2/2012

EURESCOM mess@ge

The magazine for telecom insiders

NEWS

Celtic-Plus
Smart Connected World

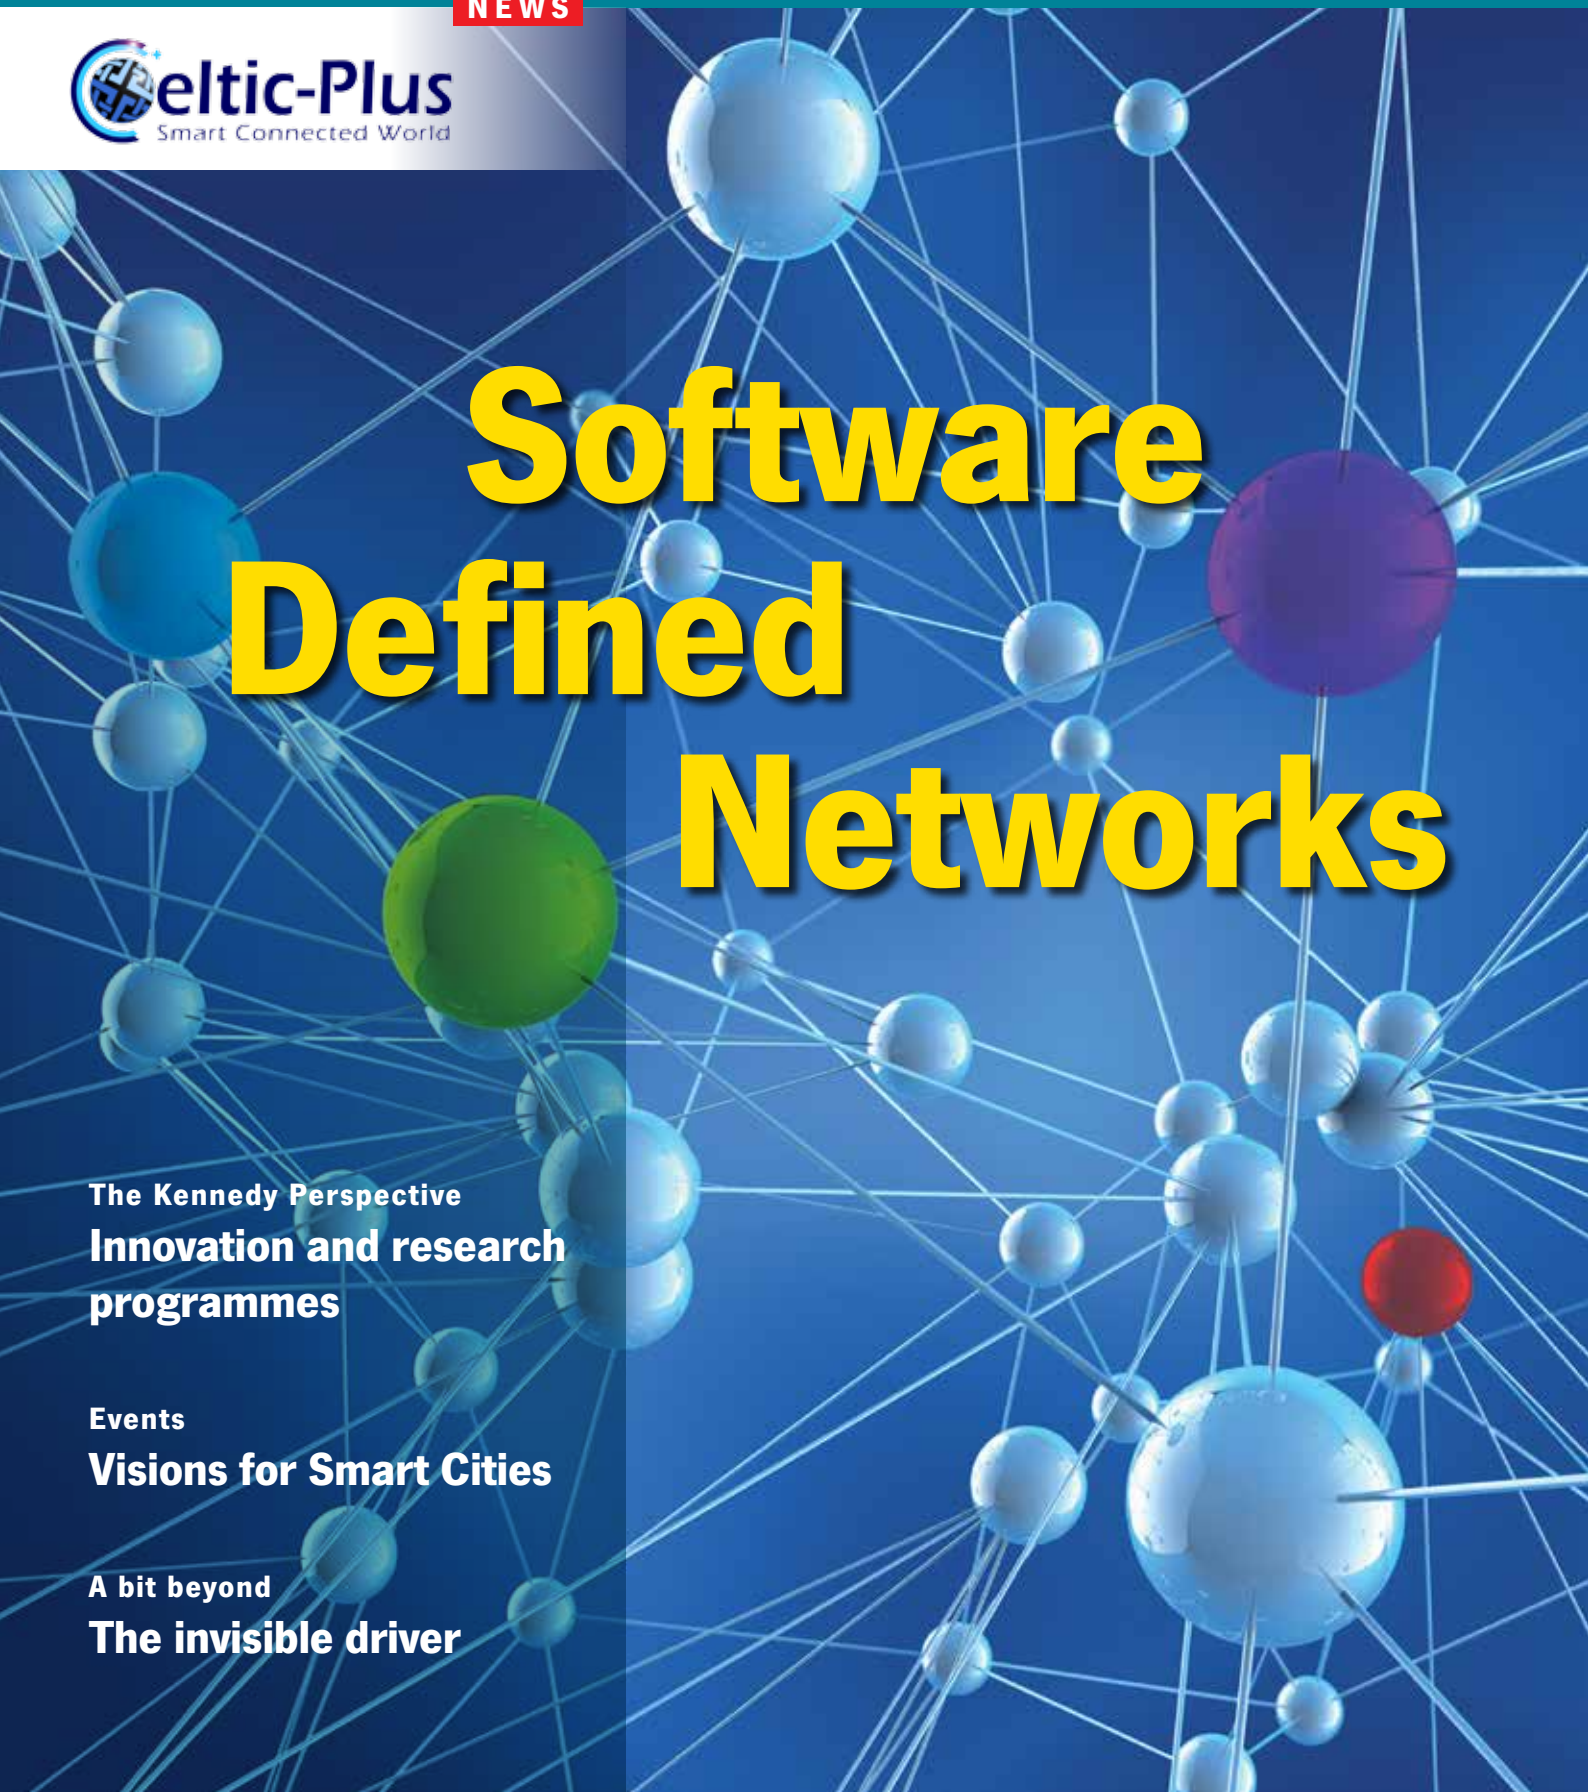

Software Defined Networks

The Kennedy Perspective
Innovation and research programmes

Events
Visions for Smart Cities

A bit beyond
The invisible driver

Net!Works Event 2012: Towards Horizon 2020!

Challenges in Communications Networks and their coverage in Horizon2020

13 November 2012 – Radisson Blu Royal Hotel, Wolvengracht 47, Brussels

Net!Works is organizing a one-day workshop in Brussels on 13 November 2012 to provide a forum for discussion of the Horizon 2020 programme. In this workshop, representatives of the European Commission, from industry and from application areas will present their perspectives on how communications networks and services can contribute to achieving the goals of their respective sector.

In Horizon 2020, communications networks and services are part of the three priorities: Excellent Science, Industrial Leadership, and Societal Challenges. Networks and services are explicitly identified as a specific topic in the priority Industrial Leadership, but are also part of the many programme sections oriented towards applications and societal challenges.

The workshop will:

- point out what role is currently foreseen for ICT and communication networks in Horizon 2020
- address potential funding opportunities
- present speakers from the European Commission, the European Parliament, and from industry, who will give their views on the relevance and importance of networking technologies and services for the European economy and society
- offer workshop participants the opportunity to raise their voice and make their perspective visible by joining the discussion groups

The workshop will collect views of the participants on topics which could be included in the detailed work plan for the communications networks and services area of the programme and identify arguments for their inclusion. The output will be fed into a document to be submitted into the Horizon 2020 process to selected target groups.

Registration and further information:
www.networks-etp.eu

Funding Opportunity for European R&D Projects

Celtic-Plus Autumn Call for Proposals – Deadline: 15 October 2012

Celtic-Plus is a EUREKA Cluster dedicated to realising the vision of a smart connected world through an industry-driven R&D programme. There are two calls per year, in spring and in autumn, with a total funding of up to 100 million euro. The funding is orchestrated via the Celtic-Plus programme and provided by Public Authorities from 47 EUREKA member countries.

Eligible topical areas

Get Connected
- Infrastructure and connectivity aspects
- Fixed/Wireless, optics, energy-efficiency
- Network architecture, autonomic networks

While Connected
- End-to-end services and applications, like
  - Digital home, digital enterprises
  - Digital City (incl. digital school, digital transport)
  - E-Health
  - Security, privacy, identity

Future Internet relations
- Complement Future Internet (FI-PPP) programme by:
  - Making the Internet a high-quality service platform
  - Introducing the 'Celtic-Plus Use-Case Factory':
    - Extending the programme by additional use-cases not covered in the FI-PPP programme
    - Contributing to Future Internet capacity building and test cases/platforms

Green-Internet relations
- Consider environmental issues in ICT
- Encourage better energy efficiency
- Consider Smart Grid, Water management & ICT
- Develop multi-disciplinary approach

User friendly call process

The Celtic-Plus programme gives proposers the opportunity to submit proposals twice in the year – Spring Call and the Autumn Call. **For the Autumn Call the deadline will be 15 October 2012**. Celtic-Plus has an accelerated one-stage call process to ensure the shortest possible time between proposal submission and start of selected projects. Celtic-Plus proposals should be complete and clearly present the technical objectives, timescales, participants, manpower, and expected results. These proposals are evaluated by independent evaluators and the proposals meeting the required standards will be retained and given the Celtic-Plus label. To be eligible for funding, project partners need to be located in EUREKA member countries.

Further information

Please visit the Celtic-Plus website at www.celticplus.eu for call details and the Celtic-Plus Purple Book for details on the R&D priorities of Celtic-Plus. For further information, please contact Heinz Brüggemann, director of the Celtic Office, at brueggemann@celticplus.eu

Celtic-Plus
Smart Connected World

www.celticplus.eu

Dear readers,

Software Defined Networking, short: SDN, is promising to be the next big thing in the evolution of communication networks. The virtualisation of network resources offers significant benefits. While Europe is still researching, testing and discussing, Google has already deployed the world's largest production Software Defined Network based on OpenFlow.

In view of the dynamic progress in the area of SDNs, particularly OpenFlow, the editorial team of Eurescom mess@ge considered the time right for dedicating a cover theme to SDN. In this issue, we will provide an overview on what Software Defined Networking is, feature some European research efforts in this area and provide a glimpse on the R&D challenges and opportunities that lie ahead.

The first article in the cover theme provides an overview on SDN. The ensuing article by Takashi Egawa shares an insider's view on the current status of SDN standardisation. In an exclusive interview, Alex Galis from University College London shares his insights on the future of Software Defined Networks.

Two of our cover theme articles are dedicated to OpenFlow. In the first, Dimitra Simeonidou and Reza Nejabati from the University of Essex describes their pioneering work on OpenFlow for high capacity transport networks. Spyridon Tompros from the OpenLab project describes an architecture for integrating OpenFlow functionality in experimental networks, which is under development in OpenLab, based on previous work from Panlab, a project coordinated by Eurescom.

Our next cover theme article highlights the importance of Software Defined Networking for innovation on the Internet through the flow processing platform of FP7 project CHANGE, which is presented by its coordinator, Adam Kapovits from Eurescom. Finally, Karsten Oberle and Manuel Stein from Alcatel-Lucent Bell Labs Germany present results of the IRMOS project on Virtual Service Networks.

This issue also includes a variety of further articles on different, ICT-related topics, including a report on the Future Internet Assembly in Aalborg as well as a report on the EU-Japan workshop which also took place within the Future Internet Week in Aalborg. See also the "The Kennedy Perspective" on innovation and research and the "A bit beyond" article on autonomous cars. I hope you enjoy reading our magazine.

My editorial colleagues and myself would appreciate your comments on the current issue as well as suggestions for future issues.

Milon Gupta
Editor-in-chief

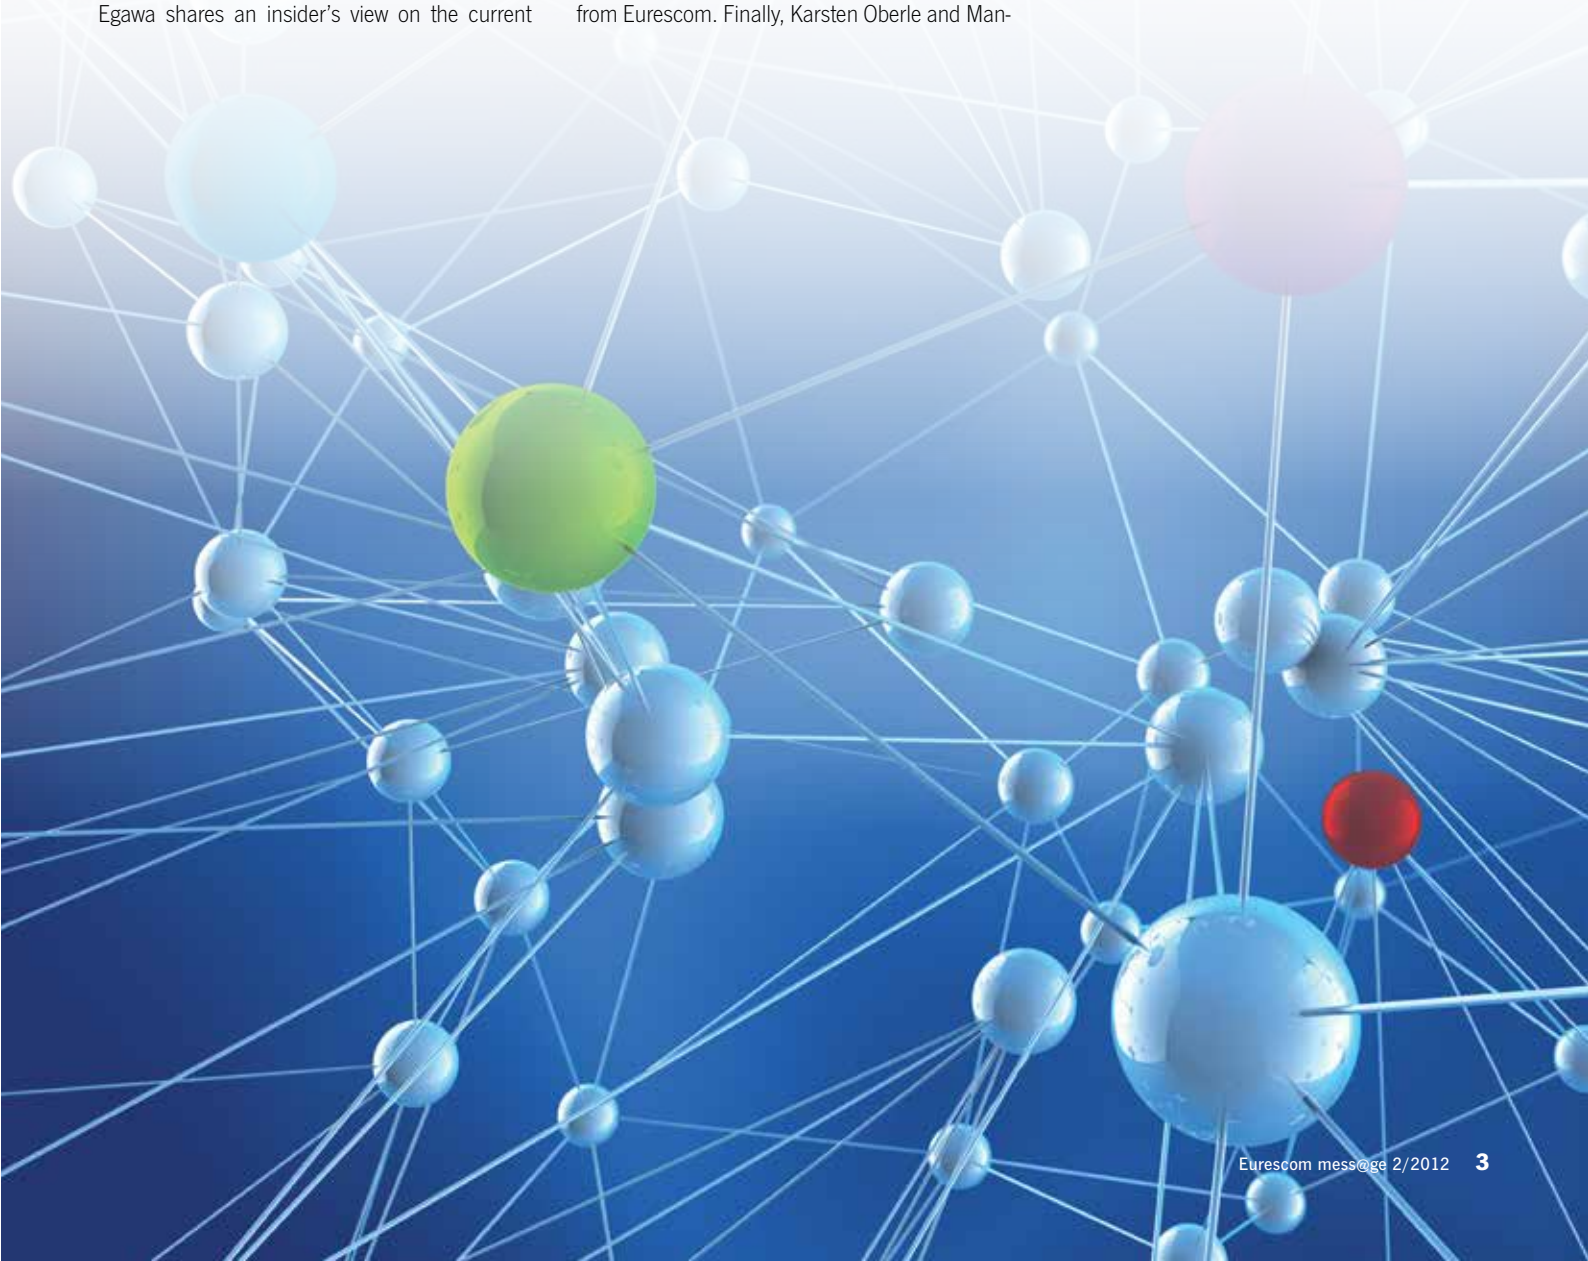

Events calendar

3 July 2012
Celtic-Plus Proposers' Day
Held in conjunction with the Future Network and Mobile Summit
Berlin, Germany
http://www.celticplus.eu

4 – 6 July 2012
Future Network and MobileSummit
Berlin, Germany
http://www.futurenetworksummit.eu/2012

31 August – 5 September 2012
IFA
Berlin, Germany
http://b2b.ifa-berlin.com/en

6 – 11 September 2012
IBC 2012
Amsterdam, The Netherlands
http://www.ibc.orgen

16 – 18 October 2012
NEM Summit 2012
Istanbul, Turkey
http://www.nem-initiative.org

17 – 19 October 2012
eChallenges e-2012 Conference and Exhibition
Lisbon, Portugal
http://www.echallenges.org/e2012

6 – 7 November 2012
FIRE engineering workshop
Ghent, Belgium
http://www.ict-fire.eu/events/fire-engineering-workshop.html

13 November 2012
Net!Works Event 2012: Towards Horizon 2020!
Brussels, Belgium
http://www.networks-etp.eu/meetings-activities/general-assembly/ga8/ga8-info.html

15 – 16 November 2012
FOKUS FUSECO Forum 2012
Berlin, Germany
http://www.fokus.fraunhofer.de/en/fokus_events/ngni/fuseco_forum_2012/index.html

3 – 14 December 2012
World Conference on International Telecommunications (WCIT)
Dubai, United Arab Emirates
http://www.itu.int/en/wcit-12

Sn@pshot

Printable Origami robot

An insect-like robot printed and designed using the new process being developed to revolutionize the way robots are developed. The robot could be used for exploring areas inaccessible to humans.

Photo: Jason Dorfman, CSAIL/MIT

Contents

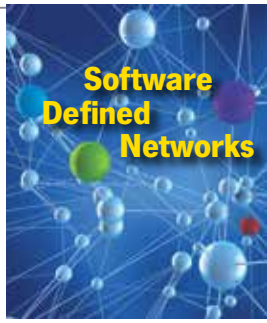

NEWS

Imprint

EURESCOM mess@ge, issue 2/2012 (June 2012)
ISSN 1618-5196 (print edition)
ISSN 1618-520X (Internet edition)

Editors: Milon Gupta (editor-in-chief), Peter Stollenmayer, Anastasius Gavras, Uwe Herzog

Submissions are welcome, including proposals for articles and complete articles, but we reserve the right to edit. If you would like to contribute, or send any comments, please contact:
Eurescom mess@ge · Wieblinger Weg 19/4 · 69123 Heidelberg, Germany
Phone: + 49 6221 989-0 · Fax: + 49 6221 989-209 · E-mail: message@eurescom.de

Advertising: Luitgard Hauer, phone: +49 6221 989-405, e-mail: hauer@eurescom.eu
Eurescom mess@ge is published three times a year. Eurescom mess@ge on the Web: www.eurescom.eu/message

Software Defined Networks

An overview

Anastasius Gavras
Eurescom
gavras@eurescom.eu

Adam Kapovits
Eurescom
kapovits@eurescom.eu

Software Defined Networks, short: SDNs, promise to solve many issues of today's network through virtualisation. In this article, we will provide an overview on what SDNs are, the networking issues they address, the players pushing SDNs as well as the open issues.

Issues of today's networks

Communication networks have become a critical component of all infrastructures in society. However, the industry and their designs have only managed to keep pace with the ever growing requirements by adopting point solutions, which have serious shortcomings. Thus, real networks became exceedingly complex, intransparent and inflexible with an ever increasing number of distributed protocols and the use of closed and proprietary interfaces within. In this environment by now it is too difficult, if not impossible, for network operators, third parties including researchers, and even vendors to innovate or innovate fast enough to meet their customer requirements.

The net result is that firstly, networks continue to have serious known problems with security, robustness, manageability, mobility and evolvability that have not been successfully addressed; secondly, their capital costs have not been reducing fast enough and operational costs have been growing, putting excessive pressures on network operators; and thirdly, network operators find it difficult to introduce new revenue generating services on their expensive infrastructures.

The concept of software defined networking (SDN) emerged in the last couple of years as the way to reap the benefits of virtualisation and overcome the issues cited above.

Those of us, who have been dealing with networking long enough to remember active and programmable networks might have a déjà vu feeling about software defined networks. The similarity in concepts is certainly striking. As active networks never really reached the market, let us see how software defined networks fare.

Definition and background of SDN

SDN is a new approach to networking, and its key attributes include: separation of data and control planes; a uniform vendor-agnostic interface called OpenFlow between control and data planes; a logically centralized control plane with an open interface; network control and management features in software; and slicing and virtualization of the underlying network.

The figure shows how OpenFlow and Software Defined Networking differ from the current networking mainstream.

In a broader context, software defined networking is part of and belongs to the network virtualisation framework, described by the ITU-T recommendation Y.3011, which was drafted by the ITU-T focus group on future networks.

Drivers of SDN

The Open Networking Foundation (ONF, www.opennetworking.org), which is standardising OpenFlow and thus represents the most important driving force behind SDN, was formed in 2011. Six companies who own and operate some of the largest networks in the world founded ONF – Deutsche Telekom, Facebook, Google, Microsoft, Verizon, and Yahoo! The board of ONF consists of the representatives of these companies as well as professor Nick McKeown from Stanford University and professor Scott Shenker from UC Berkeley, who were instrumental in setting up ONF. In the meantime, NTT also joined the board, and the number of ONF members grew to 70, including industrial heavyweights and start-ups alike.

Virtualisation as a key concept

Resource virtualisation is the fundamental underlying concept. It is used in many areas of information and communications technologies, as it promises a higher utilisation of expensive resources. The trend to virtualisation, supported by the emergence of Service Oriented Architecture (SOA), enabled the inception of practically everything as a service, including Network as a Service.

In the networking area there is currently a high interest in understanding how to define, deploy, manage and dynamically adapt network properties by applying software engineering principles. The objective is to enable functional programmability and elasticity of network provisioning combined with virtualisation of connectivity and other associated resources such as computing and storage resources. And indeed, this is the direction taken by the ITU-T focus group that has defined future networks as full service networks.

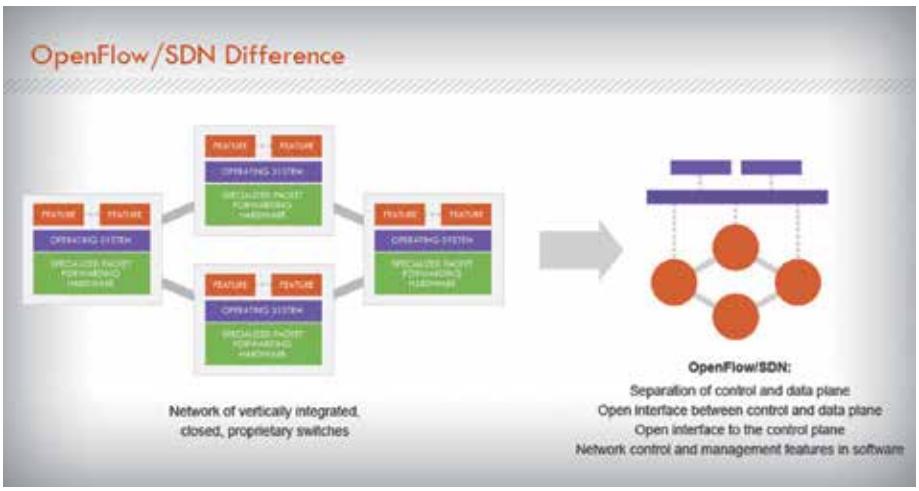

Figure: Software Defined Networking (SDN) and OpenFlow (Source: www.opennetsummit.org/why.html)

Current status

Currently SDN and OpenFlow receive a lot of attention. The latest event of the Open Networking Foundation, the Open Networking Summit in Santa Clara, California, in April 2012, was completely sold out and attracted about 900 participants. However, SDN and OpenFlow are on the right track to become more than just the next hype.

Google not only sits on the Board of the ONF, but it provided an important and rare insight into its operations giving a presentation at the Open Networking Summit and clarifying that it already runs all its internal data-centre interconnections on SDN, which it calls the G-Scale network. It took Google less than 18 months to convert its infrastructure into an SDN based one, using custom 10 Gbit/s switches with 128 ports built by Google, based on commodity components, as the necessary hardware was not yet available on the market.

The adoption of SDN and OpenFlow allows Google "centralised traffic engineering" with close to 100% utilisation of resources – 40% is considered to be good by current industry practices. Urs Hölzle, Senior Vice President Technical Infrastructure at Google, cited operational benefits, reduced network management and administration efforts and easier roll-out of new functionalities and features as main benefits of SDN and OpenFlow. For example big data moves can be planned to the last detail by simulating everything offline without touching the live network.

Although Google's G-Scale network is currently by far the world's largest production OpenFlow network, the technology is also used in other business networks, such as the one by Genesis Hosting Solutions and Nippon Express Corp.

Challenges for SDN

Even if the concept is appealing, history has taught us that nothing comes for free. SDNs will hide the complexity of the network to customers and promise to reduce operational expenses for operators. Nevertheless, the cost is hidden in the implementation of the necessary capabilities to offer these services. In particular, research has a number of unsolved topics on its agenda that need to be addressed in the near term.

Among the most important topics is the management of resources across borders or administrative domains and the strict isolation of traffic and data in shared resources, to name a few. In this context, please note that even the Google implementation is relatively simple, as it is done within one administrative domain, so there are no cross-domain issues.

Concepts and prototypes exist to address cross-domain management issues. They have to be trialled, potentially at large scale, in order to gain the necessary confidence for production deployments. This is the reason that apart from research work on OpenFlow and SDNs in a more general sense, large-scale testbed environments have been deployed worldwide.

Conclusion

The networking community embraces Software Defined Networks at high speed, driven by a number of factors, such as cost, the explosion of traffic on the Internet, and a high demand for customised service networks. Foresight studies have formulated high and very divergent requirements for future networks. These requirements originate from applications needing very high-capacity service networks, from mass market applications that are latency sensitive, and from applications that have to satisfy the requirements of cyber-physical systems.

The needed flexibility and the features of full service networks are a precondition for a long term sustainable information infrastructure that supports society. Time will tell, to what extent SDNs can do the job.

SDN standardisation is rapidly arising

Takashi Egawa
NEC Corporation
t-egawa@ct.jp.nec.com

Software Defined Networking (SDN) has been attracting various standardisation activities. The current movement of SDN comes from the emergence of OpenFlow, promoted by the Open Networking Foundation (ONF). ONF is energetically standardising OpenFlow and is gradually expanding its target to relevant areas. There are also other Standards Development Organisations (SDOs) that have started to take part in SDN activities. ITU-T SG13 is developing an overview to clarify the image of telecom SDN. IRTF is discussing to establish a Research Group to investigate necessary research items. Since SDN is a far-reaching activity, it will attract even more SDOs in the future.

SDN is gaining momentum, though the image and scope of SDN is not always the same among people, as was the case of the Internet or the Cloud. This means there are many issues to standardise, and various SDOs have started their work.

ONF and OpenFlow

The Open Networking Foundation (ONF) is an organisation which aims to commercialise and standardise SDN and its important component, OpenFlow. It was established in March 2011 by 6 operators with 17 member companies. One year on it has 70 members and is working energetically for innovating networks like a Silicon Valley start-up.

According to ONF, the ICT industry is becoming more and more reliant on software, which means standardisation should become more software-friendly, or should be developed like software. ONF's first output has been an OpenFlow switch specification that enables SDN control software to access and to manipulate a forwarding plane of network devices such as switches and routers. The specification has been updated in a short time period since ONF was established. Its latest version, 1.3.0, was approved in May 2012. Through its version-ups, the specification introduced various new features, such as IPv6 support, extensible expression and per-flow meters for QoS support.

ONF's second output has been the OF-Config protocol that enables to configure and manage OpenFlow switches and their ports/queues through, e.g., assigning controller to switches, capability discovery and tunnel configuration. It has been updated to version 1.1 in four months since the release of version 1.0. The OF-Test for interoperability testing and performance benchmarking is to be finalized in July 2012. It will be used at ONF Plug-Fests.

ONF started its activity from OpenFlow switch specification and closely related issues. They are gaining momentum, in particular in data centres, and now ONF is expanding its work areas. It is currently investigating how to make OpenFlow and legacy switches co-exist. It is also investigating through use cases the applicability of SDN in other areas, e.g., networks of enterprises, campuses or telecom carriers. And it is exploring the appropriate level of standardisation of an API between SDN control plane and applications.

ITU-T SG13

ITU-T Study Group 13, a group for telecom network architecture, has been investigating the networks of the future through the Focus Group on Future Networks (2009-10) and one of its groups, Question 21 (Q.21). They approach future networks from the top, and developed Recommendation Y.3001 that describes high-level objectives and design goals of future networks, and Recommendation Y.3011 that describes network virtualization framework. SDN is considered promising to realize these Recommendations, and in February 2012 Q.21 started to develop an SDN framework to clarify the SDN concept, problem space and terminology.

IRTF/IETF

IRTF/IETF is also paying attention to SDN. Current playground is IRTF, and the scope of SDNRG (SDN Research Group) is being discussed. One proposal contains a classification of SDN models and their relationship to existing IETF work, network description languages, abstraction and APIs, etc., though the focus is rather on research aspects, not standardisation directly.

Outlook

The activities outlined above are herald of SDN standardisation. SDN is a very generic and far-reaching concept. Their properties such as control-data plane separation, centralized network intelligence, and network abstraction for applications have a long history in telecom carrier networks, for example. Recapping their essence will inevitably lead to the involvement of many SDOs. The SDN concept also covers not only the physical layer, optical or core networks, but also applications and terminals, because in SDN, the network becomes a peripheral of terminals, finally. This is another factor for SDO involvement. Various SDOs in the ICT industry will have to contribute in their respective degree.

Further information is available on the ONF website at www.opennetworking.org

The future of Software Defined Networks

Interview with Alex Galis from University College London

Software Defined Networks, or short: SDNs, are one of the latest technology hypes in telecommunications. It remains to be seen to what extent SDNs will revolutionise the networks. In order to get a better picture of SDNs' relevance and impact, Eurescom mess@ge editor-in-chief Milon Gupta interviewed networks expert Alex Galis, Visiting Professor at University College London. Professor Galis has served as a Vice Chair of the ITU-T Focus Group on Future Networks. His numerous publications include a book on the first generation of the SDN technology: "Programmable Networks for IP Service Deployment" (A. Galis et al. 2004).

Alex Galis

What is your definition of Software Defined Networks?

Galis: I prefer to consider SDN in the context of what I call the Software Driven Network, short: SdNET. An SdNET is a network architecture and self-control infrastructure for interoperability of heterogeneous wireline and wireless networks, supporting in-bound software driven features and their optimal integration with the connectivity layers.

The network control infrastructure encompasses distinct network and service environments, including Network Clouds respectively Virtual Networks, G4 controllers, Open Flow enabled routers, which are known as Software Defined Networks, and Network of Smart Objects. It is a programmable network infrastructure that supports in-bound software driven features, which can be instantiated on-demand, based on the changing requirements and resource constraints.

Features include, for example the utilisation of shared and integrated connectivity, storage and processing resources, and a qualified access mechanism to a set of network embedded resource-facing services. Other features include functional programmability and elasticity as well as integrated content and data access.

Are SDNs really something new or just existing concepts combined under a new name?

Galis: SdNETs and SDNs have their roots in programmable and active network technology and standards developed a decade ago and in particular in the decoupling of forwarding and control planes with open interfaces to connectivity resources. However, SdNETs, and to some extent SDNs, are aiming at a novel integration and use all connectivity, and storage processing resources under new management and control systems for provisioning of on-demand networking and servicing.

What are the main benefits of Software Defined Networks?

Galis: SdNETs, and to some extent SDNs, have a couple of benefits. They empower, for example, the network with in-bound software driven features. This enables a large number of software-driven features in networking, allowing cost-effective, rapid deployment of a network and service infrastructure on demand.

Furthermore, SDNETs and SDNs support the increase in the demand of network services and networked applications as well as the increase in users' expectation with regard to efficiency, cost, and personalisation. They also enable new business opportunities by exploiting the virtual service-provisioning role as well as the efficient integration of new wire and wireless technologies and legacy systems.

Finally, SDNETs and SDNs enable network operators to efficiently manage the network infrastructure through explicit OPEX enabled by software driven features.

Which major challenges and open research questions need still to be addressed?

Galis: SdNETs, and to some extent SDNs, pose a number of challenges. An important challenge is the development of open control systems made from primitives and/or exiting control elements with open APIs and in-bound software driven features over virtualized and/or existing heterogeneous connectivity, storage and processing resources.

Other challenges include the development of in-system self-management for efficient provisioning, operation and cessation of networked services, and the development of the second generation of functional programmability for dynamic and autonomous activation of in-bound software features.

Furthermore, another challenge to be tackled is the development of network infrastructures as multiple and elastic service and application execution environments.

Finally, there is the challenge of developing in-bound software driven network testbeds on a large scale.

Do you expect large-scale deployments of SDNs in the next five years?

Galis: Testbeds for SdNETs and SDNs are under development. I expect large-scale deployments in the second half of this decade.

What impact will SDNs have on network operators and suppliers?

Galis: In the next decade SdNETs, and to some extent SDNs, will uniquely enable network operators to move forward from the existing business of managing inflexible and commoditisable infrastructures towards software-enabled network features and service execution environments. Such a distinctive opportunity could significantly increase the ICT sector's direct share of the European and global GDP. Currently, the ICT sector is directly responsible for approximately 10 percent of Europe's and the world's GDP.

OpenFlow for high capacity transport networks

Dimitra Simeonidou
University of Essex, UK
dsimeo@essex.ac.uk

Reza Nejabati
University of Essex, UK
rnejab@essex.ac.uk

OpenFlow is a technology enabler for the implementation of Software Defined Networking, which allows the development and deployment of innovative network applications, independent of the underlying network technology. This article describes research done at the University of Essex, which pioneered the implementation and demonstration of OpenFlow at the optical layer as means of enabling Software Defined Optical Networking.

OpenFlow – an enabler for SDN

OpenFlow (OF) has been recently proposed as a control framework that supports programmability of network functions and protocols (i.e. software defined networking or SDN) by decoupling the data plane and the control plane, which are currently vertically integrated in most network equipment (e.g., routers, switches, access points). OF allows user defined (software defined) network applications to take control of the network infrastructure and therefore paves the way for collapsing the traditional layered architecture.

In essence, OF is using an open standard protocol, which allows an external controller to communicate with a switch. It is based on flow switching and traffic control, with the capability to execute user defined routing, control and management applications in software outside the data path, in the OF controller.

In an OF-based network a centralised controller utilises the OF protocol to directly and securely manipulate the forwarding engine of any network element, such as Ethernet switches (see Figure 1). Using flow abstraction, the centralised controller is able to configure the forwarding table of the switches.

SDN at optical layer

Although the current implementations of the OpenFlow protocol are mainly addressing packet switched domains (i.e. Ethernet LANs), its extension to support circuit switched and specifically optical networks can provide a new framework for evolving carrier grade networks. Enabling SDN via OF at the optical layer will allow to decouple network control and management functionalities from the network elements and their vendors. It will also support the integration of electronic packet and optical circuit switched networks for access, metro and core network segments.

There have been several attempts and proposals to control both circuit switched and packet switched networks using the OpenFlow protocol. Recently, University of Essex in joint development work with ADVA Optical Networking have experimentally deployed an extended OF solution (i.e., OF protocol, controllers along with a sample network application) in a converged packet and optical circuit switching setup. This has showcased for the first time a full end-to-end network demonstration using commercial Reconfigurable Optical Add Drop Multiplexers (ROADMs).

An advanced OpenFlow network testbed

The High Performance Networks group at the University of Essex has built a unique and feature rich OpenFlow network testbed that enables advanced experimentation on network infrastructure and distributed applications. The testbed is currently being deployed in the EU projects OFELIA (www.fp7-ofelia.eu) and FIBRE (www.fibre-ict.eu), which showcase the potential of the facility by demonstrating experimental network applications deployed by network researchers at a global scale.

The OF infrastructure at Essex, depicted in Figure 2, includes a multiplicity of OpenFlow-enabled devices and network domains, such as Ethernet switches (campus and carrier grade) and commercial optical WDM and fibre switching equipment, ultra-high-definition video streaming applications, and general purpose servers and network attached storage to host technology pilot applications. The group is currently investigating the feasibility of extending the OpenFlow protocol to a sub-lambda (or multi-granular) optical switching platform. The Essex Island is connected to the Janet and GEANT networks which provide extensive national and international connectivity and which enable the island to function as the European hub for interconnection of OF testbeds beyond Europe and experimentation at a global scale.

Conclusion

Software defined networking enabled by OpenFlow decouples the control of the network from its data plane, and therefore it allows network designers and operators to simplify network operations and deploy innovative network applications. The High-Performance Networks group at the University of Essex has pioneered the potential of OF extensions towards circuit switching and particularly towards optical networks and demonstrated it in a large scale converged packet switched and optical circuit switched network testbed.

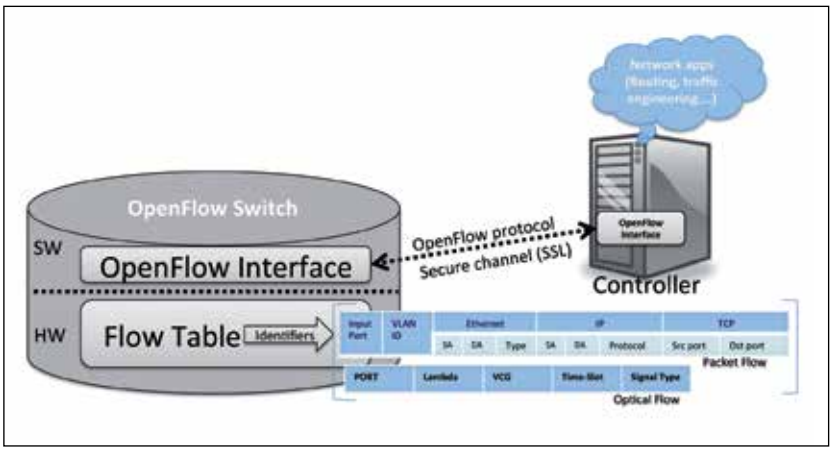

Figure 1: OpenFlow enabled resource controlled by Controller

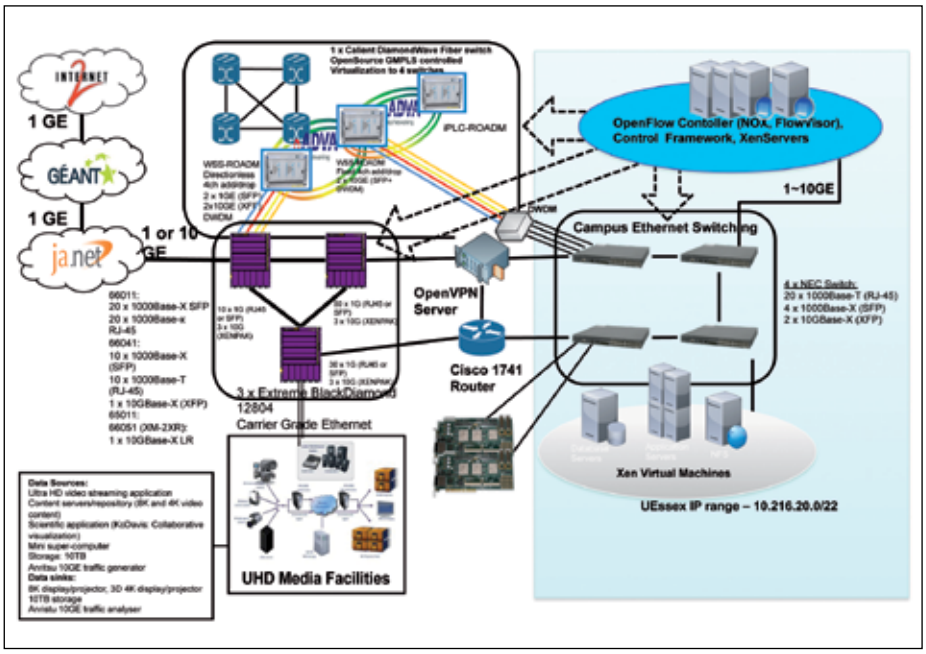

Figure 2: University of Essex OpenFlow testbed facility

Further information on OpenFlow is available at www.openflow.org. For the end-to-end network demonstration of an OF solution mentioned in the text, see the ADVA press release at [www.advaoptical.com/en/newsroom/press-releases-english/20120510.aspx

The HPN group and the OpenFlow island facility described in this article are due to move from University of Essex to University of Bristol, UK, with effect 1st of September 2012.

An architecture for integrating OpenFlow functionality in experimental networks

Spyridon Tompros
s.tompros@creativese.eu

The integration of OpenFlow in experimental networks poses some interesting architectural challenges. This article describes how these challenges have been addressed in OpenLab, an Integrating Project in the EU's Seventh Framework Programme (FP7) in the area of Future Internet Research and Experimentation (FIRE).

The OpenLab architecture

One of the goals of the OpenLab project is to enable experimentation on new protocols and services on commercial grade networks that are in the reach of every user (researcher, vendor, SME). There are two IMS networks in the OpenLab testbed offerings, one from University of Patras, Greece, and another one from TSSG/WIT, Ireland. Within OpenLab the deployment of OpenFlow over these two IMS testbeds has been scheduled. To enable experimentation over these testbeds, mechanisms from a previous project called Panlab (www.panlab.net) were adopted.

The architectural overview of Panlab is depicted in Figure 1. The main roles identified are: i) The Panlab Partner who acts as the provider of infrastructural elements. Partners are connected to the Panlab Office for offering functionality along with their respective resources to the customers. ii) The Panlab customer who utilises services provided by the Panlab office and defines federation scenarios. iii) The Panlab Office, which is the central federal entity that acts as a provider and broker among different administrative resource domains. It realises mechanisms that enable Panlab partners to be part of a federation, and Panlab customers to define federation scenarios.

Panlab's architecture introduces additional components for integrating testbeds that belong to various administrative domains, in order to become available to participate in federation scenarios. The first component is the Panlab Testbed Manager (PTM – Figure 1). PTM is responsible for accepting RESTful commands

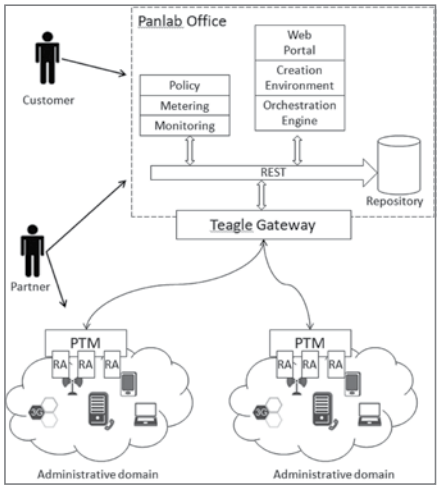

Figure 1: Architectural overview of Panlab

from the Teagle Gateway in order to configure the domain's resources. PTM implements the so-called Resource Adaptation Layer where Panlab partners "plug in" their Resource Adapters (RA). A Resource Adapter (a concept similar to device drivers, Figure 2) wraps a domain's resource API in order to create a homogeneous API defined by Panlab. Details and specifications of Panlab's components can be found at nam.ece.upatras.gr/fstoolkit.

Towards an OpenFlow integration

Figure 2 displays the current OpenFlow deployment at the University of Patras. Currently three XEN servers have been deployed where each one is capable of hosting a number of virtual machines (VMs). On each machine Openvswitch is installed, replacing Linux networking. This virtual switch is configured with two virtual bridges, each one connected to a network interface of the host. There are two networks: i) one public network for accessing each host (and eventually each vswitch) over the Internet. This is also used by experimenters for accessing machines for an experiment. ii) a data network for VM data traffic. This network is used to send traffic to and from an external VM on another host. When a VM is requested, networking XEN scripts configure this VM to have two virtual interfaces. Each is attached to virtual bridges created on the open vswitch.

To expose the testbed for experimentation (i.e. reservation, provisioning) we adopted the Panlab mechanisms. Resource adapters were developed to configure the resources of the testbed. Currently we are developing resource adapters that are able to create virtual machines to the hosts and attach their virtual interfaces to the virtual open vswitch. In addition, other resource adapters are able to configure ingress and egress policing rates and bursts, thus affecting the available QoS of applications. Finally, there is a resource adapter that gives the ability to the experimenter to configure an sflow monitoring client.

Figure 3 shows the configuration workflow. Our experiment controller called Federation Computing Interface (FCI) is used by the experimenter to request and configure resources. Each request through our Teagle Gateway is directed to our domain manager (PTM) where our resource adapters are located.

To access the testbed resources the user can use the Federation Scenario Toolkit. The user can create the scenario and configure the requested resources by expressing the experiment in the Federation Scenario Description Language (FSDL).

For better granularity of experiment control, the user can use directly the Java based FCI API. An example is given in Figure 4 where the experimenter can use a programmatic way through java to define and runtime control and monitor resources.

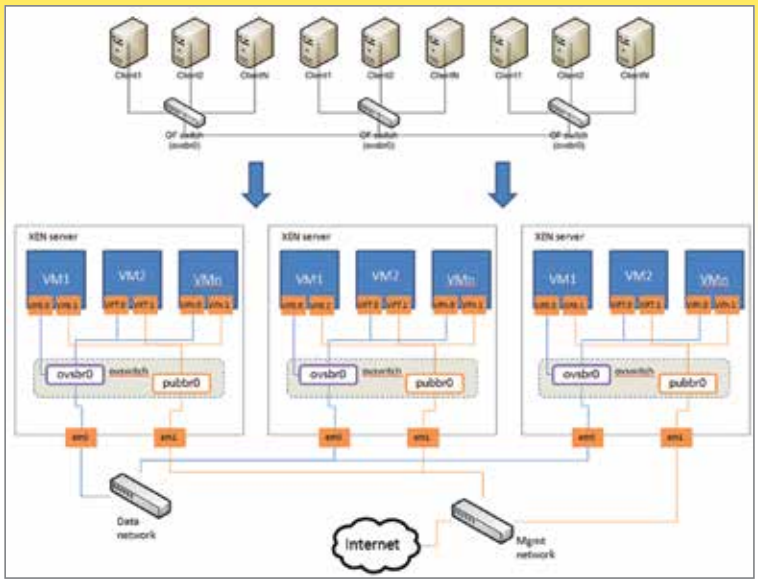

Figure 2: OpenFlow deployment at the University of Patras

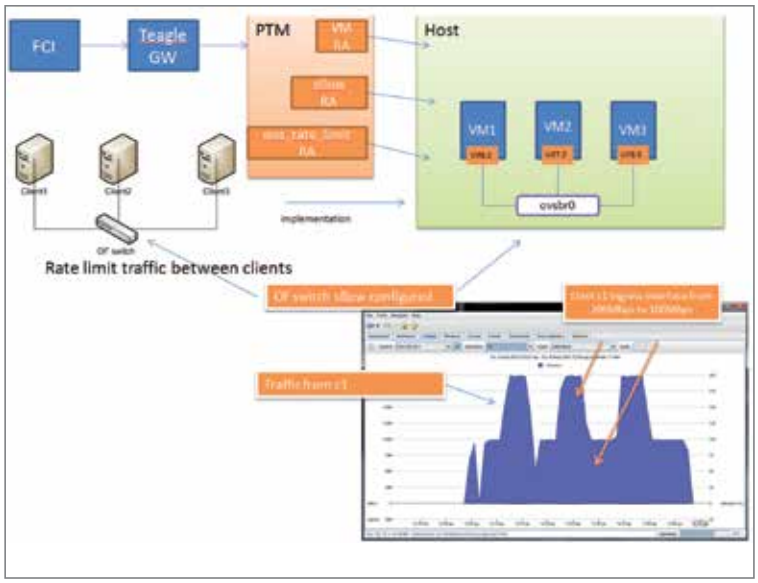

Figure 3: Enabling the testbed for experimentation

```
private void CreateScenario() {
//    Group (for grouping resources)
    ResourceGroup myGroup = FCI.getInstance().createResourceGroup("myGroup"
    //all creates
    if (createsflow){
        ResourceProxy resource_mysflow = createResource_mysflow();
        System.out.println("mysflow resource GUID: " + resource_mysflow.get
        myGroup.addResourceProxy(resource_mysflow);
    }
    ResourceProxy resource_myqos_rate_limiting = createResource_myqos_rate_
    System.out.println("myqos_rate_limiting resource GUID: " + resource_myc
    myGroup.addResourceProxy(resource_myqos_rate_limiting);

    try {

        Thread.sleep(5*60*1000);
        RemoveLimit( resource_myqos_rate_limiting );
        Thread.sleep(5*60*1000);
```

Figure 4: Controlling the experiment resources through Java API

During experimentation the experimenter can also access the switches through public IPs. The installation of user virtual machines is also automated. To access installed machines an elastic IP can be assigned to one of the requested machines, so the user can access the rest of the machines. The testbed will also be federated with other SFA enabled testbeds, including PlanetLab.

Further information on OpenLab is available at www.ict-openlab.eu

Celtic-Plus
Smart Connected World

Editorial

Dear reader,

Celtic-Plus is now in its second year, and it is now time to look in detail at the new accelerated call process with its two calls per year. As the article in this issue will show, the project set-up time has become shorter by several months, as we had envisaged. At the time of writing, Celtic-Plus has just closed its Spring 2012 call. In this first call of the year, nine proposals had been submitted. These proposals will undergo now a technical and national review process to eventually receive the Celtic-Plus label on 3rd July.

As usual we present to you some selected projects, which have recently closed or will close very soon. Since we have now also an increased participation from Korean organisations, we present to you the experience of UBRIDGE, Korea, in working in the collaborative project WiSafCar.

Enjoy reading this issue.

Heinz Brüggemann
Director Celtic Office

CONTENT

Getting faster from proposal to project

First experience from the new Celtic-Plus Call process

Heinz Brüggemann
Director Celtic Office
brueggemann@celticplus.eu

With the start of Celtic-Plus a new call process was introduced which aims to shorten the duration from proposal submission to project start by, at least, six months. From the first new call's experience, this goal is likely to be achieved.

In the first Celtic phase (2003 to 2010) the call for proposals was, like in most other clusters, based on a two phase process with a first proposal outline phase and, for those proposals which were invited, a second full proposal phase. When we discussed improvements at the start of Celtic-Plus, the question was if the advantages of maintaining a two-phase proposal process would outweigh the benefits of approaches aimed at reducing the overall call duration via a single and much shorter full proposal phase.

After discussing the pro and cons of a shorter single-phase process versus maintaining the two-phase call process, the Public Authorities and Celtic concluded to start a pilot by running two fast calls per year without considering a proposal outline phase. As a compensation of the skipped proposal-outline phase, the Office offers a first check of a project idea before elaborating a full proposal. Another important conclusion was that the Public Authorities will ask for funding applications, at least, on a provisional basis already around the proposal submission date. This early application, even without a label decision, will help to check for eligibility of a national consortium and should speed up and better synchronize the funding decisions of the involved countries.

Faster project set-up

After the first two calls in 2011, there are very strong indications that the new accelerated call process resulted indeed in a significant reduction of the project set-up time.

The average delay from first proposal submissions to the (planned) start of the project changed from an average of 12.2 months to now 5.7 months in Celtic-Plus. This means that the call duration has indeed been reduced, as envisaged, by around six months. The main reason for the reduced overall set-up duration is the fact that two full calls every six months have been introduced. But also the delay from labelling to the planned project start has been reduced by about two months, compared to average figures of the earlier Celtic projects. Note: at the moment of writing not all 2011-call projects have been finally launched, and the delay from the current planned start dates could still be longer than envisaged. This, however, should not have a strong impact on the total average duration (see figures 1 and 2).

Conclusion

The new accelerated call process has been well accepted by the project proposers and by the involved Public Authorities. The main goal of reducing the overall set-up time appears to have been reached. The new call process requires also some more flexibility from the Public Authorities, which have to accept also preliminary or provisional funding applications, even at a stage where a proposal has not yet been labelled. Many Public Authorities have found useful and pragmatic solutions to check, at least, the eligibility of proposals and involved consortia. These early checks help considerably to speed up the approval processes among the involved Public Authorities.

Even though the proposers of new projects have no longer a default possibility to receive a review feedback on their proposal outline before doing the full project definition, Celtic-Plus still offers the opportunity to have a project idea assessed by the Office, and the feedback may be used by the proposers to improve their full proposal. There is also an additional opportunity for essentially good proposals which, however, were not sufficiently elaborated and did not receive an immediate label as they could be asked to resubmit their proposal to the next call. This possibility means that, to a certain degree, a kind of two-phase submission process would still be applicable in some accepted cases.

At a recent discussion with the Public Authorities the new call process was regarded as an improvement and acceleration of the project set-up, even though a number of improvements are still needed, in particular the early contact and funding application around submission time.

The new call process will be further improved and continued in 2012.

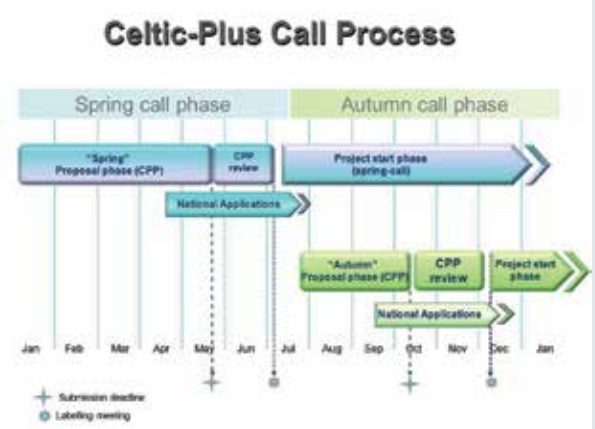

Figure 1: New Celtic-Plus call process (dates for 2012)

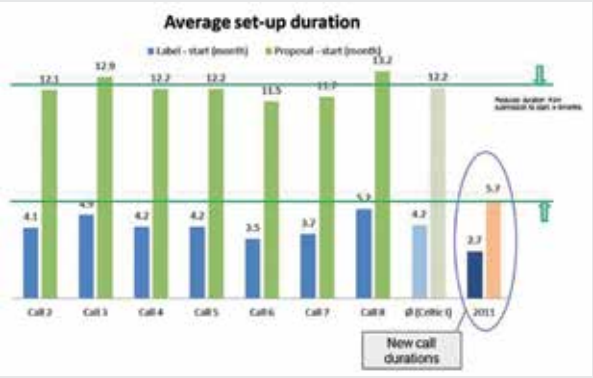

Figure 2: New versus former set-up durations (submission to start and labelling to start)

Experiences of UBRIDGE, Korea, in Celtic project WiSafeCar

KB Lee
UBRIDGE
velojazz@ubridge.co.kr

Europe and Korea share already a growing market, which has been enabled via the Korea-EU Free Trade Agreement. Collaborative research projects like WiSafeCar are the first step towards a more comprehensive and deeper commercial partnership.

UBRIDGE's core business consists of communication terminals and solutions for vehicles. A current project is, for example, Smart Highway with WAVE technology. This is a government-driven development project regarding ITS, on which UBRIDGE is currently working in Korea. A large number of research projects on WAVE technology were already under way in Europe, and it seemed likely to have other possibilities for growing our performance through an international collaborative research project.

In 2009, Korea had become the first Asian, non-European 'associated country' member of EUREKA. UBRIDGE had been funded by KIAT, the Korean Institute for Advancement of Technology, for finding European partners. KIAT is in charge of the national funding of EUREKA and FP7.

Fortunately, UBRIDGE found the opportunity to step into the already ongoing project WiSafeCar. Carrying out the WiSafeCar project with European consortium members was a valuable practical experience to see the real potential for expanding our horizons. This cumulative experience might be the technological entry barrier, because newcomers should acquire plenty of standards to meet all demands in a short time, if they are not involved in its early stages.

The first year of WiSafeCar was for us the learning stage about the process and how to cooperate transnationally in many ways, such as documentation, meetings and so on, which we managed also with the support from project members. From the second year, UBRIDGE and Finnish partners had produced integrated results based on independent technologies. One of the things we have learned through the project is that we need to communicate as much as possible and to meet in person as often as possible. Moreover, it was the team spirit, the ability to perceive and learn that we are different, which enabled us to generate significant synergy effects. Through the project WiSafeCar, we know now that team spirit also transcends national boundaries.

The upcoming project CoMoSeF, an extension of WiSafeCar, is going to be started soon, based on mutual trust and respect as the technical collaboration demands. In CoMoSeF, UBRIDGE plans to contribute to and research on SMART HMI, which is currently an active research area in Korea as well. In addition, a business collaboration with consortium members is also discussed in detail.

Recently, the Korean EUREKA Days, organized by KIAT, have been held in Korea and in Europe. The events aimed to create a market-oriented industrial R&D network and to establish connections to look into possibilities for mutual benefits. Even so, as an SME in Korea, it is not easy to meet reliable partners and establish trusting relationships, and vice versa. We have heard of the rising number of companies participating in global R&D projects, like UBRIDGE, who had gained already successful experiences.

Conclusion

We believe that the best way to have a trustworthy partnership across continents is to build relationships through a collaborative research project. Please let me know, if you are interested in innovative technologies and if you are looking for cooperation partners in Korea. UBRIDGE would like to be the bridge between European partners and Korean partners.

WiSafeCar

Co-operative traffic safety system for challenging road conditions

Pekka Eloranta
Mobisoft Oy
pekka.eloranta@mobisoft.fi

Changing weather conditions, icy and slippery roads, accidents, congestion, roadworks, approaching emergency vehicles – all of those may cause dangerous situations, if there is no proper information to inform and warn drivers. The WiSafeCar project developed a communication system for drivers to provide them with the information needed.

The aim of the WiSafeCar system has been to offer an effective tool to prevent serious or fatal accidents. The commercialisation of the system will take place during the next few years. The results of the project are totally in line with the European level (ITS Directive and ITS Action Plan) and national strategies of the participating countries (Finland, South Korea and Luxembourg).

Increased traffic safety and fluency

According to studies, the risk of an accident, when there is snow or ice on the road, is increased by a factor of 4.1 compared to normal conditions. The climate change has also changed the road weather conditions and made driving demanding. The changes in conditions may increase rapidly. To avoid accidents and related injuries, road users need to be informed about the changes, accidents and incidents immediately, when detected.

WiSafeCar developed an effective communication system and an intelligent wireless traffic safety network between vehicles and infrastructure by using both sensor and observation data to offer secure and reliable real-time services for vehicles and road users. WiSafeCar focused on building a comprehensive, secure and reliable solution for V2I (vehicle to infra) and V2V (vehicle to vehicle) communication.

The WiSafeCar system allows vehicles to transmit traffic and weather observations to a traffic service centre, where the data are analysed and conveyed to other vehicles on the road, together with information on road conditions. The system also alerts drivers in case of imminent danger of an accident due to slippery road, etc.

Co-operative driving, where vehicles and intelligent transport systems communicate and exchange information with each other, is becoming more common. It has been estimated that such systems will become a standard feature in most vehicles between 2015 to 2030. The development work is active and rapid, particularly in Europe, the USA and Japan.

The WiSafeCar project carried out short-range data transfer testing compliant with the IEEE 802.11p standard, and long-range communication based on a mobile phone network. This kind of a "hybrid system" is based on the "Day one system" objectives of the standardisation organisations (CEN, ETSI, ISO).

The results of the project were extremely encouraging. WiSafeCar was able to pilot different kinds of services related to road weather, traffic conditions, or incidents for the drivers, based on the common WiSafeCar service and communication architecture.

The WiSafeCar project involved 9 companies and research organisations from Finland, Luxembourg and South Korea: Mobisoft, Finnish Meteorological Institute, VTT, Infotripla, Taipale Telematics, Sunit, University of Luxembourg, CRP Henri Tudor, Ubridge. The project was started in 2009 and ended in March 2012. The total budget of the project was 5.9 million euro.

Deployment and commercialisation of the WiSafeCar results have already started with traffic service development, and more efforts and results will be forthcoming over the next few years. The project consortium aims to continue the dissemination and deployment of the project results under a new Celtic-Plus project named CoMoSeF (Co-operative Mobility Services of the Future).

Conclusion

The main objectives and challenges of the WiSafeCar project were the generation of true V2I and V2V communication, efficient delivery of critical data regardless of the location, speed or presence of other vehicles, and generation of services to enhance traffic safety and efficiency. The ultimate goal was to create an intelligent communication solution for vehicles where they can deliver their own observations of traffic and weather conditions to the platform core. This information was refined, processed and delivered back to the vehicles as analyses and forecasts about road weather conditions, immediate accident or incident warnings, congestion warnings and other services.

The WiSafeCar solution has proven in the tests and piloting to have clear potential for a comprehensive vehicular communication entity, with promise of increasing traffic safety and fluency and decreasing the amount of accidents and lives lost in traffic.

Further information is available at www.wisafecar.com

GENIO

Platform for next generation home networks

Carolina Benito
Instituto Tecnológico de Aragón
cbenito@ita.es

Rafael del Hoyo
Instituto Tecnológico de Aragón
rdelhoyo@ita.es

Nowadays, some of the Internet's big challenges are the viral growth of the number of connected users, devices, services and user-generated contents, which turns the user's environment, especially the home, into an increasingly complex and hostile environment for telecommunication operators and service providers.

To address these challenges the home systems need to remain simple, accessible and transparent for the user, which requires research on technologies enabling intelligent and autonomous computing and new ways of interacting with future homes. Autonomic computing systems can manage themselves at given high-level objectives. The GENIO project integrated autonomic computing, mobility, energy efficiency, security and new interactive user mechanisms like virtual agents, to realise the future smart homes. These homes are able to detect the people inside, and to self-configure by personalising the services for each user and detecting new devices plugged to the house. They are also able to self-optimise by disconnecting lights or closing doors, if people are not present; to self-heal by controlling sensors to prevent problems related to physical and software elements; and to self-protect by identifying the current users at home, and preventing external attacks.

The Smart Home in GENIO

The GENIO project has defined the Smart Home with the following features:

Figure: GENIO smart home features and components

- Advanced self-management of the home network: GENIO applies autonomic technologies to the home network configuring, service organization, supervision, diagnosis and healing tasks, aiming at maximizing home network automation and intelligence to deal with events and alarms that may arise, taking into account the energy consumption of devices.
- Heterogeneity of the devices and their interactions: GENIO enables the interconnection and intercommunication of different devices in the home network.
- Ubiquitous access to home contents: The main aspects tackled by GENIO are the access to content from everywhere, allowing nomadic usage.
- Energy management of the home devices: GENIO integrates an energy monitoring and management component, based on the use of smart plugs and enables users to specify their preferences and get alarms.
- Personalisation: GENIO works on subscription identification solutions for personalised device configuration so that users are able to carry "their home" with them, using the same authentication mechanism in any Home Gateway, giving support to the concept "My home moves with me".
- Authentication: The project looked for a solution to guarantee authentication in efficient bandwidth utilization, allowing it from a local or a remote location.
- Reusing existing wiring in the home premises and providing hybrid networks to the home: PLC+WiFi, ONT+PLC, Coax+PLC, xDSL+PLC
- Communication with the user through an Intelligent Virtual Agent. GENIO also proposes solutions to satisfy users' demand of easy-to-interact-with services using portable and ubiquitous interfaces, an intelligent multimodal virtual agent, through voice, image recognition or natural language processing capabilities

Figure 1 shows the GENIO smart home features and different components.

Conclusion

As the number of services provided for end users at home through a broadband connection and the number of connected devices in the home network are increasing, there is a strong need to provide management solutions in the home network, which allow users to be unaware of technology. The autonomic computing implementation done in GENIO allows end users to have "plug & play", enabling them to enjoy their personalised services and mitigating the risk of not achieving a massive service adoption in the home environment because it is too complex to use them.
The management solutions allow operators and service providers to reduce the OPEX (operating expenses) of the services delivered to the users. This leads to user satisfaction and, in turn, to better loyalty.

You can find more information on GENIO at http://projects.celtic-initiative.org/genio

Kusanagi

Distributed networked multimedia platform

Philippe Martineau
Eureva
philm@eureva.fr

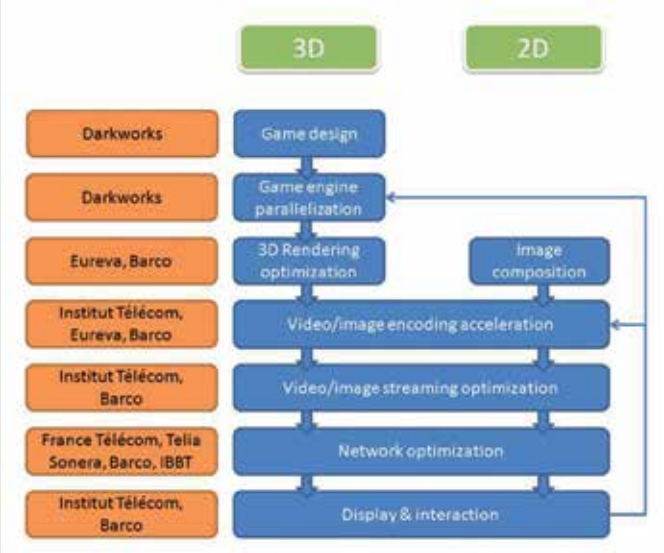

Technologies explored and KUSANAGI partners in charge

Kusanagi aims at optimizing rendering, compression, streaming, and display technologies in order to provide efficient remote rendering solutions for 2D and 3D content. The applications addressed in the project are 3D gaming and 2D ultra-high definition professional visualisation.

Various remote rendering solutions already exist, but none of them is well suited to the high requirements of 3D gaming or ultra-high definition visualisation: latency is a major problem in both cases and requires specific optimisations of the whole chain, from the rendering server to the display device. Another common goal is to use only commercial off-the-shelf hardware instead of dedicated hardware, and to concentrate the intelligence in optimised software or even firmware solutions.

Approach

To achieve its ambitious vision, the Kusanagi project has been driven by a strong consortium with extensive experience in R&D and industry leadership in blockbuster games as well as solutions for networked visualization. The consortium intelligently combines universities, research institutions, large telecom and industrial companies, SMEs and startups. In the long run, the partners of the Kusanagi project wish to be a driver of the European video game and professional imaging industries and bring strategic and technical value, in conjunction with the move from consoles and terminals to networked remote desktops and gaming services. The figure shows how the consortium is organised and who is in charge of the different technologies.

The project focused on the evolution of software architectures that will be critical to master fields of highly parallelised processing, distributed processing, high bandwidth and low latency networks and their intricate impact on content quality and speed of delivery, dataset size, optimisation complexity, open infrastructure viability, production tools, and middleware standardization efforts.

The specific aims of the project were to conduct research and perform breakthroughs around the following core topics:

1. Image composition and rendering (optimisation, distribution and tools): development of the base technology for a hardware independent rendering platform allowing efficient multi-core distribution.
2. Multimedia Compression: image/video compression processes were optimised, both from an algorithmic and an implementation standpoint.
3. Streaming / Distribution: emphasis has been put on protocol adaptability to network conditions and integration with rendering algorithms.
4. Network traffic characterization: research on network traffic patterns and optimization possibilities and Quality of Service requirements when highly interactive HD content is streamed on very low latency networks.

Achieved results

Kusanagi has been built upon existing technologies, and every link of this chain had to be optimised to provide a superior gaming or professional imaging experience. The Kusanagi prototype thus integrates all the optimised software components developed during the project and is able to compete with modern consoles and replaces existing professional visualisation solutions.

An end-to-end seamless infrastructure has been realised to publish, play and interact in real time with high-definition 2D and 3D networked multimedia content. This platform manages multiplayer gaming and relies on open-source encoding and streaming architecture standards (MPEG4 format, lossless codecs). Optimisations for real time in changing network conditions, like adaptive encoding and streaming according to network latency and bandwidth, have been studied and implemented both for wire and wireless networks. The platform offers the best and stable user experience compatible with real-time interaction constraints. Moreover, hybrid encoding and display methods for very large 2D displays have been developed to completely cope with wireless situations or ultra-high definition pictures.

Outlook

The expected impact of the Kusanagi project is to allow the gaming and professional imaging industries to progressively move from local rendering solutions to remote rendering solutions. It enables real-time visualisation applications to be deployed in Cloud infrastructures. In the gaming market, this means replacing consoles and physical data media by a game service accessed on a pay-per-use basis. In the professional imaging market, this means an increasing use of – possibly wireless – thin client devices with all computing resources in a central location.

You can find more information on Kusanagi at www.kusanagi.eu

R2D2 Networks

Enabling networks to ensure QoE for multimedia contents

David Flórez
Telefonica I+D
dflorez@tid.es

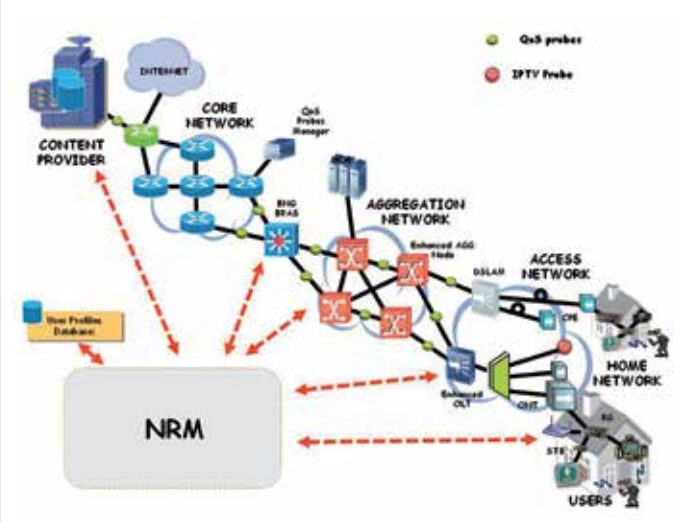

R2D2 layout

The basic goal of the R2D2 Networks project was to develop networks which are at the same time user-dependant, content-aware, and self-adaptive. These aims required the implementation of a specific module, the so called NRM (Network Resource Manager). The NRM is able to gather information from users and network elements, assess if the QoE requirements for users and contents are being satisfied, detect where in the network there could be impairments, and adopt corrective measurements.

Apart from the NRM, R2D2 envisaged an innovative network layout, as illustrated in the figure. This involved the deployment of a number of enhanced Network Elements (NE) throughout the network segments (core, aggregation, access and home) where the contents have to cross from the content provider to the user. These enhanced elements usually double as a source of feedback to the NRM and as the entry points for any corrective actions.

User dependence

In order to allow R2D2 to be user-dependant, the NRM element is endowed with an interface that enables the reception of QoE values as perceived by the users. The QoE perception is sent either by using the remote control of an R2D2 enhanced set-top box or by means of iPhone/iPad dedicated applications.

The NRM also includes a database storing a unified user profile that describes the acceptable QoE values for the user and the resource levels it has subscribed (SLA), so that corrective actions can be taken on a user-per-user basis.

Content awareness

Some of the NE deployed in the network, typically the QoE and QoS probes, are able to perform QoE measurements over the flows being transmitted and can even trigger alarms to inform the NRM that QoE could be dropping below acceptable levels.

The NRM, on its side, can gather QoS measurements and evaluate the user's perceived QoE in order to compare it to the values stored in the unified user profile, thus triggering the required corrective actions.

Self-adaptation

The NRM does not need to rely on user feedback in order to trigger corrective actions. In a normal case, the network elements are able to trigger alarms that will launch NRM's corrective capabilities. The complete procedure is not limited to checking the network element which requested an intervention, but involves all elements in the content path in order to obtain precise QoS measurements, estimate the perceived QoE, if needed, and decide where specifically in the network corrective actions are to be adopted.

Conclusion

The benefit from the approach taken by R2D2 is threefold. First, by being user aware, it enables the user to request quality improvements automatically without making a phone call to the usually slow technical support hot lines, as well as directly targeting any solution to users' profiles. Secondly, by being content-aware, the system is able to apply the best actions for the quality requirements of the transmitted contents. Thirdly, the fact of being self-adaptive means that the system actively strives at solving any problem before the user notices it, thus improving user satisfaction.

Finally, the results of R2D2 Networks were tested and demonstrated over GPON and xDSL accesses, meaning that the entire concept is valid for diverse types of access and could be easily extended to new ones.

Further information is available at www.celtic-initiative.org/Projects/Celtic-projects/Call6/R2D2NETWORKS/r2d2networks-default.asp

Impressions from the Celtic-Plus Event in Stockholm, February 2012.

IMPRINT

Editor-in-Chief:
Heinz Brüggemann
brueggemann@celticplus.eu

Contact:
Celtic Office
c/o Eurescom GmbH
Wieblinger Weg 19
69123 Heidelberg, Germany
Tel: +49 6221 989 405
Fax: +49 6221 989 451

About Celtic-Plus

Celtic-Plus is an industry-driven European research initiative to define, perform and finance through public and private funding common research projects in the area of telecommunications, new media, future Internet, and applications & services focusing on a new "Smart Connected World" paradigm. Celtic-Plus is a EUREKA ICT cluster and belongs to the inter-governmental EUREKA network. Celtic-Plus is open to any type of company covering the Celtic-Plus research areas, large industry as well as small companies or universities and research organisations. Even companies outside the EUREKA countries may get some possibilities to join a Celtic-Plus project under certain conditions.

Celtic-Plus Smart Connected World

EUREKA

Flow processing for an innovative Internet

The CHANGE platform

Adam Kapovits
Eurescom
kapovits@eurescom.eu

A central problem of today's Internet is that its size and scope make innovation through the introduction of new core network technologies very difficult. Even minor changes only happen through the accretion of point solutions that embed knowledge in the network, optimizing today's applications at the expense of tomorrow's.

The EU FP7 project CHANGE set out to offer a solution to that problem by introducing a common concept of a flow-processing platform (i.e. Software Defined Networking) instantiated at critical points in the network. Although the platform and its interfaces are common, the processing performed must be programmable, allowing the network to evolve and support the needs of rapidly-changing applications. Such platforms can be built from commodity hardware, for example x86 servers and commodity switching chipsets, and are both scalable and powerful while retaining the flexibility to quickly introduce processing primitives.

Software Defined Networking and innovation

Support to a wide range of flow processing primitives and easy deployment are key to support speedy innovation and avoid to be locked into the applications of today. Essentially this means that anything more than simple packet forwarding should be a software function, allowing quick deployment. In this vision the flow processing platforms play an important role enabling the Internet to reason about flows and to enhance the processing that flows receive in a manner that enables innovation rather than stifles it.

Contrast this with the current Internet, where flow processing is almost always performed in special-purpose boxes sold by vendors to solve a specific problem.

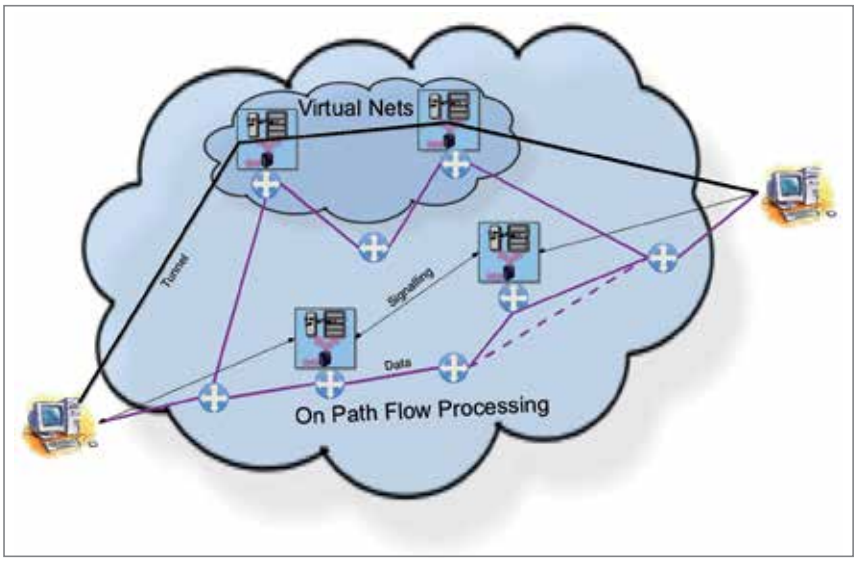

Figure: Internet level flow processing

Deployment scenarios

There are two conceptually different deployment scenarios for such flow processing platforms, which, however, are most likely to be used in combination in reality (see figure):

1. Programmable virtual networks: In this scenario flows can be classified and processed in one flow processing platform, then sent to another flow processing platform for further processing, and so on. Typically this will be done by tunnelling traffic across the Internet. Such flow-level virtual networks give operators and application writers great flexibility in controlling how their traffic is forwarded and where it is processed. This is in stark contrast to the current Internet architecture, where networks process traffic based only on destination address prefixes.

2. On-path flow processing: In this other scenario traffic traversing the network using conventional IP forwarding needs to be processed as a flow at certain points in the network; examples are firewalls and traffic shapers. The kind of flow processing platforms we envisage can provide this conventional functionality, but the real benefits come when such flow processing platforms can communicate. This allows applications to express their requirements, networks to express their constraints, and a much more flexible approach to enabling access control restrictions to be taken.

Potential application areas

Potential application areas and scenarios of such SDN platforms include:

- Virtualized ISP
- Dynamic network troubleshooting and re-locatable maintenance
- Targeted, on-demand network monitoring,
- Dynamic intra and inter-network traffic shaping
- Workable, targeted QoS
- Shippable attack mitigation.

Conclusion

FP7 project CHANGE is working on a solution to help dramatically reduce network costs by adopting flow definition and processing as primitives, leading to the goal of Network (processing) as a Service (NaaS), and will devise an architecture for innovation, performing network processing on commodity hardware to speed up the evolution of the Internet.

Further information on EU FP7 project CHANGE is available at www.change-project.eu

IRMOS

The Network and the Cloud

Karsten Oberle
Alcatel-Lucent Bell Labs
Germany
Karsten.Oberle@alcatel-lucent.com

Manuel Stein
Alcatel-Lucent Bell Labs
Germany
Manuel.Stein@alcatel-lucent.com

The objective of IRMOS (Interactive Real-time Multimedia Applications on Service Oriented Infrastructures) has been to enable, deploy and execute real-time applications over a potentially distributed Cloud infrastructure, where processing, storage and networking need to be combined and delivered with guaranteed levels of service quality.

While IRMOS worked out a complete Cloud solution covering aspects of Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS) and Infrastructure-as-a-Service (IaaS), this article focuses on networking aspects of IaaS.

The rationale for network virtualisation

Current Cloud offers providing IaaS are not adequate to meet the requirements of demanding services such as interactive real-time services. For those services the network is a key element to provide connectivity between components of distributed, interactive real-time applications with the required level of end-to-end Quality of Service (QoS) as expressed in Service Level Agreements (SLAs). The network QoS is mainly defined by parameters such as bandwidth, jitter or delay.

In order to serve those requirements within a data center (DC), especially on wide area networks (WAN) connecting multiple data centers in different locations, the concept of a fully managed network virtualisation framework has been developed.

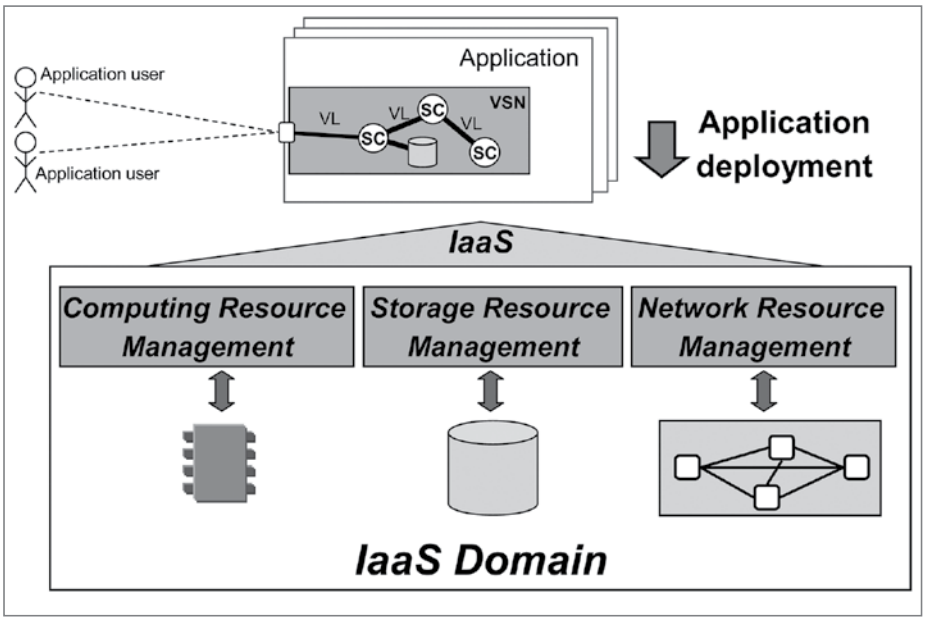

Figure 1: Infrastructure as a Service and Virtual Service Network

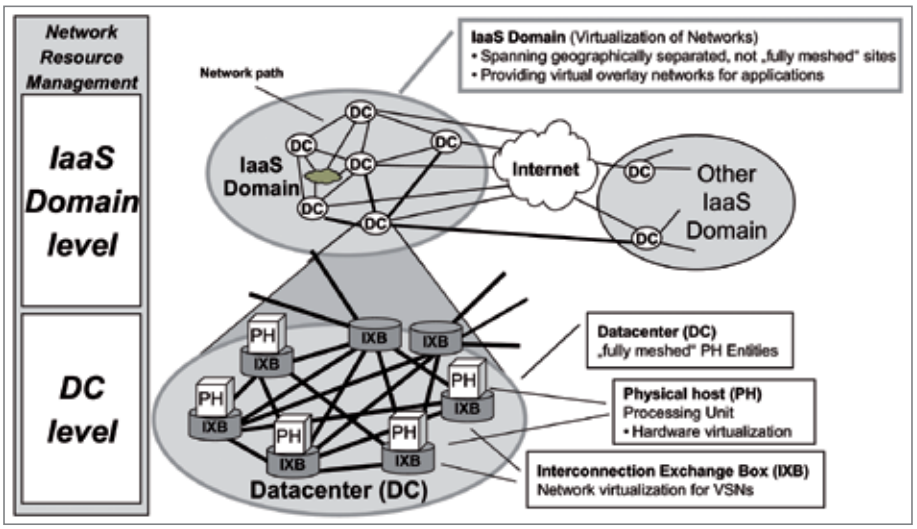

Figure 2: IaaS domain hierarchy view

Virtual Service Network

The individual application requirements need to be specified and then satisfied "on demand" by the platform. The model describing a distributed application including specifications of its required resources is called a Virtual Service Network (VSN). A sample VSN is shown in Figure 1, where the model represents resource parameters such as network, storage or computing through Virtual Links (VL) between Service Components (SC).

The performance parameter and resource requirements of the application are described as meta-information in the VSN description (VSND). It is part of a SLA between a PaaS provider and an IaaS domain of a certain IaaS provider and is valid for the lifetime of the deployed VSN.

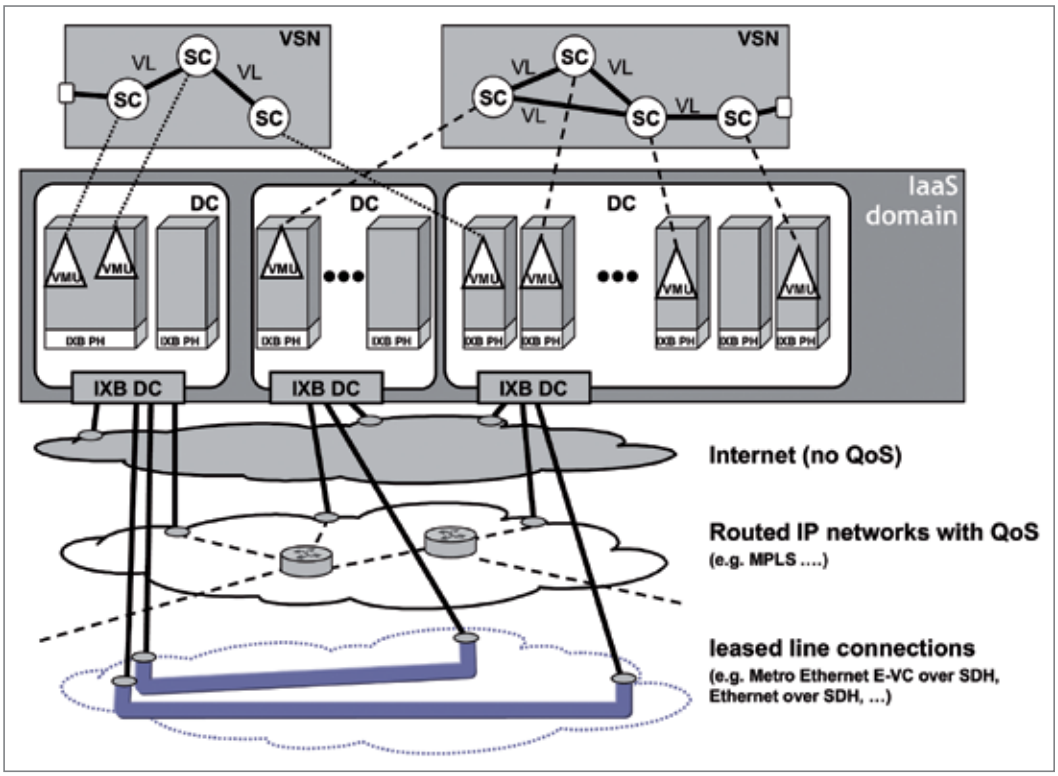

Figure 3: Deployment of services on the IaaS

The network management

As indicated in Figure 2, an IaaS domain represents an IaaS provider realm consisting of one or more data centers (DC). A DC itself consists of physical hosts (PH), storage capability and a high-performing network infrastructure.

One of the project's main research tasks targeted the management of network resources within an IaaS domain while providing QoS guarantees on a shared transport medium for multiple concurrent customer applications simultaneously – even for those having real-time requirements. The developed IRMOS network resource management conciliates all network-related application requirements and prevents network resource allocations from interfering with each other. The resulting VSN deployments represent applications executing in full isolation from each other. The network resources including already assigned resources are continuously monitored so that corrective measures are applied as required, e.g. policing or rerouting VL traffic flows by reconfiguring the network virtualization infrastructure in order to guarantee the required level of QoS.

The introduced two-level hierarchical organization (domain-level and DC-level) of network resource management alleviates the management complexity by letting each DC autonomously keep the individual QoS guarantees for deployed VLs of network-based applications.

The network resource usage is supervised and limited where required, concerning both network resources between and within DCs, to ensure that applications keep performing fully isolated from each other.

The VL management itself is performed by the IRMOS network virtualization (NV) functionality, i.e., the ISONI eXchange Box (IXB) as depicted in Figure 3. The NV is located in each PH, in the IXB PH, and on the DC border (the IXB DC).

At execution time all SC instances run within virtual machine units (VMU); their message traffic is encapsulated and forwarded by the IXB according to the interconnection model. In order to select the optimal network route and type (e.g. best effort Internet, routed IP networks or leased lines) as indicated in Figure 3, all VL need to be modelled with the corresponding network resource attributes (bandwidth, delay, jitter, etc.).

Conclusion

The presented overall approach is an essential part of a complete IaaS solution with guaranteed QoS on network level for multiple concurrently deployed interactive real-time applications with different QoS requirements.

Further information is available at www.irmosproject.eu
The relevant network virtualization deliverables can be downloaded at the project website (WP7 deliverables).

Innovation and research programmes

David Kennedy
Director of Eurescom
kennedy@eurescom.eu

Many research programmes are crying out for innovation – but do they know what they are asking for? If you ask people in the ICT sector about innovation you get many different answers – the most common ones seem to accept innovation as new products or services. Some say innovation is the process of finding new products or services, and others say it is the novelty value, the bringing of changes or even a form of competitive advantage. So which is it?

According to the dictionary

The dictionary defines innovation as the noun derived from the verb to innovate. And to innovate is defined as "to make changes in something established, especially by introducing new methods, ideas, or products".

Now this definition gives us something to work with. If we are expected to innovate, we are expected to make changes in something established. The original concept for a new product is that it does something better than the existing products – the standard business management exercise to build a better mouse trap comes to mind here.

When we apply this to the technology and services in the wider ICT context, we see that there are many opportunities for rich and exciting innovation. However, if we are creating something new, are we innovating or are we inventing? You could logically argue that one of the new Apps populating the smart phones around the world is innovation, using the programme environments available, or it is invention, as it does something that was not done before. I go for innovation here, as I see it as using the emerging capabilities of the clever devices and the modern infrastructure in combination to do something different.

Now, if we are expected to justify future research projects in terms of innovation, we need to determine what new aspects do we bring into the discussion each time we launch a new project.

The FP7 EARTH project has just been honoured for finding new ways to dramatically reduce the power consumption of 4G/LTE base station equipment. This is clearly innovation, under our current definition, as it is doing something much better than it was previously done.

"Your proposal is really innovative, but I'm afraid we can't accept it, because our evaluators don't like it."

Innovation is not new

In 1903 in his book "The laws of imitation", Gabriel Tarde defined the innovation-decision process as a series of steps that include:

- First knowledge
- Forming an attitude
- A decision to adopt or reject
- Implementation and use
- Confirmation of the decision.

The aspect that is becoming more and more emphasised in our business oriented world is the "implementation and use".

We can quickly agree that innovation is the creation of better or more effective products, processes, services, technologies. But the interesting metric is seeing which ideas are accepted and adopted by markets, governments, and society as new ways of doing things. This places more emphasis on the definition that innovation refers to the use of a new idea or method.

If we define adoption and use as part of the innovation cycle, then our use of innovation in research programmes must encompass finding better ways of doing things to demonstrate and stimulate the uptake of these "better" ways – and this is why I needed to stress the wider definition of innovation.

Stimulating innovation needs innovation

I foresee an impending contradiction in European research programmes where we seek to have real "innovation". We want more innovative proposals for research projects, and we expect innovation in the uptake and use of the "innovations".

However, the contradiction will emerge when we try to stimulate all this compound innovation, but we use outdated and hopelessly flawed methods for identifying the projects that will achieve the innovations required. From Einstein we should have learned that you cannot keep doing the same thing and expect different results.

In the most recent call of FP7 I can cite numerous examples where the evaluation process was not able to identify the difference between innovation and invention. Many innovative proposals were rejected, as the innovation, which involved enhancing and exploiting known technologies in novel ways to achieve significant gains, was not recognised in favour of projects with a pure research focus. The amount of innovation needed to get a technology from its emerging state to being actually usable is huge and usually underestimated.

I now call on the powers that be to reflect some real innovation, in terms of improvements in the processes. It must be recognised that real innovation does actually mean reinventing the wheel, it means improving it so it can go faster, quieter, more efficiently, whatever. It is time we had real innovation in the project selection process so that we can select the right projects and not just the best proposals. Look at the potential impact.

Horizon 2020 is promoted as a significant investment to keep European industry competitive. Being innovative is a competitive advantage. However, no substantial innovation will be stimulated unless the processes change. Can we change?

Visions for Smart Cities

Future Internet Assembly in Aalborg

Milon Gupta
Eurescom
gupta@eurescom.eu

The Future Internet Assembly 2012 was held on 10 and 11 May in the city of Aalborg under the Danish EU Presidency as part of the Future Internet Week. About 500 researchers and EU officials came to North Jutland to talk about "Smart Cities and Internet of Things", the theme of the event. For those who had visited the previous Future Internet Assembly (FIA) events, the conference did not offer any surprises or new insights – it was business as usual. That said, the event offered an excellent opportunity for social networking across research areas and projects.

Opening session

The opening session, chaired by conference organiser Neeli Prasad from Aalborg University, was dedicated to Smart Cities. Presenters of the first session included Thomas Kastrup-Larsen, Alderman of Health and Sustainable Development, City of Aalborg, Finn Kjærsdam, President of Aalborg University, Ramjee Prasad, Director of Center for Teleinfrastruktur, Mario Campolargo, Director of DG INFSO Emerging Technologies and Infrastructures, and Morten Østergaard, Danish Minister for Science, Innovation and Higher Education. This was followed by a keynote speech by Thierry Van Landegem, Vice President at Alcatel-Lucent, on ICT infrastructure as a key enabler of Smart Cities. The next keynote by Reinhard Scholl, Deputy Director of ITU, focused on collaboration for Smart Standards as enablers of Smart Cities and the Internet of Things.

A Round Table, chaired by Reinhard Scholl, expanded on various aspects of Smart Cities. Panelists included Thomas Kastrup-Larsen from the City of Aalborg, Asta Fog Larsen, a citizen of Aalborg, Jose M. Hernandez-Munoz from Telefonica, Rolf Hapel, Director Citizen Services and Libraries from the Danish city of Aarhus, Inigo de la Serna, Mayor of the Spanish city of Santander, and Markkula Markku, Advisor at Aalto University in Finland. The panel discussion was somehow very predictable and provided few new insights to the Smart Cities discussion, so the ensuing signing of the Smart City Manifesto was the highlight of the morning.

Keynote speech by Thierry Van Landegem from Alcatel-Lucent at the opening session of the Future Internet Assembly

Hands-on FIRE demo

Those who preferred to get their hands, or at least their eyes, on real research results related to Future Internet Research and Experimentation (FIRE), had the opportunity to visit the Hands-on FIRE demonstration, which was organised by the FIRE STATION project. The demonstrations included, for example, topics like Integrating Smart & Green Buildings with the Web (HOBNET project), Optical network control through OpenFlow (OFELIA project), OpenFlow and P2P integrated testing (OpenLab project), and Federation of Future Internet platforms (NOVI project), to name just a few. The demonstrations showed a high degree of substantial achievement by these projects.

Sessions

Thursday afternoon and Friday morning were packed with parallel sessions. The topics focused mainly on Smart Cities or the Internet of Things. Other session topics included Clouds, open platforms for innovation, standardisation, and online gaming. At one of those sessions, Dave Carter from the Manchester Digital Lab (MadLab) demanded infinite bandwidth and zero latency. This looks like a real challenge for European researchers and the underlying business models. At the venue, the Aalborg conference centre, the participants could only dream of infinite bandwidth, as the WLAN provided suffered from significant load. The contrast between the vision of perfect connectivity and the harsh reality of shared limited network resources is nowhere to be felt more than at large Future Internet events, like the Future Internet Assembly.

Hands-on FIRE demonstration

Outlook

In contrast to previous years, the FIA in Aalborg will be the only official FIA event in 2012. The FIA community had expressed concerns about the frequency of events, so the organisers have taken steps to transform the FIA into an annual event. This could help to make future FIA events more stimulating, as there will be more substantial progress on many issues between events.

Further information about the Future Internet Assembly in Aalborg is available at http://fi-aalborg.eu and at http://www.future-internet.eu

If you wish to learn more about what is going on in European Future Internet research, a good place to start is the new FIA book:

The Future Internet. Future Internet Assembly 2012: From Promises to Reality. Lecture Notes in Computer Science, Vol. 7281. Springer 2012. ISBN 978-3-642-30240-4

EU-Japan Workshop on Experimental Future Internet Research in Aalborg

On 9 May 2012, an EU-Japan Workshop on Experimental Future Internet Research at the Future Internet Week was held at the Future Internet Week in Aalborg, Denmark. The workshop attracted about 30 researchers from Japan and Europe.

It was co-organised by FP7 project FIRE STATION (http://www.ict-fire.eu/home/firestation.html) and the Japanese research organisation NICT (http://www.nict.go.jp).

The workshop focused on 'OpenFlow and its applications' as well as 'Global Cloud infrastructure'. The experts from the EU and Japan discussed collaborative activities in the Future Internet area with a focus on experimental R&D over federated testbeds in Japan and the EU. This continued the dialogue from the 4th EU-Japan Symposium on the "New-Generation Network and the Future Internet", held in Tokyo, Japan, on 19 January 2012.

The first session on 'OpenFlow and its applications' was chaired by Nozomu Nishinaga from NICT. The session featured on the Japanese side presentations on the RISE testbed, Virtual Network Services, and Open Source software activities for deploying OpenFlow. The European presentations covered the OFELIA testbed, the NOVI project, and OpenFlow and P2P integrated testing within the OpenLab project.

The second session on 'Global Cloud infrastructure' was chaired by Anastasius Gavras from Eurescom. The Japanese presentations in this session included the StarBED project on the emulation of disaster scenarios, the WIDE Cloud project and the challenges for large distributed Clouds as well as the AIST Inter-Cloud Framework. On the European side, presentations featured the BonFIRE multi-site Cloud prototype facility, experimental challenges for Cloud security, and the federation of the wireless facilities of OpenLab with wired networks and the Cloud.

Japanese and European researchers at the EU-Japan workshop in Aalborg.

The workshop has been of particular interest for both Japanese and European researchers in view of the envisaged coordinated EU-Japan call for research projects in this area. The call will be opened and closed in the second half of 2012. The projects are envisaged to start in April 2013, with a project duration of 36 months.

Further information, including the presentation slides, is available at www.ict-fire.eu/events/eu-japan-workshop-aalborg.html

News in brief

Computer skills in the European Union

There is still a wide gap regarding individual computer skills and usage in the European Union, according to data by Eurostat, the statistical office of the European Union. In the EU27, the share of individuals aged 16–74 who have used a PC in 2011 varies between 50% in Romania and 96% in Sweden. The group of countries at the top in regard to PC usage also includes Denmark, Luxembourg and the Netherlands (all 94%). At the bottom Bulgaria (55%) and Greece (59%) are only slightly ahead of Romania.

In addition to geographical differences in computer usage and skills, there are also significant differences related to age. In 2011, more than 75% of those aged 16-74 in the EU27 had used a computer, while this share was 96% amongst those aged 16-24. The share of individuals in the EU27 having written a computer programme was 10% amongst those aged 16-74 and 20% amongst the younger age group.

http://europa.eu/rapid/pressReleasesAction.do?reference=STAT/12/47&format=HTML&aged=0&language=EN&guiLanguage=en

New SIM card format for slimmer phones

ETSI has standardized a new form factor (4FF) for the SIM card, 40% smaller than the smallest current design.

At its 55th meeting held on 31 May and 1 June 2012 in Osaka, Japan, ETSI's Smart Card Platform Technical Committee agreed the new fourth form factor (4FF) for the Universal Integrated Circuit Card (UICC), better known as the SIM card.

Today's SIM card designs take up a significant amount of space inside a mobile device. This space has become more and more valuable in handsets, as they deliver an ever increasing number of features.

The 4FF card will be 12.3 mm wide by 8.8 mm high and 0.67 mm thick. According to ETSI, the 4FF card can be packaged and distributed in a way that is backward compatible with existing SIM card designs. The new design will offer the same functionality as all current SIM cards.

The SIM is the most successful smart card application ever. A SIM card is used to securely associate a mobile device with a customer account, preventing fraud and ensuring that calls are correctly routed to customers. It is an essential security feature of mobile networks and is integrated into every GSM, UMTS and LTE device. Over 25 billion SIM cards and derivatives have been produced so far, and the industry continues to issue over 4.5 billion SIM cards each year.

The new form factor was adopted by industry with the involvement of major mobile network operators, smart card suppliers and mobile device manufacturers. German ICT news portal heise.de reported that a fight between Apple on one side and Nokia, Motorola and RIM on the other side had preceded the agreement, which seems to have turned in favour of Apple. The Apple competitors had proposed a solution that was more oriented towards the shape of a MicroSD card. The decision had originally been planned for an ETSI meeting in March, but was postponed due to the controversial debate.

The new design will be published in due course in ETSI's TS 102 221 specification, freely available from the ETSI website.

www.etsi.org/WebSite/Newsandevents/2012_06_New_SIM_Card_Format.aspx

Printable on-demand robots

A new research project in the United States aims to reinvent how robots are designed and produced. The project started in Spring 2012 and is funded by a 10 million dollar grant from the National Science Foundation (NSF). The project led by the Massachusetts Institute of Technology (MIT) aims to develop a desktop technology that would make it possible for the average person to design, customize and print a specialized robot in a matter of hours.

The five-year project, called "An Expedition in Computing for Compiling Printable Programmable Machines," brings together a team of researchers from MIT, the University of Pennsylvania and Harvard University, and is funded as part of the NSF's "Expeditions in Computing" programme. The project leader is professor Daniela Rus from the MIT Computer Science and Artificial Intelligence Laboratory (CSAI).

Professor Rus and her team plan to create a platform that would allow an individual to identify a household problem that needs assistance and then go to a local printing store to select a blueprint from a library of robotic designs. The next design step is to customise an easy-to-use robotic device that could solve the problem. Within 24 hours, the robot would be printed, assembled, fully programmed and ready for action.

Today, it takes years to produce, program and design a functioning robot. The process is expensive; it involves hardware and software design, machine learning and vision as well as advanced programming techniques. The envisaged solution would automate the process of producing functional 3-D devices and allow individuals to design and build functional robots from materials as easily accessible as a sheet of paper.

This project aims to dramatically reduce the development time for a variety of useful robots, opening the doors to potential applications in manufacturing, education, personalized health care and even disaster relief.

Currently, project researchers are exploring several research areas: developing an application programming interface for simple function specification and design; writing algorithms that would allow for control of the assembly of a device and its operations; creating an easy-to-use programming language environment; and designing new, programmable materials that would allow for automatic fabrication of robots.

So far, the research team has prototyped two machines for designing, printing and programming, including an insect-like robot that could be used for exploring a contaminated area and a gripper that could be used by people with limited mobility.

The concept of printable robots is not new. In 2011, researchers from the Robotics Lab at Carlos III University of Madrid presented their Mini-Skybot, a mobile robot aimed for educational purposes. It is 3D-printable, fully Open-Source and designed exclusively with Open-Source tools.

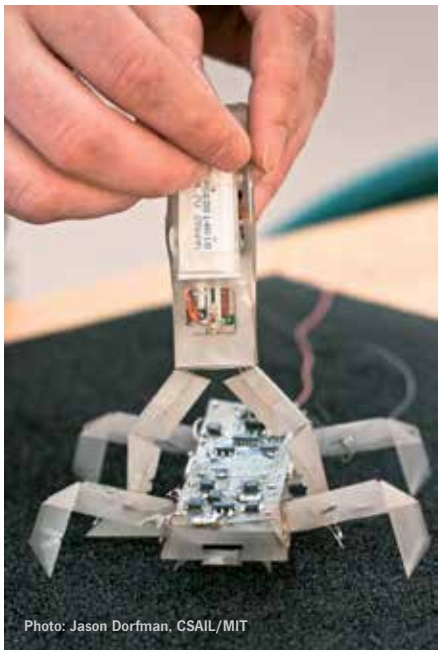

Printable Origami gripper

Project website: http://ppm.csail.mit.edu
On MiniSkybot: http://www.iearobotics.com/wiki/index.php?title=Mini-Skybot_v1.0

Restructuring at the EC: DG INFSO becomes DG CONNECT

As of 1 July 2012, the European Commission's Directorate-General Information Society and Media, short: DG INFSO, will have a new structure and change its name to Directorate General for Communications Networks, Content and Technology; short: DG CONNECT.

According to the DG's director, Robert Madelin, the new name better represents the range of topics the DG is active in. The rebranding is part of the DG's restructuring. According to Mr Madelin, the new structure "better aligns the work of the DG with key EU policies for the coming decade: ensuring that digital technologies can help deliver the growth which the EU needs".

Neelie Kroes, Vice President of the European Commission, responsible for the Digital Agenda, explained the motivation for the new structure with the need for adapting to the challenges of the next ten years. She expects that the restructuring will make the DG more flexible, less fragmented, and better capable of linking policy and research areas.

http://ec.europa.eu/dgs/information_society/connect_en.htm

Upcoming review of International Telecommunication Regulations

On 3–14 December 2012, the first World Conference on International Telecommunications (WCIT-12) will take place in Dubai, United Arab Emirates. The conference is organised by the International Telecommunication Union (ITU), the United Nations specialized agency for information and communication technologies. Conference participants will review the International Telecommunication Regulations, short: ITRs, which can be considered the basis of today's connected world. ITU intends to build on the success of the existing ITR treaty and pave the way for realising the vision of a fully inclusive information and networked society over the next decade.

Background on ITRs

The ITRs set the rules for global communications in 178 countries. The ITRs were agreed in 1988 at the World Administrative Telegraph and Telephone Conference in Melbourne, Australia, and came into force in 1990.

Treaty-level provisions are required for worldwide networks and services. The ITRs set out principles for ensuring that networks can connect with each other smoothly, and that international services will be offered in a fair and efficient manner.

They comprise ten articles dealing with such matters as cooperation among national administrations, giving priority to emergency telecommunications, and how to calculate the charges for traffic exchanged between carriers in different countries. The ITRs laid the foundation for privatisation, competition and deregulation that created the conditions for growth in ICTs, including the Internet, we see today.

Challenges for WCIT-12

Since 1988, the telecoms landscape has dramatically changed. Much of the sector has been privatised and liberalised. There is increasing use of networks and applications based on the Internet Protocol (IP), which leads to the growing importance of cyber-security. Technological convergence has blurred the distinction between voice and data traffic, and even inexpensive mobile phones are now in fact sophisticated computers. Data volumes are growing much quicker than the infrastructure needed to carry them.

Companies like Google and Facebook offer popular web-based services that use increasing amounts of network capacity without necessarily generating larger revenues for the companies that provide the infrastructure. This leads to the challenging question how the new infrastructure is to be expanded to cope with demand, and who should pay for its expansion and its use.

According to ITU, there is consensus that the ITRs must be adapted to match the rapid development of the ICT sector. Differing proposals have been put forward on how best to do this, but all agree that there must be international cooperation. Governments and the private sector will play complementary but distinct roles, i.e. governments establish regulatory frameworks, and the private sector provides the investment.

The preparatory process for the World Conference on International Telecommunications is led by a working group of ITU's Council that is open to ITU's 193 Member States and the more than 550 Members of ITU's three Sectors – Radiocommunication, Telecommunication Standardisation, and Telecommunication Development. These Sector Members – private companies, scientific or industry organizations, financial and development institutions, and other organizations dealing with telecommunication matters – include some of the world's largest ICT firms.

In addition, input has come from the more than 200 associates and academic institutions involved in ITU's activities, as well as many other entities that have participated in the series of information sessions and regional preparatory meetings held around the globe in 2012. Further information sessions are planned this year.

Proposed changes or additions to the ITRs cover a wide range of topical areas, including:

- Human right of access to communications
- Security in the use of ICTs
- Protection of critical national resources
- International frameworks
- Charging and accounting, including taxation
- Interconnection and interoperability
- Quality of service
- Convergence.

Contribution by ETNO

One of the major contributors in the process is ETNO, the voice of Europe's leading e-communications services and network providers. ETNO has submitted to the ITU proposals for changes of the ITRs providing for a new sustainable model for the Internet based on commercial agreements between undertakings. The ETNO contribution proposes a number of changes to selected articles of the ITRs, namely those dealing with international network (article 3), international telecommunications services (article 4) and, relatedly, to the definitions (article 2).

ETNO calls for a new IP interconnection ecosystem that provides end-to-end Quality of Service delivery, in addition to best effort delivery, enabling the provision of value-added network services, to both end-customers and over the top (OTT) players and content providers, and a reflection of the value of traffic delivery over network infrastructures.

Moreover, the contribution states that in order to ensure an adequate return on investment in high bandwidth infrastructures, operating agencies shall negotiate commercial agreements to achieve a sustainable system of fair compensation for telecommunications services.

It remains to be seen to what extent ETNO and the telecoms network operators will be able to convince delegates at WCIT-12 against the strong pressure of big Internet service providers like Google.

www.itu.int/en/wcit-12

The invisible driver

Autonomous cars are driving up fast

Milon Gupta
Eurescom
gupta@eurescom.eu

Sooner than you think, human drivers will be a phenomenon of the past. Within the next 10 years, the first mass produced autonomous cars will hit the road. It seems that the evolution of automobiles is entering a new phase – 240 years after the first, steam-driven car by Nicolas-Joseph Cugnot and about 130 years after Carl Benz invented today's automobile powered by a combustion engine.

Origins of self-moving vehicles

Autonomous cars will give the word automobile a new meaning, which is actually going back to the origin of the word. It is derived from the ancient Greek word "autós", meaning "self", and the Latin "mobilis", "movable"; so automobile literally means a vehicle that moves itself. The idea for an autonomous car goes back to the 1939 World's Fair, when Norman Bel Geddes presented his Futurama exhibit, which depicted electric cars powered by circuits embedded in the roadway and controlled by radio.

The first serious attempts for building self-driven vehicles started in the 1980s in Europe and the US. In Germany, Ernst Dickmanns and his team at Bundeswehr University Munich designed a vision-guided Mercedes-Benz robotic van, while in the US the DARPA-funded Autonomous Land Vehicle (ALV) achieved the first road-following demonstration that used laser radar.

First license for a self-driven car

After years of continuous development, it seems that in 2012 autonomous cars have reached a level of maturity which makes the market appearance of self-driven vehicles only a matter of years.

In May 2012, the state of Nevada issued the first-ever license to an autonomous car, a Toyota Prius modified with Google's experimental driverless technology. In June 2011 the state of Nevada had passed the first law worldwide allowing the operation of driverless cars on public roads. Florida, Hawaii, Oklahoma, and California are now also considering the legalisation of autonomous cars.

European developments

In Europe, major car manufacturers are planning to launch semi-autonomous and fully autonomous cars in the near future.

In 2011, a consortium including Volvo has begun to develop an almost-autonomous 'road train' system for highways in an EU-funded project called SARTRE (Safe Road Trains for the Environment – www.sartre-project.eu). It is expected that the 'road train' system could be integrated in cars by 2020.

Volkswagen is currently testing a "Temporary Auto Pilot" (TAP) system that will allow a car to drive itself at speeds of up to 130 km/h on the highway. In 2011, Mercedes-Benz announced its 2013 S-Class will feature an autonomous driving system for speeds of up to 40 km/h. In 2012, Audi announced plans to introduce a new autonomous driving system at speeds of up to 60 km/h, which is called 'Traffic Jam Assistant'. The system is expected to be built into the 2013 Audi A8.

Car makers in the US and China are pursuing similar plans for semi-autonomous and fully autonomous cars. In 2011, Alan Taub, GM's vice president of global R&D, stated that the company is planning to release semi-autonomous cars by 2015, and fully autonomous cars by 2020.

Benefits

Proponents of autonomous cars claim a plethora of economic, environmental, and societal benefits. In regard to safety, they expect fewer crashes, due to the autonomous system's increased reliability compared to human drivers. The economic benefits include increased roadway capacity and reduce traffic congestion, due to the reduced need of safety gaps and the optimisation of travel routes. Furthermore it is suggested that self-driving cars would alleviate parking scarcity, as cars could drop off passengers, park far away where space is not scarce, and return as needed to pick up passengers. In regard to the environment, experts expect reduced oil consumption and air pollution, due to better traffic flow and the removal of then redundant safety features.

Google's driverless car on a test drive.

Road train coming down the hill.

The promise to users is that autonomous cars would relief them of driving and navigation chores and remove any constraints on their state. This means, occupants of self-driven cars could be infants unable to reach the steering wheel as well as elderly citizens with poor eyesight. Drunken people could 'drive' an autonomous car now without being at risk of causing an accident and being caught by the police.

Acceptance

With all these benefits, you should expect everyone to be enthusiastic. Recent surveys show that this may only be true for some people. According to an online survey of 2,006 consumers in the US and the UK conducted by Accenture in 2011, 49 percent of all respondents said they would be comfortable using a driverless car.

According to a survey done by J.D. Power and Associates with 17,400 vehicle owners in 2012, more than a third (37 percent) of all survey responders initially said they would be interested in purchasing a fully self-driving car. The number of willing car buyers dropped to 20 percent once they learned the technology would cost an additional 3,000 US dollars. At the additional cost of 3,000 dollars, 25 percents of male respondents were willing to pay for a fully autonomous vehicle, while only 14 percent of women wanted the feature.

Outlook

It appears that autonomous cars are unstoppable. Fast technological progress and the need for more intelligent transport solutions on overcrowded roads will drive their development and adoption in the coming decade. Like the disappearance of the chauffeurs in the early 20th century, this will have significant economic and societal impacts; take for example the possible disappearance of professional drivers. Considering these impacts rather sooner than later is advisable, in order to make the transition easier.

EuresTools

Effective tools for successful FP7 projects

Reporting & Controlling | **Information & Knowledge** | **Communication & Interaction**

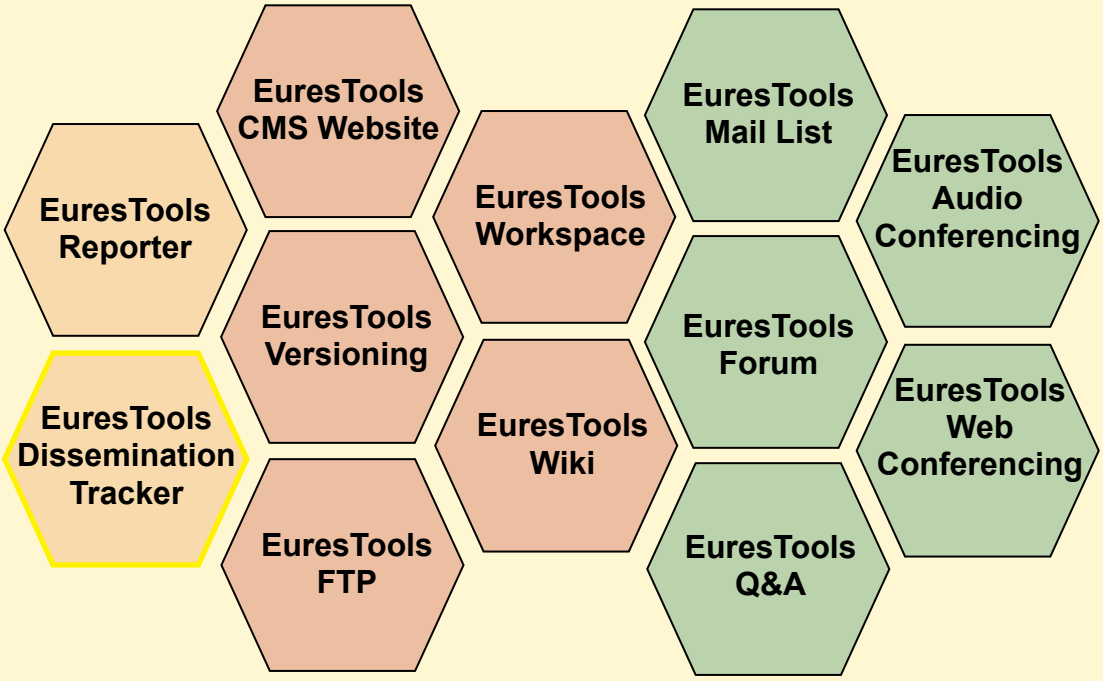

Manage dissemination activities effectively – with EuresTools® Dissemination Tracker

EuresTools Dissemination Tracker makes it easy to manage all of your project's dissemination activities, from approving a draft dissemination activity to summarising all dissemination activities in the Final Project Report.

Main benefits of EuresTools® Dissemination Tracker:

- **Easy entry** of all dissemination items by authors
- Authors can **keep track** of the status (e.g. accepted, rejected, published)
- **Effective approval** procedure for papers and other dissemination items
- Project managers can **easily retrieve** all required tables for internal reports or reports to the Commission, like periodic or Final Reports
- **Special reports**, like the annual questionnaire by DG INFSO, can be filled-in directly via the Tracker

The Tracker's reporting facility is very flexible and can respond to any changes of the European Commission's reporting requirements.

Main benefits of EuresTools®

- **Flexible:** EuresTools® provide full flexibility through modular, Cloud-based tools – just use the tools you need.
- **Comprehensive:** the EuresTools® modules cover all needs of a research project. The tool suite is continuously updated reflecting the needs of customers.
- **Intuitive:** the user interfaces of all modules are very intuitive, making it easy to enter and access data.

Free offer for the proposal phase

Eurescom offers all FP7 project proposals involving Eurescom Members, and selected other proposals, the basic EuresTools modules for free in the proposal phase. This offer is contingent on you using EuresTools in your project, if your proposal succeeds. Projects involving Eurescom Members receive a 20 percent discount on the regular price.

Contact us to get more information on EuresTools –
e-mail: **services@eurescom.eu**

www.eurescom.eu/EuresTools

EURESCOM mess@ge

The magazine for telecom insiders

Get your free subscription of Eurescom mess@ge at www.eurescom.eu/message

EURESCOM

European Institute for Research
and Strategic Studies
in Telecommunications GmbH
Wieblinger Weg 19/4
69123 Heidelberg, Germany
Tel.: +49 6221 989-0
Fax: +49 6221 989 209
E-mail: info@eurescom.eu
Website: www.eurescom.eu

Innovation through Collaboration

Eurescom is the leading organisation for managing collaborative R&D in telecommunications. Our mission is to provide efficient management and support of R&D projects, programmes, and initiatives for our customers. We offer more than two decades of experience in managing large-scale, international R&D for major industry players, the European Commission, and EUREKA Cluster Celtic-Plus. What distinguishes Eurescom is the combination of a secure, reliable infrastructure for collaborative work, a large European network of experts, and internationally outstanding project management skills.

QR code to the online edition of Eurescom mess@ge