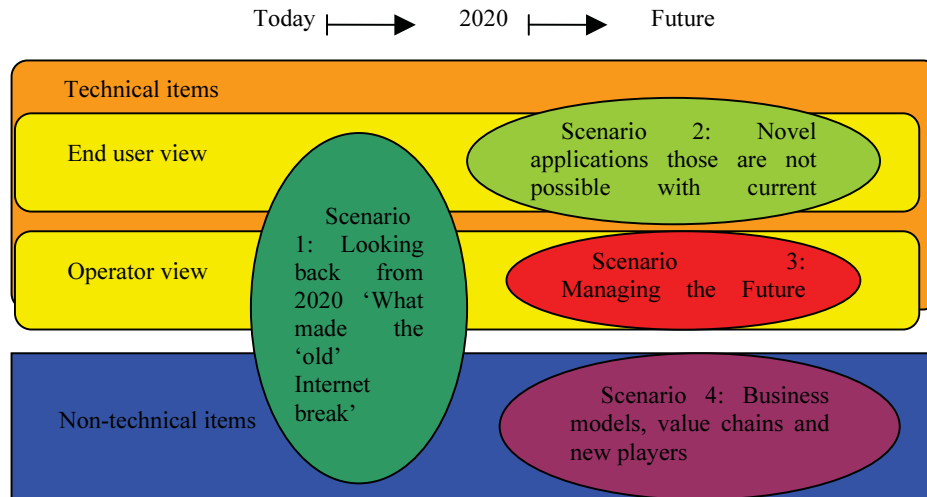
The general frame for the scenarios can be summarized as follows:

4WARD addresses the worldwide potential telecommunication market in 2020

Time frame for upcoming technologies: it should be possible around 2015, it should be widely used in 2020

Maintain an end-to-end view with respect to service, usage, business and technology development

Within this section four scenarios are described which focus on different aspects of the possible future evolution of the Internet. The following figure gives an overview on the potential areas of the scenarios.

**Fig. 1.** Scope of 4WARD scenarios and temporal horizon

The rest of the section presents the main conclusions of the different scenarios. The scenario 1 “Looking back from 2020 ‘What made the ‘old’ Internet break” outlines which technical and non-technical developments will have been decisive for the understanding that the smooth evolution of the existing Internet concepts will no longer be applicable in the communication world. This includes the analysis of infrastructure problems, innovation restrictions and the limitations in economic incentives.

The second scenario “Novel applications that are not possible with current Internet” identifies and evaluates from end user view which challenges will be posed from conceivable new applications to the Internet and how they overstrain the existing Internet concepts. This includes enablers for more user orientation, mobility support and augmentations. Some examples are networks that fit perfectly to users’ likes, dislikes, preferences, and so on, even if users temporarily use someone else’s terminals, the integration of the real world and the Internet, the potential of having a better support of non-continuous access to the Internet and asynchronous communication and the services for individual’s life kernel (SILK), e.g. for health monitoring, control of house appliances, personal identification and interaction and how these services are not supported by today’s Internet infrastructure. Sometimes they are possible only ‘in principle’, but wide-spread adoption is not possible due to complexity or scalability issues, lack of usable devices or other restrictions.

Scenario 3 “Managing the Future Internet - Benefits for operators” concentrates on network management issues which come up with the broadening of the traditional one-stop-shop operator to an environment with several partly competing, partly collaborating network operators and a multitude of service providers. Major themes covered are the blurring boundaries between operators and other players in a future Internet, the growing complexity of infrastructure and services and the associated need to find new ways of network/service management, the new capabilities provided to operators, based on innovative future Internet technologies. The separate document D.4-1 [17] details the problems and presents more in depth results.

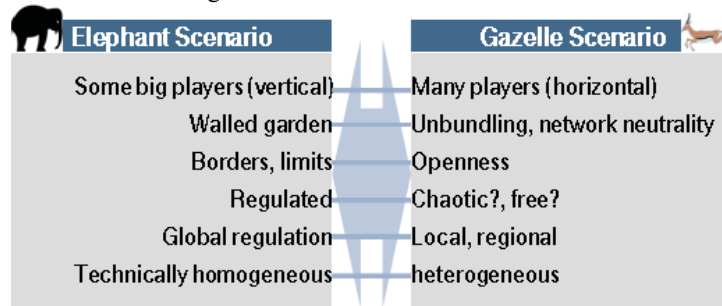
#### **Business models, value chains and new players**

The last scenario focuses on the non-technical aspects of the Future Internet. It evaluates the impact of social, economic and political trends on the telecom business to work out the most

decisive elements which will govern the future business environment. The most important questions are:

1. Will the Internet arena be dominated by a limited number of big players or is it, on the contrary, more feasible that a multitude of specialized small companies will satisfy the increasing demand for individual services?
2. Will centralisation (e.g. big server farms) or decentralisation (peer-to-peer networks) determine the direction of future developments?
3. What will be the main inhibition of growth: regulative intervention, compatibility problems of technical solutions or a mismatch in market power?
4. How can the global usage and accessibility of the Internet be assured under different market environments without global regulation?
5. Will heterogeneity in technology accelerate or retard technical innovation? Is the coexistence of multiple heterogeneous platforms (may be operating on the same physical system but separated by virtualisation) a good alternative?

First answers on these questions have led to two opposite borders, called elephant and gazelle scenario. The figure below shows the borders of the scenario framework.



**Fig. 2.** Extreme Scenarios characterized by six drivers with uncertain development

## 2.2 Use cases

Based on the four scenarios, a set of different use cases was developed. Each use case covers a precise technical or business related topic. In order to reduce the complexity in the detailed analysis, a set of mandatory business topics were defined. These are the following:

General description including a basic technical overview (provide a basic overview how the use case could be implemented in the future)

Business player (including the analysis of a value proposition, the target market description, principle revenue model, resources and processes as well as the potential impact of the 4WARD technology, and the three non-technical aspects as described above)

Customer (focus is on the potential value and the impacts from technical and non-technical area)

Rivalry (based on the idea of Porter's Five Forces [15] and its analysis principles for competitors, suppliers, substitutes, new entrants and buyers. The model's limitations are understood and it has been extended to overcome these limitations).

After several evaluation rounds, five major use cases remained for detailed analysis:

1. Enhanced connectivity in the user communication space
2. End-to-End Quality of Service
3. Virtualisation
4. Identity Provider
5. New ways of information delivery

The idea of “Enhanced connectivity in the user communication space” is to separate the service connection from the network and holding an interaction between both layers. For example, a user is allowed to freely shift his communication from one terminal to another one without losing the connection (e.g. in case of power failure of a terminal or arriving at a location with a better suited or connected terminal). Furthermore, enhanced security and the communication representation (the bit streams) will adapt to the connectivity and device constraints of the new terminal. This could change the service usage in principle and enable customers integrated service mobility in the future.

In the use case “End-to-End Quality of Service”, the focus lies on the interprovider connections at the example of Quality of Service (QoS). This includes aspects like assuring QoS service levels by leveraging network state aware routing, provisioning of connections with guaranteed QoS levels across borders of different providers. The business view is the reduction of overprovisioning of networks and the associated optimisation of costs as well as the possibility for better service delivery and potential increase in revenues.

With (network hardware) virtualisation, the idea of virtualisation in the server area is applied for network optimisation. The implications could change the whole business logic in today’s telecommunication environment and adapt to the more open Internet approach. This implies changes in the ecosystem containing now the three potential players Infrastructure Provider, Virtual Network Provider and Virtual Network Operator. In addition, other areas are interfaced like service delivery and usage or regulation. Key target is the optimisation of costs and possibility to retain profitability in the infrastructure business.

Exploration of the business opportunity arising from a combination of 4WARD concepts Network of Information and Generic Path is the focus of the use case “New ways of information delivery”. It evaluates the prospects of deploying respective technologies, trying to identify a number of services they will enable in a Future Internet and who will be the stakeholders for these services.

The two use cases “virtualisation” and “New ways of information delivery” are analysed in more detail and documented in [16].

### 2.3 Migration Issues

4WARD has taken a research approach that is called “clean slate approach” that means that from research point of view we start the research as if the Internet does not exist. In the end we will have to take a migration approach that is applying results from research into the real network.

This can basically be done in 3 ways.

One is to incrementally enhance the existing internet paradigm by adding extension to present protocols and functions without compromising current implementations; an example could be Mobile IP. By this there will be no fundamental change of Internet and the problems with the current Internet will remain.

Another approach is to make use of overlay or underlay techniques which have been used for many years in traditional telecom as well. An example of overlay techniques could be SIP for VOIP or different access technologies like Ethernet or radio as examples of underlay techniques and hereby placing functionality either on top of or below the current Internet. Although this approach solves more problems than the previous one, we are still faced with that this approach is relying on the functionality of the current Internet. There is a risk for more overhead but also fragmentation due to that applications might need to implement same or similar functions per application rather than using an underlying common support functions.

A third option would be making use of network virtualization techniques and by this separating the network into virtually independent “slices”. This means that based on a common

physical infrastructure one can have several network architectures operating in parallel, and which for example could interoperate at gatewaying points. This would allow for having networks serving different needs e.g. sensor networks, enterprise networks or even public networks, and where some could be based on new technologies.

There is of course another possibility to deploy a completely new network in parallel with the current Internet. We do not believe that this is a viable commercial option due to the immense success of Internet and its installed base. For this reason we will not pursue this approach further.

### 3 Key Components for a Future Internet

To realise those scenarios and business propositions, we have to develop a consistent set of key components for a Future Internet that we present here in the following subsections.

#### 3.1 The architecture framework of a Future Internet

To enable innovation and rapid deployment of new networking solutions, the development of new architectures suitable for a specific environment (e.g. a LAN or a new type of radio access network or a specialised application), should be facilitated and the reuse of common components made possible. We develop a new **architecture framework** that must be able to accommodate changes in business and technology environments. Such agility is emerging in the software area with service oriented architectures and design patterns. We plan to extend and generalize these approaches, and develop an architecture framework by which different network architectures, which are tailored for various purposes and environments, can be derived and implemented. The aim is to end up with lean and dedicated instantiations of network architectures that remain **interoperable and evolvable**.

The interoperability that has been solved naturally by IP becomes a concern without the universal presence of design principles. Without such principles, it will in the long run be hard to interconnect and to interoperate the resulting networks. The design principles need to express aspects and properties that pertain to naming, addressing, routing, QoS, (self-) management, security, as well as overall performance objectives. Given the coexistence of heterogeneous network environments and different (and still unknown) technologies, it is very important to carefully analyse gatewaying principles for interconnecting networks having implemented different network architectures. It is likely that a modular and scalable approach to gatewaying should be considered.

The design of an architecture framework started with defining common requirements as well as a set of invariants. They must generally concern the performance objectives, scalability, extensibility, as well as the consistency and coherency of communication systems throughout the lifetime of the architecture framework. Implicit invariants usually emerge by overloading functions intended for other purpose(s), making the adaptation/replacement of these functions impossible. Indeed, according to Ahlgren et al. [6], if invariants are not explicitly defined, the design will be deficient in the long term, despite its superficial flexibility. The properties and aspects that, for instance, a specific sensor network and a MAN, or any other network of the future, will have in common, still need to be identified and investigated. Through the architecture framework it should be possible to instantiate, e.g. a very light-weight network architecture suitable for low-energy networks, with a very limited set of features implemented. Similarly, one should be able to instantiate a network architecture suitable for a MAN, for example, with built-in features such as security, privacy, QoS, and mobility.

Reconciling such diverse aspects as discussed above will be a challenge. Thus, explicit invariants, principles, properties, and design patterns shall be carefully designed into the

architecture framework. They are, by definition, the specific characteristics that determine the options as well as limitations for how network architectures can develop and evolve over time. The first results on architecture framework can be found [18].

### 3.2 Moving from networking of nodes to networking of information

The traditional role of networking has been to interconnect remotely located devices like computers or telephones. This function is increasingly recognised to be ill-adapted and inadequate for the information-centric applications that currently generate the vast majority of Internet traffic.

In 4WARD Networking of Information (NetInf) we take a different approach. Instead of the traditional *node-centric* paradigm, we adopt an *information-centric* paradigm. In this paradigm, the communication abstraction presented to applications is based on the transfer of application data objects instead of end-to-end reliable byte-streams as used today.

The current semantic overload of the IP-address as both node identifier and locator, indicating the current point of attachment in the network topology, is replaced by a clear separation of information self-certifying object identifiers and locators. Several models for abstracting the location and focusing on networking between (mobile) hosts have been proposed, (e.g. [7], [3], [9], [10]). 4WARD builds on this prior work and by taking it one step further; we are able to design a networking architecture where mobility, multihoming and security is an intrinsic part of the network architecture rather than add-on solutions. It also allows users to gain increased control over incoming traffic enabling new possibilities for defending against denial of service attacks. The self-securing property also intrinsically facilitates possibilities for effective content protection and access rights management.

The increasing number of overlays created for the purpose of information dissemination (e.g., Akamai CDN, BitTorrent, Joost) clearly shows the need for an information-centric approach. These overlays massively distribute information and move the load away from any central server, scaling automatically to any group size. 4WARD integrates much of the functionality of these overlays, including caching. This is done in a common and open information networking service that integrates networking and storage and is generalised for use by applications.


4WARD extends the networking of information concept beyond "traditional" information objects (e.g., web pages, music/movie files, streaming media) to conversational services like telephony, and store-and-forward services like email. Special attention is paid to how services can be made to work in an environment with a heterogeneous and disruptive communication infrastructure. Furthermore, we investigate how networking of information can extend to include real world objects, and by this enabling new types of services. Our initial results in the NetInf area can be found in [20].

### 3.3 Networking with the Generic Path

The construction of a path as a sequence of relaying nodes in a network takes currently place on multiple layers. In fact, a port of an IP router in the backbone will today typically be connected to an SDH or Ethernet layer on top of an optical layer. GMPLS has been introduced as control plane for multi-layer networks [11]. Here, for the first time, the otherwise lower-layer agnostic IP routers may perceive the notion of a changeable topology, leading away from pure overlay networks with separate control to an integrated and possibly distributed management of data transport. Our approach is to define the notion of a "Generic Path", able to efficiently realize "networking of information" by exploiting cross-layer optimization and multiple network paths.

We define a Generic Path (GP) as “means to organize the accessibility of a sufficient number of parts or copies of information objects stored in a group of hosts.” See also [21]. Incorporating the paradigm of *information-centric networks* means that a GP is actually hiding the physical location of information objects. Wherever chunks or copies of information are stored, the GP takes care of delivering it to a sink.

Because cross-layer information is available, new transmission techniques can be used inside a GP. This is especially interesting for the introduction of network coding into fixed and wireless networks. Here, multipath routing needs to be combined with specific capabilities of nodes (e.g., bit-wise XOR of two frames).



**Fig. 3** The Generic Path as a hull organizing data transport over multiple paths and layers.


Generic Paths can thus be seen as a “hull” that is filled with information. One advantage of this concept is that mobility of information objects and hosts becomes conceptually equivalent and is dealt with by the GP internally.

There are a number of open questions that are addressed inside 4WARD to bring this concept to reality: routing and interaction of generic paths, the control plane for network coding, enhancement of mobility management and, importantly, the definition of a generic path API that allows the instantiation of a GP similar to today’s sockets. The GP thus appears as the fundamental information channel in the future Internet that is sufficiently flexible to adapt to different requirements and available network technologies. A description of the GP mechanism can be found in [22].

### 3.4 Towards lean and innovative networks through virtualization

To introduce clean slate solutions like information-centric networking and generic paths we have to allow them to coexist with existing and other new approaches. Virtual networks can enable new protocols and architectures to be deployed independently without disruptions. Virtualization has been used in test-bed environments and is now being proposed as the basis of commercial networks (see also e.g. [10]). Virtual networks are ideally suited to allow the coexistence of different network architectures, legacy systems included. Virtualization is thus not only the enabler for the coexistence of multiple architectures, but it is also a smooth path for the migration towards evolutionary approaches. The goal is hence to develop a systematic and general approach to network virtualization. The virtualization of individual resources is the basis of the framework as depicted in Fig. 4 below.



**Fig. 4** Virtualization Framework

While the virtualization of many types of resources, such as servers and links, is well-known and already widely used today, we aim for a generalized approach that allows the use of a broad diversity of resources with higher flexibility and security. Virtualization of both wireless and wireline resources is expected to play a key role in the Future Internet. In particular, the secure, flexible, and efficient exploitation of wireless spectrum resources and wireless infrastructure is expected to significantly improve cost-effectiveness and utilization of expensive wireless infrastructures.

Virtualization allows an evolution of communication technology while largely reusing deployed infrastructure; thereby it reduces the economic barrier for technical evolution. It further provides a general framework for *network sharing*: providing different networking services of different network service providers on a common physical infrastructure. This is particularly beneficial in network domains where the deployment costs per network user are predominant and an obstacle for frequent technology replacement.

A key concern for owners of infrastructure resources and the operators of virtual networks using these resources is security and trust. The virtualization framework must ensure the protection of the physical resources, as well as the strict isolation of concurrent virtual networks from each other. Furthermore, virtualization may significantly change the business environment for infrastructure owners and operators' business models and incentives for use in a commercial setting need to be carefully considered. Our draft approach to Virtualisation can be found in [19].

### 3.5 InNetworkManagement: A new network management paradigm

The diversity of technologies and business models envisioned in previous sections can only be supported in operative networks if adequate management functions are integrated to initiate and maintain the network infrastructure. Management capabilities in current networks typically reside outside the network. Research has focused on solutions for self-management but so far these are mainly algorithms solving specific problems. Most of these solutions lack scalability, imply considerable integration costs with central management stations and – most important – are not suitable to cope with the complexity and dynamicity of tomorrow's networks.

In order to address these issues, the 4WARD project follows a new paradigm to network management. The basic concept of the new paradigm that we call *InNetworkManagement* is (1) to have network management functions as *embedded* 'default on' management capabilities of



network devices and (2) to allow these devices to interact in a peer-to-peer fashion to enable network-wide management functions. We envision management functions as inseparable capabilities of the device and the network itself. This leads to a novel, strongly decentralized architecture where management operations are localized in the network components. As a consequence, faults can be identified more quickly and isolated using cross-layer techniques, and control loops can be enforced more efficiently than in traditional management architectures. Benefits from this approach are to access embedded functions to cope with diverse technologies, different business models and the rich mix of services instead of adding complex management systems into the networks. We believe that InNetworkManagement is particularly beneficial in large-scale, dynamic network environments. A number of these issues have been described in [17].

The new embedded management functions are accessed through a *management plane inside the network* that organises itself and automatically adjusts to different network sizes and configurations. It executes a set of distributed, self-stabilising protocols for monitoring and control, enabling a range of self-management functions inside the network. This is accomplished first of all through the definition of models of interactions between network components and the inclusion of self-organizing algorithms inside network devices. Secondly, the behaviour and objectives of the network as a whole is studied and modelled. This includes outer control loops between different components and operators' interfaces to support network-wide processes, including monitoring of aggregated states and policy enforcement.

The development of protocols for the management plane can draw on current research on the computability of distributed functions under cost constraints, sensor networks and probabilistic protocols for distributed systems [13]. However, application to network management calls for progress beyond this research, in order to take into account the particular constraints regarding operational overhead in the management plane, the richer functionality of management operations and the potentially large number of concurrently executing management functions. Therefore, a systematic analysis of network working conditions is required, in order to assess the effectiveness of management operations in different situations and scenarios. Such an analysis identifies the trade-offs for management operations, including protocol overhead, accuracy of monitored data, timeliness of self-healing actions and frequency of self-optimization actions, which should become tuneable in real-time. (See [12] for an example). In addition, mechanisms are developed that provide control over the relationship between decision quality and the cost of achieving a specific level of situation awareness in the management plane. Our concept of INM is described in [23].

## 4 Conclusions

Considerable research effort is clearly necessary to address the challenges raised by the design of a Network of the Future. This effort is currently underway with many Future Internet activities across the world. The main thrusts of 4WARD, *a new architectural design*, the *information-centric paradigm*, the *generic path network virtualization and embedded self-management*, provide candidate solutions, which, after careful evaluation, should be appropriately incorporated into the architecture of the Future Internet. A major issue will be the integration of the various approaches within a common architecture framework. Results of this work are expected over the coming years.

## 5 References

- [1] R. Tafazolli (ed.); Technologies for the Wireless Future: Wireless World Research Forum (WWRF), Volume 2; Wiley; 2006

- [2] N. Feamster, L. Gao, J. Rexford; How to lease the Internet in your spare time; ACM SIGCOMM Computer Communications Review, Vol. 37, No.1, January 2007
- [3] Stoica, D. Atkins, S. Zhuang, S. Shenker, and S. Surana. Internet Indirection Infrastructure. In Proceedings of ACM SIGCOMM 2002, April 2002. Pittsburg, USA.
- [4] D. Clark, K. Sollins, J. Wroclawski, and R. Braden. Tussle in Cyberspace: Defining Tomorrow's Internet. In Proceedings of ACM SIGCOMM 2002, August 2002. Pittsburgh, USA
- [5] P. Nikander, J. Arkko, B. Ohlman, "Host Identity Indirection Infrastructure (Hi3)", IETF Draft draft-nikander-hiprg-hi3-00.txt, December 2004
- [6] B. Ahlgren, M. Brunner, L. Eggert, R. Hancock, S. Schmid; Invariants – A New Design Methodology for Network Architectures; SIGCOMM 2004 Workshop on Future Directions in Network Architecture (FDNA'04), Portland, OR, USA, August 2004
- [7] V. Jacobson, M. Mosko, D. Smetters, J.J. Garcia-Luna-Aceves; Content-Centric Networking ; Whitepaper Describing Future Assurable Global Networks, January 2007
- [8] R. Moskowitz, P. Nikander; Host Identity Protocol (HIP) Architecture; Internet Engineering Task Force RFC 4423, May 2006
- [9] H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, M. Walfish; A Layered Naming Architecture for the Internet; ACM SIGCOMM 2004, Portland, OR, USA, August 2004
- [10] B. Ahlgren, J. Arkko, L. Eggert and J. Rajahalme, "A Node Identity Internetworking Architecture", 9th IEEE Global Internet Symposium, Barcelona, Spain, April 28-29 2006
- [11] Tomkos, et al., "Performance Engineering of Metropolitan Area Optical Networks Through Impairment Constraint Routing", IEEE Communications Magazine (OCS) pp. s40-s47, August 2004.
- [12] Gonzalez Prieto, R. Stadler: "A-GAP: An Adaptive Protocol for Continuous Network Monitoring with Accuracy Objectives", IEEE Transactions on Network and Service Management (TNSM), Vol. 4, No. 1, June 2007.
- [13] Giridhar and Kumar: "Towards a Theory of In-Network Computation in Wireless Sensor Networks," IEEE Communication Magazine, April 2006.
- [14] Global Environment for Network Innovations; geni.net
- [15] Michael E. Porter: The five competitive forces that shape strategy, Harvard Business Review, January 2008, Vol. 86 Issue 1, p78-93.
- [16] 4WARD Deliverable<sup>1</sup> D-1-1: First Project-wide Assessment on Non-technical Drivers,
- [17] 4WARD Deliverable D-4.1 Definition of scenarios and use cases
- [18] 4WARD Deliverable D-2.2 Draft Architecture Framework
- [19] 4WARD Deliverable D-3.1.0 Virtualisation Approach: Concepts
- [20] 4WARD Deliverable D-6.1 First NetInf Arch Description
- [21] 4WARD Deliverable D-5.1 Architecture of a Generic Path
- [22] 4WARD Deliverable D-5.2.0 Description of Generic Path Mechanism
- [23] 4WARD Deliverable D-4-2 In-Network Management Concept

---

<sup>1</sup> All 4WARD deliverables can be found at <http://www.4ward-project.eu>

## The Publish/Subscribe Internet Routing Paradigm (PSIRP): Designing the Future Internet Architecture

Sasu Tarkoma, Mark Ain, Kari Visala

Helsinki Institute for Information Technology (HIIT), P.O. Box 9800, 02015 TKK, Finland

{Sasu.Tarkoma, Mark.Ain, Kari.Visala}@hiit.fi

**Abstract.** Despite its success, the Internet is suffering from several key design limitations, most notably the unification of endpoint locators and identifiers, and an imbalance of power in favor of the sender of information. The unfavourable consequences are that the full range of possibilities offered by the Internet may not be fully realized and trust in its proper operation has been significantly weakened. In this paper, we introduce the Publish/Subscribe Internet Routing Paradigm (PSIRP) and present an architectural redesign of the global Internet based on an information-centric publish/subscribe (pub/sub) communication model. Through its application of pub/sub communications and efficient network design emphasizing end-to-end trust, we believe that the PSIRP-reengineered Internet may resolve many of the problems plaguing the current Internet and provide a powerful and flexible network infrastructure with a high degree of resiliency.

**Keywords:** Future Internet, information-centric, publish/subscribe, scoping

### 1 Introduction

Since its conception, the Internet has experienced tremendous growth, ever increasing traffic and new applications, including voice and video, while still retaining its original architecture drafted almost 40 years ago. The main guiding principle for the design of the Internet was the end-to-end principle [Sal1984].

Blumenthal et al. [Blu2001] identify a number of challenges for the end-to-end principle: operation in an untrustworthy Internet, more demanding applications, the rise of third party involvement, ISP service differentiation, and less sophisticated users. Moreover, one of the most notable issues in the current Internet is the imbalance of powers in favor of the sender of information, who is overly trusted. The network accepts anything that the sender wants to send and will make a best effort to deliver it to the receiver. This has led to increasing problems with unsolicited traffic (e.g. spam e-mail) and distributed denial of service (DDoS) attacks.

The *publish/subscribe (pub/sub) paradigm* has been proposed as a remedy to the problems facing the current Internet. In pub/sub networking, senders “publish” what they want to send and receivers “subscribe” to the publications that they want to receive. In principle, no one receives any material to which they have not explicitly expressed an interest by way of subscription.

One can observe that many widely used Internet applications already are essentially publish/subscribe in nature. For example, distribution of software updates is currently performed in a poll/unicast fashion that is clearly non-optimal. Instead, subscribing to the updates that are needed and distributing them via multicast, caching etc. would be much easier and more efficient from the point of view of network resource usage.

The PSIRP project will redesign the entire Internet architecture from the pub/sub point of view, taking nothing (not even IP) for granted. PSIRP's work will focus on the intersection of security, routing, wireless access, architecture design, and network economics, in order to design and develop efficient and effective solutions. In such a new Internet, *multicast* and *caching* will replace unicast and cache-free data fetching operations, while *security* and *mobility* will be embodied directly into the foundation of the architecture rather than added as after-thoughts.

The new pub/sub-based internetworking architecture aims to restore the balance of network economics incentives between the sender and the receiver and is well suited to meet the challenges of future information-centric applications and use modes. To our knowledge, this type of application of pub/sub communication models has not been tried before.

This paper is structured as follows: Section 2 provides an overview of the concept of information-centric networking and general philosophy behind PSIRP; Section 3 covers the architectural components, entities, processes, network services, and their functionalities; Section 4 discusses the prototype implementation and application considerations, and finally Section 5 contains concluding remarks.

## 2 PSIRP Background

We aspire to change the routing and forwarding fabric of the global inter-network so as to operate entirely based on the notion of information (associated with a notion of *labels* to support fabric operation) and its surrounding concerns, explicitly defining the *scope* of the information and directly addressing information (via *rendezvous identifiers*) as opposed to addressing physical network endpoints. The envisioned operation on information is in sharp contrast to the current endpoint-centric networking model. The current end-to-end model of IP networking requires that both the relevant data and explicitly-addressed network locations be known in order to transparently stream information between two endpoints. Our model emphasizes information-centric operation: data pieces are explicitly addressed through identifiers and labels serving as high-level designations/resolvers to lower-level schemas, and scoping mechanisms that can define information inter-networks and relationships within a global information taxonomy. As such, information is embedded immediately into the network and it is the only effective element in need of direct user-manipulation; the physicality of the network (i.e. endpoint locations) need not be known directly.

Another important aspect of the PSIRP architecture is that it is receiver-driven. We take the approach that the receiver has control over what it receives and we cascade this approach throughout the core of the PSIRP *component wheel* and the multiple

operational elements within the PSIRP architecture. A receiver must *elect* to join (i.e., subscribe) to an identifier before it can receive any information. Sending (i.e., publishing) as well as receiving operations are thus decoupled between the senders and the receivers in both space and time. Hence, PSIRP not only intends to move the functionality of many existing publish/subscribe systems (e.g., [Eug2003b]) onto the internetworking layer but also base the entire communication, throughout the architecture, on this paradigm.

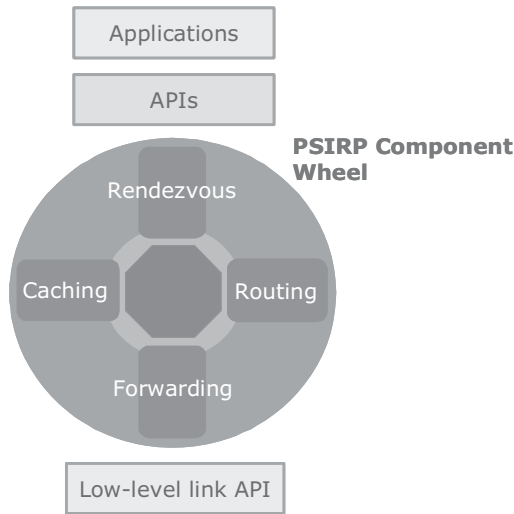
### 3 PSIRP Conceptual Architecture

This section presents the PSIRP conceptual architecture, defining the key entities and processes of system and their attributes. The PSIRP conceptual architecture consists of three crucial parts, namely the protocol suite architecture (called the component wheel), the networking architecture, and the service model.

#### 3.1 Component Wheel

The PSIRP conceptual architecture is based on a modular and extensible core, called the *PSIRP component wheel*. The architecture does not have the traditional stack or layering of telecommunications systems, but rather components that may be decoupled in space, time, and context. The idea of such a layer-less network stack has been proposed before, for example, in the Huggle architecture [Hag2007]. The novelty of the PSIRP proposal is to use publish/subscribe style interaction throughout the conceptual architecture, and thus support a layer-less and modular protocol organization. This organization is primarily achieved through the efficient structuring of information identifiers and their interactions amongst network elements, offering ample flexibility for future expansion.

Figure 1 presents an outline of the conceptual architecture with the PSIRP component wheel in the middle. Above the wheel, we have APIs that facilitate accessibility to and implementation of different networking features that are available in the system. The figure illustrates the typical components needed in the wheel for inter-domain operation: *forwarding*, *routing*, *rendezvous*, and *caching*.



**Fig. 1.** PSIRP component wheel.

### 3.2 Network Architecture

We can view the global network of information as an *acyclic graph* of related pieces of data, each identified and scoped by some identifiers. In the PSIRP architecture, identifiers define the relationships between the pieces of information in our system on the different levels, such as the application or networking level. With this in mind, we propose the following classes of identifiers:

- *Application identifiers (AId)*, used directly by publishers and subscribers.
- *Rendezvous identifiers (RId)*, used to bridge higher level identifiers with lower layer identifiers.
- *Scope identifiers (SId)*, used to delimit the reachability of given information.
- *Forwarding identifiers (FId)*, used to define network transit paths and transport publications across networks.

A rendezvous identifier is implicitly associated with a well-defined (but not necessarily fixed) *data set*, consisting of one or more publications. The data sets may also have associated *metadata*, which may include, e.g., scoping information and other useful information either for ultimate receivers or for network elements.

We also consider metadata that is understood within the network itself. Such network-level metadata might be concerned with how the communication for a particular data item may be conducted. This network metadata may be found as soft state within the network, carried as part of the communication header, or referred to by separate identifiers. Such functions may include access control, flow control, error notification, congestion notification, etc.

PSIRP necessitates the existence of a *scoping* mechanism as to limit the reachability of information. Scoping information is associated with a publication, determining the elements of the rendezvous system that act on published data and therefore defines the information network that the information belongs to. A publication may be associated with one or more scopes.

Scoping can be seen to represent a logical equivalent to the concept of *topologies* (such as link-local or site-local) in the endpoint-centric IP world. Given the information-centrism of our architecture, however, a scope is naturally attached to every item of information that is fed into the network (although we can consider the case of “no scope” being attached to a data item as being equivalent to the notion of “localhost” in the IP world - in other words, the information would not leave the local computer). In effect, scoping allows for building *information inter-networks* as opposed to topological inter-networks.

Scopes define a powerful concept that can construct social relations between entities, representing consumers and providers of information, and the information itself. This is illustrated in Figure 2, where certain information (e.g. a picture) is available to family and friends, while other information is merely visible to colleagues. Each scope is attached with a governance policy, represented as metadata, which may include (amongst other things) authentication information for potential receivers of the information.

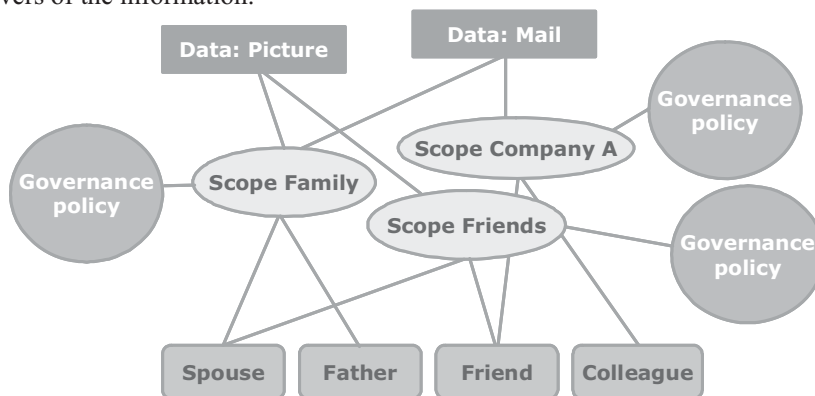


Fig. 2. Concept of scope.

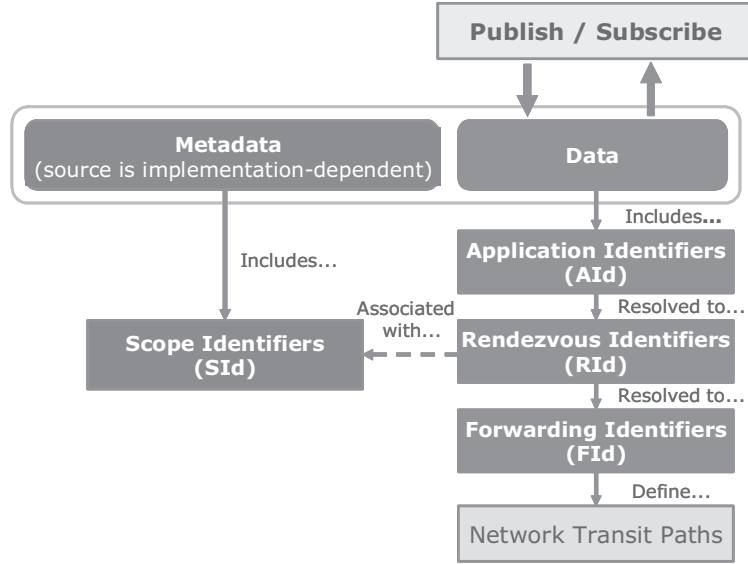
Scopes can be easily re-constructed, removing certain parties from the scope, adding new publications to the scope, and assigning information to a new scope. This ability works towards our stated ambition to enable networks to reflect social structures and information grouping in a dynamic fashion.

The publisher/sender interface supports publication of data. Each publication has an associated flat label (i.e. the rendezvous identifier), and an optional metadata part.

A subscriber initiates a receiver-driven communication through a rendezvous identifier, specified in an act of subscription. Similar to the publisher, the subscriber can specify additional metadata surrounding the request.

### 3.3 Functional Entity Relationships

Figure 3 illustrates the relationships between the key entities of the PSIRP architecture.



**Fig. 3.** Key entities of the conceptual architecture.

Typically, data is associated with one or more application identifiers and one or more scopes. Each application first resolves *application identifiers (AId)* into rendezvous identifiers.

A *rendezvous identifier (RId)* represents the network level identity of a publication and is associated with policy-compliant data dissemination graphs for publication delivery, both in the local domain (intra-domain) and between domains (inter-domain). The rendezvous identifiers are chosen from within a large enough set to provide a probabilistically feasible likelihood of uniqueness without a central allocation authority.

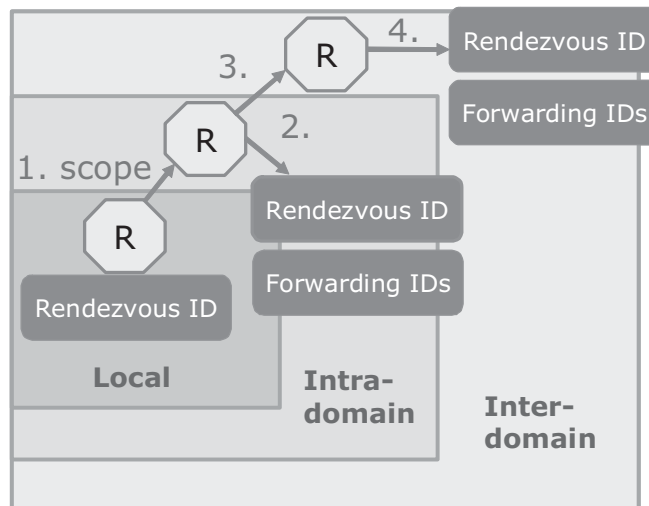
A given application then hands the rendezvous identifiers to the network, using the scopes to properly map each rendezvous identifier to one or more *forwarding identifiers (FId)*, both within a domain and between domains. It is then the responsibility of the rendezvous functions to find suitable data transit and delivery paths in the network and denote them with forwarding identifiers. The breadth of reference of FIds is variable, potentially limited to single hops or dynamically expandable to encompass full multicast trees. This relatively open structuring scheme allows concurrent use of FIds to support flexible routing mechanisms based on source routing, anycast, multicast trees etc.



### 3.4 Rendezvous and Routing

Rendezvous is the process of resolving rendezvous identifiers into forwarding identifiers within a given scope. The scope determines the part of the rendezvous system that is used by the network. The three simplistic, topology-oriented cases, reflecting the current usage, are *link-local*, *intra-domain*, and *inter-domain* scopes. The multiple scopes of rendezvous are depicted in Figure 4. However, we expect that future applications will use more semantically-based scopes, implementing, e.g., scopes based on social networking structures.

Due to its importance in policy enforcement and defining (often user-created) information scopes in various situations, the rendezvous system constitutes a relatively well-defined environment where tussle is likely to commence [Cla2002]. The rendezvous system is therefore a policy-enforcement point in the architecture and a mechanism for supporting freedom of choice for network end points. Similar rendezvous functionality has been used in many distributed systems, for example the HIP [Mos2008] [Egg2004], IP multicast [Dee1998], i3 [Sto2002] and Hi3 [Nik2004], FARA [Cla2003], PASTRY [Row2001], and HERMES [Pie2004].



**Fig. 4.** Network architecture with rendezvous.

Rendezvous state is created to allow the subscriptions and publications to meet within a specified scope. Subscriptions and pre-publication notifications (or *advertisements*), possibly included in metadata or data-correlation notifications, may be utilised by the rendezvous system to create partial forwarding state from publishers towards subscribers. When a publication becomes active, i.e., when there is both an actively sending publisher and one or more active subscribers, the rendezvous systems are used to complete the forwarding path by mapping the rendezvous identifier to intra-domain and inter-domain forwarding identifiers. This late mapping can be used to implement dynamic policies, such as inter-domain routing and caching.

The rendezvous system ensures that neither traffic policies nor publication/subscription policies and scopes are violated. We will cover the efficiency, policy, and incentive issues related to the inter-domain rendezvous in detail in our forthcoming publication. Replication is used as the main method to achieve resilience against failures in the rendezvous system.

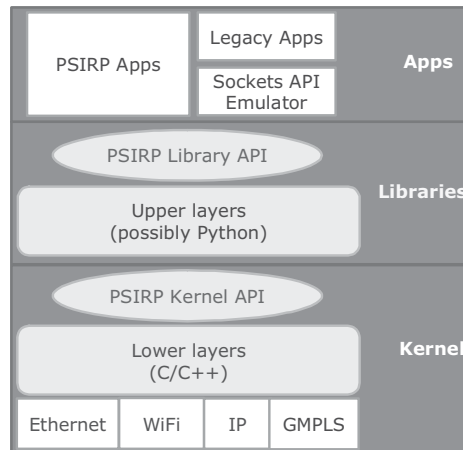
Routing processes in PSIRP are categorized as either intra-domain or inter-domain. Intra-domain routing is concerned with delivery within an administrative domain. Inter-domain routing pertains to data delivery in the global network, typically spanning several domains. FIDs specify the policy compliant paths on the level of domains which makes them more tolerant of router failures along the path.

A subset of the forwarding routers may store cached copies of publications for faster access and time-decoupled multicast lessening the burden of the publisher. The publications are persistently stored by the publishing nodes and the network state can be considered to be fully a soft state that can be recovered in the case of a failure.

Both the rendezvous process and the payload forwarding can be protected by cryptographic means to provide integrity and authenticity of information on packet level as described in [Lag2008].

## 4 Prototype Implementation

The implementation work has focused on an intra-domain implementation of the PSIRP architecture based on Ethernet. The modular prototype implementation structure is illustrated by Figure 5. The implementation and experimentation work is currently on-going.



**Fig. 5.** Implementation structure of the PSIRP prototype.

The core PSIRP prototype will be implemented in two layers:

1. The lower layer will mainly reside in kernel space for performance reasons, providing whatever functionality is deemed to be critical enough to justify its inclusion into the kernel. Lower-level protocols such as Ethernet, Wifi, and IP are included for legacy support; GMPLS offers advanced label-switching

and traffic engineering functions that positively compliment PSIRP's technical ambitions.

2. The upper layer(s) will reside exclusively in user space for reasons of flexibility, providing whatever functionality is deemed necessary to support writing PSIRP applications, supplementing the functionality provided by the lower layer.

## 5 Conclusions

This paper presents the first results of the architectural design process within the PSIRP project. It can only be seen as the first step in the desired clean slate redesign of the Internet. We envision a process of bottom-up lessons learned and top-down rationalization, the first results of which are visible in this report.

Following this methodology, the conceptual architecture and components presented in this paper are only part of our progress so far in the project. The clarification of concepts, presented in the design considerations, as well as the formulation of new questions pushing forward our future development, are central to the background work we have made. Hence, many of the issues addressed in this paper, such as identifiers, the concept of scope, rendezvous, caching, forwarding and transport as well as our inter-domain routing solution, will see further progress in the near future.

## Acknowledgements

This work has been funded by the EU FP7 PSIRP project (Contract INFSO-ICT-216173). The paper is based on the work of the PSIRP architecture and implementation teams, including Pekka Nikander (LMF), Dirk Trossen (BT), Trevor Burbridge (BT), András Zahemszky (ETH), Jarno Rajahalme (NSN), Dmitrij Lagutin (TKK-HIIT), Mikko Särelä (LMF), Janne Riihijärvi (RWTH), and Teemu Rinta-aho (LMF).

## References

- [Ber2001] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web," *Scientific American Magazine*, May 17, 2001, available at: <http://www.sciam.com/article.cfm?id=the-semantic-web> [Accessed on 15 July, 2008].
- [Blu2001] M. Blumenthal and D. Clark, "Rethinking the design of the Internet: The End-to-End arguments vs. The Brave New World," *ACM Transactions on Internet Technology* 2001, vol. 1, issue 1, 2001, pp. 70-109.
- [Cla2002] D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet," in *Proc. ACM SIGCOMM Conference on*

- Applications, Technologies, Architectures, and Protocols for Computer Communications*, New York, NY, Aug. 2002, pp. 347-356.
- [Cla2003] D. Clark, R. Braden, A. Falk, and V. Pingali, "FARA: Reorganizing the Addressing Architecture," in *Proc. ACM SIGCOMM Workshop on Future Directions in Network Architecture*, Karlsruhe, Germany, Aug. 2003, pp. 313-321.
- [Dee1998] S. Deering, D. Estrin, D. Farinacci, M. Handley, A. Helmy, V. Jacobson, C. Liu, P. Sharma, D. Thaler, and L. Wei, *Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification*, IETF RFC 5201, June 1998.
- [Egg2004] L. Eggert and J. Laganier, *Host Identity Protocol (HIP) Rendezvous Extension*, IETF RFC 5204, April. 2008.
- [Eug2003b] P. T. Eugster, P. A. Felber, R. Guerraoui, and A. M. Kermarrec, "The Many Faces of Publish/Subscribe," *ACM Computing Surveys (CSUR)*, vol. 35, issue 2, pp. 114-131, 2003.
- [Hag2007] Huggle partners, "Deliverable D1.2: Specification of the CHILD Huggle," *HAGGLE*, Aug. 2007, available at: [http://www.huggleproject.org/deliverables/D1.2\\_final.pdf](http://www.huggleproject.org/deliverables/D1.2_final.pdf) [Accessed on 15 July, 2008].
- [Jac2006] V. Jacobson, "If a Clean Slate is the Solution What Was the Problem?," *Stanford "Clean Slate" Seminar*, Feb. 2006, available at: <http://cleanslate.stanford.edu/seminars/jacobson.pdf> [Accessed on 15 July, 2008].
- [Kat2006] S. Katti, H. Rahul, W. Hu, D. Katabi, M. Medard, and J. Crowcroft, "XORs in The Air: Practical Wireless Network Coding," *ACM SIGCOMM Computer Communication Review*, vol. 36, issue 4, pp. 243-254, 2006.
- [Lag2008] D. Lagutin, "Redesigning Internet - The packet level authentication architecture," licentiate's thesis, Helsinki University of Technology, Finland, June 2008.
- [Mos2008] R. Moskowitz, P. Nikander, P. Jokela, and T. Henderson, *Host Identity Protocol*, IETF RFC 2362, Apr. 2008.
- [Nik2004] P. Nikander, J. Arkko, and B. Ohlman, "Host Identity Indirection Infrastructure (Hi3)", in *Proc. Second Swedish National Computer Networking Workshop (SNCNW 2004)*, Karlstad, Sweden, 2004.
- [Per2002] A. Perrig, R. Canetti, J. D. Tygar, and D. Song, "The TESLA Broadcast Authentication Protocol," *CryptoBytes*, vol. 5, no. 2, Summer/Fall 2002, pp. 2-13.
- [Pie2004] P. R. Pietzuch, "Hermes: A Scalable Event-Based Middleware," doctoral dissertation, Computer Laboratory, Queens' College, University of Cambridge, Feb. 2004.
- [Psi2008] M. Ain, S. Tarkoma, D. Trossen, P. Nikander (eds.), "PSIRP Deliverable D2.2: Conceptual Architecture of PSIRP Including Subcomponent Descriptions", available at <http://www.psirp.org/>.
- [Psi2008b] D. Trossen (ed.), "PSIRP Vision Document", available at <http://www.psirp.org/> [Accessed on 15 July, 2008].
- [Row2001] A. Rowstron and P. Druschel, "Pastry: Scalable, Decentralized Object Location and Routing for Large-Scale Peer-to-Peer Systems," in *Proc. of Middleware 2001*, Heidelberg, Germany, Nov. 2001, pp. 329-250.
- [Sal1984] J. H. Saltzer, D. P. Reed, and D. D. Clark, "End-to-End Arguments in System Design," *ACM Transactions on Computer Systems*, vol. 2, issue 4, pp. 277-288, Nov. 1984.
- [Sto2002] I. Stoica, D. Adkins, S. Zhuang, S. Shenker, and S. Surana, "Internet indirection infrastructure," in *Proc. 2002 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*, Pittsburgh, PA, Aug. 2002, pp. 73-86.

## Management Architecture and Systems for Future Internet Networks

A. Galis<sup>1</sup>, S. Denazis<sup>2</sup>, A. Bassi<sup>2</sup>, P. Giacomini<sup>2</sup>, A. Berl<sup>3</sup>, A. Fischer<sup>3</sup>, H. de Meer<sup>3</sup>, J. Strassner<sup>4</sup>, S. Davy<sup>4</sup>, D. Macedo<sup>5</sup>, G. Pujolle<sup>5</sup>, J. R. Loyola<sup>6</sup>, J. Serrat<sup>6</sup>, L. Lefevre<sup>7</sup>, A. Cheniour<sup>7</sup>

<sup>1</sup> University College London U.K. {a.galis@ee.ucl.ac.uk},

<sup>2</sup> Hitachi France {spyros.denazis, alessandro.bassi}@hitachi-eu.com, {yrz@anche.no}

<sup>3</sup> University of Passau Germany {andreas.berl, andreas.fischer, demeer}@uni-passau.de

<sup>4</sup> Waterford Institute of Technology Ireland {jstrassner, sdavy}@tssg.org

<sup>5</sup> Lip6 France {daniel.macedo@rp.lip6.fr, Guy.Pujolle@lip6.fr}

<sup>6</sup> UPC Spain {jrloyola,j.serrat}@tsc.upc.edu

<sup>7</sup> INRIA France {laurent.lefevre, abderhaman.cheniour}@ens-lyon.fr

**Abstract** — This paper presents a new autonomic management architectural model consisting of a number of distributed management systems running within the network, which are described with the help of five abstractions and distributed systems: Virtualisation, Management, Knowledge, Service Enablers and Orchestration Planes. The envisaged solution is applicable to the management design of Future Internet as a service and self-aware network, which guarantees built-in orchestrated reliability, robustness, context, access, security, service support and self-management of the communication resources and services.

**Keywords** — Service and Self-aware Network Management, Autonomicity, Virtualisation, Management Plane, Knowledge Plane, Service Enablers Plane, Orchestration Plane.

### 1. INTRODUCTION AND FRAMEWORK

Networks are becoming service-aware. Service awareness means not only that all content and service logic pertaining to a service are delivered but also that all business or other service characteristics (e.g. QoS, SLAs) offer are fulfilled and the network resources are optimally used in the service delivery. In addition, the network's design is moving towards a different level of automation, autonomicity and self-management. The envisaged solution for Future Internet is a service and self-aware network, which guarantees built-in orchestrated reliability, robustness, mobility, context, access, security, service support and self-management of the communication resources and services. We suggest a transition from a service agnostic Internet to service- and self-aware Internet managing resources by applying Autonomic principles as depicted in Figure 1. In order to achieve the objective of service-aware and self-aware networking resources and to overcome the ossification of the current Internet, the Autonomic Internet (AutoI) project [1][2][3][4][17] aims to develop a self-managing virtual resources that can span across heterogeneous networks and that supports service mobility, security, quality of service and reliability. In this overlay network, multiple virtual networks could co-exist on top of a shared substrate with uniform

control. One of the main research challenges for designing a new service-aware network is the inclusion of capabilities such as self-awareness, self-network knowledge and self-service knowledge. These capabilities are used to facilitate continuous tuning of the networks, adaptation to unpredictable conditions, prevention and recovery from failures and provision of a dependable environment.

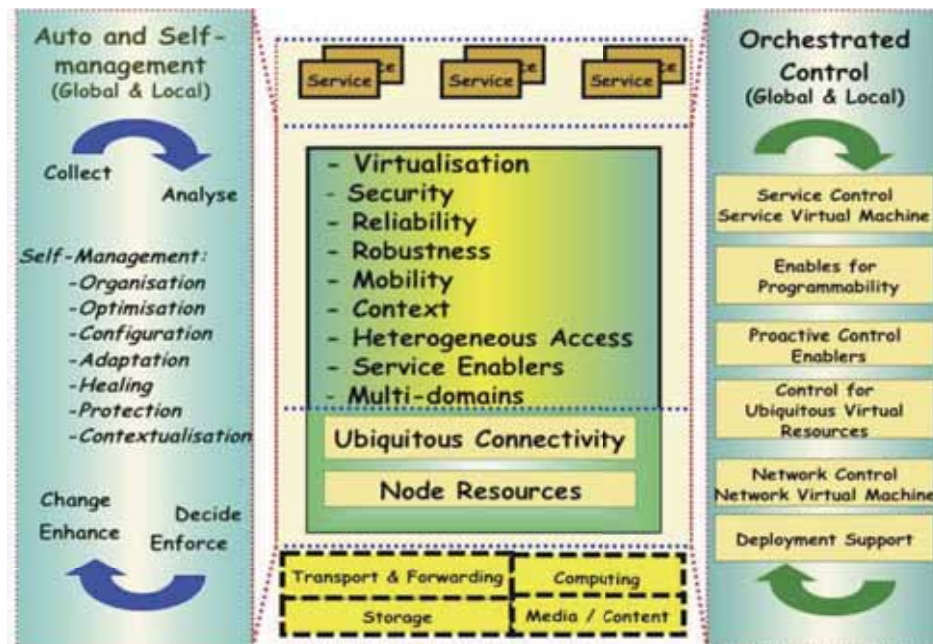


Fig. 1 - Autonomic Internet

To meet the above needs and to support the design of the self-managing virtual resources overlay that can span across heterogeneous networks where multiple virtual networks co-exist on top of a shared substrate with uniform control we proposed a new autonomic management architectural model. It consists of a number of distributed management systems within the network, which are described with the help of five abstractions and distributed systems - the OSKMV planes: Orchestration Plane (OP), Service Enablers Plane (SP), Knowledge Plane (KP), Management Plane (MP) and Virtualisation Plane (VP). Together these distributed systems form a software-driven control network infrastructure that will run on top of all current networks and service infrastructures.

Figure 2 depicts the network and management resources as distributed in the OSKMV planes. The OSKMV planes are new higher-level artefacts to make the Future Internet more intelligent with embedded management functionality, including self-knowledgeable, self-diagnosing and ultimately fully self-managing. At one level the OSKMV planes gather observations, constraints and assertions, and apply rules to these to generate service-aware observations and responses. They are embedded on network hosts, devices, attachments and servers.



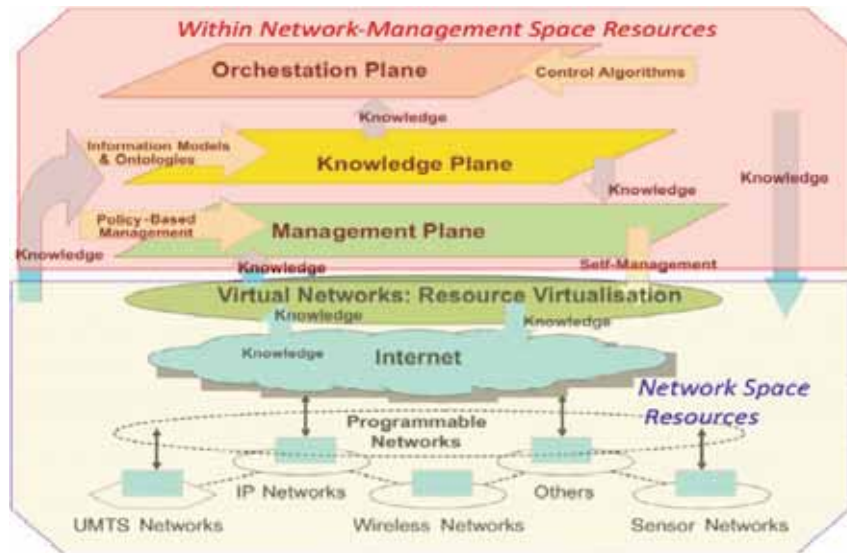


Fig. 2 - AutoI Planes

### 1.1 Architectural Model Overview

AutoI framework, as presented in the figure 3, consists of a number of distributed management systems described with the help of the OSKMV planes. Together these distributed systems form a software-driven network control infrastructure that will run on top of all current networks and application service physical infrastructures to provide an autonomic virtual resource overlay.

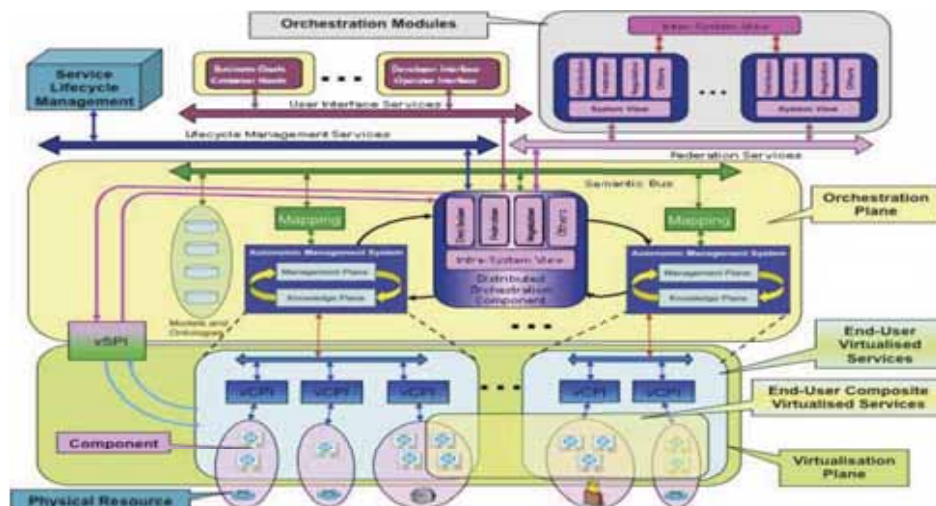


Fig. 3 – AutoI Framework

## 1.2 Orchestration Plane

The purpose of the Orchestration Plane is to govern and integrate the behaviours of the system in response to changing context [14] and in accordance with applicable business goals and policies. It supervises and it integrates all other planes' behaviour insuring integrity of the Future Internet management operations. The Orchestration Plane can be thought of as a control framework into which any number of components can be plugged into in order to achieve the required functionality. These components could have direct interworking with control algorithms, situated in the control plane of the Internet (i.e. to provide real time reaction), and interworking with other management functions (i.e. to provide near real time reaction).

The Orchestration Plane would supervise the optimisation and the distribution of knowledge within the Knowledge Plane to ensure that the required knowledge is available in the proper place at the proper time. This implies that the Orchestration Plane may use either very local knowledge to deserve a real time control as well as a more global knowledge to manage some long-term processes like planning. The Orchestration Plane would host one or more Autonomic Management Systems (AMSs) and it is made up of one or more Distributed Orchestration Components (DOCs), and a dynamic knowledge base consisting of a set of models and ontologies and appropriate mapping logic. Each AMS represents an administrative and/or organisational boundary that is responsible for managing a set of devices, subnetworks, or networks using a common set of policies and knowledge. The AMSs access a dynamically updateable knowledge base, which consists of a set of models and ontologies. A set of DOCs enable AMSs to communicate with each other. The Federation bus enable the Orchestration Plane to be composed with other Orchestration Planes.

The Orchestration Plane acts as control workflow for all AMSs ensuring bootstrapping, initialisation, dynamic reconfiguration, adaptation and contextualisation, optimisation, organisation, closing down of AMSs. The Orchestration Plane provides assistance for the Service Lifecycle Management, namely during the actual creation, deployment, activation, modification and in general, any operation related to the application services or management services.

### *Autonomic Management System*

In the current Internet, the data, control, and management, planes are bound together and often use the same links. For example, TCP connection setup control messages and SNMP management messages use the same links as the data messages. This has at least three drawbacks: i. the data plane is limited to packet-oriented data; ii. the design of each of the three planes is unduly complicated, and iii. inherent security risk exist, since it is relatively easy to get at control and management data by simply attacking the data plane path. A key advantage of the AutoI architecture is that it can provide a programmable mix of isolation and sharing of network resources. Each AMS consists of a management plane and a knowledge plane, as well as interfaces to a dedicated set of models and ontologies and interfaces to one or more Distributed Orchestration Components. Mapping logic enables the data stored in models to be transformed into knowledge and combined with knowledge stored in ontologies to provide a context-sensitive assessment of the operation of one or more virtual resources. Another set of interfaces enables framework services, such as directory services, naming, federation, and others, to be used by the AMS.



An AMS collects appropriate monitoring information from the virtual and non-virtual devices and services that it is managing, and makes appropriate decisions for the resources and services that it governs, either by itself or in collaboration with other AMSs, as explained in the next section. The DOC provides a set of framework network services [15]. Framework services provide a common infrastructure that enables all components in the system managed by the Orchestration Plane to have plug-and-play and unplug-and-play behaviour. Applications compliant with these framework services share common security, metadata, administration, and management services. The DOC enables the following functions across the orchestration plane:

- Federation: each AMS is responsible for its own set of virtual and non-virtual resources and services that it governs as a domain. Federation enables a set of domains to be combined into a larger domain.
- Negotiation: each AMS advertises a set of capabilities (i.e., services and/or resources) that it offers for use by other components in the Orchestration Plane.
- Distribution: the DOC provides communication and control services that enable tasks to be split into parts that run concurrently on multiple AMSs within an Orchestration Plane, or even across multiple Orchestration Planes.
- Governance: each AMS can operate in an individual, distributed, or collaborative mode. It collects appropriate monitoring data in order to determine if the virtual and non-virtual resources and services that it governs need to be reconfigured. Business goals, service requirements, context, capabilities and constraints are all considered as part of the decision making process.
- Autonomicity: AMSs acting, as individual entities are responsible for managing their resources and services, and send status messages to other AMSs
- Views: i. Intra - Future Internet System View provides an overall, composite view of the system as seen by the components within a given Orchestration Plane; ii. Inter - Future Internet System View provides an overall, composite view of Orchestration Planes that are collaborating, as in a multiple domain system.

### 1.3 Service Enablers Plane – Life Cycle Management

The AutoI architecture defines a common mechanism for the identification and specification of a business problem, along with the specification and development of a deployable solution, which enables services to be built and deployed. This is necessary due to the changing conditions in which a service is provisioned; hence, the service and in particular management service must be lifecycle managed as it evolves and responds to these changes. The Service Enablers Plane (SP) consists of functions for the automatic (re) deployment of new management services, protocols as well as resource – facing (i.e. QoS functions) and end-user facing service. It includes the enablers to allow code to be executed on the network entities. The safe and controlled deployment of new code enables new services to be activated on demand. This approach has the following characteristics: i. Service (re)deployment is taking place automatically and allows a significant number of new services to be offered on demand; ii. Special management functions and services can be easily enabled locally for testing purposes before they are automatically deployed network-wide; iii. Eases the deployment of network-wide protocol stacks and management services; iv. To enable secure but controlled execution environments; v. An automatic decision making infrastructure that guides the deployment of new tested network services; vi. Optimised resource utilization of the new services and the system.

#### 1.4 Knowledge Plane

AutoI introduces a focused functionality knowledge plane, consisting of models and ontologies, to provide increased analysis and inference capabilities; its purpose is to provide knowledge and expertise to enable the network to be self-monitoring, self-analyzing, self-diagnosing, and self-maintaining or -improving. AutoI's KP brings together widely distributed data collection, wide availability of that data, and sophisticated and adaptive processing or KP functions, within a unifying structure that brings order, meets the policy [8], scaling and functional requirements of a global network, and, ideally, creates synergy and exploits commonality of design patterns between the many possible KP functions. The main KP components are an information and context service plus models and ontologies, which enable the analysis and inferencing capabilities. Knowledge extracted from information/data models forms facts. Knowledge extracted from ontologies is used to augment the facts, so that they can be reasoned about. The information and context service provides: i. information-life cycle management (storage, aggregation, transformations, updates, distribution) all information and context in the network and addresses the size and scope of the Internet; ii. responsiveness to requests made by the AMSs; iii. triggers for the purpose of contextualisation of AMSs (supported by the context model of the information model); iv. support for robustness enabling the KP to continue to function as best possible, even under incorrect or incomplete behaviour of the network itself; v. support of virtual networks and virtual system resources in their needs for privacy and other forms of local control, while enabling them to cooperate for mutual benefit in more effective network management.

#### 1.5 Management Plane

The Management Plane consists of AMSs, which are designed to meet the following design objectives and functionality: i. Embedded (Inside) Network functions: The majority of management functionality should be embedded in the network and it is abstracted from the human activities. As such the AMSs will run on execution environments on top of virtual networks and systems, which run on top of all current network and service physical infrastructures.; ii. Aware and Self-aware functions: It monitors the network and operational context as well as internal operational network state in order to assess if the network current behaviour serve its service purposes.; iii. Adaptive and Self-adaptive functions: It triggers changes in network operations (state, configurations, functions) function as a result of the changes in network and service context.; iv. Automatic self-functions: It enables self-control (i.e. self-FCAPS, self-\*) of its internal network operations, functions and state. It also bootstraps itself and it operates without manual external intervention. Only manual/external input is provided in the setting-up of the business goals; v. Extensibility functions: It adds new functions without disturbing the rest of the system ((Un)Plug\_and\_Play/ Dynamic programmability of management functions & services); vi. Simple cost functions: Minimise life-cycle network operations' costs and minimise energy footprint.

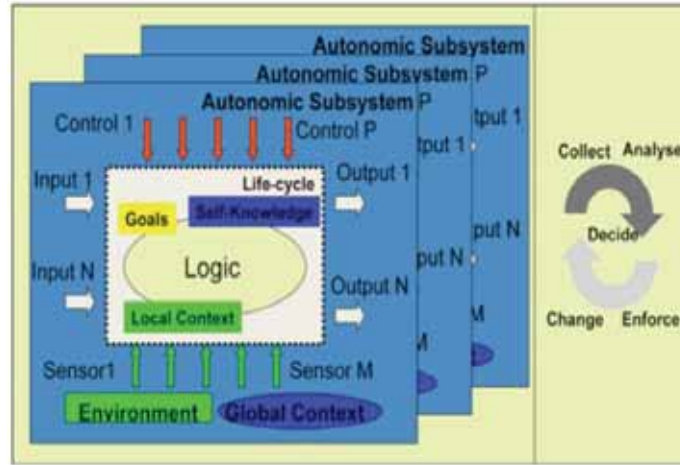


Fig. 4 – Control Loops

In addition the Management Plane, as it governs all virtual resources, is responsible for the optimal placement and continuous migration of virtual nodes and virtual servers into hosts (i.e. physical nodes and servers) subject to constraints determined by the Orchestration Plane. The AMSs are design to follow the autonomic control loops depicted in the figure 4.

### 1.6 Virtualization Plane

One of the key requirements that differentiate AutoI from other efforts is its emphasis on virtualisation (i.e., the abstraction) of resources and services, which are controlled by the MP. Virtualisation hides the physical characteristics of the computing resources being used from its applications and users. AutoI uses platform virtualisation to provide virtual services and resources. Platform virtualisation separates an operating system from its underlying platform resources; resource virtualisation abstracts physical resources into manageable units of functionality (e.g., the concept of a virtual router, where a single physical router can support multiple independent routing processes by assigning different internal resources to each routing process).

The virtualisation plane consists of software mechanisms to treat selected physical resources as a programmable pool of virtual resources that can be organised by the Orchestration Plane into appropriate sets of virtual resources to form components (e.g., increased storage or memory), devices (e.g., a switch with more ports), or even networks. The organisation is done in order to realise a certain business goal or service requirement. Two special interfaces, as shown in Figure 5, called the Virtualisation System Programming Interface (vSPI) and the Virtualisation Component Programming Interface (vCPI) are under development.

The vSPI is used to enable the Orchestration Plane to govern virtual resources, and to construct virtual services and networks that meet stated business goals having specified service requirements. The vSPI contains the system-view of the virtual resources that a particular Orchestration Plane governs, and is responsible for orchestrating groups of virtual resources in response to changing user needs, business requirements, and

environmental conditions.

The vSPI is responsible for determining what portion of a component (i.e., set of virtual resources) are allocated to a given task. This means that all or part of a virtual resource can be used for each task, providing an optimised partitioning of virtual resources according to business need, priority, and other requirements. Composite virtual services can thus be constructed using all or part of the virtual resources provided by each physical resource, as shown in Figure 5. The vSPI monitors system-level status of the virtual resources that it governs. This is different than the vCPI, which monitors “micro-level” status of the virtual resources that it configures. For example, the vSPI is responsible for informing the AMS that a particular virtual resource is ready for use, whereas the vCPI is responsible for informing the AMS that a particular virtual resource has been successfully reconfigured.




Fig. 5 – Virtualisation Interfaces

Each physical resource has an associated and distinct vCPI. This enables the AMS to manage the physical resource, and to request virtual resources to be constructed from that physical resource by the vCPI of the Virtualisation Plane. The AMS sends device-independent commands via the vCPI, which are translated into device-specific commands that reconfigure the physical resource. The vCPI also provides monitoring information from the virtual resources back to the AMS that controls that physical resource. It is responsible for providing dynamic management data to its governing AMS that states how many virtual resources are currently instantiated, and how many additional virtual resources of what type can be supported.

## 2. RELATED WORK

*Autonomic Computing and Communications* - A comprehensive state of the art review in Autonomic computing and communications is provided in [16]. The purpose of autonomic computing and networking [5][6] is to manage complexity. By transferring more manual functions to involuntary control, additional resources (human and otherwise) are made available to manage higher-level processes. One difference between autonomic computing and autonomic networking is that the latter must cope with and coordinate multiple control mechanisms, such as those used by different types of networks, which the former usually does not consider.

*Knowledge Plane* – The Knowledge Plane [9] was proposed as a research objective to build “a fundamentally different sort of network that can assemble itself given high level

instructions, reassemble itself as requirements change, automatically discover when something goes wrong, and automatically fix a detected problem or explain why it cannot do so.” The Knowledge Plane approach organised functionality into two planes. The data plane was responsible for packet forwarding; the Knowledge plane was a new construct “that builds and maintains high-level models of what the network is supposed to do, in order to provide services and advice to other elements of the network.”

*Foundation Observation Comparison Action Learn Reason (FOCALE)* - This architecture [7] is equally appropriate for legacy devices and applications as well as for next generation and cognitive networks.

*Inference Plane* - This approach [10] was created to solve some of the problems with the Knowledge Plane [9]. The Inference Plane is a coordinated set of intelligent decision-making components that represent the capabilities of the computing elements being controlled, the constraints placed upon using different functions, and the context in which they are being used.

*Ambient Networks* - AN is an EU sponsored project ([www.ambient-networks.org](http://www.ambient-networks.org)), which envisaged the development of a software-driven network control infrastructure for wireless and mobile networks. The concept of the Ambient Control Space (ACS) is introduced to encompass all control functions in a certain domain, which can be used as a control plane overlay to integrate and interoperate seamlessly any existing networks [11] [12].

*Autonomic Network Architecture* - ANA is an EU sponsored project ([www.ana-project.org](http://www.ana-project.org)), which aims at exploring novel ways of organizing and using networks beyond legacy Internet technology. ANA should be regarded as an architecture for autonomic devices.

*4WARD* is an EU sponsored project ([www.4ward-project.eu](http://www.4ward-project.eu)), which aims to make the development of networks and networked applications faster and easier, including an in-network management approach.

*Programmable Network Management* - Programmable networks techniques allow software components or programs to be activated or injected into network components, which run customised tasks on the data passing through these components. They are especially attractive in the realisation of real-time activation, adaptations and thus, the implementation of real-time capable control loops. Full review of the state of the art in programmable networks and their application to management of networks is in [13].

*FIND* (Future Internet Design) is a new long-term research program of the National Science Foundation - [www.nets-find.net](http://www.nets-find.net)). The philosophy of the program is to help conceive the future by momentarily letting go of the present - freeing our collective minds from the constraints of the current state.

*GENI* (Global Environment for Network Innovation Program - [www.geni.net](http://www.geni.net)) is a research program addressing serious problems facing to-day Internet: inadequate security, reliability, manage-ability and evolvability.

### 3. CONCLUSION

This paper presents the needs for a self-managing virtual resource overlay that can span across heterogeneous networks that can support service mobility, security, quality of service and reliability as part of Future Internet. In this overlay network, multiple virtual networks co-exist on top of a shared substrate with uniform control. In support of the

design of such overlay a new management architectural and system model for Future Internet, which is under development in the AutoI [1] project, is presented. It consists of a number of distributed management systems within the network, which are described with the help of five abstractions and distributed systems - the OSKMV planes: Virtualisation Plane (VP), Management Plane (MP), Knowledge Plane (KP), Service Enablers Plane (SP) and Orchestration Plane (OP). Together these distributed systems form a software-driven control network infrastructure that run on top of all current network and service infrastructures.

## ACKNOWLEDGMENT

This work was undertaken in the context of the Autonomic Internet FP7 project [1][2][3][4], which is partially financed by the European Union. We are acknowledging the constructive comments and discussions with Lefteris Mamatas, Meriem Abid, Zohra Boudjemil, Jean-Patrick Gelas, Odysseas Koufopavlou, Zeinab Movahedi, Martín Serrano, and Hubert Zimmermann.

## REFERENCES

- [1] Autonomic Internet Project <http://ist-autoi.eu/autoi/>
- [2] Cheng, L., Galis, A., Mathieu, B., Jean, K., Ocampo, R., Mamatas, L., Loyola, Serrat, J. R., Berl, A., de Meer, H., Davy, S., Movahedi, Z., Lefevre, L., - "Self-organising Management Overlays for Future Internet Services"- IEEE Manweek 2008 /MACE 2008; 22-26 Sept. 2008, Samos, Greece/  
<http://www.manweek.org/2008/index.php?lang=en>
- [3] Berl, A., Fischer, A., de Meer, H., Galis, A., Loyola, J.R.-" Management of Virtual Networks"- IEEE Manweek 2008/ EVGM 2008; 22-26 Sept. 08, Samos Greece/  
<http://www.manweek.org/2008/index.php?lang=en>
- [4] Bassi, A., Denazis, S., Galis, A., Fahy, C., Serrano, M., Serrat, J., -"Autonomic Internet: A Perspective for Future Internet Services Based on Autonomic Principles" - ManWeek 07/MACE 1-2 Nov. 07, San José, California, USA; <http://magellan.tssg.org/2007/mace/mace.php>
- [5] Kephart, J.O., Chess, D.M."The Vision of Autonomic Computing", IEEE Computer", January 03 [research.ibm.com/autonomic/research/papers/](http://research.ibm.com/autonomic/research/papers/)
- [6] IBM, "An Architectural Blueprint for Autonomic Computing" June05, [ibm.com/developerworks/autonomic/library/ac-summary/ac-blue.html](http://ibm.com/developerworks/autonomic/library/ac-summary/ac-blue.html)
- [7] Strassner, J., Agoulmine, N., Lehtihet, E., - "FOCALE – A Novel Autonomic Networking Architecture", ITSSA Journal, Vol. 3, No 1, May 2007, pp 64-79, ISSN 1751-1461
- [8] Davy, S., Jennings, B. Strassner, J., -"Application Domain Independent Policy Conflict Analysis Using Information Models", NOMS 2008, Salvador Bahia, Brasil, 7-11 April 2008
- [9] Clark, D., Partridge, C., Ramming, J., and Wroclawski, J.,-"A Knowledge Plane for the Internet", Proceedings ACM SIGCOMM 2003 Conference on Applications, Technologies, Architectures and Protocols for Computer Communications, Karlsruhe, Germany, pp. 3-10

- [10] Strassner, J., Foghlú, M. Ó, Donnelly, W., Agoulmine, N.,- “Beyond the Knowledge Plane: An Inference Plane to support the Next Generation Internet”, Proceedings of First International Global Information Infrastructure Symposium (GIIS), 2-6 July 2007, Marrakech, Morocco.
- [11] Mathieu, B., Song, M., Galis, A., Cheng, L., Jean, K., Ocampo, R., et all ”Autonomic Management of Context-Aware Ambient Overlay Networks” - IEEE ChinaCom, 15-17 August 07, Shanghai, [www.chinacom.org/2007](http://www.chinacom.org/2007)
- [12] Mathieu, B., Song, M., Galis, A., Cheng, L., Jean, K., Ocampo, R., Lai, Z., Brunner, M., Stiernerling, M., Cassini, M. -”Self-Management of Context-Aware Overlay Ambient Networks - 10th IFIP/IEEE IM 2007; 21-25 May 2007, Munich, Germany; <http://www.im2007.org>
- [13] Galis, A., Denazis, S., Brou, C., Klein, C.-”Programmable Networks for IP Service Deployment” ISBN 1-58053-745-6; pp450, June 2004; Artech House Books
- [14] Raz, D., Juhola, A., Serrat, J., Galis, A., ”Fast and Efficient Context-Aware Services” ISBN 0-470-01668-X; pp250, March 2006; John Wiley & Sons, Ltd.
- [15] Lefèvre, L., “Heavy and lightweight dynamic network services”- The 7th International Symposium on Autonomous Decentralized Systems, Chengdu, Jiuzhaigou, China, April 05.
- [16] Curran, K., Mulvenna, M., Galis, A., Nugent, C. -”Challenges and Research Directions in Autonomic Communications” - International Journal of Internet Protocol Technology (IJIPT) - Vol. 2 No. 1; Jan 07; SSN (Online): 1743-8217- ISSN (Print): 1743-8209; [www.inderscience.com/browse/index.php?journalID=144&year=2007&vol=2&issue=1](http://www.inderscience.com/browse/index.php?journalID=144&year=2007&vol=2&issue=1)
- [17] Davy, S., Fahy, C., Griffin, L., Boudjemil, Z., Berl, A., Fischer, A., de Meer, H., Strassner, J. - “Towards a Policy-based Autonomic Virtual Network to support Differentiated Security Services” — International Conference on Telecommunications and Multimedia (TEMU) 2008; July 16 – 18, 2008; Ierapetra, Crete, Greece. <http://www.temu.gr/2008/>



## **Towards a Future Internet: Node Collaboration for Autonomic Communication**

Tanja Zseby, Thomas Hirsch, Michael Kleis, and Radu Popescu-Zeletin

Fraunhofer FOKUS, Berlin, Germany

{tanja.zseby, thomas.hirsch, michael.kleis,  
radu.popescu-zeletin}@fokus.fraunhofer.de

**Keywords:** Future Internet, Situation Awareness, Collaboration Strategies

**Abstract.** The Internet today is a complex agglomerate of protocols that inherits the grown legacies of decades of patchwork solutions. Network management costs explode. Security problems are more pressing than ever, as organized crime discovers its value. The application and user demands on the Internet are increasing with mobile technologies and media content on the rise, all the while the number of participating nodes is equally boosting. As a direct consequence the recently triggered research on concepts for the future Internet has to cope with a high complexity at network layer and significance in mission critical service infrastructures of society. As part of this effort, the research field of autonomic communication (AC) aims at network self-management and self-protection, following the autonomic computing paradigm invented by IBM. We argue that the collaboration of network nodes provides a valuable way to address the corresponding challenges. After an in-depth analysis of the problem space, we outline in this paper the advantages and challenges of collaboration strategies in deployment. We present the Node Collaboration System (NCS) developed at Fraunhofer FOKUS for the experimental investigation of collaboration strategies and show how the system can be used in a simple setting for network self-protection.

### **1 Introduction**

When the Internet was designed, application requirements were low compared to today's standards. Mobility of network components was no issue and neighbor nodes were assumed trustworthy.

The world has changed. Nowadays, many communities and businesses rely on the Internet and demand mobility, quality of service, authorization, accounting and more to support application demands. Moreover, criminals control large parts of the Internet. A wide variety of attacks on network nodes, servers and end systems endanger the operation of components at their will and urgently call for secure solutions. Complexity and significance of critical networks let the costs for network administration explode.

Although the Internet Protocol (IP) is still the common denominator for communication in the Internet, we observe a growing patchwork of protocols deployed to serve

the needs of the increasing number of applications and different underlying communication technology. About 120 working groups within the Internet Engineering Task Force (IETF) standardize protocol after protocol to fulfill those demands. IP is still the dominant protocol but only on the networking layer for the data plane. Already on transport layer heterogeneity is growing. UDP and TCP used to be prevalent. But new protocols like the Stream Control Transmission Protocol (SCTP) and the Datagram Congestion Control Protocol (DCCP), that mix desired properties of both protocols, gain significance. Furthermore many adaptations to TCP have been proposed to address TCP problems in fixed networks and wireless environments.

On the control plane, complexity is even larger. A wide variety of control protocols to support IP (ICMP, ARP, DHCP, etc.), to provide routing (BGP, OSPF, ad hoc and multicast routing, etc.), QoS (RSVP, DiffServ) and security features (SSL, TLS, AAA) is required to operate today's Internet.

The new IP version IPv6 has been defined to cope with several problems of IPv4, most notably with the upcoming IP address shortage. However, due to the difficulties of migration and legacy support of IPv4, providers and users are switching but slowly to the new technology, and huge efforts are needed to organize co-existence of IPv4 and IPv6.

Future problems can be foreseen: Embedded devices in household and daily life applications become Internet aware. The representation of critical networks in telecommunication, health and government onto IP networks progresses quickly. Worldwide Internet connectivity is increasing and the considerations on green IT have imposed new requirements.

Future Internet initiatives address current problems and future demands of the worldwide network. Approaches span from evolutionary proposals, that target incremental changes to the existing Internet, to revolutionary ideas that plan to design a new Internet from scratch, a process dubbed clean slate design. In this paper we provide a short overview of currently proposed solutions for the future Internet. We argue that the immense administrative costs and the demands for security present the most challenging issues in future networks. We focus on the research field of autonomic communication, which provides a framework to realize self-management and self-protection of network components. Our contribution is the investigation of collaboration strategies to improve such techniques. We introduce and compare different collaboration strategies and show with an example how collaboration helps to realize self-protection.

## 2 Future Internet Trends

Future Internet research is supported by several programs in Europe, US and Asia. In the US research on future Internet and the provisioning of facilities for large scale experiments is funded by the Global Environment for Network Innovations (GENI) and the Future Internet Design (FIND) program.

The European Union also funds several projects on future Internet research and has recently started projects for the establishment of federated testbeds to support experimental research, such as Onelab [1] and PanLab [2]. Several governments support such activities with national funding. In Japan and Korea similar activities can be ob-

served (e.g., AKARI in Japan [3], Future Internet Forum in Korea [4]). The common differentiation between revolutionary and evolutionary paradigms is followed in these programs.

Inspired by IBM's autonomic computing paradigm [5] Fraunhofer FOKUS started in March 2004 an initiative to establish a new research field called autonomic communication as basis for enabling self-properties like self-management and self-protection in future communication networks [6]. Under the lead of Fraunhofer FOKUS a consortium of partners from industry and academia founded the Autonomic Communication Forum (ACF). Now the ACF has become a standardization body that standardizes the theoretic foundations of autonomic communication [7].

In 2005 the EU started a research program on Situated and Autonomic Communication (SAC) to bring forward research in the area. In 2006 four integrated projects started under this program. Although the goals were quite ambiguous and aimed at organizing communication networks with absolutely new paradigms, today results of these projects have not only materialized in sophisticated concepts in paperwork but also in running code. In July 2008 the first public release of the Autonomic Networking Architecture core (ANAcORE) has been released. The ANAcORE substitutes the traditional fixed protocol stack with flexible functional blocks that are concatenated by the network on demand to serve various application needs in a highly heterogeneous underlying network. The concept used in the ANAcORE is called *functional composition*.

An approach to share resources in the future Internet between applications and user groups with different requirements is the concept of *network virtualization*. Virtualization concepts are already used in operating systems to share resources among different tasks. Network virtualization can be seen as a radical advancement of the concept of overlay networks. Overlay networks nowadays already allow to build application- or community-specific structures on top of the current Internet. Virtualization tries to bring this idea further down into the network and generate separated networks by assigning network resources in routers and lines to slices for applications or user groups. With this each of the separate networks can serve the specific needs of the application or community. Virtualization was also proposed as solution to share resources for large scale experiments in distributed testbeds. GENI is following this approach for experimental research. The main challenge for virtualization is the management and conflict-free assignment of resources to different groups. Both, functional composition and virtualization require decision-making based on application demands and current network situation. For this we see Situation Awareness and collaboration strategies as essential building blocks. Other approaches have been motivated by the inherent problem of addressing in the current Internet. Currently IP addresses are assigned to hosts. They serve as identifier and locator at the same time. This leads to problems especially in mobile environments and with multi-homing. The Host Identity Protocol (HIP) and the Locator Identifier Split Protocol (LISP) are evolutionary approaches to split locator and identifier. More radical approaches propose to go towards a *Network of Information*. It is based on the idea that the main interest of users is to get access to information in the Internet. Therefore the proposal is to address information and services instead of hosts. Basic concepts for addressing information are already known for file sharing and content-delivery networks.

In the following, we focus on decision-cycles within the network required to achieve self-management and self-protection. We describe how to achieve Situation Awareness and use collaboration strategies to provide the mentioned self-properties. Where other future Internet approaches require decision-making, we show how concepts from Autonomic Communication may be adopted for the network protection and management of resources in such environments.

### 3 Situation Awareness

In this paper we use Situation Awareness to denote the process of perception of relevant conditions and events. Self- and context awareness are the perception of conditions within and surrounding the actor. In communication networks, Situation Awareness is a pre-requisite to make sound decisions; thus to establish autonomic communication principles. The situational view is the summary of all perceptible conditions and events. Situation Awareness is established on the one hand by observing and analyzing network node behavior and information flows in the direct neighborhood of an entity. On the other hand, collaboration is necessary to provide information on remote events. The situational view provides the basis to decide, based on the current state of the network. If perfect Situation Awareness is achieved, i.e. all important factors for the decision are known and processed with respect to the decision-makers goal, the decision is evident (see Fig. 1). Nevertheless, this ideal case is usually not achievable due to missing information, resource or time constraints. Usually it is necessary to make decisions without perfect Situation Awareness, i.e. with some degree of uncertainty about the situation, in order to invoke actions in time. Situation Awareness can be subdivided into three levels:

- **Perception**: Build the situational view by collecting relevant information.
- **Inference**: Understand the interplay of conditions, as well as other actors patterns and plans.
- **Prediction**: Predict future events and actions.



Fig. 1. Context processing [29]

Implementing Situation Awareness is not a simple task. Network events are extremely dynamic and difficult to be perceived, interfered or predicted. Hence the view of the situation needs to be constantly updated. The utopistic ideal would be to gain a complete picture of the network, and be able to process it; Observe every packet, at every network node, and fully analyze it. Then we could perfectly direct the traffic to avoid congestion and detect even sophisticated application-level attacks. However, since this utopia requires at the very minimum equal processing powers as the rest of the network, we simply cannot measure everything everywhere.

We have to deal with resource limitations. Processing power, storage, transmission capacity and speed are limited. More dramatically, as network supervision is only a support function for network operation, they should not influence network performance at all. Their costs should not exceed costs for network operation itself. Moreover, the overwhelming amount of result data we could retrieve with specialized measurement hardware has to be processed. Resource limitations are grave in terms of transmission and processing power. They are even worse in wireless networks of small embedded devices and low bandwidth. We postulate the following requirements that a system should fulfill in order to establish Situation Awareness:

**Cope with resource constraints** The amount of data traffic carried by the Internet each day has increased dramatically over the last decades. A deceleration of this trend is not in sight. Technologies that allow higher data rates increase not just the amount of data that can be measured but also the quantity of measurement results needing to be processed, stored or transferred per time unit. Approaches to cope with resource constraints are the usage of dedicated hardware (e.g. [8], [9]), the improvement of algorithms for packet processing (e.g., [10], [11]) or the use of data selection techniques ([12], [13]).

**Change viewpoints** The ability to change viewpoints is extremely valuable for establishing Situation Awareness. In order to generate a good picture of the current situation, it is useful to have the option of zooming in or out. The capability to re-configure network observation tasks provides the basis for *adaptive measurement* techniques and is a pre-requisite for resource and accuracy control. Adaptive measurements can be used to tune towards events of interest by changing observation points, classification rules, or aggregation and selection techniques on demand (e.g., [14], [15]).

**Respect privacy concerns** The need to respect privacy concerns is often in contradiction with the desire to get as much information as possible. Privacy concerns need to be respected but they do constrain data collection and information sharing. Fraunhofer FOKUS investigates in privacy-aware monitoring methods in the EU project PRISM [16].

**Cooperate to share information** Sharing information is the prerequisite for learning from others. If one node observes strange or suspicious behavior it is useful to see whether other nodes have observed similar phenomena. If a node is infected by a virus or a worm that spreads within a network, it is worthwhile checking whether neighbor nodes or neighbor networks have experienced similar events in the past. If this is the case, information from neighbors can help to analyze the event, select appropriate countermeasures or nip it in the bud.

As a consequence of above considerations and described challenges we consider collaboration as one of the key enablers for Situation Awareness. Because of this fact we will describe different collaboration strategies in section 4. Since information sharing already is a form of collaboration we further elaborate on this in section 4.1. Fraunhofer FOKUS investigates methods for an efficient and flexible establishment of Situation Awareness in the EU project 4WARD [17].

## 4 Collaboration Strategies

Decision-making requires information on which the decision can be based. However, information in a distributed system cannot be gathered without consent. Hence, collaboration methods are required for information collection and decision-making.

Sharing resources is another benefit of collaboration. Finally, where information is provided by collaboration, privacy can be largely guaranteed by the information provider. We therefore argue that a participative information collection system is one way to handle the previously mentioned challenges.

But collaboration does not come for free. It requires a communication infrastructure, an incentive to cooperate, and means to trust the behavior of other nodes. A solution for collaboration for network protection should scale and is subject to timing requirements from the decision process. Further challenges include the processing of the information from multiple sources, resilience against involuntary inconsistencies and malicious communication, and reaching agreement and conclusive actions for joint decision making.

### 4.1 Collaboration for Information Sharing

The correlation of observations at multiple observation points in the network is essential to get a networkwide view and is further required to calculate network specific metrics as one-way delay or to gather information about internet routing. Existing tools face the challenges of clock synchronization, and the efficient correlation of packet events at different observation points (e.g., [20], [21], [22]). A challenging combination of data selection techniques with multi-point measurements ensures that the same packet is selected at different points. Hash-based selection techniques are proposed in [21] and [13] that aim at an emulation of random sampling to apply statistical methods for accuracy assessment.

Several information sharing strategies help to improve a nodes Situation Awareness to support the decision process. In [23] a system is proposed where neighboring nodes may be searched for specific intrusion detection events. More general *Context information* helps to extend the network centric view. Such information covers data from different OSI Layers such as geo-location, user behavior or external events. Information from network services (e.g. DNS or AAA server) can further improve management and defense strategies [24]. For an example how to model context information we refer to [18].

To share information among network operators is a more difficult challenge. It can help to better identify bottlenecks to track the origin of a failure and to isolate the source

of the problem. It is extremely valuable for network protection since attack patterns can be learned from neighbors, and the path may be traced to the origin of the attack. But privacy and secrecy concerns make sharing of network data difficult. It can reveal information about network structure, users or vulnerabilities to competitors or potential attackers. Another collaboration is the delegation of analysis tasks that helps to make use of free resources, either centrally controlled or decentralized. Data analysis tasks may be shared between entities; strange and suspicious patterns can be forwarded to dedicated analysis components for further inspection. Commercial Intrusion Detection Systems, such as the Cisco Anomaly Detector, take a first step towards specialization of network components within a domain. In their system, anomalous traffic detected by the Anomaly Detector in the network is forwarded to a more specialized DDoS Mitigation component, the Cisco Guard [25]. Sharing information also requires standardized interfaces. In January 2008 the IETF standardized the IP Flow Information Export Protocol (IPFIX) [26] for exporting flow information from routers and network probes. This protocol can be also used for exporting packet information or derived data. In section 5.1 we will illustrate how we use this protocol for the FOKUS Node Collaboration System to share information with neighbor nodes.

## 5 Collaboration Principles for the Future Internet, a Case Study

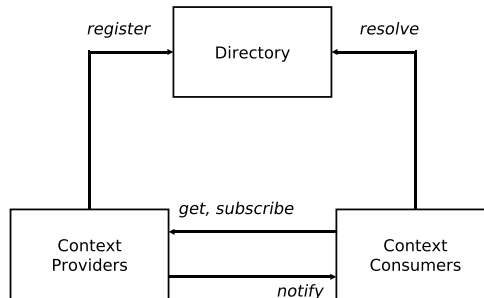
This section presents prototypes of the awareness and collaboration approaches discussed in the previous sections. They serve as a research platform for the investigation of collaboration methods and constraints.

### 5.1 Node Collaboration System (NCS)

Fraunhofer FOKUS has developed a Node Collaboration System (NCS) for the investigation of collaboration strategies for self-management and self-protection. The establishment of Situation Awareness is supported by information sharing among nodes. The system for accessing arbitrary network context information is depicted in Fig. 2. Each node can serve as context provider and provides information to other nodes. Nodes cannot only offer information that they generated by local measurements but also information that they generated by processing of information from other sources. Information is modeled as abstract Context Data Objects. In order to make information accessible by other nodes, context providers register their available Context Data Objects with a directory.

The context providers indicate the context they can provide and the means to access it (e.g. push or pull model, protocols supported by the context provider, etc.). The actual information remains with the context provider. The directory only contains references to its location using Unique Context Identifiers (UCIs). The UCI of a context data object can be seen as simple strings, which offers similar functionality as Unique Resource Identifiers (URIs) in HTTP protocol. It is also possible to register information that does not yet exist, but may be generated by the context provider, for example by invoking local measurements or processing information from other context providers.





**Fig. 2.** Context management (picture taken from [28]) .

The decision-making process that needs to access information from other nodes acts as context consumer and can retrieve the context location and means to access the information with a request to the directory. Since the investigations on collaboration strategies for decision-making are our main focus and not the context distribution, we currently work with a centralized directory for simplicity. It is possible to later distribute the directory over multiple nodes using common name service algorithms.

We currently consider the following collaboration strategies for making joint decisions for network protection.

**Election of a Leader** In this approach the cooperative decision problem is on-demand reduced to a centralized case, but has the ability to flexibly assign who becomes this central node. For ad hoc networks the authors of [27] propose to combine clustering with a periodic leader re-election in each cluster. The actual monitoring is performed by the cluster members which propagate prefiltered monitoring results to the leader for analysis. If the cluster leader detects an intrusion it can coordinate the response.

**Voting** A less central approach are voting systems. For the case of sensor networks, [31] describes a voting system that can be realised without a priori knowledge about node behavior. Each sensor is able to observe its neighbors activities and defines majority sensor behavior as "normal", based on its local view. If one of its neighbors shows abnormal behavior, the observing sensor starts a voting process and presents evidence to its neighbors. Intruders are identified by a majority of responses.

**Emergent Behavior** The authors of [32] study an emergent behavior based collaborative information processing strategy to address the cooperative monitoring problem. Using a model of ants colonies, the actual monitoring results are translated into a pheromone concentration. Thus a path of intrusion in the sensor network can be identified by its pheromone concentration.

The means to transfer information from context provider to context consumer depend on the available transport methods at both entities. NCS provides a proprietary solution to provide efficient transport of context objects but also supports the transport of context objects by the IP flow information export protocol standard [26] if this is supported at the nodes. In this case the context provider acts as IPFIX exporter and the context con-

sumer as IPFIX collector. Integration with other commonly used protocols like SNMP is possible. Due to the flexibility of the IPFIX protocol it can be used without modifications. It is only necessary to add additional information elements for the required context objects.

Nodes can double as context consumer and provider. They may for instance take over pre-processing steps for a joint decision and report their local decision.

For the valuation of information a valuation library is provided. Valuation functions are running at each node that participates in the process. Valuation results are a specific form of local decisions and can be offered by a context provider as context objects. For the decision-making we are currently using policies expressed as list of rules. I next step we also consider to integrate DENng[19], which provides an information model enabling the design and the desired operation of an autonomic network and/or autonomic network elements. The model is positioned by the ACF and by the models chief architect, the ACF chair Prof. John Strassner as the major standardisation target of the ACF.

The collected valuations from other nodes can be weighted for instance based on by the capabilities, importance or history of the node (e.g. the "opinion" of a AAA server may counts more than that of a new ad hoc node in the network). Then the decision-making process can use the valuations for instance in a voting process as shown below. It then generates a joint decision based on the local decisions of the majority. As shown above, other collaboration strategies are possible. The decision-making process then triggers the execution of the decision typically by invoking low level network functions (e.g. blocking or prioritization of traffic, re-routing, etc.).

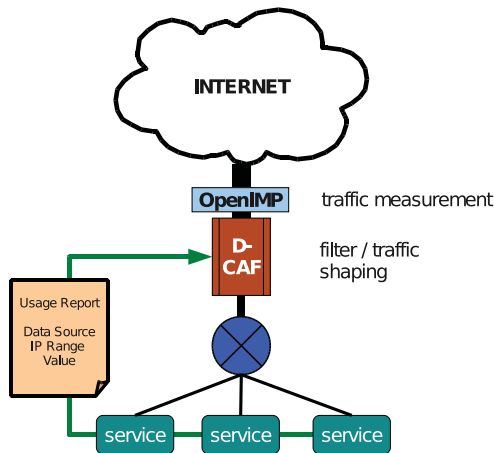
## 5.2 D-CAF: collaboration for Self-protection

The theoretical aspects of collaboration described above are implemented in the FOKUS distributed context-aware firewall (D-CAF). It specifically makes use of the valuation library of the FOKUS Node Collaboration System. In the following we present a common Internet scenario and how it can be addressed by collaboration.

Protecting ones services and ones bandwidth against misuse is a difficult task. Today's intrusion detection and prevention mechanisms are often insufficient, or restrictive for the legitimate users. This is due to two causes: First, it is virtually impossible to discern a malicious Denial of Service (DoS) attack from a sudden burst of legitimate users, a so-called flash crowd. Secondly, Intrusion Detection systems simply do not have the resources to analyze traffic with the same detail as the applications the traffic is addressed to. Thus, smart attacks may always slip past the defenses. To address these problems, we present D-CAF: a distributed context-aware firewall, which selectively limits the access to resources of the protected system by evaluating usage reports from protected services, and reacts accordingly.

The task of intrusion detection requires the analysis and classification of user behaviour on all protocol levels. Common state of the art Intrusion Detection Systems fail at this task, for the very same measurement challenges of complexity, amount of information, encryption and privacy. The alternative to monitoring the network therefore is, to profit from the collaboration of network services.

A web server is designed to parse HTTP data, to analyze user behaviour across network sites, and to detect positive and negative events, such as login attempts, system and database load, and processing of orders. It is therefore the right component to generate reports about user behaviour. In the D-CAF system, a lightweight library is available in several programming and web scripting languages. It allows to send a simple report of user identifier (IP address) and rating. It can be easily integrated in any existing application or website.




**Fig. 3.** D-CAF Information Flow.

In Fig. 3 we show the flow of information in the D-CAF firewall. The network is abstracted as such: A number of services is connected to the Internet via one link. On this link we place the FOKUS measurement platform OpenIMP [30] and the D-CAF system. The firewall receives information on the amount of observed traffic (total and per user) from the measurement system. In a first phase of normal operation, users connect from the Internet with our services as they normally would. This will generate positive and negative valuations of the users by the services, which we map to a numeric range of  $(-1;1)$ . These ratings are transmitted to the D-CAF firewall.


During this phase, the information from all services is only collected, and weighted according to the importance and specialization of the services. The summary of ratings will provide the firewall with a differentiated valuation of all users which have used the services in the past. An example of such a summary is shown in Fig. 4. The chart displays the aggregated subjective value of each IP address for the services. Both single IP addresses or whole address ranges may be valued in the range  $(-1;1)$ .

The next phase happens, when a DDoS attack is launched against one of the services protected by the firewall. This event is easily detectable by the surge of traffic reaching the service. We therefore define a simple traffic threshold which indicates whether the protected services is operating normally or whether it is about to be overloaded. The



**Fig. 4.** D-CAF Usage Report.

firewall will take action, once the observed traffic reaches this threshold. If it is exceeded, it will begin to filter those users with the least favourable ranking. This filtering process continues until the remaining traffic is contained by the threshold. In Fig. 5, we exemplify the process: Given the previous ratings from Fig. 4 (bars) and the detected traffic per IP address (line), the algorithm can filter the worst rated IP addresses (blocks) and calculate an estimate of the traffic reduction thus attained.



**Fig. 5.** Unfiltered traffic example.

The system thus reacts to attacks by prioritizing the users which have shown to behave correctly in the past - this would typically include paying customers and users with active accounts. Unknown users and systems which have misbehaved in the past are more prone to be filtered. Note that even though legal users may be filtered in the process, action is only taken when the system is overloaded. No action would imply the breakdown of the service for all users. After a pre-defined time the filter rules will be removed to be able to react to changing traffic patterns.

Finally, the firewall is distributed, as its very simple valuation semantics allow it to exchange information about the IP addresses with other similar firewalls. A complete snapshot of all the valuations in one firewall can be sent to other D-CAF instances in the same or remote networks. This is then again considered to be a subjective report from a specialized application.

## 6 Conclusion

We surveyed approaches to handle the challenges of the future Internet and point out the demand for self-management and self-protection. We show that current and expected future complexity of networking leads to loss of measurability, due to the amount of information, its distribution and its protection by the owners. This resulting challenge can best be handled by facilitating collaboration strategies in the complex network. Collaboration leads to sharing of resources, sharing of information, and owner consent on protected data sharing. We identify the challenges in collaboration and decision making in a widely distributed group, and presented several collaboration strategies for various requirements. We present our Node Collaboration System (NCS) designed to investigate collaboration strategies for self-management and self-protection and show in an example implementation how the system can be used to achieve network self-protection by node collaboration. As part of future work we will evaluate different collaboration strategies by utilizing the FOKUS Node Collaboration System. Based on the requirements of future network scenarios the best performing schemes will be used to develop a platform for Information Sharing and Decision Making which serves as an enabler for Situation Aware Networking.

## References

1. EU Project OneLab. <http://www.one-lab.org>
2. A. Gavras and H. Bruggemann and D. Witaszek. Pan European Laboratory for next generation networks and services, March 2006. Testbeds and Research Infrastructures for the Development of Networks and Communities, 2006. TRIDENTCOM 2006.
3. <http://akari-project.nict.go.jp/eng/conceptdesign.htm>
4. <http://fif.kr/>
5. IBM. An architectural blueprint for autonomic computing. white paper, IBM, 2006.
6. <http://www.autonomic-communication.org/projects/acca/index.html>
7. <http://www.autonomic-communication-forum.org/>
8. C. Fraleigh, C. Diot, B. Lyles, S. Moon, P. Owezarski, D. Papagiannaki, and F. Tobagi. Design and deployment of a passive monitoring infrastructure. *Lecture Notes in Computer Science*, 2170:556+, 2001.
9. L. Deri. nprobe: an open source netflow probe for gigabit networks. In *In Proc. of Terena TNC2003*, 2003.
10. G. Iannaccone, C. Diot, I. Graham, and N. McKeown. Monitoring very high speed links. In *ACM Internet Measurement Workshop*, 2001.
11. A. Kumar, J. Xu, J. Wang, O. Spatscheck, and L. Li. Space-code bloom filter for efficient per-flow traffic measurement. In *Infocom*, 2004.
12. N. Duffield. Sampling for passive internet measurement: A review. In *Statistical Science*, Volume 19, pp. 472–498, 2004.
13. T. Zseby, M. Molina, N. Duffield, S. Niccolini, and F. Raspall. Sampling and Filtering Techniques for IP Packet Selection RFC 5475, February 2009.
14. B.-Y. Choi, J. Park, and Z.-L. Zhang. Adaptive Random Sampling for Load Change Detection. *SIGMETRICS Perform. Eval. Rev.*, 30(1): pp. 272–273, 2002.
15. C. Estan, K. Keys, D. Moore, and G. Varghese. Building a Better NetFlow. In *SIGCOMM*, 2004.
16. EU Project PRISM. <http://www.fp7-prism.eu/>

17. EU Project 4WARD. <http://www.4ward-project.eu/>
18. J. Strassner, S. Samudrala, G. Cox, Y. Liu, M. Jiang, J. Zhang, S. v. Meer, M. Foghlu, and W. Donnelly. The Design of a New Context-Aware Policy Model for Autonomic Networking. In *Proceedings of the 2008 international Conference on Autonomic Computing*, 2008.
19. J. Strassner. Introduction to DENng for PanLab II. ACF, 2008-2009, tutorial given 21.01.2009 in Fraunhofer FOKUS.
20. I. D. Graham, S. F. Donnelly, S. Martin, J. Martens, and J. G. Cleary. Nonintrusive and accurate measurement of unidirectional delay and delay variation on the internet. In *INET*, 1998.
21. N. G. Duffield and M. Grossglauser. Trajectory sampling for direct traffic observation. In *SIGCOMM*, pp. 271–282, 2000.
22. T. Zseby, S. Zander, and G. Carle. Evaluation of building blocks for passive one-way-delay measurements. In *Proceedings of Passive and Active Measurement Workshop (PAM 2001)*, April 2001.
23. T. Gamer, M. Scharf, M. Schöller. Collaborative Anomaly-based Attack Detection. Proceedings of 2nd International Workshop on Self-Organizing Systems (IWSOS 2007), p. 280-287, Springer, English Lake District, Sep 2007.
24. T. Zseby, E. Boschi, N. Brownlee, and B. Claise. IP Flow Information Export (IPFIX) Applicability. RFC 5472, Feb. 2009.
25. <http://www.cisco.com/en/US/products/ps6235/>
26. B. Claise. Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information. RFC 5101 (Proposed IETF Standard), Jan 2008, Internet Engineering Task Force, <http://www.ietf.org/rfc/rfc5101.txt>.
27. Yi An Huang and Wenke Lee. A cooperative intrusion detection system for ad hoc networks. In *SASN '03: Proceedings of the 1st ACM workshop on Security of ad hoc and sensor networks*, pages 135–147, New York, NY, USA, 2003. ACM.
28. D. Witaszek, and J. Tiemann. Context Dissemination System: Requirements, Architecture and Ability to Support Measurement Results. *Technical Report TR-2008-0130*, Fraunhofer FOKUS.
29. J. Tiemann, and D. Witaszek. Context Coordination and Dissemination System - Architecture and Basic Implementation. *Technical Report TR-2008-0303*, Fraunhofer FOKUS.
30. M. Lutz. <http://www.ip-measurement.org/openimp/index.html>
31. Fang Liu, Xiuzhen Cheng, and Dechang Chen. Insider attacker detection in wireless sensor networks. *INFOCOM 2007. 26th IEEE International Conference on Computer Communications. IEEE*, pages 1937–1945, 6-12 May 2007.
32. Soumya Banerjee, Crina Grosan, Ajith Abraham and P.K. Mahanti. Intrusion Detection on Sensor Networks Using Emotional Ants. IN *International Journal of Applied Science and Computations*, USA, Vol.12, No.3, pp.152-173, 2005.

# Creating a viable Evolution Path towards Self-Managing Future Internet via a Standardizable Reference Model for Autonomic Network Engineering

Ranganai Chaparadza<sup>1</sup>, Symeon Papavassiliou<sup>2</sup>, Timotheos Kastrinogiannis<sup>2</sup>,  
Martin Vigoureux<sup>3</sup>, Emmanuel Dotaro<sup>3</sup>, Alan Davy<sup>4</sup>, Kevin Quinn<sup>4</sup>, Michał  
Wódczak<sup>5</sup>, Andras Toth<sup>6</sup>, Athanassios Liakopoulos<sup>7</sup>, Mick Wilson<sup>8</sup>

<sup>1</sup>Fraunhofer FOKUS Institute for Open Communication Systems, Berlin, Germany.

<sup>2</sup>Institute of Communications and Computer Systems (ICCS), Athens, Greece.

<sup>3</sup>Alcatel-Lucent, Nozay, France

<sup>4</sup>Telecommunications Software & Systems Group, Waterford Institute of Technology,  
Ireland.

<sup>5</sup>Telcordia Poland Sp. z o.o., Poznań, Poland.

<sup>6</sup>Ericsson, Stockholm, Sweden.

<sup>7</sup>Greek Research and Technology Network (GRNET), Athens, Greece,

<sup>8</sup>Fujitsu Laboratories of Europe, United Kingdom

**Abstract.** Clearly, whether revolutionary/clean-slate approaches or evolutionary approaches should be followed when designing Future Multi-Service Self-Managing Networks, some holistic Reference Models on how to design autonomic/self-managing features within node and network architectures are required. Why Reference models?: (1) to guide both approaches towards further architectural refinements and implementations, and (2) to establish common understanding and allow for standardizable specifications of architectural functional entities and interfaces. Now is the time for *harmonization* and *consolidation* of some ideas emerging (or achieved so far) from both approaches to Future Internet design, through the development of a common, unified and “standardizable” Reference Model for autonomic networking. This paper presents this vision. We also present the design principles of an emerging Generic Autonomic Network Architecture (GANA)—a holistic Reference Model for autonomic networking calling for contributions. We describe different “instantiations” of GANA that demonstrate its use for the management of a wide range of both basic and advanced functions and services, in various networking environments.

**Keywords:** *pre-Standardization through an Industry Specification Group (ISG), a call for Specifications, Self-managing Networks, Future Internet, Autonomic Network Architectures.*

---

<sup>1</sup>[Ranganai.Chaparadza@fokus.fraunhofer.de](mailto:Ranganai.Chaparadza@fokus.fraunhofer.de), <sup>2</sup>[papavass@mail.ntua.gr](mailto:papavass@mail.ntua.gr), [timothe@netmode.ntua.gr](mailto:timothe@netmode.ntua.gr),  
<sup>3</sup>[martin.vigoureux@alcatel-lucent.fr](mailto:martin.vigoureux@alcatel-lucent.fr), [emmanuel.dotaro@alcatel.fr](mailto:emmanuel.dotaro@alcatel.fr), <sup>4</sup>[adavy@tssg.org](mailto:adavy@tssg.org), [kquinn@tssg.org](mailto:kquinn@tssg.org),  
<sup>5</sup>[mwodczak@telcordia.com](mailto:mwodczak@telcordia.com), <sup>6</sup>[andras.toth@ericsson.com](mailto:andras.toth@ericsson.com), <sup>7</sup>[aliako@grnet.gr](mailto:aliako@grnet.gr), <sup>8</sup>[Mick.Wilson@uk.fujitsu.com](mailto:Mick.Wilson@uk.fujitsu.com)



## 1 Introduction

The two basic ways to address the management challenges of the Future Internet could be either evolutionary (incremental) or revolutionary (clean slate). It is a requirement rather than a desire, to develop and test in large-scale environments, an enhanced, flexible and robust intrinsic management approach. The vision presented here is motivated by the EC-funded EFIPSANS-FP7 project [1] which is one of large-scale research projects that is seeking to create a viable roadmap for the evolution of today's networking models, paradigms and protocols (in particular IPv6 protocols) towards the self-managing Future Internet. The rest of this article is organized as follows. We briefly present the rationale behind the call for contributions to the development of a **Standardizable Reference Model** for autonomic network engineering that should be used as a guide for creating an **Evolution Path** towards the Self-Managing Future Internet. We then present a holistic evolvable Reference Model for Autonomic Network Engineering emerging from the EC-funded EFIPSANS-FP7 project [1] called the Generic Autonomic Network Architecture (GANA), emphasizing on the self-management aspects within node/device and network architectures in Future Internet and calling for further developments and contributions from diverse ideas from both revolutionary/clean-slate and evolutionary approaches to Future Internet design. Then, different instantiations of the GANA approach are presented, which demonstrate its use for the management of a wide range of functions and services, including both basic network services such as routing and monitoring, as well as enhanced ones such as mobility and Quality of Service (QoS) management. Finally, we give conclusions and an insight into further work in Section 9.

## 2 The Vision of a Self-Managing Future Internet

The vision of the Future Internet, is of a self-managing network whose nodes/devices are designed/engineered in such a way that all the so-called traditional network management functions, defined by the FCAPS management framework (Fault, Configuration, Accounting, Performance and Security) [2], as well as the fundamental network functions such as routing, forwarding, monitoring, discovery, fault-detection and fault-removal, are made to automatically feed each other with information (knowledge) such as goals and events, in order to effect feedback processes among the diverse functions. These feedback processes enable reactions of various functions in the network and/or individual nodes/devices, in order to achieve and maintain well defined network goals. In such an evolving environment, it is required the network itself to help detect, diagnose and repair failures, as well as to constantly adapt its configuration and optimize its performance. Looking at **Autonomicity and Self-Manageability**, we see that autonomicity (i.e. control-loops and feed-back mechanisms and processes, as well as the information/knowledge flow used to drive control-loops) [3], becomes an enabler for self-manageability of networks. As such, even the FCAPS functions become diffused within node/device architectures, apart from being part of an overall network architecture—whereby traditionally, a distinct management plane is engineered separately from the other functional planes of the

network. Since even the management functions become inherent functions of the fundamental node/device architectures, it means that the functional planes of a self-managing network, would need to be (re)-defined and re-factored (refer to [4]). New concepts, functional entities and their associated architectural design principles that facilitate Self-Management at different levels of node/device and network functionality and abstractions, are required.

### 3. An initiative of an Industry Specification Group (ISG):“Autonomic Network Engineering for the Self-Managing Future Internet” (AFI) has been established in ETSI

ETSI has recently launched the initiatives of Industry Specification Groups (ISGs). An Industry Specification Group (ISG): **Autonomic Network Engineering for the Self-Managing Future Internet (AFI)** has just been established in ETSI by the EU-funded FP7-EFIPSANS project [1]. For more information on the AFI\_ISG including the “Rationale” and “*Terms of Reference*” of the AFI\_ISG, envisaged liaisons with the likes of IETF, 3GPP, NGMN, TMF, Autonomic Communication Forum (ACF)[12], etc, we refer to <http://portal.etsi.org/afi> [14]. Through the AFI, we are calling for *Contributions to the Definition and Specifications of a Unified Common Reference Model for Autonomic Network Engineering for the Self-Managing Future Internet* i.e. the further development of *detailed Specifications* of all the issues we have identified as requiring *detailed specifications in the GANA Reference Model*. An **Evolution Path** can be created that starts with the current IPv6 and produces Extensions to IPv6 towards IPv6++(refer to [13]) and other types of network architectural extensions such as cross-layering as necessitated by the GANA Reference Model for engineering Autonomic/Self-Managing Networks. In EFIPSANS, some ideas on Extensions to IPv6 are now emerging as early draft IPv6 Extension Headers (new IPv6 protocols that complement existing IPv6 protocols), protocol Options in the Extension Headers that support the notion of Options, extensions to the “*management interfaces*” of some protocols that ensure enriched autonomic control of the protocols by associated autonomic Decision-Making-Elements (DMEs), and network architectural extensions such as cross-layering, etc. Examples of IPv6 protocol extensions being proposed by EFIPSANS include ICMPv6++ for advanced control information exchange, ND++ for advanced Auto-Discovery, DHCPv6++ for advanced Auto-Discovery, some recommendations for Extensions to protocols like OSPFv3, and some newly proposed Extension Headers, etc.

### 4. The Emerging GANA, as an Evolvable holistic architectural Reference Model for Self-Management within node/device and network architectures

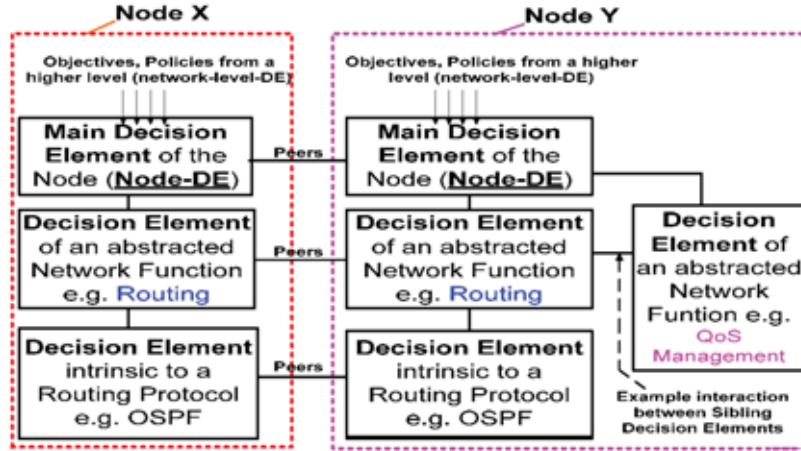
The adopted Generic Autonomic Network Architecture (GANA) [4], sets the fundamental principles and guidelines that need to be followed towards realizing our

vision of the Self-Managing Future Internet, and does not intend to provide any specific solution or implementation. In contrast to any other of today's best known approaches, including clean-slate approaches (both pure and non-pure) such as 4D [5], ANA [6], CONMan [7], a Knowledge Plane for the Internet [8], FOCAL [9,3], a Situatedness-based Knowledge Plane for autonomic networking [10], the approach adopted here introduces Autonomic Manager Components for different abstraction levels of functionality, which are designed following the principles of *Hierarchical*, *Peering*, and *Sibling* relationships among each other within a node/device or network. Moreover, these components are capable of performing autonomic control of their associated Managed-Entities, as well as co-operating with each other in driving the self-managing features of the Network(s). Among GANA objectives is to address the following problems and issues: **(1)** Complexity—by defining some abstractions for autonomic/self-management functionality at *four hierarchical levels* as described later; **(2)** How to ensure that the decision-making-processes for autonomicity (self-management behaviours) within a node/device and the network as a whole, are conflict-free; **(3)** Capturing the kind of perspectives offered to end-users or operators of autonomic/self-managing networks, such as the interfaces that are meant to allow humans to define network-level objectives that govern the operation of an autonomic (self-managing) network under the control of an administrative domain. In GANA, *four levels of abstractions* for which DMEs, MEs and Control-Loops can be designed, are described below (following a bottom up approach).

**Level-1:** Self-manageability issues may be associated with some implementation of a single network protocol (whether monolithic or modular). This level is the lowest level of abstraction of functionality in GANA and is associated with the manifestation of control-loops, as depicted in Fig. 1.

**Level-2:** The concepts of a Control Loop, Decision-Making Element, Managed-Entity(ies), as well as the related self-manageability issues may be associated with a higher level of abstraction than a single protocol (Fig. 1). This means that the aspects of Autonomicity/Self-management may be addressed at the level of “abstracted networking functions” (or “network functions”) such as *routing, forwarding, mobility management, QoS management, etc.* At such a level of abstraction, what is managed by an assigned DME are a group of protocols and mechanisms that are collectively wrapped by what we may call a Function Block or Functional Block, and are considered to belong to the functionality of the abstracted networking functions e.g. all routing protocols and mechanisms of a node become managed by a Decision-Making-Element (Routing\_Management\_DE) assigned and designed to manage only those protocols and mechanisms. This level of abstraction allows us to talk about autonomicity of self-managing properties at this particular level of abstracted network function e.g. autonomic routing, autonomic forwarding, autonomic QoS management, autonomic mobility management, in the node/network. We call the DEs operating at this level, the “Functions-Level” DEs.

**Level-3:** On a higher level of autonomic networking functionality than the level of “abstracted networking functions” of a node/network, the concepts of a Control-Loop, Decision-Making Element, Managed-Entity(ies), as well as the related self-manageability issues may be associated with a system (node) as a whole.



**Fig. 1.** Examples of Hierarchical, Peering, Sibling Relationships and Interfaces of DEs in GANA, calling for Specifications

Fig. 1 illustrate that at this level of self-management (autonomic) properties, the lower level Decision-Making-Elements operating at the level of abstracted networking functions become the Managed Automated Tasks (Managed-Entities) of the main Decision-Making-Element (DME) of the system (node). This means that the node's main DME has access to the "views" exposed by the lower level DMEs and uses its overall knowledge to influence (enforce) the lower level DMEs to take certain desired decisions, which may in turn further influence or enforce desired behaviours on their associated Managed-Entities, down to the lowest level of individual protocol behaviour. A "Sibling" relationship simply means that the entities are created or managed by the same upper level Decision-Making-Element (DME/DE). This means that the entities having a sibling relation can still form other types of peer relationship within the autonomic node or with other entities hosted by other nodes in the network, according to the protocol defined for their needs to communicate with other DEs.

**Level-4:** The next level of self-manageability (autonomicity) after the "node level" described above, is the "network level". There may exist a logically centralized Decision-Making-Element or isolated decision plane/cloud such as the one proposed in the 4D network architecture [6] that knows (through some means) the objectives, goals or policies to be enforced by the whole network. The objectives, goals or policies may actually require that the main (top-level) DMEs of the nodes of the network covered by the centralized DME or plane export "views" such as events and state information to the centralized DME or plane. This may happen in order for the centralized DME to influence or enforce the DMEs of the nodes to take certain desired decisions following specific network policies that may in turn have an effect of inductive decision changes on the lower level DMEs of individual nodes i.e. down to protocol level decisions. A distributed network-level Control-Loop may be implemented following the above set-up, while another case of implementing a distributed Control-Loop would involve the main Decision-Making Elements of nodes working co-operatively to self-organize and manage the network without the presence of a logically centralized DME or an isolated decision plane that manages the whole network (i.e. the possibility for performing "in-network" management).

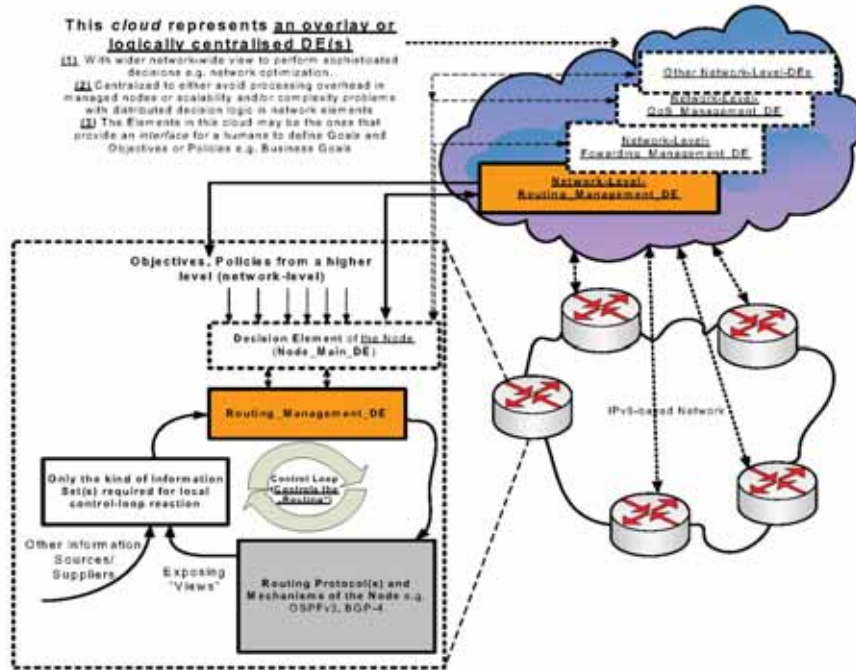


Fig. 2. Autonomicity as a feature in Routing Functionality in a IPv6 based network

## 5. The Instantiation of GANA for Routing and Autonomicity in Fixed Networks

The Routing Functionality (Function) of nodes in a fixed IPv6 based network and the network as whole can be made autonomic by making diverse Routing Schemes and Routing Protocol Parameters employed and altered based on network-objectives, changing network context and the dynamic network views in terms of events, topology changes, etc. Fig. 2 depicts how the routing behaviour of a device and the network as a whole can be made autonomic. Two types of Control-Loops are required for managing/controlling the routing behaviour. The first type is a node-local control loop that consists of a Routing\_Management\_DE embedded inside an autonomic node e.g. a router. The local Routing\_Management\_DE is meant to process only that kind of information that is required to enable the node to react autonomically (according to some goals) by adjusting or changing the behaviour of the individual Routing protocols and mechanisms required to be running on the node. The Routing\_Management\_DE reacts to “views”, such as “events or incidents” exposed by its Managed Entities (MEs)—the Routing protocols or mechanisms. Therefore, the Routing\_Management\_DE implements the *self-configuration and dynamic reconfiguration* feature specific to the routing functionality of the autonomic node. It is important to note that due to scalability, overhead and complexity problems that arise with attempting to make a Routing\_Management\_DE of a node process huge information/data for the control loop, a logically centralised Decision Element(s), may be required, in order to relieve the burden. In such a case, a network-wide



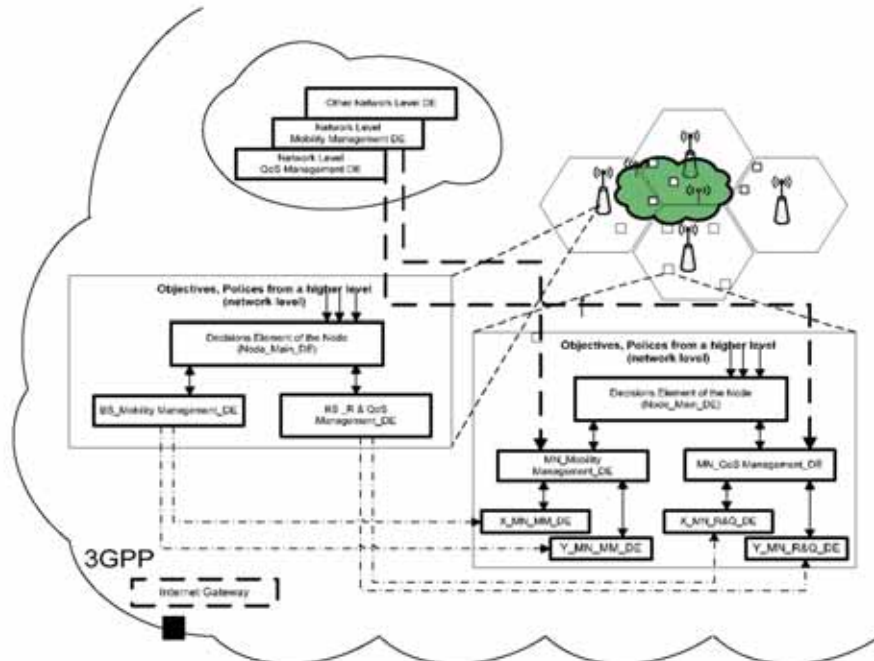
Control Loop is required in addition to the node-local control (with both types of loops focussed on controlling/managing the routing behaviour in an autonomic way). Therefore, both types of control loops need to work together in parallel via the interaction of their associated Routing\_Management\_DEs (one in the node and another in the realm of the logically centralised network overlay decision making elements). The node-scoped (node-local) Routing\_Management\_DE focuses on addressing those limited routing control/management issues for which the node needs to react fast and autonomously. At the same time however, it listens for control from the network-level Routing\_Management\_DE that has wider network-views and dedicated computational power, and is able to compute routing specific policies and new parameter values to be used by individual routing protocols of the node, based on the wider network-views it knows, and disseminate the computed values/parameters to multiple node-scoped Routing\_Management\_DEs of the network-domain. The interaction between the two types of Routing\_Management\_DEs is achieved through the Node\_Main\_DE of a node which verifies those interactions against the overall security policies of the node. The node-scoped Routing\_Management\_DE also relays the “views” such as “events or incidents” to the network-level Routing\_Management\_DE for further reasoning.

## **6. The Instantiation of GANA for Mobility and QoS Management and Autonomicity in Heterogeneous Wireless Networks**

Based on GANA’s design principles, we describe an overall autonomic driven mobility and QoS management architecture, which adopts current mechanisms, methodologies and protocols, enhanced with autonomic behaviours. Figure 3 illustrates the fundamental DEs and their corresponding interactions, that allows us to enable autonomic mobility management and QoS-driven resource allocation functionalities of devices (i.e. mobile node (MN) and base station (BS)) and thus networks, over a heterogeneous wireless environment (e.g. when two networks coexist namely X (CDMA cellular) and Y (WLAN)).

A mobile node’s Resource Allocation and QoS provisioning DE regarding network X (i.e. X\_MN\_R&Q\_DE) realises a self-adaptation mechanism – with respect to QoS-aware self-optimization – in terms of a node’s local control loop that: a) constantly monitors a user’s services performance as well as the corresponding environmental changes, b) analyzes their current status with respect to QoS requirements and, c), reacts to QoS triggering events towards optimizing its services performance. In accordance to the DEs’ hierarchy in GANA, on the one hand a node’s X\_MN\_R&Q\_DE controls node’s resource allocation and QoS provisioning protocol concerning network X, while on the other hand it is controlled (i.e. is a managed entity) by node’s QoS\_Management\_DE.

A mobile node’s QoS\_Management\_DE is responsible for controlling the corresponding protocols’ X\_MN\_R&Q\_DEs, which exist for each one of the available networks in the node’s locality, when it has multimode capabilities, by enabling advance autonomic functionalities regarding overall node’s current services,



**Fig. 3.** Autonomicity as a feature in Mobility Management & QoS Management in Heterogeneous Wireless Networks.

such as: a) optimal available networks' – requested services' assignment, b) node's QoS-related available resource allocation prioritization, and finally c) steering overall node's QoS-aware behaviour by complying to overall network policies imposed by Network Level QoS Management DEs. The interaction between the two types of QoS DEs (i.e. node's QoS\_Management\_DE and Network Level QoS Management DE) is achieved through Node\_Main\_DE which verifies those interactions against the overall security policies of the node. In an X-type network cell base station (e.g. eNodeB), BS\_R&Q\_DE enables autonomic call admission control (CAC) and QoS-aware resource allocation mechanism, by realizing optimal self-adapting radio-resources (e.g. power and rate) allocation procedures that simultaneously satisfy various and often diverse users' services QoS prerequisites, residing at X cell's base stations. Towards achieving the above goal, each base station's BS\_R&Q\_DE interacts with the corresponding currently attached nodes' X\_MN\_R&Q\_DEs (i.e. peering Des). Finally, neighbouring base stations' BS\_R&Q\_DEs of various co-located wireless networks collaboration allows the realization of proficient joint resource allocation and load balancing mechanisms.

Towards enabling autonomic nodes seamless mobility capabilities over a heterogeneous wireless environment, the following autonomic functionalities (i.e. DEs) are introduced in each of the networks components. A mobile node's Mobility Management DE for network X (i.e. X\_MN\_MM\_DE) controls and enhances with self-adaptation attributes node's horizontal handoff, vertical handoff, and mobile IPv6 functionalities. Moreover, since when a mobile node is roaming over a heterogeneous wireless environment can be simultaneously attached to more than one access wireless



networks at its locality, its corresponding X\_MN\_MM\_DEs for each of the available X networks' are controlled and managed by an upper in the hierarchy of GANA-based DE, namely Mobility Management DE (i.e. MN\_Mobility Management\_DE belonging at Functions-Level). A node's MN\_Mobility\_Management\_DE introduces autonomicity in mobile node's or corresponding services' advanced mobility functionalities such as, multihoming (i.e. in terms of assigning different services to different access networks), multi-connection (i.e. in terms of splitting the data of one application across multiple connections) and dynamic alteration of the networks that a node is currently attached to at the event of QoS-triggering affairs. The latter is achieved via interacting with the corresponding node's QoS\_Management\_DE. Moreover, a node's MN\_Mobility\_Management\_DE is steering the overall node's mobility behaviour by complying nodes actions with the overall network policies imposed by the Network Level Mobility Management DEs through its Node\_Main\_DE.

## **7. The Instantiation of GANA for Traffic Monitoring and Autonomicity in Fixed Networks**

Autonomicity as a feature of Traffic Monitoring, coupled with Quality of Service (QoS) management functions of an ingress edge router are at focus within this instantiation of GANA. The objective of QoS control at the ingress within a DiffServ domain is to ensure traffic admitted to the network is appropriately classified, policed and shaped to ensure QoS targets imposed on traffic will be maintained as traffic passes through the network.

The configuration of network monitoring protocols and mechanisms can be managed through a dedicated Traffic-Monitoring-DE designed to operate inside a node. The monitoring information collected by the monitoring protocols is required for driving the behaviours of diverse DEs and some pure MEs of a node, and should be of the minimum level of accuracy required by the requesting network functions and applications i.e. pure MEs and/or DEs of a node. Therefore, as traffic and network conditions change, monitoring protocols and mechanisms require to be constantly re-configured by the Traffic-Monitoring-DE to ensure certain requirements and goals are satisfied, as necessitated by the requirements from DEs and MEs of the node(s).

This instantiation focuses on developing autonomic features for the QoS\_DE and its associated Admission Control AC\_ME that is considered to belong to the lowest level of MEs in GANA, and for the Traffic-Monitoring-DE and its associated specialized Traffic Measurement ME (TM\_ME) called the Effective Bandwidth measurement entity of the ingress edge router, also considered to belong to the lowest level of MEs in GANA and being autonomically controlled by the Traffic-Monitoring-DE. There is a dependency relationship between the two of the lowest level MEs according to GANA (Fig. 4), and as conditions change within the network, each of the two lowest level MEs can be re-configured by their associated DEs to operate in an optimal manner in the face of these changes. The autonomic behaviour instilled within the ingress edge node is the ability to control the admission of traffic into the network, while maintaining a high degree of confidence in the admission

decisions under varying traffic conditions. This is provisioned by an interaction between the Traffic-Monitoring\_DE and the QoS\_DE where the traffic monitoring requirements of the AC\_DE change and the associated TM\_ME must be re-configured dynamically by the Traffic-Monitoring\_DE to suit.

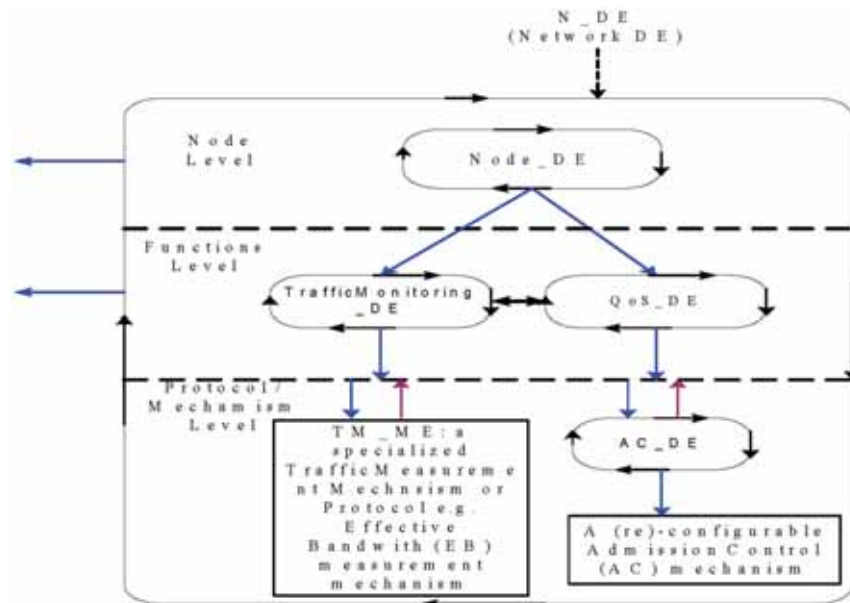


Fig. 4. Autonomicity as a feature of Traffic Monitoring Functionality

## 8. The Instantiation of GANA for Auto-Configurations in MANETS

Auto-configuration is one of the key aspects of IPv6-based Future Internet. This especially holds true for tactical environments where one can envisage multiple groups of mobile nodes forming MANETs on the move. These groups might merge or further split as well as encounter specific faulty conditions (Figure 5). The role of auto-configuration is then not only to enable an efficient address assignment scheme but also provide certain capabilities making it feasible for the network to survive as a whole. The measure of the level of survivability might be defined e.g. as the ability to offer basic services such as routing. For this purpose the auto-configuration entity must interact with some other entities including the ones responsible for fault-management, resilience and survivability as well as routing. This puts sophisticated requirements on the architecture of such an autonomic network and that is where GANA comes into play. In particular it is envisaged that the aforementioned different entities are instantiated by their corresponding Decision Elements (DEs) that interact among themselves and control specific Managed Entities (MEs). In particular, the Fault-Management-DE (FM\_DE) is responsible for fault diagnosis, localization, isolation and removal. This DE analyses information regarding the current situation in

the network and not only tries to resolve the existing problems but, what is more, based on different symptoms makes an attempt to infer what other problems might be coming. This information is really crucial because the Resilience and Survivability DE (RS\_DE) may exploit it for the purposes of MANET reorganization so it becomes ready to survive both the existing situation and its negative consequences. For this purpose there might arise e.g. the need one of the groups of the nodes to split in two parts so each of them joins another neighboring group. As a consequence the network might become more resilient but on the other hand, one needs to keep in mind that this network should be still in a position to offer other basic functions such as routing. This requirement suggests that autonomic decisions need to be taken by different DEs depending on the current situation in the network. The Auto-configuration DE (AC\_DE) is then not only responsible for the optimum deployment but also provides quick and efficient IPv6 address configuration so duplicate addresses are efficiently avoided. It exploits information delivered by RS\_DE on the basis on the FM\_DE and also interacts with the Routing DE (RT\_DE). The purpose of this interaction is that on the one hand the AC\_DE, being aware of the requirements pertaining to routing, may try to make decisions that include taking into account these requirements. On the other hand, even if there are no specific requirements but it is possible to offer reliable routing, the interaction between RT\_DE and AC\_DE might result in a better overall network robustness as regards the network itself (a group of nodes) and the services it offers. All the aforementioned DEs are assumed to operate at the *Functions-Level* of GANA's hierarchy of DEs, but one could still try to define their mutual hierarchical relations when viewed from AC\_DE perspective. Information exchange among them is assumed to be performed with the aid of some specially designed new IPv6 Extension Headers.




Fig. 5. Tactical MANET scenario

## 9. Conclusions and Future Work

In this paper, we presented the idea of a new initiative concerning the establishment of an Industry Specification Group (ISG): “**Autonomic Network Engineering for the Self-Managing Future Internet (AFI)**” within ETSI. [14] presents the rationale behind this new initiative. Further information regarding the developments related to this initiative including invitations for participation to the activities of the AFI\_ISG can be found in [14]. Therefore, this paper also serves to communicate the initiative to the wider community. In this paper, we also presented an emerging holistic reference model for autonomic network engineering (GANA), as a fundamental enabler for

self-management within node and network architectures. GANA should be seen as a common reference model that can benefit both the evolutionary and revolutionary approaches towards Future Internet design, with both approaches contributing to its further development. Such *harmonized contributions* of consolidated ideas and concepts from both revolutionary/clean-slate approaches and evolutionary approaches, to the further development of GANA, e.g. from results from multiple research projects (past and future), can be achieved only through the AFI\_ISG.

## Acknowledgement

This work has been partially supported by EC FP7 EFIPSANS project (INFSO-ICT-215549).

## References

1. EC funded- FP7-EFIPSANS Project: <http://efipsans.org/>
2. The FCAPS management Framework: ITU-T Rec. M. 3400.
3. B. Jennings, S. van der Meer, S. Balasubramaniam, D. Botvich, M.O. Foghlu, W. Donnelly, J. Strassner, "Towards Autonomic Management of Communications Networks," IEEE Communications Magazine, Vol. 45, pp. 112-121, Oct. 2007.
4. R. Chaparadza, "Requirements for a Generic Autonomic Network Architecture (GANA), suitable for Standardizable Autonomic Behaviour Specifications of Decision-Making-Elements (DMEs) for Diverse Networking Environments," to appear in International Engineering Consortium (IEC) in the Annual Review of Com., Volume 61, Dec. 2008.
5. A. Greenberg et al, "A clean slate 4D approach to network control and management," ACM SIGCOMM Computer Comm.Review, vol. 35(5), pp.41-54, 2005.
6. European IST FP6 ANA(Autonomic Network Architecture) Project: <http://www.ana-project.org/>
7. H. Ballani, and P. Francis, "CONMan: A Step Towards Network Manageability," ACM SIGCOMM Computer Comm.Review, vol. 37(4), pp.205-216, 2007.
8. D. D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski, "A Knowledge Plane for the Internet," in Proc. of ACM SIGCOMM 2003, Karlsruhe, Germany, Aug. 25-29, 2003.
9. J. C. Strassner, N. Agoulmine, and E. Lehtihet, "FOCALE – A Novel Autonomic Networking Architecture," In proc of the Latin American Autonomic Computing Symposium (LAACS), Campo Grande, MS, Brazil 2006.
10. T. Bullot et al, "A situatedness-based knowledge plane for autonomic networking," International Journal of Network Management, John Wiley & Sons, 2008.
11. IBM article: Understand the autonomic manager concept: <http://www-128.ibm.com/developerworks/library/ac-amconcept/>.
12. Autonomic Communication Forum (ACF): <http://www.autonomic-communication-forum.org/>
13. Chaparadza R.: Evolution of the current IPv6 towards IPv6++ (IPv6 with Autonomic Flavours). Published by the International Engineering Consortium (IEC) in the Review of Communications, Volume 60, December 2007.
14. AFI\_ISG: Autonomic network engineering for the self-managing Future Internet (AFI): <http://portal.etsi.org/afi>

## A Scalable, Transactional Data Store for Future Internet Services<sup>\*</sup>

Alexander Reinefeld, Florian Schintke, Thorsten Schütt, Seif Haridi

{reinefeld, schintke, schuett}@zib.de

Zuse Institute Berlin and onScale solutions

and

haridi@kth.se

Royal Institute of Technology, Sweden

**Abstract.** Future Internet services require access to large volumes of dynamically changing data records that are spread across different locations. With thousands or millions of distributed nodes storing the data, node crashes or temporary network failures are normal rather than exceptions and it is therefore important to hide failures from the application. We suggest to use peer-to-peer (P2P) protocols to provide self-management among peers. However, today's P2P protocols are mostly limited to write-once/read-many data sharing. To extend them beyond the typical file sharing, the support of consistent replication and fast transactions is an important yet missing feature.

We present *Scalaris*, a scalable, distributed key-value store. *Scalaris* is built on a structured overlay network and uses a distributed transaction protocol. As a proof of concept, we implemented a simple Wikipedia clone with *Scalaris* which outperforms the public Wikipedia with just a few servers.

### 1 Introduction

Web 2.0, that is, the Internet as an information society platform supporting business, recreation and knowledge exchange, initiated a business revolution. Service providers offer Internet services for shopping (Amazon, eBay), online banking, information (Google, Flickr, Wikipedia), social networking (MySpace, Facebook), and recreation (Second Life, online games). In our information society, Web 2.0 services are no longer just nice to have, but customers depend on their continuous availability, regardless of time and space. A typical trend is illustrated by Wikipedia where users are also providers of information. This implies that its underlying data store is updated continuously from multiple sources.

---

<sup>\*</sup> This work was partly funded by the EU projects SELFMAN under grant IST-34084 and the EU project XtremOS under grant IST-33576.

How to cope with such strong demands, especially in case of interactive community services that cannot be simply replicated? All users access the same Wikipedia, meet in the same Second Life environment and want to discuss with others via Twitter. Even the shortest interruption, caused by system downtime or network partitioning may cause huge losses in reputation and revenue. Web 2.0 services are not just an added value, but they must be dependable. Apart from 24/7 availability, providers face another challenge: they must, for a good user experience, be able to respond within milliseconds to incoming requests, regardless whether thousands or millions of concurrent requests are currently being served. Indeed, scalability is a key challenge. In addition to scalability and availability any global service to be affordable, somehow requires the system to be self managing (see sidebar).

**Availability** is the proportion of time a system is in a functioning condition. More formally, availability is the ratio of the expected value of the uptime of a system to the aggregate of the expected values of up and down time. Availability is often specified in a logarithmic unit called “nines” which corresponds roughly to a number of nines following the decimal point. “Six nines”, for example, denote an availability of 0.999999, allowing a maximum downtime of 31 seconds per year.

**Scalability** refers to the capability of a system to increase the total throughput under an increased load when resources are added. A scalable database management system is one that can be upgraded to process more transactions by adding new processors, devices and storage, and which can be upgraded easily and transparently without service interrupt.

**Self Management** refers to the ability of a system to adjust to changing operating conditions and requirements without human intervention at runtime. Self Management includes self configuration, self healing and self tuning.

Our Scalaris system, described below, provides a comprehensive solution for self managing and scalable data management. Scalaris is a transactional key-value store that runs over multiple data centers as well as on peer-to-peer nodes. We expect Scalaris and similar systems to become an important core service of future Cloud Computing environments.

As a common key aspect, all Web 2.0 services have to deal with concurrent data updates. Typical examples are checking the availability of products and their prices, purchasing items and putting them into virtual shopping carts, and updating the state in multi-player online games. Clearly, many of these data operations have to be atomic, consistent, isolated and durable (so-called ACID properties). Traditional centralized database systems are ill-suited for this task, sooner or later they become a bottleneck for business workflow. Rather, a scalable, transactional data store like Scalaris is what is needed.

In this paper, we present the overall system architecture of Scalaris. We have implemented the core data service of Wikipedia using Scalaris. Its scalability and self-\* capabilities were demonstrated in the IEEE Scalable Computing Challenge

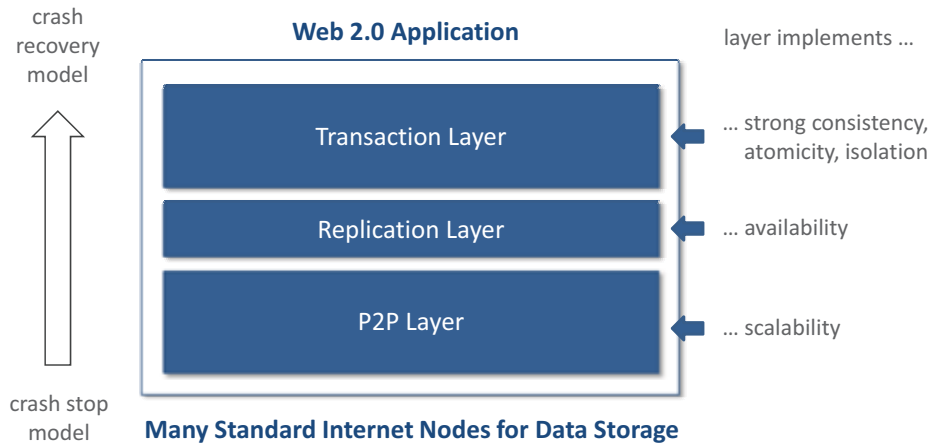


Fig. 1: Scalaris system architecture.

2008, where Scalaris won the 1<sup>st</sup> price ([www.ieeetcsc.org/scale2008](http://www.ieeetcsc.org/scale2008)). Talks on Scalaris were given at the the Google Scalability Conference 2008 [19] and the Erlang eXchange 2008.

The paper is organized as follows. The following Section provides an overview on Scalaris' system architecture, Section 3 describes its self-management features and Section 4 gives further details on the implementation. In Section 5 we demonstrate how Scalaris can be used for implementing Web 2.0 services. As a proof-of-concept, we have chosen a simple Wikipedia clone; performance results are given in Section 6.

## 2 Scalaris

As part of the EU funded SELFMAN project we set out to build a distributed key/value store capable of serving thousands or even millions of concurrent data accesses per second. Providing strong data consistency in the face of node crashes and hefty concurrent data updates was one of our major goals.

With Scalaris, we do not attempt to replace current database management systems with their general, full-fledged SQL interfaces. Instead our target is to support transactional Web 2.0 services like those needed for Internet shopping, banking, or multi-player online games. Our system consists of three layers:

- At the bottom, an enhanced structured overlay network, with logarithmic routing performance, provides the basis for storing and retrieving keys and their corresponding values. In contrast to many other overlays, our implementation stores the keys in lexicographical order. Lexicographic ordering instead of random hashing enables control of data placement which is necessary for low latency access in multi-datacenter environments.



- The middle layer implements data replication. It enhances the availability of data even under harsh conditions such as node crashes and physical network failures.
- The top layer provides transactional support for strong data consistency in the face of concurrent data operations. It uses an optimistic concurrency control strategy and a fast non-blocking commit protocol with low communication overhead. This protocol has been optimally embedded in the overlay network.

As illustrated in Fig. 1, these three layers together provide a scalable and highly available distributed key/value store which serves as a core building block for many Web 2.0 applications as well as other global services. The following sections describe the layers in more detail.

## 2.1 P2P Overlay

At the bottom layer, we use the structured overlay protocol Chord<sup>#</sup> [17,18] for storing and retrieving key-value pairs in nodes (peers) that are arranged in a virtual ring. This ring defines a key space where all data items can be stored according to the associated key. In our case we assume that any key is an arbitrarily long string of characters, therefore the key space is infinite. Nodes are placed at arbitrary places on the ring and are responsible for all data between their predecessor and themselves. The placement policy ensures even distribution of load over the nodes. In each of the  $N$  nodes, Chord<sup>#</sup> maintains a routing table with  $O(\log N)$  entries (fingers). In contrast to traditional Distributed Hash Tables (DHTs) like Chord [21], Kademlia [12] and Pastry [15], Chord<sup>#</sup> stores the keys in lexicographical order, thereby allowing range queries, and control over the placement of data on the ring structure. To ensure logarithmic routing performance, the fingers in the routing table are computed in such a way that successive fingers in the routing table jump over an exponentially increasing number of nodes in the ring. This finger placement will yield uniform in-/out-degree of the overlay network and thus avoids hotspots.

Chord<sup>#</sup> uses the following algorithm for computing the fingers in the routing table (the infix operator  $x . y$  retrieves  $y$  from the routing table of a node  $x$ ):

$$finger_i = \begin{cases} successor & : i = 0 \\ finger_{i-1} . finger_{i-1} & : i \neq 0 \end{cases}$$

Thus, to calculate the  $i^{th}$  finger, a node asks the remote node, listed in its  $(i-1)^{th}$  finger, for the node at which its  $(i-1)^{th}$  finger refers to. In general, at any node, the fingers at level  $i$  are set to the neighbor's finger at the preceding level  $i-1$ . At the lowest level, the fingers point to the direct successor. The resulting structure is similar to a skiplist, but the fingers are computed deterministically without any probabilistic component and each node has its individual exponentially spaced fingers. The fingers are maintained by a periodic stabilization algorithm according to the above formula.

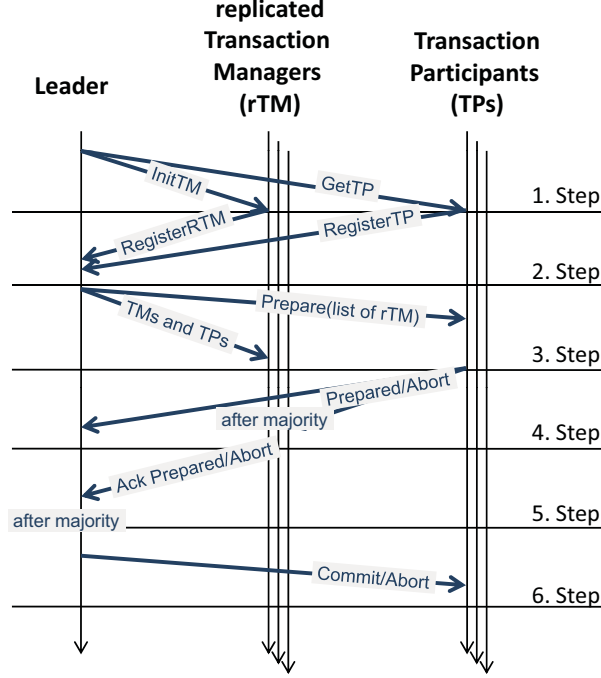


Fig. 2: Adapted Paxos used in Scalaris.

Compared to Chord [21], Chord<sup>#</sup> does the routing in the *node space* rather than in the *key space*. This finger placement has three advantages over that of Chord: First, it naturally works with any type of keys as long as a total order over the keys is defined, and second, finger maintenance is cheaper [17], requiring just one hop instead of a full logarithmic search (as in Chord). To support logarithmic routing performance in skewed key distributions while nodes are arbitrarily placed in the key space—which we have to in our scenario—the third and probably most important difference becomes our trump card: the incoming routing links (fingers) will still be evenly distributed across all nodes. This prevents nodes from becoming hot spots and ensures continuous progress when routing.

## 2.2 Replication and Transaction Layer

The scheme described so far provides scalable access to distributed key/value pairs. To additionally tolerate node failures, we replicate all key/value pairs over  $r$  nodes using symmetric replication [5]. Basically each key is mapped by a globally known function to a set of keys  $\{k_1, \dots, k_r\}$  and the item is replicated according to those keys. Read and write operations are performed on a majority

of the replicas, thereby tolerating the unavailability of up to  $\lfloor (r-1)/2 \rfloor$  nodes. This scheme is shown to provide key consistency for data lookups under realistic networking conditions [20]. For repairing the replication degree of items, nodes have to read the missing data from a majority of replicas. This is necessary to guarantee strong data consistency.

The system supports transactional semantics. A client connected to the system can issue a sequence of operations including reads and writes within a transactional context, i.e. *begin trans ... end trans*. This sequence of operations are executed by a local transaction manager TM associated with the overlay node to which the client is connected. The transaction will appear to be executed atomically if successful, or not executed at all if the transaction aborts.

Transactions in Scalaris are executed optimistically. This implies that each transaction is executed completely locally at the client in a read-phase. If the read phase is successful the TM tries to commit the transaction permanently in a commit phase, and permanently stores the modified data at the responsible overlay nodes. Concurrency control is performed as part of this latter phase. A transaction  $t$  will abort only if: (1) other transactions try to commit changes on some overlapping data items simultaneously; or (2) other successful transactions have already modified data that is accessed in transaction  $t$ .

Each item is assigned a version number. Read/write operation works on a majority of replicas to obtain the highest version number. A Read operation selects the data value with highest version number, and a write operation increments the highest version number of the item.

The commit phase employs an adapted version of Paxos atomic commit protocol [9], which is non-blocking. In contrast to the 3-Phase-Commit protocol used in distributed database systems, the Paxos commit protocol still works in the majority part of a network that became partitioned due to some network failure. It employs a group of replicated transaction managers (rTM) rather than a single transaction manager. Together they form a set of acceptors with the TM acting as the leader.

The commit is basically divided into two phases, the validation phase and the consensus phase. During the validation phase the replicated transaction managers rTM are initialized, and the updated data items together with references to the rTM are sent to the nodes responsible for the data items in a Prepare message. These latter nodes are called transaction participants TPs.

Each TP proposes ‘prepared’ or ‘abort’ in a fast Paxos consensus round with the acceptor set. As each acceptor collects votes from a majority of replicas for each item, it will be able to decide on a commit/abort for the whole transaction. For details see [13,20]. This scheme favors atomicity over availability. It always requires a majority of nodes to be available for the read and commit phase. This policy distinguishes Scalaris from other distributed key-value stores, like e.g. Dynamo [3].

### 3 Self-Management

For many Web 2.0 services, the total cost-of-ownership is dominated by the costs needed for personnel to maintain and optimize the service. Scalaris greatly reduces the operation cost with its built-in self\* properties:

- *Self healing*: Scalaris continuously monitors the hosts it is running on. When it detects a node crash, it immediately repairs the overlay network and the database. Management tasks such as adding or removing hosts require minimal human intervention.
- *Self tuning*: Scalaris monitors the nodes' workload and autonomously moves items to distribute the load evenly over the system in order to improve the response time of the system. When deploying Scalaris over multiple data-centers, these algorithms are used to place frequently accessed items nearby the users.

These protection schemes do not only help in stress situations, but they also monitor and pro-actively repair the system before any service interruption might occur. With traditional database systems these operations require human interference which is error prone and costly. When using Scalaris, fewer system administrators can operate much larger installations compared to legacy databases.

### 4 Implementation

Implementing distributed algorithms correctly is a difficult and tedious task, especially when using imperative programming languages and multi-threading with a shared state concurrency model. The resulting code is often lengthy and error-prone, because large parts of the code deal with shared objects [22] and with exception handling such as node or network failures.

For this reason, message passing as in the *actor model* [7] is becoming the accepted paradigm for describing and reasoning about distributed algorithms [6]. Scalaris was also developed according to this model. The basic primitives in this model are actors and messages. Every actor has a state, can send messages, act upon messages and spawn new actors.

These primitives are easily mapped to Erlang processes and messages [1]. The close relationship between the specification and the programming language allows a smooth transition from the theoretical model to prototypes and eventually to a complete system.

Our Erlang implementation of Scalaris comprises eight major components with a total of 11,000 lines of code: 7,000 for the P2P layer with replication and basic system infrastructure, 2,700 lines for the transaction layer, and 1,300 lines for the Wikipedia infrastructure. Each Scalaris node is organized into the following components:

- The *Failure Detector* supervises other peers and notifies subscribers of remote node failures.

- The *Configuration Store* provides access to the current configuration and allows modifications of various system parameters.
- The *Key Holder* stores the identifier of the node in the overlay.
- The *Statistics Collector* collects statistics and forwards them to central statistic servers.
- The *Chord<sup>#</sup> Node* component is composed of subcomponents for overlay maintenance and overlay routing. It maintains, among other things, the successor list and the routing table. It provides the functionality of the structured overlay layer.
- The *Database* stores the key-value pairs of this node. The current implementation uses an in-memory dictionary, but disk store based on DETS or Mnesia could also be used.
- The *Transaction Manager* runs the transaction protocols.
- The *Replica Repair* maintains the replication degree of items.

The processes are organized in an Erlang OTP supervisor tree. When any of the slaves crashes, it is restarted by the Erlang supervisor. When either of the Chord<sup>#</sup> Node or the Database component fails, the other is explicitly killed and both are restarted to ensure consistency. This is equivalent to a new node joining the system.

## 5 Deployment: Wikipedia on Scalaris

As a challenging benchmark for Scalaris, we implemented the core of Wikipedia, the "free encyclopedia, that anyone can edit". Wikipedia runs on three sites. The main one in Tampa is organized in three layers, the proxy server layer, the web server layer, and the MySQL database layer. The proxy layer serves as a cache for recent requests, and the web server layer runs the application logic and issues requests to the database layer. Wikipedia handles about 50,000 requests per second, from which 48,000 are cache hits in the proxy server layer and 2,000 are processed by the database layer. The proxy and the web server layers are embarrassingly parallel and therefore trivial to scale. From a scalability point of view, only the database layer is challenging.

Our implementation uses Scalaris to replace the database layer. This enables us to run Wikipedia on geographically distributed sites and to scale to almost any number of hosts, as shown in the evaluation section. Our Wikipedia implementation inherits all the favorable properties of Scalaris, such as scalability and self management.

Instead of using a relational database, we map the Wikipedia content to our Scalaris key/value store [14]. We use the following mappings, using prefixes in the keys to avoid name clashes.

	key	value
<b>page content</b>	title	list of Wikitext for all versions
<b>backlinks</b>	title	list of titles
<b>categories</b>	category name	list of titles

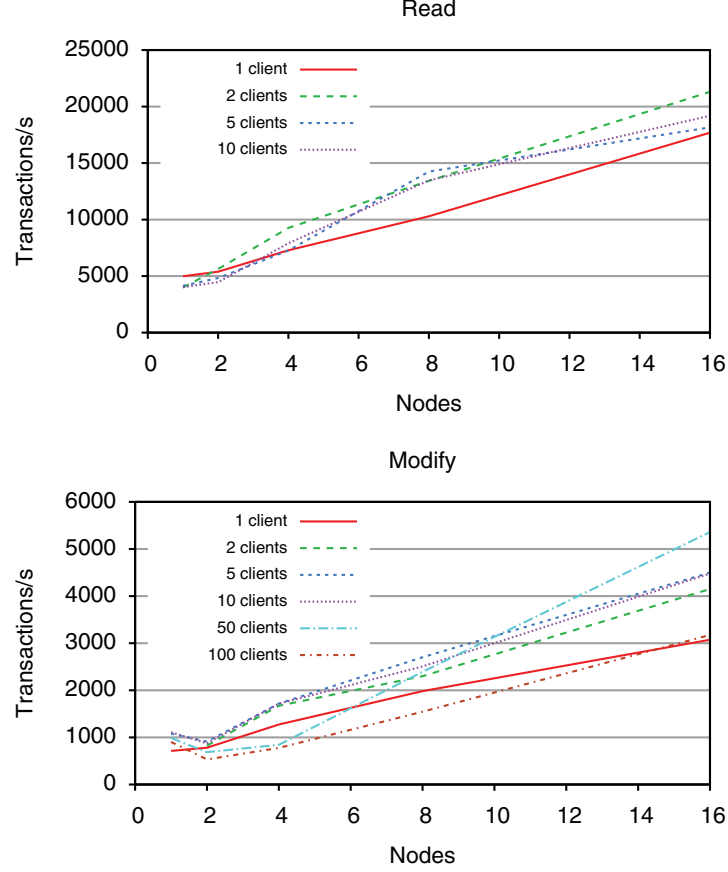


Fig. 3: Performance of Scalaris: (a) Read operation, (b) Modify operation for different numbers of local threads and cluster sizes.

On a page update a transaction across all affected keys (content, backlinks, and categories) and their replicas is triggered.

## 6 Evaluation

We tested the performance of Scalaris on an Intel cluster up to 16 nodes. Each node has two Quad-Core E5420s (8 cores in total) running at 2.5 GHz and 16 GB of main memory. The nodes are connected via GigE and Infiniband; we used the GigE network for our evaluation.

On each physical node we were running one multi-core Erlang virtual machine. Each virtual machine hosted 16 Scalaris nodes. We used a replication degree of four, that is, there exist four copies of each key-value pair.

We tested two operations: a *read* and a *modify* operation. The *read* operation reads a key-value pair. The *modify* operation reads a key-value pair, increments the value and writes the result back to the distributed Scalaris store. To guarantee consistency, the read-increment-write is executed within a transaction. The read operation, in contrast, simply reads from a majority of the keys.

The benchmarks involved the following steps:

- Start watch.
- Start  $n$  Erlang client processes in each VM.
- Execute the read or modify operation  $i$  times in each client.
- Wait for all clients to finish.
- Stop watch.

Figure 3 shows the results for various numbers of clients per VM (see the colored graphs). In the read benchmarks depicted in Fig. 3.a, each thread reads a key 2000 times while the modify benchmarks in Fig. 3.b modify each key 100 times in each thread.

As can be seen, the system scales about linearly over a wide range of system sizes. In the read benchmarks (Fig. 3.a), two clients per VM produce an optimal load for the system, resulting in more than 20,000 read operations per second on a 16 node (=128 core) cluster. Using only one client (red graph) does not produce enough operations to saturate the system, while five clients (blue graph) cause too much contention. Note that each read operation involves accessing a majority (3 out of 4) replicas.

The performance of the modify operation (Fig. 3.b) is of course lower, but still scales nicely with increasing system sizes. Here, the best performance of 5,500 transactions per second is reached with fifty load generators per VM, each of them generating approximately seven transactions per second. This results in 344 transactions per second on each server.

Note that each modify transaction requires Scalaris to execute the adapted Paxos algorithm, which involves finding a majority (i.e. 3 out of 4) of transaction participants and transaction managers, plus the communication between them. The performance graphs illustrate that a single client per VM does not produce enough transaction load, while fifty clients are optimal to hide the communication latency between the transaction rounds. Increasing the concurrency further to 100 clients does not improve the performance, because this causes too much contention. Note that for the 100-clients-case, there are actually  $16 \cdot 100$  clients issuing increment transactions.

Overall, both graphs illustrate the linear scalability of Scalaris.

## 7 Summary

Scalaris provides a scalable and self managing transactional key-value store. We have implemented Wikipedia using Scalaris. Its scalability and self\* capabilities were demonstrated in the IEEE Scalable Computing Challenge 2008, where Scalaris won the 1st prize.



Compared to other data services, Scalaris has significantly lower operating costs and is self-managing. Scalaris and similar systems will be an important building block for Web 2.0 services and future Cloud Computing environments.

While Wikipedia served here as a first demonstrator to show the potential of Scalaris, we envisage a large variety of commercial Web 2.0 applications ranging from e-commerce and social networks to infrastructure services for maintaining server farms. The Scalaris code is open source (scalaris.googlecode.com).

## Acknowledgements

Many thanks to Nico Kruber, Monika Moser, and Stefan Plantikow who implemented parts of Scalaris. Also thanks to Ali Ghodsi, Thallat Shafaat, and Joe Armstrong for their support and many discussions.

## References

1. J. Armstrong. *Programming Erlang: Software for a Concurrent World*. Pragmatic Programmers, ISBN: 978-1-9343560-0-5, July 2007.
2. R. Baldoni, L. Querzoni, A. Virgillito, R. Jiménez-Peris, and M. Patiño-Martínez. *Dynamic Quorums for DHT-based P2P Networks*. NCA, pp. 91–100, 2005.
3. G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: Amazon’s Highly Available Key-Value Store. *Proceedings of the 21st ACM Symposium on Operating Systems Principles*, Oct. 2007.
4. JJ Furman, J. S. Karlsson, J. Leon, A. Lloyd, S. Newman, and P. Zeyliger. Mega-store: A Scalable Data System for User Facing Applications. *SIGMOD 2008*, Jun. 2008.
5. A. Ghodsi, L. O. Alima, and S. Haridi. Symmetric Replication for Structured Peer-to-Peer Systems. *3rd Intl. Workshop on Databases, Information Systems and P2P Computing*, 2005.
6. R. Guerraoui and L. Rodrigues. *Introduction to Reliable Distributed Programming*. Springer-Verlag 2006.
7. C. Hewitt, P. Bishop, and R. Steiger. A Universal Modular ACTOR Formalism for Artificial Intelligence. *IJCAI*, 1973.
8. A. Lakshman, P. Malik, and K. Ranganathan. Cassandra: A Structured Storage System on a P2P Network. *SIGMOD 2008*, Jun. 2008.
9. L. Lamport. The Part-Time Parliament. *ACM Transactions on Computer Systems* 16(2): 133–169, 1998.
10. L. Lamport. Fast Paxos. *Distributed Computing* 19(2):79–103, 2006.
11. M. M. Masud and I. Kiringa. *Maintaining consistency in a failure-prone P2P database network during transaction processing*. Proceedings of the 2008 International Workshop on Data management in peer-to-peer systems, pp. 27–34, 2008.
12. P. Maymounkov and D. Mazieres. Kademlia: A peer-to-peer information system based on the XOR metric. *IPTPS 2002*, Mar. 2002.
13. M. Moser and S. Haridi. Atomic Commitment in Transactional DHTs. *1st Core-GRID Symposium*, Aug. 2007.

14. S. Plantikow, A. Reinefeld, and F. Schintke. Transactions for Distributed Wikis on Structured Overlays. *18th IFIP/IEEE Distributed Systems: Operations and Management (DSOM 2007)*, Oct. 2007.
15. A. Rowstron and P. Druschel. Pastry: Scalable, distributed object location and routing for large-scale peer-to-peer systems. *Middleware 2001*, Nov. 2001.
16. Scalaris code: <http://code.google.com/p/scalaris/>.
17. T. Schütt, F. Schintke, and A. Reinefeld. Structured Overlay without Consistent Hashing: Empirical Results. *GP2PC'06*, May 2006.
18. T. Schütt, F. Schintke, and A. Reinefeld. A Structured Overlay for Multi-Dimensional Range Queries. *Europar*, Aug. 2007.
19. T. Schütt, F. Schintke, and A. Reinefeld. Scalable Wikipedia with Erlang. *Google Scalability Conference*, Jun. 2008.
20. T.M. Shafaat, M. Moser, A. Ghodsi, S. Haridi, T. Schütt, and A. Reinefeld. Key-Based Consistency and Availability in Structured Overlay Networks. Third Intl. ICST Conference on Scalable Information Systems, June 2008.
21. I. Stoica, R. Morris, M.F. Kaashoek D. Karger, and H. Balakrishnan. Chord: A scalable peer-to-peer lookup service for Internet application. *ACM SIGCOMM 2001*, Aug. 2001. Concepts, Techniques, and Models of Computer Programming
22. P. Van Roy and S. Haridi. Concepts, Techniques, and Models of Computer Programming. MIT Press, March 2004.

## Future Internet in Home Area Networks: Towards Optical Solutions?

Roberto Gaudino<sup>1</sup>, Daniel Cardenas<sup>1</sup>,  
Martial Bellec<sup>2</sup>, Benoit Charbonnier<sup>2</sup>, Noella Evanno<sup>2</sup>, Philippe Guignard<sup>2</sup>,  
Sylvain Meyer<sup>2</sup>, Anna Pizzinat<sup>2</sup>,  
Ingo Möllers<sup>3</sup>, Dieter Jäger<sup>3</sup>

<sup>1</sup>Dipartimento di Elettronica, Politecnico di Torino, Italy,

[roberto.gaudino@polito.it](mailto:roberto.gaudino@polito.it), [daniel.cardenas@polito.it](mailto:daniel.cardenas@polito.it)

<sup>2</sup>France Telecom - Orange Labs R&D, France,

[martial.bellec@orange-ftgroup.com](mailto:martial.bellec@orange-ftgroup.com), [benoit.charbonnier@orange-ftgroup.com](mailto:benoit.charbonnier@orange-ftgroup.com),

[noella.evanno@orange-ftgroup.com](mailto:noella.evanno@orange-ftgroup.com), [philippe.guignard@orange-ftgroup.com](mailto:philippe.guignard@orange-ftgroup.com),

[sylvain.meyer@orange-ftgroup.com](mailto:sylvain.meyer@orange-ftgroup.com), [anna.pizzinat@orange-ftgroup.com](mailto:anna.pizzinat@orange-ftgroup.com)

<sup>3</sup>Zentrum für Halbleitertechnik und Optoelektronik, Universität Duisburg-Essen, Germany

[ingo.moellers@uni-due.de](mailto:ingo.moellers@uni-due.de), [dieter.jaeger@uni-due.de](mailto:dieter.jaeger@uni-due.de)

**Abstract.** Future Internet Access technologies are supposed to bring us a very performing connection to the main door of our homes. At the same time, new services and devices, as for example digital Audio-Video (AV) terminals (such as HDTV videos) and their increased use will require data transfers at speeds exceeding 1 Gbps inside the home at the horizon 2012. Both drivers lead to the deployment of a high-quality, future-proof network inside homes, to avoid a somehow ironic, but indeed possible situation, in which the Home Area Network (HAN) becomes the actual bottleneck of the full system. In this paper we review the requirements for next-generation HAN, showing that this environment may end up taking advantage of optical cabling solutions as an alternative to more traditional copper or pure wireless approaches.

**Keywords:** Home Area Network (HAN), Plastic Optical Fiber (POF), In-House-Network, Optical Communication, Network Architecture

### 1 Introduction

Broadband connection to each apartment is today becoming a commodity in most countries. Over the past year, the number of broadband subscribers in the Organization for Economic Co-operation and Development (OECD) increased 24% from 178 million in June 2006 to 221 million in June 2007, corresponding to an increase of the broadband penetration rates in the OECD from 15.1 to 18.8 subscriptions per 100 inhabitants [1]. These connections to the final users have shown in this last decade a constant increase in performance; Fiber-To-The-Home (FTTH) enabling now even higher speed services [2]. Hundreds of Mbps per user are reasonably reachable in the coming future, backed up by emerging standards such as Gigabit Passive Optical Network (GPON) [3]. FTTH is indeed today widely deployed

in countries like Japan and Korea, while in Europe and in the US several companies are developing detailed commercial plans for a mass deployment of FTTH. It is thus possible to envision in the near future a scenario in which a large amount of users will be offered a very high performance connection up to the “main door” of their homes.

Consequently, new high data rate services such as TV programs or Video on demand over xDSL are made available to the end-users and are already a commercial success [4]. New mass storage devices are on the market, such as media renderers and servers, and promise to make the digital experience even more exciting for instance with High Definition Television (HDTV) and upcoming 3D television. Digital mass storage devices gain more success to the home every day. These devices, whose pre-standardization is ongoing for instance within Wireless World Research Forum (WWRF) [5], Digital Living Network Alliance (DLNA) [6] or Home Gateway Initiative (HGI) [7], offer not only demodulation of digital broadcast programs, access to remote services by operator's networks, but also high connectivity to end devices such as TVs, home cinema or PCs. To enable the use of these devices, the trend is that it shall be possible to use them everywhere at home with high data rate connectivity to transfer content either from remote servers or between end-devices sparsely distributed everywhere at home. Moreover, end user devices are fitted with high speed interfaces to easily transfer all types of multimedia supports. In the coming future these trends will certainly make the UBB-HAN (Ultra Broad Band - Home Area Network) a convergence arena where these devices and services will have to interwork at home and in continuity to the operator's network.

This paper proposes a study of how the network inside the apartment or HAN will be implemented and deployed.

Regarding its performance, the HAN should have at least the same “quality” as the access connection. Regarding its possible architectures, some specific requirements are identified that make it quite different from the Access Network, namely:

- It should be very easily installable, possibly by unskilled technicians or even by the final user himself (“Do It Yourself” approach) and easily reconfigurable: while the Access Network, once deployed is not expected to be reconfigured for tens of years, the HAN cabling should be easily changed and adapted to suit the changing needs of the end-user.
- It should have very high data rates reaching at least 1 Gbps, even higher than the access connection speed, so to distribute very high quality video services inside the apartment.
- Newer applications may also require very low latency: highly interactive services may benefit from ultra short latency times (e.g. online gaming, online remote storage, applications using thin client terminals, telepresence...). As a reference, the latency for VDSL is greater than 16 ms, while for FTTH GPON for instance, it is below 200  $\mu$ s.

The previous discussion puts in evidence that a high speed and high Quality of Service backbone is necessary to connect together the end-devices in the HAN. In spite of its ease of deployment and rapid customer adoption a no new wire approach (power line communications or radio backbone) might not be able to achieve such goals. This paper focuses on fiber optic solutions, as it is today (2009) investigated by several large EU research projects, such as ICT-ALPHA [8]. Wireless technologies are not excluded; on the contrary, we have in mind the co-existence between a very

high capacity fiber “back-bone” inside the house, potentially capable of matching the bit-rates that the access network will provide in the long term, i.e., the Gbps data rates, and wireless solutions for any application with bit rates typically below 100 Mbps and, in particular, measurement, control and internet-of-things devices. Competition is not between optical and wireless solutions, since wireless is and will always be present inside the house. The “competition” we want to address in this paper is actually among different “wired” solutions. To build the HAN backbone, fiber optic offers some significant advantage over copper. Besides the well known higher bandwidth, one of the key points for which national operators in several countries are studying fiber inside the house, and in particular POF, is that many national regulations allow fibers to be deployed inside power ducts, while this is forbidden for unshielded twisted pair (UTP) copper cables for safety reasons. This apparently minor detail may turn out to be fundamental in existing houses, where the electrical power ducts can be re-used to run HAN fibers.

Additionally, only a fiber based HAN could be able to offer architecture flexibility, multi format transmission and protocols and robustness to medium or long term evolution of requirements. Furthermore, we keep in mind that users have developed a strong preference for a wireless end-connectivity: such a high bit rate wireless connectivity can be assured by means of radio over fiber techniques over the optical backbone.

The paper is organized as follows. After defining the specific technical UBB-HAN requirements, we will review the three different fiber types that are available for deployment of a high speed optical backbone, namely:

- Glass Optical Fibers (GOF) in their different types: Single Mode Fiber (SMF) and Multi Mode Fiber (MMF)
- Plastic Optical Fibers (POF).

Then we proceed to present several types of optical cabling and solutions that may meet the HAN requirements. Furthermore, we propose four different physical and topological architectures for the optical HAN; for each of them we show how it could be deployed, the constraints in terms of the optical system, the fiber that could be used and the transparency to services. Finally, we draw some conclusions and compare the proposed solutions with other more “traditional” approaches.

## 2 HAN Technical Requirements

We introduce in this section the technical requirements that the HAN should meet in terms of data rates and latency, considering not only today available applications and services, but also trying to provide a longer term view. The considerations presented here come from the analysis carried out inside France Télécom R&D Réseaux d'Accès (RESA) group and in the frame of the ICT-ALPHA project [8]. First of all, high data rate needs definition. Thanks to a Dynamic Bandwidth Assignment function, about 400 Mbps access speed to the telecom network can be envisaged for GPON FTTH access, relayed by more than 1 Gbps at home. Ultra small latency time, less than 1 ms round trip time can be offered thanks to simple framing and coding of short packets. This performance has to be equalled by the HAN in case of low latency services.

We have defined six kinds of profiles/classes of services for the traffic flows that have to be transmitted on the HAN, associated to different applications, and we have shown here the expected characteristics:

1. Internet browsing: A downlink variable bit rate flow up to 100 Mbps, associated to an uplink variable bit rate flow for TCP acknowledgements (< 4% of data). The next major application development in the Internet's evolution is Web3D, i.e. a system of linked interactive 3-D and 2-D environments including everything from user-specific, private applications like immersive learning simulations to virtual worlds open to anyone who wants to join. Web3D will require bit rates up to 1 Gbps [8].
2. VoIP: two fixed bit rate flows of several tens of kbps, with delay constraints lower than 200 ms depending on the end to end transmission distance.
3. Videophony and videoconferencing: two variable bit rate flows from 100 kbps to 10 Mbps, with delay constraints of 10 ms to 100 ms depending on the end to end transmission distance for the call. A future development is immersive video conferencing based on Ultra HDTV that could require up to 640 Mbps (compressed) [8].
4. File sharing (data transfer over TCP, using peer to peer): A downlink variable bit rate flow up to 1 Gbps (or anyway up to the maximum speed given by the available access network connection) associated to an uplink variable bit rate flow for TCP acknowledgements (< 4% of data).
5. Video and audio broadcasting: fixed bit rate (2 to 50 Mbps for video, up to 100 kbps for audio) with uplink low bit rate channel for RTP signals and fast channel for ARQ. This class includes also IPTV and Video on Demand (VoD). Innovative future options in this class are immersive TV (as ultra HDTV), stereoscopic (3D) TV and free viewpoint TV. These services will require bit rates up to 1 Gbps in compressed format [8].
6. Intra Home Communication: all connections internal to the home between 2 end-devices using the HAN as a cable connecting them like security, sensor and control applications. To be future-proof, very high bit rate link (up to 2 Gbps) should be deployed, for instance to feed a display with un-compressed video. Here a jitter constraint less than some 100µs is requested.

In addition to these six profiles, other emerging services are the online virtual environments, health/monitoring services and remote technical services (as remote residential backup, remote home monitoring, network watchdog, etc.) that are not very demanding in throughput for the moment [8].

In designing the next-generation HAN, attention has also to be paid to express coverage issues, not only in terms of "geographic" terms but also in terms of (semi-) static or dynamic scenarios.

- When moving from one room to the other, an end-device can be seen as either nomadic or mobile, transposing set of issues of cellular networks coverage everywhere at home. The difference between mobile and nomadic is that in the first case one can access to the services when moving from a room to another without the need of establishing a new session, conversely for nomadism it is needed to establish a new session when moving.
- As stated before, end devices such as HDTV shall be connected by some means to the HAN, raising up the question of "UBB static coverage" as well,

- End devices from different manufacturers have also to interwork either directly or via the HAN.
- When several services run simultaneously, some of them may tolerate transfer disruption or delay whereas others cannot. Quality of Service (QoS) is another dimension for consideration to guarantee quality of simultaneous services.

**Table 1:** Expected HAN specifications for the different types of digital traffic. Empty cells denote relaxed constraint.

Profile type	Throughput	Delay end-to-end (one way)	Jitter	Packet loss	TCP	Mobility
<b>Internet browsing</b>	1-100 Mbps (up to 1 Gbps for Web3D)		<10 ms	None (BER<10 <sup>-8</sup> )	Variable return flow less than 4% of data	Nomad within the HAN, but no mobility
<b>VoIP</b>	2 constant flows of a few tens of kbps (up to 96 kb/s IP)	High constraints: ~10 to 100ms < 150 ms	<20 ms	<10 <sup>-3</sup>	UDP	Mobility throughout the house
<b>Videophony and video conferencing</b>	2 flows from 128 kbps up to 10 Mbps (640 Mbps for video conferencing)	10 to 100 ms	<10 ms	<10 <sup>-5</sup>	None	Mobile or Nomad within home
<b>Files downloading (8 to 20 gigabytes), video for instance</b>	Variable high data rate 1 Gbps to prevent TCP from congestion		<10 ms	None (BER <10 <sup>-8</sup> )	Variable return flow for TCP ACK (<4% of data)	Nomad within home
<b>Video and audio broadcasting</b>	CBR flow: 2 to 50 Mbps for video, few kbps for audio Up to 1 Gbps for free viewpoint TV	Real time constraints <400ms to synch	Packet shift <1 ms	<10 <sup>-5</sup>	CBR or VBR with a low data rate	Nomad within home
<b>Intra Home Comm. DVI/HDMI</b>	Several hundreds of Mbps up to 2 Gbps	<400 ms	<1 ms		None	One single room

Table 1 summarizes all these characteristics. A key role to meet all these specifications will be played by the “Home Gateway”, i.e., the “box” that will interface the HAN to the Access Network. As it can be observed in the table the throughput requirements are up to 1 Gbps for already existing services and will surely increase with the development of new services. It is thus evident that a high speed backbone is necessary in the HAN and several solutions are possible. The ICT-OMEGA project deals with no new wires solutions (radio, power line communications and free space optics) [9]. Here, we focus on a solution with new wires. In this case two choices could be made: either copper cables or optical fiber. Even if copper cables can allow bit rates in the order of 1 Gbps, they have some disadvantages with respect to optical fibers as: the overall dimensions, aesthetic impact, electromagnetic issues, lifetime expectancy, upgradeability, flexibility, support of multiple services and formats and easiness of installation. For all these reasons, we will focus on solutions for a future HAN with optical backbone.

On top of the digital services above, the deployment of high speed wireless services is seen as key to the ease of use of the multiservice HAN, indeed wireless terminal connectivity is expected. Wireless services include cable TV, FemtoCells (indoor mobile telephone coverage) and high speed wireless technologies for Ethernet like Ultra-Wideband (UWB). All these services use at the RF level signals that are



“analog” in nature, at least in the sense that they cannot be carried directly by digital baseband modulation. Optical cabling solutions offer the possibility for semi-transparent transport of these signals by using Radio over Fiber (RoF). These analog services will induce a special set of requirements on the HAN:

1. Cable TV or Community Antenna TV (CATV): complete analog CATV spectrum distribution. This service is a HAN only service.
2. FemtoCells: same requirements as for mobile (Universal Mobile Telecommunications System (UMTS), High Speed Downlink Packet Access (HSDPA), Long Term Evolution (LTE)) distribution. This analog service may be digitized in the gateway and aggregated to the digital flow. Main requirements are symmetrical data rates of 512kbps and max. 200 ms delay.
3. Wireless home coverage using RoF: some next generation HAN architectures envision the use of next generation high speed wireless technologies. For instance, Ultra-Wideband (UWB) and 60GHz radio according to IEEE 802.15.3 can be used to cover each individual room with an extremely high-performance link. Clearly, this imposes the same requirements as the UWB protocol [10].

In addition to the mentioned technologies for Web, personal and entertainment communication applications, fixed and mobile network units for HAN sensing and control applications will gain more and more interest [11]. While these applications do not depend on high data rates and low latency times, the reliability of the network will be of utmost importance. Hence optical fibers, well known for their immunity against electromagnetic interference (EMI), can meet the requirements of these networks whereas wireless technologies like ZigBee (IEEE802.15.4) may offer mobility with eventually higher failure rate.

One of the most important issues is the interoperability and convergence of all services, devices and therewith also universally convergent protocols of the HAN [12]. As FTTH is heading towards Ethernet based GPON, Ethernet protocol seems to be the most promising standard in the upcoming HAN architectures [13], [14].

### 3 Optical technologies for the HAN backbone

Several optical media could be used to deploy the architectures described above and enable high performance HAN and so provide many broadband services. These media are classified into three main categories: standard single mode fibers, multimode fibers (silica and polymer) and microstructured fibers. Each category has specific properties in terms of attenuation loss, wavelengths to use, bandwidth, macro bending loss, connectivity issue and maturity of the technology; a comparison is necessary to assess the choice of the best trade-off between fibers, architecture and, above all, overall cost. We discuss the available options in order of increasing performance.

**POF.** With regard to polymer fibers, they are classified into two categories: PMMA step index fibers (SI-POF, 980  $\mu\text{m}$  core diameter and 1 mm fiber diameter) and perfluorinated graded index fibers (PF GI-POF, 120  $\mu\text{m}$  core diameter and 500  $\mu\text{m}$  fiber diameter). PMMA step index fibers with transmission windows in the visible wavelength range show the advantage of having a great ease of installation and connection. High Speed LEDs can be used with at the same time eye safe installation. Recent developments in an EC funded project POF-ALL show high data rates of 1Gbps over 100m and 100Mbps over more than 200 m of standard SI-POF [15].

However, graded-index fibers, more difficult to connect because of a smaller core diameter, show improved optical characteristics (lower attenuation loss and higher bandwidth especially in the infrared region) which are very close to the ones of MMF. A transmission of 40Gbps over 50m of PF GI-POF has already been shown [16].

**MMF.** Glass multimode fibers have a graded index profile and a 50 or 62.5 $\mu$ m diameter. Over the last years, their bandwidth over distance has been improved and so they become a potential medium to use in a private network to support 10 Gigabit Ethernet applications over a few hundred meters, and can thus be considered future-proof at least in terms of data rates. This solution is interesting because it enables cheap sources like VCSELs to be used and leads to an easier and lower cost connection compared to single mode fibers.

**SMF.** Standard single mode fibers, compliant with ITU-T G.652.D recommendation, are widely used in transport and access networks. They present very interesting properties in transmission (low attenuation loss, useable on a large wavelength window if the water peak is controlled and no practical limitation in terms of bandwidth over distance). For indoor cabling, fibers require in addition very low bending loss at small bending radius. For SMF used in FTTH, an optimization has been sought in order to achieve low bending and splice losses simultaneously. A new recommendation ITU-T G.657 has been defined at the beginning of 2006 to describe low bending loss single mode fibers for FTTH. Although bend optimized fibers have been developed, the connectivity issue of coupling precision of small core (9 to 10  $\mu$ m) SMF lead to high installation costs.

**Other Types.** Over recent years glass and polymer microstructured fibers e.g. Hole-Assisted Fibers (HAF) have attracted increasing attention because they offer unique optical properties and design flexibility. These fibers contain an arrangement of holes running along the fiber length and can present very interesting properties with regards to bending loss that can reach values lower than 0.1 dB/turn for a 10 mm winding diameter compared to 40 dB/turn of standard fiber. Although this technology is much less mature compared to standard single mode fibers, the performance of hole assisted fibers in terms of attenuation loss can be very interesting especially in the home context where transmission distances are short.

**Table 2. Characteristics of various kinds of optical media for HAN**

	Silica SMF	Silica MMF	POF	HAF
<b>Wavelengths to use</b>	1300 to 1600 nm	850 to 1300 nm 1240 to 1550 nm	850 to 1300 nm GI-POF 650 nm for SI-POF	C-band
<b>Attenuation loss (dB/km)</b>	<0.4 from 1300 to 1600 nm <0.25 @1550 nm	<3 @ 850 nm <0.8 @1300 nm	<40 for GI-POF <160 for SI-POF	<0.5 @ 1550 nm
<b>Bandwidth</b>	No limit available	1 Gbps over 1 km 10 Gbps over 300 m	40Gbps over 50m PF GI-POF 100Mbps over 200m and 1Gbps over 100m SI-POF	No limit available
<b>Macrobending loss</b>	$\varnothing=30$ mm, < 0.25 dB/10 turns @1550nm	$\varnothing = 75$ mm, < 0.5dB/100 turns @ 850 & 1300 nm	$\varnothing \geq 40$ mm GI-POF $\varnothing \geq 60$ mm SI-POF	$\varnothing = 15$ mm, < 0.1 dB/10 turns
<b>Connection</b>	<0.1dB splicing <0.3dB connect. The operations require specific set-ups	Same values than SMF, coupling is simpler due to large core diameter	For GI-POF, more simple coupling (< 1 dB) For SI-POF, "DIY" concept (< 2 dB)	More difficult connection because of the holes but losses < 0.5 dB

**Comparison of fiber based technologies.** Each solution described above presents specific characteristics in term of attenuation loss, bandwidth, and connectivity. These ones are summarized in table 2 in order to provide some information about the kind of fiber which is more adapted to the architectures (PtP, PtM, PON...) that will be presented in the next section. Today the single mode fiber is the most mature fiber technology used in access, metro and other optical telecommunication backbone networks. For the other fiber types, a survey activity is necessary to compare various solutions and architectures. For our study of the home cabling solutions, two cases must be considered: new homes and existing homes.

In new homes, a pre-cabling solution is proposed with the installation of tubes or ducts at the time of the building construction (e.g. in the walls). Afterwards, a small optical cable can be pushed or blown on demand through these tubes. In this configuration, all fiber types (POF, MMF, SMF, HAF) can be used because the installation constraints are not strict (mostly regarding bending loss requirements).

In existing homes the situation is different. The trend is to develop "Do It Yourself" (DIY) cabling solutions not requiring the intervention of an external technician. These solutions must provide all infrastructure components required to simply and safely install a residential home network. The POF technology is well adapted to fulfil these requirements. The PF GI-POF fiber is also suitable, because the fiber termination is fast, the cable is easy to handle and can be installed in an environment with corners and doors. Regarding to the use of Silica fibers in HAN, the connection issue remains the most difficult point although some works are in progress to develop field mounting connectors and installation technicians are being trained as well as field adapted termination tools are developed for FTTH.

New fields of optical communication systems for HANs are free space optical communication systems. White light that is inherently used for illumination can be modulated with the data; recent developments show that this technology may offer an alternative to radio transmission systems like IEEE 802.11a/b/g/n for in house networking [17]. Data rates of 40 Mbps with OOK and 100 Mbps with subcarrier multiplexing were achieved over very short distances.

## 4 HAN Architecture

The set of specifications shown in Section 2 are, at least for the most advanced scenarios, very challenging to be met, and clearly show the rationale toward optical cabling inside the house. This chapter analyzes the possible architectures of the HAN focusing specifically on optical fiber solutions.

Starting from a quite general description, the HAN should connect the Home Gateway to several "access points" (AP) scattered around the house. The AP is defined here as a "plug" through which the HAN services can be accessed, in order to provide suitable connectivity throughout the house. The usual figure is at least one AP per used room (lounge, kitchen, home office, bedrooms), leading to approximately a minimum of five AP per house. As a maximum, a significantly higher number of AP, up to some tens, can be necessary covering each room and giving extra ease of use in the most connected rooms (e.g. home office). The maximum cabling distance is typically considered to be 50 m.

#### 4.1 Point-to-point (P2P) solutions (active star)

Today, most home networks are based on a home gateway typically placed at the apartment entrance, where the external access network is terminated. Terminal premises are connected to this gateway through point-to-point links (or wireless links). In this case, the gateway acts as a central switch, and the resulting star architecture is indicated here as P2P. The simplest configuration is when all data are digital and encapsulated into the same protocol, as shown in Fig. 1. For instance, the gateway can be an Ethernet switch or an IP router.

Today, in residential areas, the most used cabling system is multimedia copper cabling (UTP/STP) with RJ45 outlets. The progressive introduction of optical cables in the HAN backbone can be made by




Fig. 1. P2P traditional architecture

proposing a hybrid wiring solution containing either a hybrid cable (a copper cable and an optical cable placed inside a same sheath) or a copper cable juxtaposed with an empty tube in which a small optical cable can be pushed or blown on demand in the future. The last solution is more economical today because the overcharge due to the optical fiber is pushed back in the future. As well, the choice of fiber type is delayed and potentially, the optical cable can be upgraded in the future by one with more bandwidth if necessary.

In this simplest application, the key advantages given by optical cabling are: very high available bandwidth, complete immunity to electromagnetic interference, total protection against eavesdropping.

Preliminary versions of home-gateways with optical ports (typically POF) are currently being deployed on the market, for instance by Avago/Infineon and Siemens.

**Network performances for an active star architecture.** As mentioned above, an active star is based only on point to point links, with no major issue concerning optical budget or system bandwidth, since a link only transports the data flow dedicated to one device. If all fiber types may be used in an active star, it is relevant to focus on POF technology, as this will be the first fiber entering the home, thanks to its easy implementation related to material and geometrical properties. In the same time, this technology exhibits the poorest performances, compared to other fiber types. Nevertheless, offering 100Mbps bit rate over about 100m with standard modulation, its performances match the requirements for the first generation services in the home network. We have to keep in mind some limitations that could appear in a real cabling, with additional losses due to unavoidable fiber bends. Table 3 provides some result examples obtained while evaluating performances of bi-fiber POF media-converters. It shows the variation of the reach of the link versus the number of fiber bends. The average transmitted power and receiver sensitivity were respectively -4.87 dBm and -27.20 dBm, giving an average optical budget of 22.33 dB. Fiber attenuation @ 650 nm and additional loss per bend were 0.28 dB/km and 0.9 dB, respectively. Even with 10 bends, the reach remains compatible with lengths required in the HAN.

**Table 3. Reach vs. number of fiber bends with 100 Mbps duplex POF media-converters**



Reach vs. number of fiber bends using standard SI-POF					
0 bends	2 bends	4 bends	6 bends	8 bends	10 bends
98.7 m	90.7 m	82.8 m	74.8 m	66.9 m	58.9 m

To show the importance of the fiber bends, other tests were carried out with mono-fiber media-converters. Their advantage is a very simple use, with only one fiber for the two transmitting directions. In the other side, as these devices integrate splitters or multiplexers, the optical budget is reduced, and we measured an average value of only 12.73 dB. With no bends, an acceptable reach of 56 m was obtained, decreasing to about 16 m with 10 bends, which is quite insufficient.

#### 4.2 MultiPoint architectures

To increase flexibility in the HAN, more transparency is required. Optical technologies provide solutions with multipoint transparent architectures.

**1xN "PON – like" architecture.** The idea here is to transpose access PON concepts to the HAN context, as shown in Fig. 2. The gateway is then a Home Network Optical Line Termination (HN-OLT). Premises are connected using Home Network Optical Network Termination (HN-ONT). All PON mechanisms can be reused. Real time traffic and best effort traffic coexist with a good QoS, as PON allows resource reservation. Increasing the number of AP is very easy: PON are presently sized with up to 128 ONT. A major advantage is the optical transparency as services can be transmitted in various formats (RoF, CATV ...) using one or more wavelengths in overlay. The main issue today is to reduce the cost of PON equipments, taking into account the relaxed specifications compared to access context (reduced distance, tolerable increased failure rate etc...).

**Fig. 2.** 1xN PON-like architecture**Fig. 3.** NXN "LAN – like" architecture

**NxN "LAN – like" architecture.** Solutions derived from LAN and Ethernet context can be used. All premises can be connected to an optical bus based on a NxN splitter (see Fig. 3), using optical transceivers. The gateway is on the same hierarchical level than other premises and the network manages itself. The advantage is the simplicity of the Medium Access Control (MAC) e.g. CSMA/CD. The drawback is that this type

of network has been designed for best effort traffic and no QoS is currently considered. But, with short transmission distances, the limited number of AP - compared to the LAN context and considering the increasing speed of transceivers - real time and best effort traffic can be transported simultaneously. All previous remarks on transparency, in order to implement services in various formats on different wavelengths in overlay, are also relevant in this configuration.

**Broadcast & Select CWDM architecture.** This configuration, shown in Fig. 4, is based again on a NxN coupler. A pair of fibers (TX and RX) connects one coupler to each AP. Each TX fiber is connected to all RX fibers through this coupler. This fully transparent architecture uses CWDM technology to pile different network layers in parallel without interference between each other. Different topologies (P2P, 1xN, NxN), different protocols and different formats can coexist. At any AP, an optical filter (e.g. CWDM thin film filter) is used to separate or aggregate traffic to the WDM spectrum on the wavelength specific to the desired service. Note that all services (i.e. all wavelengths) are accessible at each AP using the correct filter and the filters can be cascaded to access several services at the same AP.




Fig. 4. Broadcast & Select CWDM architecture

**Network performances for multipoint architectures.** For the "PON-like" and the "LAN-like" architectures, 1xN or NxN splitters are required. In addition, for the Broadcast & Select CWDM solution, CWDM technology has to be implemented. For these reasons, SMF appears as the best choice, since all the required components are available for this technology. With SMF, the capacity in terms of bit rate can be considered as quasi unlimited for HAN applications. The major issue could be related to the optical budget associated to each application (digital, terrestrial TV, RoF).

For basic "PON-like" and "LAN-like" solutions, i.e. without other services in overlay on additional wavelengths, the optical losses are quite similar. With a 1x16 or a 16x16 splitter, taking four connectors and a fiber length of 100m into account, total losses were estimated to about 14.8dB. With a 1x32 or a 32x32 splitter, the total losses are then equal to 18.8 dB.

For the "PON-like" and "LAN-like" architectures with additional services in overlay, or for the Broadcast & Select CWDM solution, it is necessary to insert wavelength filters, for example add & drop modules, at least at the receiver side, with an additional loss of about 1.7 dB (including connectors). Total losses are then about 16.5 dB with a 1x16 or a 16x16 splitter, or 20.5 dB with a 1x32 or a 32x32 splitter.

Various systems were used to experiment all these configurations:

- a PON system, 2.5 Gbps downstream, 1 Gbps upstream
- 100 Mbps Ethernet LAN prototypes cards, implementing CSMA/CD protocol to work on an optical bus (the passive star)



- Point to point digital links, using SFP at 1 Gbps
- Terrestrial TV transmission equipment
- RoF system based on UWB with two 5 m aerial radio links.

Table 4 provides the optical budget values for the different applications, in the conditions of our lab demonstration.

**Table 4. Optical budgets for the different tested applications**

Applications	Optical budget
PON-like	25.2 dB
LAN-like	30.46 dB
Point-to-Point	28.42 dB
Terrestrial TV	20.8 dB
RoF	27 dB (10 dB antenna, two 5m aerial radio links)

Some comments to sum up the main results:

- For basic "PON-like" and "LAN-like" applications, we observe that both systems work in a 1x32 or 32x32 configuration with a comfortable margin, allowing a significant increase of the number of ports of the splitter.
- These conclusions are still true within a CWDM context, "PON-like" and "LAN-like" with additional applications on overlay, or Broadcast & Select CWDM solution. The margin is enough to implement a splitter with at least 64 ports.
- The same conclusions apply for 1 Gbps point to point link, and RoF transmission with a 10 dB antenna configuration: both are compatible with a CWDM context with a 32 port splitter.
- The most critical applications remain terrestrial TV and RoF transmission. Implemented in overlay, the optical budget is very close to the transmission limit (even below). The use of a 16 port splitter may be preferable for home networks.

## 5 Conclusions

With the multiplication of Digital Mass Storage devices in the home, it is now required to access and transfer vast amounts of data from any location in the home. A subjacent requirement for the HAN is to support data transmissions comparable to hard-disk to hard-disk transfer rates inside a single computer. Throughputs up to 1 Gbps have to be reached for that with, at the same time, high reliability for sensor networks and control applications. Moreover, higher rates will be necessary for uncompressed HDTV streams but these should be supported by short fixed links and not necessarily by the whole network. Lastly, analogue or quasi-analogue services could be added on top of these requirements to comply with the end-user requirements for a wireless end-connectivity enabling flexibility and even mobility within the house. Furthermore, the increase in available access rates makes it possible for applications requesting large transfer rates to run in parallel within the same home thus pushing higher the data rate demands for the HAN.

In this context the use of optical fiber to provide an UBB-HAN backbone for Future Internet is seen as key to the continued development of communications and infotainment. The deployment, architecture and fiber type used for such a backbone will be strongly influenced by economical factors, end-user requirements and installation constraints. Two main possible directions for UBB-HAN are identified as:



- Plastic Optical Fiber infrastructure to provide a point to point architecture suitable for already constructed houses where the installation of a new cable can be performed by the user itself. This infrastructure will be suitable for the performance required in a medium term scenario.
- Silica Fiber infrastructure for new houses where the optical cables will be installed in ducts running in the walls at construction time. The architecture will have the possibility to evolve from point to point offering sufficient performance again for a medium term scenario to fully transparent multipoint to multipoint network able to respond to any future bandwidth requirements. For this scenario, SMF is believed to be the best choice as it guarantees the long term suitability of the network and it benefits from the economy of scale and experience gained from current FTTH deployments.

Other solutions are also considered using silica fiber type resilient to bending loss suitable for retrofitting existing buildings or higher bandwidth plastic optical fiber which could be suitable for backbone deployments in new buildings, but these less technically mature solutions will need further survey and experience.

**Acknowledgments.** The work was carried out with the support of the BONE-project ("Building the Future Optical Network in Europe" Network of Excellence) and the ALPHA-project ("Architectures for Flexible Photonics Home and Access networks", Integrated Project). The OMEGA-project ("Home Gigabit Access", Integrated Project) contributed to the section on HAN technical requirements. These projects are funded by the European Commission through the 7th ICT Program.

## References

1. OECD, „OECD Communications Outlook 2007“, ISBN: 978-92-64-00704-8, July 2007 (<http://www.oecd.org/sti/telecom/outlook>)
2. Very High Speed pilot program ("Fiber To The Home") [http://www.francetelecom.com/en/financials/journalists/press\\_releases/CP\\_old/cp060117.html](http://www.francetelecom.com/en/financials/journalists/press_releases/CP_old/cp060117.html)
3. GPON 2.4 G [http://www.itu.int/newsarchive/press\\_releases/2003/04.html](http://www.itu.int/newsarchive/press_releases/2003/04.html)
4. MaLigneTV. Digital TV on ADSL. <http://www.malignetv.fr/>
5. WWRF SIG 4 "Home and enterprise networks" <http://www.wireless-world-research.org/?id=92>
6. Digital Living Network Alliance <http://www.dlna.org/home>
7. Home Gateway Initiative <http://www.homegatewayinitiative.org/>
8. ICT-ALPHA project deliverable D1.1p "Specification of services for access, mobile and in-building networks", available at <http://www.ict-alpha.eu>
9. <http://www.ict-omega.eu>
10. Standard ECMA-368, Geneva, 1st ed. Dec.2005
11. Akyildiz, I.F.; et al.," A survey on sensor networks", IEEE Comm. Mag., Vol. 40, Issue 8, Aug. 2002
12. Miller B. A. et al., "Home Networking with Universal Plug and Play", IEEE Comm. Mag., Vol. 29, Issue 12, Dec. 2001
13. Green, P.E.,"Fiber to the Home: The Next Big Broadband Thing", IEEE Comm. Mag., 42/9, Sep. 2004
14. Topalis, E.,et al., "A generic network management architecture targeted to support home automation networks and home Internet connectivity", IEEE Transactions on Consumer Electronics, Vol. 46, Issue 1, Feb. 2000
15. Gaudino R. et al., "Status and Recent Results from the POF-ALL EU Project: Toward Improved Capacity and New Application of Large-Core POF, Proceedings of Int. Conf. on POF, pp. 79-84, Turin, Italy (2007)
16. Schöllmann, S. et al., „First Experimental Transmission over 50 m GI-POF at 40 Gb/s for Variable Launching Offsets", ECOC 2007
17. Grubor J. et al., "Wireless High-Speed Data Transmission with Phosphorescent White-Light LEDs", ECOC 2007

## DICONET: future generation transparent networking with dynamic impairment awareness<sup>\*</sup>

Ioannis Tomkos<sup>1</sup>, Yvan Pointurier<sup>1</sup>, Siamak Azodolmolky<sup>1</sup>, Michael Eiselt<sup>2</sup>,  
Thierry Zami<sup>3</sup>, Radoslaw Piesiewicz<sup>4</sup>, Chava Vijaya Saradhi<sup>4</sup>, Matthias  
Gunkel<sup>5</sup>, Uri Mahlab<sup>6</sup>, Ming Chen<sup>7</sup>, Yabin Ye<sup>7</sup>, Mario Pickavet<sup>8</sup>, Maurice  
Gagnaire<sup>9</sup>, Emmanouel Varvarigos<sup>10</sup>, Josep Solé Pareta<sup>11</sup>, Reza Nejabati<sup>12</sup>,  
Yixuan Qin<sup>12</sup>, Dimitra Simeonidou<sup>12</sup>

<sup>1</sup> Athens Information Technology (AIT)

<sup>2</sup> ADVA AG Optical Networking

<sup>3</sup> Alcatel-Lucent Bell Labs France

<sup>4</sup> Center of REsearch And Telecommunication Experimentations for NETworked  
communities (Create-NET)

<sup>5</sup> Deutsche Telekom AG/T-Systems

<sup>6</sup> ECI Telecom

<sup>7</sup> Huawei Technologies Deutschland GmbH

<sup>8</sup> Interdisciplinair Instituut voor Breedband Technologie, VZW (IBBT)

<sup>9</sup> Telecom ParisTech

<sup>10</sup> Research Academic Computer Technology Institute (RACTI)

<sup>11</sup> Universitat Politècnica de Catalunya (UPC)

<sup>12</sup> University of Essex

**Abstract.** Transparent networks are widely seen as the prime candidates for the core network technology of the future. These networks provide ultra high speed end-to-end connectivity with high quality of service and failure resiliency. A downside of transparency is the accumulation of physical impairments over long distances, which are difficult to mitigate using physical-layer techniques only, and the novel challenges in fault detection/localization. We present here the DICONET project, a set of techniques and algorithms implemented at multiple layers, culminating with the physical implementation of a transparent optical network on a testbed. DICONET consists of a set of impairment-aware network management algorithms, such as routing and wavelength assignment, monitoring, failure localization, rerouting, all integrated within a unified control plane, which extends known solutions to include the impairment-awareness of the underlying layers.

### 1 Introduction

With ever increasing bandwidth needs, spurred by the emergence of increasingly bandwidth demanding applications such as e-science, e-health, and high-definition video-on-demand and video broadcasting, the future Internet will need

<sup>\*</sup> Corresponding author: Ioannis Tomkos, Athens Information Technology, 0.8km Markopoulou, Peania 19002, Greece. Email: itom@ait.edu.gr

an ultra high-speed backbone. In addition to the large raw available bandwidth, the future Internet is expected to offer such services as high resiliency to hardware failure, the possibility to request and be granted the utilization of resources with guaranteed quality of service (QoS), for instance, in terms of signal quality, or in terms of guaranteed available bandwidth. The current infrastructure is called “opaque” because signals are regenerated by electronic devices at every node. This makes ultra-high bandwidth and QoS guarantees/management difficult, not scalable, and cost-ineffective. A departure from opaqueness is offered by the possibility to switch signals in the optical domain, rather than in the electrical domain, using so-called “optical crossconnects” (OXCs). In transparent optical networks, where light is switched in the optical domain, data is carried over pre-established circuits called “lightpaths”, consisting of a route and a wavelength. The transition from opaque to transparent networking, however, is not possible for all backbone networks. Indeed, transparency implies the transmission of signals over very long distances with no electrical regeneration. Physical impairments accumulate over such distances (potentially thousands of kilometers for very large, continental-sized networks), making error free transmission difficult or impossible to achieve. To overcome this issue, it is possible to regenerate signals at a small number of sites, thereby increasing the total distance that can be spanned by lightpaths. Such networks where regeneration is present at certain nodes only is a compromise between transparent networks and opaque networks are called semi-transparent “managed reach” optical networks.

Transparency (full or partial) in optical networks eliminates the electronic bottleneck, thereby allowing ultra-high datarates in core networks in a cost effective fashion, and the utilization of the circuit-switched technology is an enabling component for many traffic engineering techniques aiming at providing end-to-end QoS and high resiliency. The evolution of networks from opaqueness to transparency requires new hardware, such as the transition from electrical switches to OXCs, but also novel higher layer techniques and protocols to operate and manage the network in order to guarantee that the benefits of transparency for the end-users (QoS, resilience) are actually attained despite the adverse effects of the physical layer. In addition to impairment accumulation, transparent networks make failure localization difficult: indeed, it is not possible to know which equipment on a path is responsible for a fault that is detected by standard electronic hardware in an end-to-end fashion. Electrical regeneration at all nodes permits to isolate faults to the link/node where the fault occurred. Such fault isolation, and, in some sense, mitigation, inherent to opaque networks, is removed in transparent networks. Efficient fault detection and mitigation is needed to achieve high resilience purposes and meet Service Level Agreements (SLAs).

The two aforementioned issues inherent to transparency — enhanced physical impairments and change of paradigm needed in failure localization — have prevented operators to deploy transparent networks, despite the benefits in both CAPEX and OPEX. The DICONET project (Dynamic Impairment Constraint networking for transparent mesh Optical NETworks) project brings answers to these open issues [1].




Fig. 1. DICONET vision.

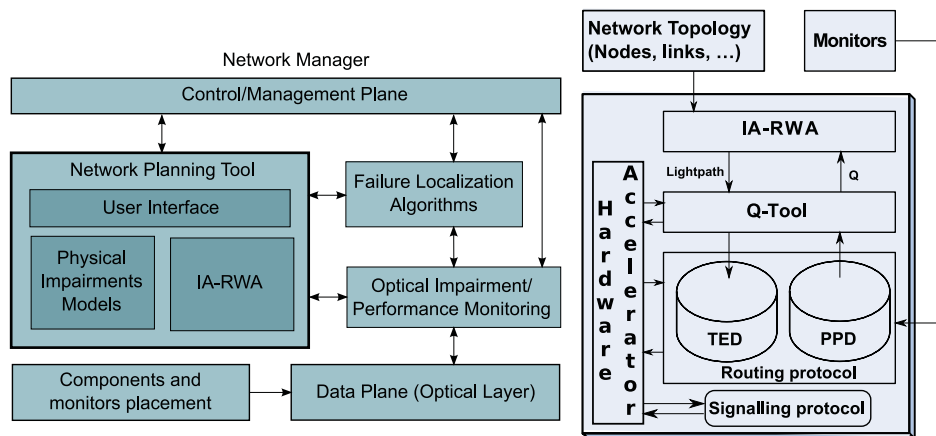


Fig. 2. DICONET components.

Fig. 3. Control plane architecture: overview.

As depicted in Fig. 1, we propose to extend the core optical networks intelligence to the data plane on the optical layer using a crosslayer approach. To achieve this, a network planning tool is responsible for integrating all the information needed for the control plane to make decisions as is shown in Fig. 2. For instance, the network planning tool gathers the topology and monitoring data to compute Q-factors (a measure of signals Quality of Transmission, QoT) on which control plane routing and wavelength assignment (RWA is the process of finding a lightpath for a network utilization demand) decisions are based. The network planning tool also deals with complex situations where monitoring information is missing and must be estimated from past or other measurements.

The techniques and algorithms developed within the DICONET project are validated using simulations and experiments — mainly using a small-scale testbed mixing real and emulated all-optical nodes. In particular, simulations are carried on over two realistic topologies for which all necessary physical parameters are known: the “Deutsche Telekom” topology, a country-sized (diameter: 800 km) network of 14 nodes and 23 bidirectional links, and the “GEANT-2” topology, a continental-sized (diameter: 7000 km) network of 34 nodes and 52 bidirectional links. Although the Deutsche Telekom topology can be used to simulate fully transparent networks, the GEANT-2 topology is too large to be fully transparent, and hence is used to validate algorithms specific to semi-transparent scenarios such as regenerator placement.

In the remainder of this article, the elements revolving around the network planning tool are described: physical layer-related techniques and tools in Section 2, Impairment-Aware RWA in Section 3, the fault management techniques 4, and the control plane 5. Future work that goes beyond DICONET is outlined in the concluding section.

## 2 Physical Layer modeling, monitoring, estimation

In current backbone networks, where signals are regenerated by optical-electrical-optical converters at every node, achieving error-free transmission (e.g., Bit-Error Rate:  $\text{BER} < 10^{-5}$  before Forward Error Correction) is done by proper link engineering. In future generation transparent core optical networks, lightpaths traverse several links without regeneration and physical impairments accumulate over potentially very long propagation distances. Thus, real-time or near real-time monitoring of the physical impairments that have an impact on signals’ BER is needed, in order to provide the control plane the adequate information to *detect failures or equipment quality degradation/aging*, ensure that *Service Level Agreements* in terms of signals’ quality of transmission are met, and to make appropriate *RWA or rerouting* decisions to mitigate physical layer impairments.

Relevant physical impairments for 10Gbps transparent optical networks with standard fiber, dispersion map and grid spacing were shown to be the interplay between chromatic dispersion and self-phase modulation, amplifier noise, and multichannel nonlinear effects (crossphase modulation, XPM, and four wave mixing, FWM). Multichannel impairments (XPM and FWM) cause the QoT of

different lightpaths to be interdependent, since establishing a new lightpath may increase the impairments seen by other already established lightpaths to increase. As will be seen shortly in Section 3, this interdependence renders physical-layer only network design cost-ineffective and paves the way for more advanced, cross-layer techniques, to mitigate physical layer impairments. Although relevant for different architectures (for instance, with tighter grids or higher bitrates), node crosstalk resulting from optical single leaks at nodes can be ignored with the standard 50 GHz ITU grid spacing for 10Gbps NRZ-modulated signals. At 40Gbps, polarization mode dispersion (PMD) has to be accounted for.

We combine two complementary approaches to know the quality of the transmission seen by lightpaths — which can be lightpaths already established in the network, or candidate lightpaths selected by a RWA algorithm. Hardware monitors, placed at strategic locations, inform the control plane about the current state of the network. We classify monitors into the following two types: Optical Impairment Monitors (OIM) are deployed at the *link level* and allow for efficient failure localization and lightpaths transmission quality *estimation*, while Optical Performance Monitors (OPM) give *measurements* of *end-to-end* lightpaths transmission quality. To achieve efficient and useful monitoring, it is necessary to deploy the following OIM equipment: optical power, optical signal to noise ratio (OSNR), chromatic dispersion and PMD monitors. While optical power monitors are typically implemented by default at all optical nodes, the other OIMs are expensive equipment which can only be deployed at a few locations. In the case of OSNR, however, it was shown recently that, by adding power monitors at non-standard locations (e.g., inline amplification sites) and by using appropriate estimation method, the information gained with the additional power monitors is sufficient to accurately estimate OSNR. End-to-end lightpaths' quality of transmission will be monitored through Q-factor monitors — the so-called “Q-factor”, defined as the ratio  $(\mu_1 - \mu_0)/(\sigma_0 + \sigma_1)$  between the mean difference  $(\mu_1 - \mu_0)$  over the sum of the standard deviations  $(\sigma_0 + \sigma_1)$  of sampled “0” and “1” symbols after photodetection. Q-factor monitoring is generally expensive, especially because the clock recovery step needed to perform synchronous sampling to obtain the relevant quantities  $\mu_0, \mu_1, \sigma_0, \sigma_1$  at the optimal sampling time. We have proposed novel asynchronous Q-factor monitoring techniques that bypass the need for clock recovery [1].

In some cases, no OPM information is available but the BER of a lightpath needs to be known. This can be due to the lack of OPM hardware at a specific location, or because the BER of a lightpath that is not yet established (hence does not physically exist). In such cases, *monitoring* is not possible and *estimation techniques* have to be employed. In particular, we use a “QTool” to estimate Q-factors for lightpaths not monitored or established. “QTool” relies on well-known physical layer models [2] and is fed with the OIM information. If OIM data needed by QTool is missing, it is interpolated from available data and analytical models. The QTool is an important part of a transparent network architecture as the lightpath establishment module relies on it to estimate a priori the Q factor of candidate lightpaths. Underestimation of Q may lead to rejec-



tion of lightpaths which Q factor is actually acceptable, while overestimation may lead to acceptance of lightpaths which Q factor is actually too low to guarantee error-free transmission. In addition, overdesigning is needed to counter adversarial effects of physical parameters measurement/estimation errors. We have shown that an uncertainty in powers (which are in turn needed to compute Q factors) above 0.5 dB leads an explosion in the amount of the resulting needed overdesigning (e.g., in terms of regenerators) by network designers [3].

### 3 Impairment aware lightpath routing

In transparent networks, data is transmitted over “lightpaths”, the combination of a route and a wavelength. Because all-optical wavelength conversion is still experimental, once a signal is launched over a channel, it has to remain on the same channel (wavelength) from end-to-end, or from electrical regenerator to electrical regenerator in the case of semi-transparent networks. Such constraint, unique to all-optical networks, is called “wavelength continuity constraint”. RWA refers to the problem of finding a route and a wavelength for each demand in a set; the demands can be static (e.g., known in advance, for long-term capacity allocation) or dynamic (lightpath demands arriving at the network management system potentially randomly, e.g. for e-science, e-health, content delivery networks, and other very high-speed applications). Even if the demand set is known, as in the static case, the RWA problem is known to be NP-complete [4], warranting the search for heuristics. In the case of transparent networks, the situation is made more complicated by the accumulation of physical impairments as signals propagate through the links. Indeed, mechanisms must be devised to ensure that signals’ quality remain above a predetermined threshold to guarantee error-free communication even as physical impairments (potentially originating from interaction with other signals as in XPM and FWM) accumulate. When lightpaths are blocked because its QoT would be insufficient, or because establishing it would cause the QoT of other lightpaths to become insufficient, “QoT blocking” occurs. The QoT of a signal depends not only on the signal’s path, but also on the existence of other signals in the network through nonlinear effects (XPM/FWM). This further adds to the complexity of the RWA problem in transparent networks. RWA algorithms that take into account physical layer information/impairments to make a decision are called impairment-aware RWA (IA-RWA) algorithms.

Although much research has been devoted to the online routing case where demand arrivals and terminations are dynamic (see the survey [5], where we compared and categorized more than 80 RWA algorithms), static or offline RWA where the lightpaths’ demand set is known in advance by the network operator has been far less studied.

In particular, in [6], we solved the offline IA-RWA problem using a linear programming (LP) formulation. Although the RWA problem is originally an integer, not linear, programming problem, we used computationally efficient LP techniques and obtained integer solutions by proper piecewise linearization of



the cost function and rounding techniques. Physical layer impairments are accounted for directly within the programming formulation by adding two sets of constraints, one seeking to minimize lightpaths' lengths and the other to minimize interference between copropagating lightpaths. In another approach, all physical constraints were integrated in a single set of constraints by converting the impact of all effects into a single "noise" parameter. It was shown through simulations that considering physical impairments within the LP formulation decreases sharply demand blocking rate compared with the case where physical impairments are evaluated as a final check. The single-parameter formulation decreases blocking rate even further.

Although LP is a valuable tool to compute (near-)optimal solutions to the IA-RWA problem, it is computationally expensive to run and faster heuristics should be considered too. In [7], we developed such heuristics, based on a preprocessing step where demands are ordered taking into account their expected resource consumption: demands using more capacity, which are expected to be more difficult to accommodate, are allocated first. Since the algorithm relies on simple heuristics, it is computationally efficient and we showed through simulations that blocking rate was decreased when using this pre-ordering heuristic. The algorithm was also adapted to the case where some lightpaths needed protection.

## 4 Failure localization

By design, transparent nodes (OXCs) do not decode signals that traverse them and component failure can only be detected in an end-to-end, as opposed to local, fashion. This makes failure localization difficult in transparent networks. At the same time, failure recovery and localization is very important in networks where a single link can carry dozens of wavelengths modulated at 10-40Gbps each. It was shown in a recent study that CAPEX gains of shared (e.g. 1:1) protection over dedicated (e.g. 1+1) path protection are negligible in transparent optical networks [8]; however, it should be kept in mind that shared protection incurs higher OPEX than dedicated protection.

Single failure detection and recovery is a well-known topic in transparent optical networks, but the multiple failure case was far less studied. This problem was shown to be NP-complete and heuristics are needed. We will propose a cross-layer failure detection algorithm that is able to account for the differences at the physical layer of the various network components and effects of potential effects (e.g., in terms of failure propagation) to effectively determine the origins of detected failures. Another algorithm has been proposed to correlate multiple failures locally at any node and to discover their tracks through the network. To identify the origin and nature of the detected performance degradation, the algorithm requires up-to-date connection and monitoring information of any established lightpath, on the input and output side of each node in the network. This algorithm mainly runs a localization procedure, which will be initiated at the downstream node that first detects serious performance degradation at an arbitrary lightpath on its output side. Once the origins of the detected failures

have been localized, the network management system can then make accurate decisions to achieve finer grain recovery switching actions. Failure detection relies on the knowledge of the network's physical layer state, which is addressed in the following section through control plane mechanisms.

## 5 Control Plane

Monitors, nodes, network management system communicate through a so-called "control plane", a set of protocols which makes interactions between all elements in the network possible. The control plane is ultimately responsible for lightpath establishment, tearing down, rerouting, failure detection, dissemination of hardware monitoring information, and traffic engineering in general. Several protocols well-adapted to circuit-switching architectures at the wavelength granularity already exist; however, none incorporates physical layer characteristics. For this reason, we propose to extend well-known protocols to include optical layer characteristics in control planes. DICONET will evaluate and implement two distinct GMPLS-based control planes: a centralized control plane and a distributed control plane, both relying on the building blocks depicted in Fig. 3.

The IA-RWA module makes RWA decisions based on QTool computations. The QTool computes Q factors based on information contained in two databases: the "traffic engineering database" (TED), which contains resource availability information while the "physical parameters database" (PPD) contains impairment-related information.

In the centralized approach, all computations are done by a dedicated process called "Path Computing Element" (PCE) on a dedicated server. In Fig. 3, the PCE groups the IA-RWA and QTool blocks as well as the hardware acceleration module, which is presented in detail further in this section. The control plane learns monitoring information via extensions of the routing protocol (OSPF-TE) or with SNMP queries. Standard RSVP-TE is used for the signaling.

In the distributed approach, an instance of the control plane runs on every node. A link-state routing protocol, OSPF-TE, can be used to disseminate the monitor information to all nodes. By design, information in the TED or PPD can be outdated as OSPF-TE takes a positive time to converge. Signaling is done by a modified version of RSVP-TE that is able to account for impairments in real time, as a lightpath is being established [9]. This is done to counter any stalled or outdated information in the local TED/PPD.

These two solutions are investigated in depth to determine which is best suited to transparent core optical network. The centralized approach suffers from the lack of scalability and the single point of failure issues, but is able to compute optimal paths as all relevant and current data is known by the control plane. The distributed approach is more scalable but may make sub-optimal choices especially if the data it bases its decisions on is outdated.

Signaling and routing protocols of the GMPLS control plane stack are complex with many messages, parameters and procedures. Frequent updates are needed to account for the dynamic network state. Thus, current implemen-

tations of GMPLS control plane protocols (signaling and routing) are purely software-based. The software-based implementation can handle a complex control plane and in a flexible fashion. However, the performance of the software implementation can degrade dramatically when the network size, connectivity and number of calls or connection requests increase. GMPLS implementations in software are rarely capable to handle large number of requests in a few milliseconds. Considering that the connection setup time per optical switch is in the order of milliseconds, this bottleneck of control plane can have adverse effect on the network performance. The DICONET control plane uses extended GMPLS to facilitate IA-RWA and fault localization which makes the software stack even more complex and computationally intensive compared to standard GMPLS implementations. Example of these functions include impairment-aware forwarding and path selection and fault localization/detection, and in particular online physical layer impairment processing like the QTool. Therefore, to improve performance of the control plane, a hardware implementation of some of computationally intensive control protocol procedures can be envisioned. The main objective is to overcome the complexity of the control plane stack by implementing only time critical procedures of the DICONET control protocols in Field Programmable Gate Arrays (FPGAs) with embedded network processor in the form of a control protocol hardware accelerator. The hardware can potentially perform control protocol procedures up to 1000 times faster than the equivalent software based approach. Some of the protocols mentioned in the three architectures needs extensions to standard protocols (e.g., RSVP-TE, OSPF-TE) and standardization efforts for an impairment-aware control plane has already started [9, 10].

## 6 Future work

DICONET is a broad project which aims at making concrete many requirements of the core optical network of the future, focusing on the interplay between the physical layer, which effects cannot be ignored at the core networks scale, and the upper layers, by using an adapted network planning tool and control plane. The project has started in January 2008 and will end in June 2010. By that time, a working implementation of the techniques developed will be running over a testbed mixing real and emulated OXCs. The so-called network planning tool will integrate all developed techniques and a cross-layer control plane will be implemented and running. Key functions will be implemented into FPGA, a demonstration will take place, control plane extensions will be standardized and the economic viability of the DICONET approach to future generation optical networks will be proved with a techno-economic study. Because we can leverage the strong complementarity between academic and industrial partners, it is expected that the prototype can be used as a starting point for industry-grade development and deployment of future generation optical networks.

Although the DICONET project addresses challenging issues, such as creating a control plane for a new kind of network and developing failure localization

algorithm, it is only a first step towards even more complex architectures. Indeed, even when transparency is achieved and the issues addressed by DICONET are solved through cross-layer design for core networks, the segmentation between access and core networks, and the isolation between core networks operated by different managing entities, means that QoS from end-user to end-user is currently impossible to achieve. Therefore, the research presented here in the context of core networks should be extended to provide true end-to-end connectivity and transparency or partial transparency, through appropriate technical and standardization efforts.

## 7 Acknowledgments

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement 216338.

## References

- [1] DICONET: Dynamic Impairment Constraint network for transparent mesh optical NETworks <http://www.diconet.eu>.
- [2] Pachnicke, S., Reichert, J., Spalter, S., Voges, E.: Fast analytical assessment of the signal quality in transparent optical networks. *J. Lightwave Technol.* **24** (February 2006)
- [3] Zami, T., Morea, A., Leplingard, F., Brogard, N.: The relevant impact of the physical parameters uncertainties when dimensioning an optical core transparent network. In: *Proc. European Conference on Optical Communications (ECOC)*. (September 2008)
- [4] Zang, H., Jue, J., Mukherjee, B.: A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks. *Optical Networks Magazine* **1** (January 2000)
- [5] Azodolmolky, S., Klinkowski, M., Marin, E., Careglio, D., Solé Pareta, J., Tomkos, I.: A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks. *Elsevier Computer Networks* (in press).
- [6] Manousakis, K., Christodouloupoulos, K., Varvarigos, E.: Impairment-aware offline RWA for transparent optical networks. In: *Proc. IEEE Conference on Computer Communications (INFOCOM)*. (April 2009)
- [7] Azodolmolky, S., Pointurier, Y., Angelou, M., Solé Pareta, J., Tomkos, I.: An offline impairment aware RWA algorithm with dedicated path protection consideration. In: *Proc. IEEE/OSA Optical Fiber Communication conference (OFC)*. (March 2009)
- [8] Staessens, D., Colle, D., Pickavet, M., Demeester, P.: Path protection in WSXC switched networks. In: *Proc. European Conference on Optical Communications (ECOC)*. (September 2008)
- [9] Martinelli, G., Zanardi, A.: GMPLS signaling extensions for optical impairment aware lightpath setup (February 2008) IETF Internet draft.
- [10] Bernstein, G., Lee, Y., Li, D.: A framework for the control and measurement of wavelength switched optical networks (WSN) with impairments (October 2008) IETF Internet draft.

## From Software Services to a Future Internet of Services

**Marco Pistore and Paolo Traverso**

**Massimo Paolucci and Matthias Wagner**

Fondazione Bruno Kessler  
38100 Trento, Italy  
[pistore,traverso]@fbk.eu

DOCOMO EuroLabs  
80687 Munich, Germany  
[paolucci,wagner]@docomolab-euro.com

**Abstract.** Research in Service Oriented Computing has been based on the idea that software applications can be constructed by composing and configuring “software services”, i.e., software utilities that can be used but that are not necessarily owned by consumers. A key aspect has however been dramatically underestimated in this research, namely the fact that – in most cases – software services are software components that provide electronic access to “real services” (e.g., a software service for travel booking allows us to access the actual service behind it, namely “the possibility of traveling”). Our claim is that the “Internet of Services” should focus on real services, rather than software services. In particular, we investigate the new role of Internet, which is a supporting infrastructure in the case of software services, but becomes a key enabler for real services, offering a unique capability to communicate in real time changes in real services and allowing for immediate reactions by service consumers. In the paper, we illustrate the project we are undertaking to demonstrate that Internet can become the service delivery platform of the future. We illustrate, in particular, the research challenges this vision produces in the areas of service usage, representation, engineering, and delivery, as well as the results we have already achieved.

### 1. Introduction and Motivations

So far, research on Service Oriented Computing [1] has been based on the idea that software applications can be constructed by composing and configuring “**software services**”, i.e., software utilities that can be used but that are not necessarily owned by users [2]. This idea has been posing interesting novel challenges for research, in that **software services are no longer under control of developers**. New paradigms are being investigated for the description of services and for the negotiation of Service Level Agreements between providers and users; novel methodologies and tools are being defined for engineering service-based systems – i.e., for selecting, composing and configuring software services – as well as for monitoring, managing and supporting the adaptation of services.

In the research on Service Oriented Computing, a key aspect has however been dramatically underestimated, namely the fact that in most cases software services are software components that provide an electronic access to “**real services**”. That is, we use a software service such as a travel booking service since it allows us to access the actual service behind it, namely the possibility of traveling. **The characteristics of**

**real services are often very different from those of the corresponding software services.** For instance, the duration of a software service (i.e., the time for booking a travel) is limited with respect to the duration of the real service (i.e., the actual travel). Software services are static and accessible anywhere and anytime, while the actual services are dynamic and context dependent, since they happen in the real world. Software services are rather independent and the constraints and conflicts among them are limited, while the actual services the user expects are usually related, and hence heavily constrained (it is easy to book two conflicting travels, but it is not possible to exploit both of them).

As a consequence, the concepts used to describe services and the approaches for “composing and configuring” them are radically different from those proposed for software services. It is insufficient to provide a technical description of the functional and non-functional aspects of services (interfaces, behavior, quality, security, and so on). The representation of real services must be based on a set of **key assets** that the services represent for their consumers and providers [3]. So, for instance, in the organization of a travel, aspects related to the duration and to the cost of the trip are much more relevant than the travel booking software service, since time and money are key assets we are well aware of. Such key assets have also the capability of providing those links between services that are neglected by software services. Indeed, assets such as time, money, social relations are fundamental in the organization of our lives, and are transversal to all the services we may exploit for our work, social life, free time, and so on. A similar situation also occurs in a business context, where the characteristics of the actual services (delivery, production, stocking ...) are strictly related to the key assets of the company, and are hence much more relevant than the software interfaces that encapsulate them.

Our claim is that research on **Internet of Services should focus on real services** and on the key assets these services relate to. Indeed, Internet has a marginal role for software services – it is a convenient infrastructure for publishing, discovering, and executing software components. **Internet is instead a key enabler for “real” services**, in that it offers a unique capability to communicate to the user in real time the dynamicity of services and of their context (i.e., the cancellation of a flight can be immediately notified to the user), and to allow the user to react immediately to this dynamicity (i.e., re-scheduling the travel or looking for a train service). The effect will be similar to that in the traditional Internet, where the changes to the Web pages are seen immediately by everyone who connects to the Web: namely, the Internet of Services will allow users to live in and react to a world that changes at a higher and higher speed.

In this paper, we illustrate the project we are undertaking to demonstrate that Internet can become the service delivery platform of the future. The key ideas of the project, which we will describe in this paper, are: (1) to focus on real services instead of software services; (2) to describe these “real” services exploiting a small number of “core assets” that capture key concepts for service user and provider; and (3) to exploit the architecture of the Internet of Web pages we all know as a reference for the architecture of the future Internet of Services. The structure of the paper is as follows. In Section 2 we describe a motivating scenario based on [3]. In Section 3 we describe the project’s objectives and approach. Finally, Section 4 concludes the paper.

## 2. Motivating Scenario: User-Centered Services

Mobile phones are becoming an essential tool in our life. They not only act as phones and media players, but more fundamentally they give us access to variety of services that we need in everyday life. A simple catalogue of such services includes those for travel activities (e.g., navigation and map services, ticket booking via mobile, SMS notifications of flight delays), social networking, personal assistance, entertainment, and so on. These services can be used to help us in managing a wide range of situations both in our private life and related to our business activities.

For example, even now it is possible to use the mobile phone to get information about movies by interacting with active posters equipped with RFID or 2D bar codes [4], to access Web services for booking movie tickets, and to exploit telco services for payment. We can store a movie event in the calendar and set up a reminder for it. We can share the information about the movie with the people in our contact list using telco services. And, finally, we can use navigation services to route us to the cinema.

Another example of the capabilities of mobile phones is when we receive an e-mail with an invitation to business meeting. We can save also this event in our calendar. If the meeting takes place in a different city or country, we can access Web services for organizing the trip (plan the itinerary, book and pay the travel tickets, make a reservation in a hotel, and so on). We can use the mobile phone to monitor and detect problems with our travel (for instance, Lufthansa provides a service that sends SMS notifications when our flights are being delayed or cancelled) and react to them (e.g., by re-scheduling the flight, or by informing the other participants of the meeting that we will not be able to attend). Finally, if we successfully reach our destination, we can receive information on how to reach the venue of the meeting, e.g., using local public transportation.

These scenarios rely on a set of already available services and applications (e.g., calendar, communications, map and navigation, context tracking, payment) that refer to different domains (traveling, personal activity management, entertainment, and so on), and show the possibilities offered by their composition. Unfortunately, while the number of available services is rapidly growing, each service is narrowly directed to solve a specific need of the user, and no attention is paid to how services may work together or to how the user may utilize these services in combination. As a consequence, at the moment the user is alone when he faces the problem of their composition.

- First, the user has to deal with different services, and consequently, with their specific interfaces, formats, and protocols.
- Second, the user has to manage by hand the composition of these services and information flow across them, copying data among services and “orchestrating” by hand the execution of these services.
- Third, the user has to continuously ensure the consistency of the information used by different services: if the user is not able to go to the movie, he has to take care of propagating the effects of this decision by means of relevant services, i.e., removing the event from the agenda, sending an alert to the friends, canceling the ticket on-line, and so on.



- Finally, it is entirely up to the user to identify conflicts and overlaps among the different personal activities that the user is managing through the calendar. For instance, the cinema and business trip scenarios, which we have presented independently, may be strongly interconnected: if the cinema and the business meeting are scheduled for the same day, it might be impossible for us to participate to both of them, or we may need to organize the travel to the meeting taking into account the constraints due to the cinema.

To solve these problems, we need to be able to compose and integrate a wide set of heterogeneous services, allowing the user to be in control, to be alerted when conflicts and inconsistencies emerge, and to be able to decide among alternative solutions. At the same time, we need to abstract away the technical aspects of the service implementations and the differences between the protocols and data formats.

We remark that, in the scenarios just described, there are of course software services – for buying tickets, for planning trips, for booking flights, and so on – as well as other “technical” services available on the mobile phone – for sending SMS, and emails, for managing the calendar and the appointments, and so on. The focus is however not on these software/technical services, but on the “real” services; that is, the scenarios are about being able to watch a movie, to attend a business meeting, to travel to this meeting, as well as about managing the overlaps and relations between all these activities.

### 3. Objectives and Approach

In order to investigate the vision of Internet of Services described in the introduction, and to address the scenario described in Section 2, FBK has recently launched a research project, which is being executed in collaboration with DoCoMo EuroLabs and other key partners in academics and industry. The ultimate goal is to demonstrate that Internet can become the (reference) service<sup>1</sup> delivery platform, and to realize all the necessary building blocks necessary for this demonstration. More concretely, the project pursues three kinds of objectives: (1) Theories, methodologies, and techniques that support different key aspects of Internet of Services. (2) A prototype platform that implements and integrates these theories, methodologies, and tools. (3) An experimentation of these results on key application domains.

The development of novel **theories, methodologies, and techniques** will be based on two main concepts, which we will better describe in the next sub-sections. First, the development of the Internet of Services will be based on a novel model of services which is based on “key assets”. Second, the Internet of Services that we will build will have a structure resembling that of the Internet of Web pages.

#### 3.1. Asset-Based Service Composition

Service composition in the Internet of Services requires managing a wide variety of services belonging to different domains – in Section 2 we have given some evidences

<sup>1</sup> From now on, by services we mean “real” services as described in the introduction.

of this. It is hence necessary to find the right concepts that allow for (1) describing all these services despite their high heterogeneity, and (2) allowing the user to be in control of the usage and composition of these services.

Addressing these two challenges at the same time is far from trivial. For instance existing Semantic Web Service approaches such as OWL-S [5] or WSMO [6] allow for managing a wide heterogeneity of services, but these techniques have been designed to provide a complete description of the services, of all their possible usages and of all their effects; as a consequence, the description of services they provide is too technical for allowing the user to keep control of the composition – higher level “semantic” concept are needed. Techniques like mash-ups [7] or Yahoo pipes [8] are nearer to what we need, but also in this case they are based on concepts such as data- and control-flows, which are still too technical and, once again, more adequate for software services than for real services.

Our claim is that, in order to be able to manage real services and to deal with the complexity described in Section 2, the concepts to describe the services should not relate to technical properties (inputs, outputs, data-flow, control-flow, and so on), but rather on the user assets that are affected by the services. In [3] we proposed the usage of a small set of core “**user assets**” for describing and organizing services. In particular we identified four resources:

- **Time**, representing the temporal relation of user’s activities, as well as the overlaps and conflicts among these activities;
- **Location**, representing the (current and perspective) location of the user, the availability of services in these locations, as well as the necessity of traveling or moving to exploit services;
- **Social relations**, representing other parties (family, friends, colleagues...) involved in the user activities;
- **Money** and other values, representing costs and assets involved in the user activities.

These assets are at the same time *enablers* for the usage of services and *constraints* for their exploitation. Indeed, being in a given location allows us to access services (e.g., entertainment services) that are available at that location, but it also forces us to organize travels in order to reach (services available only at) different locations. Our social network constraints our activities (e.g., obligations with our family reduce our freedom in undertaking other activities), but also provides help in case of problems (e.g., family or friends can help with activities we would not be able to undertake alone).

The mobile phone already provides well know applications for managing these resources, namely calendars, maps, contact lists, e-wallets, and so on. These applications, which are rather intuitive and easy to use also for a non-technical user, can be used to expose services to the user, and to let the user control the composition.

Consider for instance, the scenarios described in Section 2. Once the movie event has been stored in the agenda, we can use the calendar as the entry point for accessing additional actions and services that we can perform with the movie (see Fig. 1(a)), such as buying the ticket, looking to the trailer, plan the trip to the cinema, share the event, e.g., with our wife. When the beginning of the movie approaches (see Fig. 1(b)), the calendar in the mobile can alert us, navigate us to the cinema, or allow us to call a taxi. Similarly, when we add the business meeting to the calendar, the trip

to the location of the meeting can be planned and booked through additional services accessible from the calendar. Moreover, the system can detect that there is an overlap between the travel and the movie (see Fig. 1(c)) and signal this problem to the user. If the user resolves this problem by moving the movie to a new date (see Fig. 1(d)), the system can take care of updating the reservation, and, in case this event is shared with our wife or with some friends, of notifying them the update, e.g., via SMS.



**Figure 1. Service composition using the calendar.**

This example shows how the time resource can be used to link and compose highly heterogeneous services such as those necessary for buying movie tickets, scheduling meetings, planning trips; and how the calendar can be used as a front end for controlling and orchestrating this composition. Similarly, we could use the other core assets – locations, social relations, money/values – for defining and identifying the corresponding front-end tools – maps, contact lists, e-wallets – for managing these compositions.

We believe, however that the applicability of core assets as driving concepts for service description and composition in the Internet of Services is of broader applicability than the user-centered services available through a mobile phone. These concepts can indeed be applied also in business scenarios, even if in these cases the “assets” will be different, and specific to the specific business domain. So, for

instance, in case of emergency management, a map of the territory provides the main view for planning and running emergency services – environmental sensors, reports on emergency events, location of emergency units, movement plans, and so on can all be represented, supervised, and “orchestrated” on this map. In the case of a financial company, the key assets and the main view to the business will probably be related to balances and charts of accounts; in case of manufactory, the production chain may offer the main view to the core assets. And so on.

### 3.2. Conceptual Architecture for the Internet of Services

In order to understand how to build the future Internet of Services, we will use as reference the current Internet of Web pages. More precisely, we will concentrate on four aspects that, in our opinion, contributed in a fundamental way to the success of the Internet we know, and we will investigate how these four aspects can be replicated in the Internet of Services. As shown in Figure 2, these four key aspects of the Internet of Web pages are:

- a powerful, easy-to-use **browser**;
- flexible tools for page editing and **content creation**;
- infrastructure tools (such as **Google**) for easing the access to the pages; and
- tools for modeling the content and knowledge in the Web pages (the **semantic web**).

In the following, we describe the corresponding aspects we foresee in the Internet of Services.

- **Asset-Driven Service Modeling.** In our vision, asset-driven service modeling will play the same role that semantic web plays in the classical Internet of Web pages [9]. Indeed, as shown in Subsection 3.1, assets capture the key aspects of services for provider and for consumer. Moreover, in our approach, assets are the glue among the different components on the Internet of Services. For this reason, a coherent and flexible modeling of assets is a fundamental element in the project. Novel methodologies and tools are needed to support the modeling of the key assets of services; it is also necessary to support the modeling and understanding of the processes and knowledge that these services need, and to associate these processes and knowledge to the service assets [10]. The methodologies and tools that we intend to adopt will allow the different actors involved in the provision and exploitation of services (e.g., service providers, end users, domain experts, knowledge engineers) to collaboratively work towards the modeling of asset-driven services. This, in spite of the different knowledge engineering and modeling skills of the different actors [11].
- **Service Delivery Infrastructure.** By “service delivery infrastructure” we mean all the mechanisms and tools that are between the service provider and the service consumer. In the classical Internet, this infrastructure includes for instance tools like Google, MySpace, YouTube, which facilitate finding contents relevant for us, or making our content easier to find. In the Internet of Services, the infrastructure tools need to address different purposes, the most important being **service composition**. In the case of software services, service composition is a design task performed by the service engineer (see, for instance, the approach described

in [12]). In the case of real services, the composition is dynamic, context-aware, user-centric, and asset-based; as a consequence, it stops being a design task, and becomes the most relevant functionality the delivery infrastructure should provide. New methodologies, techniques and tools are needed for this novel service composition which lives in the delivery infrastructure, and new reasoning services are needed which are able to take into account the heterogeneous knowledge used by service. In addition to service composition, the delivery infrastructure should support all the other operations necessary to use services. Namely, it should provide mechanisms and tools for the **enactment, monitoring, adaptation, management** of the delivered services.

- **Service Front End and User Interaction.** The goal is to design the equivalent of the browser for Internet of Services, that is, to develop a tool that allows the end user to access and exploit services. This requires the development of **intelligent interfaces** which support novel interaction paradigms, in particular in the mobile and situated settings; these interfaces will use the concept of service assets to drive the interaction of the user with the services. However, the effectiveness of the user interface depends not only on the quality of the technologies. The **user acceptance** of these technologies, in terms of usability, appropriation, personal and social impact, is as important, specifically if the goal is to deliver “real” services. We will exploit a value-based evaluation process [13] to assess user acceptance; in particular, we will assess the importance of using assets for modeling services for user acceptance.

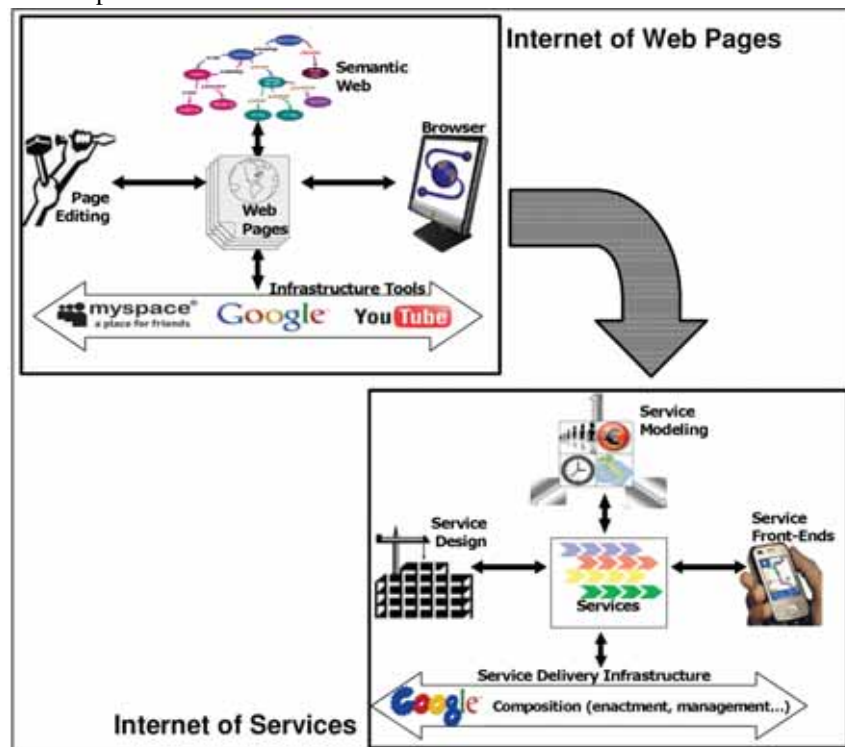


Figure 2. From the Internet of Web Pages to the Internet of Services.

- **Service Design and Engineering.** In the classical Internet, producing pages requires tools that are much more sophisticated than HTML editors: content managers, dynamic pages, Web application, Wikis, have been developed over time. Similarly, for the Internet of Services, there is a need for novel methodologies and techniques for designing and for managing the life cycle of the “real” services, as well as of the software components that encapsulate them. In particular, in order to be easy to monitor and to adapt to the specific user requirements and usage contexts, novel design and engineering principles and methodologies are needed, which should rely upon a deep understanding of user assets and requirements [14].

### 3.3. Integration and Evaluation

In addition to investigating these research challenges, the project will develop an **integrated platform** for the Internet of Services. This platform addresses two different requirements, namely to integrate the results of the different research lines just described in a coherent system, and to allow for a practical demonstration of the feasibility of the project vision.

Finally, the results of the project will be evaluated by applying the platform in two pilots based on different **application scenarios**. The first scenario is in the **user centered** domain, which concerns all applications where the services are used and combined according to the personal needs of (a certain class of) end-users. In this case, the person and her/his life is the center, the key assets are personal assets (like her/his personal agenda or social relations), and the services are organized and composed based on problems and opportunities related to such assets. The second scenario is in the **business centered** domain, where the focus is not the single user, but the definition of a complex business ecosystem based on the exchange of services among different types of actors. With respect to the previous one, in this scenario we not only have heterogeneous services, but we also have different actors with different key assets in which real services must be represented.

## 4. Concluding Remarks

In this paper, we introduced a project that we launched recently on the Internet of Services. This project relies on the key idea that, while most of the research in Service Oriented Computing has concentrated on “software” services, the future Internet of Services should focus instead on the “real” services that are behind this software. This shift requires re-thinking the current approaches for service representation, service engineering, service delivery, and service usage. In order to make it possible to address the new challenges deriving from this shift, in the project we adopt two hypotheses. First, we assume that a small set of core assets will be able to describe the key features of these “real” services, as well as to drive the service composition process, and to “orchestrate” the execution of these services. This hypothesis is based on previous results [3] obtained in the case of user-centric services accessible through

the mobile phone, which show that “time” can be used as key asset in this context, and that the calendar can be used by the user to control the service composition. The second hypothesis is that the conceptual architecture of the Internet of Services will be similar to that of the Internet of Web pages, and hence we can learn from the key achievements in the latter in order to direct the research in the former.

The project is in its early stages, and most of the objectives and challenges described in Section 3 will be addressed in our future work during the five years of the project. During all our work, the validity of our hypothesis and the correctness and effectiveness of the results will be checked and challenged through the two application scenarios described in Section 3, but also through a continuous interaction and collaboration with the broader community of researchers and practitioners interested in shaping the future Internet of Services.

## References

1. Papazoglou, M.; Traverso, P.; Dustdar, S.; Leymann, F.: *Service-Oriented Computing Research Roadmap*. 2006.
2. NESSI Strategic Research Agenda. Vol. 1. *Framing the Future of the Service Oriented Economy*. 2006.
3. Kazhamiakin, R.; Bertoli, P.; Paolucci, M.; Pistore, M.; Wagner, M.: *Having Services "YourWay!": Towards User-Centric Composition of Mobile Services*. FIS 2008.
4. Paolucci, M.; Broll, G.; Hamard, J.; Rukzio, E.; Wagner, M.; Schmidt, A.: *Bringing Semantic Services to Real-World Objects*. Int. Journal on Semantic Web and Information Systems 4, 2008.
5. *OWL-S Home Page*, <http://www.daml.org/services/owl-s/2003/>.
6. *Web Service Modeling Ontology (WSMO)*, <http://www.wsmo.org/TR/d2/v1.2/>.
7. Blake, B.; Nowlan, M.: *Predicting Service Mashup Candidates Using Enhanced Syntactical Message Management*. SCC 2008.
8. *Yahoo Pipes*, <http://pipes.yahoo.com/>.
9. *The NEON project*. <http://www.neon-project.org/>.
10. Di Francescomarino, C.; Ghidini, C.; Rospoche, M.; Serafini, L.; Tonella, P.: *Reasoning on semantically annotated processes*. ICSOC 2008.
11. Ghidini, C.; Rospoche, M.; Serafini, L.; Kump, B.; Pammer, V.; Faatz, A.; Zinnen, A.; Guss, J.; Lindstaedt, S.: *Collaborative Knowledge Engineering via Semantic MediaWiki*. I-SEMANTICS 2008.
12. Marconi, A.; Pistore, M.; Traverso, P.: *Automated Composition of Web Services: the ASTRO Approach*. IEEE Data Eng. Bull. 31(3), 2008.
13. Not, E.; Leonardi, C.; Mennecozzi, C.; Pianesi, F.; Zancanaro, M.: *Beyond Usability: a New Frontier for User-Centered Design of "Future Internet" Services*. FIS 2008.
14. Penserini, L.; Perini, A.; Susi, A.; Mylopoulos, J.: *High Variability Design for Software Agents: Extending Tropos*. ACM Transactions on Autonomous and Adaptive Systems 2(4), 2007.



## Multi-level SLAs for Harmonized Management in the Future Internet<sup>★</sup>

Wolfgang Theilmann<sup>1</sup> and Luciano Baresi<sup>2</sup> <sup>★★</sup>

<sup>1</sup> SAP Research – CEC Karlsruhe

Vincenz Priessnitz Strasse 1, 76131 Karlsruhe, Germany

wolfgang.theilmann@sap.com

<sup>2</sup> Dipartimento di Elettronica e Informazione Politecnico di Milano

P.zza L. Da Vinci, 32 - 20133 Milano, Italy

baresil@elet.polimi.it

**Abstract.** The Future Internet is about to fundamentally change social and economic interactions at a global scale. The integrated access to people, media, services, and things will enable new styles of interaction at unprecedented scale, flexibility, and quality. However, this also calls for a well-defined and sound approach for management and governance that allows for clear harmonization and translation of issues across domains and layers. This paper presents a proposal that aims to blend management and governance issues at business, software, infrastructure, and network level, and introduces a multi-level SLA management approach to bridge these issues across different layers. It also sketches some insights on management and governance practise and requirements in various industrial domains.

### 1 Introduction

The integrated access to people, media, services, and things, provided by the *Future Internet*, will enable new styles of societal and economic interactions at unprecedented scale, flexibility, and quality. The Future Internet, through the metaphor of *Internet of Things*, will provide location independent, interoperable, scalable, secure, and efficient access to a coordinated set of services [3], but such broad vision demands for a sound and well-defined approach for management and governance. This approach has to harmonize and bridge the various views

---

<sup>★</sup> The research leading to these results is partially supported by the European Community's Seventh Framework Programme ([FP7/2001-2013]) under grant agreement n. 216556.

<sup>★★</sup> On behalf of the SLA@SOI consortium [7] which includes T. Ellahi, H. Li, W. Theilmann (SAP), F. Torelli (Engineering), J. Kennedy (Intel), M. Alvarez, A. Castro, J. Lambea, S. Aleman (Telefonica Investigacin y Desarrollo), C. Kotsokalis (University of Dortmund), M. Trifu, C. Momm (Research Centre Karlsruhe), A. Marconi, M. Pistore (Fondazione Bruno Kessler), L. Baresi, E. Di Nitto, S. Guinea (Politecnico Milano), G. Spanoudakis (CITY University), R. Perrot, T. Harmer (Queens University Belfast), G. Pipan (XLAB), G. Armellin (GPI) and M. Evenson (eTel).

and layers of the Future Internet following the subsidiary principle: as many issues as possible should be dealt with locally, while as few issues as possible are to be managed in a more integrated way.

This multi-layered view requires that supplied services<sup>3</sup> be managed coherently. Service Level Agreements (SLAs) and policies are becoming common means for doing this. SLAs specify the conditions under which services are provisioned, but current management frameworks typically only focus on single service interfaces. They neither use SLAs for managing the implementation and delivery of services, nor they recognize/support the fact that many services may be composed of lower-level services, involve third-party providers, and rely on a possibly complex business/IT stack [2]. While SLAs are routed in the respective customer requirements, policies are provider-specific means to express constraints and rules for their internal operations. These rules may be independent of any particular customer.

The paper presents the proposal of the EC project SLA@SOI to integrate and blend management and governance at the different levels. The core idea is to use SLAs as primary means to express management and governance concerns and to share and translate these concerns across different views, perspectives, and layers of the Future Internet.

The main innovation of the framework will be the integration of the following features: (1) standardized models for SLA descriptions at the different layers, (2) an automated e-contracting framework, (3) systematic grounding of SLAs from the business level down to the infrastructure, (4) methods and tools for multi-layer SLA management, including planning, optimization, and provisioning, (5) methods and tools for monitoring and accounting services and SLAs through standardized interfaces, (6) exploitation of virtualization technologies at infrastructure level for SLA enforcement.

The remainder of this paper is organized as follows. Sections 2-5 introduce the most important management and governance issues at business, software, infrastructure, and network level, respectively. Section 6 provides details on how the multi-level SLA management can contribute to a harmonized and holistic approach across these layers. Section 7 provides insights on the practise and requirements of different industrial domains, and Section 8 concludes the paper.

## 2 Business Level

Nowadays, service providers (like Google, Amazon, and Facebook) must mediate between the business challenges enabled by network and IT convergence and users demanding for more and more new value-added services [8].

The ideal solution would be the emergence of a new service *marketplace*, formed by the convergence of Internet, media, and telecommunication industries. This new marketplace should support the exposition of multiple heterogeneous services, from different providers and industries, control and guarantee their quality, and maximize customers and revenues. New services could be

<sup>3</sup> In this paper, we consider services as means to deliver value to customers.

added by composing those supplied by different providers, and the marketplace should also help automatically establish and manage business agreements and relationships among the participants, and expand marketing possibilities beyond advertisement-based business models (e.g., subscriptions or pay-per-use models). The end-to-end manageability of the service lifecycle—from creation to monetization—becomes a key issue, and must work seamlessly also for those services that span different providers.

The use of SLAs is indispensable for this new marketplace to emerge. An SLA may specify for instance the levels of availability, serviceability, performance, operation, or other attributes (e.g., billing) and even the penalties when the agreement is violated [10]. At this level, it is worth distinguishing between: *supplier/partner SLAs*, which include the conditions under which a third-party service can be contracted and reused by other service providers, and *customer SLAs*, which include the conditions under which the service can be contracted by a customer.

Service conditions are formalized using SLA templates (e.g., based on WS-Agreement [4]) and the production of these templates must become an intrinsic part of service development. SLA templates contain default values specified during the service development phase. Once a service provider or customer enters into negotiations to contract a service, the SLA template is used as a blueprint for the SLA [9]. If a service depends on other services, supplied by different providers, customer SLA templates must take into account these dependencies and ensure the consistency of SLA parameters and values.

### 3 Software Level

Management and governance at this layer primarily consist of two aspects: a centralized or federated configuration management database (CMDB), populated with the live state of the service and software landscapes, and a set of processes carrying out the activities of the service lifecycle. These processes are executed in accordance with the information available in the CMDB. Additionally, these processes produce further information being pushed to the CMDB for satisfying the requirement of predictive autonomic management.

The CMDB holds information, SLAs, rules, policies, and various models to capture diverse aspects of the software and service landscapes. More in detail, it comprises:

- The **software landscape**, which contains software component models based on standards like SCA (Service Component Architecture). These artifacts are extended with non-functional properties of the service/software components.
- The **software configuration models**, which contain information about the deployment-time configurations of the various software components.
- The **service landscape**, which offers a comprehensive information model containing detailed descriptions of the various elements: services, SLAs associated with them, their software components, and the required execution

infrastructure (e.g., middleware, application servers, and databases). This landscape also considers the relationships between these elements and those of the business and infrastructure layers.

- The **operational rules and policies**, which are consulted by the processes during the execution and operation phase of the service lifecycle. These rules and policies ensure that the service landscape be in a consistent state and adheres to set business policies.

The processes are carried out in various phases of the lifecycle in accordance with the information fetched from the CMDB. These processes can operate only at this level or work together with the other levels to ensure the fulfillment of multi-layer SLAs. These processes include:

- **Deployment/Provisioning** activities to identify the required set of software components essential for delivering the services. Moreover, some capacity planning and sizing procedures help identify the required setting of logical resources (lower levels) to ensure the service levels as constrained by the SLAs. These activities rely on the SoA models defined in the software landscape.
- **Configuration** activities to focus on the identification and setting of “knobs and switches” used during the later phases of the lifecycle for SLA management. These activities are performed by utilizing the configuration models from the software landscape and information derived from higher-level SLAs.
- **Monitoring** activities to help configure knobs and switches. These are the most critical processes for management and governance. Monitored information is stored onto the CMDB for historical purposes and for the sake of autonomic predictive management.
- **Adaptation/Change Management** activities to analyze the information gathered from monitoring and take appropriate actions if the current configuration could lead to possible SLA violations. Effective governance requires proper change management procedures; the service landscape models can facilitate change impact analysis to understand the implications before performing the actual changes.

## 4 Infrastructure Level

This level allows the provider to advertise and manage the infrastructure, reserve and provision it as per negotiated SLAs, and stage the software images. Management at this level centers on the infrastructure landscape: what virtual machines (VMs) are provisioned where, and what commitments and reservations have been made into the future. Governance is concerned with higher-level business rules defined by the infrastructure provider. These policies can span areas such as security, energy-consumption, data privacy and isolation, business continuity, and geographical and legal constraints. Both management and governance are being enabled by a policy-driven architecture to allow the overall system to

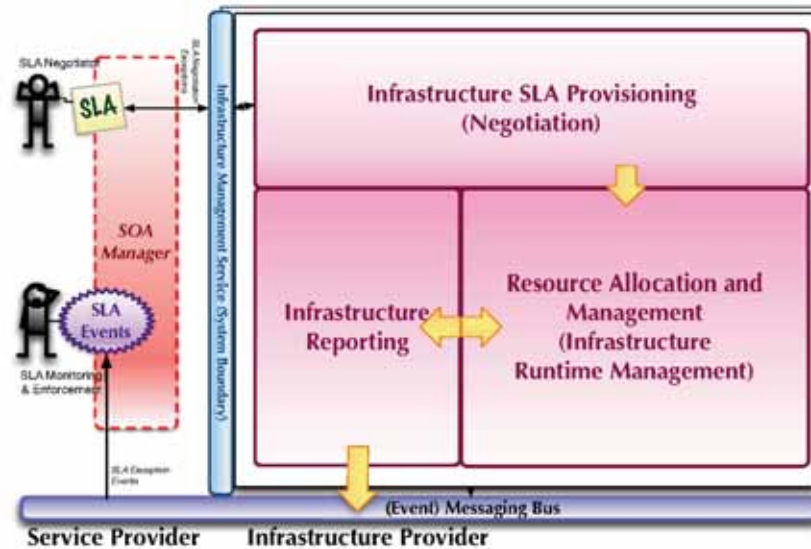


Fig. 1. Infrastructure Management Architecture.

be kept under control despite the potentially diverse needs of the individual SLAs.

Figure 1 provides a high-level view of this architecture. **Infrastructure SLA provisioning** includes the functionality that is required prior to actual resource allocation. It translates requests from potentially high-level infrastructure requirements originating from the software-level to technical requirements relevant to the infrastructure level. It also maps these to appropriate bundles of resources, and verifies that the infrastructure is available to accommodate the actual requests. Policies at this stage could mandate provisioning VMs on physical machines dedicated only to that customer, or mandate particular failover resources being reserved in parallel with the main request.

**Resource allocation and management** performs the actual provisioning. Hardware, either physical or virtual, is provisioned by either internal or external resource providers. This functional component includes an *autonomic management* capability to identify if and when re-provisioning should occur to properly accommodate defined policies, or SLA conditions that are in danger of being violated. Once the infrastructure is provisioned, appropriate event logging, correlation, and monitoring is instantiated in the **Infrastructure reporting** component to allow all relevant conditions or events to be identified and escalated, either internally for immediate disposition, or externally in case of, for example, the unavoidable violation of an SLA. Again, policies are used extensively to define what information needs to be escalated internally and to where, notwithstanding the additional external commitments agreed in the SLAs.

## 5 Network Level

Nowadays, the underlying network must allow users to stay connected permanently from different environments (e.g., when at home, or at work) with a wide range of devices. This scenario must be paired with the growth of the telecommunication sector, with the evolution of both mobile communications (e.g., UMTS and HSDPA) and fixed broadband access (e.g., ADSL and FTTH). Future network infrastructures will support the convergence and interoperability of both network technologies. The ultimate goal is to optimize the transmission of voice, data, and media to and among users, no matter of their locations or devices. In the more immediate future, this convergence means that a single service can be executed through and switched between different wired and wireless networks.

Historically, the services provided to users have been tightly coupled with specific technologies and networks, but this solution is not suitable to cover the needs of the above scenario. Services and networks must be decoupled, but we must prepare the network to cover this need. Network and service providers are studying how to evolve their infrastructures to new architectures like NGN (New Generation Network). The new concept of network management must move from vertically integrated services (silos-oriented service management) to horizontally integrated services, with many common capabilities available at the same time.

Once networks and services are decoupled, we can start organizing a service-based business environment. This new multi-domain/multi-provider environment is challenging to network management systems, which now should provide a complete supervision of the services through different domains and different networks. In addition, it is mandatory not to forget customers and their perception of offered services. In this scenario SLAs play a key role. The agreements between customers and service providers must specify the conditions under which the service is provided and consumed. These parameters must be built up from the parameters of the individual networks or domains. The agreement becomes the basis for monitoring the network through the various layers.

## 6 Multi-level SLAs for Harmonized Management

Currently the integration of the different management activities described so far is mostly based on human activities. For instance, experts are usually able to recognize if a degradation in performance is due to a wrong software configuration or to insufficient hardware resources. In the former case, they can solve the problem by using their expertise, while in the latter case, they raise the issue to the infrastructure manager to solve it. If the problem is not solved directly by the infrastructure manager, the issue is raised to the business level to decide, for example, the activities and systems with higher priority when assigning resources.

As the complexity of applications grows, this human-based process becomes more and more challenging. Difficulties increase when applications cross the boundaries of a single organization, and when independent human-based and

automated services are integrated over the Internet [1]. Approaches and tools for automating the integration of management activities become mandatory.

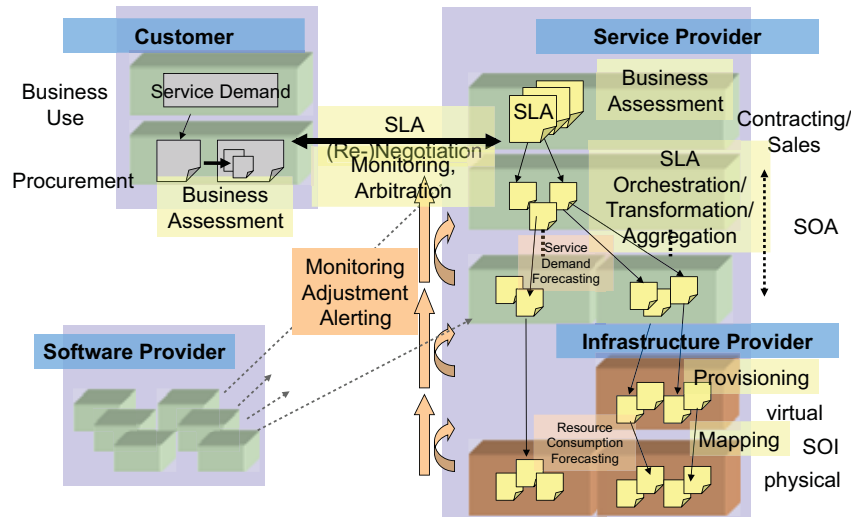
The proposal presented in this paper is based on a conceptual and a software framework for the harmonized management of all application levels based on the unifying concept of Service Level Agreement (SLA). Several SLA management frameworks have been presented in the past, but they typically focus on the specification of processes and models, giving no support to automation [6]. Moreover, they focus on the level of single service interfaces and do not recognize/support the fact that many applications are composed of a hierarchy of services at different levels [2]. Composition of services is one of the basic assumptions of the SoA approach [5], nevertheless current relevant methodologies only pay attention to the composition of software services without considering neither their business-level representations nor the infrastructure on which these services are provisioned.

Our proposal applies the concept of SLA in a uniform way, not only to software level services, but to all levels of the application stack. Usually SLAs are the result of a negotiation between a customer and the service provider. The proposal regards each application level as a service provider for the higher levels and as service consumer for the lower ones. Each provider tries to fulfill its policy rules to satisfy internal requirements, while negotiating with the potential consumers to satisfy external requirements.

SLAs and negotiation work for both human and automated services, for internal and external use, and at all levels to improve the management of IT applications. The possibility that applications be distributed among specialized providers—in a more reliable and collaborative way—will encourage the use of the SoA model at global level, allowing for a more efficient use of all kinds of available resources.

Figure 2 gives a simplified overview of the SLA management process. As today's business systems typically consist of complex layered systems, user-level SLAs cannot be directly mapped onto the physical infrastructure. Services might be composed of other more fundamental services that could be even provided by external parties. Consequently, a stepwise mapping of higher-level SLA requirements onto the lower levels, and the aggregation of lower-level capabilities into higher-level aggregations, is crucial for grounding user-level SLAs onto the infrastructure. This vertical information flow must carefully reflect service interdependencies as well as the originating business context. In addition to SLAs, the vertical information flow also covers monitoring, tracking, and accounting data and must support brokering and negotiation processes at each layer. As shown in the figure, the overall SLA management process may include different stakeholders, namely customers, service, and infrastructure providers, and also various business steps such as business assessment, contracting, and sale. The overview is intentionally simplified in the sense that no service chains are visualized. Such chains would represent the cases where service providers rely on additional external providers.





**Fig. 2.** Envisaged interactions of SLA stakeholders.

A service provider offers services with differentiated, dependable, and adjustable SLAs and can negotiate concrete SLAs with (individual or groups of) customers in an automated fashion. This business goal imposes requirements on software providers (to provide components with predictable non-functional behaviour), on infrastructure providers (to support SLA-aware management of resources), and also on the service provider (to translate and manage SLAs from business level along the IT stack down to the network). Of course, complete business value chains can be easily composed on top of this setup.

## 7 Industrial Practise and Requirements

This section briefly presents the industrial practice in the financial, e-Government, and ERP domains, which are used as case studies in the project.

A financial service provider must comply with management restrictions. The general financial service regulator typically stipulates conditions in which sharing of data may take place on an individual server or server site. For example, trading and customer accounts data cannot be held or processed on the same server. The regulator may specify that certain portfolio data may only be held and/or processed within county boundaries, geographic regions or ranges, or in fact any combination of these. For example, US government pension data may only be held within the US and within 50 miles of New York.

Individual financial service customers may have their own management restrictions as far as precise security infrastructures are concerned depending on the nature of their business and the service they are availing of. For example, a back test prediction service may not require security encryption on data trans-

mission of historical data, but may require full authentication/authorization access controls to profit prediction services and security encryption of transmitted results. Financial customers may also require that their data are held and processed on an exclusive server or even site to ensure that data sharing does not happen with other customers.

The design and management of e-Government services for citizens, such as the provision of medical and social assistance to elderly or disabled people, is a complex process that involves several actors with different roles and expectations. The citizen, who is the service beneficiary, expects efficient health and social services for addressing his/her own specific need. The governance, which is the service customer, has the necessity of monitoring and analysing the costs and quality of the health and social system as a whole. The different organizations that contribute as services providers (e.g., hospitals, nursing homes, private companies and charities) have their specific business goals to fulfill and need to negotiate with the governance the quantity and quality levels of provided services (e.g., number of medical treatments offered per day by an hospital).

The fundamental role of SLAs in e-Government service provisioning has been already recognized in the context of the so called G2G (government-to-government) services. Their adoption is, in contrast, still very limited in scenarios such as the medical and social care, which also require G2B (government-to-business) and G2C (government-to-citizen) service provisioning. The particular challenge in this context is that SLAs are not only based on market rules, but they are most often driven by “social” agreements between public bodies and citizens. As a consequence, the SLA negotiation (both between public bodies and citizens and between public bodies and private service providers) is different than in market-oriented domains. Another challenge of this domain is that it requires an integration of human-based services (e.g., home care, medical assistance at home, and transport services) with IT services; the underlying service oriented infrastructure is hence not only a technological infrastructure, but also a social and organizational one.

ERP services typically constitute the core of the IT architecture. They tend to be long-lasting and typically require lengthy implementation phases. However, also ERP services are more-and-more offered in an on-demand fashion, both to allow customers to have a dynamic consumption with low capital expenses but also to support the flexible construction of more complex business value chains.

Management and governance in these environments lead to significant challenges in terms of interoperability of solutions, but also in a proper negotiation and resolution of possibly conflicting policies. Furthermore, policies often do not just apply to a single service but to a complete system of services that somehow constitute a separate business context. Current practise in this area is still rather primitive. SLAs are partially formally expressed at the business level, however never translated or correlated with other SLAs on service/system layers. Policy support is mainly realized for local perspectives (e.g., for the application administrator of a set of software services) with poor interlinkage to other roles, views or organizations.

## 8 Conclusions and Future Work

This paper discusses the proposal of the EC project SLA@SOI [7] to integrate and blend management and governance at the different levels of the IT stack. The core idea of this paper is to use SLAs as primary means for expressing management and governance concerns and to share and translate these concerns across the different views, perspectives and layers of the Future Internet. The main goal of the project is to provide an SLA management framework that allows for consistent specification and management of SLAs in a multi level environment. The framework is designed for integration into different service-oriented infrastructures and will be evaluated within various complementary industrial case studies.

## References

1. Chesbrough H. and Spohrer J.. A Research Manifesto for Services Science. Communications of the ACM, 2006. 49(7): p. 35-40.
2. CoreGRID. Using SLA for Resource Management and Scheduling - A Survey, TR 0096. August 2007, [www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0096.pdf](http://www.coregrid.net/mambo/images/stories/TechnicalReports/tr-0096.pdf).
3. Future Internet Assembly. Bled Declaration on Future Internet. April 2008, [www.future-internet.eu/index.php?id=47](http://www.future-internet.eu/index.php?id=47).
4. Open Grid Forum. Web Services Agreement Specification, March 2007, [www.ogf.org/documents/GFD.107.pdf](http://www.ogf.org/documents/GFD.107.pdf).
5. Papazoglou M.P. and Van Den Heuvel W.J. Web Services Management: A Survey. IEEE Internet Computing 9 (2005) 58-64.
6. Schaaf T. Frameworks for Business-driven Service Level Management: A Criteria-based Comparison of ITIL and NGOSS. Proceedings of BDIM 2007, 65-74, Munich, Germany.
7. SLA@SOI consortium. SLA@SOI project (IST- 216556; Empowering the Service Economy with SLA-aware Infrastructures). [www.sla-at-soi.eu](http://www.sla-at-soi.eu)
8. tmforum. Service Delivery Framework Overview (TR139 Version 2.0, July 2008)
9. tmforum. SLA Handbook Solution Suite v2.5 [www.tmforum.org/page30754.aspx](http://www.tmforum.org/page30754.aspx)
10. Wikipedia [en.wikipedia.org/wiki/Service\\_level\\_agreement](http://en.wikipedia.org/wiki/Service_level_agreement)

## The Service Web: a Web of Billions of Services

John Domingue<sup>1</sup>, Dieter Fensel<sup>2</sup>, John Davies<sup>3</sup>, Rafael González-Cabero<sup>4</sup>, Carlos Pedrinaci<sup>1</sup>

<sup>1</sup>Knowledge Media Institute, The Open University, Milton Keynes, UK

[j.b.domingue@open.ac.uk](mailto:j.b.domingue@open.ac.uk), [c.pedrinaci@open.ac.uk](mailto:c.pedrinaci@open.ac.uk)

<sup>2</sup>University of Innsbruck, Technikerstraße 21a, Innsbruck, Austria

[dieter.fensel@sti2.at](mailto:dieter.fensel@sti2.at)

<sup>3</sup>Next Generation Web Research, IT Futures Research Centre, BT, Ipswich, UK

[john.nj.davies@bt.com](mailto:john.nj.davies@bt.com)

<sup>4</sup>Atos Research & Innovation, Atos Origin, Madrid, Spain.

[rafael.gonzalez@atosresearch.eu](mailto:rafael.gonzalez@atosresearch.eu)

**Abstract.** SOA4All, a collaborative European research and development project, is pioneering advanced web technology that will allow billions of parties to expose and consume IT services online. Four complementary technical advances are being integrated to create a coherent and domain-independent service delivery platform. Service-oriented architectures and service-orientation principles are being used to support the development of complex services based on distributed and reusable components. Web principles and technology are used to provide an underlying infrastructure that allows the integration of services at a world wide scale. Web 2.0 is used to structure human-machine cooperation in an efficient, user-adapted and cost effective manner. And semantic technology is used to enhance service discovery, composition and execution.

**Keywords:** Service Web, Semantics, SOA, Web 2.0, WSMO

## 1 Introduction

The software industry has been experiencing a dramatic shift from products to services [1]. From a technical perspective, service-orientation is increasingly being adopted to allow applications to be created flexibly by linking loosely-coupled components, typically over a network. From a business viewpoint, fierce competition between vendors is forcing significant reductions in the price of software products. Companies are being compelled to embrace services to add value to product offerings and establish new revenue streams. Furthermore, the appearance of the internet as a disruptive platform has given rise to new business models such as Software as a Service (SaaS) and advertising-based revenue models (e.g., Google). In the light of these trends, business experts are suggesting that consideration should be given to ‘servicising’ products to provide added value, and ‘productising’ services so that they can be delivered more efficiently and at lower costs [1].

Large organisations such as Verizon already have systems based on approximately 1,500 web services [2]. However, the web contains only around 27,000 web services based on the Web Services Definition Language (WSDL) [3]. In consequence, service-oriented architecture (SOA) is largely still an enterprise-specific solution exploited by, and located within, large corporations. Because it minimises costs and maximises the potential market, the fully-automated delivery of services over the web appears to be a potential ‘silver bullet’ for delivering IT services.

We envisage that the combination of semantic technologies and SOA will lead to the creation of a ‘service web’—a web that allows billions of parties to expose and consume billions of services seamlessly and transparently, where all types of stakeholders – from large enterprises to SMEs and singleton end users – engage as peers, consuming and providing services within a network of equals. However, as was highlighted in [4], SOA will not scale without:

- properly incorporating principles that made the web scale to a worldwide communication infrastructure;
- significant mechanisation of service lifecycle activities (which includes location, negotiation, adaptation, composition, invocation and monitoring as well as service interaction requiring data, protocol and process mediation); and
- a balanced integration of services provided by humans and machines.

In a service-oriented world, services must be discovered and selected based on requirements, then be orchestrated and adapted or integrated. If a web interconnecting billions of services is to be realised, the way this is done must clearly be both scalable and manageable.

In this paper, we present the principles and rationale behind the EU Seventh Framework project ‘Service Oriented Architecture for All’ (SOA4All) [5] based upon the integration of SOA, semantic, Web and Web 2.0 technologies. After describing how these ‘technological pillars’ will be integrated in SOA4All, we indicate how our approach would be instantiated within a telecoms application.

## 2 SOA and Service-orientation Principles

At the heart of our approach are SOA and service-orientation that embody a number of key principles [4], which we describe below.

***Standardized Service Contract Principle.*** In order to make the description of service capabilities understandable to any interested party, the properties of a service – the service contract – should be compliant with some design standard. The service contract may include information to identify the services such as a URL; functional properties, such as the type of the input/output parameters; and non-functional properties, such as Quality of Service (QoS). Standardisation supports the global interpretability of services, resulting in an increase in the predictability of the service behaviour. The ability to predict the future behaviour of a service is a key mechanism to achieve scalability, since it facilitates the evaluation of the necessary computational

resources required to enact a specific service. This mechanism enables the intelligent provisioning of resources.

**Abstraction Principle.** The abstraction principle dictates that the details of software artefacts that are not required for effective use should be hidden. Therefore all the information necessary to invoke the service is contained in the service contract. Conversely, all the knowledge of the underlying logic and platform technology should be buried completely. The abstraction principle enables replaceability which, when combined with fault isolation and recovery as outlined in [6], enhances scalability.

**Reusability Principle.** The reusability principle states that the functionality provided by services must be as domain- and context-independent as feasible, facilitating reuse [7]. As a direct consequence of the application of this principle, the logic of a service should be highly generic – that is, independent from its original usage scenario. The reusability principle is a key enabler for SOA infrastructures, since it makes possible the creation of huge libraries of domain-independent services that leverage the construction of new complex context-dependent services.

**Autonomy Principle.** The autonomy principle states that services should be able to carry out their processes independently from outside influences. The only way to affect the results of a service should be through the modification of the input parameters as specified in the service contract. Service autonomy increases reliability and, more importantly, predictability and fault isolation. As set out in [6], this leads to improved system scalability.

**Statelessness Principle.** The statelessness principle dictates that services should minimise resource consumption by deferring the management of state information when necessary [7]. This notion of statelessness has been taken to the extreme in the REST architectural style [8], which has also been successfully applied to SOA in recent years. Because state maintenance is one of the most resource consuming tasks in computer science, conformance with the statelessness principle is vital if the entire SOA infrastructure is to be scalable.

**Discoverability Principle.** The discoverability principle states that, to enable discovery by interested parties, services should be annotated with metadata. This principle is closely related to the standardised service contract principle.

**Composability Principle.** The composability principle requires services to be identified as effective composition participants, regardless of the size and complexity of the composition [7]. From a bottom-up perspective, this allows simpler services to be combined into larger services (c.f. reusability). Looked at top-down, service composition is an effective way to tackle the complexity of certain types of processes (c.f. abstraction). Because the ability to create new services easily is a key pre-requisite for the widespread take up of SOA, the composability principle is a core element within the definition of a service web.

Despite the original aim to support inter-organisational processes on the basis of SOA, it still remains almost exclusively an intra-organisational solution. To support the service web, we believe the principles underlying SOA must be revisited and extended with others coming from the web, semantic technology and Web 2.0 systems. We consider each of these in turn in the next three sections.

### **3 Enhancing SOA with Web Principles**

The service web will lead to a global, dynamically-changing environment of services accessible for third-party usage. Services will undergo many changes within this environment, as resources currently do on the web, and the rate of churn will be very high. Users and resources will appear, disappear, and change location. Resources that initially are available free will transform to pay-per-use and vice versa. And services may occasionally be blocked, out of service, or inspected for antitrust violations. As a result, we believe that the transformation of SOA into an architecture comprised of billions of services requires the embodiment of the principles which made the web such a successful platform for the worldwide sharing of content.

Among the main principles underlying the web, its distributed and open nature is most relevant. The web is essentially a collection of distributed resources which, however, transparently appears to the user as a single entity where anybody can provide and consume resources. Distribution enables scalability but, when it comes to supporting the execution of processes, the absence of central control over many of the resources and services involved means other mechanisms must be provided to cope with changing conditions. Research in service-oriented computing already recognises the need for dynamic and adaptive processes within enterprises [9]. The service web will exacerbate this to a point where the existence of adequate means for self-configuring, self-adapting or self-healing processes will determine its success. In addition, aspects like the discovery and composition of services and the distributed and open nature of the service web will make it a very valuable yet technically challenging environment.

Another principle underlying the web is the use of a vendor-neutral interoperability platform that supports, and is based on, the integration of heterogeneous and proprietary solutions. A web of billions of services will require an advanced infrastructure that properly supports the integration of data and processes independently of their format, protocol or location. In fact, we foresee the creation of constellations of service ecosystems that span institutions, communicating seamlessly using the service web as a common interoperability platform.

Finally, a key feature of the web lies in the central role that users play. The ability to maintain this role in the service web will determine the market for services, which will in turn affect the uptake of service technologies on a web scale. For users to continue to play a central role as the service web evolves and grows, simple yet fully-fledged and customisable solutions must be created.

### **4 Enhancing SOA with Semantic Technologies**

Standards for describing web services currently use syntactic (XML-based) notations such as WSDL. These descriptions are not amenable to applying automated reasoning, so specialist workers must be involved throughout the web service lifecycle. This causes numerous problems, the most significant of which is the lack of scalability. It simply isn't possible to maintain millions, let alone billions, of services to cope with environmental and context changes, discover new services, compose



them or support their adaptation at runtime through human effort alone. Researchers studying the semantic web, semantic web services and more traditional artificial intelligence have shown it is possible to automate some of these tasks to a significant extent (see, for instance, work in OWL-S [10], WSMO [11], WSDL-S [12], semantic execution environments [13, 14] or planning [15]). Based on their findings, we decided to use semantic technologies as a core pillar of the service web.

By providing semantic descriptions of service interfaces and capabilities, the way is paved for simplifying and further automating the discovery of suitable services. In particular, services can be discovered based on the capabilities they offer rather than the low-level messages exchanged. Our approach to discovering suitable services is based on the notion of *goal* set out in WSMO. On the one hand, goals provide means for users to model explicitly their requirements (what to do); on the other, they serve as a natural abstraction over web services (how to do it). The notion of goal therefore enables the reusability of existing services while reducing the complexity for managing billions of services by providing additional layers of abstraction for guiding the discovery process [2].

This very distinction between goals and web services also enables what is usually referred to as late binding – that is, service selection at run-time [13, 14] – which has proven a useful approach for the management and execution of services [16]. From the management perspective, specifying processes in terms of goals provides further robustness and maintainability to the process models avoiding the creation of ‘spaghetti solutions’ that deal with the specifics of different service providers’s implementations and using instead a single entry point. From an execution perspective, the infrastructure is given the possibility to choose the ‘best’ service taking into account contextual factors.

To achieve the level of flexibility required for the service web, we decided to extend the heuristic classification-based engine we use to dynamically select trusted services [17] to support more general contextual data such as location and QoS. To cope with the high dynamism and unreliability of the web – that services may disappear, for example – additional adaptive capabilities are required. To deal with changing conditions while also ensuring that service-level agreements are met, processes must be self-optimising. And to deal with unreliable services, they must also be self-repairing. To meet these requirements, we make use of:

- semantic descriptions of functional and non-functional properties of services;
- parametric design techniques [18] for supporting the (re)configuration of services based on contextual factors; and
- event-based opportunistic reasoning techniques based on the use of semantic spaces [19, 20].

By reducing the complexity for creating new services out of existing ones, full support for the creation of services within the service web will play a very important role in the uptake of service technologies. We have therefore decided to approach the creation of services through both manual and (semi-) automated processes. As is explained in more detail in the next section, this provides better support to users creating services. The (semi-) automated creation of atomic services will also be supported by intelligent recommenders that analyse service documentations and suggest semantic annotations. Furthermore, to ensure scalability, the creation of


composite services will be supported by applying parametric design over process templates. This avoids the computational complexity of planning techniques [21].

Finally, throughout their lifecycle, services will benefit from the presence of formal semantic descriptions that enhance interoperability between humans and machines as well as between machines and other machines. Conceptual descriptions will bridge the gap between the user and the service descriptions and support the integration of information through semantic web technologies such as RDF(S) [22].

## 5 Enhancing SOA with Web 2.0 Technologies

The adoption of Web 2.0 technologies and principles within the service web will accelerate the take-up of service technologies and make it easier to integrate the user within the overall architecture as an extremely valuable source of information and computational power. To achieve this goal, the semantics we use are kept lightweight. User interfaces are enhanced with recommender systems that make them more intuitive to use, and the distinction between providers and consumers is blurred.

The use of lightweight semantic descriptions reduces computational complexity, making it easier for users to achieve their tasks. In particular, we use WSMO-Lite [23] and MicroWSMO [24] for the annotation of WSDL and RESTful web services. These semantic descriptions will allow efficient and scalable reasoning about services and will reduce the complexity for users. Note that, in this sense, the project intends to be protocol-neutral, an important feature given the split between RESTful and WSDL-based web services. Figure 1 shows the usage of each as reported by ProgrammableWeb [25], a website that maintains a comprehensive directory of web Application Programming Interfaces (APIs) [26] and the protocol(s) they support.

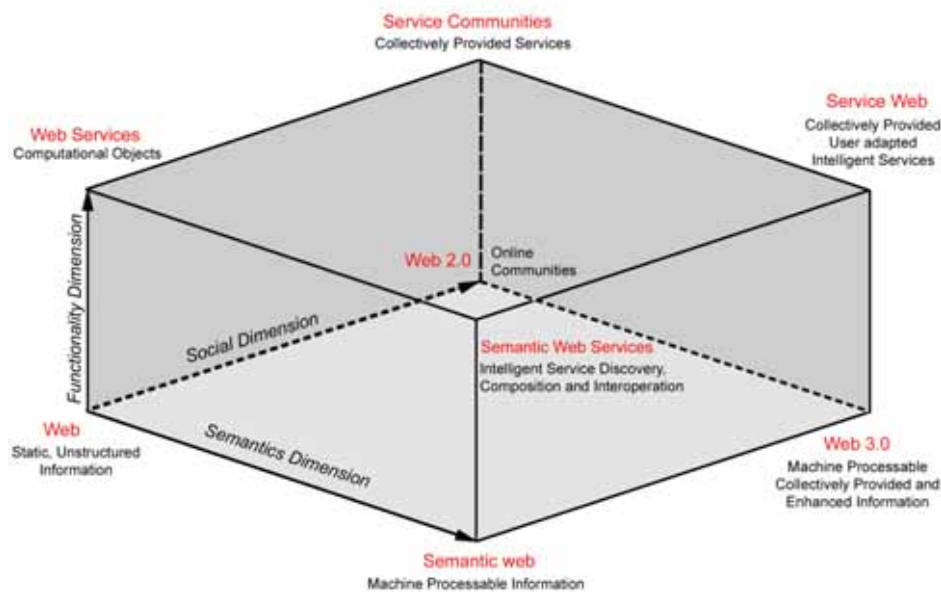


**Fig. 1.** Protocol Usage of Published Web APIs.

Users will be given the means to create lightweight service annotations by making use of recommender systems, but also through intuitive user interfaces such as mash-ups or pipes. Currently applied to data, these Web 2.0 solutions will be adapted to support the definition of simple service compositions and their publication. Through

this simple mechanism, people will co-operate in building online communities around services. Simultaneously, such services will interoperate with one another to offer more sophisticated functionalities to the users in a largely automated way. Services will be combined in increasingly complex mash-ups that offer functions to support users both at home and work, helping them to perform their daily activities. This approach will blur the distinction between service consumers and providers, shifting the paradigm to one where users are active and are therefore contributing to the service web, rather than just being consumers of services. By incorporating human interaction and cooperation in a comprehensive fashion, tasks such as service ranking and mediation, that would otherwise be computationally expensive or even infeasible, can be addressed. In our view, the quality of the services available to the end-user can be increased significantly if it is transparent whether the ‘engine’ that abstracts services is completed by humans or machines.

Figure 2 summarises the impact of the four technological pillars discussed above on the service web of the future. By combining the semantic and social dimensions, Web 2.0 and semantic technology, we arrive at Web 3.0. By combining semantic technology and web services, we create semantic web services. And by combining Web 2.0 and web services, we enable the creation of service communities. Finally, as shown in Figure 2, when we combine all the advantages and enhancements offered by the four pillars, we arrive at the service web.



**Fig. 2.** Service Web technological pillars.


## 6 Application in Telecommunications: Telco 2.0 and Beyond

In this section, we will look at the potential impact of the emerging service web on the telecommunications sector. We start by setting out the business background, showing how the drive to an open service ecosystem is increasingly seen as inevitable and the implications for the telecommunications sector and beyond. We then explain the potential benefits gained by the deployment of SOA4All technologies in a specific telecommunications public SOA environment, BT's Ribbit services platform.

### 6.1 Business Background and the Drive to Open Service Ecosystems

With the increasing tendency of service providers of all types to publish services via the web and the emergence of Web 2.0 technologies, traditional telecoms companies (telcos) are being forced to evolve. Indeed, Figure 3 reveals that the telecommunications sector is among the leading areas for publication of functionality via the web. The key technological trends which demand telcos' immediate response come from Web 2.0 developments. Webcos – companies adopting Web 2.0 principles in their business models – are able to respond to changing demands and expectations in the marketplace by innovating at multiple levels. Their products are distinguished by the following key characteristics:

- **Web as platform:** The platform is no longer a specific server or application, but exists on the web ('in the cloud') and is characterised by the publication of capabilities rather than vertical applications (for example, the exposure by Google and Amazon of APIs via the web).
- **Architecture of participation (harness collective intelligence):** designed to encourage users to take part, to share, to customise, to connect and even to participate in future product design (for example, user generated content – Flickr, YouTube, Delicious, mySpace, eBay, Amazon, social networks and user participative sites such as Wikis, blogs, Amazon reviews). The wisdom of crowds is a phrase used to refer to the harnessing of large numbers of end-user views and insights via these technologies.
- **Mash-ups:** as mentioned in the previous section, lightweight and rapid service/product composition in an open service ecosystem, leading to a larger number of applications developed more quickly and able to serve niche markets.
- **Long tail:** monetising the demand from the large number of highly diverse potential customers with non-typical requirements. Traditionally, telcos have focussed their efforts on a much smaller number of products and services addressing a larger market. In the Web 2.0 world, telcos can no longer ignore the revenues generated from the long tail.



**Fig. 3.** Top 15 categories of Published Web APIs (source [25]).

Ribbit [27] is a wholly-owned US-based BT subsidiary and a voice-oriented web company ('webco'). Ribbit currently allows developers to consume voice services accessible via Adobe's Flex and Flash. The services that are exposed at the moment include voice calls, call routing, call management, third-party call control, voice and text messaging, speech-to-text, VoIP and contact management, with more to come in the near future. Currently, Ribbit users require detailed technical knowledge of the Flex and Flash programming languages in order to be able to access, combine and use web services.


In the SOA4All project, the use of contextual knowledge will support both the composition and provisioning of services in a customised manner. Contextual knowledge, includes, for example, knowledge about end-users and their preferences as well as about the previous experience of other users with respect to use of particular services. Using automation, we plan to shield service users from the complexities of creating such knowledge. BT will also take advantage of semantic descriptions of services in building (semi-) automated provisioning, composition, consumption and monitoring tools. The next-generation platform we envisage will also enable inclusion of third party services. In addition, it will address several key issues for BT's transformation, including: reducing time to market; enabling integration of third party services into BT's portfolio; increasing New Wave revenue and extending BT-wide SOA to the public web.

The possibility of increased competition from the 'Over The Top' (OTT) service providers – service and content providers that do not own the network they use – highlights a risk that telcos could become 'disintermediated' from the digital supply chain, consigned to the role of a commodity 'dumb' pipe provider. Webcos typically use advertising-based business models, whereas telcos collect revenues through usage-based billing. As these two sectors converge, the challenge for telcos is to reconcile the two different business models, finding ways to generate revenue from advertising, while continuing to offer billable services where appropriate. Finally,

other changes are apparent such as the rise in virtual social networking, the roll-out of alternative access networks such as WiMax and the emergence of enterprise mashups that use Web 2.0-style applications alongside more traditional IT assets.

Considering all these aspects, by appropriately positioning themselves in the Web 2.0 world, telcos will continue to evolve and transform themselves from mere 'dumb' pipe providers to providers of 'smart' pipes (connections backed by QoS guarantees and service level agreements), platforms that support an open service ecosystem and a range of telco and third-party applications that run on such platforms. Telcos will not only need to create new services to address the needs of the long tail, but also allow third-party service providers to make use of telcos' underutilised operations and business support systems capabilities to create new service offerings, thereby creating new revenue streams.

Figure 4 shows the long tail, representing untapped new business revenues, and the short head representing the conventional telco business revenues. New growth opportunities will arise from externalising current capabilities, allowing both the telco itself and third-party developers to address long tail demand by combining core services into more complex applications.



**Fig. 4.** Leveraging the long tail.

Telcos have historically been used to controlling the entire value chain from core network to end user. As discussed above, the long tail business model involves letting the market innovate, by using third-party developers to define how new services will be generated. Because this involves opening up the network, some loss of control is inevitable. In the future, the technologies employed will be telecom web services, representing an amalgam of telco services and internet services, enabling the telco to access long tail revenues via innovative niche applications from third party developers: the telco is taking the role of an aggregator.

## 6.2 Application of SOA4All Technology

The partners in SOA4All are investigating how BT – and, potentially, other telcos – can apply the technology being developed in the project’s various core research areas to deliver next-generation web-based open services platforms. Below we describe the application of each technology area in turn.

### Web 2.0

User-generated and maintained communities are an essential part of the case study, as well as encouraging ease of use and a low barrier to entry in utilising SOA. SOA4All technology based on Web 2.0 principles helps encourage this, specifically to:

- encourage users to work together to create new and innovative uses for BT services by providing an appropriate Web 2.0 community environment;
- provide facilities for users to share information about services and applications;
- allow users to tag, rate and comment on services for improved service discovery; and
- improve the usability of SOA, and facilitate the creation and management of new composed services, in a lightweight manner, with a low barrier to entry.

Related to the availability of Web 2.0-style tools for developers and users, the use of contextual knowledge will be key in supporting both the customised discovery and composition of services. The complexity underlying this will, however, be automated, so the process will remain hidden from the service user.

Service composition will use contextual knowledge about both the user domain, (for example, knowing the current user is an ISP targeting rural areas) and community knowledge (for example, knowledge of services tending to be used by similar ISPs) to suggest and customise solutions based on prior experience. As a result, the user will obtain more specific and more adapted proposals, which will simplify the selection of suitable services. This can be achieved by pre-defining each constituent service with a number of options, by parameterising certain functionality, and by providing a context-adapted service discovery and selection process according to pre-identified context dimensions.

Service provisioning will also be enhanced by the use of contextual knowledge. This will be achieved at run-time by integrating execution information, together with contextual knowledge such as user preferences and trust relationships (for example, which service providers does a given end-user trust). To do this, the service execution platform will integrate the information populated by the monitoring infrastructure with the contextual knowledge base, and will seamlessly invoke the context parameterisation engine as the need arises.

### Semantics

BT will take advantage of the semantically-enabled improvements in provisioning, consumption and monitoring tools that are offered by SOA4All technology. Semantic descriptions for web services and goals will be built by BT and non-BT users in a straightforward way using purpose-built tools, allowing them to make available their services, or to discover and use the services of others, more efficiently. Similarly,



semantics will also be important to enable new composition tools that will enable the creation of more complex services at lower cost.

Ontologies and semantic descriptions will form the basis for data, process and service representation in SOA4All. In the context of applying SOA4All technology to the future web-based telco platforms, telecommunications domain ontologies will also be used [28]. One example is NGOSS (Next Generation Operations Systems and Software), an industry-wide specification developed by the TeleManagement Forum [29] the purpose of which is to organise and guide the design and development of next generation operation systems in the telco domain.

NGOSS contains a set of frameworks, high-level architecture and methodology. It consists of four frameworks related to the different levels of looking at business:

- Business process framework (Enhanced Telecom Operations Map – eTOM)
- Enterprise-wide information framework (Shared Information and Data model – SID)
- Applications framework (Telecom Applications Map – TAM )
- Systems Integration framework (Technology Neutral Architecture – TNA)

Ontologies have been developed that capture telecommunication sector knowledge from NGOSS's standards. The SID model contains domain concepts related to market, product portfolio, customer, services, resources, the enterprise and supplier/partner, as well as common business terms called Core Business Entities (CBE) which are captured in the CBE Ontology (CBEO). The eTOM map defines a set of functional areas which serves as a reference classification for the business goals a process fulfils and which are captured in the business goals ontologies (BGO). The TAM map defines the typical IT systems map of telecommunication companies, and serves as a reference classification for a company's services map.

### **Services**

As mentioned above, BT's current offering consists of Ribbit's services including voice calls, call routing, call management, third-party call control, voice and text messaging, speech-to-text, Voice over Instant Messaging (VoIM) and contact management, with more to follow shortly. As more services are made available, they will be included on the next generation platform.

In addition, there are also other business-to-business gateways and APIs that BT uses for interaction with its customers and suppliers, such as broadband provisioning [30] repair and diagnostic services. SOA4All technology also has the capability to describe these services in the same semantic framework (see, for example, [31]), allowing them to be combined with Ribbit services.

The third and final class of services are third-party services. One of the main aims of the case study is to promote the uptake of Ribbit by offering tools to encourage innovation in using and combining BT's services. SOA4All technology will make it easier not only to consume BT's services but also to combine them with other people's services to make new and interesting applications.

## 7 Summary

SOA4All will help to realise a world where a massive number of parties expose and consume services via advanced web technology. The outcome of the project will be a comprehensive framework and software infrastructure which will integrate complementary advances into a coherent and domain-independent worldwide service delivery platform. To achieve such a scalable and widely adopted infrastructure and framework, SOA4All stands on four main principles comprising: SOA, semantics, the web and Web 2.0. Our technologies will be proven in a number of real-world applications including the telco scenario outlined above.

From our point of view, we believe that the successful integration of semantic web and service-oriented technologies will form the main pillar of the software architecture of the next generation of computing infrastructure. We envision a transformation of the web from a web of static data to a global computational resource that truly meets the needs of its billions of users. Computing and programming will be positioned within a services layer thus putting problem solving in the hands of end-users through a truly balanced, cooperative approach.

## Acknowledgements

This work has been supported by the EU co-funded IST project SOA4All (FP7-215219).

## References

1. Cusumano, M.A.: The Changing Software Business: Moving from Products to Services. *Computer* 41 (2008) 20-27
2. Stollberg, M.: Scalable Semantic Web Service Discovery for Goal-driven Service-Oriented Architectures. Austria (2008)
3. Seekda home page. <http://seekda.com> (2008)
4. Benjamins, R., Davies, J., Dorner, E., Domingue, J., Fensel, D., Lopez, O., Volz, R., Wahler, A., Zaremba, M.: Service Web 3.0. DERI (2007)
5. SOA4All: SOA4All web site. <http://www.soa4all.eu> (2008)
6. Armstrong, J.: Making Reliable Distributed Systems in the Presence of Software Errors. (2003)
7. Erl, T.: SOA Principles of Service Design. Prentice Hall (2007)
8. Fielding, R.T.: Architectural Styles and the Design of Network-based Software Architectures. (2000)
9. Papazoglou, M.P., Traverso, P., Dustdar, S., Leymann, F.: Service-Oriented Computing: State of the Art and Research Challenges. *Computer* 40 (2007) 38--45
10. Martin, D., Burstein, M., J., H., Lassila, O., McDermott, D., McIlraith, S., Paolucci, M., Parsia, B., Payne, T., Sirin, E., Srinivasan, N., Sycara, K.: OWL-S: Semantic Markup for Web Services. <http://www.daml.org/services/owl-s/1.0/owl-s.pdf> (2004)
11. Fensel, D., Lausen, H., Polleres, A., de Bruijn, J., Stollberg, M., Roman, D., Domingue, J.: Enabling Semantic Web Services: The Web Service Modeling Ontology. Springer (2007)

12. Akkiraju, R., Farrell, J., Miller, J., Nagarajan, M., Schmidt, M.-T., Sheth, A., Verma, K.: Web Service Semantics - WSDL-S. <http://www.w3.org/Submission/WSDL-S/> (2005)
13. Norton, B., Pedrinaci, C., Domingue, J., Zaremba, M.: Semantic Execution Environments for Semantics-Enabled SOA. *it - Methods and Applications of Informatics and Information Technology Special Issue in Service-Oriented Architectures* (2008) 118--121
14. Fensel, D., Kerrigan, M., Zaremba, M. (eds.): *Implementing Semantic Web Services: The SESA Framework*. Springer (2008)
15. Traverso, P., Pistore, M.: Automated Composition of Semantic Web Services into Executable Processes. *3rd International Semantic Web Conference (ISWC 2004)* (2004) 380--394
16. Pedrinaci, C., Brelage, C., van Lessen, T., Domingue, J., Karastoyanova, D., Leymann, F.: Semantic Business Process Management: Scaling up the Management of Business Processes. *Proceedings of the 2nd IEEE International Conference on Semantic Computing (ICSC) 2008*. IEEE Computer Society, Santa Clara, CA, USA (2008)
17. Galizia, S., Gugliotta, A., Pedrinaci, C.: A Formal Model for Classifying Trusted Semantic Web Services. *3rd Asian Semantic Web Conference (ASWC)*, Bangkok, Thailand (2008)
18. Motta, E.: *Reusable Components for Knowledge Modelling. Case Studies in Parametric Design Problem Solving*, Vol. 53. IOS Press (1999)
19. Pedrinaci, C.: Knowledge-Based Reasoning over the Web. (2005)
20. Krummenacher, R., Simperl, E., Fensel, D.: Towards Scalable Information Spaces. In: Piskac, R., van Harmelen, F., Zhong, N. (eds.): *New Forms of Reasoning for the Semantic Web*, Vol. 291. CEUR-WS.org (2007)
21. ten Teije, A., van Harmelen, F., Wielinga, B.J.: Configuration of Web Services as Parametric Design. In: Motta, E., Shadbolt, N., Stutt, A., Gibbins, N. (eds.): *EKAW*, Vol. 3257. Springer, Whittlebury Hall, UK (2004) 321--336
22. Brickley, D., Guha, R.V.: *RDF Vocabulary Description Language 1.0: RDF Schema*. (2002)
23. Vitvar, T., Kopecky, J., Viskova, J., Fensel, D.: WSMO-Lite Annotations for Web Services. In: Hauswirth, M., Koubarakis, M., Bechhofer, S. (eds.): *Proceedings of the 5th European Semantic Web Conference*. Springer Verlag, Berlin, Heidelberg (2008)
24. Kopecky, J., Vitvar, T., Gomadam, K.: *MicroWSMO. Conceptual Models for Services Working Group* (2008)
25. Musser, J.: ProgrammableWeb: 1000 web APIs. <http://blog.programmableweb.com/2008/11/03/1000-web-apis> (2008)
26. ProgrammableWeb: Web 2.0 API directory. <http://www.programmableweb.com/apis/directory> (2008)
27. Ribbit home page. <http://www.ribbit.com> (2008)
28. Munoz, H., Pérez, N., Evenson, M., Szczekocka, E., Ksiezak, H., H.Kupidura, Belecheanu, R., Krysteva, P., Roman, D., Todorova, P., López, O., Cicurel, L.: Deliverable 8.1: Telecom Vertical Domain and Process Ontologies. *SUPER (IST 026850)* (2007)
29. Forum, T.: The TeleManagement Forum home page <http://www.tmforum.org> (2008)
30. BT Wholesale, B2B Gateway. [http://www.btwholesale.com/pages/static/Applications/Orders/B2B\\_Gateway.html](http://www.btwholesale.com/pages/static/Applications/Orders/B2B_Gateway.html) (2008)
31. Duke, A., Richardson, M., Watkins, S., Roberts, M.: Towards B2B Integration in Telecommunications with Semantic Web Services. *Second European Semantic Web Conference, ESWC 2005, Crete, Greece* (2005) 710--724

## User-Centric Future Internet and Telecommunication Services

Carlos Baladrón<sup>1</sup>, Javier Aguiar<sup>1</sup>, Belén Carro<sup>1</sup>, Laurent-Walter Goix<sup>2</sup>, Alberto León Martín<sup>3</sup>, Paolo Falcarin<sup>4</sup>, Jürgen Sienel<sup>5</sup>

<sup>1</sup>Universidad de Valladolid

Campus Miguel Delibes, Camino del Cementerio s/n, 47011 Valladolid, Spain.

+34 983 423 704

{cbalzor, javagu, belcar}@ribera.tel.uva.es

<sup>2</sup>Telecom Italia

Via Reiss Romoli 274, 10148 Torino, ITALY

laurentwalter.goix@telecomitalia.it

<sup>3</sup>Telefónica I+D S.A.U.

Emilio Vargas 6. 28004 – Madrid, SPAIN

Tel: +34 913374000 Fax: + 34 913373966

alm@tid.es

<sup>4</sup>Politecnico di Torino

Corso Duca Degli Abruzzi 24, Torino, Italy

paolo.falcarin@polito.it

<sup>5</sup>Alcatel-Lucent Deutschland AG

70435 Stuttgart, Germany

Juergen.Sienel@alcatel-lucent.de

**Abstract.** This paper analyses the current service creation trends in telco and Web worlds, showing how they are converging towards a future Internet of user-centric services embracing typical telco capabilities. The OPUCE platform is presented as the next step towards this integrated, user-centric future: a platform which offers intuitive tools for graphical service creation aimed at individuals with no specific skills in computer science or programming and a service-oriented execution environment capable of a seamless interoperation of Web Services and telco applications based on operator-owned infrastructure. The OPUCE platform is compared to existing mashup creation tools to show its advantages and how it could be used to implement a converged and open service marketplace for the Future Internet.

**Keywords:** Service Creation, User-centric, Service Delivery Platform, Web 2.0. mashup, Web Service.

### 1 Introduction

The whole Information and Communication Technologies (ICT) world is facing a revolution influenced mainly by the so called Web 2.0 [1], where users take control of

their own experience in the Internet not only consuming contents, but also creating them. The Web 2.0 has reached the majority of the digital society, on every facet of the modern life. Its revolution is one of the most impacting paradigm shifts in the ICT history, and is called to be a key driver for the Future Internet challenges.

The Web 2.0 is heading towards a global service ecosystem, and a good example is the mashup phenomenon, where small applications are built by coordinating remote Web services. In order to allow user-centric mash-up creation, some online tools already exist to allow easy and intuitive composition of Web services for early adopters, like Yahoo Pipes (<http://pipes.yahoo.com>), or general public, like Microsoft Popfly (<http://www.popfly.com>). Users can compose with these tools their own simple mash-ups, in some cases without requiring very high computer science and ICT skills. The vision of the Internet of Services is starting to become real thanks to these tools. In the nowadays Internet, users are able to provide their own content. In the future Internet, they will be able to provide their own services.

Analyzing current telecommunications innovations and market trends, it is easy to realize that there is a clear approach to synergize with Information Technologies (IT) and their new user-centric Web 2.0 philosophies. A real convergence is starting to take place, allowing for the first time an easy development and management of value-added merged services by using Telco and IT infrastructures and technologies, but considering the social benefits brought by Web 2.0. This will foster new business models for an integrated market, evolving from the traditional walled-garden of the operators to the common open one seen in the Internet, with new business actors and a full revision of the former value chain.

As such, the new Internet paradigms are being also adopted by key Telco players. This merger is lowering down barriers for SMEs and empowering users that will not need to deploy their own infrastructure to become Service Providers and share revenues. The goal is to empower the users, enrich their experience and satisfaction, and be able to face the increasing competence in the telco field posed by IT companies that are entering the communications market.

In order to achieve a complete convergence, Telco companies need to complete two stages: first, offer their services through the internet, seamless integrated with existing information services. And second, to let end-users take control of their communication experience building their own communication services. This new landscape is sometimes referred as Telco 2.0.

This paper presents how the telco world is moving towards this Telco 2.0 paradigm, opening their infrastructures, employing Web Services technologies to wrap their applications and allowing them to interact in the Internet, and replacing the old monolithic service creation approach by user-centric Service Creation Environments inspired in the existing Internet mashup editors. As a specific example, the OPUCE platform is presented: a prototype environment where end-users are able to create, share, manage, enjoy and consume their highly-personalized, converged Internet and Telco services. It offers intuitive graphical tools for service creation, lifecycle management and community support, and an open execution environment where telecommunication services can seamlessly interact with Internet-based Web Services.

This paper is organized as follows. Section 2 discusses the two trends that are giving birth to a converged, end-user driven Web of Services: the user-centric service creation in the Internet and the horizontal network of the telco operators which allows easy and open access to network functionalities. Section 3 presents the OPUCE platform, an environment to allow user-centric creation of converged services. Section 4 compares

the OPUCE platform with existing user-centric service creation in the Internet. And finally Section 5 exposes the conclusions.

## **2 Current Trends**

### **2.1 User Centric Service Creation**

Web 2.0 is user-centric. Currently, the user is mostly assuming the role of content generator through Web sites like Youtube, Flickr or Wikipedia, but as the Web of services is approaching, end-users would eventually end in control of creation and management of their own services. And as specialized knowledge is required to perform these tasks, suitable automation tools should be developed to guide non-expert individuals.

User-centric service creation is entering the Web in the form of Mashups, small Web applications made by the composition of two or more Web services. But it is still necessary to have specific knowledge in order to build them, and because of that, tools are required to let non-expert individuals jump into the mashup phenomenon. Two of the flagship enterprises of the ICT world, Yahoo! and Microsoft, have recently released their own Service Creation Environments (SCEs) for mashup composition: Yahoo! Pipes and Microsoft Popfly. These environments are tools in which the creation of services is made by graphically orchestrating a set of basic functionalities offered by the environment. Boxes representing these functionalities (operations implemented in the environment and remote services accessible through the Web) are dragged and dropped, and then inputs and outputs linked, so it is not necessary to write any kind of code to develop a new application.

However, these environments for mashup creation are limited because only simple services involving basic IT capabilities can be built. In order to follow the convergence trends with communications, the next step is the integration with Telco 2.0, which probably requires true SOA implementation (unlike most current mashup creation environments) to allow proper interoperation [2].

Some initiatives, such as [3], have developed user-oriented creation/execution environments for converged ICT functionalities, but they basically fall short to be “true Web 2.0” because, first, they are heavily tied to a technology and thus hard to integrate with services out of the platform defined; and second they only offer tools for creation and execution, while service management implies a lot more processes like deployment, sharing, publishing or adaptation. A platform aimed at offering true user-oriented service management has to offer automation tools to perform all this steps.

### **2.2 Service Creation in Telecommunications**

Telco operators have been facing problems because for years their main sources of revenues have been pure transport services. The increasing number of competitors due to the liberalization of markets, the spreading perception of transport services as a commodity and the telco services offered by the Internet industry, are pushing operators to enter new markets in order not to live only on those pure transport services.

As such, operators are moving away from their old business models. They are taking advantage of the infrastructures they own by offering value-added services difficult to provide without control over the networks.

The old vertical network infrastructure is depicted in Fig. 1 included an independent service stack for each access network. It presented some problems when used to deliver high-level services: it required independent implementations of the logic, deployment procedures and servers for each access network and service. The answer to these problems, the horizontal network [4] depicted in Fig. 2, abstracts the network capabilities by means of a middleware known as Network Services Layer used to virtualize the network resources underneath. In this approach only one implementation of the service logic is enough, because the new layer adapts it to the network elements implied. Additionally, access to network capabilities could be granted to external third parties, resulting in an “open service marketplace” [4]: third parties can provide services without the need of deploying their own networks, and the operator gets additional incomes from the usage of its networks and widens its service catalogue.




Fig. 1. Vertical network scheme.

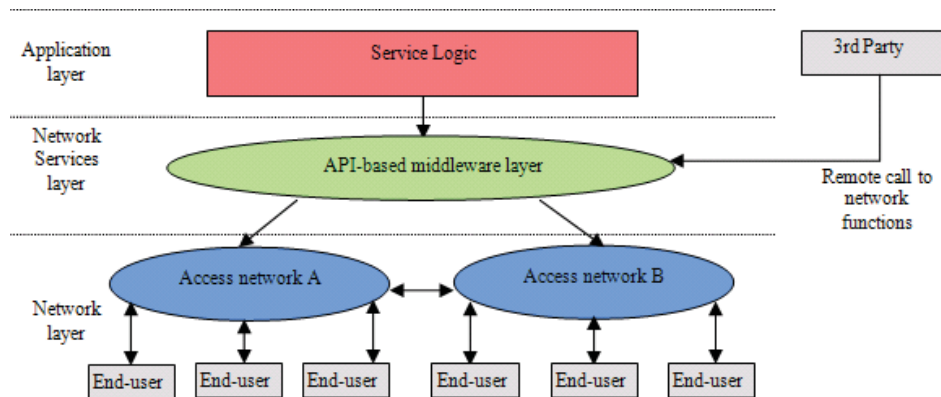


Fig. 2. Integrated horizontal network.

Several industrial solutions have been designed to implement the Network Services layer, mostly in the form of APIs (Application Programming Interfaces), such as JAIN

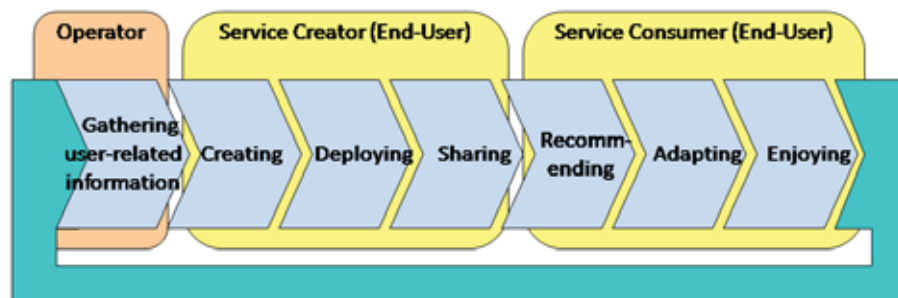


and Parlay [5] (which historically derives from the TINA effort). Additionally, innovative SOA approaches (such as Parlay X) are being designed in order to allow remote calling of network capabilities in the form of loosely coupled Web services. One of the most groundbreaking initiatives in this area is the Web21C SDK project of British Telecom [6]: Network capabilities of the BT infrastructure have been exposed in the form of Web Services, so Web applications are able to make use of them in exchange of a fee.

The development of environments (an infrastructure plus a set of high level tools for service composition [7] on top of the Network Services layer) to allow fast and easy creation and management of services is a priority of the current ICT community. An example is the prominence given to this issue in the sixth Framework Program (FP6) of the European Commission, where several funded projects are pushing research in open and integrated network infrastructures. The Mobile Service Platform (MSP) cluster and projects such as SPICE (Service Platform for Innovative Communication Environment), SMS (Simple Mobile Services), LOMS (Local Mobile Services) or OPUCE, deal in one way or another with the development of this kind of open service platforms.

### 3 OPUCE Platform for User-Centric Service Creation and Management

FP6-34101 OPUCE (Open Platform for User-centric service Creation and Execution, <http://www.opuce.eu>) is an Integrated Project (IP) of the sixth Framework Programme of the European Commission inside the Information Society Technologies (IST) priority.



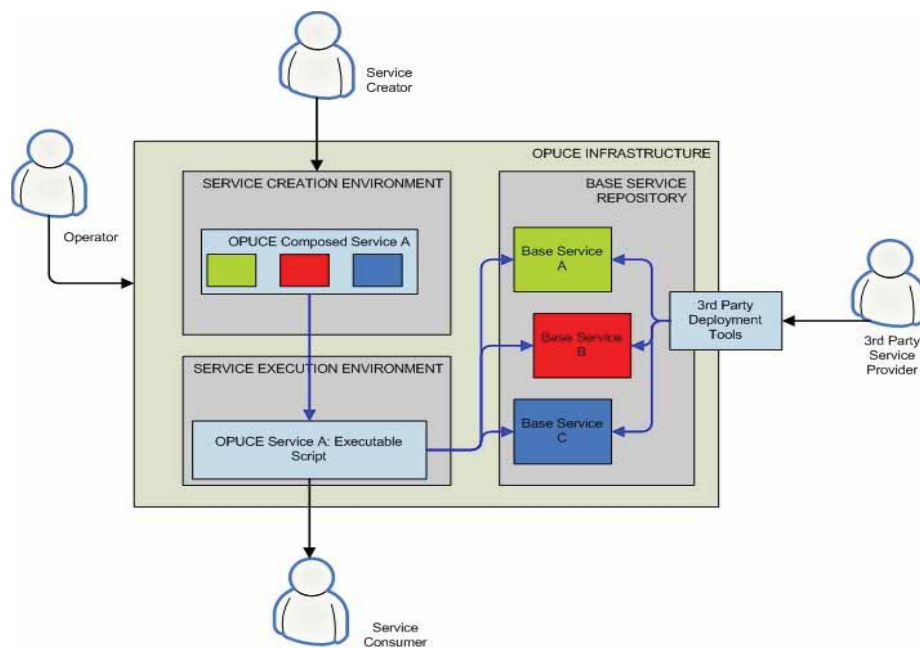
**Fig. 3.** OPUCE service lifecycle.

Its aim is to develop a Platform for service creation and management totally centred on end-users, so a non-expert individual could be completely in control of the lifecycle of a service. In the OPUCE context, this lifecycle has been defined to be comprised by the seven steps depicted in Fig. 3. All these stages are user-centric (controlled by the end-user, either assuming the service creator or service consumer role), except the first one, data gathering. It is performed by the operator because of obvious security and integrity reasons, and involves first the retrieval of information about users (their profile, preferences and context) to allow personalization of services and platform

experience, such as service recommendation; and second, the monitoring of all service lifecycle steps to retrieve usage history data.

The approach taken to service creation is similar to the one employed in current user-driven mashup-creation tools in the Web, like Yahoo Pipes: The environment offers a palette of basic functionalities, such as “Retrieve Location” or “Make Call”, called Base Services in the OPUCE context, that creators should orchestrate defining a workflow. For instance, a creator could link a “receive e-mail” Base Service with a “send SMS” Base Service to form a composed service which sends an SMS to a given number whenever an email is received.

The Platform is also designed to be accessible by third party service providers in order to implement innovative “open service marketplace” business models: the operator provides the infrastructure, third parties insert basic functionalities (Base Services), active service creators get minor revenues and consumers pay the former three for usage of services and networks. The role of these four actors is represented in Fig. 4.



**Fig. 4.** Main actors in OPUCE and simplified architecture.

Base Services are applications provided by the Platform or authorized third parties, wrapped as Web services following the OPUCE Base Service description model [8]. This allows them to be implemented in a variety of technologies, and then still be able to interact with the Platform execution environment.

Each Base Service exposes a set of properties and actions, and fires a set of events. Composition is made by linking events with actions. For example, a base service could fire the event “When-SMS-is-received” and another one expose the action “Make-video-call”. They could be linked to produce a service that automatically initiates a

video call whenever a SMS is received. Fig 5. shows an example of service composition in the OPUCE Service Creation Environment.

Properties are used to configure Base Services and share information between them. The “Make-video-call” Base Service would have a property “callee” that could be configured to be always a fixed phone number, or instead retrieve it from the property “SMS-body” of the “When-SMS-is-received” Base Service.

The output of the SCE is a service description [8], an XML document that stores all the information required by the platform about the new service: graphical layout in the SCE, semantics, service logic, etc.

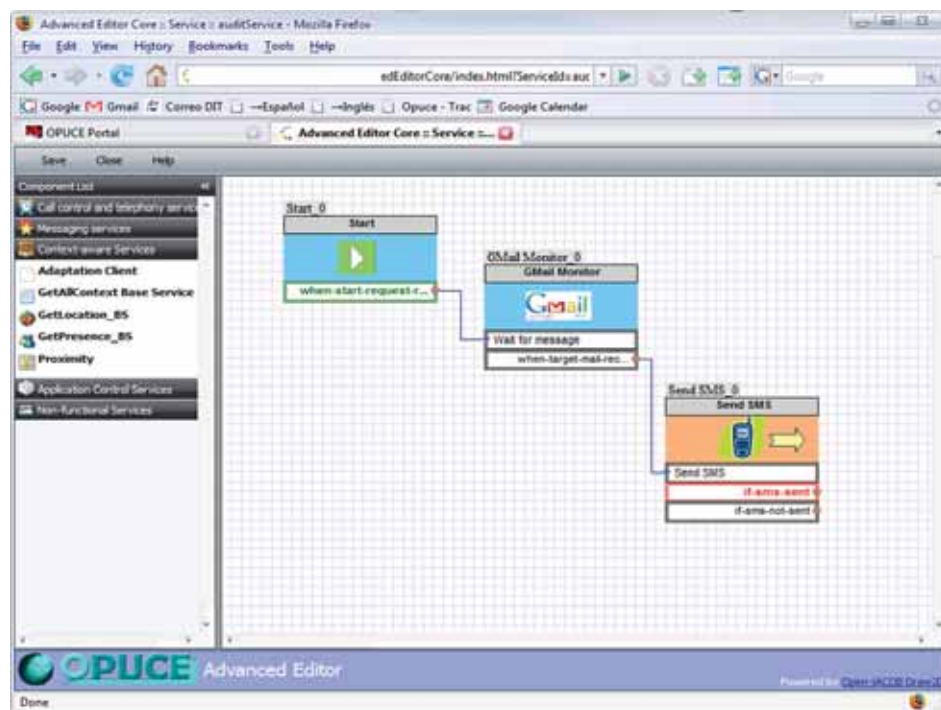


Fig. 5. OPUCE Service Creation Environment.

The OPUCE Platform Service Execution Environment [9] is two-tier: the first layer is a BPEL orchestrator which coordinates the execution of Base Services, and the second layer a set of low level execution containers which run those Base Services.

Whenever a service is started, a new instance is created inside the orchestrator, obtaining the BPEL script from the service logic repository. The orchestrator evaluates the script and invokes the appropriate base services in whichever execution environment they are located, forwarding user-related properties obtained from the OPUCE User Information Database, or creation time specified constants obtained from the service description. The events thrown by base services are captured by an event handler and forwarded to the orchestrator to determine the next action to invoke.

Each Base Service is deployed in an appropriate execution container (JAIN SLEE, J2EE, .NET, etc.), but all of them expose standard Web Service interfaces. The logic

inside the Base Service can implement a network function, accessing the operator's core infrastructure, or simply be a proxy for any external Web Service in the Internet. While these Web Services could be directly invoked by the orchestrator, the Base Service acting as a proxy performs security and accounting operations, relieving the external Web Service provider from these functions.

When a third party develops a Base Service, the OPUCE Base Service Manager tool could be used to wrap it as an OPUCE Base Service and deploy it in the appropriate execution container.

Additionally, the OPUCE platform offers a set of support tools: a Web Portal to act as a front-end towards end-users, a Service Lifecycle Manager [9] to automate difficult tasks such as service deployment, a Service Advertiser [8] to allow easy discovery of services and a User Information Manager to manage user profiles and context data.

#### **4 Value-Added Features of OPUCE Platform and Related Works**

User-driven mashup creation seems to be a very promising trend. Among other lesser options, three of the IT giants -Yahoo, Microsoft and Google- offer their own solutions in the form of Pipes, PopFly and GMashEd respectively. The OPUCE platform is based on the same principles than these tools, but applies them to the telco domain. The comparison of OPUCE against these tools could help to properly allocate it in the current user-driven mashup creation landscape. Table 1 compares the main characteristics of these four tools.

The very first identity sign of OPUCE, and its main advantage over existing mashup editors, is that it offers communication capabilities to be added to the created services. OPUCE could be used to create applications involving phone calls, multiconferences, SMS sending, etc. Additionally, as OPUCE is telco-driven, the great amount of information the operator keeps about the user (location, accounting, presence, etc.) could be employed to build context-aware highly-personalized services, while IT-driven tools can offer typical but limited Web 2.0 personalization (i.e.: through login).

Traditionally, telco applications have followed an asynchronous approach and IT applications a synchronous one. The OPUCE composition model is asynchronous to allow interaction with communication services, and at the same time integration with synchronous applications is granted by the platform middleware.

OPUCE services interact and are implemented with standard Web Services technology. This means that any other application in the Internet wrapped as a Web Service could be easily plugged into OPUCE. While current mashups usually employ other light-weight technologies, it seems that for professional applications true Web Service based SOA is the most suitable option [2] [10] [11] [12] because of a greater robustness and versatility, and that is the reason behind the choice of SOAP over the light-weight RSS used by other mashup editors for data exchange. If mashups end leaving the domestic domain towards a professional use, OPUCE is in a better position to fulfill the requirements of enterprises.

Finally, OPUCE is oriented towards building a viable market solution for the integration of communications in the Web, and as such it represents a step further in the migration of operators' business towards value-added converged services. Because of that, specific marketing systems such as billing support, advanced advertising or tools for third party service providers are included.

**Table 1.** Comparison between OPUCE and other mashup creation tools.

	OPUCE	Microsoft PopFly	Yahoo Pipes	Google GMashEd
<b>Telco capabilities</b>	YES	NO	NO	NO
<b>Information exchange format</b>	SOAP	RSS	RSS	RSS
<b>Graphical editor</b>	YES	YES	YES	NO
<b>Output/Execution</b>	OPUCE Platform	SilverSphere	RSS	RSS
<b>Composition model</b>	Asynchronous, event based	Synchronous	Synchronous	Synchronous
<b>Service personalization</b>	Advanced, context-aware personalized services	Limited	Limited	Limited
<b>Billing</b>	Flexible accounting system for all platform layers	Some 3 <sup>rd</sup> party blocks may require a fee	Not expected	Not expected
<b>Advertising/sharing</b>	Context-aware notifier	Web-based sharing	Web-based sharing	Web-based sharing
<b>Basic block provider</b>	3 <sup>rd</sup> parties, operator	End-users, 3 <sup>rd</sup> parties, Microsoft	Yahoo	N/A

## 5 Conclusions

After analyzing the current trends in the IT and telco worlds, it seems clear that the future Internet will be an Internet of converged services where end-users will be able to create their own integrated services packing information and communication capabilities. The Web and telco operators are moving towards an integration that will eventually result in an open service marketplace involving all actors.

After the OPUCE project is finished, the prototype OPUCE Platform could be considered to be at an alpha stage, fully operational but not tested to support high numbers of users. Although some open tests have been conducted in conjunction with OMF (Open Móvil Forum) where small amounts of independent prosumers have participated in web-based mashup creation contests, the platform is still not ready to be deployed as a commercial product.

However, the OPUCE platform packs several important advances towards the future integrated, user-centric Internet. It offers tools to completely automate all stages of the service lifecycle in order to allow non-expert end-users to create, manage and execute services without requiring any specific knowledge. It brings the concept of Web 2.0 user-driven mashup creation to the telco world, also applying SOA technologies in order to promote a future global converged Web of ICT services and allow heterogeneous coexistence of different technologies. And it also promotes the implementation of new “open service marketplace” business models, allowing external third parties to use the OPUCE Platform to provide additional functionalities by paying the operator accordingly.

Therefore, OPUCE could be considered as a step forward towards the integration of IT and telco worlds, offering an environment where both worlds can interoperate

seamlessly embracing the latest Web 2.0 user-centric trends. It offers relevant technical answers to key questions such as the interoperation of heterogeneous telco resources, how to provide an intuitive environment to let non-experts build services on their own or how to automate difficult tasks such as deployment and service management.

### Acknowledgements.

The work presented in this paper is executed as part of the OPUCE project and partly funded by the European Union under contract IST-034101. OPUCE is an Integrated Project of the 6<sup>th</sup> Framework Programme, Priority IST.

### References

1. Schroth C., Janner T.: Web 2.0 and SOA: Converging Concepts Enabling the Internet of Services. *IT Professional*, vol. 9, no. 3, pp. 36--41 (2007).
2. Pollet T., Maas G., Marien J., Wambecq A.: Telecom services delivery in a SOA. In: 20th International Conference on Advanced Information, networking and Applications, vol. 2, 18--20, (2006).
3. Glitho R. H., Khendek F., De Marco A.: Creating value added services in Internet telephony: an overview and a case study on a high-level service creation environment. *IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and Reviews*, vol. 36, no. 4, pp. 446--457, (2003).
4. De Serres Y., and Hegarty L.: Value added services in the converged network. *IEEE Communications Magazine*, vol. 39, no. 9, pp. 146--154, (2001).
5. Moerdijk A. J., Klostermann L.: Opening the networks with Parlay/OSA: Standards and aspects behind the APIs. *IEEE Network*, vol. 17, no. 3, pp. 58--64, (2003).
6. Web21C SDK project, <http://web21c.bt.com>.
7. Licciardi C. A., Falcarin P.: Technologies and guidelines for service creation in NGN. In: 8th ITU-IEEE International Conference on Intelligence in next generation Networks (ICIN '03), (2003).
8. Baladrón C., Aguiar J., Carro B., Sienel J., Trapero R., Yelmo J. C., Del Álamo J. M., Yu J., Falcarin P.: Service Discovery Suite for User-Centric Service Creation. In: *Service Oriented Computing: a look at the Inside (SOC@Inside'07) workshop*, Vienna (2007).
9. Cipolla D., Sienel J. et al.: Web Service Based Asynchronous Service Execution Environment. In: *Workshop on Telecom Service Oriented Architectures*, to appear as Springer Lecturer Notes on Computer Science, Vienna (2007).
10. Belaunde M., Falcarin P.: Realizing an MDA and SOA Marriage for the Development of Mobile Services. In: 4th European Conference on Model Driven Architecture (ECMDA08), Berlin (2008).
11. Falcarin P., Venezia C.: Communication Web Services and JAIN-SLEE Integration Challenges. in *Journal of Web Services Research* v. 5(4), IGI-Global, (2008).
12. Pautasso C., Zimmerman O., Leymann F.: Restful web services vs. big web services: making the right architectural decision. In: 17th international conference on World Wide Web (2008), 805--814.
13. Baresi L., Miraz, M.: A distributed approach for the federation of heterogeneous registries. In: 4th international conference on Service Oriented Computing (ISOC 06), (2006).
14. Tim O'Reilly. "What Is Web 2.0". O'Reilly Network, <http://www.oreillynet.com/pub/a/oreilly/tim/newa/2005/09/30/what-isweb-20.html>.



## Design for Future Internet Service Infrastructures

B. Rochwerger<sup>1</sup>, A. Galis<sup>2</sup>, D. Breitgand<sup>1</sup>, E. Levy<sup>3</sup>, J. A. Cáceres<sup>4</sup>, I. M. Llorente<sup>5</sup>, Y. Wolfsthal<sup>1</sup>, M. Wusthoff<sup>3</sup>, S. Clayman<sup>2</sup>, C. Chapman<sup>2</sup>, W. Emmerich<sup>2</sup>, E. Elmroth<sup>6</sup>, R. S. Montero<sup>5</sup>

<sup>1</sup>IBM Haifa Research Labs - Israel {rochwer, davidbr, wolfstal}@il.ibm.com

<sup>2</sup>University College London, U.K. {a.galis, s.clayman}@ee.ucl.ac.uk; {c.chapman, w.emmerich}@cs.ucl.ac.uk

<sup>3</sup>SAP Research – Israel and U.K. {eliezer.levy, mark.wusthoff}@sap.com

<sup>4</sup>Telefónica I+D - Spain {caceres@tid.es}

<sup>5</sup>Universidad Complutense de Madrid - Spain {llorente, rubensm}@dacya.ucm.es

<sup>6</sup>Umeå University- Sweden {elmroth@cs.umu.se}

**Abstract.** This paper presents current research in the design and integration of advance systems, service and management technologies into a new generation of Service Infrastructure for Future Internet of Services, which includes Service Clouds Computing. These developments are part of the FP7 RESERVOIR project and represent a creative mixture of service and network virtualisation, service computing, network and service management techniques.

**Keywords:** Service Computing, Service and Network Management, Virtualisation, Service Infrastructure

### 1. Background and Motivation

At present a number of fundamental concepts and systems, including: grid and service computing, virtualisation, networking, service and network management are being developed separately. This paper argues for the integration of such systems into a new type of Service Infrastructure for Future Internet of Services.

Virtualisation has re-emerged as a gripping method for reducing service lifecycle costs and for increasing physical resource utilization. The main idea of all virtualisation techniques is the introduction of a logical structure between the physical resources and the computational processes. Virtualisation itself takes many forms. The most commonly known form of virtualization is "System virtualisation", also referred to as server virtualisation, is the ability to run multiple heterogeneous operating systems on the same physical server. With server virtualisation a control program (commonly known as "hypervisor" or "virtual machine monitor") is run on a given hardware platform and provides an environment for one or more other computer environments (virtual machines). Each of these virtual machines, in turn, runs its respective "guest" software, typically an operating system, which runs just as if it were installed on the stand-alone hardware platform. Additional forms of virtualisation include "storage virtualisation" and "network virtualisation", which give the ability to present a logical view of the storage and network resources respectively, which is different than the underlying physical resources.



Several research efforts have investigated the use of virtualisation with grid environments, and these can be largely classified as either virtual machine management on grid or grid-like virtual machine management. When combined with grids, virtualisation technologies suffer from several shortcomings. These shortcomings include: limitations on where and when a virtual machine can run (e.g., a Xen virtual machine cannot run on a VMware hypervisor); when, where, and how a virtual machine can be relocated (e.g., relocation can take place within an IP subnet and between hosts with shared storage); limitations on the performance of the virtual machine; overly-complex administrative interfaces; lack of mechanisms to meet pre-defined SLAs; and lack of adequate security.

Grid computing [1-4] is a powerful paradigm for running ever-larger workloads and services; in grids, many heterogeneous computing, network and storage resources are connected across administrative boundaries, and service providers share and exploit the infrastructure across nodes to run their services.

The Service Cloud Computing paradigm for hosting web-based services (i.e. Amazon Elastic Compute Cloud (EC2) [5] or Google's App Engine [6]) aims to facilitate the creation of innovative internet scale services without worrying about the computational infrastructure needed to support these services. However, these new "cloud computing infrastructure providers" have a scalability problem of their own. That is, what warranties can a single hosting company give to ensure that resources will always be available? In fact, no single hosting company can create a seemingly infinite infrastructure capable of serving increasing number of online-services, each having massive amounts of users and access at all times from all locations. It is only by partnering with each other that infrastructure providers can achieve the economies of scale needed to provide a seemingly endless compute utility.

What makes cloud computing [7-10] different is that recent developments in IT such as fast adoption of virtualisation technology servers [11-13], as well as adoption of Software as a Service [14-16] as an alternative method for delivering functionality to both individuals and companies, has finally created an opportunity for a global service computing utility. The provision of Software as a Service, requires businesses to monitor the behaviour of these services.


Business Service Management (BSM) is the application of service management principles to manage the Service Levels for a business function. The purpose of BSM is to ensure that the ICT infrastructure will support the business functions by meeting requirements that are set in Service Level Agreements (SLAs). A Service Level Agreement is an agreement, or contract, between a service provider and a service consumer, in which expectations are set for the level of service to be provided by the infrastructure. The SLA is specified with respect to availability, performance, and other measurable objectives. Some of the key challenges in BSM involve SLA management.

The combination of these concepts and systems into a new computing paradigm called Service-Oriented Computing [17-19] will foster new and advanced services presented as software components exposed through network-accessible, platform and language independent interfaces. This will enable the composition of complex distributed applications from loosely coupled components. Service-Oriented Computing (SOC) carries the visionary promise of reducing software complexity, reducing costs, expediting time-to-market, improving reliability, and enhancing accessibility of consumers to both government and business services.

The paper is structured as follows: section 1 presents the motivation for a new generation of Service Infrastructure, section 2 provides the new model Service Oriented Infrastructure (SOI), section 3 describes the main functions and requirements envisaged by the SOI and section 4 concludes the paper and gives some further work.

## 2. Model for Service Oriented Infrastructure

RESERVOIR has a new and unique approach to Service Oriented Computing. In the RESERVOIR model, there is a clear separation between service providers and infrastructure providers. Service providers are the entities that understand the needs of particular business and offer service applications to address those needs. Service providers do not need to own the computational resources needed by these service applications, instead, they lease resources from an infrastructure provider. The infrastructure provider owns and leases a computing cloud, which provides the service provider with a seemingly infinite pool of computational resources. The cloud is capable of giving resources to many service providers.



**Fig. 1** - Service applications are executed by a set of VEEs

This computing cloud is made up of cooperating Reservoir Sites, which own and manage the physical infrastructure on which service applications execute. To optimise resource utilisation, the computational resources within a site are partitioned by a virtualisation layer into virtual execution environments (VEEs). The VEEs are fully isolated runtime environments that abstract away the physical characteristics of the resource and enable sharing of the physical hardware. The virtualized computational resources, alongside with the virtualisation layer and all the management enablement components, are referred to as the VEE Host.

A service is a set of software components, which work collectively to achieve a common goal. Each component of a service application executes in a dedicated VEE. The running VEEs are placed on the different VEE Hosts within the site. In some cases, it can be possible to migrate VEEs to different sites, according to automated placement policies that govern the site (see Fig. 1). The VEEs are represented by squares in this figure. The VEEs can be seen distributed across the Virtual Execution Environment Hosts (VEEHs), which make up the computing cloud. The VEEs for a particular service application may all be collocated in the same VEEH (i.e. as in service application 1), or may spread across VEEHs within the same site (i.e. as in

service application 2), or may even spread across sites (i.e. as in service application 3). As long as SLA is maintained, the service is unaware of the actual location of its VEEs.

Service providers deploy applications in the computing cloud by passing a service definition manifest to a single infrastructure provider. This manifest includes: i) a list of the functionally distinct component types that make up the application; ii) The functional requirements for each component type, characterized by a reference to a master image, which is a self contained software stack (OS, middleware, applications, data and configuration) that fully captures the functionality of that component type; iii) Component grouping instructions, which are the requirements and constraints referring to a group of heterogeneous components so that they are treated as a single allocation entity; iv) Component topology instructions, that is how the different components types are related to each other and what are their inter-dependencies; v) Capacity Requirements, that is how much memory or cpus are needed; vi) Elasticity Rules, which are set of rules that express how the application scales. These specify the capacity (resource requirements) of each application component instance, as well as the number of instances of a particular component type, and how they can be dynamically adapted to properly satisfy the requirements of the application, while at the same time minimize cost; vii) Service Level Objectives (SLOs), that is the rules to ensure the service levels for business function are maintained viii) a Monitoring Specification, which specifies which elements of the application can send data for the elasticity rules and the SLOs.




Fig. 2 – Reservoir Site Management Stack & Cross Layer Functionality

The layers of the architecture for a RESERVOIR system are presented next. The **Service Manager** is responsible for the instantiation of the service application by requesting the creation and configuration of VEEs for each service component in the manifest. In addition, the Service Manager is responsible for ensuring SLA compliance by monitoring the execution of the service applications and executing the elasticity rules. That is, adjusting the application capacity either by adding or removing service components and/or changing the resource requirements of a

particular component according to the load and measurable application behaviour. Within each Reservoir Site, the resource utilization is monitored and the placement of VEEs is constantly updated to achieve optimal utilization. Similarly, the execution of the service applications is monitored and the capacity is constantly adjusted to meet the requirements specified in the manifest. These on-going optimizations are done without human intervention by the Reservoir site management stack (see ).

The **Virtual Execution Environment Manager (VEEM)** is responsible for the placement of VEEs into VEE hosts. It receives requests from the Service Manager to both create and resize VEEs, and it also finds the best placement for these VEEs in order to satisfy a given set of constraints (set by the Service Manager). The VEEM optimizes a site total utility function, i.e., VEEM is free to place and move the VEEs anywhere, even on remote sites, as long as the placement is done within the constraints such as VEE affinity and anti-affinity, security (never place VEEs from competitors together), and cost. In addition to serving local requests, the VEEM is the component in the system that is responsible for the federation of remote sites.

The **Virtual Execution Environment Host (VEEH)** represents a virtualized resource that can host a certain type of VEEs. This abstraction is needed to enable the separation of the logical algorithmic processing of the system from the actual plumbing, i.e., VEEM issues generic commands to manage the lifecycle of VEEs, and VEEHs are responsible for translating these commands into commands specific to the virtualisation platform abstracted by each VEEH. For example one type of a VEEH can be a physical machine with the Xen hypervisor controlling it, whereas another type can be a machine with the necessary software to host Virtual Java Service Containers (VJSC). In addition to translation functionality VEEHs are responsible for adding to the virtualisation platform they encapsulate the necessary hooks to meet the advanced requirements of different use cases.

### 3. Future Internet Service Infrastructure Functions

This section presents an overview of the main functions of the Reservoir Infrastructure that are important for Future Internet Service provision. The functions were determined through the analysis of many use-cases and from the general research direction set for Service Oriented Computing, for Cloud Computing, and for Future Internet of Services. These functions, together as a set, but not necessarily per individual requirement, define what distinguishes Reservoir from earlier virtualisation technologies and what the Reservoir project brings to the Future Internet efforts.

The main functions for a Future Internet Service Infrastructure are defined as:

**Separation between Infrastructure and Services:** Reservoir as a virtualisation infrastructure for services, enforces a clear separation between service providers and infrastructure providers. A Reservoir infrastructure provider will own and manage the computational, networking, and storage resources necessary to host arbitrary service applications. The infrastructure will provide functions and management facilities, which allow dynamic mapping of service components to the physical computational, networking and storage resources. In particular, service components should be opaque

to the infrastructure providers and the service provider can deliver components that could contain virtually arbitrary software stacks.

**Extensibility:** At all layers Reservoir would support a minimum set of capabilities, yet provide for the extensibility of capabilities using a Plug-and-Pay and Unplug-and-Pay fashion. For each layer the capabilities are presented.

- **Service Manager (SM)** capabilities include: (i) request VEEs for a service; (ii) SLA repositories and management; (iii) Monitor SLA commitments at all levels; (iv) Monitor Context changes at all levels; (v) maintenance of Service related metrics; (vi) trigger and manage migration/ configuration/ contextualisation of service components as function of changes in context and/or SLA; same or multiple domains; (vii) assurance management; (viii) accounting and billing management; (ix) service life-cycle management; (x) performance management; (xi) open interfaces to service portals; (xii) open business policies framework for service and infrastructure providers relationship management (i.e. pay-per-use business model management).

- **Virtual Execution Environment Management (VEEM)** capabilities include: (i) abstraction of execution environments; (ii) supports the execution of services; (iii) Virtual Machines; (iv) service containers; (v) dynamic provisioning, supervision and re-allocation of VEEs; (vi) service-driven policy engine; (vii) open interfaces to control and monitor VEEs

- **Virtual Execution Environment Host (VEEH)** capabilities include: (i) interface with different virtualisation technologies. Reservoir would provide an abstract interface that is agnostic to the virtualisation technology; (ii) partition and management of physical resources in VEEs; (iii) open interfaces to VEEMs.

**Multi-Site Operation:** Reservoir provides the capability of sharing resources for the execution of a service application across multiple sites and multiple administrative domains (that are operated and managed by the same or different authorities). A Reservoir system is actually a federation of sites that cooperate for the optimal execution of service applications.

**Service Orientation:** Reservoir is about efficient provisioning and optimal management of service applications. We have determined that:

- a service application may be a complex entity and therefore its provisioning should be automated and streamlined.
- a service application is an infinite computation that is characterized by a workload that changes (sometimes dramatically) over time. This is in contrast to job scheduling in High-Performance Computing where resources are allocated to jobs with a finite duration.

Resources, therefore, should be allocated in reaction and in proportion to the changing workload.

**Virtualisation Technology Independence:** Reservoir should support different VEEH virtualisation technologies. Namely, Reservoir would provide an abstract interface that is agnostic to the virtualisation technology.

**Security:** Reservoir should provide a seamless, comprehensive, and flexible security scheme that operates consistently across dynamically changing layers. This embedded security shall be characterized by: (i) trustworthy operation; (ii) robustness and resilience under attack and mishap; (iii) protection and privacy of user and service

information and assets; (iv) protection and privacy of identity and location and accountability; (v) all relevant service information should stay on some specific trusted domains; (vi) confidentiality and privacy for services and data (using ciphers) will be maintained; (vii) when two or more sites are cooperating, a trust-relationship is created. This includes especially data for authentication, authorization and accounting.

**Accountability:** The mechanisms that enable accountability of the service components are encapsulated in the Utility Computing Cost Model. In particular, this should provide an efficient, reliable, and secure way to collect and manage accounting data to support the different business cases and “pay-per-use” schemes. It should also support accounting across multiple sites and manage accounting for migrating components. Accountability denotes the need for the various interfaces, that Reservoir supports, to be capable of conveying or collecting accounting information. For the purpose of accounting and billing, Reservoir provides metering functions of the use of the virtual and physical resources. The granularity of the metering must at least reflect the information need for producing bills, in accordance with the accountability requirements.

### 3.1 Virtualisation Technology Assumptions

The assumptions regarding the VEEH underlying the virtualisation technologies are made explicit so that they can be treated as requirements for providers of virtualisation technologies.

**Virtualisation Overhead:** The underlying assumption is that virtualisation adds an overhead of at most 10% to the end-to-end performance of the application. Otherwise, the use cases make no sense to begin with. The most stringent consequence is that the throughput of a virtualized setup of the application does not degrade by more than 10% compared to a native setup running on the same hardware, and all this without compromising the end user experience.

**Migration Performance:** Live migration degrades the performance of the migrated application as it consumes CPU and I/O resources while preparing for the migration (especially for stateful applications). In addition, live migration requires an actual downtime of the application during the actual migration. The application relies on complex stack that is sensitive to timeouts at different levels (e.g., network, database). The performance degradation and the downtime should not have adverse affect in terms of too many aborted user transactions.

### 3.2 VEEM & VEEH Infrastructure Requirements

This section presents the requirements for the VEE Hosts, the VEE Host Interface, VEEM and the VMI. These components are largely unaware of the service semantics associated with the components they execute and manage.

**Adaptive Resource Allocation:** Reservoir should enable dynamic changes of the physical resources allocated to a VEE in a case the VEE requires additional or different capacity (e.g. CPU capacity, Memory capacity, I/O bandwidth etc) provided that the underlying physical system is able to serve the requirements for more or



different capacity. If the resource consumption in a site is oversized, automatic downsize of the consumed resources in order to limit costs should be initiated. These dynamic changes should be transparent to the service consumer.

**Elastic Arrays of VEEs:** Reservoir should support dynamic control of the number of identical VEEs for the purpose of adapting this number according to the load, for example. The relevant service application manifest must indicate this potential elasticity. In particular, all VEEs in the array share the same master image as specified in the manifest. The dynamically launched VEE should transparently join its already running siblings in the sense of serving some of the workload. It should be possible to stop a running VEE that was defined as part of an array without disrupting the service of the overall service application. It is the responsibility of the Service Provider to implement correctly a service component that is array-enabled.

**Warm Images:** The execution of a service component may require complex preparation steps (e.g. retrieving data from the backend). Rather than doing the preparation separately for each dynamically launched component, it should be possible to create an image that creates a warm VEE with all its context (e.g., warm caches, established connections), configuration and state.

**Migration:** Reservoir supports live migration of a virtual system to another pool of physical resources (i.e. computational, networking and storage resources). The live migration is performed of service components while maintaining state (of the components itself and also any impending data exchanges). Migration could also be performed on a suspend/resume mode with minimal service disruption. The migration capability is performed transparently to the service applications, which run on the virtual system. Reservoir should support the commonly practiced migration scenarios. Namely:

- Migration of groups of VEEs in order to optimise the utilization of physical resources in order to save power and to manage power in any period.
- Migration of groups of VEEs in order to facilitate massive lifecycle and maintenance scenarios (e.g. install patch, physical resource /driver upgrades, etc).
- Completion of migration of certain components within a specific time frame.
- Request response time as observed by the end-user must not exceed a limit.
- Seamless migration, without downtime, of groups of virtual systems when physical resources or the hypervisor would require maintenance activities to be performed (e.g. install patch, physical resource upgrade, driver upgrades, etc).

The Reservoir-specific dimensions for migration are:

- Migration can be across sites.
- The grouping of migrated VEEs should reflect their membership in and the structure of the relevant service applications.
- Migration of groups of virtual systems in order to reduce/optimize the number of physical resources or save power or to better manage power in any period.

### 3.3 Migration and Elasticity Transparency

The following requirements specify that behaviour (as observed externally and between its components) of a service application does not change as a result of elastic starting and stopping and migration its components. First, the following assumptions should be stated explicitly: (1) In most cases, it can be assumed that storage can be (logically) shared between the origin and destination hosts of the migration; (2) The



service application has a built-in elasticity capability. That is, the application is capable of dynamically adding and removing components while running. Under these assumptions, the relevant requirements are:

- The networking and storage views of VEEs are kept intact when VEEs are migrated. That is, if a VEE accessed particular storage device and communicated with particular network entity, the same view is preserved in spite of migrations.
- A VEE that is cloned in a VEE array inherits the storage and networking view of the other array members.
- The public end-points of the service application are kept intact in spite of VEE migration and cloning.

**Cost-Based Optimization:** The resource allocation optimization should be driven by a configurable cost-model. The cost-model should approximate the relative anticipated latencies associated with the different allocations options. The cost model, for example, should factor the cost of cross-site migrations.

**Autonomous Local Optimizations:** A VEEM should be able to autonomously improve and optimize the utilization of the site local resources regardless of the service-level monitoring and optimization. That is the VEEM should be able to take advantage of opportunities of free local resources on its own right as long as SLAs are not violated.

**Management Interface for Standardization:** Reservoir should expose a management API that hides the details of the virtualisation technology. The goal is to standardize this API. It should be possible to compose API primitives in scripts for the streamlined automation of more complex tasks in the data centre.

**Plug-Ins:** The VEEM should expose a plug-in architecture, such that various implementations of the API can be created. Moreover, a VEEM should be able to manage hosts of different virtualisation technology using the API.

### 3.4 Service Management Requirements: Service Definition Manifest

A service application may be composed of inter-related components. It is required to completely specify the application components and setup using a declarative language in the form of a manifest. The manifest should specify, for example, all the images, the storage configuration, the database content and the relevant applicative configuration. Moreover, the application should be provisioned as a single logical unit. The manifest should support the encapsulation of various service components as images for the rapid initial provisioning as well as the rapid dynamic adaptation of service applications. Moreover, the manifest should enable the automation of provisioning and management of the service application.

**Template-Based Provisioning:** It should be possible to use the service manifest as a template for easily provisioning instances of the application. The template must allow for instance-specific components due to instance-specific configuration and customization. This requirement is important for multi-tenant scenarios where the template should be parameterized by tenant.

**Flexible Virtualisation Configurations:** A virtualisation configuration maps the components specified in the service manifest to physical hosts. The manifest should support expressing flexible virtualisation configurations of the applications in order to satisfy various performance and TCO requirements. In particular, it should be possible to specify dependencies and starting order among components. It should be possible to

specify sharing of a component by more than a single service application. For example, a DBMS service component may be shared by multiple service applications.

**Resource Consumption, Management and Enforcement:** Reservoir provides facilities for monitoring, management and enforcement of physical resource consumption. Through isolation the degree of resource consumption can be controlled including control of greedy services.

**Conflicts Resolution and Avoidance:** Service components may require certain resources from the system to be allocated statically, e.g. a certain port number. To resolve conflicts between different services the service components may be executed isolated in virtual systems while sharing physical resources.

## 4. Conclusions

RESERVOIR's research on virtualisation, service computing, network and service management both enables and unifies some of the emerging trends in Service Oriented Computing and the Future Internet. This paper presents work in progress (Reservoir project started work in February 2008) for the definition and integration of such systems into a new generation of Service Infrastructure. Full design, realisation, and evaluation of the RESERVOIR Infrastructure will be completed in the next 2 years.

## Acknowledgments

This work was undertaken in the context of the Reservoir FP7 project [20], which is partially financed by the EU. The Reservoir consortium consists of International Business Machines Haifa Research Lab (HRL), Telefónica Investigación y Desarrollo (TID), Centre d'excellence en technologies de l'Information et de la communication (CETIC), University College of London (UCL), Universidad Complutense de Madrid (UCM), Elsas Datamat (ED), Sun Microsystems (Sun), Thales, Università della Svizzera Italiana (University of Lugano) (USI), Umeå University (Umeå), SAP Research, University of Messina (UniMe), OGF.eeig

## References

- [1] "The Information Factories" - Wired Magazine, Issue 14.10, October 2006  
[http://www.wired.com/wired/archive/14.10/cloudware.html?pg=1&topic=cloudware&topic\\_set=](http://www.wired.com/wired/archive/14.10/cloudware.html?pg=1&topic=cloudware&topic_set=)
- [2] "Reflections on Cloud Computing" - Irving Wladawsky-Berger's Blog, March 2008,  
<http://blog.irvingwb.com/blog/2008/03/reflections-on.html>
- [3] "Understanding Cloud Computing" - Wallis Paul, Keystones and Rivets, February 2008,  
<http://www.keystonesandrivets.com/kar/2008/02/cloud-computing.html>
- [4] "The Big Switch – Rewiring the World from Edison To Google"- Nicholas Carr, published by W. W. Norton, January 2008
- [5] "Amazon Elastic Compute Cloud" -Amazon EC2 web site,

- <http://www.amazon.com/gp/browse.html?node=201590011>
- [6] "What is Google App Engine"  
<http://code.google.com/appengine/docs/whatisgoogleappengine.html>
  - [7] "The MAC system : The computer utility approach" - R. M. Fano, in IEEE Spectrum, vol. 2, pp. 5644, January 1965.
  - [8] "What is the Grid? A Three Point Checklist" - Ian Foster, 2002,  
<http://www-fp.mcs.anl.gov/~foster/Articles/WhatIsTheGrid.pdf>
  - [9] "The Globus Toolkit" - <http://www.globus.org/toolkit/>
  - [10] "The Open Grid Services Architecture"- <http://www.globus.org/ogsa/>
  - [11] "Server Virtualisation: Doing More with Less"- Leon Erlanger, Inforworld Report, Sept 06, [http://www.infoworld.com/article/06/09/11/37FEvirtcaseserv\\_1.html](http://www.infoworld.com/article/06/09/11/37FEvirtcaseserv_1.html)
  - [12] "Xen and the Art of Virtualisation" - P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt and A. Warfield, in Proceedings of the 19th ACM Symposium on Operating Systems Principles, Bolton Landing, NY, USA
  - [13] "The VMware" Web Site [www.vmware.org](http://www.vmware.org)
  - [14] "Turning Software into a Service" - Mark Turner, David Budgen, Pearl Brereton, Computer, vol. 36, no. 10, pp. 38-44, Oct., 2003
  - [15] "The Different Faces of IT as Service" I. Foster, S. Tuecke - [www.ggf.org/documents/Diff\\_Faces\\_foster.pdf](http://www.ggf.org/documents/Diff_Faces_foster.pdf)
  - [16] "The salesforce" web site <http://www.salesforce.com/>
  - [17] "Service Oriented Computing" –A Research Roadmap, M. Papazoglou, P. Traverso, D. Schahram, F. Leymann, B. Kraemer, Dagstuhl seminar 2006  
<http://drops.dagstuhl.de/opus/volltexte/2006/524/>
  - [18] "Service-Oriented Computing: State of the Art and Research Challenges" -Mike P. Papazoglou, Paolo Traverso, Schahram Dustdar, Frank Leymann: IEEE Computer 40(11): 38-45 (2007)
  - [19] "Service-Oriented Computing: a Research Roadmap" - Mike P. Papazoglou, Paolo Traverso, Schahram Dustdar, Frank Leymann: Int. J. Cooperative Inf. Syst. 17(2): 223-255 (2008)
  - [20] "Reservoir project" [www.reservoir-fp7.eu](http://www.reservoir-fp7.eu)

# Above the Clouds: From Grids to Service-oriented Operating Systems

Lutz SCHUBERT<sup>a,1</sup>, Alexander KIPP<sup>a</sup> and Stefan WESNER<sup>a</sup>

<sup>a</sup>HLRS – University of Stuttgart, Germany

**Abstract.** Over recent years, resource provisioning over the Internet has moved from Grid to Cloud computing. Whilst the capabilities and the ease of use have increased, uptake is still comparatively slow, in particular in the commercial context. This paper discusses a novel resource provisioning concept called Service-oriented Operating Systems and how it differs from existing approaches of Grids and Clouds. The proposed approach aims for making applications and computers more independent of the underlying hardware and increase mobility and performance. The base architecture and functionality will be detailed in this paper, as well as how such operating systems could be deployed in future workspaces.

**Keywords.** operating systems, network, grid, service oriented architecture, many core, mobile grids, distributed computing

## Introduction

The Grid concept is understood in this paper as a set of distributed resources integrated in a fashion allowing remote execution of processes and applications with different requirements towards the underlying resources, such as computational power. Next to execution of tasks, it is also often considered a space for storing vast data amounts. These types are often referred to as *computational* or *storage* Grids. Depending on whether execution / storage is supposed to be stable non-regarding changes in the Grid structure (i.e. failure of individual resource nodes), we speak of *managed* grids.

Though these do not comprise the full usage scope of Grids, it nonetheless reflects the base commonality of most use cases, namely the capability to provide (managed) resources in a fashion that they can be used remotely with little to no impact on the local execution of processes, respectively extending the capabilities of local resources.

Cloud computing has recently emerged from this movement as a means to provide in particular computational resources, even though storage clouds are gaining in popularity too. From the Grid perspective, Clouds are only passive resources in the sense of nodes in the computational or data grid case – they do not offer the enhanced capabilities of Virtual Organisations or similar, nor do they easily plug in into existing grid structures (such as GT4). Typically cloud and similar resource providers expose their own proprietary APIs which imply that the resources are used in a more manual fashion than originally envisaged by the Grid. From this point of view, clouds are only extended single Grid nodes. However they can be seen as an intermediary step to enabling dynamic outsourcing in a form that extends local resources.

---

<sup>1</sup> Corresponding Author: Lutz Schubert, HLRS – High Performance Computing Centre, University of Stuttgart. Nobelstr. 19, 70569 Stuttgart, eMail: schubert@hlrs.de

This paper will examine the current problems in resource provisioning and uptake in commerce, as well as proposing a means to overcome these problems (section 1 & 2). We will present the notion of Service Oriented Operating Systems (SOS) that integrate dynamic resource pools and execute processes across these (section 3 & 4). We will then compare this model with current related approaches (section 5) and conclude by examining the relevance of resource fabrics for the future internet.

## 1. From Grids to Clouds and their limitations

Early Grid solutions had been introduced to overcome the limitations of computing systems and integrate geographically dispersed resources into a metacomputer [19]. This concept had been further extended to allow for integration of more heterogeneous resources. Prominent realisations of this are the EGEE [20] and DEISA [21] infrastructures. While they are well perceived and used within the scientific community their uptake for business or general purpose applications remains quite limited.

The move from proprietary solutions towards Service Oriented Grids [22] [23] has increased uptake within the community as critical aspects such as commercial level security, exploitation of Web Service standards and frameworks decreased the entry level border for programmers and companies significantly and is further developed utilizing semantic technology towards Service Oriented Knowledge Utilities [24]. Grids realising cross-organisational collaborations remain still comparatively complex and competing standards for key aspects make investment in this technology risky.

Cloud computing addresses the major limitations of Grid solutions and convinces by its (seeming) simplicity. They pick up the concept of Virtual Hosting Environments [23] [2] and realise them in an efficient and easy to use fashion. In order to be commercially viable, outsourcing approaches like Clouds have to take extreme security and legal precautions, as typically the data transported to and computed on external resources is sensitive to the respective customer. The complexity of outsourced computational resources was one of the reasons to make the approach of grid- / web service based Virtual Organisations [14] so attractive for enterprises.

Existing Cloud solutions only partially address such business requirements so far and clearly attract users without specific requirements e.g. related to the confidentiality of their data. The realization of Grids over Clouds (cf. Figure 3) would enable such combined features but require additional mechanisms to bridge the gap between the Grid and the Cloud environment.

Current approaches e.g. by Amazon [5] or SalesForce' [4] suffer in particular from their proprietary APIs and their clear disjunction from local machines: whilst network storage can already be easily integrated into the user's operating system given that it is static, resources, services and applications on a network level cannot. Even though the web service approach would allow for this integration due to its standard interface approach, usage is not easy and typically comes at the cost of speed.

Web Services and Grid put the user in control of the actions and integrations – something that very few users are willing to put up with, let alone have the knowledge and capabilities to do so. For low level resources this is furthermore not even required.

Service-oriented Operating Systems treat resources as what they are: resources, i.e. they are treated by the OS in the same way as local devices. As we will show in the discussion section, even web based applications effectively similar to computational resources with very limited capabilities (namely execution of the respective process).

## 2. Resource Fabrics: Building up the Service-oriented Operating System

For a SOS, the network builds a mesh of potential resources, similar to P2P where the nodes build up the network. The main OS instance will identify remote resources and their capabilities – ranging from speed, availability etc. to type and functionality.

Such a resource mesh consists of all machines available in the network that run the base infrastructure (see below) and integrates them *virtually* into a local environment. The Service-oriented Operating System therefore treats the web environment as a *resource fabric* where individual providers are subject to constant changes.

A “resource fabric” consists of resources exposed over the internet and made available on a low Operating System level. Grid and Cloud provide and integrate their resources on higher levels, close to the actual application layer. Note that even though clouds can host virtual machines (and thus OS) they are still subject to the underlying infrastructure and require a full operating system from the user.

Implicitly, classical approaches suffer from all according deficits and usage & control are completely up to the user. Extending capabilities implies reprogramming and deployment is restricted by the middleware (cf. Figure 3, left). Infrastructure and operating system even often clash, so that additional steps are required to compensate for the “deficiencies”, ranging from security to slow message processing. In other words, efficient usage of resources on the upper layers requires in-depth expertise.

### 2.1. Virtualizing the Environment

The Service-oriented approach aims to overcome these obstacles by integrating resources right on the OS level – the fabric thus becomes generally transparent. In other words, which resources, applications and devices can be used does not only depend on the local deployment, but also on the machines available in the fabric.

Service-oriented OS require an intermediary layer that virtualises the hardware underneath it. This is already a common approach in most Operating Systems, such as the Hardware Abstraction Layer (HAL) [7], in order to reduce resource management overhead. Classical OS however treat networking on a higher layer with only the network card represented on the hardware layer, so as to account for different manufacturers. The SOS extends this concept by introducing cross network virtualisation which integrates resources across the web. The operating system therefore makes use of double blind invocation techniques (cf. [1], [2]), which renders resources and their location more independent from the calling infrastructure (Figure 1). This concept was already introduced by the IST project TrustCoM [3].

Double blind invocation acts on both sender and recipient side to enhance control, i.e. applications invoke hardware non-regarding their interface and only with respect to their functionality (printer, storage, computation etc.) – this goes beyond the scope of mere low-level devices and can circumscribe applications too. The Virtual Environment layer will identify appropriate resources and providers on basis of the required capabilities – note that the OS is *not supposed* to take higher-level aspects, such as user profiles, into consideration at this layer, as these are application-specific. During usage, interaction with the resource will either take place locally without intermediary messaging or across the network. In both cases, the recipient will not directly expose the resource, but through a reduced OS instance that grants secure access control and routing capabilities, i.e. it is completely up to the provider, which resources and services to expose or where to host them. Thus, the recipient’s infrastructure can span a

local network. This is principally identical to cloud providers such as Salesforce which expose a common interface to a dynamic and fully managed cloud farm, thus reducing the user's overhead to select and maintain appropriate resources.

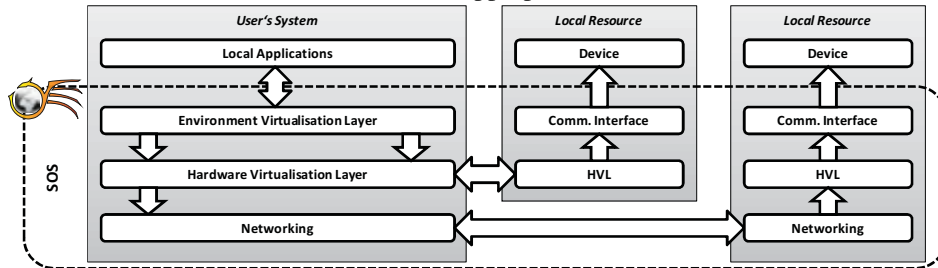


Figure 1: virtualization using double blind invocation in SOS

The communication interfaces also allow introduction of standard messaging structures, thus granting standard manufacture-independent access. With the remote instance hosting a reduced OS, the main instance can exploit the capabilities whilst the remote instance can take care of actual hardware interaction. As device drivers are the main source for Operating System crashes, this approach improves stability significantly, and reduces the risk of incompatibilities between different drivers accordingly.

## 2.2. Sub Instances

We can distinguish different types of resources and hence sub-OS instances, or micro Kernels, that can be integrated (cf. Figure 2):

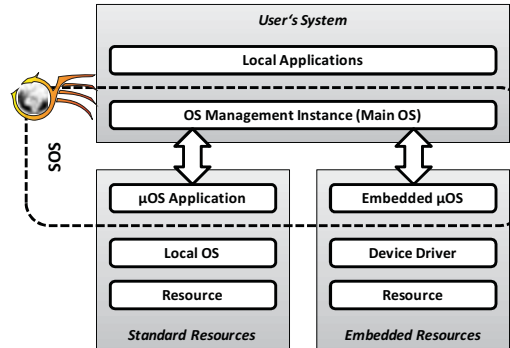


Figure 2: main kernel instances

(1) in classical, standalone mode, the OS will make use of local resources which do not require their own kernel, though this is possible with new hardware architectures and would grant additional flexibility at run-time – this however exceeds the scope of this paper and will be addressed in a future publication.

(2) dedicated resources which are only intended for usage in the SOS context which require embedded micro kernels that take over messaging and management, as well as incorporation of the driver. All access to these resources are routed via the same instance, leading to straight forward request queuing. More intelligent queues and self-management capabilities are specific to the device and may extend the kernel.

(3) resources used in multiple contexts and in particular used locally by the owner for his / her own purposes, need to provide the extended functionality *on top of* existing



operating systems, i.e. in the form of a standalone kernel wrapping as an application. The standalone kernel thereby allows designation of specific resources and applications for provisioning without requiring additional technical know-how: as the SOS operates on the code basis and not on the application level, dedicated resources of any type will become available to the main instance under the restrictions specified by the user. This allows hosting of applications equally to computational and storage resources.

### 2.3. Integrating Clouds and Grids

It is obvious that the virtualisation layer acts as a reduced Grid infrastructure incorporating classical resource management strategies. Notably, SOS does *not* provide the same support as current Grid related approaches, as these are not relevant for base resource infrastructures but add extended functionalities on top of it. Bearing the end-user in mind, a usable system should not exceed the capabilities on the cost of usability, but should realise all requirements in a stable and easy to use fashion. The system is thereby designed open enough to allow for flexible extensions on higher levels, i.e. integrating research results in the area of QoS, semantic annotation and reasoning etc.



Figure 3: layer model without (left) and with (right) Service oriented Operating Systems

Currently, usage and research of Grid technologies in particular suffer from the lack of available, performing and stable infrastructures that are easy to use. With the SOS a new resource model is announced that incorporates main features relevant for cross-internet resource integration and usage in an efficient and managed fashion.

Along the same line, Clouds typically provide managed computational resources over a common interface, and are thus considered dedicated resources by the SOS. Accordingly, the operating system can principally make use of cloud computing and storage given the right interface, i.e. the right extension of the micro kernels. This would allow easy integration and hence easy exploitation of these capabilities.

With Service-oriented Operating Systems, the classical layer model (Figure 3, left) would change to the one depicted on the right in Figure 3 – note that Clouds themselves act *on top* of the network but would be integrated below the network (on resource level) from the end-user perspective acting on SOS.

## 3. The Service Oriented Operating System Concept

Service oriented Operating Systems would execute applications and processes *across* the resource fabric – i.e. as opposed to classical distributed operating systems or common grid middleware, the SOS concept makes use of resources on code / data level.

In other words, not only are services hosted by different providers, but more importantly, the executing code is distributed across the fabric, making use of specific conditions according to requirements.

This has the particular advantage that the *usability* of the system, i.e. the look & feel is adapted to the user's needs whilst extending the capabilities of the machine beyond the local resources. The operating system thus allows for new business models, where not only access to services is granted on a per-use basis, but also where whole code segments can be shipped and executed under license restrictions and limited to usage time, payment etc. This paper will *not* investigate the business aspects though for the sake of the underlying technical details (see e.g. [1]).

In order to achieve cross fabric code execution, SOS aggregates and virtualises the whole fabric as a distributed, annotated virtual memory for the execution layer – in other words the application, process or service runs in its own virtualised environment similar to a hypervisor, yet with the difference that all execution codes basically share the same environment *across* distributed resources and thus can exchange data easily.

The virtualisation technique bases on the same double blind invocation methodology introduced above: each application is hosted in its own virtual memory environment with the Service oriented OS mapping invocations between different such execution environments according to the respective setup and requirements.

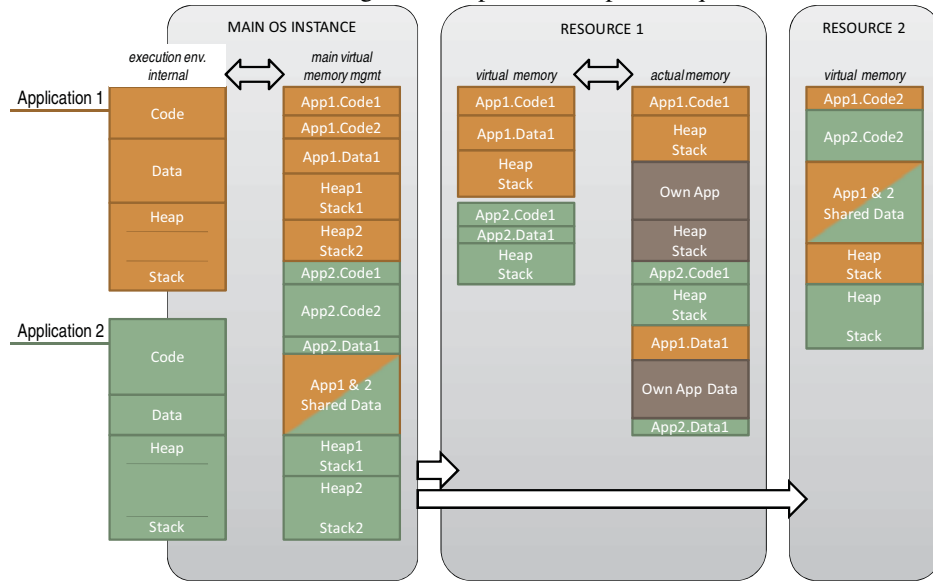


Figure 4: example distribution of applications across the resource fabric

Figure 4 depicts the virtual execution environment of each application and its actual mapping across the resource fabric: it can be noted that from the main OS instance (i.e. end-user side) all applications *appear* local, even though the actual code is distributed across resources according to execution specifics.

### 3.1. Code Segmentation

Even though processes typically perform a uniform functionality, they are nonetheless signified by a series of sub-routines and methods with individual requirements towards

infrastructure and hosting environment. One can thus e.g. distinguish between user interface and actual execution code in the background which may be idle for multiple cycles, when no specific user input requires actions. Obviously, this depends highly on the type of application, but it clearly shows that different code parts of one and the same application may have different requirements towards its hosting infrastructure. Notably, there are already commercial applications which address similar issues [8].

Service oriented Operating Systems can apply different methods to identify the code segment requirements: the principle approach bases on annotating the virtual memory in order to identify usage frequency and dependencies between segments, thus allowing assessments of couplings and basic infrastructure requirements. It also gives a rough indication of usage frequency, and thus of the user profile. Note that the accuracy of the segmentation and the according requirements increases over time, so that initial estimations may lead to performance losses. This is principally compensated through a “pessimistic” approach where requirements are initially estimated to be higher than the usage figures indicate. The details of such an annotated virtual memory management are beyond the scope of this paper and will be published separately.

Depending on these indicators, code blocks may be moved to remote low performance computers, to devices near the user with little delays in communication or somewhere in-between. The overall OS architecture thereby makes no distinction between different connection modes of the resources, i.e. the lower layers take care of the actual communication non-regarding the upper layers’ requests. All resources are assigned with additional information related to their capabilities, such as connection bandwidth, response delay, computational speed etc. Information about these resources is gathered partially through ad-hoc publication, partially through run-time evaluation.

With the application requiring fast available (i.e. highly connected) resources, local ones will be preferred over remote ones, whilst e.g. computational intense processes with little interaction requirements may select remotely hosted resources that are only accessible via the internet. In other words, the lower layer of the OS takes care of communication modes relevant for the upper execution layer.

This makes the system transparent to the resources and thus also capable of hosting many-core platforms. As will be discussed below, this does not make the system more performing with respect to individual applications, but makes better use of resources when multiple threads are executed, in particular multiple applications in parallel.

### 3.2. Code Execution

Even with Quad-Core processors being more and more common, the typical use case is signified by multiple applications running in (pseudo)parallel and not with individual applications being executed in parallel threads. Accordingly, code execution is in its essence still sequential, with a main process steering the execution of individual applications and processes on different cores.

The Service-oriented Operating System is no exception to this rule: it is not proposing new methods to automatically parallelize code, even though the execution of the *overall* process may be distributed across multiple resources within the fabric. In other words, the basic behavior is identical to a many-core platform with the addition that the actual application or process may be broken up into sub-processes and executed remotely *according to the overall sequential execution*.

As illustrated, during “normal” pseudo-parallel execution of services and applications on e.g. a quad-core PC, each core is responsible for a dedicated set of

processes between which it constantly switches according to the multitasking concepts (Figure 5, left). In the case of SOS, the applications may be distributed across the resource fabric and executed effectively sequentially, whilst each processor switches between local and remote processes like in the classical case (Figure 5, right).

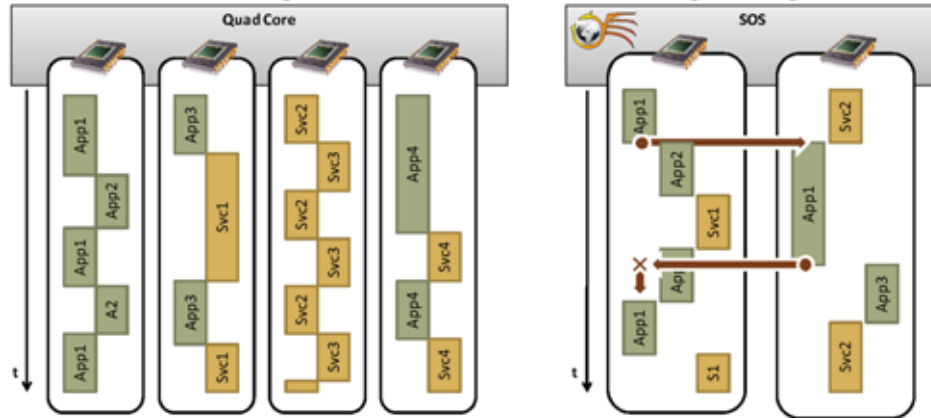


Figure 5: pseudo-parallel execution on a quad core PC (left) and execution distributed between two cores in the resource fabric (right)

One of the major problems during distributed execution does not consist in the distribution of the threads, but in handling conflicting invocations between currently active threads, i.e. maintaining the priority whilst distributing computational power between processes (cf. Figure 5, red arrows).

#### 4. Critical Issues & Solution Approaches

Obviously, there are more critical issues involved in distributed code execution in the way presented here, than “just” scheduling between processes on a single node or core. Within this section, we will briefly discuss some of the major issues. It has to be kept in mind that Service-oriented Operating Systems such as presented here are not fully realized as yet, but still in an experimental research stage. It will also be noted that some of the issues have been addressed before, though in most cases not focusing on the full scale of such a distribution and execution model.

##### 1. Stability and Reliability in Dynamic Networks

Resource fabrics are subject to potentially high volatility, since resources may become inaccessible or fail. This applies equally to network provided resources, as to nodes in a HPC cluster. Since in the SOS case, *essential* execution code may be hosted on a remote resource, losing connection to one of these resources will hence lead to the according process failing on a global level. It has to be noted though that the processes’ main instance (the “provider”) will have the full code available prior to its distribution. The information lost is related to the process *state* (or context) at a given time.

Where such resource loss can be predicted in advance, graceful shutdown is a simple means to circumvent process failure, as this allows moving the according code (or data) blocks to other OS instances in time. In many cases, resource loss cannot be predicted however and additional means need to be provided. The classical way of compensating node loss in grid-like environments consists in dedicated replicated

execution [13] or checkpointing [12]. However, checkpoints are stored only in specific intervals so as to reduce impact on execution – thus, failures may result in backtracking. In particular for short code blocks, this approach would lead to too high disruption.

SOS proposes to regard each application context change as a low-level checkpoint in which all execution information is stored in a highly available memory (RAM). Less available memories will lead to significant latencies during execution. Extended virtual memory management allows not only replicated memory across the web (with the according latency) but also asynchronous updating of memory with remote resources (without latency). The latter approach is effectively comparable to checkpointing on OS layer – however, with the Virtual Memory Manager being extensible, additional methods for delta propagation etc. can be applied to reduce the time-window between checkpoints, making them effectively identical to the actual execution timeframe.

## 2. Consistency in Distributed Shared Memory

Shared memory machines suffer from significant consistency issues comes to concurrent access to the same memory location. This is particularly true for parallel processes and applications that share memory (i.e. data). Numerous literature on this issue exists from both distributed systems perspective, as well as for concurrent thread execution and principally no further additions are needed here (see e.g. [17] [18] [25]).

Note that SOS is not a system for automatic parallelisation so that the issue is not of direct concern for the concept. Consistency is however an implicit issue, in particular for processes sharing data access, although in these cases, consistency maintenance is mostly up to the developer. There nonetheless needs to be a means to assess the state of data and to prevent essential failures due to concurrent access (read / write protection). As SOS shifts full code contexts between processing nodes, data is effectively replicated in multiple locations so that no single view exists.

The approach chosen by SOS consists in the hierarchical nature of the Virtual Memory Manager: as each executing resource hosts its own manager which obviously relates to different physical (local) addresses, the consistent view between the individual managers consists in an application specific virtual memory environment which implies a (virtual) data heap and stack. Consistency across these process specific views is maintained by the main OS' memory manager which maps all processes and memory locations across the fabric. With each execution shift, the "global" view on the memory state is updated to reference to the local view of the according executing resource, respectively considered blocked, if strong consistency protection is to be applied.

## 3. Scheduling

Scheduling is a classical issue in the area of distributed and grid computing [15] [16] – up to now, no general scheduling mechanism has been developed which is able to fulfill all relevant criteria. With SOS, this issue becomes even more critical: remote hosts may feed more than one main OS instance, so that the concurrent needs are not as easily evaluated as within a single main OS instance, in particular when the remote instance hosts a main instance for its own purposes competing over available resources.

As there is no general solution to this issue, SOS proposes the most straightforward approach, allowing more scenario specific approaches on a higher level. Effectively, this approach is basing on a fair-share scheduling discipline which takes an assigned priority as an adjustable weight into account. Priority weights are decided foremost by the resource owner and secondary by the process owner – as such, a resource owner can assign local processes with higher priorities than individual hosted

instances. Priority weighing may be applied to reduce the risk of malevolent assignments by “greedy” consumers.

#### 4. Communication Latency

Communication latency is already critical within the close proximity and high bandwidth of nodes within a single cluster, even more so does the comparatively high latency of the web affect system execution [25] [26]. However, as plenty of web applications have already shown, this impact is *perceived* differently, depending on the type of application [27]. As SOS does not primarily aim at execution of applications with high performance computing needs, but instead at extending the resource scope for average application developers and users, perceived average performance is more relevant than optimizing communication paths, which is critical in HPC processes. Notably, dedicated HPC developers will explicitly specify the requirements of the individual threads and code blocks so that SOS will not have to identify “most suiting” resources according to own assessments.

In other words, primary concerns are maintenance of user interactivity so that no latency is perceptible, and increment of the *overall* performance in comparison to local execution. *Optimization* is thereby currently no concern. Accordingly, the impact of network latency depends completely on the type of application and its distribution across the web (cf. below).

#### 5. Security

With remote instances hosting multiple processes from different main instances, security becomes a crucial concern: though the virtual memory manager creates exclusive instances, there is currently no way to protect the data from the resource owner or malicious attacks from other instances. Nonetheless, there exist numerous approaches to network security and data protection mechanisms that we do not consider this a specific concern of the operating system architecture as yet.

#### 6. Code & Memory Block Relationships

“Live” code block assessment as described in section 3.1 is accompanied with the risk of frequent code block shifts during execution when new resources become available or the assessment changes during remote execution. Also, competition with new processes may lead to constant reallocation which obviously impacts on efficiency considerably. An outstanding issue to solve consists hence in the right cut-off between the gains from data shifts versus the loss of maintaining code blocks at the same site.

### 5. Conclusions

Bringing distributed capabilities down to the operating system is not a new approach as such: MPI is a means to write distributed (parallel) applications [9], and e.g. the IP project XtremOS moves essential grid functionalities onto the OS level [10], thus increasing performance and reducing complexity. Further to this, the .NET Framework 3.5 brings essential networking and distribution features directly to the application layer [11] – future Windows versions may incorporate these features into the OS layer.

All these indicate a clear trend towards shifting the management overhead further away from the user and building up a distributed resource environment on OS / process execution level, so basically realizing a first step towards the resource fabric. However,

none of these approaches distributes actual execution code dynamically across different device types, thus exploiting the full resource fabric and in particular building up the future operating system that allows for dynamic low level resources and n-core support.

We have described in this paper how future OS models will realise a shift from application layer resource sharing to low-level resource webs, called “resource fabrics”. Such operating systems will not only shift a great amount of management overhead away from the user, thus making resource and application sharing more attractive, but it will also improve performance through reduction of messaging overheads.

These future models will enable flexible resource integration and distributed execution of code as originally envisaged by the Grid community (see e.g. [28]). Recent development in the area of Service Oriented Architectures, as well as the advances in computer and network architectures have now reached a stage where this original vision becomes feasible in a complete new fashion, namely the resource fabric.

Future operating systems building upon this model will thus significantly increase the dynamicity and stability of the main execution core, as it will become more independent of the resource fabric.

## Acknowledgments

The work presented in this paper bases on results from collaborations and evaluations of the research projects TrustCoM (<http://www.eu-trustcom.com>), CoSpaces (<http://www.cospaces.org>) and BREIN (<http://www.gridsforbusiness.eu>) and has been partly funded by the European Commission’s IST activity of the 6th Framework Programme. This paper expresses the opinions of the authors and not necessarily those of the European Commission. The European Commission is not liable for any use that may be made of the information contained in this paper.

## References

- [1] Schubert, L., & Kipp, A., 2008. Principles of Service Oriented Operating Systems. *in print*.
- [2] Haller, J., Schubert, L., & Wesner, S., 2006. Private Business Infrastructures in a VO Environment. In P. Cunningham, & M. Cunningham *Exploiting the Knowledge Economy - Issues, Applications, Case Studies*. p. 1064-1071.
- [3] Wilson, M.; Schubert, L. & Arenas, A. (2007), 'The TrustCoM Framework V4', <http://epubs.cclrc.ac.uk/work-details?w=37589>. Accessed January '08
- [4] Salesforce, 'force.com - platform as a service', <http://www.salesforce.com>. Accessed June '08
- [5] Amazon, 'Amazon Elastic Compute Cloud (Amazon EC2)', <http://aws.amazon.com/ec2>. Accessed June '08
- [6] Hinchcliffe, D. (2008), 'The Next Evolution in Web Apps: Platform-as-a-Service (PaaS)', Technical report, Hinchcliffe Company, <http://bungee-media.s3.amazonaws.com/whitepapers/hinchcliffe/hinchcliffe0408.pdf>. Accessed August '08
- [7] Englander, I. (1996), *The Architecture of Computer Hardware and Systems Software: An Information Technology Approach*, Wiley.
- [8] Moskowitz, J. (2008), 'Making SoftGrid Apps Work On the Road', *Windows IT Pro* August '08.



- [9] Gropp, W. (2000), Using MPI: Portable Parallel Programming with the Message-passing Interface, MIT Press.
- [10] Cortes, T. et al (2008), 'XtreemOS: a Vision for a Grid Operating System', Technical report, XtreemOS, Project no. IST-033576, <http://www.xtreemos.eu/publications/research-papers/xtreemos-cacm.pdf>. Accessed September '08
- [11] Microsoft, 'Microsoft .NET Framework 3.5', <http://www.microsoft.com/net/>. Accessed November '08
- [12] Elnozahy, E.; Alvisi, L.; Wang, Y. & Johnson, D. (2002), 'A survey of rollback-recovery protocols in message-passing systems', *ACM Computing Surveys* **34**(3).
- [13] Birman, K. P. & Joseph, T. A. (1989), Exploiting replication in distributed systems, in Sape Mullender, ed., 'Distributed Systems', ACM, New York, NY, USA, pp. 319-367.
- [14] Saabeel, W.; Verduijn, T.; Hagdorn, L. & Kumar, K. (2002), A Model for Virtual Organisation: A structure and Process Perspective, in 'Electronic Journal of Organizational Virtualness', pp. 1-16.
- [15] Doulamis, N.; Varvarigos, E. & Varvarigou, T. (2007), 'Fair Scheduling Algorithms in Grids', *IEEE Trans. Parallel Distrib. Syst.* **18**(11), 1630-1648.
- [16] Phan, T.; Ranganathan, K. & Sion, R. (2005), Evolving Toward the Perfect Schedule: Co-scheduling Job Assignments and Data Replication in Wide-Area Systems Using a Genetic Algorithm 'Job Scheduling Strategies for Parallel Processing', Springer Berlin / Heidelberg, pp. 173-193.
- [17] Adve, S. V. & Gharachorloo, K. (1995), 'Shared Memory Consistency Models: A Tutorial', Technical report, Rice University ECE.
- [18] Mosberger, D. (1993), 'Memory consistency models', *SIGOPS Oper. Syst. Rev.* **27**(1), 18-26.
- [19] Catlett, C. & Smarr, L. (1992), 'Metacomputing', *Communications ACM* **35**(6).
- [20] Laure, E. et al (2004), 'Middleware for the next generation Grid infrastructure', (EGEE-PUB-2004-002)
- [21] DEISA Primer, <http://www.deisa.org/files/DEISAPrimer-V1-1.pdf>, Accessed August '08
- [22] Surridge, M.; Taylor, S.; Roure, D. D. & Zaluska, E. (2005), Experiences with GRIA - Industrial Applications on a Web Services Grid, in 'E-SCIENCE '05: Proceedings of the First International Conference on e-Science and Grid Computing', IEEE Computer Society, Washington, DC, USA, pp. 98-105.
- [23] Dimitrakos, T. et al (2003), An Emerging Architecture Enabling Grid-based Application Service Provision, in '6th IEEE International Enterprise Distributed Object Computing Conference (EDOC 2003)'.
- [24] Jefferey, K. & Snelling, D. (2004), 'Next Generation Grids 2', Technical report, EC, [ftp://ftp.cordis.lu/pub/ist/docs/ngg2\\_eg\\_final.pdf](ftp://ftp.cordis.lu/pub/ist/docs/ngg2_eg_final.pdf). Accessed May '08
- [25] Delaney, D.; Ward, T. & McLoone, S. (2006), 'On Consistency and Network Latency in Distributed Interactive Applications: A Survey - Part I', *Presence: Teleoperators & Virtual Environments* **15**(2), 218-234.
- [26] Delaney, D.; Ward, T. & McLoone, S. (2006), 'On Consistency and Network Latency in Distributed Interactive Applications: A Survey - Part II', *Presence: Teleoperators & Virtual Environments* **15**(4), 465-482.
- [27] Ruddle, A.; Allison, C. & Lindsay, P. (2001), Analysing the latency of WWW applications, in 'Computer Communications and Networks', pp. 116-121.
- [28] Foster, I. & Kesselman, C. (1998), Computational grids, in 'In VECPAR', Morgan Kaufmann, pp. 15-52.

## Future Internet: towards context information brokering

M. Oskar van Deventer<sup>1</sup>, Paul Tilanus<sup>1</sup>, Mike Schenk<sup>1</sup>, Eelco Cramer<sup>1</sup> and Joost Adriaanse<sup>1</sup>

<sup>1</sup> TNO Information & Communication Technology

P.O. Box 5050, 2600 GB Delft, The Netherlands

{oskar.vandeventer, paul.tilanus, mike.schenk, eelco.cramer, joost.adriaanse}@tno.nl

**Abstract.** *Future Internet* is about combinations of communication, content and context services. Whereas the former two have achieved already a reasonable state of maturity, context information services are still at their infancy: at most stand-alone applications with limited on-line-or-off-line presence information. The critical success factor is still missing, namely context federation, which is the exchange of context information between different application, services and providers.

This article investigates how context services could be successfully federated by following the same pattern that many other information and communication services have successfully followed in the past. First *Context Information Aggregators* and later *Context Information Brokers* play a critical role in addressing the market need for federated context information services.

This article highlights challenges that have to be overcome to make this vision come true. If Europe takes the lead in overcoming these challenges, Europe can become a flourishing ground for a new context-brokering industry.

**Keywords:** Future Internet, Federation, Interconnection, Context, Presence.

### 1 Introduction

Future Internet is more than fast video communication and sharing of content. An important aspect is social networking (“Web 2.0”) and the context of the people that are networking: where are they, what are they doing, how do they feel, what do they have to share, how do they want to be reached. Ideally, your friends can easily obtain a full insight in your context, in a way that conforms your own wishes and policies while giving you a minimum of hassle, that respects your privacy, that is aggregated from multiple sources and over multiple service providers, and that is visualized with the best possible user experience for your friends. Currently, context information is rather limited and cumbersome in use. At most, your friends can see whether you are logged-in to a specific service. Moreover, they first need to log in to the same service from the same service provider to get that information.

Federation is a powerful concept in Information Society Technologies. The global coverage of telephony and Internet could only be achieved by the federation of the

networks of many telephony and Internet service providers. In order for context information to become as ubiquitously available as telephony or Internet, context federation is needed.

This article investigates how federation has developed in successful multi-operator services in the past. A pattern emerges that is subsequently applied to context services. This pattern starts with monolithic service providers, like pre-1970 telephony service providers. Stimulated by the arrival of pseudo federation providers, often identified as “parasitic” by the monolithic incumbents, the market need for federation is addressed. Over time, pseudo federation evolves into symbiotic federation and full federation, with a major role for Context Information Aggregators and later also Context Information Brokers.

This article does not intend to provide a complete solution for context federation. It is however our belief that federation at service level in general and context federation in particular is a key challenge in the development of the Future Internet. If Europe takes the lead in overcoming these challenges, it is well placed to become a flourishing ground for a new federated context-brokering industry.

Section 2 highlights the federation pattern. Section 3 applies the federation pattern to context information services. Section 4 presents the challenges associated with this vision. Section 5 concludes with a European view on the context information market.

## 2 The federation pattern, from monoliths to full federation

Service federation is an important feature of modern telecommunications. It refers to a situation where several service providers work closely together to provide a seamless service to their combined group of end-users. Prime examples of federation are international telephony and GSM roaming. In both cases, the end-user only has a formal relationship (the subscription) with its own operator, but is able to access users and resources in the domain of other operators. Technical and business agreements between all involved operators ensure that the end-users experience a seamless service, as if it were offered by one global telecom operator.

The above example is probably the most mature type of service federation possible. It has the characteristics of a political federation, as is defined below:

*“A federation (Latin: foedus, covenant) is a union comprising a number of partially self-governing states or regions united by a central (“federal”) government. In a federation, the self-governing status of the component states is typically constitutionally entrenched and may not be altered by a unilateral decision of the central government.”*

Source: Wikipedia

In the context of federation for ICT services, the “federal government” is often a standard-setting body, such as: ITU, IETF or GSMA.

## 2.1 Types of service federation

The following important types of service federation are distinguished in this paper:

- Communication service federation
- Content service federation
- Context service federation

*Communication service federation* is about the establishment of service sessions with two or multiple end-points, where each of these end-points is in a similar role and may have immediate reactions to the actions of other endpoints. Examples are phone/video calls and instant messaging.

*Content service federation* is about the delivery of content from content providers to consumers. The end-points in the established service sessions play different roles. Architectures for content federation are currently elaborated by the EUREKA CELTIC RUBENS project [1].

*Context service federation* is about the distribution of context information between providers and consumers. Context information may be (semi) static (e.g. address information) or (highly) dynamic (e.g. location, mood, availability). The issues with context federation are standardization of the various contexts, privacy management and the vast amount of data involved. Context federation can be used to enrich and link content and communication services.

These types of service federation are not mutually exclusive. A single service, offered across domains, may include several types of service federation. The types of federation are named after the group of services where they are most important.

## 2.2 The four stages of service federation

Full service federation is thought to be the most mature stage in the development of ICT services. However, when looking at the developments in ICT services other stages can be identified, some of them offering a quite acceptable substitute for federation. These stages and their characteristics are described below; see also Figure 1.





Figure 1: The four stages of service federation.

### 2.2.1 Monolithic: Service Providers competing for market share

The monolithic stage is believed to be the first stage in the development of services. Each Service Provider is striving to offer a unique service and wants a large market share. Because of the fierce competition, the providers are not concerned with working together to offer a seamless service to their combined user base.




**Figure 2: Monolithic Service Provider competing for market share.**

The most obvious current example is Instant Messaging (IM). There are several IM providers such as ICQ, Google Talk and Windows Live Messenger (WLM). They all want to achieve the highest market share and lure away end-users by creating a community feeling. An individual end-user that wants to use IM with a buddy that is connected to another IM provider will usually try to influence that buddy to “join the club” and subscribe to their IM service. As a result, end-users with buddies in several communities are forced to subscribe to several IM service providers and implement the interworking themselves (by installing several IM clients or reverting to multi-headed clients such as Pidgin<sup>1</sup>, AdiumX<sup>2</sup> or Miranda<sup>3</sup>).

### 2.2.2 Pseudo (“parasitic”) federation: disruptive innovation

The pseudo federation stage is entered when a service becomes successful, with several monolithic Service Providers offering a similar service. End-users will grow weary of the fact that they have to implement interworking themselves and are inclined to divert to any provider that offers a solution for their predicament. This is the time when the Pseudo Federation Provider arises.



**Figure 3: Pseudo-federation, disruptive third parties addressing a market need.**

<sup>1</sup> <http://www.pidgin.im/about/>

<sup>2</sup> <http://trac.adiumx.com/wiki/AboutAdium>

<sup>3</sup> <http://www.miranda-im.org/about>


Figure 3 explains the concept of pseudo (“parasitic”) federation. A Pseudo Federation Provider connects to the existing interfaces offered by each of the monolithic Service Providers and then provides a new interface towards the end-users. End-users that find interworking important will start obtaining the service from the Pseudo Federation Provider, while other users will stay with the existing Service Provider. As a result, the Pseudo Federation Provider is in direct competition with the Service Providers and is regarded as parasite by these monoliths.

In the context of Instant Messaging (IM), Palringo<sup>4</sup>, Fring<sup>5</sup>, Nimbuzz<sup>6</sup> and many others can be regarded as the first Pseudo Federation Providers. However, these have not yet gained much momentum, as apparently not enough end-users are bothered by the lack of federation in IM.

Pseudo-federation is difficult to block for the monoliths because the required technology is rather simple. They can try to make life difficult for the Pseudo Federation Provider, but they cannot go very far in this. Because the Pseudo Federation Provider uses the same interface as the regular end-users, barriers implemented here will also drive away the regular users.

### 2.2.3 Symbiotic federation: maturing of the market

The next stage is symbiotic federation. From a technical point of view, it is similar to pseudo federation. However, the monolithic Service Providers have started to accept the fact that there are Pseudo Federation Providers and they also start to see the value of those parties. Especially for niche markets they will refer their end-users towards the Symbiotic Federation Provider.



**Figure 4: Symbiotic Federation Providers becoming accepted players.**

Paypal<sup>7</sup> is an example of a Symbiotic Federation Provider. A web shop that wants to provide its users with the ability to pay with credit cards can either obtain the services from Paypal or from each of the individual credit card companies. However, with Paypal the web shop owner requires only one connection and corresponding contract instead of hooking up with each of the credit card companies. Likewise, the credit card companies are not that interested to directly deal with each individual small web shop and will refer a small starting web shop to Paypal.

<sup>4</sup> <http://www.palringo.com/features>

<sup>5</sup> [http://www.fring.com/fring\\_is/what\\_is\\_fring/](http://www.fring.com/fring_is/what_is_fring/)


<sup>6</sup> [http://www.nimbuzz.com/en/about/nimbuzz\\_for\\_you](http://www.nimbuzz.com/en/about/nimbuzz_for_you)

<sup>7</sup> <https://www.paypal-media.com/aboutus.cfm>

### 2.2.4 Full federation, towards global service coverage

Full federation is federation as we know it from telephony. It does not matter from which service provider the end-user obtains the service; it can interwork with all of the service providers resulting in global coverage for the service. There is no intermediate federation provider.

Important in full federation is standardization; each service provider offers a similar interface to both end-users as well as the other service providers.



**Figure 5 Full federation; integrating the Service Provider and Federation Provider roles.**

## 3 Applying the pattern to context services: towards brokering

As described, the federation pattern is market-driven. It is the user that ultimately requires federation. Regulators often play a key role in this process.

This section applies the federation pattern to context information services. Figure 6 sketches the exchange of context information at a conceptual level. Following the IETF terminology [2], the “presentity” is the identity to which the presence or context information relates. In its most basic form, a presentity would be the identity of a device or application, e.g. a mobile phone or PC being on-line or off-line. In a more advanced form, the presentity could refer to a person with multiple devices and accessing multiple applications simultaneously. A watcher is the identity that requests and consumes presence or context information.



**Figure 6: Conceptual representation of a context information service.**

Figure 7 applies the four stages of the federation pattern to context services. The following business roles are distinguished.

- CIP: Context Information Provider. This is the Service Provider for whom at least some form of context information (e.g. simple on-line-off-line presence information) is part of the service offering.
- CIA: Context Information Aggregator. This is the Pseudo Federation Provider whose core business is the aggregation of context information aggregation.



- CIB: Context Information Broker. This is the Symbiotic Federation Provider whose core business is the brokering of context information aggregation. The following subsections explain the model of Figure 7 in more detail.




Figure 7: Applying the pattern to context services.

### 3.1 Monolith: Context Information Providers have no interest in federation

Currently, presence is not a stand-alone service. Context information is usually delivered as part of a larger service, like instant messaging and voice/video communication services. Examples of such communication services are Skype<sup>8</sup>, Windows Live Messenger<sup>9</sup> and Google Talk<sup>10</sup>. Users can see whether other users of the same service are on-line or not, and that is about it.

The context information is available for users to see whether other users are available for communication through the particular service. These services are typically paid by advertisements [3]. The main driver for players in this market is to attract more customers, getting a larger market share and obtain more advertisement revenues, see Figure 8. Notice that context information is not the core business of parties acting as Context Information Provider.

Even though it would be more practical for users to use a single service to communicate with any other users, the advertisement revenues of the service providers provides no incentives to offer such communication interconnection, let alone the sharing of context information.




Figure 8: Context Information Provider's main driver is market share.

<sup>8</sup> <http://about.skype.com/>

<sup>9</sup> [http://en.wikipedia.org/wiki/Windows\\_Live\\_Messenger](http://en.wikipedia.org/wiki/Windows_Live_Messenger)


<sup>10</sup> <http://www.google.com/talk/about.html>

### 3.2 Context Information Aggregators fill in an upcoming the market demand

Context information aggregation is a first step to address the market demand for context information services. Two types of presence aggregation can be distinguished: watcher aggregation and presentity aggregation:

- *Watcher aggregation* is the aggregation of presence information from multiple sources offered/tailored to watchers. Watchers can attribute different sources of presence information to a single presentity, and watch the aggregated presence information through a single application.
- *Presentity aggregation* is the aggregation of presence information from multiple sources about a single presentity, which is typically a person in the case of presentity aggregation. The presence information could be about availability to communicate through different channels (fixed phone, mobile phone, instant messaging, SMS), about the location/speed of the presentity or about the activities of the presentity. It depends on the configured policies what information specific watchers may see and at what level of aggregation and detail. Presentity federation is complex, as it requires identity federation [4].

As monolithic Context Information Providers have no interest in federation, this offers a business opportunity to external Context Information Aggregators to deliver context information aggregation services, see Figure 9. Initially, such a service would focus on watcher aggregation, as it is easier to implement by a third party. The Jabber network<sup>11</sup> can be seen as an example of a Context Information Aggregator. Context Information Providers perceive Context Information Aggregators as “parasitic” as they use the Context Information Provider’s capabilities without respecting the Context Information Provider’s business model.



**Figure 9: Market demand for aggregation addressed by third parties, Context Information Aggregators.**

### 3.3 Symbiotic federation: a role for Context Information Brokers


Symbiotic federation is a next step to address the market demand for context information services. This will induce changes in the business model for context information. Currently this information is “for free”, paid for through advertising following the classical Internet business model. Over time, other business models will

<sup>11</sup> <http://www.jabber.org>

arise. Watchers and/or presentities may recognize the value of good context information and become willing to pay for the publication and/or reception of this information. Alternatively, the costs related to context information may be subsidized by the paid communication services that they induce in a similar way that “free” SS7 signaling is being paid for by phone calls. In this alternative, “presence” has become tomorrow’s dial tone [5][6].

With the changing business models, the relationship between “parasitic” Context Information Aggregators and the “incumbent” Context Information Providers will evolve into a more symbiotic one. Upcoming Context Information Brokers would collect and aggregate context information on behalf of presentities, enabling users to provide a self-controlled presentity-aggregated view to watchers.

Context Information Brokers would play a complementary role to the Context Information Providers in a rich context-information ecosystem, with the Context Information Providers focusing on communication services and providing context information to Context Information Brokers with context information, see Figure 10.




**Figure 10: Context Information Brokers aggregating and brokering context information in association with Context Information Providers.**

### 3.4 Full interconnection-based context federation.

Full interconnection-based context federation would in theory be the final step in the evolution of presence into context services. This step assumes the arrival of pure Context Information Providers to whom context information services is core business. Context Information Providers use direct interconnection to share context information with other Context Information Providers, see Figure 11.

We consider such a scenario unlikely, as presence and context services are typically supplementary to communication services. So, even if the communication services themselves are interconnected, the brokering of context information may remain the field of specialized Context Information Brokers.



**Figure 11: Full federation for Context Information Providers.**

## 4 Challenges

This section presents the challenges that need to be tackled to turn the vision of the last section into reality. These topics should be addressed in a research program focusing on context federation in the Future Internet.

### 4.1 Current developments on service federation

The traditional telecom operators have applied a closed and central federation model for many years. With the newly evolved VoIP services, we see new and open federation models being applied once VoIP islands need to be interconnected, mainly because of the different business models applied (see for example the IETF SPEERMINT initiative [7]). As the requirements of specific communication services differ, there is also a need for different federation models. Although still unclear, it is expected that different communication federation models will coexist [8].

A currently important subject for context service federation is performance. Although current protocols standards are sufficient to support small-scale presence and messaging environments, support of large scale and federated environments is lacking. Even in new IETF drafts on the optimization of SIP as presence and messaging protocols, we see doubt on applicability of the current SIP protocol in large multi domain implementations [9]. Possible solutions are proposed where the client uses a different protocol than the central servers [10] or modifications are made to SIP [11].


Furthermore we see that federation of the identity management context service, implemented over different organisations, is becoming more and more important [12].

### 4.2 Business role and model challenges

Business roles and changing business models are a major challenge for context information brokering. The emergence of lucrative context information business will depend on successful business models, either evolutionary or disruptive. As is stated in [13] federation is the next best thing as seen from a user's perspective but there is often no incentive for the business parties involved to give up their direct link with the end-user. While the referred paper focuses on identity federation, it is our belief that similar issues will arise with any kind of federation in the Future Internet. Current services in the Internet are provided by monoliths with business revenues based on the number of end-users they serve. They are unlikely to give up their direct user contact without a fight.

Therefore, it is likely that federation will be slowly introduced and an analysis should be made of the characteristics of the business roles in the different market evolution stages. It should be studied to what extent the different forms of presence aggregation (presentity aggregation versus watcher aggregation) relate to similar or

distinct business roles. It should also be studied how business roles could gradually evolve and which business roles can only emerge in a disruptive way, see Figure 12.



**Figure 12: Elaboration of business roles and associated technical interfaces**

Also business models should be further elaborated, starting with a value analysis. In an advertisement-based business model the value is in the watched advertisements, and the focus is on stimulating the use of the service. In a watcher aggregation model, the value is in the context information itself and the aggregation of this information, which implies that the model should fuel on revenue streams related to watchers. In a presentity aggregation model, it is the presentity that uses context information to publish his reachability for communication services, which implies that the model may be fueled from communication services revenues (“dial-tone of the future”).

Finally, migration and evolution of business models should be studied. For instance, commercial television channels used to be paid from advertisement revenues and broadcast networks would share advertisement revenues with television stations.

### 4.3 Architecture and traffic management challenges

Defining technical solutions for aggregation services and context information brokering implies the design of one or more architectures.

Traffic analyses will form an important input to the architecture design. An analysis should be made on the flow of context information, including realistic estimates on amounts of context information traffic, aggregated and not-aggregated and traffic matrices. Different traffic scenarios may be worked out depending on the use of the context information. For example, a semi-continuous stream of context information may cause serious congestion in mobile networks.

Based on traffic analyses, recommendations can be made on a high-level architecture. For instance, a client-server-based architecture would aggregate context information in servers, which has the advantage of central processing, limited context information traffic and simpler (hence cheaper) end terminals. On the other hand, a peer-to-peer-based architecture would limit the centralized operation to a minimum,

most likely only group list management, and have the context information be exchanged directly between end devices without any direct interference of the context information broker.

One or more functional architectures should be worked out in terms of services supported, elementary functions, functional elements, data model, reference points/interfaces, protocol flows and protocols used. Part of the analysis is the question which interfaces require standardization, either from a multi-vendor perspective of the technical interfaces between different business roles. Reuse of existing standards, possible modifications and profiling, and making contributions to standardization bodies are part of this challenge.

#### **4.4 Privacy protection and policy management challenges**

Privacy protection and policy management are essential aspects of any context information ecosystem. An end user should be able to rely that his precious context information is treated confidentially and only be disclosed to selected watchers.

A first step is a thorough privacy threat analysis. An evaluation should be made which different types of context information there are and what levels of sensitivity should be distinguished. For example, simple on-line-or-off-line information would in most cases be much less sensitive than accurate GPS location information, let alone health sensor information. One analysis would be primarily from the presentity point of view: who is entitled to watch which context information, which person is currently actively watching. Another analysis would be from the service provider perspective, which does not want its user information be “data-mined” by potential competitors. Finally, an analysis from a context information broker perspective is applicable, focusing establishing trust as key value, similar to the role of e.g. a bank.

Directly related to privacy protection is policy management. A user wants to control which other users can watch what types of information. For example, direct colleagues may access most information in one’s electronic agenda, whereas customers would be restricted to aggregated in-office-or-out-of-office information. Such policies are typically coded in XML [15]. An analysis should be made on the types of policies required from the end-user perspective. Following the “80/20” rule, a limited number of “default policies” should be defined.

### **5 Conclusion: Europe should become a flourishing ground for a new context-brokering industry**

This article explains the four-stage federation pattern associated with the maturing of communication, content and context services. By applying the pattern to Context Information services, it is shown that the market for context information is still at its infancy, and that there are major business opportunities for Context Information Aggregators and Context Information Brokers.

In the past, Europe has shown its force in the federation of mobile communication services. Stimulated by European collaborative projects and European-scale

standardization, technologies like GSM and UMTS have reached maturity. Federation has been the critical success factor here.

The market, architectures, technology and standards for context information services are at a similar infant stage as GSM and UMTS respectively in the early 1980's and 1990's. Europe has the opportunity to take the lead in overcoming both technical and business model challenges. If Europe takes this lead, then Europe can become a flourishing ground for a new context-brokering industry.

## References

- [1] EUREKA CELTIC RUBENS, Rethinking the Use of Broadband access for Experience-optimized Networks and Services  
<http://www.celtic-initiative.org/~pub/Project-leaflets/Webquality/rubens-lq.pdf>
- [2] IETF RFC 2778, A Model for Presence and Instant Messaging; Day, et al, February 2000.  
<http://www.ietf.org/rfc/rfc2778.txt>
- [3] Business models on the web; Professor Michael Rappa.  
<http://digitalenterprise.org/models/models.html> - Advertising
- [4] Liberty Alliance  
<http://www.projectliberty.org/liberty/about>
- [5] Presence in AIM: Evolution of the Online Dial Tone; Kevin Farnham, November 2006.  
[http://dev.aol.com/blog/kevinfarnham/2006/11/presence\\_in\\_aim](http://dev.aol.com/blog/kevinfarnham/2006/11/presence_in_aim)
- [6] Presence: The Dial Tone for Internet Communications; Peter Saint-Andre, Director of Standards Jabber, March 2008.  
<http://ecomconf.com/2008/xmpp-peter-andre.php>
- [7] IETF SPEERMINT draft-ietf-speermint-architecture-07; R. Penno, D. Malas, S. Khan, A. Uzelac, M. Hammer, 03-11-2008.  
<http://www.ietf.org/internet-drafts/draft-ietf-speermint-architecture-07.txt>
- [8] Analysis of VoIP interconnection evolution; Mika Lahti, Helsinki University of Technology, Master's Thesis, February 14, 2008.  
<http://www.tml.tkk.fi/~anttiyj/Lahti-VoIP.pdf>
- [9] IETF SIP/SIMPLE draft draft-houri-simple-interdomain-scaling-optimizations-00.txt; A. Hour, V. Singh, H. Schulzrinne, S. Parameswar, E. Aoki, 01-07-2007  
<http://tools.ietf.org/html/draft-houri-simple-interdomain-scaling-optimizations-00>
- [10] Presence Interdomain Scaling Analysis for SIP/SIMPLE; A. Hour et al, October 2008.  
<http://tools.ietf.org/html/draft-ietf-simple-interdomain-scaling-analysis-05>
- [11] IETF SIP/SIMPLE draft draft-ietf-simple-view-sharing-02; J. Rosenberg, S. Donovan, K. McMurphy, 03-11-2008.  
<http://www.ietf.org/internet-drafts/draft-ietf-simple-view-sharing-02.txt>
- [12] Identities federation and Identity Providers; Miska Laakkonen, Helsinki University of Technology, 2008.  
<http://www.tml.tkk.fi/Opinnot/T-109.7510/2008/fed.pdf>
- [13] Identifying Patterns of Federation Adoption; Heather Hinton and Mark Vandenwauver in ISSE 2006 – Securing Electronic Business Processes, Vieweg 2006.
- [14] Interdomain Presence Scaling Analysis for the Extensible Messaging and Presence Protocol; P. Saint-Andre, January 2008.  
<http://xmpp.org/internet-drafts/draft-saintandre-xmpp-presence-analysis-03.html>
- [15] Presence is as presence does (sheet no. 17); Robert Sparks.  
<http://isoc.nl/activ/2008-FutureOfPresence/RobertSparks-FutureOfPresence.pps>



## **S-Cube: Addressing Multidisciplinary Research Challenges for the Internet of Services\***

Elisabetta Di Nitto<sup>1</sup>, Dimka Karastoyanova<sup>2</sup>, Andreas Metzger<sup>3</sup>, Michael Parkin<sup>4</sup>,  
Marco Pistore<sup>5</sup>, Klaus Pohl<sup>6</sup>, Fabrizio Silvestri<sup>7</sup>, Willem-Jan Van den Heuvel<sup>8</sup>

<sup>1</sup> Politecnico di Milano, 20133 Milano, Italy, dinitto@elet.polimi.it

<sup>2</sup> University of Stuttgart, 70569 Stuttgart, Germany, karastoyanova@iaas.uni-stuttgart.de

<sup>3</sup> University of Duisburg-Essen, 45117 Essen, Germany, andreas.metzger@sse.uni-due.de

<sup>4</sup> Universiteit van Tilburg, 5037 AB Tilburg, Netherlands, m.s.parkin@uvt.nl

<sup>5</sup> Fondazione Bruno Kessler, 38100 Trento, Italy, pistore@fbk.eu

<sup>6</sup> University of Duisburg-Essen, 45117 Essen, Germany, klaus.pohl@sse.uni-due.de

<sup>7</sup> Consiglio Nazionale delle Ricerche, 00185 Roma, Italy, fabrizio.silvestri@isti.cnr.it

<sup>8</sup> Universiteit van Tilburg, 5037 AB Tilburg, Netherlands, W.J.A.M.vdnHeuvel@uvt.nl

**Abstract.** The Service Oriented Architecture (SOA) is increasingly adopted by industry as a paradigm for building distributed software applications. Yet, the SOA has currently several serious limitations and many crucial service issues are not addressed, including, for example, how to establish, monitor and enforce quality in an end-to-end fashion, as well as how to build service-based applications that proactively adapt to dynamically changing requirements and context conditions. This paper provides an overview of the service research challenges identified in S-Cube, the European Network of Excellence on Software Services and Systems. S-Cube strives to address those challenges by bringing together researchers from leading research institutions across diverse disciplines. The S-Cube researchers are joining their competences to develop foundations and theories, as well as novel mechanisms, techniques and methods for service-based applications, thereby enabling the future Internet of Services.

**Keywords:** Service-based Applications, Service Oriented Architecture, Engineering, Design, Adaptation, Monitoring, Quality

### **1 Motivation**

Software services are self-contained, platform-agnostic computational elements, which can be flexibly and dynamically composed to create complex service-based applications. The functionality provided by a service ranges from answering simple requests to executing sophisticated processes requiring peer-to-peer relationships between multiple service consumers and providers. For the service consumer, a software service represents functionality that can be invoked through the service interface. The actual software that implements this functionality is executed, maintained

---

\* The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement 215483 (S-Cube).

and owned by the service provider. Thus, software services take the concept of ownership to the extreme: Not only the development, quality assurance, and maintenance of the software is under the control of third parties, but the software can even be executed and managed by third parties [2].

The Service Oriented Architecture (SOA) is increasingly adopted by industry as a paradigm for building distributed service-based applications [3][6][7]. According to IT analyst Forrester Research, 67% of the largest enterprises were using SOA-based implementations by the end of 2006 and nearly 70% of those indicated that they intended to increase their use of it [1]. These facts make services technology of paramount importance to the European software and telecommunications industry.

Currently, the common practice for developing service-based applications (SBAs) following the SOA paradigm distinguishes between three functional layers [2]:

- **Service infrastructure:** This layer supports describing, publishing and discovering services and provides the run-time environment for the execution of SBAs. It provides core functionalities for service communication (e.g., SOAP), service description (e.g., WSDL), as well as capabilities for service discovery (e.g., UDDI).
- **Service composition and coordination:** This layer supports the aggregation of multiple (individual) services into service compositions (e.g., using BPEL). Service compositions can in turn be offered to service clients, used in further service compositions and eventually be composed to service-based applications.
- **Business process management (BPM):** This layer provides end-to-end visibility and control over all parts of a long-lived, multi-step business process that spans multiple organizations and can involve human actors. BPM provides mechanisms for expressing, understanding, representing and managing an organization in terms of a collection of business processes realized in a service-oriented fashion.

When setting out to build innovative software services and service-based applications of the future, relying on the current layers of the SOA will not suffice. In this paper we elaborate on the issues that are still unsolved and outline the importance of interdisciplinary research to address them. Consequently, this paper provides an overview of the key challenges in Section 2. Then, Section 3 motivates the need for interdisciplinary research, and how S-Cube – the European Network of Excellence on Software, Services and Systems – addresses this need. Section 4 introduces and illustrates the S-Cube research framework. Section 5 concludes the paper.

## 2 Research Challenges for the Internet of Services

As has been observed in [11][10], many important challenges for building future service-based applications are still to be resolved. For the key areas shown in Figure 1, those challenges are summarized below.

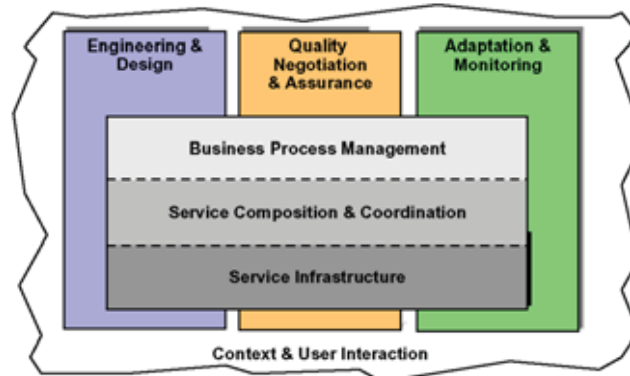
### 2.1 Engineering and Design

Designing service-based applications shows some peculiarities. Such applications are built by composing services which may be already built and running when the appli-

cation is deployed. This enables a bottom-up development approach based on the idea of searching for services and building around those that are identified as suitable.

Research is thus required in high-level declarative language concepts for the specification of services that allow lay and experienced users and other stakeholders to express their views and requests in terms of “what is needed” rather than “how it is needed”. One direction which could be followed is expressing the requests of the stakeholders at the intentional level, i.e., as high-level business requirements or goals.

Services that compose a service-based application may not be under the control of the organization that is operating the application. This results in the need for defining proper service contracts (such as SLAs) and quality assurance techniques (cf. Section 2.3). Additionally, it must be planned for runtime self-adaptation of the application (cf. Section 2.2) in the case component services become unresponsive or show behaviours that are not acceptable for the application. In general, we say that service-based applications need to be designed to be ready for adaptation. In fact, adaptation cannot be completely specified in advance due to the incomplete knowledge about the interacting parties as well as the application context. In the literature, there are some approaches focusing on design for adaptation. However, a coherent approach that integrates the initiatives of various areas (requirements, design, testing, deployment and operation) to create a proper life-cycle is still to come.



**Fig. 1.** Areas relevant for service-based applications

## 2.2 Adaptation and Monitoring

SBA run in dynamic business environments and address constantly evolving requirements. These applications should hence be able to adequately identify, and react to various changes in the business requirements and application context. These challenges make monitoring and adaptation key elements of modern SBA functionality.

With respect to adaptation and monitoring, the state-of-the-art approaches (e.g., see [8]) need to be extended in several directions. Firstly, a broader perspective is needed on *what*, *how*, and *when* we may monitor and adapt to accommodate to changes and deviations. For instance, a deviation detected through run-time monitoring of a single specific execution of a SBA may trigger an adaptation that can be achieved through the (automated or semi-automated) evolution of the whole application. Furthermore,

adaptation decisions, applied for instance by the user or maintainer of the SBA, can be learnt and transformed into adaptation patterns which can then be exploited to simplify and drive these decisions in the future.

Secondly, the definition of monitoring and adaptation itself must be extended: To this end, *monitoring* should subsume all the techniques and tools that allow for identifying, detecting, and even predicting critical events and situations (cf. Section 2.3). In this way, for instance, online testing techniques can be exploited as monitoring tools if they are used for predicting possible execution problems [5]. The same holds for *adaptation*: all the facilities for modifying the application regardless the timing and the effect count for the adaptation problem.

Thirdly, various research disciplines, different application domains (ranging from B2B to user-centric systems), as well as different functional SBA layers need to be considered during adaptation and monitoring. On the one hand, this allows reusing ideas and approaches from existing fields; e.g., using data and process mining techniques for post-mortem business process monitoring in order to gather information about SBA evolution relevant for the adaptation of the latter. On the other hand, only such an integration makes the cross-layer adaptation and monitoring possible in the first place, providing ways to reveal and accommodate to the changes in those elements of the SBA architecture that have an impact on the other layers. Indeed, this is not possible in the current SOA approaches, where the monitoring and adaptation facilities at different layers are considered in isolation.

### 2.3 Quality Negotiation and Assurance

To provide the desired end-to-end quality of globally distributed service-based applications, the dynamic agreement and assurance of quality becomes a key issue. This requires that not only quality aspects are negotiated and agreed, but also that those are checked during run-time. In a service-based application, different kinds of quality attributes are important [12]: Quality of Service (QoS; e.g., performance, availability), Quality of Experience (QoE; e.g., usability and trust), Quality of Business (QoBiz; e.g., revenue, profit), and Quality of Information (QoI; e.g., accuracy, completeness, relevancy). There is thus a strong need for methods that address quality attributes in a comprehensive and cross-cutting fashion across all layers of a service-based application. Specifically, end-to-end quality provision implies that the dependency between different kinds of quality attributes must be understood. For instance, the interrelation between the fulfilment of different QoI attributes on the infrastructure layer, the satisfaction of QoE on the service composition layer and the achievement of business value (QoBiz) at the BPM layer (cf. Section 2.4) is an open issue.

Further, to address dynamic adaptations of service-based applications, a growing need for automating the negotiation of quality attributes (e.g., stipulated by SLAs) can be observed. However, this issue requires considering user interaction and experience issues that may impact on the negotiation itself. This aspect calls for a multidisciplinary effort in which technology researchers will interact with researchers addressing user interaction issues.

Given the change of the life-cycle for service-based applications (cf. Section 2.1), quality assurance techniques that can be applied at run-time become essential. There-

fore, standard and consolidated “off-line” software quality assurance techniques (like testing and analysis) need to be extended to be applicable while the application operates (“online techniques”).

Finally, to support the vision of pro-active adaptation (cf. Section 2.2), novel quality prediction techniques need to be devised. Depending on the kind of quality attribute to be predicted, these can range from ones that built on traditional techniques to ones that exploit modern technologies of the Future Internet. As an example for the first case, “correctness” or “performance” (QoS) could be predicted by building on techniques similar to online testing [5] or run-time model analysis [4]. As an example for the latter case, “usability” of services (QoE) could be predicted by extending existing principles of reputation systems.

## 2.4 Business Process Management (BPM)

Business Process Management (BPM) is the activity associated with modelling, designing, deploying, monitoring and managing information technology aligned to meet the goals of an organisation and its customers [9]. BPM provides entire life-cycle management for multiple business processes that together contribute to the success of a business. Thus, from the BPM perspective of the service network described above, there is a need to define the activities that achieve business goals like lowering costs whilst increasing market share, profits and customer satisfaction.

Currently, there is a gap between business strategy, BPM and business models and their implementation in SBAs. Therefore, the objective of the BPM research area in S-Cube will be to develop fundamental new concepts in service engineering that can drive service implementation, configuration, operation and management from business models and their goals. This requires investigation into, for example, new process languages to enable the reuse of existing service compositions, choreographies, communication and service interaction patterns, the mechanisms of business transactions, collaboration and decision-making within service networks and the verification and demonstration of the compatibility of business process orchestration with respect to compliance with regulation.

As shown in Figure 1, BPM sits above the service composition and co-ordination layer (cf. Section 2.5) that provides functions exposed as services for use in business processes. Thus, integral to this research will be the investigation of how unanticipated changes in the service composition and co-ordination will be dealt in an agile, automated and transparent manner with ‘new generation’ BPM that provides business activity monitoring (BAM) through the measurement of KPIs and business critical events (cf. Section 2.2).

In summary, and to paraphrase [12], BPM is a natural complement to the techniques of service composition and co-ordination and a mechanism through which an organisation or business can apply and utilize service networks to achieve business goals. S-Cube plans to bring the often-fragmented research of these two areas together through the investigation of mechanisms and models that correlate KPIs with SLAs and business processes (cf. Section 2.3).

## 2.5 Service Composition and Coordination

Current research results in the field of service composition are designed in isolation from the BPM layer and the service infrastructure layer. While such an approach reduces complexity by abstracting away details from these layers, it does not sufficiently tackle all problems that need to be addressed in a concrete application domain. Therefore, we observe a gap between the requirements of the BPM layer (cf. Section 2.4) and the service compositions that implement them, in particular with respect to the key performance indicators specified on the BPM layer; i.e., the service compositions are not designed such that they can guarantee the desired KPI values.

Additionally, the KPIs on the BPM layer may evolve over time, which needs to be propagated to the service composition layer. Due to the separation of research, this adaptation on the BPM layer currently cannot be propagated to the service compositions, and moreover, the service compositions cannot adapt themselves to meet the modified requirements from the BPM layer. The service compositions require additional support from the service infrastructure, in particular in terms of discovery and selection of services complying to the overall quality attributes of the service composition and not only with the quality requirements of individual tasks. Therefore, we identify the need for the creation of service and service composition models involving quality characteristics and behavioral features. These models will reflect the inherent relationship among the BPM layer and the service compositions.

Based on the models and languages for service compositions, mechanism for service composition adaptation are needed, which are driven by quality attributes and by the requirements of the BPM layer and which are influenced by the service infrastructure. Such mechanisms will inevitably influence the service composition models, i.e., the mechanisms will be supported by the models for service compositions and enabled by corresponding language elements. The mechanism will enable adaptation mechanisms (cf. Section 2.2) which will be identified as necessary for SBAs and which will depend on the technology used to implement the service composition models. For example, for process-based compositions, such adaptations may be realized in terms of control flow changes (i.e., deletion of tasks, inclusion of tasks, etc.).

To enable the monitoring of service compositions (cf. Section 2.2), an event model for event notifications is expected to provide information related to the execution status of individual tasks and about the quality attributes.

## 2.6 Service Infrastructure

Service infrastructures will need to be scalable and of high performance in order to support for the execution of future service-based applications. Traditional infrastructures have been thought, mainly, to support enterprise applications. This idea has to be extended in order to support the execution of large-scale multi-enterprise service-based applications, which form complex service networks (cf. Section 2.4). This will require the effort of diverse communities like: high performance computing, grid computing, service oriented computing, cloud computing, etc.

In particular from grid computing, the main contributions are expected in the area of self-\* infrastructures. Self-\* includes self-healing, self-optimizing, and self-protecting [8]. Those self-\* properties, in fact, have to be enforced both at a local and

a global level. At a local level, self-\* capabilities allow services to react to sudden changes in the infrastructure state. At a global level, self-\* mechanisms trigger changes that will be propagated to the application.

Future infrastructures have to support effective and efficient service discovery through service registries, which could exploit novel mechanisms of peer to peer architectures. Those have shown – in other contexts (e.g., file sharing systems) – to be a good choice in case of highly dynamic environments. Also, historical information about how services have performed (cf. Section 2.3) could be used to improve the effectiveness of service registries. In response to a query for a service, QoE factors can be taken into account to select the (set of) best service(s) to propose for being included in the application.

Furthermore, novel SOA infrastructures should be designed to include services that are offered through the Internet via Web 2.0.

### 3 Multi-disciplinary Research in S-Cube

Section 2 has highlighted that many service research activities are fragmented and, as a result, each research community concentrates mostly on its own specific research techniques, mechanisms and methodologies. Thus, the proposed solutions are not aligned with or influenced by activities in related research fields.

In order to address the challenges introduced above, a holistic view and approach to services research is thus required. To this end, S-Cube, the European Network of Excellence on Software and Services ([www.s-cube-network.eu](http://www.s-cube-network.eu)), aims to establish a unified, multidisciplinary, vibrant research community. S-Cube is funded for a period of four years by the European Community's 7<sup>th</sup> Framework Programme. In S-Cube, over 70 researchers and over 50 Ph.D. students from 16 institutions, pursue the following objectives:

- Defining a broader research vision and perspective to shape the software-service based Internet of the future.
- Re-aligning, re-shaping and integrating research agendas of key European players from diverse research communities to achieve a long-lasting foundation for steering research and for achieving innovation at the highest level.
- Inaugurating a Europe-wide program of education and training for researchers and industry to create a common culture and impact on the future of the field.
- Establishing a proactive mobility plan to enable cross-fertilisation between research communities.
- Establishing trust relationships with industry (e.g., via NESSI) to strengthen Europe's industrial competitiveness.

To reach the above objectives, S-Cube members jointly carry out the following activities:

- **Integration Activities:** Integration activities tackle fragmentation and isolation of research by different means: (1) The *S-Cube Knowledge Model* will capture terminology and competences of S-Cube members and their research, thereby enabling understanding and eliminating the duplication of research efforts. (2) The *Distributed Service Laboratory* will be established as a research infrastructure to provide access to state-of-the-art collaboration facilities. (3) S-Cube's *program of*

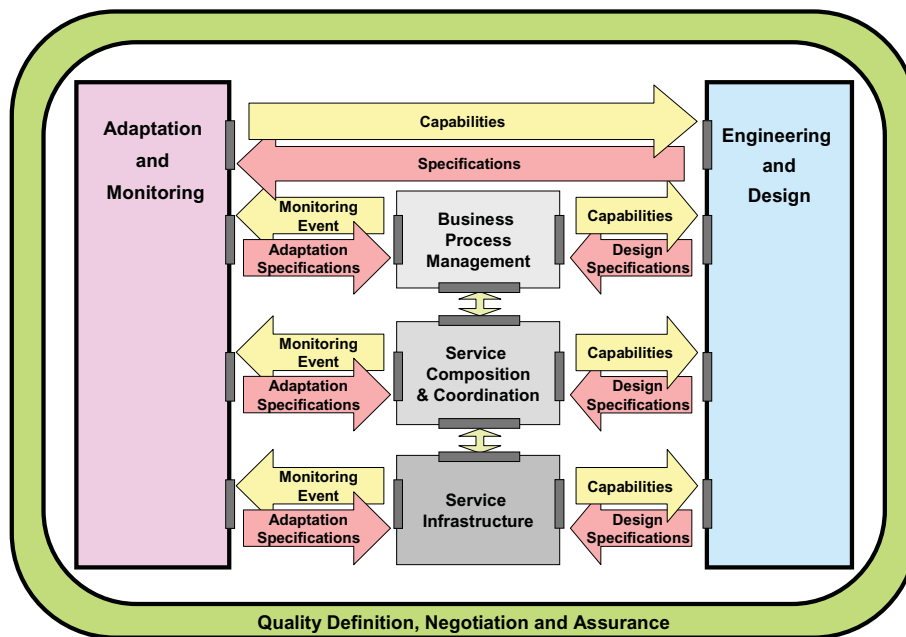


*education and training*, together with the *mobility programme*, will lead to cross-fertilisation of knowledge and durable research integration.

- **Joint Research Activities:** Work in S-Cube will be guided by the S-Cube research framework, which will be introduced in Section 4.
- **Spreading of Excellence Activity:** This activity will ensure a broad dissemination of research results and includes the organisation of international conferences, specialised workshops and summer schools, as well as a European Ph.D. programme.

## 4 The S-Cube Research Framework

The S-Cube research framework (see Figure 2) guides the joint research activities of S-Cube. In general, the framework distinguishes between principles and methods for engineering service-based applications and the technologies (or mechanisms) which are used to realize those applications. Principles and methods address cross-cutting issues like adaptation and monitoring, as well as quality definition, negotiation, and assurance. Technologies support specific requirements of the individual layers and provide capabilities to the cross-cutting principles and methods.



**Fig. 2.** The S-Cube research framework

What makes the S-Cube research framework unique when compared to the traditional “layered” way of developing service-based applications (see Section 1) is that the framework systematically addresses cross-cutting service issues. Further, the framework sets out to make the knowledge of the functional layers (which currently is mostly hidden in languages, standards, etc.) explicit in order to avoid overlaps and to

identify gaps in research. Finally, the framework is designed to handle the complexity of developing and adapting service-based applications.

To this end, the elements of the S-Cube research framework are defined by following a clear separation of two concerns:

**Concern 1: Technologies and local principles & methods:** The three horizontal layers of the framework are, similar to the traditional SOA layers, responsible for techniques and methods which are applicable locally within the layers. Also, concrete service technologies fall under the responsibility of the layers.

- The *service infrastructure* layer provides a high-performance execution platform supporting adaptation and monitoring of SBAs. The platform provides a set of core services, like search engines and virtualisation services to the other layers.
- The *service composition and coordination* layer focuses on novel service composition languages and techniques. Especially it provides mechanisms to adapt and monitor service compositions.
- The *BPM* layer addresses modelling, designing, deploying, monitoring and managing service networks to meet the goals of an organisation and its customers through the correlation and analysis of KPIs from the service composition and coordination layer with business processes.

**Concern 2: Overarching / cross-cutting principles, techniques and methods:** In addition to the local principles and methods, principles and methods falling into the following key cross-cutting aspects are addressed:

- The *engineering and design* aspect includes all issues related to the life-cycle of services and SBAs. This includes principles and methods for identifying, designing, developing, deploying, finding, applying, provisioning, evolving, and maintaining services, while exploiting novel technologies from the functional layers. An example is exploiting future service search engines for bottom-up SBA design.
- The *adaptation and monitoring* aspect includes all concerns with respect to the self-adaptation behaviour of a SBA. Specifically, this comprises languages and methods for defining adaptation goals and different adaptation strategies, which are triggered by monitoring events. An example for an adaptation technique that falls into the responsibility of this aspect is a strategy that correlates the monitoring events across the functional layers, thereby avoiding conflicting adaptations.
- The *quality definition, negotiation and assurance* aspect involves principles and methods for defining, negotiating and ensuring quality attributes and SLAs. Negotiating quality attributes requires understanding and aggregating quality attributes across the functional layers as well as agreeing on provided levels of quality. To ensure agreed quality attributes, techniques which are based on monitoring, testing or static analysis (e.g., model checking) are employed and extended by novel techniques exploiting future technologies (like the Web 2.0). Additionally, techniques for ensuring the quality of the actual adaptations are relevant here.

For each element of the framework, interfaces are defined that describe the capabilities that are provided by one element of the framework to another element, resp., the capabilities required by one element from another. As an example, one interface of the service composition and coordination layer defines which kinds of monitoring events (cf. Figure 2) are provided for the adaptation strategies defined in the adaptation and monitoring aspect.

## 5 Conclusions

The innovation required for devising theories, mechanisms and methods for making the next generation of services and service-based applications become reality, cannot be delivered by any research group in isolation. It requires the synergy and integration of a variety of research communities including but not limited to Grid Computing, Service Oriented Computing, Software Engineering, Business Process Management, and Human Computer Interaction. To this end, S-Cube, the European Network of Excellence on Software Services and Systems, brings together major European research institutions to jointly devise the scientific foundations for future service technologies and methods. The results of S-Cube will thus equip the organizations of the future with the capabilities to develop and evolve innovative software services and service-based applications.

**Acknowledgements:** We cordially thank all S-Cube members. Their contributions to the state of the art deliverables have been an excellent input for this paper.

## References

- [1] –. Survey Data Says: The Time For SOA Is Now. Forrester Research, Inc. (2006)
- [2] Di Nitto, E.; Ghezzi, C.; Metzger, A.; Papazoglou, M.; Pohl, K.: A Journey to Highly Dynamic, Self-adaptive Service-based Applications. *Automated Software Engineering*. 15 (2008) 313–341
- [3] Erl, T.: *Service-oriented Architecture*. Prentice Hall (2004)
- [4] Gallotti, S.; Ghezzi, C.; Mirandola, R.; Tamburrelli, G.: Quality prediction of service compositions through probabilistic model checking. In *Proceedings 4th International Conference on the Quality of Software-Architectures, QoSA 2008, Karlsruhe, October 14-17, 2008*. LNCS 5281, Springer (2008) 119–134
- [5] Hielscher, J.; Kazhamiakin, R.; Metzger A.; Pistore, M.: A Framework for Proactive Self-Adaptation of Service-based Applications Based on Online Testing. In *Proceedings ServiceWave 2008 conference, Madrid, December 10-13, 2008*, LNCS 5377, Springer (2008) 122–133
- [6] Josuttis, N.: *SOA in Practice: The Art of Distributed System Design*. O'Reilly (2007)
- [7] Kaye, D.: *Loosely Coupled: The Missing Pieces of Web Services*. RDS Press (2003)
- [8] Kephart, J.; Chess, D.: The vision of autonomic computing. *IEEE Computer*. 36(1), (2003) 41–50
- [9] Papazoglou, M. and Leymann, F.: *Business Process Management*. IEEE Encyclopedia of Software Engineering, John Wiley & Sons (2008)
- [10] Papazoglou, M.; Pohl, K.: Report on Longer Term Research Challenges in Software and Services. Results from two workshops held at the European Commission, with contributions from Boniface, M.; Ceri, S.; Hermenegildo, M.; Inverardi, P.; Leymann, F.; Maiden, N.; Metzger, A.; Priol, T. European Commission, <http://www.cordis.lu> (2008)
- [11] Papazoglou, M.; Pohl, K.: S-Cube: The Network of Excellence on Software Services and Systems. In *At Your Service: An Overview of Results of Projects in the Field of Service Engineering of the IST Programme*. MIT Press - Information Systems (2008)
- [12] van Moorsel, A.: Metrics for the Internet Age: Quality of Experience and Quality of Business. In: *Proceedings 5<sup>th</sup> Int'l Workshop on Performability Modeling of Computer and Communication Systems, Erlangen, September 15-16, 2001*, Arbeitsberichte des Instituts für Informatik, Universität Erlangen-Nürnberg, (2001)

## Survey on P2P Overlay Streaming Clients

Alexandro SENTINELLI<sup>1(a, b)</sup>, Luca CELETTTO<sup>(a)</sup>, Damien LEFOL<sup>(c)</sup>,  
Claudio PALAZZI<sup>(d)</sup>, Giovanni PAU<sup>(e)</sup>, Theodore ZAHARIADIS<sup>(f)</sup>, Ahola JARI<sup>(g)</sup>

<sup>a</sup>Advanced System Technology, STMicroelectronics, Agrate Brianza (MI), Italy

Tel. +39 039 603 {7600 | 7488} {alexandro.sentinelli | luca.celetto}@st.com

<sup>b</sup>UCLA, CA, USA Tel. +1 310 206 3212, gpau@cs.ucla.edu

<sup>d</sup>Livestation, 36-38 Hatton Garden, London EC1N 8EB, UK

Tel. +44 (0)20 7405 1444, damien.lefol@livestation.com

<sup>e</sup>Università di Padova, Padova, Italy Tel. +39 049 827 1426, cpalazzi@math.unipd.it

<sup>f</sup>Synelixis Ltd., Chalkida, Greece, Tel+30 22210 61309, zahariad@synelixis.com

<sup>g</sup>VTT, Tampere, Finland Tel. +358 20 722 3334 Jari.Ahola@vtt.fi

**Abstract.** Peer-to-peer (P2P) streaming systems grow in numbers and potential and several commercial products are already competing. Internet home users – through the diffusion of xDSL connections – represent the potential market of IPTV channels that Content Generators may distribute at reduced costs. This work describes the state of the art of P2P streaming clients and poses some questions about the end-user perspective in heterogeneous networks. To this aim, a representative set of experiments has been performed on a popular P2P system. The client offers live streaming content from some European broadcasters, start-up delay is a few seconds and the user satisfaction rank is good. The trend moves toward solutions that try to optimize the whole network stack, pursuing flexibility in terms of user needs and system requirements. This work is aimed at focusing on the key-drives in the design of P2P streaming clients.

**Keywords.** P2P Streaming, MDC, SVC, layered video coding.

### Introduction

In TV Broadcasting the acronym of P2P (peer-to-peer) is often perceived like the panacea of cost balance sheets just by exploiting the virtue of P2P scalability. There are, however, a lot of tradeoffs that need to be observed, especially for high quality stream.

Investigation evolved toward new approaches since the first Coolstreaming and relatives [3] generating new products offered to consumers (e.g., Zattoo [1]). At this state of the art, we might say those mainly depend on the type of video content and the platform environment (network infrastructure, the rendering device). An important social event, (ex. the soccer world cup), brings with it strict technological constraints such as the start-up delay, the video resolution and so on. Such constraints may become more flexible if the user is watching the news, weather, a music TV show, or an unpopular event. Moreover the streaming platform determines different user needs such as the resolution of the display, the cost of the network access (wired/wireless), the

---

<sup>1</sup> The work has been partially funded by the European Commission under the projects ICT-214063/SEA ([www.ist-sea.eu](http://www.ist-sea.eu)) and the P2PNext (<http://www.p2p-next.org/>). Finally, the authors would also like to thank the Media Delivery Platforms (MDP) Cluster for reviewing the paper.

computational power of the user device. In essence, user needs and expectations have a huge impact on the protocol design. Nowadays several commercial products, like the one chosen for our case study, offer the same content at different qualities (thus, bitrate) to satisfy different sets of users. The scenario can be very heterogeneous and involves a variety of fields and competences. For instance, an interesting synergy may overcome cross-layered coding techniques (SVC/MDC) in P2P network that use multiple tree schemes. The synergy produces better results than the state-of-the-art technology since such solutions allows distributing the same content to a larger portion of the overlay (SVC) or makes the overlay more flexible to network congestions or channel zapping behaviour (MDC). In the next section we overview related work in the P2P streaming literature. Then we show a set of experiments on a new successful client. Section 3 describes the qualitative synergy between SVC and MDC in P2P systems. Finally, in Section 4, conclusions are drawn.

## 1. Related Work

The power of P2P is derived from several advantages in terms of robustness, reconfigurability and scalability.

From the broadcaster point of view, the P2P approach permits to serve a large audience without the need of additional resources. From the user point of view the P2P improves the visual experience by delivering video content and allows publishing own content with less (in some cases without) infrastructure costs and overcoming bandwidth/processing load bottlenecks. P2P streaming systems strive to optimize three important metrics: i) start-up delay (i.e. the time from when the user first tunes on the channel to when the video is visible), ii) end-to-end delay (i.e. the delay between the content originator and the receiver, also known as playback delay), and iii) playback continuity index (i.e. the counter of frames rendered in the right order by the player). Most of the systems may be classified based on the type of distribution graph they implement: mainly *tree* and *mesh*, though a lot of hybrid solutions have been implemented already. Tree-based overlays implement a tree distribution graph, rooted at the source of the content. In principle, each node receives data from a parent node, which may be the source or a peer. If peers do not change too frequently, such a system requires little overhead; in fact, packets can be forwarded from node to node without the need for extra messages. However, in high churn environments (i.e. fast turnover of peers in the tree), the tree must be continuously destroyed and rebuilt, a process that requires considerable control message overhead. As a side effect, nodes must buffer data for at least the time required to repair the tree, in order to avoid packet loss. Mesh-based overlays implement a mesh distribution graph, where each node contacts a subset of peers to obtain a number of chunks. Every node needs to know which chunks are owned by its peers and explicitly pulls the chunks it needs. This type of scheme involves overhead, due in part to the exchange of buffer maps between nodes (nodes advertise the set of chunks they own) and in part to the pull process (each node sends a request in order to receive the chunks). Thanks to the fact that each node relies on multiple peers to retrieve content, mesh based systems offer good resilience to node failures. On the negative side, they require large buffers to support the chunk pull, as large buffers are needed to increase the chances of finding the missing chunks in the playback sequence. In the following, we begin with a brief overview of popular mesh-based systems and then focus on tree-based ones.

### 1.1. Mesh-based Systems

*BitTorrent's* technology, after the success as a file-sharing P2P system, has been applied to streaming applications: now the client must meet the playback deadline. New nodes register and receive the addresses of the *trackers*, which track which nodes have downloaded a piece of content. When the node contacts the peers advertised by the tracker, the node receives a map of the chunks of data they own and are able to share. At this point, based on various heuristics (e.g., bandwidth, delay), the node selects a subset of peers and requests chunks from them.

*PPLive* is a proprietary popular mesh [7] video streaming client. In order to relax the time requirements, to have enough time to react to node failures, and to smooth out the jitter, packets flow through two buffers, one managed by *PPLive* and the second by the media player. Two types of delay can be identified: i) the interval between channel selection and media display (10 to 15 s) and ii) the playback time, required for fluent playback in spite of the jitter (10 to 15 s extra). The time lag between nodes may range up to about one minute, which is unacceptable for some popular events (i.e. neighbours screaming "GOAL" even just 2 seconds before you!). Nevertheless, *PPLive* has proven to perform remarkably, for instance, on January 28, 2006, *PPLive* delivered a popular event in China, hosting over 200 K users, at data rates between 400 and 800 Kbps.

*SopCast* builds a mesh too [3], but enables easily anyone with an ordinary broadband connection to broadcast their own contents. Start-up delay can be from 15 seconds up to 2 or 3 minutes: it strongly depends by the type of content streamed, because that determines the size of the overlay, thus the availability of upload resources.

*DONet* (or *Coolstreaming*) is another very successful P2P streaming system implementation [8]. This system works similarly to *PPLive* for features such as registration, peer discovery and chunk distribution. At the opposite from *PPLive*, its creators published a lot of information about the internals of their scheme. As a peculiar feature, *DONet* implements an algorithm that chooses to download first the chunks with the least number of suppliers. In case of ties, *DONet* chooses the chunks owned by nodes with the largest bandwidth.

Differently from the aforementioned schemes, with *Anysee* nodes participate in building the mesh network but they do not pull chunks from other peers [9]. Every node in the mesh keeps an active path for data and a set of backup paths, in case the active path fails to deliver within certain time constraints. Furthermore, this scheme introduces the concept of inter-overlay optimization by involving all nodes in improving the global performance. For instance, it uses the spare bandwidth capacity of the nodes that are receiving CNN to help those nodes that are receiving NBC. Smaller buffers are then required compared to chunk-based schemes.

### 1.2. Tree-based Systems

As one of the first examples of end system multicast targeting video stream applications, the system described in [4] proposes to build a mesh topology that connects the participating nodes by selecting the links based on round-trip-time (RTT) estimates between nodes. On top of this mesh, a source rooted minimum delay tree is built and used for media delivery. This solution has been implemented and tested with conferencing applications and is the underlying technology of real systems such as ESM (End System Multicast).

Nice [5] is another tree-based solution designed for low-bandwidth, data streaming applications with a large number of receivers. Based on RTT information exchanged among hosts, this solution builds a hierarchy of nodes; in this structure, nodes keep detailed knowledge of peers that are close in terms of hierarchy and coarse knowledge of nodes in other groups. No global topological information is needed.

### 1.3. Multiple Trees Systems

A tree-based system, designed to limit end-to-end delay, tends to have a large number of leaf nodes that do not contribute to the overall performance of the system generating unfair sharing of network resources among nodes. *Splitstream* [6] fixes this problem by building multiple trees, where a node can be a leaf in all trees but one. Data, divided into stripes, are propagated using a different multicast tree for each stripe. A receiver, that wishes to attain a certain quality of service by receiving a certain number of stripes, joins the trees that correspond to those stripes.

Other advanced schemes such as *CoopNet* [10] and *ChunkySpread* [11] proposed to mitigate the strong dependency of a peer on all its ancestors in architectures based on a single tree. They are typically designed to work with more advanced video encoding techniques. For example, *CoopNet* uses Multiple Description Coding (MDC), which encodes a media stream into multiple independent descriptions. It constructs multiple independent multicast trees, one for each substream. A peer can improve its media quality by joining more multicast trees under the constraint of its download link capacity. More importantly, the departure of one ancestor in a multicast tree does not severely degrade the media quality of a peer, since it can still receive the majority of substreams from other multicast trees.

These hybrid schemes (tree vs. mesh) tend to get the best features from the two approaches: robustness to high churn rate (mesh network) and a better efficiency (tree-based) in terms of traffic overhead through a more ordered distribution of requests. Some of the above protocols represent also an interesting example of cross-layer optimization between the application and network layer. Layered video coding techniques, in fact, merge conceptually with the idea of splitting the main content in several sub-streams, fitting the variety of user needs in heterogeneous environments.

## 2. Case Study

For our case study we choose a live streaming client with good ranks from streaming-community forum and business-tech reviews.

The client delivers video and audio content from both popular and unpopular European broadcasters. Some channels are available in 2 resolutions, to meet the user preferences and/or needs.

### 2.1. Setting

We used a HP laptop, Centrino processor, 512 DD RAM, Windows XP operating system with a commercial ADSL connection.



What we measure in terms of statistics and performances is only related to the network traffic generated by the P2P client. Since the software does not allow our node to be the source of the video stream, observations have been performed only at client side. The main tools of this case study are Wireshark (Ethereal), a network analyser and Dumeter, a bandwidth monitor able to give a quick overview of the ongoing traffic. Other minor tools are a stopwatch, a bandwidth shaper, PacketPlotter to export Ethereal traces on Windows Excel.

## 2.2. Network Traffic

Exporting with PacketPlotter Ethereal trace Figure 1 we get a shot of the download traffic per IP address for a short session.




Figure 1. Download traffic volume. 800kbps stream.

The download rate is higher at the start-up but then is always stable at the same bitrate  $B$  even when the node changes supplier (from IP address supplier 72.55.146.237 to 72.51.38.218). The traffic volume grows nicely linearly and the content streams fluent and smooth. Though, it is just remarkable that there's no upload for the majority of channels (thus it works more like a Client-Server model). It follows the same chart for the popular streaming client PPLive (Figure 2). Solutions are both performing but designers faced eventually different constraints. PPLive is a pure P2P client where the infrastructure relies just on peers. The other client is a commercial product that has to deliver high quality live content at a remarkable bitrate (800 vs. 400 kbps for PPLive). At the moment there is no commercial Telecom provider able (or intending) to host a pure P2P network offering such a per-user bitrate. In video streaming, either live or VoD (Video On Demand), the P2P approach is not negligible to reduce costs, but servers are still needed.




Figure 2. PPlive client - Download traffic volume; 400kbps stream.




Figure 3. Start-up for the same channel at different bitrates.

### 2.3. Surplus Bitrate at the start-up

One of the biggest issues of P2P clients is indeed represented by the start-up delay.

As a matter of fact, it represents the responsiveness of the system from a user perspective. With respect to other popular P2P systems, this client never passes 10 seconds before getting a fluent stream. We tried to understand what happens just monitoring the traffic at the beginning of the streaming session, thus we clearly remarked a higher bitrate, approximately 1.5 Mbps (Figure 3). We can measure such interval  $I$  before the step down to the bitrate of the stream. The video actually starts after  $\sim 5$  sec, but it keeps downloading at 1.6 Mbps for a time interval depending on the bitrate of the channel. In Table 1 we see the aforementioned values after setting.

Table 1 Start session for different channels (VLC cache has been set to only one second).

Chan.	Bitrate (kbps)	Start-up Delay (sec)	(*)Interv. Higher bitrate (sec)	(**) Initial bitrate (kbps)
1a	400	2.3	15.00	1550.00
1b	800	7.0	40	1550.00
2	450	3.6	14	1550.00
3	400	7.0	15	1550.00

This is possible only if the server delivers the stream with an end-to-end delay bigger than the start-up delay perceived at client side (it also means that the stream cannot be pure live). Physically, we have two flows to the buffer, one in ( $f_1$ ) that accumulates the stream (the down process), one out ( $f_2$ ) that empties the buffer (at bitrate  $B_2$ ). In Figure 4 the green *not* overlapped (bitrate  $B_1$ ) area corresponds to the

surplus stream that the client has downloaded at the start-up. As any streaming client, the surplus covers a sort of guard interval to smooth the bursty nature of Internet traffic (or in case of temporarily network congestions) and to guarantee an ordered sequence of data chunks (especially when the node has more than one father).




Figure 4. Traffic surplus at start-up.

The solution here is as simple as efficient. The server stores at least  $(t_2 - t_0)$  “live” content, which can be considered as relatively popular. If the user is not able to check the “reality” of the content the end-to-end delay loses importance. Instead, the start-up delay  $(t_1 - t_0)$  was moved back until a few seconds. Other client’s start-up delay can be up to 1 minute, unsustainable for a commercial application. This performance has been achieved through the use of servers carefully dimensioned to the overlay size. P2P helps, but it represents just a contribution.


### 3. Heterogeneous environment

The heterogeneity of the network scenario determines as well different sets of users.

The popular channels of our client are available at two resolutions independently encoded. Such solution does meet the user requirement and may still exploit the virtue of P2P systems. However the selection of one fixed quality can be restrictive in conditions with varying bandwidth availability (frequent event in shared LAN, wireless,...). It is possible to improve this approach by adapting the quality stream on the fly, but we must ensure continuous playback by keeping the buffer not-empty. This is possible only if *several streamlets* are being downloaded in parallel. Starting with the lower quality channel reduces the start-up delay (Cf. Table 1) and switching to higher quality once enough buffering is done can significantly improve user experience. This solution ensures continuous playback, but is wasteful of bandwidth. SVC coding, instead, can provide different qualities from only one stream, and brings an interesting optimization at the network layer.


SVC is a layered encoding technique developed by the JVT committee to meet the requirements in heterogeneous scenarios. As an extension compatible with the already existing AVC/H.264, SVC makes possible, for an Internet video provider, to generate and store a single version of the video, maintaining the ability to deliver HD to premium customers and SD version content to client with less capable connections. This emerging standard is particularly suitable for IP networks where network fluctuations are frequent and unpredictable. The H.264/SVC allows an adaptation that is as easy as dropping some of the information that is packed in Network Adaptation Layer Units (NALU), whose first bytes give the information about the scalability layer they belong to; in other words the down-scaled bitstream is extracted from the main one with a sort of “cut and paste” mechanism. Even when the loss of compression

efficiency due to scalability is taken into account, SVC improves user experience compared to streamlet. This is evident in Figure 5, where a SVC stream containing four layers is compared to four independent streamlets of similar quality. If comparing only the best quality streamlet to the SVC stream containing all layers, the bitrate of SVC is around 30% higher, but this is more than offset by the gain of flexibility and saving of bandwidth. Moreover, SVC brings an interesting and unexpected synergy if used in P2P environments [2]. Although SVC loses a bit in compression compared to a simulcast like approach, the latter does not fully exploits the virtue of P2P systems.



**Figure 5.** Switching from one quality to another: Streamlet (top) vs SVC (bottom).

If the broadcaster delivers two different qualities, in the previous solution the two classes of users cannot share the base layer because the streams are independent (Figure 6, top).




**Figure 6.** P2P overlay networks: independent video encoding (top) vs. SVC (bottom).

Through SVC (Figure 6, bottom) we get a common content shared in a much bigger overlay: the two sub-overlays become an overlay embracing the whole one plus a smaller one delivering only the enhancement layer. This means that, at least for the base layer, the research of good candidates is faster because every peer can share his own content and resources. The degree of cooperation increases and the load of requests is better distributed. This type of approach is also very well suited for commercial application based on heterogeneous p2p networks. These applications usually rely on a mix of servers or CDN backbone and p2p for distributing content. The


backbone can then be used to insure that the base layer is delivered to all peers, and the peering is used to distribute enhancement layers. This enables a low start-up delay as the client connects directly to the server without waiting to find peers and the base layer stream is low bitrate. Once peers are located, the quality of the stream is improved by increasing the number of layers received. Relying only on the p2p network to distribute the base layer can be a risky strategy as without this layer no video can be decoded.

A different scenario and advantages manifest themselves when the platform streams content is encoded with Multiple Description Code (MDC). The improvement, in order of relevance, is represented by resilience to missed chunks and flexibility in the description assignment. In the case of MDC based on spatial redundancy, for instance, having multiple descriptions means distributing the information of spatially closed pixels to sub-streams that will be routed through different paths and shared in different sub-overlays.



**Figure 7.** Loss resilience comparison between classical approach (top) and MDC (bottom).

When the client misses one description's chunk the effect at the end-user side is a picture with a few missing dots distributed over a wide and redundant spatial area. These black pixels then are covered with the information coming from the other descriptions: this is possible because the sub-stream descriptions can be decoded independently one from the other.



**Figure 8.** Flexibility of MDC in P2P streaming.

The stream loses in compression efficiency because descriptions are redundant of spatial information. On the other hand the resilience to loss outperforms the classical

approach with a unique description: Figure 7 shows the effects and differences in error recovery from the end-user point of view. MDC offers also an interesting flexibility at network layer just in P2P scenarios. In Figure 8 we show an environment with a couple of description D1 and D2 and a critical (and typical) event where network flexibility is an important feature to pursue. The peer  $P_n$ , which receives the description D2, is not anymore supplied by his father at the edge of the overlay. The reason can be his father's zapping through channels or network congestion affecting the unique download link. Looking at the whole topology we observe that  $P_n$  may easily recover another description from the sub-overlay sharing D2. Since MDC is not based on hierarchical layers, the peer can change sub-overlay but still perceiving the same quality of the stream. Therefore, the peer  $P_n$  changes description's (D2 to D1) request and migrates to another sub-overlay, but still keeping the same quality of the original stream.

#### 4. Conclusion

The aim of this work is to understand the state of the art of P2P streaming design and particularly to show the importance of the user needs in heterogeneous environments.

Our case study points out the importance of the user experience and differentiation of system requirements. The user, depending on the type of content, may relax his expectations about the end-to-end delay but is still sensitive to responsiveness. We also have described aspects of cross-layer investigation offered by using layered video coding (i.e., SVC and MDC) techniques in p2p scenarios. The optimal design incorporates a flexible implementation able to adapt to constraints efficiently and dynamically. The user point of view represents one of the key-drives for this new investigation approach where cross-layers and user satisfaction metrics are still to be further analyzed, optimized and, most likely, discovered.

#### References

- [1] [www.zattoo.com](http://www.zattoo.com)
- [2] T.-C. Lee, P.-C. Liu, W.-L. Shyu, C.-Y. Wu, *Live Video Streaming Using P2P and SVC*, IFIP/IEEE - MMNS 2008, ), Samos Island, Greece, Sep 2008.
- [3] A. Sentinelli, G. Marfia, S. Tewari, M. Gerla, L. Kleinrock, *Will IPTV Ride the P2P Stream?* in IEEE Communication Magazine, vol. 45, no. 6, Jun 2007.
- [4] Y. H. Chu, S. G. Rao, H. Zhang, *A Case for End System Multicast*, in Proc. of ACM SIGMETRICS 2000, Santa Clara, CA, USA, Jun 2000.
- [5] S. Banerjee, B. Bhattacharjee, C. Kommareddy, *Scalable Application Layer Multicast* in Proc. of SIGCOMM 2002, Pittsburgh, PA, USA, Aug 2002.
- [6] M. Castro, P. Druschel, A.-M. Kermarrec, A. Nandi, A. Rowstron, A. Singh, *Splitstream: High-bandwidth Multicast in Cooperative Environments*, SOSP 2003), Bolton Landing, NY, USA, Oct 2003.
- [7] X. Hei, C. Liang, J. Liang, Y. Liu, K. W. Ross, *A Measurement Study of a Large-scale P2P IPTV System*, IEEE Transactions on Multimedia, vol. 9, no. 8, Dec 2007.
- [8] X. Zhang, J. Liu, B. Li, T. Yum, *Coolstreaming/DONet: A Data-driven Overlay Network for Efficient Live Media Streaming*, in Proc. of the 24th IEEE INFOCOM, Miami, FL, USA, Mar 2005.
- [9] X. Liao, H. Jin, Y. Liu, L. Ni, D. Deng, *Anysee: Peer-to-Peer Live Streaming*, in Proc. of the 25th IEEE INFOCOM, Barcelona, Spain, Apr 2006.
- [10] V. Padmanabhan, H. J. Wang, P. A. Chou, *Supporting Heterogeneity & Congestion Control in P2P Multicast Streaming*, IPTPS 2004, San Diego, CA, USA, Feb 2004.
- [11] V. Venkataraman, K. Yoshida, P. Francis, *Chunkspread: Heterogeneous Unstructured Tree-Based Peer-to-Peer Multicast*, ICNP 2006, Santa Barbara, CA, USA, Nov 2006.

## Content Adaptation Issues in the Future Internet

THEODORE ZAHARIADIS<sup>1\*</sup>, CATHERINE LAMY-BERGOT<sup>2</sup>,  
THOMAS SCHIERL<sup>3</sup>, KARSTEN GRÜNEBERG<sup>3</sup>, LUCA CELETTO<sup>4</sup>,  
CHRISTIAN TIMMERER<sup>5</sup>

*1 Synelixix Solutions Ltd*

*2 Thales Communications S.A*

*3 Fraunhofer HHI*

*4 STMicroelectronics S.A*

*5 Klagenfurt University*

**Abstract:** Future Media Internet is envisaged to provide the means to share and distribute (advanced) multimedia content and services with superior quality and striking flexibility, in a trusted and personalized way, improving citizens' quality of life, working conditions, edutainment and safety. Based on work that has taken place in projects ICT SEA and ICT OPTIMIX, and the Media Delivery Platforms Cluster of projects, we try to provide the challenges and the way ahead in the area of content adaptation.

**Keywords**—Future Media Internet, Adaptation, Scalable Video Coding

### 1. Introduction

The Internet has evolved and changed the way we work and live. End users of the Internet have been confronted with a bewildering range of media, services and applications and with technological innovations concerning media formats, wireless networks, terminal types and capabilities. In the near future these numbers are expected to rise exponentially. Moreover, it is envisaged that the Future Media Internet will provide the means to share and distribute (advanced) multimedia content and services with superior quality and striking flexibility, in a trusted and personalized way, improving citizens' quality of life, working conditions, edutainment and safety.

In this evolving environment, new transport protocols, new multimedia encoding schemes, cross-layer and in-network adaptation, machine-to-machine communication, rich 3D content as well as community networks and the use of peer-to-peer (P2P) overlays are expected to generate new models of interaction and cooperation. Furthermore, this will enable the support of enhanced Perceived Quality of Service (PQoS) and innovative applications “on the move”, like virtual collaboration environments, personalized services/media, virtual sport groups, on-line gaming, and edutainment. In this context, the interaction with content combined with

---

\* Contact information: Dr. Th. Zahariadis, Synelixix Solutions Ltd, 10 Farmakidou Av, Chalkida, GR34100, Greece, Email: zahariad@synelixix.com



interactive/multimedia search capabilities across distributed repositories, opportunistic P2P networks and the dynamic adaptation to the characteristics of diverse mobile terminals are expected to contribute towards such a vision.


Based on work that has taken place in the projects ICT SEA<sup>†</sup> and ICT OPTIMIX<sup>‡</sup> and the Networked Media Unit: Media Delivery Platforms (MDP)<sup>§</sup> cluster of projects, we try to provide in the following an overview of the challenges and the way ahead in the area of content adaptation. The remainder of this paper is organized as follows. Section 2 introduces a content-aware (access) network architecture. The means for cross-layer adaptation for enriched PQoS is described in Section 3. The main challenges we established for cross-layer adaptation are highlighted in Section 4 and Section 5 concludes the paper.

## 2. Content-aware Access Network Architecture

Even in the near future, the access network (even the evolved one) will remain the weaker part of the network. Moreover, in Peer-to-Peer (P2P) networks the end-to-end path may be unknown or time variant. Thus, it is desirable to have as much information and adaptation at the lower layers (up to the network layer) as possible, along with scalability functionality coming with the media codecs. Certain functions such as content caching in the network, content adaptation and cross-layer optimization would certainly need knowledge of the network conditions and characteristics.

In order to overcome this problem, wherever applicable in the network architecture, we introduce intelligent media/network-aware nodes. Within SEA we have introduced two types of content-aware edge devices/Media Aware Network Element (MANE):

- Home Media Gateway (HMG), located at the edge of the home environment and
- Network Media Gateway (sNMG) at the edge of the access networks, e.g., the 3GPP Service Architecture Evolution (SAE)



**Figure 1:** The concept of P2P overlay architecture


<sup>†</sup> ICT-214063/SEA (SEAmless content delivery) [www.ist-sea.eu](http://www.ist-sea.eu)

<sup>‡</sup> ICT-214625 OPTIMIX (Optimisation of Multimedia over wireless IP links via X-layer design) [www.ict-optimix.eu](http://www.ict-optimix.eu)

<sup>§</sup> Media Delivery Platforms Cluster, [mdpcluster@synelixis.com](mailto:mdpcluster@synelixis.com)

In general content-aware MANEs can offer multimedia storage, dynamic content adaptation and enriched PQoS by dynamically combining multiple multimedia content layers from various sources. Moreover, as they have knowledge of the underlying networks, this information on the network conditions/characteristics can be provided to and utilized by cross-layer control mechanisms and adapt the multimedia streams to the next network in the delivery path. This is an extremely important point for low bandwidth but with guaranteed QoS mobile networks as well as for the broadband but best effort P2P topologies. Introducing the content-aware nodes at the edges of the networks also enables us to realize a Peer-to-Peer (P2P) overlay topology as shown in Figure 1. Given content protection and management is in place, network operators and service providers may offer value-added streaming services with remarkable PQoS. Moreover, individuals may produce their own (real-time) content and make it publicly available to a larger audience, without having to rely on a specific, expensive networking infrastructure. In this environment, video streaming scalability, resilience and PQoS may be increased, as multiple sources may stream video segments, enriching the content on the fly, either at the network and/or at the end-user terminal.

The sHMG and the sNMG are key components in the SEA network architecture and play an important role in the cross-layer control and adaptation. However, this is not the only approach including an adaptation engine: the OPTIMIX project is working on a similar architecture including adaptation modules, whose role is to adapt the transmitted stream based on the current available QoS information (*quality feedbacks* which can include channel state information, packet error rate at various layers, retransmission rates, video quality, etc.) from the upcoming link. More precisely, the adaptation module will allow to transcode the stream based on requirements produced by the control algorithms (enforced by the controllers at application and base station levels). For scalable streams such as SVC ones, it will read and parse the stream to extract interesting portions of it, and for non-(sufficiently) scalable streams, it will introduce real transcoding of the stream, and not only parsing and cutting of it. Depending on the stream features, the adaptation module will be able to change the spatial, temporal resolutions or the data rate in an efficient manner. Typically, a temporally hierarchical stream [9] will allow temporal downgrade without error propagation, while a normal stream may result in prediction errors when downgrading is performed. Finally, another foreseen option of the adaptation module is the introduction of extension features, such as ciphering capability.



**Figure 2:** SEA adaptation network architecture

Envisioning such adaptation features in the communication chain nodes (MANEs), we may foresee a number of adaptation scenarios taking into account the final terminal capabilities (ranging from laptops to mobile phones). In order to optimize adaptation and increase the number of available scenarios, we extend the previously introduced

SEA and OPTIMIX adaptation network architecture in a more general format as shown in Figure 2. In this view, we assume that in the path from the Content Provider to the terminal, we may have  $N+1$  Adaptation Engines (AE) with  $N$  as the number of core networks. Each engine is responsible for adapting the video stream to the next network in the path, i.e.,  $AE_i$  adapts the video stream to the characteristics/capabilities of Network  $i$ , always taking into account the final terminal capabilities and user requirements.

### 3. Cross-layer Adaptation for enriched PQoS

One of the main challenges in Future Internet audio/visual communication will be the ability to provide a sustainable end-to-end quality as indicated by the user (PQoS), throughout the entire duration of the service delivery. Offering QoS-based services involves interactions, not only among a number of entities along the service delivery chain, but also across different layers. To coordinate effective adaptation and mapping of QoS parameters at service, application and network layers, cross-layer interactions are required. The objective of this adaptation and interaction is to find a satisfactory QoS trade-off, so that each end-user's service can be supported with available network resources. In this chapter, we highlight a very important issue in streaming multimedia: the cross-layer adaptation issues in order to achieve an enriched PQoS.

#### 3.1. Cross-Layer Control/Optimization/Adaptation

During the last couple of years, it has been shown that adaptation techniques limited to adaptation within a single layer are deficient in providing global optimal settings for the system. In contrast, cross-layer approaches have been extensively discussed in recent research literature for its viability for providing better performance than traditional layered architecture. Cross-layer approaches increase interaction among different layers to exploit the inherent characteristics of underlying network to maximize the utility (e.g., QoS) and reduce the cost (e.g., battery life, bandwidth). The involvement of multiple layers in cross-layer adaptation is important otherwise various mechanisms available at different layer are likely to counteract each other's effect. Although cross-layer designs emerged as a by-product of recent proliferation of wireless networks having totally different properties from wired networks, it offers various opportunities for heterogeneous environment, where a variety of application types, network technologies and terminal capabilities are utilized. Initial motivation to work on cross-layer issues was primarily derived from following reasons:

- Wireless networks are characterized by high bit error rate due to fast fading, co-channel interference and shadowing. To overcome these issues, different layers can cooperate to make transmission more resilient to noise.
- Effective network condition estimation requires parameters from multiple layers, e.g., packet loss ratio, BER, SNR etc. Network condition estimation is necessary to increase utilization and reduce cost
- Low efficiency of transport protocols over wireless networks due to their inherent characteristics is also a reason for the consideration of cross-layer design.
- Heterogeneity of applications, terminals and networks require more rigorous adaptation mechanisms. Especially, in the context of multimedia services, content adaptation is absolutely necessary due to enormous dependencies arising from

heterogeneity. Cross-layer adaptation can play a key role in handling such multiplicity of dependencies.

- Cross-layer adaptation can assist smooth transition from best effort to QoS.

### 3.2. Adaptation Control

In the proposed adaptation architecture, multiple adaptation engines can use information gathered from various layers. Thus, the AEs (see Figure 2) need to coordinate their adaptation information and decisions by proper signalling. The following sections describe how advanced media attributes can be signalled using format specific extensions to SDP and how MPEG-21 elements can be used for media adaptation.

#### 3.2.1. Signaling advanced content attributes over the Internet

Today's Internet streaming systems utilize the Session Description Protocol (SDP) for session declaration or negotiation in the context of other IETF protocols such as the Real Time Streaming Protocol (RTSP) for controlling point-to-point multimedia streaming sessions, the Session Announcement Protocol (SAP) for indicating multicast multimedia streaming sessions, or the Session Initiation Protocol (SIP) for negotiation of multi-directional conversational multimedia sessions. A session description consists of both session wide information and media description sections. For the purpose of adaptation, the media description provides crucial information which can be evaluated by different network elements. Transport of layered media offers opportunities for adaptation by selective manipulation of the different layers, if the attributes of each layer are described in the SDP.

In the following we show an example for specific media signalling for SCV. **Table 1** shows an example where a server offers a multi-session transmission with up to three potential media sessions. Lines 1 to 7 describe the session. The attribute specified in line 7 declares a group having decoding dependencies which contains the media sessions "1", "2" & "3" identified by the "mid" attributes assigned in lines 15, 21 and 26, respectively, to the media description blocks shaded with different colours in **Table 1**. Additionally, each media description is associated with one or more payload type (PT) numbers in lines 8, 16 and 22. Dependencies are given per payload type for all layers except the base layer of an SVC stream or the base view of an MVC stream by an "a=depend:" attribute line [7]. A detailed description for each media session is given by the format specific attributes in lines starting with "a=fmtp:", followed by the PT number to which the line applies. The parameters used in the example are all optional and apply specifically to media of MIME types "H264-SVC" or "H264":

**Table 1:** SDP example describing Scalable Video Coding content

Line #	SDP text
1	v=0
2	o=alice 2890844526 2890844526 IN IP4 192.0.2.12
3	s=SVC Scalable Video Coding session
4	i=SDP is a multi-session offer
5	c= IN IP4 192.0.2.12
6	t=0 0
7	a=group:DDP 1 2 3

Line #	SDP text
8	m=video 20000 RTP/AVP 96 97 98
9	a=rtpmap:96 H264/90000
10	a=fmtp:96 profile-level-id=4d400a; packetization-mode=0; mst-mode=NI-T; sprop-parameter-sets=Z01ACprLFicg,aP4Eag==;
11	a=rtpmap:97 H264/90000
12	a=fmtp:97 packetization-mode=1; mst-mode=NI-TC; sprop-operation-point-info=<1,2,0,1,4d400a,C80,B0,90,80,100> sprop-parameter-sets=Z01ACprLFicg,aP4Eag==;
13	a=rtpmap:98 H264/90000
14	a=fmtp:98 packetization-mode=2; mst-mode=I-C; init-buf-time=156320; sprop-operation-point-info=<1,2,0,1,4d400a,C80,B0,90,80,100> sprop-parameter-sets=Z01ACprLFicg,aP4Eag==;
15	a=mid:1
16	m=video 20002 RTP/AVP 99 100
17	a=rtpmap:99 H264-SVC/90000
18	a=fmtp:99 packetization-mode=1; mst-mode=NI-TC; sprop-operation-point-info= <2,3,1,0,53000c,1900,160,120,C0,200> sprop-parameter-sets=Z01ACprLFicg, Z1MADEsA1NZYWCWQ,aP4Eag==,aEvgRqA=,aGvgRiA=;
19	a=rtpmap:100 H264-SVC/90000
20	a=fmtp:100 packetization-mode=2; mst-mode=I-C; sprop-operation-point-info= <2,3,1,0,53000c,1900,160,120,C0,200> sprop-parameter-sets=Z01ACprLFicg, Z1MADEsA1NZYWCWQ,aP4Eag==,aEvgRqA=,aGvgRiA=;
21	a=mid:2
22	a=depend:99 lay 1:96,97; 100 lay 1:98
23	m=video 20004 RTP/AVP 101
24	a=rtpmap:101 H264-SVC/90000
25	a=fmtp:101 packetization-mode=1; mst-mode=NI-T; sprop-operation-point-info= <3,3,1,1,53000c,1900,160,120,100,400> sprop-parameter-sets=Z01ACprLFicg, Z1MADEsA1NZYWCWQ,aP4Eag==,aEvgRqA=,aGvgRiA=;
26	a=mid:3
27	a=depend:101 lay 1:96,97 2:99

- *profile-level-id*: profile and level of the contained AVC, MVC or SVC bitstream;
- *packetization-mode*: specifies whether NAL units are sent one per RTP packet, or whether NAL units may be aggregated either in a non-interleaved or interleaved manner;
- *mst-mode*: specifies for multi-session transport (MST) the way bitstream re-assembly is supported, e.g., relying on RTP timestamps or using cross-session decoding order numbers (CS-DON)
- *sprop-operation-point-info*: one or more vectors of 10 values describing the operation point(s) included in that media session. If present, the hexadecimal values specify: *layer-ID*, *temporal-ID*, *dependency-ID*, *quality-ID*, *profile-level-*

*ID*, *avg-framerate*, *x-resolution*, *y-resolution*, *avg-bitrate*, *max-bitrate*. A valid vector contains at least the triplet *temporal-ID*, *dependency-ID*, *quality-ID*;

- *sprop-parameter-sets*: Sequence and Picture Parameter Sets of the H.264, MVC or SVC stream.

As can be seen from the example above, the client can choose to request either one, two or three layers of an SVC stream, depending on its capabilities or user preferences, the choice being based either on the profile and level it supports or on more specific values included in the *sprop-operation-point-info* parameter, e.g., frame rate, resolution or bit rate. For media sessions 1 and 2, it can also choose different payload format numbers according to the packetization modes it supports.

Using SDP for the media signalling and adaptation control is by nature a media specific solution. In order to use a wider range of future media codecs within an adaptation context, the next section presents the generic approach of MPEG-21 Multimedia Framework.

### 3.2.2. MPEG-21 Multimedia Framework

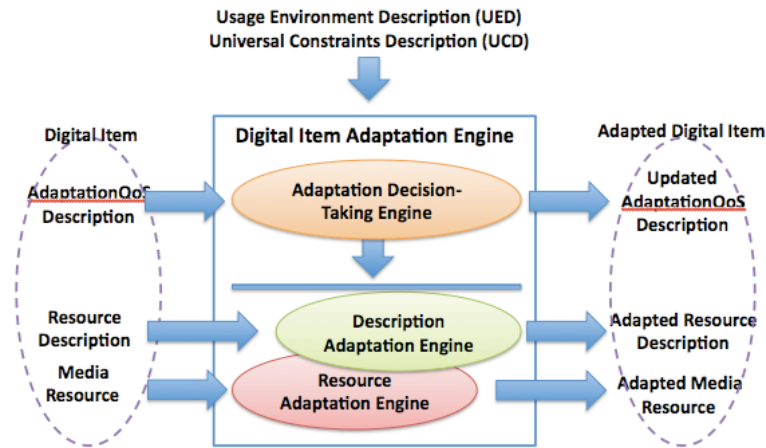
A comprehensive framework that deals with all these issues is MPEG-21 [3]. All parts of MPEG-21 address a distinctive set of requirements, which allow implementers of the standard to design and implement a system or application that goes beyond simple multimedia content delivery in an interoperable way. The MPEG-21 standard provides the transaction of Digital Items among Users. A **Digital Item** is a structured digital object with a standard representation and metadata. A **User** is defined as any entity that interacts within this framework or makes use of Digital Items. A Digital Item can be thought as a virtual structured digital container of media resources and metadata. It can include resources of any media type: audio, video, text, images, and so on. Metadata is the related information for the entire DI or part of a DI which provides semantic support.

Besides the interoperability, digital content needs to be adapted to various transmission channels and terminal devices for delivery. Digital Item Adaptation (DIA) can be achieved by applying various approaches such as adaptation at the server side, at the intermediate proxy or at the terminal. We list here all the relevant requirements exposed to the terminal side from these adaptations:

- 1) **Device independence adaptation:** From a terminal's perspective, terminal-independence adaptation is usually employed. *User Environment Description (UED)* is the key of this approach. It includes descriptive information related to user characteristics, (e.g., user information and user preferences), terminal capabilities (e.g., codec capabilities and display capabilities), network characteristics (e.g., available bandwidth, delay, and error), and natural environment characteristics (e.g., location and time).
- 2) **Content dependence adaptation:** such approach relies on the coding scheme which provides scalability. Particularly, in the case of SVC, it has achieved temporal, spatial and quality scalabilities co-existing in a single bit stream. This allows video adaptation at bit stream level. Such benefit outperforms other coding schemes as it increases the adaptation flexibility. For example, if a terminal is limited by certain constraints, e.g., computing memory or power, and its decoder can support SVC, there is no need of intermediate adaptation, since the receiver can perform the adaptation itself by discarding the relevant Network Abstraction Layer

(NAL) Units that convey enhancing layers. However, in this case the enhancement layers that are dropped at the decoder are delivered to the terminal for nothing and, thus, bandwidth is wasted.

- 3) **Adaptation by quality constraints:** to achieve optimal parameter settings under certain constraints imposed by terminals and/or networks for QoS management, Adaptation QoS (AQoS) is provided to assist the adaptation engine for decisions. AQoS specifies the relationship among various constraints, feasible adaptation operations satisfying these constraints, and associated qualities. AQoS can be used together with *User Constraints Description (UCD)* to acknowledge the adaptation.



**Figure 3:** Digital Item Adaptation.

Figure 3 shows a DIA engine incorporating the above-mentioned features but it should be noted that the actual implementation of adaptation engine is outside the scope of the standard. DIA specifies syntax and semantics of the description formats that steer the adaptation. The adaptation put forward several requirements (and possible approaches for solutions at the same time) for a terminal: 1) UED/UCD functional modules needs to be integrated in terminals; 2) a media decoder with support for the media resources' codec (e.g., SVC); 3) terminal and network QoS management for AdaptationQoS need to be provided.

The concept of MPEG-21-enabled cross-layer adaptation can be described [5]:

1. **Cross-Layer Model (XLM):** provides means for describing the relationship between QoS metrics at different levels – i.e., PQoS, ApQoS, and NQoS – and layers – i.e., according to the well-known ISO/OSI reference model.
2. **Instantiation of the XLM by utilizing MPEG-21 metadata:** Description formats (i.e., tools) as specified within MPEG-21 Digital Item Adaptation are used to instantiate the XLM for a specific use case scenario, e.g., Video-on-Demand. In particular, the *Adaptation QoS (AQoS)* description tool is used as the main component to describe the relationship between constraints, feasible adaptation operations satisfying these constraints, and associated utilities (qualities).
3. **Cross-Layer Adaptation Decision-Taking Engine (XL-ADTE):** The XL-ADTE is the actual subsystem which provides the optimal parameter settings for media



resource engines according to the XLM by processing the metadata compliant to MPEG-21 DIA.

Within the end-to-end multimedia *delivery chain*, the network QoS may be measured on an aggregated level and mapped to PQoS of individual streams [6].

#### 4. Challenges in Cross-layer Adaptation

The concept of cross-layer design sounds persuasively appealing. However, the successful experience of layered architecture burdens the adoption of a cross-layer approach. Currently, the research community is endeavouring the following challenges:

1. Cross-layer adaptation of the complete network infrastructure is very intricate due to handling enormous dependencies possibly in real time. A flexible architecture with proper interfacing between the layers is inevitable.
2. Cross-layer design breaks the layers and hence a clean isolated implementation of different protocols is no longer possible. Each cross-layer approach affects a complete system. An analysis of these effects becomes difficult.
3. The effects of coexistence of different cross-layer interactions are to be observed in terms of system performance, maintenance and scalability. Analysis is further perplexed if different types of cross-layer optimizations are deployed across an end-to-end delivery chain.
4. Global metrics are required that maximize the utility (e.g., QoS) and minimize the cost (e.g., battery life) under various constraints by efficiently prioritizing layers' local optimization criteria.
5. Optimization of cross-layer parameters is a complex multivariate problem with various constraints derived from QoS guarantees, available bandwidth, power consumption, etc. Apart from the essential requirement of computational efficiency, the highly dynamic nature of wireless networks demands a rapid convergence of the solutions.
6. It has to be evaluated where the actual control of a cross-layer adaptation should be located. Without a central control, different cross-layer adaptations might counteract each other. Different candidates include a separate coordinator or a particular OSI layer.
7. Cross-layer adaptation simulations are generally more complex than traditional network simulations. Hybrid approaches combining network simulation tools, hardware support and analytical approaches are usually required.
8. Not even a single cross-layer proposal has been tested comprehensively under real world traffic scenarios and hence QoS, power consumption and scalability of these approaches are yet to be gauged deterministically.
9. The assurance of fairness is yet an un-promised reality by cross-layer design.
10. As cross-layer designs break the well-established layers of the ISO/OSI model, interoperability issues arise also which are not considered as major at the moment but will emerge once these designs will find their ways into products.

#### 5. Conclusions

In the Future Internet, new formats for multimedia content will evolve and emerge. From today's AVC, AAC, MP3, and early instantiations of SVC, the media delivery

platforms will accommodate the carriage of a wide range of the above formats as well as SVC, MVC, a multitude of audio and gaming-friendly formats, H.265, MPEG/Laser and other surprising industry standards and ad-hoc media formats. All this while striving to be on the one hand content-agnostic, yet applying network intelligence, achieved through intimate content awareness, for the purposes of traffic shaping, PQoS, security, reliability and more, a tough challenge. Furthermore, the prevalence of virtual and parallel personalized worlds, coupled with progressively changing virtual characters, adds a dimension of complexity tricky to contain and to scale up.

Finally, existing networks' Cross Layer Control (CLC) and adaptation provides significant improvements in the PQoS under specific networking and transmission conditions. However, further research is required especially in the case of P2P topologies, where the physical infrastructure may be an arbitrary, timely varying combination of links belonging to different networks. Moreover, CLC schemes are required to face the network and terminal heterogeneity and take advantage of new (3D) advanced coding and delivery schemes by proposing network abstraction mechanisms, able to model the underlined end-to-end paths, describe the functional dependencies and determine the optimum adaptation of the multimedia resources.

**Acknowledgment:** The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2013) under grant agreement n°214063 SEA (SEAmless content delivery, [www.ist-sea.eu](http://www.ist-sea.eu)) and n° 214625 OPTIMIX (Optimisation of Multimedia over wireless IP links via X-layer design, [www.ict-optimix.eu](http://www.ict-optimix.eu)). The authors would also like to thank the Media Delivery Platforms (MDP) Cluster for reviewing the paper.

## References

- [1] H. Schwarz, D. Marpe, and T. Wiegand, "Over-view of the scalable video coding extension of the H.264/AVC standard", *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 17, no. 9, Sep. 2007, pp. 1103-1120
- [2] Aljoscha Smolic, Karsten Mueller, Philipp Merkle, Christoph Fehn, Peter Kauff, Peter Eisert, and Thomas Wiegand, "3D Video and Free Viewpoint Video – Technologies, Applications and MPEG Standards", *Proceedings of International Conference on Multimedia and Expo (ICME 2006)*, Toronto, Canada, pp. 2161-2164, July 2006.
- [3] S. Devillers, C. Timmerer, J. Heuer, and H. Hellwagner, "Bitstream Syntax Description-Based Adaptation in Streaming and Constrained Environments", *IEEE Transactions on Multimedia*, vol. 7, no. 3, pp. 463- 470, June 2005.
- [4] F. Pereira, J.R. Smith, A. Vetro, (eds.), "Special Section on MPEG-21", *IEEE Transactions on Multimedia*, vol. 7, no. 3, Jun. 2005.
- [5] V. Srivastava and M. Motani, "Cross-Layer Design: A Survey and the Road Ahead", *IEEE Communications Magazine*, vol. 43, no. 12, December 2005.
- [6] B. Shao, M. Mattavelli, D. Renzi, M. T. Andrade, S. Battista, S. Keller, G. Ciobanu, and P. Carvalho, "A multimedia terminal for adaptation and end-to-end QoS control", *IEEE International Conference on Multimedia & Expo (ICME 2008)*, Hannover, July 2008.
- [7] T. Schierl, S. Wenger, "Signaling media decoding dependency in Session Description Protocol (SDP)", work in progress, Internet Engineering Task Force (IETF), Multiparty Multimedia Session Control (mmusic), Oct 2008, <http://tools.ietf.org/html/draft-ietf-mmusic-decoding-dependency-04>
- [8] S. Wenger, Y.-K. Wang, T. Schierl, A. Eleftheriadis, "RTP payload format for SVC video", work in progress, Internet Engineering Task Force (IETF), Audio Video Transport Group (avt), November 2008, <http://tools.ietf.org/html/draft-ietf-avt-rtp-svc-15>
- [9] C. Bergeron, C. Lamy-Bergot, G.Pau, and B. Pesquet-Popescu, "Temporal Scalability through Adaptive M-Band Filter Banks for Robust H.264/MPEG-4 AVC Video Coding", *EURASIP Journal on Applied Signal Processing*, vol. 2006, Article ID 21930, 2006.

# QoE and \*-awareness in the Future Internet

Fidel LIBERAL <sup>a,1</sup>, Jose-Oscar FAJARDO <sup>a</sup> and Harilaos KOUMARAS <sup>a,b</sup>

<sup>a</sup> *University of the Basque Country, Spain*

<sup>b</sup> *NCSR Demokritos, Greece*

## Abstract

Future Internet will have to cope with yet unknown terminals and services (even users), in a number and heterogeneity never seen before. So, flexibility or adaptability will be considered as one of the most important design principles. This flexibility will demand different kinds of *awareness* both in the ends and in every node in the service supplying chain, while targeting users' satisfaction as the final goal of any management process. Although virtualization and "everything is a service" approaches seem to be promising foundations to guarantee this flexibility, Future Internet will be built upon real world mobile wireless network technologies, so that cross-layers issues and quality constraints will persist. Quality of Experience (QoE) could play a significant role there, since it could provide an unified metric, isolating users from low level details or complex NQoS definitions. We will also show an example of how user-aware network tuning mechanisms are able to provide similar users' QoE with lower resources consumption and therefore propose that QoE and \*-awareness were considered in the Future Internet design from the very beginning.

**Keywords.** QoE, NQoS, awareness, network intelligence

## Introduction

During the past years researchers have shifted the focus on the deployment and growth of the Internet, from an initial technology-driven approach to a user needs-driven one. This user-centered approach has resulted in several proposals aimed at bringing *awareness* to the network beyond the bit-pipe service-neutral network paradigm.

This *awareness* reflected the need for future networks **capable of coping not only with technological challenges related to performance, but also with users' preferences, location or context**. These networks will be built upon different access technologies and would deal both with network performance issues, service-specific constraints and even characteristics of the content (such as its type of content, codec, or the dynamics of the information represented).

These multiple needs lead to different research topics that have been widely studied in latest years:

---

<sup>1</sup>Corresponding Author: Fidel Liberal -ADAMANTIUM FP7 Project Grant no. 214751-, University of the Basque Country (UPV/EHU), ETSI de Bilbao, 48013 Bilbao, Spain; E-mail: fidel.liberal@ehu.es.

Jose Oscar Fajardo E-mail: joseoscar.fajardo@ehu.es

Harilaos Koumaras E-mail: koumaras@iit.demokritos.gr

- QoS-aware networks (including \*-constrained routing protocols, traffic differentiation and TE schemes, QoS brokers...)
- Ambient networks
- Location based services (including multihoming aspects, location based CDNs...)
- Self-managed, -learning, -organizing networks, technologies, radio interfaces...
- Content-aware networks
- Network-aware content and services

Most of the management schemes proposed in these different user-centered research areas usually share a common target: all management activities are devoted to guaranteeing *quality*. Since these proposals usually focus on a single technology, managed performance issues (i.e. Key Performance Indicators -KPIs-) are technology dependent. However, at the end, actual user satisfaction or *Quality of Experience (QoE)* will depend on several factor related not only to these “simple” network performance issues but also to more complex non-technical ones, such as content characteristics, users expectations and their particular context. Future networks should consider user satisfaction as the final end and, therefore, should be able to handle every single parameter in all aforementioned domains (content, network performance, services, users preferences...) that has an impact on this satisfaction.

On the other hand, quite surprisingly, network infrastructure seems to be no longer a constraint for e2e services. In fact, its structure, including network nodes and links, has already begun to dissolve into the “*everything is a service*” paradigm by means of **virtualization**. These virtualization proposals somehow admit the **incapability of a single protocol suite or network architecture to cope with the great variety of different services and users requirements**. So, instead of trying to provide a good solution for all, virtualization aims at providing “users” with the means for building the network that best fulfills their particular requirements.

Nevertheless, although virtualization seems to be a promising foundation for future networks, at the end there will be a real (certainly mobile and wireless) transmission technology behind all virtual networks. Then, even with virtualized links and network nodes, cross-layer issues will appear and demand richer and more accurate definitions of network behavior, far beyond traditional simple NQoS parameters (i.e. capacity, delay, jitter, losses), and more closely related to specific aspects of the services to be deployed (see [1] and [2] for examples of other type of definitions).

Future Internet will have to handle all these new service requirements, so that flexibility or adaptability should be considered as the first design principle. This flexibility will demand **different kinds of awareness both in the ends of the communication and in every node in the service supplying chain with users’ QoE as the final objective**.

The rest of the document is organized as follows: In Section 1, today’s initiatives around including *awareness* in networking technologies will be analysed. In Section 2, we will examine two important drawbacks of current proposals aimed at developing more intelligent networks, namely technology dependence and cross layer issues. Then, the role of Quality of Experience (QoE) management will be described. The case study of VoIP over  $\geq 3G$  accesses will therefore show in Section 3 the importance of handling QoE. This analysis will motivate the conclusions in Section 4, that will state the need for considering content and service characteristics, together with user preferences in the design and operation of the core of a QoE-aware Future Internet.

## 1. Awareness in Today's Internet

Following the maxim “those who cannot remember the past are condemned to repeat it” we have analysed prior research proposals while trying to draw a coherent picture of network related roadmaps and different visions of *awareness* in Today's Internet.

We have carried out an intensive survey of EC funded research projects in Europe over the past 10 years. This roadmap should not be seen as a proposal toward non-disruptive design principles for the Future Internet (vs. the clean-slate approach that is gathering momentum among the research community). Instead, we just have tried to identify unsolved research hot topics that have been faced with nowadays technologies but that will still determine the design of the Future Internet. Neither it was an exhaustive statistical exercise, but an attempt to figure out the big numbers behind *awareness*.

In the introduction of this document we identified 6 different areas of the so called *awareness*. Based on these 6 different research areas we have selected and classified 66 research projects that covered one or more of these areas, within successive Framework Programs (from FP4 to FP7, in IST/ICT areas). The basic information of these projects is publicly available using the search tools in the EC CORDIS website [3].

The first result of our analysis was that funding associated directly to some kind of awareness-related topic has been increased dramatically over the past years (see Figure 1), showing the growing interest on research areas related to bringing awareness to the network. Besides, in our survey we have only considered those projects that explicitly addressed these topics in their summarized description (and/or keywords). So, there would be many other networking projects that, although focused on different research topics, considered also awareness as a secondary target within Integrated Projects or Networks of Excellence. We should also notice here that statistics for year 2008 and later show a decrease because they do not take into account future Calls and proposals currently under evaluation.

The growth is not equal in all the research areas (see Figure 2), since some of them have appeared recently and some others have been surrounded by *buzzwords* that have changed over the years (while the main concept has remained more or less the same). In fact, since these 6 different research areas are inter-related, the overall contribution in Figure 1 provides a clearer view of time evolution of *awareness* in Today's Internet.

An inspection of the results and proposals within analysed projects lead to some well know issues:

- **QoS-aware proposals** quite often use some kind of QoS brokering systems on top of traditional network management protocols and resource managers (see for example [4] for DAIDALOS QoS architecture). Most of these initiatives claim that complex QoS request mechanisms are usually not standardised. As a result, nowadays there seems to be no working global scale NQoS management framework. We could point out different reasons to explain this lack of success, such as the traditional scalability issues of QoS management systems or that different connectivity providers take part along the service provision path (so, administrative or business model issues rather than technical ones). The need for per flow marking and scheduling lead also to well-known performance problems while, at the end, handled NQoS parameters do not ensure that final QoE will satisfy users' needs. So, the *quality* of the content delivered to users due to network transmission effects would be still an open issue.




Figure 1. Total budget of *awareness* related projects

- Regarding **Ambient Networks**, associated projects introduce research challenges regarding intelligent handover between very different access technologies based in different criteria (such as signal strength, coverage, terminal type, user preferences,... -see for instance [5]-) while providing seamless connectivity. Once more, the actual situation is that there has been no standardized global scale deployment of such kind of solutions and that, most of them, are deeply technology dependent (i.e. handover between all possible combination of different radio access technologies).
- **Location based services** are generally focused on two different planes: on one hand, location awareness as an input for the logic behind end-to-end services (such as the service that suggest you the best restaurant closest to your location). On the other, as a support tool for some routing/handover decision mechanism (i.e. as in location based mobility management -i.e. see [6] for WINNER IST project results-). However most protocols' inner structure does not provide fields or mechanisms to include any location information yet.
- **Self managed/organized systems** address a large variety of different research fields, from MANETs to sensors networks or cognitive radios. The use of different kind of algorithm, such as simulated annealing, genetic algorithms, bayesian reasoning or neural networks, in this type of proposals aims at optimizing general network performance parameters in an automated way (see [7]). Self-managed systems are a rather new and promising research field that, although still not too mature, will definitively play a significant role in the design of the Future Internet.
- **Content-aware networks** try to behave according to the specific content delivered (including the content itself, used codec, packetization scheme, transport protocol, etc...). In order to do so, some kind of source coding (or intelligent edge marking) is needed in order to allow an efficient handling of multimedia flows throughout the different networks nodes. So, some common frameworks to define content characteristics have been defined (i.e. MPEG7, MPEG21 Digital Item Adaptation, or specific ones such as those in [8,9]). However the specific treatment applied to each flow is usually based on particular effects of each technology into end to end transmission (i.e. interaction between VoIP calls and low level UMTS RLC procedures to be seen in Section 3).



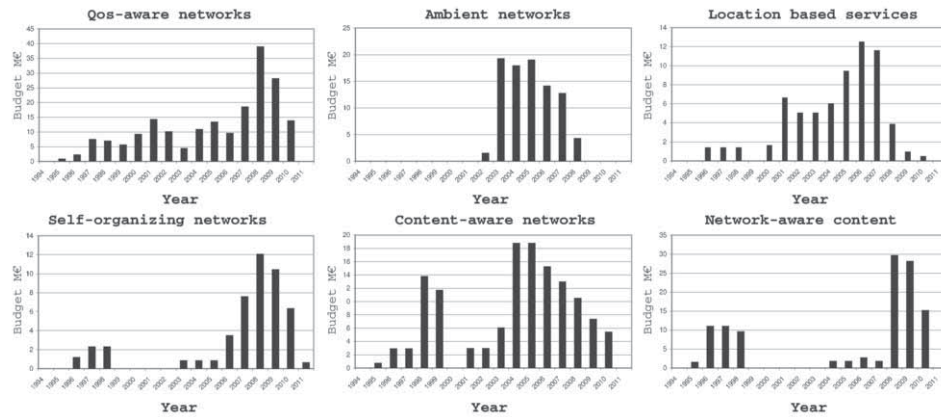


Figure 2. Budget evolution of *awareness* related projects

- Finally the term **Network-aware content** is used with those multimedia services (or more precisely codecs) that adapt their behavior according to particular network conditions. A higher content quality increases the QoE under perfect network conditions, but requires more QoS resources. Thus, upon network impairments, a lower target content quality may result on better service experience. This concept is currently being introduced into multimedia services. The Adaptive Multi-Rate (AMR) codec, standardized by the 3GPP in 1998, implemented eight codec modes at different data rates and, consequently, different initial listening quality. Under network degradations, VoIP services are able to decrease the target bitrate in real-time to cope with the new QoS constraints and enhance users' experience. A similar approach is being adopted by the joint Video Team (JVT) for networked video services. The Scalable Video Coding (SVC) extension endows the H.264/MPEG-4 AVC video compression standard with the capability to divide a video flow in a set of substreams, each of them providing different estimated QoE levels, and with the particular property that the reception of different substreams may contribute to the content quality additively (see for example [10], that shows results from DANAe IST project).

## 2. Technology dependence, cross layer issues and the role of QoE

In previous Section we have reviewed the outputs of several R&D projects and identified troubles to be still faced. As an overall conclusion, current “independent layers”-based structure of Internet protocols and their technological-only approach make it difficult to provide users with the different service characteristics they demand. So there are several common aspects related to technology dependence and cross-layer issues that will have to be considered in the design of the Future Internet.

First, there's no doubt that seamless mobile connectivity will be built upon several different access technologies. With Today's technology, even with full IP access networks, some patches must be used in order to solve handovers between each pair of



technologies, both from the mobile terminal's and internetworking technologies' point of view (regardless many efforts of IETF, 3GPP, ITU or UMA standardization initiatives). Most of needs for *awareness* related to handovers are the result of a tight coupling between content delivery capabilities and network constraints.

Similarly, different other alternatives try to adapt content to actual network capabilities by recoding and/or protocol adaptation. In any of these cases researchers have to deal with a lot of low level interactions between different layers, which result on poor e2e performance even when typical averaged NQoS parameters in each layer are apparently above acceptable thresholds.

Therefore, Future Internet will have to provide mechanisms to ensure that required complex *quality* demands are satisfied. This complexity is usually specified now by different metrics, associated with the low level parameters in the underlying technology. Since the objective of any network is providing users with multimedia content with enough quality for them to be satisfied, QoE could be used as a final single metric associated to the specific service and technology independent. For example, most users have already identified some multimedia formats as "enough quality". Except from advanced users, few of them bother about MP3 codec rate or DIVX/mpeg4 encoding scheme, framerate or number of processing steps (if you can burn it into a CD it is "good"). So, they have assimilated that nearly any MP3 or divx film fulfills their requirements. At the same time, clock speed has suddenly disappeared from microprocessors names and advertisements, replaced by other performance benchmarks. **Future Internet should be able to provide this kind of confidence to users. They should be provided with multimedia contents with the QoS required for them to be fully satisfied**, regardless all the low level technical details that the network intelligence will have to deal with.

In order to do so, a lot of research has been focused on proposing cross-layer adaptation techniques for the latest audio and video encoding standards. The overall aim of all the cross-layer adaptation concept is to provide QoS continuity across different layers of the delivery chain. More specifically, the research interest has been focused on the impact of each layer involved in the provision process (i.e. Service, Application and Network Layer) on the perceptual quality level of the finally delivered service by defining and correlating the various *quality*-related metrics of each layer. Regarding the mapping between the various discrete *quality* layer (i.e. QoE/ApQoE/NQoS), Table 1 provides an example of the representative metrics of each level, which will be used in the mapping process, for Video delivery systems:

At the **Service layer** the critical metric is the user satisfaction. The QoE evaluation will give service providers and network operators the capability to minimize storage and network resources by allocating only those resources needed to preserve a specific level of user satisfaction. At the **Application layer**, given that during the encoding/compression process of the initial content the quality is degraded by the appearance of specific artifacts, the values of the Application QoE (ApQoE) parameters (i.e. bit rate, resolution) determine the finally achieved QoE. Thus, the various encoding parameters must be considered as significant metrics of the application layer, since they have a straightforward impact on the deduced QoE level. If additional transmission problems are considered due to limited available bandwidth, network congestion etc... they will be also should be also considered as metrics at the ApQoE layer. At the **Network layer** NQoS related metrics (i.e. Packet Loss Ratio, Packet Loss scheme and Packetization

scheme) are used in an objective aspect, trying to determinate the impact of all low level interactions into final e2e NQoS achieved.

Service QoS Level	Application QoS Level	Network QoS Level
User Satisfaction	Decodable Frame Rate	Packet Loss Ratio
QoE level	Decoding Threshold	Packet Loss Scheme
Terminal Specifications	Encoding Parameters	Packet Size

Table 1.: Example of metrics at different layers for video

Similarly, in VoIP communications, the QoE is mainly determined by the following characteristics:

- **Session establishment delay.** In mobile data networks, the most relevant delays to be taken into account are the radio bearer set-up time and the performance of the session signalling protocol.
- **Interactivity.** The feeling of interactivity in conversational services is determined by the round-trip delay at user level. If this time increases over a threshold, both users could not coordinate when to speak or to remain listening.
- **Listening quality.** The primary factor determining the listening quality level is the fact that words are understandable. Otherwise, the purpose of the communication would not be fulfilled. The quality of received voice is mainly determined by the digitalization and codification processes, and the possible loss of voice frames in the transmission.

Regardless the type of service considered, in order to deploy QoE-driven network performance management systems we will have to evaluate the relationships between technical and user perceptions dimensions, which are greatly affected by the service conditions. For example, the content codification method has a great impact on quality perception results, since different codecs show different resiliency to frame losses. Additionally, the user device type and configuration is also to be considered. The same network performance values could result on different QoE levels depending e.g. on the device buffering capacity, the screen resolution or the processing capacity. However, QoE related considerations have been mostly incorporated into QoS management systems as upper thresholds for every individual performance metric leading to an over-provisioning of resources for some users and under-provisioning for others.

In addition to a better resource planning, the user- and QoE-awareness is a critical factor in the Future Internet for overcoming possible network degradation states. The reaction to network degradations performed by the current QoS management model is based on the set of pre-established actions for the affected class of service regardless special characteristics of each flow. Yet, a QoE-driven QoS management mechanism would try to maximize the general QoE level by taking into consideration specific content characteristics, such as the specific codec and FEC characteristics or loss patterns, in a specialized way.

As a result, the Future Internet will not only benefit of a QoE-driven management in terms of a higher capacity, but will be able to mitigate the effects of network impairments in a more optimal way.




Figure 3. MOS scores at different CoS and configurations

### 3. Case Study: VoIP and $\geq 3G$ data access

The evolution of radio access technologies makes us think on a Future Internet with plenty of itinerant users launching resource-greedy multimedia-enabled services. Thus, besides the performance variability inherent to the radio transmission technologies, one of the hot topics is how currently proposed access and backhaul networks will cope with the resulting volume of variable data rates.


The expected evolution of users and services in the Future Internet requires a more specialized and personalized QoS management, based on keeping the accurate QoE levels. On one hand, the perceptual schemes of mobile users are not the same to the traditional fixed-access case, resulting on different tolerance thresholds. On the other, Internet access through radio technologies is becoming more and more usual even for non-mobile users. Thus, the user-awareness can not be performed just based on the connectivity, but also the location and other contextual factors have to be considered.

As cited previously, the current resource management model based on aggregating traffic flows of similar QoS requirements into classes of services involves several deficiencies. For example, the QoS model proposed by the 3GPP recommends that VoIP services should be treated as Conversational class, associated to a set of maximum values for different network metrics.


Then, the transmission delay for the UMTS Bearer Service of a Conversational class is recommended to be kept below 100ms. Following the E-model [11], up to 100ms of one-way delay can be considered negligible for the conversation quality, while an increase from 100ms to 200ms corresponds to a perceptual degradation of 0.1 in the MOS scale. Therefore, it shall be individually considered if the increase in resource consumption is in correspondence to users' service experience.

The general trend is to consider that currently deployed UMTS networks, mainly based on the Background class, are not sufficient for an accurate provisioning of conversational services, due to the variable delays. However, suitable service and device configuration can do the best out of this type of networks, allowing users to reach similar QoE levels to the QoS-enabled solution based on live network measurements. Even more, under certain conditions, it could be preferable to provide a mobile VoIP service based on Background class. Not only the economical perspective shall be considered, that will lead to a user-awareness, but also the power consumption is to be taken into account, introducing the device-awareness.

In order to evaluate how low level QoS affect end users QoE we have carried out several simulations. The proposed scenario is based on the simulation of a VoIP service pro-



**Figure 4.** VoIP MOS=f(buffer sizes) SDU Discard Timer = 250ms



**Figure 5.** Empirical CDF

vision over a Packet-Switched UMTS service. Since the aim is to evaluate the combined impact of the application-level and UMTS RLC-level parameters, several simulations have been run with different configuration combinations. The VoIP dejittering buffer size has been configured to 40ms, 60ms, 80ms and 100ms. Additionally, the RLC SDU recovery function has been modified from UM to AM with the SDU Discard Timer values configured to 250ms, 500ms and 1000ms. Figure 3 shows the MOS scores obtained for those configurations that resulted on the optimal performance for different combinations. Figure 4 and Figure 5 shows a more detailed study of the case where RLC is configured to AM with a SDU Discard Timer of 250ms. In this case, it can be observed that the best resulting configuration depends on the analysis. Figure 4 represents the values of the obtained MOS score through time. Analysing the mean value of this variable, it results that the 60ms buffer size obtains a better performance. Figure 5 show the Empirical CDF for the same traces. It is remarkable that for the area of lower MOS scores, a buffer size of 40ms results on a better performance.

As a result, we can see how **the management of low level network parameters can result in very different users' QoE**. So, users would get **"better" services with equivalent (or even lower) network resources consumption** with some kind of **QoE-aware management mechanism that took into account both technology, service and users constraints**. Since so many cross-layer interaction still exists Future Internet must be able to address and provide complex definitions of QoE requirements and capabilities.

#### 4. Conclusions

In this work we have analysed the role of QoE-targeted *awareness* in the design of the Future Internet.

By carrying out an intensive analysis of R&D projects during the last 10 years in Europe we have identified 6 different research areas around *awareness* with open issues in Today's Internet. In order to face associated challenges most of the proposals have aimed at bringing some kind of *awareness* to the network. However, since associated intelligence has not been incorporated in the design of Today's Internet from the beginning, proposed "patches" generally lack of global scale adoption.

Furthermore, even when virtualization seems to be a promising approach in order to define a flexible Future Internet, cross-layer and technology dependence problems still arise since, after all, future networks will be built upon real mobile wireless access networks. Far beyond typical technical only NQoS demands, users' satisfaction should

be addressed as the final target for any network management mechanism. In order to do so, Quality of Experience (QoE) could be used as the final metric in order to guide the design process of Future Internet while isolating users from all low level details and complex NQoS metrics.

We have shown the relevance of our QoE approach by analysing the results of a comparison between VoIP services over different UMTS accesses and how carefully selected low level parameters could lead to equivalent users' QoE with lower resources consumption. So, in both simulated cases users' will not notice any difference and network will cope with its responsibility of providing users with highest *quality*.

Finally, Future Internet should be designed by incorporating mechanisms to provide service-, user-, content-, terminal- and network- aware capabilities targeted at guaranteeing users' QoE in a flexible and service-dependent way, beyond the bit pipe approach. This will demand not only more intelligence in network nodes but also content and user preferences describing languages and world scale NQoS management schemes (including QoE-aware, service dependent and cross-layer request mechanisms and evolutioned inter-provider SLAs).

The research leading to these results has received funding from the European Community's Seventh Framework Programme FP7/2007-2013 under grant agreement n<sup>o</sup> 214751/ /ICT-ADAMANTIUM/.

## References

- [1] J. Jin and K. Nahrstedt. QoS Specification Languages for Distributed Multimedia Applications: A Survey and Taxonomy. *IEEE MULTIMEDIA*, pages 74–87, 2004.
- [2] G. Dobson, R. Lock, and I. Sommerville. QoSOnt: a QoS Ontology for Service-Centric Systems. In *Software Engineering and Advanced Applications, 2005. 31st EUROMICRO Conference on*, pages 80–87, 2005.
- [3] CORDIS. Ec cordis website: <http://cordis.europa.eu/>.
- [4] G. Carneiro, C. Garcia, P. Neves, Z. Chen, M. Wetterwald, M. Ricardo, P. Serrano, S. Sargento, and A. Banchs. The DAIDALOS Architecture for QoS over Heterogeneous Wireless Networks. *Proceedings of the 14th IST Mobile and Wireless Communications Summit, June, 22, 2005*.
- [5] R. Ocampo, L. Cheng, Z. Lai, and A. Galis. ContextWare Support for Network and Service Composition and Self-adaptation. *LECTURE NOTES IN COMPUTER SCIENCE*, 3744:84, 2005.
- [6] C. Mensing, E. Tragos, J. Luo, E. Mino, and G.A. Center. Location Determination using In-Band Signaling for Mobility Management in Future Networks. In *Personal, Indoor and Mobile Radio Communications, 2007. PIMRC 2007. IEEE 18th International Symposium on*, pages 1–5, 2007.
- [7] M. Conti, S. Giordano, G. Maselli, and G. Turi. MobileMAN: Mobile Metropolitan Ad Hoc Networks. *LECTURE NOTES IN COMPUTER SCIENCE*, pages 169–174, 2003.
- [8] E. Exposito, M. Gineste, L. Dairaine, and C. Chassot. Building self-optimized communication systems based on applicative cross-layer information. *Computer Standards & Interfaces*, 2008.
- [9] S. De Zutter, M. Asbach, S. De Bruyne, M. Unger, M. Wien, and R. Van de Walle. System architecture for semantic annotation and adaptation in content sharing environments. *The Visual Computer*, 24(7):735–743, 2008.
- [10] M. Wien, R. Cazoulat, A. Graffunder, A. Hutter, and P. Amon. Real-Time System for Adaptive Video Streaming Based on SVC. *Circuits and Systems for Video Technology, IEEE Transactions on*, 17(9):1227–1237, 2007.
- [11] I. Rec. G. 107-The E Model, a computational model for use in transmission planning. *International Telecommunication Union*, 2003.

## A Future Perspective on the 3D Media Internet

Petros Daras<sup>1</sup> and Federico Alvarez<sup>2</sup>

<sup>1</sup> Informatics & Telematics Institute, 6th km Charilaou-Thermi Rd,  
PO BOX 60361, 57001 Thessaloniki, Greece

<sup>2</sup> Universidad Politécnica de Madrid. E.T.S.Ing. Telecomunicación, Ciudad Universitaria s/n  
E-28040, Madrid, Spain

[daras@iti.gr](mailto:daras@iti.gr), [fag@gatv.ssr.upm.es](mailto:fag@gatv.ssr.upm.es)

**Abstract.** The Internet is a living, dynamic “wizard” who is constantly and rapidly evolving, reshaping and transforming and consequently it is changing our social and economic world. However, Internet was designed and primarily used by scientists for networking research and for exchanging information. The current form of Internet cannot efficiently serve future needs raised by the exponential increase of the available cross-modal content and ensure consistency and correctness in terms of media storage, distribution and consumption, along with options to navigate, interact, search and retrieve this content. In this paper we make an attempt to identify the main barriers posed by the Current Internet and analyse major issues and challenges which is expected to be the pillars of the Future 3D Media Internet.

**Keywords:** Future Media 3D Internet, Future Content Network Architectures, Immersive Environments.

### 1 Introduction

For many years, the Internet was primarily used by scientists for networking research and for exchanging information. Remote access, file transfer, and e-mail were among the most popular applications, and for these applications the datagram model works well. The World Wide Web (WWW), however, has fundamentally changed the Internet. It is now the world's largest public information network. Many applications—such as video conferencing, Web searching, electronic media, discussion boards, e-commerce and Internet telephony—have been developed at an unprecedented speed. Similar to the explosion of textual content in the Internet, one can easily observe a dramatic increase of network-based audiovisual material (networked media) that has been produced by professional and amateur users.

This explosion has been triggered by the widespread availability of broadband home access, digital recording devices, improved modelling tools, advanced scanning mechanisms as well as display and rendering devices. Today, over one billion of users access the Internet on regular basis, more than 100 million users have downloaded at least one (multi)media file and over 47 millions of them do so regularly, searching in more than 160 Exabytes of content [1]. In the near future these numbers are expected



to exponentially rise. It is expected that the Internet content will be increased by at least a factor of 6, rising to more than 990 Exabytes before 2012, fuelled mainly by the users themselves. Moreover, it is envisaged that in a near- to mid-term future, the Internet will provide the means to share and distribute (new) multimedia content and services with superior quality and striking flexibility, in a trusted and personalized way, improving citizens' quality of life, working conditions, edutainment and safety [2].

In the longer term, the exponential increase of the user generated multimedia content and the number of mobile users will raise many new challenges. In this respect, the Future Media Internet will not simply be a faster way to go online. It will be redesigned so as to overcome current limitations and to address emerging trends including: new network architectures, content and service mobility, diffusion of heterogeneous nodes and devices, mass digitisation, new forms of (3D) user centric/user generated content provisioning, emergence of software as a service and interaction with improved security, trustworthiness and privacy.

In this evolving environment, machine-to-machine communication (including RFIDs and audiovisual sensor networks), rich 3D content as well as community networks and the use of peer-to-peer (P2P) overlays are expected to generate new models of interaction and cooperation. Furthermore, they will be able to support new innovative applications "on the move", like virtual collaborative environments, personalised services/media, virtual sport groups, on-line gaming, edutainment, etc. In this context, the interaction with content combined with interactive/multimedia search capabilities across distributed repositories, opportunistic P2P networks and the dynamic adaptation to the characteristics of diverse mobile terminals are expected to contribute towards such a vision. On the other hand, advances in scalable video coding and 3D video processing, dynamically adapted to the network conditions will give rise to innovative applications such as massive multiplayer mobile games, digital cinema and in virtual worlds, placing new types of traffic demands and constraints on mobile network architectures.

In this paper an attempt is made to collect, review and evaluate the related existing technology with the aim to give research directions towards redesigning the Future Internet, having in mind the obstacles posed by the Current Internet, and the provisioned new applications. Thus, the main contribution of this work is to present, in a coherent way, a view on how 3D Media Internet is likely to look like and which could be its main pillars. Furthermore, the long-term vision is given containing possible research directions which is expected to contribute to the realisation of the Future 3D Media Internet.

The paper is organised as follows: In Section 2 the limitations and barriers of the Current Internet (CI) are analysed while in Section 3 the provisioned characteristics of the Future 3D Media Internet are given. In Section 4 the long-term vision along with possible research directions are sketched and finally, in Section 5 conclusions are drawn.



## 2 Limitations and barriers of the Current Internet

Since the main parts of the CI were developed 30 years ago for serving research demands (host-to-host communications), it is obvious that it cannot be used with the same efficiency today where new demanding applications rise.. What follows is an analysis of the obstacles and limitations of the CI, with respect to media and network, which put current and future uses at risk [3].

### 2.1 Limitations with respect to content

Nowadays we are rapidly moving from a mainly textual-based to a media-based Internet, where rich audiovisual content, 3D representations, virtual and mirror worlds, serious games, lifelogging applications, etc. become a reality. In this environment it is obvious that there is a need to support the *experience* in form of real enjoyment of these media, in the sense of having true *interaction* with both the people and the media. This lack mainly happens due to the restrictions imposed by current limitations in network reliability in terms of bandwidth, as a physical constraint, and the consequent delays that are imposed.

This plethora of media clearly generates the need for algorithms and tools to ease its manipulation (i.e. automatic or semi-automatic media annotation, indexing of media in large databases, visualisation of media in heterogeneous terminals, etc.) and efficient generation, sharing, search and retrieval of media based on both content and context.

Although popular attention has highlighted the phenomena of end-user generated content, it is commonly overlooked that in the majority of existing examples, user activity has typically been bootstrapped and aggregated around a pool of existing, professionally produced, and often commercial content (e.g. YouTube). This reinforces the role of professional content producers and distributors and, at the same time, requires rethinking the role of users and enabling technologies: though users are not the sole or primary source of content, they are active co-creators and participants who need to be integrated in co-creation and exploitation within the media value chains.

Of great importance are also tools able to provide real collaborative environments so as to boost the productivity of the “creative” community and bring together professional and amateur media creators. Nowadays, still in the minds of successful media producers there exist two separate worlds: the amateurs collaborate in the Web2.0 and the professionals compete in the commercial media world. Thus, there is a clear need for a change in the professional work practices and in attitudes and novel networked working environments in future networks in order to collaborate in the professional creative media sector in order to compete on the global content market.

The CI is characterized by a disembodied and non-multimodal access to content. Interaction with content lacks inaction and immersive participation, apart from very specific cases, usually intrusive and not easily accessible to a wide range of users. The role of sound can be strategic to improve immersivity, to enable inaction in the interaction with content and embodied content processing in Future Internet (FI) applications. Sound has an important role in our life: “embodied sound media” and

applications can contribute to improve both novel A/V services (e.g. experience-centric and participative music applications, e.g. future active listening applications) and support applications for industry and also to improve cultural and social (e.g. health, elderly) situation in EU. The lack of embodiment in current Internet could be faced by enhanced support of multimodality, including sound, haptics, visual, gestural, physiological, toward a deeper exploitation and integration of communication and interaction through the physical, non-verbal, full-body channels.

The support to social and emotional communication among users and communities is another aspect that is lacking in the CI. There are initiatives (e.g. W3C, MPEG) to code emotional cues, and to define use cases significant for future developments. However, they are mostly based on the current paradigm of research on emotions, based on emotional labels and explicit modeling of affective states. It is suggested here a widening of paradigms, aimed at empowering the role of emotional interaction in future Internet scenarios: this consists of facing the modeling of non-verbal subtle communication channels, i.e. strongly related to empathy, engagement, synchronization, or, mentioning Japanese culture, “kansei information processing” [4]. As an example, let us consider the emotions and empathy communication between two or more humans standing still in a waiting room (e.g., waiting for a medical treatment). Even without exchanging a single word, they exchange a lot of non-verbal information (emotional, empathic, etc.). This subtle, and at the same time very powerful, aspect of communicating through empathy and emotion should be tackled, modeled, and exploited in future Internet applications, achieving more effectiveness and immersiveness and contact among humans. This also contributes to social networks issues in e2e platforms, and raises important ethical issues. In short, the focus here is towards more “sensible” (and therefore effective in experience-centric and participation of users) Future Internet applications.

To sum up, FI related applications will involve both professional and user generated content able to be used for enhancing interactivity and social and emotional communication between users. To this end, it is obvious that the Future Media which will emerge will require significant research efforts for processing, manipulation and (re)use. This is expected to lead to a new Future Media era in the body of FI.

## **2.2 Limitations with respect to content delivery networks**

The limitations of the current content delivery networks regarding the FI are not based only on the limited capacity or the lacking of IPv6 deployment, but more importantly on the introduction of a new convergent structuring of the networks for 3D media delivery. These content oriented networks should provide adaptation, contextual services and users social and personal network focus, rather than only the service provider orientation.

But which limitations are really in the field today, for content delivery networks? On the one side we can consider the pure network aspects for providing good Quality of Services (QoS) to the final users such as reliability, recovering, convergence time of routing protocols, interruption of the service, etc. It is clear that new applications, based on Future Media, such 3D IPTV, 3D Telepresence or 3D Media applications in general, will need higher bandwidth networks, in mobile or nomadic environments.

Besides, QoS based on dynamic bandwidth allocation or based on Service Level Agreements (SLAs) are simple options when the Perceived Quality of Service (PQoS) comes into play, or the Quality of Experience (QoE). In these situations no simple metrics or mechanisms can be used. But looking further from these simple conclusions, there are limitations which are not only related to pure bandwidth or address limitations.

Furthermore, the users are nowadays creating contents, such as collaborative and user-generated content regardless of the network used (wired, wireless) or the device to produce and visualise the content. The current networks are not adapted to the user-centric media characteristics (especially in the mobile networks) and only some efforts have been devoted (e.g. P4P networks) to address this need. The evolution of the content delivery network should allow new ways of collaborative and 3D applications, also being user context-aware and content aware. Context has been referred in many cases as only as the parameters of the network in the user premises or the user location. User context-awareness should fail more in the QoE of the user combined with the user consumption context and user-generated content creation profiling, isolated and in communities, for a context based on the user context and societal environment, more than extracted from “raw technical network parameters”.

Concerning the inter-domain routing, it is currently determined primarily by business relationships between providers and not by technical aspects of the path – e.g. shortest or widest path. As for the multicast protocols, they enable the dissemination of many to many endpoints in a network efficient way. Even though multicast protocols have been deployed by many providers, a lot of problems exist that hinder the inter-domain multicast deployment. One example is the multicast Peer-to-Peer (P2P) IPTV applications which can solve some limitations of the current unicast P2P networks for media services.

Some other limitations are related to content security, rights and trustworthiness of the content networks, which should evolve in the sense of providing enhanced security mechanisms without disturbing the content delivery and affecting less the system as it currently happens. Anyway this will be a long time discussion, as the users not always agree with securing or protecting content.

In summarising, nowadays the network is used mainly to transport information irrespective of the content's carried characteristics using network protocols to distribute the content. However, FM3DI needs a coupling of network and 3D content features to adapt, enrich and optimal distribute the content to the users and their contexts.

### 3 Characteristics of the Future Media 3D Internet

Taking into account the aforementioned limitations of the CI we present in this section some of the possible characteristics of the FM3DI with respect to both its content and network characteristics.

The content of the FM3DI could be:

- *Intelligent*: able to be adapted to the users with respect to their preferences (personalisation), devices (terminals) and access networks. In order to allow

for a good user experience regarding the media content, this content should be adapted to the user. It should be possible for the user to personalize the media objects by annotating, modifying or creating and sharing it the way they consider appropriate.

- *3D and haptic*: able to be used in many future applications such as realistic virtual/mirror worlds' creation, human representation (avatars), etc. The future media content will be fundamentally 3D (real-time or not) including visual, sound, and other sensorial features such as haptics; it will be able to convey pressure, forces, vibrations and/or motions to the user, as well as physiological or emotional user's state.
- *Interactive* for all different terminals (PC, Set-Top Box, mobile, etc.): The user should be able to interact with the media objects by modifying and/or render them using multiple views and perspectives. Real-time interactivity with other users through the media will be required in order to achieve the maximum level of collaboration.
- *Live or real time (live recording, live performing)*: the most attractive media types tend to be preformed or generated in real-time circumstances. Therefore, FIM would need to facilitate live multimodal media, such as video, events in virtual worlds and live music performances to users and in addition, enable collaboration in distributed environments.
- *Cross modal*: Future media would need to be intuitively inter-linked and accessible. Therefore, they need to support cross modal approaches to media creation, retrieval and consumption. Just as the humans easily identify a song with a film, or smell with particular environment and time, FIM needs to inherently facilitate cross modality of the content and its tasks.
- *Iteratively and cooperatively negotiated in communities of professionals and amateurs*: The content of FIM should not be limited to professional producers creating for consumers, but will be created iteratively and cooperatively in negotiations across multiple communities of professionals and amateurs.
- *Publicly opened and controversial*: FM3DI should not be closed, but open for public participation and even be supportive of establishing communities across controversial issues and incorporating stakeholders with conflicting interests.
- *Collaboratively edited/filtered*: In order to have media professionals making maximum use of the internet, the media content should be edited/filtered/written/manipulated in a collaborative way.

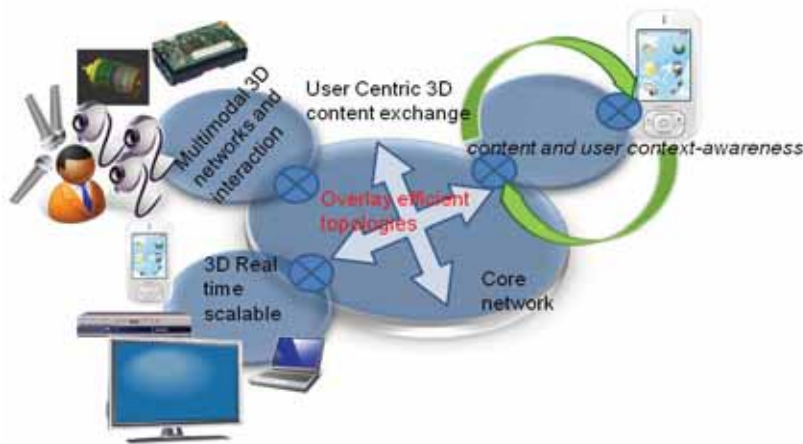
Furthermore the content delivery of the FM3DI network should:

- *Rely on Content/service-centric network*: The CI is service location centric, the main drive of the delivery of information relies on the IP addresses of source and destination of a communication, where the IP address acts as locator and identity of the counterparts. A FI should be content and user-centric, stressing the importance of the content and the user versus the service location.
- *Be able to Transport 3D multimodal media*: The FM3DI will rely on 3D multimodal content and should provide mechanisms for signaling and

describing multimodal 3D media objects (including haptic, tactile and smell information) seamlessly and adapted to the presentation terminals.

- *Integrate real scalable and self-adaptable mechanisms for heterogeneous devices*: More and more specialised and multimodal devices will appear in the market. The *FM3DI content networks* should be able to adapt the content to the user and device characteristics coping with scalability from low resolution to real 3D video multimodal media formats and allow for creation, modification, search and sharing of the new media objects.
- *Be Real time*: Due to a strong demand for real-time quality of Future Media, it will need to deliver media in real-time throughout the whole pipeline of communication: from the source to the user, regardless of the network architecture. The delivery of real-time applications with the proper network parameters required regarding delay, latency, jitter, is vital in order to enable collaborative editing and creation of media objects or real-time participation and interaction in events.
- *Network content and user context-aware*: User context should go beyond network parameters at user premises to real content aware-networks which can provide real-time adaptation and user context personalization of 3D Media heterogeneous services ranging from simple 3D IPTV to real 3D (multimodal) telepresence. User social profiling, consumption, interaction and historical usage should be combined with QoE for producing real Future Internet services.

The network vision described above, is depicted in Fig. 1



**Fig. 1.** Representation of the FM3DI network

The FM3DI network should combine the characteristics mentioned above to both network and content to produce a real “Content Centric Network” enabling on-the-fly content enrichment and adaptation of the content to the network, the user and its context, without disturbing the normal content delivery. In summarising, the Future Network should avoid being only a mean of information transport but should also

support FM3DI applications and services which can enable new 3D personalised experiences to the users, adding a support to future concepts after 3D services will become a reality.

#### **4 Long term Vision of the Future Media 3D Internet and possible research directions**

The Internet is rapidly transforming into a fully fledged virtual environment that facilitates services, interaction and communication. In order to be able to predict the shape of the FM3DI we should take into account two certainties; that the technology used to build the Internet will change, and that the fundamental (human) needs that it must ultimately serve, will not change dramatically. Our needs include the following: we want to be told (or discover) *stories*, we want to share *our* stories, we need to be part of a social community, but have *identity*, we want to *discover* new ‘stuff’ for our community and we want to *play* and *escape*. Therefore, the vision of the FM3DI is to be able to realise these needs.

Towards this aim, multiple regional initiatives are currently emerging in view of defining future global networks. Japan (through the AKARI Architecture Design Project) [5] and Korea [6] have made public their ambitious initiatives, China is supporting the domain through an ambitious and integrated industrial policy, in the US the FIND and GENI programmes [7] and facility is a key contributor to the debate on the future of the Internet and with Latin America there are several ongoing initiatives for identification of opportunities for ICT collaborations [8], [9]. These initiatives are not all tackling the issue of the Internet evolution as part of their core objectives, but are certainly related to technological and socio-economic scenarios (ubiquity, connected devices) that will clearly need to be taken into account when addressing the Internet of Tomorrow.

In Europe, the Future Internet Research and Experimentation (FIRE) [10], the Future Internet Assembly and the Future Media 3D Internet Task Force [11] are the main endeavours that are currently taken place in Europe with the goal to have a unique position in the Future Internet research.

According to the outcome of the aforementioned EU efforts [12], specific attention should be given to both main pillars of the FM3DI, namely the content and the delivery of the content. It is envisioned that the FM3DI will include *interactive/proactive autonomous characters*, where a character is any object in a 3D scene with its own opinions and suggestions, able to understand its environment and take decisions and participate in *3D social communities*, which allow people to use 3D environments to communicate and interact with each other using rich communication means similar to those used in face-to-face meetings (gaze awareness, gestures, facial expressions, correct sound direction, manipulation of social signals).

Also, the FM3DI will allow for *personalised* entertainment supporting *interactive*, and *senses* to be engaged in an *immersive* experience (participation in - or bringing theatre, movies, games) leading to the coexistence of virtual and real worlds maintaining perceptual coherence.



In order for the aforementioned vision to become true the following network characteristics should be fulfilled: higher bandwidth needs to be coupled to new traffic patterns, content adaptation in the network and the terminals that enable the availability of media for a range of heterogeneous devices, new models of content distribution that consider not only the protocols involved but also the social characterization of the nodes and the data, and new network management modules and mechanisms to provide and monitor QoE, trust and privacy (Fig 2).

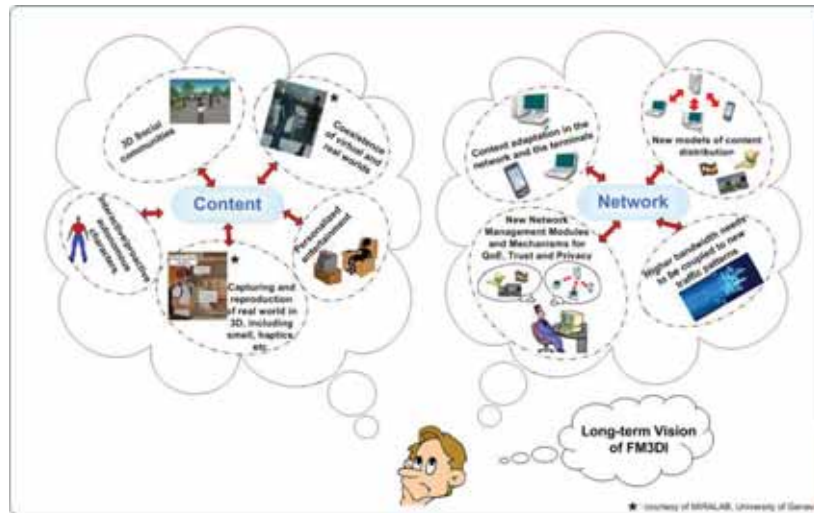


Fig. 2. Characteristics of the FM3DI [12]

These concepts require extensive research endeavours on finding optimum methods for capturing the visual appearance of the real world, including 3D/multiview, high dynamic range and high frame rate. Further, efforts should be devoted on developing rich interfaces allowing for multimodal interaction 3D navigation and strong personalization, on the extensibility, scalability, distribution and availability of the content anywhere, anytime and in any terminal and on new mechanisms for native searching.

## 5 Conclusions

The main goal of this paper was to present our vision for the 3D Media Internet with respect to its two main pillars: the Future Media 3D Internet content and the Future Internet Network Architectures needed for supporting the content advances.

After an in depth analysis of the current technologies and limitations posed by the CI we concluded that more research efforts should be devoted to realise the vision of the FM3DI for both the content and the networks. It became obvious that we need to think “out of the box”, in a more creative way so as to redesign things from a clean slate. Short-term endeavours, like for example the Next Generation Networks, seem not to be adequate for serving the demanding applications of the Future. .



Also, by reporting the envisioned research directions and by sketching the long-term vision for the FM3DI, based on the efforts of the EU and worldwide, we believe that there is a long path to be walked until reaching a real FM3DI complete development. Significant research investment and cooperation between non-EU and EU countries are needed in order a European FM3DI can be in a leading position.

### **Acknowledgments**

This work was supported by the EC funded projects VICTORY, SEA and ARENA.

### **References**

1. IST SAPIR “Large Scale Multimedia Search & P2P”, <http://www.sapir.eu>
2. Zahariadis, T., Daras, P., Laso, I.: Towards Future 3D Media Internet. In: NEM Summit 2008. Saint-Malo, France, October 13-15 (2008)
3. White paper on “Future Media Internet”, European Commission, Networked Media Unit, Information Society and Media, January 2008
4. <http://www.infomus.org/Research/Kansei.html>
5. <http://akari-project.nict.go.jp/eng/overview.htm>
6. [http://mmlab.snu.ac.kr/fiw2007/presentations/architecture\\_tschoi.pdf](http://mmlab.snu.ac.kr/fiw2007/presentations/architecture_tschoi.pdf)
7. [http://www.nets\\_find.net](http://www.nets_find.net) and <http://www.geni.net/office/office.html>
8. <http://www.solar-ict.eu/>
9. <http://www.salamas.eu/>
10. <http://www.cordis.europa.eu/fp7/ict/fire>
11. <http://www.future-internet.eu/>
12. White paper on “Research on Future Media and 3D Internet”, European Commission, Networked Media Unit, Information Society and Media, October 2008

## **Towards an Architecture for a Real World Internet**

Alexander Gluhak<sup>1</sup>, Martin Bauer<sup>2</sup>, Frederic Montagut<sup>3</sup>, Vlad Stirbu<sup>4</sup>,  
Mattias Johansson<sup>5</sup>, Jesus Bernat Vercher<sup>6</sup>, and Mirko Presser<sup>7</sup>

<sup>1</sup> LM Ericsson Ireland, <sup>2</sup> NEC Europe Ltd., <sup>3</sup> SAP AG, Switzerland, <sup>4</sup> Nokia Research Centre, Finland, <sup>5</sup> Ericsson AB, Sweden, <sup>6</sup> Telefonica I+D, <sup>7</sup> University of Surrey, UK  
alexander.gluhak@ericsson.com, martin.bauer@nw.neclab.eu, frederic.montagut@sap.com,  
vlad.stirbu@nokia.com, mattias.a.johansson@ericsson.com, bernat@tid.es,  
m.presser@surrey.ac.uk

**Abstract.** Sensor and actuator networks (SAN) will play a key role in delivering information about and enable interactions with the physical world for next generations of highly autonomous and adaptive Internet services and applications. Current SAN deployments are still low in numbers and represent heterogeneous, vertically closed solutions. True world awareness can only be achieved if existing and future SANs can be integrated into a scalable real world information and interaction fabric, connecting the current Internet with the physical world. In this paper we present architectural challenges, goals and concepts on the way towards such a real world enabled Internet.

**Keywords:** Real world awareness, sensor and actuator networks, future Internet architecture, Real World Internet

### **1 Introduction**

Recent advances in research on the semantic Web and ambient intelligent technology enable more autonomous ways of interaction of computer systems with each other, with humans or with their environment [1]. By integrating semantic capabilities, context-awareness and reasoning mechanisms into the service layer of the Internet, the autonomy of computer systems or networked services can be greatly enhanced [13]. These technologies will serve as building blocks for a new generation of highly distributed services based on efficient information and autonomous machine-to-machine (M2M) interactions, reducing the need of human input.

As more decisions are pushed into the digital world without human involvement, it is very important to increase the accuracy of such autonomous decision making by providing detailed real world information. The success of these ICT-based services will be determined by the quality of available real world information. It is therefore crucial to connect the digital world with the real world in an efficient manner.

Today already a fair amount of real world information is available on the Internet, but entered and classified by humans, which does not by far capture its entire complexity. True real world awareness can only be achieved if information

concerning the physical world can be captured in an automatic fashion, ideally in real time with respect to arising demands.

In order to accomplish such a task, sensor and actuator networks (SANs) will play an important role at the edges of the Internet. SANs represent an inexhaustible resource for real world information. Ubiquitously deployed, they can capture a diverse set of physical phenomena and real world events characterized by multiple sensing modalities and also enable localized interaction [2]. The main challenge is to integrate SANs efficiently into the Internet as natural extensions of the infrastructure and make them universally accessible, while leaving room for emerging applications and not degrading any of the current services. This would enable the transformation of small distributed pockets of sensor data into a vast pool of global, coherent real world information by the means of the Internet.

In this paper we present our first attempt on designing an architecture for a Real World Internet (RWI). Section 2 introduces our vision of the RWI and analyzes its properties and connected challenges. This vision motivates a set of design goals which we present together with a high level overview of our architecture in section 3. In section 4 we take a closer look at the components of our architecture. Finally, in section 5 we aim to complement the picture by highlighting additional functionality that is necessary to realistically support the overall system and by providing an outlook on technological design choices for our architecture.

## 2 The Real World Internet

The RWI will provide an infrastructure that enables augmentation of and interaction with the physical world, without human intervention. The RWI will be composed of a large number of sensors and actuators connected to the Internet that will provide their information in an interpretable form so that it can be leveraged by other systems. Additional components will be required in the network to meet these requirements. As such, the RWI provides true real world awareness to its users and enables both passive and active interactions.

Passive interactions consist in capturing information about the physical world and making it available as contextual information. The primary enablers for this type of interaction are sensors that will feed information about real world phenomena into the RWI infrastructure. There, this information can be further processed before being delivered to interested parties. The RWI also enables active interaction with the real world that may have an impact on its state. Actuators are the primary interfaces of the RWI for this type of interaction, which can range from simple activation up to complex control loops. Both sensors and actuators will enable various forms of interactions and are expected to be deployed at the edges of the RWI. The RWI distinguishes itself considerably from the goals of the current Internet and the World Wide Web. In the following we provide a set of properties linked with the core challenges of the envisioned RWI:

*Number of devices and users* - The number of expected devices will be orders of magnitude larger than in the current Internet. It will indeed require a significant amount of devices to achieve ubiquitous coverage for different forms of interaction

modalities. Traffic will shift from mainly human centric interactions to predominantly machine to machine interactions.

This property illustrates the challenge for scalability as the RWI is expected to grow rapidly [3]. The challenge here is thus to provide an architecture for which performance does not degrade with scale, including the underlying core mechanisms such as discovery, rendezvous, access protocols and management, handling more connected devices as well as human and pure machine based users.

*Heterogeneity of edge devices* - The heterogeneity of edge devices is expected to increase significantly as these will include tiny sensor and actuator nodes and RFIDs.

With the increase of heterogeneity it is becoming significantly difficult to interact with devices in a unified manner [3]. Enabling easy convergence and interoperability of heterogeneous SAN with the Internet will be key to the success of the RWI. Mechanisms that can help are plug and play capabilities [4], which in addition can keep the threshold to enter the RWI low, giving it the ability to grow organically, much like the Internet.

*Information explosion and privacy* - The amount of available information that can be accessed and the diversity of information will increase exponentially [5]. As opposed to the current Internet, a more diverse set of information concerning the real world will be generated autonomously by a large number of edge devices. The increased number of information sources together with the ability of automatic generation will lead to an explosion of available information that may be expressed in various representations (e.g. SensorML [14]). The increasing volume of available data will however come at the expense of security and in particular privacy. The challenge here is thus to provide access to context information and actuation services in a unified manner, while ensuring adequate security properties.

*Importance of metadata* - Information generated by the real world will often make only sense when put into the right context. As available data may not only be limited to raw ones but also include processed or cached ones within the RWI, it is important to tag this information with metadata for further processing [6]. Examples of such metadata include geographic location or quality of information. Although metadata are already widely used on the current Internet, it will become a commodity for information generated within the RWI. One important observation is that the ratio of the size of metadata compared to the information item it is attached to may be larger as is it is currently on the Internet. Efficient means to associate metadata with real world information is a challenge, as well as defining a unified method of providing this metadata.

*“Freshness” of information* - The characteristics of information such as its lifetime may largely vary. Real world information is indeed expected to be much more ephemeral. Despite this ephemeral nature, caching mechanisms may be required in the RWI. The problem of data caching is not novel and has been extensively researched in context of the current Internet [7] [8]. However the sheer amount of data that can be possibly produced in real time by the real world fabric together with its nature will require new mechanisms and policies as to where to store/cache generated real world information and how long to keep it within the system. Such support for efficient information management may have to be deeply embedded in the infrastructure of the RWI, in order to avoid overload from out-dated information.

*Information flows and traffic patterns* - Information flows in the current Internet are mostly end-to-end between two or more end-points. Information flows in the RWI are more complex, often involving several end points that act as sources of information, as sinks, or as intermediate processing points that act as both. Longer-lived flows may change end points or intermediate points along the paths, in order to adapt to changing operational conditions. New types of real world interactions will generate traffic patterns that differ from those common on the current Internet. Shorter payloads and irregular patterns driven by external real world events can be expected [17]. So, existing tenets of the Internet architecture like the end-to-end principle [9] have to be rethought. Recent research on end point-imposed middleboxes [10] [11] and architectures for new naming systems [12] point the way towards further extensions in the more complex context of RWI.

*Mobility* - As in the current Internet the RWI will have to deal with the consequences of mobility of the entities inside the system. This includes supporting the mobility of end-points or whole edge networks, while also dealing with temporary disconnection. Complexity in the RWI further increases as the mobility of real world entities outside the system may directly impact the performance and mechanisms inside the RWI. The challenge is to design mechanisms and protocols that deal with the consequences of mobility. It is expected that existing research [15] [16] can solve many of the problems of node and network mobility, even in the context of SANs. The mobility of entities of interest with respect to the SANs observing or acting upon them, however, is an additional challenge neglected in existing research.

### 3 Architectural challenges, goals and concepts

The challenge of integrating the real world information into the virtual world is not new and has been addressed mainly in two research fields: SANs and Context-awareness [25]. From one side, SANs provide information about sensors, i.e., the type of information they provide, the units of measurement, etc., while Context-awareness provide information about entities of the world and their properties, e.g., like the environmental temperature of a user or the activity she is performing.

These two different ways of conceiving the world led to different architectures that provide the real world information, the usually referred as Sensor Network Frameworks and Context Frameworks. Context Frameworks deal with high-level issues related to context, like how to create, update, maintain or process context; but no or little attention is paid on how the information of the environment is gathered or how to manage the underlying Sensor Network. On the other hand, Sensor Frameworks deal with the open issues that prevent the development of remote application using heterogeneous and geographically dispersed sensor networks, but not much attention is paid to the use of the information.

The successful convergence of Sensor and Context Frameworks is still an open issue in the research community and must be addressed in the definition of a RWI architecture. Our proposal aims at providing both functionalities for the applications in a coherent way: A Context Framework build on top of a Sensor Framework that efficiently meets the requirements of the later, but independent, so multiple context

frameworks could be provided, using different context models or functionalities, sharing the same Sensor Framework.

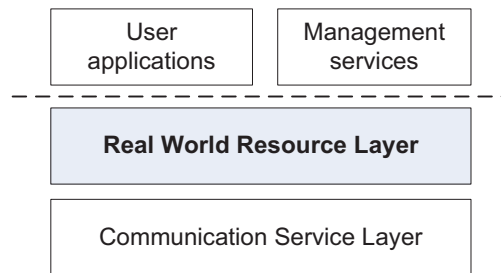
The open research challenges presented in Section 2, have motivated the definition of a set of design goals (summarized in table 1) that an architecture for the RWI should address. The description of properties has provided a comprehensive set of challenges that need to be solved to realize the RWI. These challenges motivate a set of design goals for an architecture of an RWI, that are summarized in table 1. Based on these goals an initial conceptual model of the architecture has been developed. The presented architecture concepts are the first steps of an iterative design process that will eventually lead to the refinement of the overall architecture. Starting with an initial top down specification of architecture concepts, this iterative refinement process is based on the validation of concepts obtained from feedback through bottom up experience with respective technological design choices.

**Table 1: Design goals for an RWI architecture.**

Design goal	Description
G1: Scalable Internetworking	Supporting efficient Internetworking of a large number of highly distributed service end-points acting as producers and consumers of real world context and actuation.
G2: Horizontalisation	Facilitate the horizontal reuse of sensing, actuation and processing services for a large number of applications.
G3: Privacy and Security	Protect the privacy of and offer adequate security for its participating systems and the entities being observed and acted upon.
G4: Heterogeneity	Accommodate a variety of different (technology, administrative domains) SANs at its edges
G5: Reduced Complexity	Reduce the complexity of accessing sensing and actuation services for applications
G6: Simplicity	Reduce the barrier of participation for SANs and thus facilitate deployment by ease of integration.
G7: Manageability	Permit distributed management of its participating systems and their resources.
G8: Service Differentiation	support service differentiation, in order to ensure predictable system behavior according to agreed service levels among participants despite changing system conditions
G9: Continuity	Ensure that requested services are provided with adequate quality, despite change of availability due to loss/disconnection or mobility of system entities.
G10: Evolvability	The architecture must be evolvable to withstand technological change forced upon by tussles [20] carried out by actors in the eco-system.
G11: Locality	Support local operation of the system, in case of failure or disconnection (exploit locality, self-sufficient).

The conceptual model separates the overall architecture into a real world resource layer and an underlying communication service layer, forming the connectivity

substrate of a current or Future Internet. This separation has been found useful as it allows the initial architecture design to focus on functionality without making too many assumptions on the underlying connectivity functions. The conceptual architecture of the real world resource layer can then be either mapped to a connectivity substrate based on evolutions of the current Internet for near-term deployment or incorporated into more revolutionary clean slate designs of the underlying network infrastructure.



**Fig. 1.** High level overview.

Figure 1 depicts a high level view of the conceptual architectural model. Applications and services gain unified access to real world information and interaction capabilities via interfaces provided by the real world resource layer.

In contrast to the current approach of the Internet, the architecture for the RWI aims for inclusion of system management services in a unifying manner. The abstractions required for the retrieval of management information from network elements or the configuration of those are likely to be similar to those for interacting with the sensors and actuators. For example one can imagine the reading of a state variable of a network element as reading a “virtual” sensor that observes the state of the network element or other managed objects associated with it. A similar analogy can be established between the invocation of an actuations service and the setting of some configuration parameters on a network element or the triggering of a management control loop in the network.

It is expected that most of the initial design for the Real World Internet will concentrate on the resource layer. Once the functional components of the real world resource layer are clearer and the interactions understood, detailed requirements for the connectivity service layer can be derived. It should then become apparent what functionality is expected from the connectivity substrate, in order to effectively support the information flows generated by the real world resource layer. A discussion of the requirements on connectivity service layer is beyond the scope of the work presented here.



## 4 Components of the Real World Resource Layer

In the following section we briefly present a first functional decomposition of the envisioned real world resource layer. We introduce the concept of a resource endpoint (REP) as one of the fundamental design considerations of our architecture. Building upon this concept, we detail a first set of functional components of the proposed real world resource layer. In [19] we have presented some rationale for the split of components from the principle of evolvability of our architecture. Here we concentrate on describing the functional components and the way they are expected to address our design objectives.

### 4.1 Resource concept

According to the design goals, the real world resource layer is expected to incorporate a large diversity of real world information sources and interaction capabilities into a homogeneously accessible RWI fabric. This fabric needs to enable horizontal reuse and simple access for applications and services, at the same time, keeping the barrier for participation as an element of the fabric low and also to enable the distributed management of these elements.

An essential step towards these objectives is to provide a unifying abstraction for accessing of and interaction with the elements of this real world Internet fabric. In our proposed architecture this abstraction is provided by the concept of the REP with which real world resources expose their functionality within the resource layer.

The idea of introducing an abstraction layer that hides the underlying heterogeneity by implementing a uniform access is not new in the SAN domain. For example, the Sensor Network Middleware [18] provide such abstractions at a programming level. Our notion of resource differs from others since it provides a logical abstraction for the services provided either by sensors, sensor networks, processing services, etc.

A (real world) resource within the context of the resource layer is any physical or virtual elements that can provide directly or indirectly information related to the real world or enable interaction with it. Examples of such resources are sensors and actuators of a SAN island that can provide direct information of the real world. Example resources providing indirectly real world information are processing services that may combine information with possibly intermediate processing of data (e.g. aggregation, fusion, inference etc.) from multiple resources, e.g. sensors to deliver some high level context of the real world or that realize control loops by means of other sensors and actuators.

In order to become visible for resource users, a resource needs to be exposed as a REP within the real world resource layer. A REP is a logical construct that encompasses unified interfaces for the access and the discovery of a resource.

Resources within the resource layer can be composite resources. This means they can be resource users of other REPs, in order to create a new resource that is offered as a new resource end point in the resource layer.

Typically providers of resources choose the way resources are exposed, by modeling a REP at a desired level of abstraction and defining policies for its discovery and access. For example a provider of a resource may choose to expose

temperature sensors within a building only as a “collective average of the past hour” accessible for all resource users in the resource layer, while selected resource users, e.g. administrators, may gain access to live readings of each individual temperature sensor in the building. According to the discovery policies only authorized users may actually discover that individual temperature sensor in each room may exist as REPs in the resource layer.

The concept of REP is not limited to the primary capabilities of real world resources themselves (providing real world data or actuation). The same abstraction can be applied to expose internal state and mechanisms concerning the management of network elements acting as hosts of these resources. Examples are variables and managed objects containing information about network elements or the whole network or parameters and mechanisms influencing their operational behavior.

## 4.2 Rendezvous

One of the main design principles that we applied to the real world resource layer is the one of loose coupling. Loose coupling supports several of the design goals by increasing the flexibility of interactions within the real world resource layer that is necessary for the evolvability of the system (G10) and continuity of required services (G9), at the same time reducing the complexity of interaction with the resources (G5).

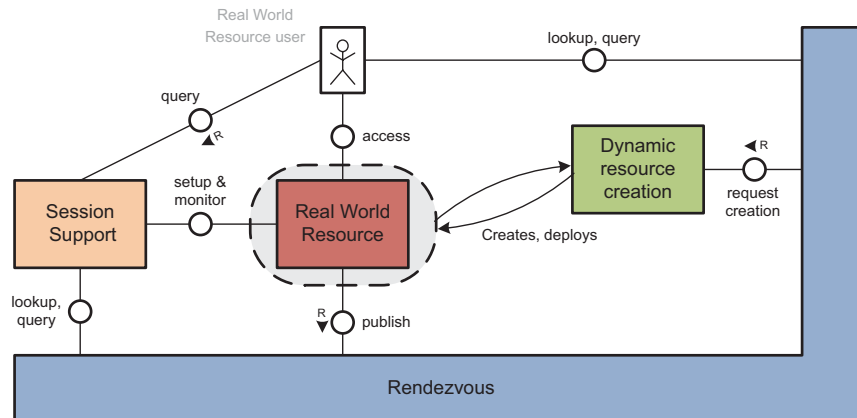
Applying the loose coupling principle requires minimizing the inter-dependencies of resources and their users within the real world resource layer. As a result resource users do not need to know the actual resource that can deliver the desired service (context information or actuation) at design time. Instead they only require knowledge of the nature of the desired service, and rely on the real world resource layer to provide an REP at runtime.

The rendezvous component acts as facilitator for loosely coupled interactions within the real world resource layer. It enables a late binding of resources, by allowing resource users to resolve desired REPs at run time. The rendezvous component represents an additional policy enforcement point in the real world resource layer, as it controls the dissemination of knowledge on resources.

Figure 2 shows the rendezvous component in relationship to other high level components of the architecture. REPs register with (or become discoverable by) the rendezvous function via their discovery interface. Resource users can query the rendezvous for a particular (set of) REP(s) and look up their locations for subsequent access.

The rendezvous component is realized by two logical functions, a resource discovery (acting as part of the logical Sensor Framework) and a semantic query resolution (acting as part of the Context Framework).

The resource discovery function provides a low level lookup service of the location of a REP.. The function consists of two sub-functions, one allowing the registration of REPs with the conceptual resource pool, and one allowing the lookup of REPs by users based on their respective access credentials.



**Fig. 2.** High level components of the real world resource layer.

The semantic query resolution function provides a high level rendezvous mechanism that enables resource users to discover adequate resources based on high-level declarative queries, based on context information. The query analyzer sub-function examines the high level declarative query and determines with the help of the task planning sub-function a (set of) resource(s) that satisfy the query. In order to complete the resolution steps, the semantic query resolution leverages the resource discovery function.

### 4.3 Dynamic resource creation

Existing resources in the Real World Internet architecture service layer may not be directly able to satisfy all requests for context information or actuation tasks. This would leave the burden on the resource user to find resources for partial answers and do the necessary composition and/or the additional processing required. In order to reduce the complexity on the resource user (G5), a dynamic resource creation component is introduced in the system.

The dynamic resource creation component is able to dynamically create and instantiate a new resource (on the fly at run-time) that is able to satisfy the query, based on code fragments or processing components obtained from a repository. It is invoked by the semantic query resolution, if no adequate resource can be determined for a high level query. This significantly reduces the complexity on the resource users, as they can simply interact with the new resources via the resource access function. At the same time, a newly created resource may be made available to other resource users in the system, contributing to the design goal of horizontalisation (G2).

### 4.4 Session support

Short-lived queries for context information and actuation tasks, e.g., one shot queries, are typically not strongly influenced by system dynamics. This assumption however is not valid for longer lasting queries, be it for continuous or event based queries. The

resource availability may often vary over longer periods of time. As a consequence a resource or the initially selected (set of) resource(s) may not be continuously available for such a query.

In order to ensure continuity of an accepted longer lasting query in the light of such system dynamics, the composition of resources may have to be adapted during the query life time (G9). One way of dealing with such situations is to introduce more autonomous intelligence at resources and resource users that allows the detection of changes in resource availability and the ability of self-healing the resource composition when required. While such decentralized approach is highly scalable it requires additional complexity at the resources and their users.

An alternative is to provide dedicated system support to deal with longer-lasting system interactions. The session support component comprises a set of functions that provide resource users with a convenient way of handling such adaptations, reducing some of the complexity required at resources and resource users' sides. A request management function keeps track of incoming queries and performs an initial resolution of the query into (a set of) appropriate resource(s) using the semantic query resolution component. If the request is longer lasting, it sets up a session between the resource user and the resource(s) possibly making use of session monitoring functions to trigger adaptation when resource availability changes in the system. The session support provides only control functionality. Resource users indeed communicate directly with respective resources, contributing to the overall scalability of the system (G1).

## 5 Completing the Picture

The previous sections have only introduced an initial set of components of our envisioned architecture. A complete architecture for a Real World Internet requires consideration of a larger set of features, in order to holistically address the envisioned design goals and requirements. Among those are the management of identities within the real world resource layer and solutions to enforce access control policies and to secure communications, which we will briefly discuss in the following.

The architecture design defines two main policy enforcement points (PEP) in order to protect the access to resources. The first one is deployed within the Rendezvous component. We consider that a policy specified by the resource provider is attached to resources when they are published to the Rendezvous component. This policy defines the set of credentials that principals should satisfy in order to be able to retrieve the contact information of the associated resource. This access control mechanism is only a privacy-preserving mechanism as it limits the number of peers able to contact resources. A second PEP deployed at resource providers' side is therefore required to protect resources themselves. It ensures that only authorized peers are able to access functionalities of resources. Solutions ranging from assertion validation [21] to certificate-based authentication mechanisms [22] can be leveraged in order to implement this second PEP.

Efficient management of identities is another concern that has to be addressed within the real world resource layer. The heterogeneous deployment of resources on a

global scale will make adaptation of a uniform identity management scheme highly unlikely. The Uniform Resource Identifier (URI) concept has proven to be a successful means to handle heterogeneous identities of resources in the World Wide Web. The use of URIs to identify resources within the Real World Resource Layer could enable us to easily integrate multiple identification management frameworks and facilitate the creation of globally unique identifiers, while providing a smooth integration into the existing Internet architecture.

In realizing the Real World Resource Layer, we argue for an opportunistic approach that takes into consideration the characteristics of the computing environment and the communication substrate in the target deployment. Our architecture has to be mapped onto specific technological realization, based on the envisioned deployment environment. For example RESTful services are an obvious design choice for the realization of the Real World Resource Layer in the current Web, while enterprise environment would use widely deployed Web Service technologies [23]. An obvious design choice for mobile operator would be based existing IMS infrastructures [24].

Looking forward, the Real World Resource Layer will be integrated with the communication substrate of the Future Internet. Keeping the entry level for enabling interaction with the real world remains low, could lead to an explosion of Real World Resource Layer compliant deployments that provide the critical mass for the emergence of the RWI.

Deploying a truly global Real World Resource Layer is a challenging task that will take years to complete and special considerations need to be paid to facilitate the evolution of the system. Starting from the initial isolated deployments, we will witness an evolutionary process that will lead towards more complex hierarchical, peer-to-peer or federated structures, before the emergence of the Real World Resource Layer will span the whole globe.

### **Acknowledgements**

The research leading to these results has received funding from the European Community's Seventh Framework Programme (FP7/2007-2011) as part of the SENSEI project under grant agreement n° 215923.

### **References**

1. D. Preuveneers and Y. Berbers, "Encoding Semantic Awareness in Resource-Constrained Devices", in *IEEE Intelligent Systems*, Volume 23, Issue 2, March-April 2008 pp 26-33
2. M. Presser, A. Gluhak, S. Krco, J. Hoeller and L. Herault, "SENSEI - Integrating the Physical with the Digital World of the Network of the Future", 17th ICT Mobile Summit, Stockholm, Sweden, 10-12 June 2008
3. D. Clark et al., "Making the World (of Communications) a Different Place", End-to-End Research Group, IRTF, ACM SIGCOMM Computer Communication Review, Vol. 35, No.2, pp. 91-96, July 2005.
4. V. Stirbu, "Towards a RESTful Plug and Play Experience in the Web of Things", 2008 IEEE International Conference on Semantic Computing pp. 512-517

5. M. Kitsuregawa, "'Socio Sense' and 'Cyber Infrastructure' for Information Explosion Era": Projects in Japan", *Lecture Notes in Computer Science*, Volume 4443/2008
6. M. Botts, G. Percivall, C. Reed, J. Davidson, "OGC Sensor Web Enablement: Overview and High Level Architecture", *Open Geospatial Consortium, Inc Whitepaper*, OGC 07-165
7. A. Vakali and G. Pallis, "Content Delivery Networks: Status and Trends", *IEEE Internet Computing*, IEEE Computer Society, pp. 68-74, November-December 2003.
8. A. M. K. Pathan and R. Buyya, "A Taxonomy and Survey of CDNs", *Technical Report*, GRIDS-TR-2007-4, The University of Melbourne, Australia, Feb. 2007.
9. J. Saltzer, D. Reed, and D. Clark, "End-To-End Arguments in System Design". 2nd International Conf on Dist Systems, Paris France, April 1981.
10. H. Balakrishnan, K. Lakshminarayanan, S. Ratnasamy, S. Shenker, I. Stoica, and M. Walfish. "A layered naming architecture for the Internet." In *Proc. of ACM SIGCOMM*, Portland, OR, USA Aug. 2004.
11. M. Walfish, J. Stribling, M. Krohn, H. Balakrishnan, R. Morris, and S. Shenker. "Middleboxes No Longer Considered Harmful." In *Proc. of OSDI 2004*, pages 215-230, San Francisco, CA, USA, Dec. 2004
12. T. Koponen, M. Chawla, B. Chun, A. Ermolinskiy, K. H. Kim, S. Shenker and I. Stoica, "A data-oriented (and beyond) network architecture.", *ACM SIGCOMM Computer Communication Review*, Vol. 37, Issue 4 (Oct. 2007), pp. 181-192.
13. T. Berners-Lee, M. Fischetti, "Weaving the Web: Origins and Future of the World Wide Web", Britain: Orion Business, ISBN 0-7528-2090-7, 1999.
14. OpenGIS Implementation Specification, "OpenGIS Sensor Model Language (SensorML)".
15. Novaczki, Bokor, Imre, "A HIP Based Network Mobility Protocol," 2007 International Symposium on Applications and the Internet Workshops (SAINTW'07), 2007
16. C. Bernardos, A. de la Oliva, M. Calderon, D. von Hugo, H. Kahle, "NEMO: Network Mobility Bringing ubiquity to the Internet access", *IEEE INFOCOM 2006*
17. Krnic J., Krco S., "Impact of WSN Applications' Generated Traffic on WCDMA Access Networks", *Proc. of the IEEE PIMRC'08*, Cannes, France, September 2008.
18. M. Wang, J. Cao, J. Li, and S. k. Dasi, "Middleware for Wireless Sensor Networks: A Survey," *Journal of Computer Science and Technology*, vol. 23, no. 3, pp. 305-326, May 2008.
19. T. Baugé, A. Gluhak, M. Presser, L. Herault, "Architecture Design Considerations For The Evolution Of Sensing And Actuation Infrastructures in a Future Internet", *Special Sessions of the 11th International Symposium on Wireless Personal Multimedia Communications - Mobility Challenges in the Future Internet*, Lapland, Finland, September 2008.
20. D. Clark, J. Wroslawski, K. Sollins and R. Braden, "Tussle in Cyberspace: Defining Tomorrow's Internet", in *IEEE/ACM Transactions on Networking*, Vol.13, 3, pp 462-475, June 2005
21. E. Maler et al., *Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML)*. OASIS, September 2003.
22. C. Adams, S. Farrell, "Internet X.509 Public Key Infrastructure: Certificate Management Protocols", *RFC 2510*, March 1999.
23. Pautasso, C., Zimmermann, O., and Leymann, F. 2008. Restful web services vs. "big" web services: making the right architectural decision. In *Proceeding of the 17th international Conference on World Wide Web (Beijing, China, April 21 - 25, 2008)*. WWW '08. ACM, New York, NY, 805-814
24. M. Strohbach, J. Bernat Vercher and M. Bauer, "Harvesting the Power of Ubiquitous Sensor Networks – A case for IMS", *WWRF21, WG2-Service Architecture*, Stockholm, Sweden, 13-15 October 2008.
25. A. K. Dey and G. D. Abowd. "Towards a better understanding of context and context-awareness" *College of Computing, Georgia Institute of Technology*, GVT Technical Report GIT-GVU-99-22, 1999.

# Roadmap for Real World Internet applications – Socioeconomic scenarios and design recommendations

Fabrice FOREST <sup>a1</sup>, Olivier LAVOISY <sup>a</sup>, Markus EURICH <sup>b</sup>, Jilles VAN GURP <sup>c</sup>  
and Duncan WILSON <sup>d</sup>

<sup>a</sup> *Université Pierre Mendès-France - UMAN Lab, Grenoble,*

<sup>b</sup> *SAP Research/ETH, D-MTEC, TIM, Zurich,* <sup>c</sup> *NOKIA Research Center, Helsinki,*

<sup>d</sup> *Ove Arup and Partners International, London*

**Abstract.** This paper emphasises the socioeconomic background required to design the Future Internet in order that its services will be accepted by its users and that the economic value latent in the technology is realised. It contains an innovative outlook on sensing aspects of the Future Internet and describes a scenario-based design approach that is feasible to roadmap the dynamic deployment of Real World Internet applications. A multi-faceted socioeconomic assessment leads to recommendations for the technology deployment and key features of the Future Internet that will globally integrate technologies like Wireless Sensor and Actuator Networks and Networked Embedded Devices.

**Keywords:** Real World Internet, Future Internet, Scenario-based Design, Socioeconomics, Business Models, Requirements

## 1. Introduction

In the context of the current endeavour to create a global architecture as well as application services for the Future Internet, the SENSEI project [1] will be considered as a case to demonstrate how a scenario-based approach and socioeconomic assessment can contribute to the design of the Future Internet (FI). SENSEI (Integrating the Physical with the Digital World of the Network of the Future) is an Integrated Project in the EU's Seventh Framework Programme, in the ICT (Information and Communication Technologies) call: The Network of the Future.

The SENSEI project is designing an architecture that aims at being a key enabler of the real world dimension of the Future Internet [2]. SENSEI's vision is to realise ambient intelligence in a future network and service environment, and to integrate Wireless Sensor and Actuator Networks (WSAN) efficiently into the Future Internet. The system developed by the SENSEI project (of the same name, i.e. SENSEI) is expected to play an essential part in transforming the existing Internet, Mobile

---

<sup>1</sup> Fabrice Forest: Social Sciences Researcher, Université Pierre-Mendès-France, UMAN Lab, BP 47, 38040 Grenoble, France; E-mail : fabrice.forest@umanlab.eu



Networks and Service Infrastructures into a future network capable of dealing with the challenging demands of a future networked society.

In this paper, the benefits of the technology characteristics in SENSEI's Future Internet framework are demonstrated through application scenarios. These scenarios are used as a roadmap for the system and architecture deployment. The application scenarios were not only used to picture a realistic vision of a Future Internet, but were also used to anticipate the evolution of the Real World Internet (RWI). The goal is to enable a Future Internet system that is designed for change, is accepted by all stakeholders involved, and creates value for the users.

To design such a system, a group of experts was gathered to analyse the socioeconomic dimensions of the Future Internet. The experts from different European universities and industrial companies in the ICT sector have backgrounds in computer sciences, engineering, business administration and management. Face-to-face workshops were organised and a wiki was set-up to produce application scenarios, to analyse the sociological perspective, and to assess the business characteristics. The results of this study are summarised in this paper.

In section 2 three application scenarios picture a roadmap for the deployment of Real World Internet services in the urban context. Section 3 explores to which extent this real world dimension of the Future Internet could be beneficial to the user and to society. Section 4 focuses on the business perspective of the dynamic deployment of the Future Internet and its applications. Based on the socioeconomic assessment of the roadmap, the concluding section 5 deduces the recommendations for the technology design and deployment.

## 2. Application scenarios for SENSEI applications roadmap

Eighteen application scenarios in eight application spaces were created by the SENSEI partners to analyse the "real world" dimension of the Future Internet [3]. The elaboration of the scenarios was continuously driven by the Future Internet vision and its key dimensions: Future Networks, Internet of Things, 3D Internet, Internet of Services and Internet of Contents [4]. The scenarios depict FI perspectives from short, to mid, to the long term. In the following case study, three scenarios are analysed to roadmap some RWI applications in the Smart City application space. The initial scenarios [3] – 1° *Smart Places*, 2° *Networked Inhabitants*, 3° *City Information Model* - present RWI applications through concrete usage situations and cover three deployment phases that describe a "now", "new", "next" time horizon.

### 2.1. Now - Smart Places

The first phase scenario describes the RWI applications that benefit to the stakeholders of a shopping mall. This is a realistic and "feasible" scenario in the short term since it involves a set of services provided to the users within a limited area. This scenario is incremental from a technological point of view since it exploits existing communication networks and trivial WSN independent solutions (CCTV, anti-theft gates, temperature, bar code based stock control, etc.). In terms of context awareness, it is based on consumers' real-time geo-positioning, physical condition and dynamic user

profile. From a business perspective, it involves the mall stakeholders federated by a neighbourhood of business interest. From societal perspective it provides services adapted to a user lifestyle that is very similar to the current lifestyle and relationships in western cities.

## 2.2. *New - Networked Inhabitants*

The second phase scenario describes RWI benefits to city-dwellers in mobile lifestyle. This scenario is innovative from the technology perspective since it involves sensing various physical phenomena through heterogeneous communication networks and horizontal exploitation of WSN (see section 5.2). It relies on the improvement of the existing infrastructure: new sensor nodes and WSN are added to enlarge the scope of the sensed phenomena with a continuous quality of sensed and processed information. From a business perspective, additional malls are connected together in order to offer a broader and ubiquitous scope of end-user services. From the societal perspective, it involves a new generation of applications that will impact deeply the users' behaviours, mobility, sense of sustainability and relationships.

## 2.3. *Next - City Information Model*

City Information Model is the most futuristic scenario from a technology perspective. It involves a global coverage and networking of the city by WSN. Horizontalisation is optimised through using and reusing non dedicated WSN for an exponential amount of applications. From a business perspective it facilitates dynamic supply and demand increase as well as just-in-time provision and the trend towards the experience economy [5] reduces stocks in stores and increases uptake of "you shop we drop" services. Micro transactions and demand for behavioural intelligence pulls data provision markets. New business actors such as the transport sector enter the network in order to provide a new generation of services to the citizens and to regulate the prices depending on energy savings and the carbon emissions related to clients' purchases. From a societal point of view, it involves a holistic vision where new types of applications influence the society and new methods for supporting users' behaviours and values are created.

## 2.4. *Scenarios road mapping*

The three phases involve different levels of societal changes, business innovation and technical feasibility. They are not discreet but show a continuous timeline which depends on the context of the actual end use.

The first phase – Now - is *evolutionary* from a societal point of view and *incremental* from technological angle since it is the least integrated: the infrastructure of a mall is used for applications dedicated to the stakeholders of this place.

The second phase – New - is more *futuristic* from the socio economics point of view and *innovative* from the technology side since it implies the deployment of connections between different and separate areas in the city and it starts to integrate different entities in extension to the shopping mall, e.g. private residential WSN infrastructures.

The third phase – Next - is the most *revolutionary* one from the society point of view because it involves holistic applications of RWI. It proposes a fully horizontal

vision of RWI applications with integration of all types of WSAN infrastructures in the city for the provision of an unlimited scope of applications. This is a *disruptive* vision compared to the existing Internet technology.

### 3. Societal analysis of the RWI scenario roadmap

The Life Cycle description of Real World Internet applications enables to analyse some requirements that a RWI system has to challenge to provide benefits to the user and society [6] through FI applications in key domains such as environment, mobility, safety, professional and industrial activities, citizenship, and ethics.

#### 3.1. The Future Internet experience

Real Word Internet will play an essential part in transforming the existing Internet, mobile networks and service infrastructures into a Network of the Future which is capable of dealing with the challenging demands of a future networked society [7]. It will change the user experience of Internet in the sense that the existing distinction between the physical and digital worlds will increasingly blur. Today's Internet will change from the distinct network, providing specific services accessible through dedicated terminals, to an Internet dissolved in the artefacts of the physical world accessible via heterogeneous networks enabling users to browse the world as they browse the Internet.

#### 3.2. Environment

RWI applications must support the reduction of the human impact on the environment. The benefits are highlighted in the *New* and *Next* scenarios and are particularly tangible in city planning, transport schemes and built environment. Beyond the application spaces, a RWI system should be designed according to sustainable constraints:

- The RWI framework should support *horizontal* use and reuse of common WSAN infrastructures to develop a variety of applications. The RWI technology must reduce the necessity of end-to-end and vertical infrastructure dedicated to the delivery of a given application. It should not therefore require as many WSANs as applications.
- RWI system architecture should be *scalable* to enable its functions to *evolve* in order to meet the future requirements of technology changes and growth. It will not be necessary to replace the system to meet the increasing demand of context information and actuation.

#### 3.3. Mobility

RWI applications will support the users' mobility [8], in particular in the city lifestyle and transport organisation. These benefits will be perceptible at short term as exposed in the *Now* scenario but will be at their optimum level at medium term as pictured in the *New* scenario. Beyond application spaces, RWI design principles should integrate mobility in their intrinsic properties:

- The RWI system must ensure the *continuity* of the services that the user needs with an adequate quality despite the user's mobility.
- The RWI framework should *reduce complexity* to enable an easy access of user applications to the sensing and actuation services that are available everywhere. The RWI framework should provide mobile users with a good level of *security and privacy* protection.

### 3.4. Safety and security

RWI applications will improve the users' safety in various activities, in particular in the transport, built environment, crisis management and healthcare domains. These benefits will be perceptible in the short term as shown in the *Now* scenario but will be at their optimum level at medium term as pictured in the *New* scenario. Beyond the application spaces, a key design goal for RWI system should be *security* in order to guarantee the quality and reliability of the context information and actuation.

### 3.5. Citizenship

Integrating the physical with the digital world addresses the socio-economic medium and long term needs that arise through the increasing demand for incorporating ICT in many services for citizenship and quality of life. RWI applications are expected to increase the sense of the community by making perceptible the side effects of individuals' behaviour. In turn, it should improve the quality of services provided to the citizens. These benefits will be perceptible at short and medium terms as it is shown in the *Now* and *New* scenarios. The *Next* scenarios depict holistic applications of WISAN at the level of the entire society that involve deep changes in the citizens' practices (e.g. transport), in the society rules and business activities.

### 3.6. Professional and industrial

RWI will support professional and industrial activities. These benefits will be perceptible at short term as shown in the *Now* scenario. RWI system must support new business opportunities and new industrial partnerships by optimising the integration of sensed and controlled physical phenomena to the Internet:

- In particular, RWI must facilitate the *horizontal* reuse of sensing and actuation services for many applications marketed by diverse business stakeholders.
- RWI *scalability* must support the growth of business needs for sensing and actuation services.
- *Privacy and security* mechanisms must provide the adequate level of trust to support the information and actuation trades between the stakeholders.
- The *reduced complexity* of accessing sensing and actuation services for the applications they develop, use or trade must encourage the business actors and newcomers to develop new generation of services.
- The RWI framework should be *evolvable* to support system upgrades and technological changes required by business tussles [9] carried out by actors.

### 3.7. Privacy and Ethics

With the integration of a real world dimension to the Internet, privacy and related ethical issues will increase [10]. Even if RWI technology integrate the appropriate mechanisms, privacy and ethics can persist as critical issues and mistrust may slow the adoption [11] despite it enabling an open and secure market space for context-awareness and real world interaction. To address this challenge, mechanisms for accounting, security, privacy and trust that provide different levels of granularity are essential for RWI participating systems. Access to context and actuation services must be trustable and secure, while the privacy of individuals and corporations are not violated. RWI system must ensure that only authorised services are able to gain access to potentially sensitive private information and critical actuation.

## 4. Business challenges of Future Internet applications

The Life Cycle description of RWI applications is a way to depict the business challenges in deploying the Internet of the Future. The Smart City roadmap can be seen as a paradigm, which highlights two points. First, applications in Smart City domain are already available now. In this respect the stakeholders are already involved in laying the enabling foundation for the RWI vision. Second, a requirement of paramount importance is to integrate the stakeholders in this “runtime” [9] design process. From that point of view the scenarios also aim at considering how the stakeholders will make the system evolving and, in sum, will co-design the system.

To enable applications like those explained for the Smart City domain the Internet will require an evolution in its applications and architecture. The evolution of the Internet goes along with the extension of its reach to mobile users and end-devices with the inclusion of next generation mobile networks. The drivers for this evolution are manifold: Advances in ICT foster the emergence of ever more players in the Internet value network. Different players as diverse as users, developers, manufacturers, operators, service and content providers, and legal bodies are involved in different RWI application spaces, all of them having different business rationales and potentially conflicting interests. Their different ambitions may lead to a permanent process of continuous refinement of the system. Users desire reliable, secure, exciting and easy-to-use services that are reasonably priced with the complexity of the underlying technologies hidden from view. Service providers aiming at making profit, try to exploit the latest ICTs to satisfy users’ needs. Sustainable economic success requires insights into both the technological development and the realization of economic value latent in ICT. ICT must create value for the customer therefore business models must prove their appropriateness as intermediaries between ICT and economic value. The deployment of an infrastructure capable for RWI services is the basis for applications that provide end-user oriented services. The ultimate goal is the creation of end-user oriented services which generate revenue.

Deployment of a vast range of RWI services implies the involvement or influence of many different stakeholders. Deployment of a particular SENSEI service may involve only one or a few. Some of the stakeholders might have direct benefit from the RWI services once they are up and running – others may not. Therefore, it is essential

that the stakeholders require deploying RWI services which do not directly benefit from the RWI service get a recompense for their engagement.

Stakeholders can generally be seen as taking on different business roles. Several key stakeholders are of importance for the development, deployment and operations of RWI services and systems.

First, basic system enablers need to provide an underlying infrastructure. The System Manufacturer constructs and assembles a self sustained system, e.g. a computer or a sensor node. The component manufacturer makes components of systems and provides those to system manufacturers. Component Manufacturers includes sensor manufacturers or computing processor manufacturers.

Second, the process of developing and deploying RWI services and the required system infrastructure for use in service “runtime” involves another set of business roles. These can be directly involved in the process, or have an indirect influence by providing conditions for what services will be needed, or how the services and systems will be provided and operated. The Application Developer designs, constructs, develops and does initial testing of RWI services. The System Deployer deploys the system in the field where it will be functionally operable after it has been appropriately configured. It relies on e.g. non-functional system requirements. Stakeholder examples include a building construction firm doing electrical or appliance installations, or an operator which deploys a targeted infrastructure for RWI service provisioning. Authority bodies could, for example amend or repeal laws, which can have an impact on services developing and deploying.

Third, after setting up the infrastructure and the development of services, several business roles are needed to operate and deliver RWI services. The Application Service Provider provides the end-user oriented services. Examples of end-user oriented services include the provision of context-aware information adapted to personal references, or building monitoring and control services. The Connectivity Provider makes connectivity available between different deployments in the system, for example wide area cellular connectivity for WSN, or connectivity to hosting servers. The Sensor and Actuation Service Broker acts as a broker between users of services and providers of services. The Content Provider gathers, organizes, and presents information that is available from different sources excluding sensors and actuators. Content Providers might specialize on a certain topics, like weather forecast, monitoring and control information, tourist information, or map databases. The WSN Service Provider delivers the services or information from one or several WSN deployments.

An understanding of the roles and interests of different stakeholders is important to provide recommendations for the architectural design of the RWI. The architecture of the RWI must be design in a way that it realizes the economic value latent in technology. The concluding design recommendations will address this issue.

## 5. Concluding design recommendations

The scenario based assumptions of the RWI's impact on society support the business and technology roadmap assessment since it is aimed at understanding the possible arising demands for a system that globally integrates WSA in the context of a Future Internet. Consequently, the architecture of a RWI application platform should be designed in a way that meets the demands of multiple players in the Future Internet: scalability, horizontalisation, privacy and security, heterogeneity, reduced complexity, simplicity, manageability, service differentiation, continuity, and evolvability [2].

### 5.1. Scalability

The architecture must be scalable to support efficient internetworking of an increasing number of highly distributed service end-points acting as producers and consumers of real world information and actuation. The key scalability requirement is to allow billions of sensors and actuators to be active simultaneously. It needs to support dense configurations of WSA with many nodes active in each. This requires efficient usage of radio spectrum to allow for hundreds or thousands of nodes to be active in the same local area and to prevent conflicts with neighbouring WSA on the same frequencies.

### 5.2. Horizontalisation

The horizontal reuse of sensing, actuation and processing services for a large number of applications must be facilitated. A key element in each step of the roadmap is the availability of a wide range of custom and personalised services coming from many vendors. Horizontalisation means that the infrastructure needs to be updated from whichever sensors and actuators are available. Similarly, the way WSA are accessed from services needs to be independent of any infrastructure specifics. In addition, WSA cannot be assumed to be homogeneous and single vendor: RWI application platform needs to provide horizontal functionality accessible to all of these services and that allows them to abstract from the vendor specifics of the infrastructure underneath, which can vary from location to location.

### 5.3. Privacy and Security

Privacy must be protected and adequate security must be offered for participating systems and the entities being observed and acted upon. The context model needs to adapt to the security context of those accessing it. It is also necessary to support resource protection at the network level. Moreover, it is required that communication between nodes in a WSA cannot be interfered with.

### 5.4. Heterogeneity

The RWI application platform architecture must accommodate a variety of different (technology or administrative domains) WSAs at its edges. The applications need to accommodate a very heterogeneous environment with WSAs and devices from many vendors, a wide variety of networking technologies, including: WLAN, 2G, 3G, 4G mobile networks, and Body Area Networks. The context model needs to abstract from



vendor specifics and provide a generalized, vendor independent way of accessing context while at the same time supporting vendor specific features as needed. The network & security infrastructure needs to interoperate across different networks. WSAAN access points need to support nodes from different vendors

#### *5.5. Reduced Complexity*

Complexity of accessing sensing and actuation services for applications must be reduced. The network has to self-configure and make it easy for service implementers to find and interact with resources in the RWI platform. To enable the diverse service ecosystem that the applications need, RWI applications must hide specifics and provide a way to implement services that hides the complexity of the RWI infrastructure.

#### *5.6. Simplicity*

The architecture must reduce the barrier of participation for WSAAN and thus facilitate deployment by ease of integration. Uniform ways to access context information should be supported. Adding resources to RWI application platform as well as accessing existing ones should be straightforward. Service creation in the RWI application platform ecosystem needs to be simple so as to stimulate maximum market adoption.

#### *5.7. Manageability*

The architecture must permit distributed management of its participating systems and their resources should the involved management authorities belong to different administrative domains. Due to the heterogeneity of the infrastructure as well as the large number of different stakeholders, management should be decentralized. A pluggable framework for sensor and actuator networking should be provided.

#### *5.8. Service differentiation*

Service differentiation should be supported to ensure predictable system behaviour in accordance with agreed service levels among participants despite changing system conditions. Context models need to take into account the vastly differing needs of services. Extensibility and management features are required to allow new differentiation services to be integrated. It should be possible to update sensor actuator nodes in the field when new differentiating services emerge. This means that depending on e.g. availability of certain sensors, networks or context information, different services are active or that services provide different levels of quality.

#### *5.9. Continuity*

It must ensure that requested services are provided with adequate quality, despite change of availability due to loss/disconnection or mobility of system entities. This means that it should be ensured that the network adjusts dynamically in case of disrupted services or networks. And finally the system should ensure that WSAAN are robust against e.g. access point and node failures.

### 5.10. Evolvability

The RWI architecture must be evolvable to withstand technological change forced upon by tussles carried out by actors in the eco-system. This means that the management framework should support adding new networks and services. In the field equipment such as sensor/actuator nodes and access points are upgradable with new software.

These design objectives are structuring the technology developments of SENSEI system [12] which is aimed at being a key enabler of Real World Internet applications. However, such scenario-based road mapping and the derived assumptions have to be assessed and consolidated with regard to the real users and business stakeholders' expectations towards Future Internet. To do so, SENSEI project will submit its scenario portfolio to the users and business stakeholders by means of field inquiries. The analysis of their feedback will enable to consolidate and to adjust the initial requirements that are presented in this paper and will be reflected in the final design steps of SENSEI system.

### Acknowledgment

This paper describes work undertaken in the context of the SENSEI project, 'Integrating the Physical with the Digital World of the Network of the Future' ([www.sensei-project.eu](http://www.sensei-project.eu)). SENSEI is a Large Scale Collaborative Project supported by the European 7th Framework Programme, contract number: 215923.

The authors would like to acknowledge the contributions of their colleagues from the SENSEI Consortium, in particular Jan HÖLLER and Richard GOLD, *Ericsson AB*, Richard EGAN, *Thalès Research & Technology (UK) Limited*, Alfonso TIerno, *Telefonica I+D*.

### References

- [1] EU FP7 project SENSEI - <http://www.sensei-project.eu/>
- [2] SENSEI D3.2 – Reference architecture, Public SENSEI Deliverable, November 2008
- [3] SENSEI D1.1 – SENSEI Scenario Portfolio, User and Context Requirements, Public SENSEI Deliverable, October 2008
- [4] J. Schwarz da Silva, "Future Internet Research: The EU framework", *ACM SIGCOMM Computer Communication Review*, Vol. 32, No2, p85-88, April 2007, download: <http://ccr.sigcomm.org/online/files/p85-v37n2p-schwarz-da-silva.pdf>, last access 07.02.09
- [5] J. Pine, J. Gilmore, *The Experience Economy*, Harvard Business School Press, Boston, 1999.
- [6] David Hausheer, Pekka Nikander et al., *Future Internet Socio-economics – Challenges and Perspectives*, FISE White Paper, Future Internet Assembly, November 2008, download: [http://smoothit.org/wiki/uploads/FISE/FISE\\_position\\_paper\\_final.pdf](http://smoothit.org/wiki/uploads/FISE/FISE_position_paper_final.pdf), last access 07.02.09
- [7] P. Mahonen et al., "The Future Networked Society", *EIFFEL Think Tank*, December 2006, download: <http://www.fp7-eiffel.eu/fileadmin/docs/EIFFEL-FINAL.pdf>, last access 07.02.09
- [8] L. Rainie, J. Quitney Anderson, *The Future of Internet III*, Elon University, Pew Internet Project, December 2008
- [9] D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden, *Tussle in Cyberspace: Defining Tomorrow's Internet*, ACM SIGCOMM, Pittsburg 2002
- [10] L. Rainie, J. Quitney Anderson, *The Future of Internet II*, Elon University, Pew Internet Project, September 2006
- [11] e-SENSEI D1.4.1, Report on evaluation of human impact, WP1, September 2007
- [12] A. Gluhak, T. Bauge, V. Stirbu, M. Bauer and M. Johansson, *System Concepts, Interactions and Technical Challenges in the SENSEI System*, ICT Mobile Summit 2008, Stockholm, June 2008

## Context-Aware Systems and Implications for Future Internet

Nigel Baker<sup>1</sup>, Madiha Zafar<sup>1</sup>, Boris Moltchanov<sup>2</sup>, Michael Knappmeyer<sup>1,3</sup>

<sup>1</sup>University of the West of England, UK

<sup>2</sup>Telecom Italia Lab, Italy

<sup>3</sup>University of Applied Sciences Osnabrück, Germany

{nigel.baker, madiha.zafar, michael.knappmeyer}@uwe.ac.uk,  
boris.moltchanov@telecomitalia.it

**Abstract.** The ubiquity of mobile devices and proliferation of wireless networks will allow everyone permanent access to the Internet at all times and all places. The increased computational power of these devices has the potential to empower people to generate their own applications for innovative social and cognitive activities in any situation and anywhere. This wireless connection is not limited to user devices, almost any artefact from clothing to buildings can be connected and collaborate. Furthermore new sensor technologies and wireless sensor networks provides environmental intelligence and the capability to sense, reason and actuate. This leads to the exciting vision of the interconnection of artefacts embedded in our real environment, forming a society of “intelligent things” and “smart spaces”. This paper discusses the main concepts and role that context-awareness and context aware systems will play in this vision and the significance for future networks and future Internet.

**Keywords:** context, context-awareness, sensor networks, future internet, sense, reason, actuate

### 1 Introduction

In the real world being aware of context and communicating context is a key part of human interaction. Context is a much richer and more powerful concept particularly for mobile users and can make network services more personalised and useful. Location and presence are examples of context based services widely deployed today. Harvesting of context to reason and learn about user behaviour will enhance the “internet of services” or “cloud computing” vision allowing services to be composed and customised according to user context. The concept of awareness and context aware applications and systems is a much more difficult proposition. Context awareness refers to the capability of an application, service or even an artefact being aware of its physical environment or situation and responding proactively and

intelligently based on such awareness. Context-aware applications, context-aware artefacts or context aware systems are aware of their environment and circumstances and can respond intelligently. The ubiquity of mobile devices and proliferation of wireless networks will allow everyone permanent access to the Internet at all times and all places. The increased computational power of these devices has the potential to empower people to generate their own applications for innovative social and cognitive activities in any situation and anywhere. This wireless connection is not limited to user devices, almost any artefact from clothing to buildings can be connected and collaborate. Furthermore new sensor technologies and wireless sensor networks provides environmental intelligence and the capability to sense, reason and actuate. This leads to the exciting vision of the interconnection of artefacts embedded in our real environment, forming a society of “intelligent things” and “smart spaces”. This will enable all sorts of innovative interactive pervasive applications. The key denominator in all these applications and systems is that awareness manifests itself from the self property of being able to sense, reason and actuate. A future internet capable of embracing this concept and delivering context aware services to users and artefacts elevates this to a pervasive sensing and acting knowledge network. This would be a network able to make decisions, actuate environmental objects and assist users.

## **2 Context and Situations**

It is a challenging task to define the word ‘context’ and many researchers grapple with the task of creating definitions. Ryan et al. [1] referred to context as the user’s location, environment, identity and time. Dey [2] defines context as the user’s emotional state, focus of attention, location and orientation, date and time, as well as objects and people in the user’s environment. Another common way of defining context was the use of synonyms. Hull et al. [3] describe context as the aspects of the current situation. These kinds of definitions are often too broad. Brown [4] defines context to be the elements of the user’s environment which the computer knows about. Perhaps the most often used definition is given by Dey and Abowd [5]. These authors refer to context as “any information that can be used to characterize the situation of entities (i.e., a person, place, or object) that are considered relevant to the interaction between a user and an application, including the user and the application themselves.” Another common way to classify context instances is to distinguish different context dimensions. Prekop and Burnett [6] and Gustavsen [7] refer to these dimensions as external and internal, and Hofer et al. [8] refer to it as physical and logical context. External (physical) dimension refers to context that can be measured by hardware sensors, for example location, light, sound, movement, touch, temperature or air pressure, whereas the internal (logical) dimension is mostly specified by the user or captured by monitoring user interactions, for example the user’s goals, tasks, work context, business processes, the user’s emotional state or social relationships.

The notion of situation is closely related to the concept of context. Zimmermann [9] defines it as “the state of a context at a certain point (or region) in space at a

certain point (or interval) in time, identified by a name". Therefore it can be considered as a structured representation of a part of the context with direct comparison to a snapshot taken by a camera. Consequently we can consider that situation can be derived from aggregating and refining types of context information. In other words, as remarked by Loke [10], situation can be viewed as being at a higher level of abstraction than context. Taken to the extreme, situation can be considered as "the complete state of the universe at an instant of time". Therefore a situation may comprise an infinite variety of contextual information. Computational and Artificial Intelligence aspects of situation have been widely explored. Henricksen presents a logical and arithmetical model based on object relations [11], Barwise and Perry have developed a situational calculus [12]. Giunchiglia follows a more philosophical approach and sees context as a "subset of the complete state of an individual that is used for reasoning about a given goal" [13]. The key point is that designers can use sensors to capture and build high level context models of parts of the real world then using these techniques recognise and reason about situations.

### 3 Context-Awareness and Adaptation

In computing literature the term context-aware first appeared in [14] (1994). The following year one of the authors, Schilit [15], describes context-aware software as adapting according to the location, identities of nearby people, objects and changes to those objects. The primary goal of a context aware application or service is to be able to change its behaviour in response to a context change. Context-aware applications, context-aware artefacts or context aware systems are aware of their environment and circumstances and can respond intelligently. Adaptation therefore is an essential element of a context-aware system [16]. It is important to note that adaptation can be described in terms of adaptive and adaptable properties. An adaptive system is one that adapts to changes in user-related or environmental situations or context with the explicit goal of automatically assisting users. In contrast adaptability empowers end-users to customise or personalise computer systems according to their individual preferences. In other words it is an adaptable system. Adaptive and adaptable systems are complementary to each other [17] and when used together increase the match between user needs and system behaviour. The property of context-awareness can be applied to all types of applications and systems and as such has been identified as an essential feature of pervasive computing. The essential aspect however is that it enables automatic proactive assistance reducing human intervention. Many context aware applications can provide this automatic assistance by using logical context alone that is stored in profiles, databases or social websites. However with the proliferation of wireless sensor-actuator networks there is an increasing interest in context-aware systems that make use of external (physical) context factors such as location, presence, temperature and light information and interact with the environment. This capability to sense, reason and actuate has the potential to imbue the property of awareness to almost any artefact or object. This leads to the vision of the interconnection of artefacts embedded in our real environment, forming a society of "intelligent things" and "smart spaces".

## 4 Supporting Context Awareness

The development of context aware applications is a complex task because of the need to accommodate for a wide variety of context types and their values, including the ones that cannot be anticipated at the time when the system is designed. In order to handle this complexity many early examples of context aware systems were designed around specific applications and domains. This approach of hard-coding mappings between all possible combinations of context values and corresponding application behaviour is impractical. More importantly it makes context aware systems difficult to extend and almost impossible to introduce new applications without considerable re-engineering to cope with new context types. It is at best extremely demanding to foresee all contexts an application may encounter during its lifetime. Consequently the approach taken is to design a flexible context infrastructure capable of adapting to different applications.

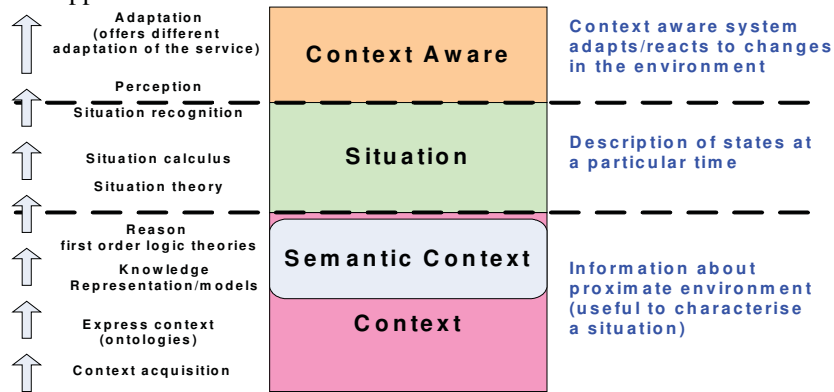


Fig. 1. Context Layers

As illustrated in Fig. 1, there are several layers of abstraction in a context-aware system and any context-aware middleware or architecture must therefore be capable of building representations and models of these abstractions. But these higher level abstractions can only be made from lower level context which requires some form of context management function. The main context management features are context acquisition, context aggregation & fusion, context dissemination, discovery and lookup. Context can be acquired from a diversity of sources from social websites, profiles, databases and physical sensors and filtered and aggregated to form higher level context. Context dissemination is the propagation of context to other entities. Context has a lifetime and must be continuously refreshed. Situation recognition entities and reasoning engines must find or lookup sources of relevant context. Similarly context sources must be able to publish or advertise the context that they have to offer. Distributed context dissemination and discovery requires considerable design effort in larger context aware systems. Establishing context quality is therefore an essential feature of any context management system.

In order to manipulate context information it must be represented in some form that is compatible with the models that will be used in the reasoning and situation recognition processes. These models could be object orientated, ontological, rule

based, logic based, based on semantic graphs or fuzzy logic sets. Expressing context using just one representation is almost impossible since the range is from the most specific, for example a temperature reading, to the most abstract, the state of happiness. Furthermore the representation must lend itself to the reasoning and inference techniques to be used such as classification, taxonomies, data mining, clustering and pattern recognition. Reasoning mechanisms allow higher level context to be deduced or situations to be recognised. Reasoning is also used to check the consistency of context and context models. Finally it must be possible to query context models and context repositories in order match the spontaneous needs of context-aware applications.


## 5 Context-Aware Systems

Engineering a large scale context aware system capable of scaling to the size of an urban cellular network and supporting smart urban type services and applications is a demanding challenge. There are essentially two main parts to any such context-aware system:

- Context Management subsystem concerned with context acquisition and dissemination
- Context Modelling concerned with manipulation, representation, recognising and reasoning about context and situations

Taking into consideration the full sense-decide-actuate cycle of context-awareness then two other subsystems can be identified:

- Sensor and sensor network subsystem: this will include logical as well as physical sensors
- Actuation and Service Composition: Once the system/application becomes aware it adapts by actuating some device/display and or some service is automatically triggered and delivered. This could well be considered as two subsystems.



**Fig. 2.** Context Aware System Layers

There is a natural hierarchy and layering of these subsystems as illustrated in Fig. 2, which captures the complete sense-actuate cycle. This is very much a knowledge layered architecture. Although there are examples that are not, quite often context-aware systems are special examples of knowledge based systems. The processing techniques in such a system range from the fairly simple such as filtering, aggregation, feature extraction to machine learning, knowledge manipulation, rule based processing, situational recognition. Further, different types of context and



situations require different types of processing techniques. Most often the context output of one process can be used as an input to another. Designing generic context-aware infrastructure that will support many different types of distributed context-aware services therefore requires a flexible model. This model must support a collaborative component based approach capable of connecting and distributing context between many context processing entities. A generic model that satisfies some of these constraints is a producer-consumer, publish-subscribe, broker model. This approach is used in project C-CAST and other projects. All context processing entities can either be Context Consumers (CC), Context Providers (CP) or a combination of both. A Context Broker (CB) is required to discover and connect providers with consumers and is responsible for distributing relevant context amongst them. This may be done by query or by publish-subscribe notification methods. Many context processing entities will be both consumers as well as providers of context. Context sources such as WSN gateways, sensor platforms and web sites offering context can be wrapped to provide the common context provider interface.

All context-aware systems must be capable of adapting to context. One of the consequences of context adaptation particularly in a Service Orientated Architecture (SOA) is that adaptation can be partly satisfied by the ability to compose and orchestrate services on the fly, based on user and or environmental context. Web services, semantic web and match making algorithms can play a role in achieving this goal. Service composition is a dynamic and flexible process, which allows for reconfiguration as the context changes. It is not only services that will adapt, displays may change; doors open and traffic lights turn red. This is the actuation part of adaptation indispensable in a future internet of things and artefacts. This raises the issue of how sensor subsystems are best connected to such an infrastructure. For the internet of content then media will be recomposed and advertisements changed. In a distributed system to achieve these outcomes is a formidable challenge. It requires coordination, control and strategic decision making entities.

## **6 State of Practice**

Building upon the context acquisition and context management models discussed in literature, numerous context management architectures have been proposed. Chen [25] presents a Context Broker Architecture (CoBrA), which is an agent-based architecture, for supporting context-aware systems in smart spaces (e.g., intelligent meeting rooms, smart homes, and smart vehicles). Other well-known examples for broker-based approaches are the SOUPA, and GAIA [26]. Many EU projects have explored context-awareness aspects, for example, SPICE [20], OPUCE [21] and MobiLife [19]. However the proposed approaches fail to completely offer a generic, scalable and flexible architecture supporting both evolving context models and evolving services and applications. Moreover, an efficient context diffusion and coherent integration into mobile communication services continues to be a challenging research area. Industrial research already addresses operations on context-aware information such as context capturing and context reasoning in [22]. In [23] Ti-Lab proposed a context management architecture, based on the producer-consumer-

broker model. An application called TeamLife was evaluated on this platform during the Venice Carnival 2008. End users shared contextualised photos on a portal and in real time they were available at several locations of the city on mega-screens [24]. When a picture was taken the application automatically collected related context data, by inquiring the Context Broker [Fig. 3], and proceeds to a seamless machine tagging of the image.

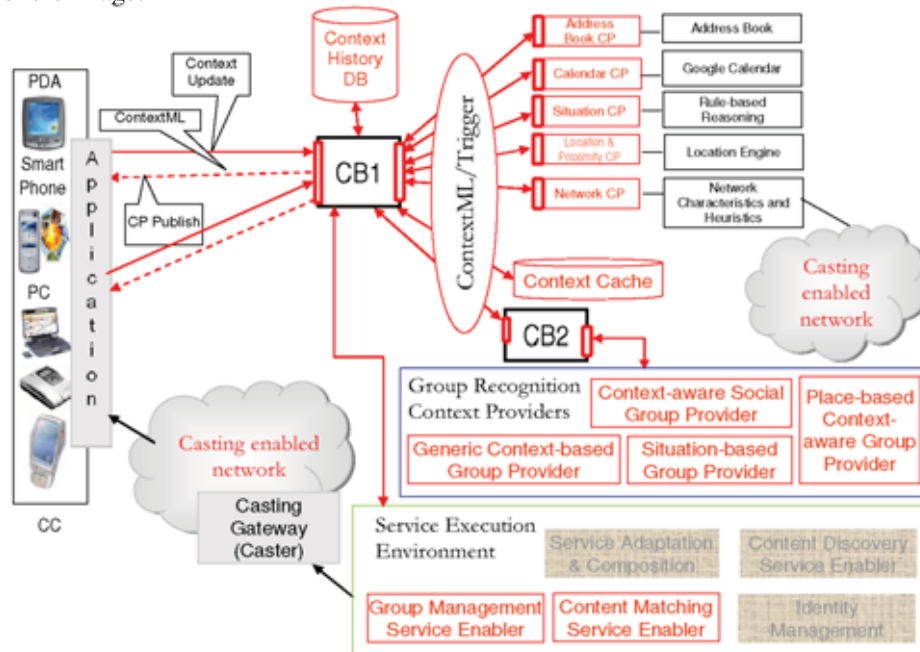


Fig. 3. C-CAST Context-aware platform

The context aware platform proposed by C-CAST is based on a distributed context broker concept and is an enhancement of TiLab's context platform presented in [24]. As illustrated in Fig. 3, the architecture currently encompasses various Context Providers: (1) Address Book CP offering access to address data and relations stored in profiles, (2) Calendar CP providing entries of Google calendars, (3) Location CP providing physical, logical location and proximity to other contextualised entities, (4) Situation CP inferring the users' situation (e.g. formal meeting) based on primitive context provided by other CP, (5) Network CP providing QoS measurements and other information related to the access network. A Context History database allows for storing obsolete context information. This can be used for autonomously learning from previous contexts and actions as well as for estimating user goals. Since C-CAST focuses on efficient group-aware content provisioning, the framework comprises several CP for dynamic ad-hoc group recognition. An event based triggering mechanism adopting the subscribe/publish paradigm is planned. In this mode the CC subscribes to context changes of interest and is asynchronously notified by the CP or by the mediating CB in case of event occurrence.

## 7 Conclusion

Context-aware applications, context-aware artefacts or context aware systems are aware of their environment and circumstances and can respond intelligently. This is the vision that those working in mobile, ubiquitous and pervasive computing are working towards. A future internet capable of embracing this concept and delivering context aware services to users and artefacts elevates this to a pervasive sensing and acting knowledge network. This would be a network able to make decisions, actuate environmental objects and assist users. There are some immense and fundamental challenges for the supporting infrastructure for scenarios of this type:

- This is a major change in the computing paradigm. These are pervasive adaptive systems that respond to things around them with no centralised authority.
- Services provided are more dynamic requiring discovery, ad-hoc composition and orchestration.
- It is a more knowledge based infrastructure that attempts to recognise situations and reason about the environment.

Setting aside the artificial intelligence aspects and concentrating on the network issues then our work in this area reinforces other researchers' experiences that:

- There are many context dissemination and discovery models and architectures. A standard architecture or middleware with well known interfaces is required. We have adopted a broker architecture style augmented with an event based publish subscribe mechanism. But most of this is middleware and perhaps not the concern of the network. Nonetheless discovery, lookup and event distributors do fall in the network domain.
- Representing, learning and reasoning about situations and circumstances in the real world is very hard. Many types of inference engines are required how do we plug them all together. Is this an integral part of the future Internet vision?
- A context aware system by definition is able to sense, adapt and respond. This may require designing in support for communication sense-decide-act loops into a new network. Furthermore if parts of the future network are context aware subsystems themselves designed with intelligent behaviour then there will be conflict. This will arise when components or subsystems make competing and conflicting decisions.
- Context leads to privacy problems. One of the most contentious issues is that gleaning of context is tantamount to monitoring people's daily activities and situations. There must be simple mechanisms for users to withdraw from such services and applications.

Context-awareness is just one feature and for complete analysis we need to understand its position in relation to the holistic concept of a network, what it does and what users expect from it. The Internet network concept was well thought out. At the network layer IP packets are routed through to the end terminal with best effort delivery; packets can arrive in the wrong order, duplicated, missing or in error. It is the transport protocol that runs only in the end user terminal that provides end to end reliability. So there is minimalist expectation of the network. Therefore IP can usually work over any type of network. It is a huge global network and works well so why

would we wish to change it? The user view of technology and expectation would be a good start. Evidence of social networking sites and blogs support the view that when users are given the capability creative and innovative user generated applications and content abound. The simple abstraction of a Web Page built over a network has been very good at this because users have a good understanding of the concept of a page and the technology is fairly intuitive. In contrast the number of users building applications and content for computers and mobile devices is much lower. There are so many operating systems (Vista/Windows {CE, XP}/Mac/Sybian/Linux), interfaces and releases that it is more difficult for the user. Most of the applications and services are owned by software companies or network operators. For the mass population it is still a fight to work with this technology. So the Clouds or Internet of Services is an interesting proposition. In the cloud computing paradigm people define what they want to run or store in the cloud and the cloud architecture does the rest. This is an ideal match for mobile devices with relatively small storage and computing power. So effectively a large part of computer operating system functionality is performed in the network. If the network can compose, orchestrate and deliver services to the user's device then the analogy is complete. The worry of file management, storage, versions and operating systems is left to the network operators. But this vision stands in complete contrast to the current Internet where all the functionality is in the end terminal by design. There is a way of composing services but this is using the Web Page abstraction and Web Services. Furthermore if when designing the Cloud vision effort was channelled into making it easier for composing user generated services then the move to this type of network is even more compelling. Open interfaces and service enablers with which to build applications available to the user community is the only way to empower user creativity and innovation. Contrast this with the situation of cellular networks today. The next step on from an Internet of Services is an Internet of Context-aware Services that incorporates an Internet of Aware-Things. The network would then contain all the entities that we have discussed in the first sections of this paper. Cloud computing or "The Network is the Computer" viewpoint is essential since large amounts of computing power are required to support machine learning, data mining, reasoning and situation recognition. It is a very exciting and challenging idea but what is the motivation? From a user point of view, if it works, then it places technology into the background and makes it more useful. Perhaps of more significance is its potential to promote sustainable living. The planet is under pressure, resources are finite, there is a drive to joined-up government, joined-up local services, transport and general infrastructure in a concerted effort to manage resources in a more sustainable manner. An intelligent network offering context-aware services can assist in this endeavour. In conclusion therefore having made this argument for the future internet vision and the role of context-awareness; is this not Web X.0 in disguise?

## **Acknowledgment**

This work is supported by the European ICT project "C-CAST" - Context Casting - (Contract-No. ICT-216462) [18].

## References

1. Ryan, N., Pascoe, J. and Morse, D. (1997) 'Enhanced reality fieldwork: the context-aware archaeological assistant', Proc. of the 25th Anniv. Computer Applications in Archaeology.
2. Dey, A.K. (1998) 'Context-aware computing: the CyberDesk project', Proceedings of the AAAI, Spring Symposium on Intelligent Environments, Menlo Park, CA, pp.51-54.
3. Hull, R., Neaves, P. and Bedford-Roberts, J. (1997) 'Towards situated computing', Proceedings of the First International Symposium on Wearable Computers (ISWC '97).
4. Brown, P.J. (1996) 'The stick-e document: a framework for creating context-aware applications', Proceedings of the Electronic Publishing, Palo Alto, pp.259-272.
5. Dey, A.K. and Abowd, G.D. (2000) 'Towards a better understanding of context and context-awareness', Proceedings of the Workshop on the What, Who, Where, When and How of Context-Awareness, ACM Press, New York.
6. Prekop, P. and Burnett, M. (2003) 'Activities, context and ubiquitous computing', Special Issue on Ubiquitous Computing Computer Communications, Vol. 26, No. 11, pp.1168-1176.
7. Gustavsen, R.M. (2002) 'Condor - an application framework for mobility-based context-aware applications', Proc. of the Workshop on Concepts and Models for Ubiquitous Computing.
8. Hofer, T. et al. (2002) 'Context-awareness on mobile devices – the hydrogen approach', Proceedings of the 36th Annual Hawaii International Conference on System Sciences.
9. Zimmermann, A. et al. (2005) 'Personalization and Context Management.' User Modeling and User-Adapted Interaction 15, 3-4 (Aug. 2005), pp. 275-302.
10. Loke, S (2007) 'Context-aware Pervasive Systems', Auerbach Publications, New York.
11. Hendricksen, K. et al. (2003) 'Generating context management infrastructure from high-level context models', Proceedings of the 4th International MDM, pp.1-6.
12. Barwise, J. and Perry, J. (1983) 'Situations and Attitudes', Cambridge, MA: MIT-Bradford.
13. Giunchiglia, F., (1992). 'Contextual Reasoning', Trento, Italy.
14. Schilit, B. and Theimer, M. (1994) 'Disseminating active map information to mobile hosts', IEEE Network, Vol. 8, No. 5, pp.22-32.
15. Schilit, B. N. (1995) 'A System Architecture for Context-Aware Mobile Computing', PhD Thesis, Columbia University.
16. Ryan, N., Pascoe, J. and Morse, D. (1997) 'Enhanced reality fieldwork: the context-aware archaeological assistant', Proc. of the 25th Anniv. Computer Applications in Archaeology.
17. Oppermann, R., (2005) 'From User-Adaptive to Context-Adaptive Information Systems', i-com Zeitschrift für Interaktive und Kooperative Medien, 4(3): pp. 4-14.
18. Context Casting (C-CAST), a Specific Targeted Research Project in European Union's ICT 7th Framework Programme, <http://www.ist-ccast.eu>
19. MobiLife, an Integrated Project in European Union's IST 6th Framework Programme, <http://www.ist-mobilife.org>
20. Service Platform for Innovative Communication Environment (SPICE), An Integrated Project in European Union's IST 6th Framework Programme, <http://www.ist-spice.org>
21. Open Platform for User-centric service Creation and Execution (OPUCE), An Integrated Project in European Union's IST 6th Framework Programme, <http://www.opuce.tid.es>
22. Lamorte L. et al. (2007) 'A platform for enabling context aware telecommunication services', Third Workshop on Context Awareness for Proactive Systems.
23. Claudio Venezia, Carlo Alberto Licciardi (2007), 'Improve ubiquitous Web applications with Context Awareness, 11th ICIN 2007.
24. Moltchanov, B. et al. (2008) 'Context-Aware Content Sharing and Casting', 12th ICIN 2008.
25. Chen, H., Finin, T. & Joshi, A. (2003) 'An Intelligent Broker for Context-Aware Systems', Adjunct Proceedings of Ubicomp 2003.
26. M. Roman et al. (2002) 'GAIA: A Middleware Infrastructure to Enable Active Spaces', IEEE Pervasive Computing, vol. 1, no. 4, 2002, pp. 74-83.

## Agreeing Upon SOA Terminology – Lessons Learned\*

Vanessa Stricker<sup>1</sup>, André Heuer<sup>1</sup>, Johannes Maria Zaha<sup>1</sup>,  
Klaus Pohl<sup>1</sup>, Stefano de Panfilis<sup>2</sup>

<sup>1</sup> University of Duisburg-Essen, 45117 Essen, Germany,  
{stricker, heuer, zaha, pohl}@sse.uni-due.de

<sup>2</sup> Engineering Ingegneria Informatica S.p.A., 00185 Roma, Italy  
depa@eng.it

**Abstract.** Building service-based systems with the Service Oriented Architecture (SOA) requires knowledge and experience from diverse domains, including user interaction, service-oriented computing, as well as service platforms and infrastructures. These domains are addressed by different research communities. Therefore, joint research activities between these communities are key to provide novel service technologies for the Future Internet. As each community uses its own language, this poses significant communication challenges. To foster a common understanding of researchers, this paper reports on the process, the results and the lessons learned in devising an agreed terminology within the context of the NEXOF initiative. This terminology is freely accessible on the Web.

**Keywords:** Services, Service-based Systems, Service Oriented Architecture, Glossary, Evolution Process, Future Internet

### 1. Introduction

Due to steady innovations in modern information and communication technologies, the ICT environment at home, at work as well as for mobile usage has changed and advanced in various ways, leading to a continuously increasing number and diversity of software services. Such software services, together with innovative service engineering methods and service technologies will thus become central to the development of the Future Internet and will shape its evolution. Devising these technologies requires knowledge and experience from diverse domains. This paper reports on the process, the results and the lessons learned in devising an agreed terminology to foster a common understanding of researchers of those various domains.

---

\* Research leading to these results has received funding from the EC's Seventh Framework Programme FP7/2007-2013 under grant agreement 216446 (NEXOF-RA).

The agreed terminology exists in web-accessible form as a glossary, consisting of a list of terms and their definitions. It is part of the NESSI<sup>†</sup> strategic project NEXOF-RA that aims at defining a SOA reference architecture (see Section 2). The glossary definition process as well as the lessons learned described in this paper reflect the work and experience of the NEXOF-RA project.

The remainder of the paper is structured as follows: In Section 2 the NEXOF-RA project is briefly introduced. Section 3 presents the goals and purpose of the NEXOF-RA glossary. In Section 4 related approaches for the definition of glossaries are discussed. Section 5 gives an overview of the NEXOF-RA glossary definition process as well as its outcomes. In Section 6 the lessons learned during the glossary definition process are presented. Section 6 summarizes the paper and highlights future research activities.

## 2. The NEXOF-RA Project

In Europe, initiatives that allow and support joint research activities regarding service-based systems include, besides others, the “Internet of Services” initiative of the European Commission, the “Future Internet Assembly” as well as the European Technology Platform NESSI (Networked European Software & Services Initiative).

NESSI aims at promoting the transformation of the EU economy into a service economy. Central to NESSI’s research efforts is NEXOF, the NESSI Open Service Framework. NESSI defines the Open Framework as “[...] an integrated, consistent and coherent set of technologies and associated methods and tools intended to:

1. provide European Industry and the Public Sector with efficient services and software infrastructures to improve flexibility, interoperability and quality;
2. master complex software systems and their provision as service oriented utilities; establish the technological basis, the strategies and deployment policies to speed up the dynamics of the services eco-system;
3. develop novel technologies, strategies and deployment policies that foster openness, through the increased adoption of open standards and open source software as well as the provision of open services;
4. fostering safety, security and the well-being of citizens by means of new societal applications, enhanced efficiency of industry and administrations, and competitive jobs.”

A first step to achieve those goals is NEXOF-RA, an integrated project (IP) whose overall ambition is to deliver a reference architecture for the NESSI Open Service Framework. This SOA reference architecture ranges from the infrastructure up to the interfaces with the end users, leveraging research in the area of service-based systems to consolidate and trigger innovation in service-oriented economies. The consortium is composed of 16 partners from eight different countries (Spain, Italy, Germany, UK, Ireland, Israel, The Netherlands and France).

---

<sup>†</sup> [www.nessi-europe.com](http://www.nessi-europe.com)



### 3. Goals and Purpose of the NEXOF-RA Glossary

Since the NEXOF-RA project unites several different communities and domains, the agreement on terms and reaching a common understanding is crucial to the success of the initiative. Although the aim is to find consensus and agree on a well-defined definition of terms several definitions might also be the result as the concise definition of a term can depend on the context in which it is used. Nevertheless, it is crucial to gain a common understanding about all those definitions and the specific contexts in which they are used within the project to avoid miscommunication.

To understand the process described in this paper, first the objectives that have been set for this glossary will be described. These objectives are mainly identified according to the different groups of glossary users. As the project includes different groups of contributors the definition of the glossary aims at addressing all of those groups for different purposes. The objectives can be summarized as follows:

- **Baseline for NEXOF-RA internal activities:** To define a common terminology the project includes the creation of a conceptual model of the reference architecture as well as the definition of a glossary. It contains and defines the concepts of the model as well as the central terms of the specification definition. The glossary should be used and referenced in all project-internal activities and deliverables.
- **Integration of results of NESSI strategic projects (NSPs):** Within the NESSI initiative several different projects have been launched which focus on different aspects of service-based systems (e.g. SOA4All, RESERVOIR, SLA@SOI). The results of their research activities should be included and considered in NEXOF-RA in order to prevent that effort is spent on topics that already have been discussed. The integration of these results also makes a common terminology desirable. The NSPs are asked to provide the aspects of their results that are relevant for the definition of a reference architecture based on the agreed terminology. Accordingly, the projects are also involved and can actively participate in the definition and agreement of the glossary definition.
- **Reference for the NEXOF community:** Besides the integration of results from existing projects NEXOF-RA also performs an “Open Construction Cycle”<sup>‡</sup> in which external contributors (e.g. the S-Cube Network of Excellence) are asked to submit their solutions to specific topics of a SOA. The work done within this process should be based on the agreed terminology covered by the glossary. Since the external contributors use the definitions and it is expected that new insight are gained by their work, they can also influence the evolution of the glossary by suggesting changes. However, they cannot actively participate in the agreement and decisions.

Accordingly, the NEXOF glossary provides a common source of reference for all NEXOF partners as well as for partners involved in the community process of defining the NEXOF reference architecture.

---

<sup>‡</sup> [http://www.nexof-ra.eu/?q=open\\_construction\\_process](http://www.nexof-ra.eu/?q=open_construction_process)

#### 4. Related Work

The literature on building glossaries is mainly focussed on automatic or semi-automatic derivation of glossary terms and their definitions. In the European Network of Excellence INTEROP, Velardi et al. describe their method for building a glossary [7]. Therein, the glossary is part of a Knowledge Map Acquisition Chain, beside a lexicon, a taxonomy and an ontology. To build the glossary, they use an automated terminology extraction algorithm that analyzes a large set of collected documents within a community or project (e.g. State-of-the-Art surveys and papers, deliverables, workshop proceedings, etc.). Thus, this approach bases on existing documents. In the case of NEXOF-RA, such an automated approach was not feasible. The glossary definition process could not rely on documents produced within the project, because the glossary definition started on day one of the project. Of course, quite a few documents exist in the different research communities. However, these have not been consolidated by taking into account the scope and the goals of NEXOF-RA.

One promising paradigm for creating collaborative knowledge is Open Authoring. This is realized by wikis [1], which allow users to create and edit any page in a web site. It is exciting in that it encourages democratic use of the web and promotes content composition by all users [1]. In the context of NEXOF-RA however, normally only one person per partner was assigned to work on the glossary. Therefore, it appeared that the community that “creates” the glossary using the wiki would not be large enough. Further, the people involved in the creation of the project had conflicting domain knowledge, which they proposed as definitions. Accordingly, the free authoring approach could have resulted in rewriting and especially overwriting terms almost arbitrarily. This mutual overwriting of definitions without tracing what knowledge about a term already has been proposed and deleted could have resulted in cycles of definitions. Although, the wiki approach always allows to make a roll-back to an earlier version this is not a step a of a constructive knowledge consolidation process. Doing a roll-back the new information would get lost instead of being consolidated with existing information. With this conflicting knowledge no collaborative knowledge creation could be performed as intended.

#### 5. The NEXOF-RA Glossary Definition Process and its Outcomes

The key idea behind the NEXOF-RA glossary definition process is to exploit elicitation techniques that are used in Requirements Engineering in order to drive the consolidation. Specifically, written surveys ([3], [4]) are performed which provide input for a consolidation workshop following the format described by Pohl [2] and Leffingwell and Widrig [6].

The NEXOF-RA glossary definition process is structured into three main phases:

- **Phase 1 - Survey of initial terms:** collect initial list of definitions of terms from the participants
- **Phase 2 - Construction and consolidation of term definitions:** build a first consolidated version and participants comment on this version as baseline for a workshop in which an agreement and consolidated version is defined

- **Phase 3 - Change and evolution:** evolve the definitions by collecting requests for changes and agreeing upon changes. This phase of the process is iterative in order to allow a step-wise evolution of the glossary.

### 5.1. Phase 1: Survey of initial terms

As mentioned above, the used approach for the survey is geared to elicitation techniques of the written survey (cf. [3] and [4]). It offers the possibility to elicit an initial set of requirements and term definitions respectively from a large number of stakeholders [5]. It also identifies new sources for term definitions, e.g. other glossaries.

Thus, phase 1 started with sending the template for defining terms to all participants. They were encouraged to fill in the template and define terms that they would like to be considered in the glossary. Each participant was asked to define terms, which are the most relevant ones for the glossary in the project context (e.g. Service, Service Level Agreement, Orchestration, Composition, etc.), and also to add new terms if needed. The mentioned terms and all additional terms had to be described in four sections:

- **Description:** Describe the term short and precise, but as detailed as necessary
- **Reference:** Reference to literature or to domain where term is used
- **Rationale:** Rationale for the definition
- **Synonyms:** List of words with the same meaning

A template was used to support the survey and the following construction and consolidation phase in different ways. First of all, it assured that all participants would fill in the information needed to construct the glossary as well as that the proposed definitions could easily be integrated. The reference section was helpful since it made clear that existing maybe commonly accepted standards or sources were referenced. Besides this, own rationales have been asked in order to strengthen the different positions during the discussions and to communicate the different considered aspects before the discussions started.

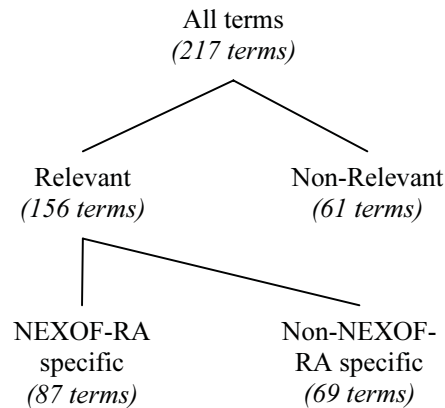
**Outcomes:** Terms considered as relevant for the glossary cover all terms that are considered as relevant in the whole NEXOF-RA project. This is the main reason for the large number of terms defined after phase 1 (over 200 terms, including multiple definitions for one term).

### 5.2. Phase 2: Construction and consolidation of term definitions

Based on the terms received from the participants, a preliminary version of the glossary was built and made accessible via the NEXOF-RA web based common working environment. While building this preliminary version, a first consolidation was performed with regard to multiple definitions and synonyms for a term by the glossary moderator. Thus, a web page for a given term was structured as follows: the first definition on a page was the current definition that is “accepted”. Alternative definitions (if existing) were listed below. Finally, there was a section for comments.

Since every contribution of a participant was included at that stage, there have been terms that were only NEXOF-RA related from the point of view of one domain.

Therefore, participants were encouraged to indicate whether a given term is NEXOF-RA relevant or not. The first separation of relevant and non-relevant terms revealed



**Figure 1** Structuring principles of the NEXOF-RA glossary

that only few of the existing terms were not relevant. In order to build a basis for discussion for the consolidation, the participants were furthermore encouraged to add comments on the definitions and in case they do not agree with a “proposed” definition to propose an alternative wording. Thereby, it was crucial that added comments could be referenced by other comments. The commenting showed that some participants have concentrated on specific terms and did not comment every definition of terms. Thereof, the decision was reached to further separate the relevant terms in two

categories: NEXOF-RA specific and non-NEXOF-RA specific terms. Specific terms include all terms, which are closely associated to the NEXOF-RA context. All other terms are non-specific. Figure 1 illustrates the decomposition.

To consolidate the proposed definitions, a consolidation workshop, according to the workshop format described by Pohl [2] and Leffingwell and Widrig [6], was organized open to all participants. During the workshop, groups of participants from similar domains discussed the proposed definitions. They had also to identify overlaps in the definitions which needed to be harmonized. Thereafter, the results were discussed and approved in the plenary. First, the specific terms were discussed and after reaching an agreement, the non-specific terms were in the focus. However, some of these terms needed some substantial rewriting. Thus, this task was postponed and executed by the designated identified expert. The rewritten definitions were reviewed by other participants.

**Outcomes:** In the project glossary 61 of the existing terms were considered of not being relevant. Out of the 156 terms that have been considered as being relevant for the NEXOF-RA project, 87 were marked as NEXOF-RA specific and 69 as non-NEXOF-RA specific. The output from this phase was distributed to all participants. Before the first version has been frozen, it was ensured that participants of the NSPs had the possibility to provide feedback to the glossary and add terms that had not been considered till this point. Afterwards, this first version of the glossary was published on the NEXOF-RA portal in order to allow external contributors to access the terminology, to use it in their contributions as well as to provide further feedback that can be incorporated into the glossary during the revision phase.

### 5.3. Phase 3: Change and evolution

As described in the previous section, the current version of the glossary has been accepted by all participants and representatives of the NSPs and it has been made publicly available. From now on, multiple changes are expected to occur regarding the existing definitions for terms included in the glossary due to newly gained insights concerning the research activities. To allow for the baseline glossary to evolve over time according to the needs identified during the project, a defined and structured change and evolution phase is implemented.

This process will be carried out multiple times in the remaining duration of the NEXOF-RA project.

A single iteration of this change and evolution phase is described as follows:

1. **Submission of Requests for Change (RfCs):** All participants and external contributors are able to send Requests for Change (RfC) to the glossary moderator. These RfCs should address defined terms in the current glossary (the current status will always be available in an *internal* wiki) as well as new terms that have not been considered. Each request should contain:
  - a revised (or new) definition of the term
  - the rationale for the request
  - a literature reference (if available)
  - the context in which the request emerged (e.g. from which project results).
2. **Agreement on changes:** Since each RfC is potentially changing a consolidated and agreed definition in the glossary, each RfC needs to be accepted by all participants as well as the representatives of the NSPs. In order to reach such an agreement, all RfCs will be made available in the wiki for comments at least one week prior to the discussion about the acceptance of these RfCs in the project plenary. During this discussion, the requester or a representative of the requester needs to justify the request.
3. **Acceptance by all participants:** After the construction of a revised and consolidated version of the glossary, it has to be accepted by all participants, possibly involving an adaptation with respect to the external contributors.

**Outcomes:** At the time of writing around 150 RfCs have been submitted and will be discussed in the next plenary meeting of the project. It is expected that several iterations of this change and evolution phase will be performed till the end of the project. This process makes the glossary a living document.

## 6. Lessons Learned

The following experiences and lessons learned have been gained during the glossary definition and consolidation process:

### 6.1. Traceability supported the consolidation of terms:

We experienced that tracing the sources of definitions as well as provided rationales has been supportive in order to assist the consolidation and agreement activities on

proposed definitions. When commenting on definitions, the rationales could be used to understand the intention of the authors as well as the aspects that should be stressed by a specific definition. In the NEXOF-RA context this knowledge helped in uncovering conflicting misunderstandings. This also saved time during the consensus finding in the glossary workshop since we could assume that the intention of a definition was known by the other participants and did not need to be clarified.

#### **6.2. Structuring the contributions was helpful for the consolidation:**

We observed that structuring the glossary improved the readability, usability and navigation of the definitions. It also improved the creation and consolidation phase. During the definition of the glossary in the NEXOF-RA context, we noticed that the level and quality of submitted terms of the survey and of comments were quite different. Furthermore, some participants tended to spend a lot of their time on commenting definitions of terms that are not part of the terms most crucial to the project. In order to avoid the mis-prioritization of terms and work effort, a first structuring of the glossary into two categories has been performed. Accordingly, the glossary is divided into two categories as described: NEXOF-RA specific and non-NEXOF-RA specific terms. NEXOF-RA specific terms include all terms which are associated to NEXOF-RA core activities. All other terms are non-NEXOF-RA-specific. This also allowed focusing on a pre-selection of most relevant and crucial terms. We think that this structuring has also been shown to be reasonable for grouping terms according to different topics. These topics could be assigned and discussed in smaller groups according to identified competences. Furthermore, we experienced that structuring the comments also supported the agreement process since it could be traced which comments are suggested by which participant. We observed that comments referred to other comments and they were used in order to ask for justification and resolution of conflicts.

#### **6.3. Defining rules guided the agreement:**

We observed that some comments that have been submitted before as well as in the glossary workshop have been repeated several times. Thus, it has been agreed that some basic rules were defined which should be used and checked during each agreement on term definitions. Such a rule was for example: “Combined terms should be avoided. If necessary, it should be a specialization of the single term”. We observed that these rules were cited several times during the workshop and thus avoided redundant discussion on a topic.

#### **6.4. Guidance has increased the homogenous quality of contributions:**

In order to ensure that the results have a certain quality and that the contributions are on a similar level of detail and thereby comparable, the survey and consolidation in the NEXOF-RA context have been guided. Accordingly, a template that explicitly defines and describes all needed information has been distributed to improve the contribution of term definitions. Furthermore, we observed that the provision of terms

that should be in the core of the submitted definitions has been a good starting point for the contributors and have led to further detailed knowledge provisioning. We expect our experiences about guidance to also be valid for the change and evolution phase. Regarding the context in which the glossary is defined, it had to be considered carefully if the definition and construction process of the glossary should be completely decentralized. In this case no authority would guide the definition process. Based on our experience the assignment of a moderator was adequate in the NEXOF-RA context in order to increase the glossary quality.

#### **6.5. Highlighting the objectives of the glossary assisted in focusing:**

A tendency that has been observed during the glossary definition process is that the participants wanted to define the “whole world” in the glossary in order to capture everything that eventual might be relevant for the project. This was reflected in the received definitions during the survey phase. In order to stay in control of the effort needed in the consolidation phase, in the glossary workshop we clearly stated and reminded of the objectives of the glossary. Accordingly, the first step in the consolidation workshop was to identify those terms that are out of scope for the glossary since they are not related to the core activities performed within the project. Thus, the boundaries have been defined and considered during the discussion and agreement of possible terms as well as during the submission of RfCs.

#### **6.6. Consolidation workshop fostered the common agreement:**

We observed in the glossary workshop that the main obstacle to the agreement on terms was the tendency to stick with a formerly adopted definition. Accordingly, in the NEXOF-RA context it was difficult for involved partners to detach from existing research contexts in order to promote the gained experience in new fields. The discussions that emerged at the workshop showed these obstacles. However, these face-to-face discussions resulted in an agreement on definitions supported by all partners.

### **7. Conclusion and Future Work**

The paper described the process of defining an agreed SOA terminology within the NEXOF-RA project. It has been discussed why a more controlled and triggered process of term collection, consolidation and agreement has been performed. The result of that process is a glossary that currently contains about 160 SOA related terms from several different domains and research communities. The glossary is made available to the whole research community on the NEXOF-RA web site.

To support its evolution the glossary provides the possibility to cope with changing terminology in a structured way. The future work will mainly focus on optimizing the way in which the change and evolution phase is performed and new or changed terms are consolidated within the broader community. This includes the dissemination of



the glossary outside the already existing NEXOF-RA and NESSI community into related research communities. The NEXOF-RA glossary is an evolving and living document which aims at being valid beyond the boundaries and the duration of the project.

## 8. References

- [1] Leuf, B.; Cunningham, W.: *The Wiki Way – Quick Collaboration on the Web*. Addison-Wesley, 2001.
- [2] Pohl, K.: *Requirements Engineering – Grundlagen, Prinzipien, Techniken*. dpunkt Verlag, 2008.
- [3] Bray, I. K.: *An Introduction to Requirements Engineering*. Addison-Wesley, Reading, 2002.
- [4] Oppenheim, A. N.: *Questionnaire Design, Interviewing and Attitude Measurement*. Pinter, London, 1999.
- [5] Zowghi, D.; Coulin, C.: *Requirements Elicitation – A Survey of Techniques, Approaches and Tools*. In: Aurum, C.; Wohlin, C. (eds.): *Engineering and Managing Software Requirements*. Springer, Berlin, Heidelberg, New York, pp. 19-46, 2005.
- [6] Leffingwell, D.; Widrig, D.: *Managing Software Requirements – A Unified Approach*. Addison-Wesley, Reading, 2000.
- [7] Velardi, P.; Navigli, R.; Petit, M.: *Semantic Indexing of a Competence Map to support Scientific Collaboration in a Research Community*. In: *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, Hyderabad, India, 2007.
- [8] W3C: *Web Services Glossary*. <http://www.w3.org/TR/ws-gloss/>. Accessed on 2008-06-16, 2004.
- [9] MacKenzie, C. M.; Laskey, K.; McCabe, F.; Brown, P. F.; Metz, R.: *OASIS Reference Model for Service Oriented Architecture 1.0. Committee Specification 1*, <http://www.oasis-open.org/committees/download.php/19679/soa-rm-cs.pdf>, 2006.

## Subject Index

a call for specifications	136	layered video coding	273
actuate	335	locality-awareness	24
adaptation	263, 283	management plane	112
architecture	79	many core	238
attitude	12	markets	1
autonomic network architectures	136	MDC	273
autonomicity	112	mobile grids	238
awareness	293	monitoring	263
business models	1, 325	multipath TCP	79
collaboration strategies	123	network architecture	91, 160
content	1	network intelligence	293
context	250, 335	network virtualisation	91
context-awareness	335	network(s)	1, 238
customization	1	NQoS	293
design	263	operating systems	238
distributed computing	238	optical communication	160
Economic Traffic Management	24	orchestration plane	112
engineering	263	P2P networks	1
evolution	12	P2P Streaming	273
evolution process	345	P2P VoD	24
federated testbeds	67	Panlab	67
federation	67, 250	Plastic Optical Fiber (POF)	160
future	12	Presence	250
future content network		pre-Standardization through an	
architectures	303	Industry Specification Group	
future Internet	1, 57, 79, 91, 102,	(ISG)	136
123, 136, 250, 325, 335, 345		pricing	1
future Internet architecture	313	privacy	1, 57
future media 3D Internet	303	providers	1
future media Internet	283	publish/subscribe	102
glossary	345	QoE	293
grid	238	QoS	1
Home Area Network (HAN)	160	quality	263
identity management	57	real world awareness	313
immersive environments	303	real world Internet	313, 325
information-centric	102	reason	335
information-centric networking	91	reference point	67
In-House-Network	160	regulations	1
Interconnection	250	requirements	325
inter-domain traffic	24	resource accountability	79
Internet	12	resource control	67
Internet governance	1	resource pooling	79
knowledge plane	112	scalable video coding	283
law	57	scenario-based design	325

scoping	102	situationawareness	123
security	47	SmoothIT architecture	24
self-management	91	SOA	203
self-managing networks	136	socioeconomic(s)	1, 12, 325
semantics	203	standardization	1
sense	335	SVC	273
sensor and actuator networks	313	teagle	67
sensor networks	335	technology	12
service and network management	227	testbed	67
service and self-aware network management	112	Trilogy project	79
service computing	227	trust	1
service creation	217	trust model	47
service delivery platform	217	user behaviour	1
service discovery	47	user identity	1
service enablers plane	112	user-centric	217
service infrastructure	227	users	1
service oriented architecture	238, 263, 345	value chains	1
service web	203	virtualisation	112, 227
service-based applications	263	Web 2.0	203
service-based systems	345	Web 2.0. mashup	217
services	1, 345	web service	217
		wireless sensor networks	47
		WSMO	203

## Author Index

Abramowicz, H.	91	Domingue, J.	ix, 203
Adriaanse, J.	250	Dotaro, E.	136
Aguiar, J.	217	Eardley, P.	79
Ain, M.	102	Eiselt, M.	173
Alvarez, F.	47, 303	Elmroth, E.	227
Azodolmolky, S.	173	Emmerich, W.	227
Baker, N.	335	Eurich, M.	325
Baladrón, C.	217	Evanno, N.	160
Baresi, L.	193	Fajardo, J.-O.	293
Bassi, A.	112	Falcarin, P.	217
Baucke, S.	91	Fensel, D.	203
Bauer, M.	313	Fischer, A.	112
Behrmann, M.	1	Fogliati, V.	1
Bellec, M.	160	Forest, F.	325
Berl, A.	112	Gagnaire, M.	173
Bernat Vercher, J.	313	Galis, A.	ix, 112, 227
Besson, L.	47	Gaudino, R.	160
Boniface, M.	1	Gavras, A.	ix, 67
Breitgand, D.	227	Giacomin, P.	112
Burness, L.	79	Girao, J.	57
Cáceres, J.A.	227	Gluhak, A.	313
Callejo, M.Á.	1	Goix, L.-W.	217
Campolargo, M.	v	González-Cabero, R.	203
Cardenas, D.	160	Grüneberg, K.	283
Carro, B.	217	Guignard, P.	160
Celetto, L.	273, 283	Gunkel, M.	173
Chaparadza, R.	136	Hancock, R.	79
Chapman, C.	227	Haridi, S.	148
Charbonnier, B.	160	Hausheer, D.	ix, 1, 24
Chen, M.	173	Hecht, F.	24
Cheniour, A.	112	Heuer, A.	345
Clayman, S.	227	Hirsch, T.	123
Courcoubetis, C.	1	Hoßfeld, T.	24
Cramer, E.	250	Hrasnica, H.	67
Daras, P.	303	Jäger, D.	160
Da Silva, J.	v	Jari, A.	273
Davies, J.	203	Johansson, M.	313
Davy, A.	136	Johnsson, M.	91
Davy, S.	112	Joosen, W.	35
de Meer, H.	112	Jorba, S.R.	1
de Panfilis, S.	345	Karastoyanova, D.	263
Denazis, S.	67, 112	Kastrinogiannis, T.	136
Desmet, L.	35	Kind, M.	91
Di Nitto, E.	263	Kipp, A.	238

Kleinwächter, W.	1	Philippaerts, P.	35
Kleis, M.	123	Pickavet, M.	173
Knappmeyer, M.	335	Piesiewicz, R.	173
Koumaras, H.	293	Piessens, F.	35
Krco, S.	ix	Pistore, M.	183, 263
Ladid, L.	1	Pizzinat, A.	160
Ladis, E.	47	Pohl, K.	263, 345
Lamy-Bergot, C.	283	Pointurier, Y.	173
Lavoisy, O.	325	Popescu-Zeletin, R.	123
Lefevre, L.	112	Presser, M.	313
Lefol, D.	273	Pujolle, G.	112
Lehrieder, F.	24	Qin, Y.	173
Leligou, H.	47	Quinn, K.	136
Levy, E.	227	Quittek, J.	91
Li, M.-S.	1	Racz, P.	24
Liakopoulos, A.	136	Reinefeld, A.	148
Liberal, F.	293	Rochwerger, B.	227
Llorente, I.M.	227	Rubina, J.M.	12
Lotz, V.	ix	Saradhi, C.V.	173
Loupis, M.	47	Sarma, A.	57
Loyola, J.R.	112	Schenk, M.	250
Lozano, D.	67	Schierl, T.	283
Macedo, D.	112	Schintke, F.	148
Mahlab, U.	173	Schubert, L.	238
Manner, J.	47	Schütt, T.	148
Martin, A.L.	217	Sentinelli, A.	273
Massacci, F.	35	Serrat, J.	112
Metzger, A.	263	Siahaan, I.	35
Meyer, S.	160	Sienel, J.	217
Mischler, D.	67	Silvestri, F.	263
Möllers, I.	160	Simeonidou, D.	173
Moltchanov, B.	335	Solé Pareta, J.	173
Montagut, F.	313	Sorge, C.	57
Montero, R.S.	227	Soursos, S.	24
Naliuka, K.	35	Spirou, S.	1
Nejabati, R.	173	Srassner, J.	112
Niebert, N.	91	Staehle, D.	24
Nikander, P.	1	Stamoulis, G.D.	24
Oechsner, S.	24	Stiller, B.	1, 24
Ohlman, B.	91	Stirbu, V.	313
Palazzi, C.	273	Stricker, V.	345
Paolucci, M.	183	Tarkoma, S.	102
Papadopoylos, K.	47	Theilmann, W.	193
Papaefstathiou, Y.	47	Tilanus, P.	250
Papafili, I.	24	Timmerer, C.	283
Papavassiliou, S.	136	Tomkos, I.	173
Parkin, M.	263	Toth, A.	136
Pau, G.	273	Trakadas, P.	47
Pedrinaci, C.	203	Tran-Gia, P.	24

Traverso, P.	183	Wilson, D.	325
Tselikis, C.	47	Wilson, M.	136
Van den Heuvel, W.-J.	263	Wódczak, M.	136
van Deventer, M.O.	250	Woesner, H.	91
Van Gorp, J.	325	Wolfsthal, Y.	227
Vangelatos, C.	47	Wünstel, K.	1, 91
Vanoverberghe, D.	35	Wusthoff, M.	227
Varvarigos, E.	173	Ye, Y.	173
Vigoureux, M.	136	Zafar, M.	335
Visala, K.	102	Zaha, J.M.	345
Wagner, M.	183	Zahariadis, T.	ix, 47, 273, 283
Wahle, S.	67	Zami, T.	173
Wesner, S.	238	Zseby, T.	123