FUTURE INTERNET ASSEMBLY 2010

Valencia, Spain, 15 – 16 April 2010

CONFERENCE REPORT

Information Society and Media DG - Directorate for Converged Networks and Services

"The Internet People"

Contents

Glossary

EFII	European Future Internet Initiative
EIT	European Institute of Innovation & Technology
FCN	Future Content Networks (an FIA Working Group)
FIA	Future Internet Assembly
FIRE	Future Internet Research & Experimentation
FISE	Future Internet Socio-Economics (an FIA Working Group)
FISO	Future Internet Service Offer (an FIA Working Group)
ICT	Information and communications technology
IP	Internet Protocol
IoS	Internet of Services
IoT	Internet of Things
KIC	Knowledge & Innovation Community
MANA	Management & Service-aware Networking Architectures (an FIA Working Group)
PPP	Public-private partnership
RWI	Real-World Internet (an FIA Working Group)
T&I	Trust and Identity (an FIA Working Group)
TCP	Transmission Control Protocol

Definition

Future Internet is a widely used term describing research activities dedicated to further developing the original internet.

Further information:
European Commission – Directorate-General Information Society
E-mail: infso-future-internet@ec.europa.eu

European Commission portal on the Future Internet:
http://ec.europa.eu/information_society/activities/foi/index_en.htm

Future Internet Assembly website: www.future-internet.eu
FIA Valencia conference website: www.fi-valencia.eu

Report prepared by conference rapporteur: Edwin Colyer, Scientia Scripta (www.scientiascripta.co.uk), with input from FIA session moderators and caretakers and European Service Network (www.esn.eu).

Executive Summary

Hints of orange blossom scented the boulevards of Valencia. And though the morning air was cool and the sun hid behind a layer of haze, more than 500 delegates involved in Future Internet R&D and policy made their way to the fifth Future Internet Assembly (FIA), taking place at the Palau des Congressos in Valencia, Spain.

Arriving early, a coffee and churros in hand, they patch into the conference WiFi network to check their emails. Some notice an unusual news story: the eruption of the Eyjafjallajökull volcano in Iceland, a massive cloud of ash spreading east and south, and the closure of UK airspace. A few notice the irony of the spoof 'Internet Express' newspaper supplied in their conference packs, dated 11 May 2024: 'Plan-it-net eruption warning prompts evacuation in Pacific', says the headline.

But they are not worried, they are ready to enjoy two excellent days of presentations, updates, discussion and – perhaps most importantly – networking. In the auditoria and seminar rooms, around drinks and paella, they are ready to share their ideas about the evolution of the information society in Europe and beyond.

The importance of the Future Internet is reflected at almost every level of European policy. And its elevated status as a fundamental enabler of modern society is reflected in its inclusion as a side event of the FP7 conference 'From Economic Recovery to Sustainability'.

"I think this kind of piggybacking on other events is a great idea," commented Luis Rodríguez-Roselló, director a.i. of Converged Services and Networks in DG-INFSO, in his closing remarks. "It brings in a larger and wider variety of people to explore this exciting new area."

"I think this kind of piggybacking on other events is a great idea… It brings in a larger and wider variety of people to explore this exciting new area"

The Future Internet Assembly is an initiative supported by the European Commission. Its purpose is to foster open interactions and cross-fertilisation across technical Future Internet domains, reaching out to whom ever has the talent. This was the fifth FIA, continuing a process to coordinate European Future Internet R&D which was launched by the Bled Declaration in March 2008.

The Future Internet PPP

Rodríguez-Roselló introduced the Future Internet Public-Private Partnership (PPP). He explained how its focus fell between the long-term visionary work of the EU's Seventh Framework Programme (FP7) for research and the immediate deployments of pre-commercial technologies and solutions covered by the ICT Policy Support Programme (PSP) of the Competitiveness and Innovation Programme (CIP).

The draft FI PPP programme, which is expected to publish its first call in summer 2010, encompasses research into the core technologies of the Future Internet architecture. But it also includes a considerable investment in exploring application pilot projects.

"The ultimate aim of FI PPP is to complement research with innovation," Rodríguez-Roselló remarked.

Building the Future: Future Internet Architecture

FIA Valencia ran several breakout sessions. Three sessions focused on the Future Internet 'building blocks' – the components required for the next generation internet – and how they should work together. Researchers and engineers are working to develop a consensus blueprint, or architecture, known as the Future Internet Reference Architecture. It is hoped that the input and discussions from

these sessions will lead to an action plan on how to achieve a common view with a roadmap on a European Future Internet Reference Architecture.

The three sessions focused on:

- **The Future Internet Reference Architecture:** What are the design goals for the Future Internet? This session presented several reference architectures and tried to draw out how they related to each other. It also explored how different architectures could be evaluated, compared and then contribute to the agreement on one 'European' reference model.

- **Concrete results:** A session to catch up on the output from several EU architecture projects involved in both evolutionary Future Internet designs and entirely new (clean slate) approaches. This session fleshed out some of the theoretical aspects of proposed architectures with some real data, but always with a view on how experimental results could inform a consensus model.

- **Standardisation:** The present-day internet only really works because most of its architecture follows internationally agreed standards. The success of the Future Internet also will rely on standardisation. But how do you get from theoretical architectures and laboratory results to agreed standards? This session heard from several standardisation bodies which explained the objectives and scope of their work. But an alternative way to disseminate results and establish standards was also proposed: open source solutions.

The architecture sessions also heard from the Fraunhofer Institute for Open Communication Systems (FOKUS) about a recently launched competition to compare Future Internet architectures. Like the well-known RoboCup, this Future Internet Tournament (http://www.fit-2010.net) encouraged researchers to investigate the performance of different architectures and hopefully find the ones that merit the most attention from the research community.

Beyond Architecture: Future Internet Applications and Socio-Economics

In parallel to the Future Internet architecture stream, FIA Valencia built on the success of FIA Stockholm and offered several sessions looking at applications and socio-economic aspects of the Future Internet. There were opportunities for delegates to hear presentations and participate in discussions and brainstorming exercises about the application of the Future Internet in the real world – how the Future Internet will be used and experienced by citizens, government and enterprise, and the services and opportunities it should be able to deliver. These sessions dealt with the big question, 'What do we want from the Future Internet?'

> ***"[The] importance of end-users was a key message throughout FIA Valencia: a bottom-up approach to Future Internet R&D must not be neglected"***

The broadening of FIA to include delegates involved in developing Future Internet applications encourages cross-fertilisation between the architecture engineers and the application developers. The aim of FIA is to get everyone talking and working together, without ever losing sight of the needs and desires of end-users.

Indeed, the importance of end-users was a key message throughout FIA Valencia: a bottom-up approach to Future Internet R&D must not be neglected. Perhaps the significant involvement of application developers at Valencia demonstrates one way in which this can be done.

Several key areas of policy and societal issues were selected as the most prominent areas where the Future Internet could have a significant and/or early impact. Several of these themes are addressed by European flagship policy initiatives and European Technology Platforms (ETPs). These

communities recognise the importance of internet-enabled solutions and are happy to engage with the architecture engineers in discussions about these topics.

- **Smart Energy:** The 20-20-20 climate change package adopted by the EU in 2009 is an extremely ambitious policy agenda to cut CO_2 emissions, increase renewable energy and reduce energy consumption. The Future Internet could help on two fronts: enabling better management of the electricity grid and delivering innovative services to help people reduce their consumption.

- **Foundations of Trust:** The Future Internet will see networks and smart technologies infiltrate almost every aspect of our lives. But their success depends on people having enough trust to use them. This session explored the ideas of trust and trustworthiness from both the technical and human perspectives, but with particular focus on end-users.

- **Enterprise:** The Future Internet will lead to Future Business – and that is not business like we know it today. This session explored how the Future Internet could help to develop new business models, foster innovation and help enterprise thrive in the post-crisis landscape.

- **Smart Health:** The demographics of Europe are changing rapidly, driven by the mobility of citizens and an ageing population. Our health care services much adapt – and soon. The Future Internet could revolutionise the delivery of care, from remote monitoring of patients with chronic conditions to assistive technologies which help to keep the elderly active and independent. This session looked at some of the innovations already in development and discussed the aspects of the Future Internet that would be important for e-health.

- **The Economics of Information:** The internet has transformed society, but why? Because it creates and delivers a wealth of information far beyond anything we have ever had before. Of course, the Future Internet must first be able to cope with billions of information exchanges. But we need more than the transfer of digital data. We must also understand this great enabler, how digital information can be used for societal and economic benefit, and how it should be governed.

- **Search:** Embracing the call for grass roots involvement this session explored how SMEs and citizens will continue to rely on an ability to search for information in the future. Keeping the user in the loop is vital, since a search engine's efficiency is maximised when if it can learn from and adapt to user preferences and behaviour. In the Future Internet, searching must take account of content and context. This session sought to analyse the search requirements for different users and identify challenges that could limit the effectiveness of searching in the future.

  "In the Future Internet, searching must take account of content and context"

- **Smart Cities:** Cities can be wonderful, dynamic and exciting places. And they are the perfect environment to put the Future Internet to the test, not least because Future Internet applications promise solutions to many urban challenges. In cities, the Future Internet can really get its hands dirty. But there is a possible research gap: smart city research explores the problem space through application pilots and experimental research; Future Internet research explores open, internet-scale infrastructure and platforms. This session brought the two groups together to share experiences and identify possible collaborations and synergies.

- **Future Internet Research & Experimentation (FIRE):** It's all very well coming up with grand plans for the next-generation internet (r)evolution, but how do you test the ideas? How can you determine which architectures will perform best at a global scale? The FIRE community

is developing experimental platforms that allow Future Internet researchers to conduct their experiments. This session heard from several FIRE projects and learned about the expansion of the community through international cooperation. It then explored how to 'federate' these facilities and tackle issues such as standardisation, brokering for research access and common management policies.

Something to Ponder

The FIA delegates also heard presentations in the plenary session that opened and closed the event:

- In a keynote presentation, Van Jacobson of the Palo Alto Research Centre suggested some relatively simple approaches to 'fix' some of the internet's problems with bandwidth and trust. He proposed that policies for buffering memory in ISP routers could be changed so that they stored, rather than continually flushed, content, according to what was most popular. “A user would typically only have to go one or two steps up the chain before they got the content they wanted,” he said, arguing that this would have a massive impact on traffic and bandwidth. Jacobson also proposes a content-centric networking (CCN) model which is based on naming content instead of naming hosts. But, how much could cached copies of original content be trusted? First of all, CCN focuses on security derived from the data, not from the communication channel that is used. Secondly, CCN security aims to bind together content (integrity), name (relevance) and publisher (provenance).

- In four short presentations in the opening plenary session, a selection of authors summarised their papers published in the FIA Valencia book. Arto Karila described a publish/subscribe model for internet content management; Sergios Soursos explained a methodology for the economic management of overlay traffic; Theodore Zahariadis picked up on the theme of the keynote address with his views on a content-centric internet architecture; and Stuart Clayman described a new framework for monitoring internet services located in 'the cloud'.

- Michael Boniface highlighted the EU-funded specific support actions (SSAs) that are involved in the different themes of Future Internet research. These SSAs aim to coordinate and cross-fertilise the numerous R&D projects running in Europe and build international connections. The SSAs will work closely with FIA, providing resource, organisational support and contributing to future agendas and programmes.

FIA Ghent

With FIA Valencia riding so high, FIA will once again piggyback on another related conference The Future Internet Conference Week, to be held in Ghent, Belgium in December 2010 (http://www.fi-week.eu). Wim De Waele of the IBBT in Belgium said that the week of activities and events promised to be the biggest Future Internet event yet to be held in Europe with more than 2000 delegates expected to attend.

Delegates were invited to contribute their own ideas for themes over the web (voting on topics has now closed but they can still be seen at http://bit.ly/fia-ghent-topics).

FIA Ghent conference website: http://www.fi-ghent.eu.

Opening Session (plenary)

15.04.10 0800: Icelandic volcano erupts. STOP. Ash and ice plume 20,000m high. STOP. Westerly wind. STOP.

*The presentations from the opening session are available on the FIA Valencia website: (*http://bit.ly/fia-valencia-opening*)*

Welcome Message

Mário Campolargo, Director of Emerging Technologies and Infrastructures at the European Commission

Key points:

- *FIA is an opportunity for researchers, industry and users to explore ideas, identify challenges and establish opportunities to find synergism in their work.*

It is a beautiful picture of the internet's success: 400 delegates from across Europe, sitting in an auditorium. Many are talking, chatting. "Ah, we meet at last!" and "Finally we meet in the flesh!" filter through the background buzz of talk.

Event those who aren't chatting are busy. With laptops open and smartphones on, they sit in Valencia, but interact with those back home, and check for updates on airspace closures in northern Europe. With a WiFi connection, distance does not put a stop to everyday activities with colleagues from across Europe and round the globe.

"FIA is providing a key opportunity to bring everyone together to share, learn and support"

The success of today's internet was plainly visible to Mário Campolargo as he welcomed the delegates and speakers and opened the meeting.

"We have a record number of registrants," he announced. "I think interest in the Future Internet topic is increasing dramatically and FIA is providing a key opportunity to bring everyone together to share, learn and support."

Campolargo outlined the programme and underlined its purpose to bring together researchers in a wide variety of activities related to the Future Internet. "Here today we have industry, R&D and most importantly users together in one place," he observed. "Let's benefit from the occasion to work more efficiently and synergistically together."

The Future Internet Public-Private Partnership (FI PPP)

Luis Rodríguez-Roselló, Head of Networked Media Systems at the European Commission

Key points:

- *Europe has a healthy, federated approach to Future Internet research*
- *FI PPP will lead Europe from research to innovation*

Following the official launch of the Future Internet Public-Private Partnership the previous day, as part of the FP7 'From Economic Recovery to Sustainability' conference, Luis Rodríguez-Roselló presented FIA Valencia with an overview of plans for Future Internet R&D.

"We are taking a holistic technological approach," he explained. "The European Commission currently funds more than 100 projects. We encompass everything to do with Future Internet architecture. But we are being careful not to focus entirely on the components; we are thinking about the application layer too, and socio-economic aspects. In Valencia, FIA brings together researchers, industry and users so that we can approach the priorities of FP7's ICT Challenge 1 more efficiently."

Rodríguez-Roselló painted a healthy picture of 'federated' European Future Internet research which encompasses evolutionary and revolutionary, clean slate approaches. Coordination between projects and international cooperation are both strong, with plans for a joint call with Brazil in the near future.

Introducing the Future Internet Public-Private Partnership, Rodríguez-Roselló explained how it fitted into the existing research landscape. It is designed to bridge the gap between the long- range, visionary research of FP7 and the more immediate deployments of pre-commercial technologies and solutions covered by the ICT Policy Support Programme of the Competitiveness and Innovation Programme. The ultimate aim of FI PPP is to complement research with innovation.

"FI PPP will lead Europe from research to innovation"

The purpose of FI PPP would be to drive R&D geared towards establishing Europe as an internet-enabled service economy to address the many challenges of today's society (for example, how to reduce Europe's carbon footprint, how to deal with and leverage the volume of data, how to deliver tailored services to citizens, etc.).

"We are at the confluence of two paradigms," Rodríguez-Roselló asserted. "Our Future Internet vision is based on two primary transformations, both radical departures from the internet of today. First, is the idea of networked objects – the Internet of Things – where embedded chips and wireless technology essentially create an entirely new network of sensors and objects. Second, is the evolution of the web-centred internet, where content is king, but built on a vast volume of data of unimaginable scale. The confluence is 'smart': a new architecture that permits platforms to be built."

The FI PPP draft programme – still open to alterations – leads Europe along an exciting road of discovery (see figure below). With emphasis on large integrated projects, the programme encompasses research into the fundamental core technologies of the Future Internet architecture. But whilst the focus is on the core platform, FI PPP will use application scenarios as a means to explore and define common technologies to develop in the platform.

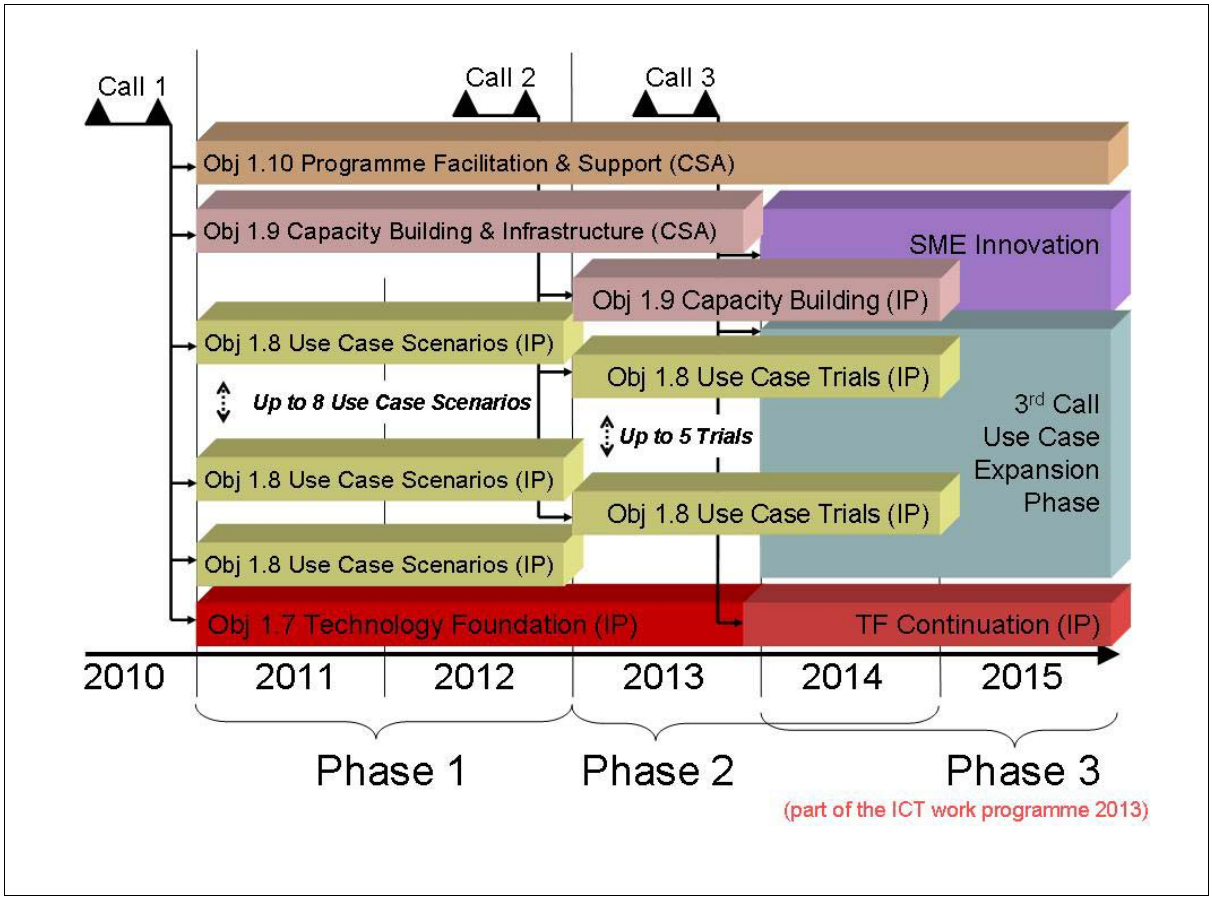

Draft FI PPP Work Programme

The FI PPP is not "business as usual" Rodríguez-Roselló concluded. "We are driving research with a major emphasis on innovation – applications that solve real world challenges. We will build on our existing strengths, but continually look for new players and perspectives beyond the typical ICT sphere. This will be user driven innovation.

Keynote Speech: Content-Centric Networking

Van Jacobson, Palo Alto Research Centre (PARC)

Key points:

- *The memory in routers could be redeployed to store content, making content delivery more efficient by minimising latency and bandwidth*
- *Content-centric networking replaces named hosts with named content – each packet of data holds the key to its integrity, relevance and provenance*
- *Strong and automatic security is an emergent property of a content-centric network*

It was such an easy brief, yet Van Jacobson failed outright. "I've been told you've had a busy week so far," he said, alluding to the packed programme of the FP7 conference "From Economic Recovery to Sustainability". "So it's my job to put you to sleep." But laptops clicked shut, handhelds were switched off, news clips of volcanic ash clouds were paused and attention turned to the platform.

And Jacobson decided to open by talking about failure – the failure of the internet today. "The basis of all our problems," he explained, "started 150 years ago with the invention of the telephone. Telephony linked two devices together so you could have a conversation connected by a wire."

But everything has changed from those early days. The telephone is replaced with the computer, copper wires have become fibre optic cables, the modern switchboard is the router. But one thing stays the same: the internet still works as if we are having conversations. It worked in the 1960s, but today networking has expanded far beyond anything anyone could have originally imagined. Today, people use the internet to watch videos. We want to look at stuff, not talk. It is the content that matters, not where or how we get it.

> ***"It is the content that matters, not where or how I get it"***

With the current model, everything is turned into bits that can be sent over a wire. Sounds OK? No way, it's broken, says Jacobson. "Think of it like this: there's a great video you want to watch so you dial up the base station, make your connection and the bits of video get send down the wire to your phone. But what if it is really popular? If you need a line for every person watching the video, that's a lot of phone lines you need to connect."

The video of Susan Boyle's audition for Britain's Got Talent is approaching 400 million views. That's 400 million phone lines, Jacobson remarked, and we simply don't have enough. Bandwidth is getting constantly squeezed, he said – to murmurs in the audience from those trying to access news sites about Eyjafjallajökull over the venue's WiFi network.

Simple Solution

But the speaker did not stop at a mere observation. He also highlighted an important social consequence of limited bandwidth: artists and the creative community are being stifled. "Anyone can turn their work into bits. I can put my work online and it doesn't matter because it is crummy and no-one is looking at it. But what if it is good? As soon as it becomes popular and millions look at it my ISP kicks me off. The artist has to relinquish control and power over their work and give it to the distributor. I'd like to redress that balance."

The answer is simple; the answer is storage, Jacobson contested, explaining that every router comes with a minimum of one gigabit of memory which could be used to store copies of content. Today, this memory is still considered evil, a legacy of the telecoms model of the internet that perceives the need for router memory as a necessary evil for handling errors. But a small change in router algorithms could remember the most requested packets rather than flushing all the traffic. In this way requests for content from a terminus do not have to go all the way to the host, only just as far as the first stored copy. "This very simple change could create very efficient content distribution with almost no effort at all," said Jacobson. "It is almost indistinguishable from today's networks, but there's a radical change. Content only goes where there's interest, you minimise average latency and total bandwidth."

Trust.com

Jacobson conceded that this replication of content raises problems for of trust and security. "If I type in to my browser the URL for today's New York Times, is what I receive really that publication?" he asked. "It hasn't come directly from the nytimes.com, so how do I know it isn't a corrupt or

tampered copy? Perhaps some other media tycoon has slipped in their own copy of the news? How do I know this is what I really want?"

He briefly outlined the Simple Distributed Security Infrastructure (SDSI) proposed by Ron Rivest that proposes a way to build very strong, content-oriented security. The idea is that security and trust can be derived from data. The data itself provides a links between integrity (content), relevance (for example, through semantic naming) and provenance (e.g. who published it and how it has since been processed).

Jacobson is working on a content-centric networking (CCN) architecture which enables the content itself to be named, signed and associated with a key. And the security of the content gets stronger if more is known about the content, Jacobson pointed out. "In the old internet model you pour something down a pipe; the pipe doesn't know what is in it. With CCN you build up evidentiary trust. Any attack has to be consistent with everything you know about the content or else it will be spotted."

> ***"It is the content that matters, not where or how I get it"***

"In a content-centric architecture like this security will be an emergent and automatic property of the Future Internet." An architecture that solves bandwidth crisis and makes content secure and trusted? Not a bad start to the future of the internet.

Discussion

One delegate asked Jacobson to explain how his **naming system** was different to the **handle system**. "Our mechanism doesn't care about names," he answered bluntly. He explained that a handle system has been incorporated into CCN, which was not trying to change naming mechanisms. "It took 15 years to deploy DNS," he observed, "and we're not going to try and change naming now. We should use names because they've been so hard to set up and will be even harder to take away. CCN has a handle system and we can distribute information based on it."

There was some discussion about the **policies for content storage**. Should content be stored on routers if it is 'boring' i.e. unpopular? Jacobson explained that the rate at which the routers 'forget' the information depends on how often it is requested. The smallest chip on an upstream router is 1GBit so it is not like a cycling modem, so content should stay stored upstream of a terminal for quite a long time. Evidently, routers in the higher tiers would cycle through content much faster, but that would not matter because most requests for popular content never reach them – it would be stored in the intermediate caches. He conceded that the caches at the edge need to be larger than the caches near the centre.

Someone asked Jacobson about the **protocols** in the CCN architecture. He replied that HTTP was wonderful for content production, but not good for content distribution these days. His proposed content-centric networking involves the CCNx transport protocol to deliver content rather than connecting hosts to other hosts.

CNNx website: http://www.ccnx.org

Another participant raised the issue of **copyrights** and digital rights management. Jacobson pointed out that in a content-centric paradigm the content is encrypted; if there are certain rights issues, then the user needs a key to unlock these rights. “Right now we try to protect content by controlling distribution,” he said. “In a content-centric network, the content is protected however it is distributed. You end up with a system that is scalable without affecting revenues.”

Paper Presentation: Publish/Subscribe for the Internet (PSIRP)

Arto Karila, Helsinki Institute for Information Technology

Arto Karila presented the paper explaining the work performed in the FP7 PSIRP project which picked up on the similar themes to Van Jacobson's keynote address. "We wanted to develop a new approach that solves many of today's problems associated with the end-to-end packet transfer of data in which the sender is always in complete control of the communication. "The current set-up leads to some of the internet's biggest problems," Karila insisted. "A sender can send out spam or launch denial of service attacks. There is a rigid protocol stack and efficient multicast is practically impossible on the internet scale."

> ***"The current set-up [with end-to-end packet transfer of data] leads to some of the internet's biggest problems"***

"At the end of the day, users are interested in content. They don't care about where exactly it is located within the network," he asserted.

Karila described an approach that takes a publish/subscribe model of networking. Put simply, information is published and users can subscribe to it. PSIRP is different because it is not an overlay solution (although there is a downloadable IP compatible overlay PSIRP solution that could be used for migration or for testing purposes). "We have built the network from scratch, with no resort to IP protocols. The PSIRP networks effectively decouple the sender and receiver in time and space, but they can both connect with the data (i.e. the publication) that sits between them. The publication is a persistent, immutable association between an identifier and the data value of the publication created by the publisher. Knowing the identifier, the subscriber can retrieve the corresponding data value using the PSIRP network. Publications cannot be changed, so it is possible for them to also be cached anywhere in the network, helping to reduce traffic and keep communication at a local level.

Identifiers in PSIRP are like metadata on today's web. Karila went on to describe the types of identifiers in PSIRP, how they are structured and how they are used by publishers and subscribers. He also outlined the four distinct parts of the PSIRP architecture: rendezvous, topology, routing and forwarding.

Over two years, the PSIRP project has implemented a prototype for the FreeBSD operating system. A PSIRP plug-in has been developed for the Firefox web browser to facilitate testing and the development of applications.

PSIRP website and downloads: http://www.psirp.org.

The full paper 'Publish/Subscribe for Internet: PSIRP Perspective' is published in the FIA Book 2010 'Towards the Future Internet – Emerging Trends from European Research'

The FIA book is available online (http://bit.ly/fia-book-2010).

Paper Presentation: A System for Economic Management of Overlay Traffic (SmoothIT)

Sergios Soursos , Athens University of Economics and Business

The efficiency of traffic management is not just about trying to 'speed things up' and improve the available bandwidth in the network. Efficiency is rooted firmly in the real world; poor traffic management can have a serious financial cost for ISPs.

Sergios Soursos described the ongoing tussle between the underlay of the network and overlays, especially peer-to-peer (P2P) applications. Often measures used by ISPs to reduce their operational costs lead to a decrease in the performance of the overlays.

> ***"SmoothIT project is working in this area to develop traffic management techniques that deliver the elusive 'triple win'"***

Recent studies have shown that the localisation of overlay traffic could help to solve this emerging problem. The SmoothIT project is working in this area to develop traffic management techniques that deliver the elusive 'triple win': where ISPs, overlay providers and end-users all benefit from optimised performance and financial gains.

SmoothIT has identified numerous solutions. For example, localisation could occur if peers in a domain have their overlay neighbours rated according to their underlay proximity. A peer sends a list of peers to the SIS (SmoothIT Information Service), which obtains from the underlay locality information about each peer on the list, then rates them. This list is returned to the requesting peer and overlay operations then favour peers in them same domain.

Another approach is to enable the ISP to download faster into its domain, then distribute the content among local peers. Soursos described the ISP-owned-Peer (IoP), controlled by the ISP and given extra bandwidth and storage to make its content available quickly to its peers. Here the IoP is part of the peer network, but with enhanced downloading capacity. Another approach involving a highly active peer (HAP) is similar, but decentralised. In this set-up, the most active peers are dynamically upgraded by the ISP and given enhanced bandwidth and connection speed.

Finally, it may sometimes be necessary to introduce some level of collaboration between P2P domains.

Soursos described the ETMS architecture proposed by SmoothIT. Initial validation and test-bed trials have been run and the concept has been proved, Soursos reported. Analysis of simulation results is ongoing (preliminary results have identified some side effects in small-sized P2P swarms).

The mechanism for promoting locality (called BGP-Loc) seems to produce a 'win' for the ISP (inter-domain download traffic was reduced by 3%) and a 'no-lose' for end-users (no degradation in the service). For the IoP and HAP mechanisms, the benefits are harder to guarantee for the ISP and depend on the cost of inter domain connections and operational expenses.

SmoothIT website: http://www.smoothit.org.

The full paper 'ETMS: A System for Economic Management of Overlay Traffic ' is published in the FIA Book 2010 'Towards the Future Internet – Emerging Trends from European Research'.

The FIA book is available online (http://bit.ly/fia-book-2010).

Paper Presentation: Towards a Content-Centric Internet

Theodore Zahariadis , Synelixis

In this presentation Theodore Zahariadis continued to explore various Future Internet options which could improve the way end-users search for and retrieve content. Like the speakers before him, Zahariadis advocated the content-centric paradigm, where the actual data, rather than its location, is the key asset.

Just in case delegates were not yet convinced that the current architecture is all but broken, the speaker reminded them of its inefficiencies and the extraordinarily long pathways that data must follow before it gets from a content server to a user's terminal. Reiterating suggestions by other speakers, Zahariadis said that the system could be improved if:

- content could be cached closer to the end-users;

- routers could identify/analyse what content is flowing through them and search engines could use this information;
- the network could identify the best path to the user;
- content could be automatically adapted to the context (optimising bandwidth for different devices).

He went on to describe an evolutionary Future Internet Architecture (EFIA) which makes use of hierarchical virtual clouds. Four levels of virtual cloud are associated with the networking, content, services and applications. The advantages of EFIA include its backward compatibility, flexibility and easy deployment.

Zahariadis also described an Autonomous Layer-Less Object Architecture (ALLOA) which he suggested could transform media from today's stream of pixels and sound to malleable content objects, related with each other in time and space. Applications of the ALLOA model could include user-friendly applications that would make it easy to build three-dimensional virtual worlds from simple two-dimensional sketches. "ALLOA could open the door to content mash-ups, the reuse of objects in different contexts and enable online collaboration to edit audiovisual content. Content objects have exciting possibilities and great value," Zahariadis remarked.

"Content objects have exciting possibilities and great value"

The full paper 'Towards a Content-Centric Internet' is published in the FIA Book 2010 'Towards the Future Internet – Emerging Trends from European Research'.

The FIA book is available online (http://bit.ly/fia-book-2010).

Paper Presentation: Monitoring Service Clouds in the Future Internet

Stuart Clayman , University College London

While the volcanic ash cloud continued its slow spread from the UK to northern European and south into France, Stuart Clayman presented research related to create quite a different sort of cloud.

"One of the features of the Future Internet," explained Clayman, "is that many of its core features and functions will be based on virtualised resources – even virtual machines that are something quite different from the actual hardware. These virtual resources come together to create what we call 'service clouds'. But how do you manage these resources that are very dynamic and volatile. Components can jump around, federate and change their virtual location at any point, even whilst running. We need a monitoring system to collect and report on the behaviour of these resources."

He continued: "Existing monitoring systems have addressed [the] monitoring of large distributed systems, but they have not addressed a rapidly changing and dynamic infrastructure seen in service clouds." He then introduced seven key elements of cloud monitoring in the Future Internet: scalability, adaptability, elasticity, federation, migration, isolation and autonomy.

Clayman briefly described some of the features and architecture of the 'Lattice monitoring framework' which has been designed with these characteristics in mind. The monitoring system itself is based on the idea of producers and consumers. Data producers collect data from probes in the system; consumers then read the monitoring data. Consumers and producers are connected via a network which distributes the collected measurements. It is possible to switch the distribution system as necessary, for example from IP multicast to an event bus to a publish/subscribe mechanism.

The Lattice framework has been tested in the Reservoir service cloud. It has been used to build a monitoring infrastructure that collects, processes, and disseminates network and system information in real time. The framework can therefore form the basis of a service cloud management system. “But there's one important thing to remember,” Clayman concluded. “You can't add monitoring afterwards – it has to be built in.”

“[T]here's one important thing to remember… You can't add monitoring afterwards – it has to be built in”

Reservoir website: http://www.reservoir-fp7.eu.

The full paper ‘Monitoring Service Clouds in the Future Internet’ is published in the FIA Book 2010 ‘Towards the Future Internet – Emerging Trends from European Research’.

The FIA book is available online (http://bit.ly/fia-book-2010).

Prize Award

Ioanna Roussaki, Nikos Kalatzis, Kevin Doolin, Nick Taylor, Guido Spadotto, Nicolas Liampotis and Howard Williams were conferred with the best poster award. Their posted presentation entitled "Self-Improving Personal Smart Spaces for Pervasive Service Provision" was voted as the best poster at FIA Stockholm.

The poster presented some of the work of the PERSIST project which seeks to drive forward pervasive systems by centring them on the portable devices that people already carry around with them. The premise of the project is that you can now get smart homes and smart offices, but on the way from home to the office there is no support for 'smart people'. PERSIST seeks to provide a personal smart space by using mobile devices and personal area networks. When a person enters a smart office, their personal smart space will negotiate with the smart office and be able to offer them the services they need.

A summary of the paper available at http://bit.ly/fia-valencia-award.


The European RTD
From E
to Susta
IV Conference on the FP7 in Spain
12-16 April 2010 / Valencia - Spain
eu
12 13 14 15 16
FP7: European Innovation
and RTD Transforming
Sectors Week

Architectures of the Future Internet

Rapporteur: John Buckley

Introduction

The current internet has evolved from its earliest beginnings around the universality of the TCP and IP protocols; it is this simplicity that has enabled an open explosion of globally reachable data, services and applications.

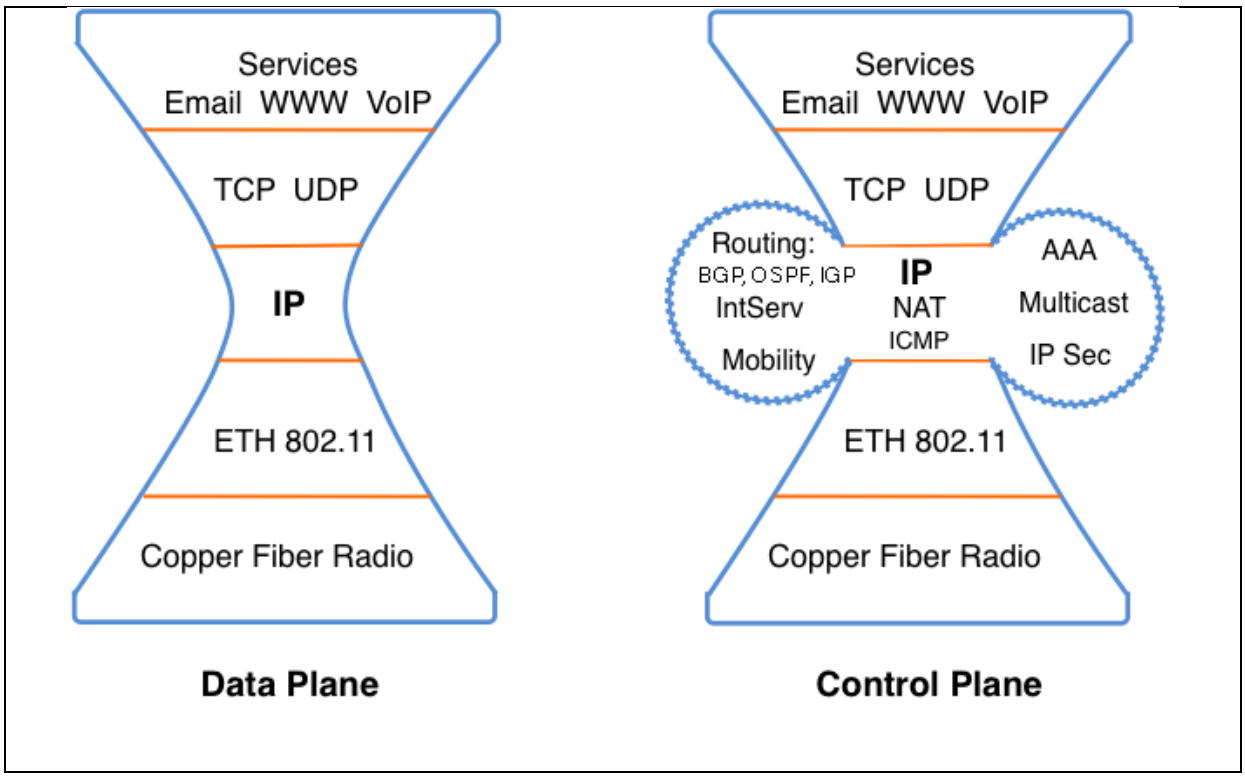

Current Internet Reference Model, picture from introductory presentation

But we have paid a price for this simplicity. The size of the internet (number of end points, quantity of data, number of services) is growing such that operational and management costs are reaching crisis point.

Moreover, a lack of various features in the current architecture presents a barrier to innovation. Mobility, in-system management, autonomic deployment of new functions and services, quality of service, trustworthiness, security and reliability are all poorly addressed. And where they are addressed, it is often by ad hoc solutions that may be partial, costly and difficult to scale.

For these reasons, over the next two decades the internet needs a complete revamp; this work to create this Future Internet needs to start within the next five years.

But the global Future Internet cannot be created by a linear process of engineering and technological design. Before work can really start, we need a reference architectural model – a kind of blueprint to describe and guide the process of internet evolution and unify key technology developments in the future.

An architecture like this provides a breakdown of the functional components and interfaces that will exist in the Future Internet. It provides an intellectual framework against which requirements and

technical solutions can be discussed; a reference model supplies a harmonising framework for ongoing R&D.

Achieving an enduring and useful reference architecture will be an extremely complex task, drawing together a wide variety of expertise and experience. It will be difficult to understand the totality of the requirements, to balance them against technical, societal and commercial constraints, all the while achieving the necessary flexibility for new applications and services to reach a wide range of users and customers.

The Future Internet will differ from the current internet in several ways. The architecture will build in several important features including:

- a higher degree of 'virtualisation' at all levels (i.e. applications, services, networks and resources);
- greater mobility;
- autonomic management, security, privacy and continuity of service;
- minimised energy consumption.

Existing FP7 projects are making tangible progress towards these goals, illustrating that the effort is productive and worthwhile.

To be effective, a consensus architectural model – the Reference Architecture – must be promoted globally through standardisation bodies. To this end, a number of standardisation initiatives have begun or are about to begin, although some people have voiced their concern that all this work may not be timely and that it is premature to hope at this stage for the necessary consensus.

At present, a number of different architectural models are being developed by different groups; these models are still being explored, discussed and refined. Position papers have also been published outlining the requirements and research challenges for Future Internet architectures. The research community continues to debate the extent to which communication, service, content, management, 'things' and resources should be integrated into architectures.

The architecture track during FIA Valencia sought to advance this work. Three sessions focused on:

- **The Future Internet Reference Architecture:** looking at proposed architectures, discussing how they could be evaluated and compared and how they could contribute to a single 'European' reference model.
- **Concrete Results:** presentations of actual results from some of the existing EU architecture projects involved in both evolutionary Future Internet designs and entirely new (clean slate) approaches.
- **Standardisation:** introductions to the current work in this field from several standardisation bodies, with ideas on how to disseminate results and contribute to Future Internet standardisation.

Together these sessions aimed to add significant momentum to collaboration on a consolidated European Reference Architecture. The aim is to present a consensus model during FIA Ghent in December 2010.

Session 1: The Future Internet Reference Architecture

Chair: Alex Galis, University College London

Objectives

The first session focused on why we need a new architecture for the Future Internet and what its design goals should be. Speakers presented several existing reference architecture models, and they tried to relate their particular architecture model to other models and to overall goals. The session also wanted to explore ways in which architectural proposals could be evaluated and compared. This purpose of the session was to take a significant step towards building consensus on a European Reference Architecture.

Presentations

• Future Internet design goals (functional and architectural metrics, cost performance metric potentially from different stakeholder perspectives)	Dimitri Papadimitriou	Alcatel-Lucent
Examples of Future Internet reference models		
• Future converged network reference model – consequence of novel access technologies	Mikail Popov	Acreo AG
• Radio networks	Panagiotis Demestichas	University of Piraeus
• RWI reference model	Srdjan Krco	Ericsson (SENSEI project)
• MANA reference model	Alex Galis	University College London
• FCN reference model	Theodore Zahariadis	Synelixis
• FISO reference model	Frederic Gittler	HP Labs
• How to evaluate and compare architectures	Kurt Tutschku	University of Vienna

*The presentations from this session are available on the FIA Valencia website (*http://bit.ly/fia-valencia-architectures*).*

Summary of Presentations

Principles for design

The global Future Internet is a complicated structure that cannot be created by a linear process of engineering and technological design, argued Dimitri Papadimitriou. A prior requirement is an agreed understanding about what the internet has to do and the basic functions it must incorporate. These must address wider concerns than function and performance, also satisfying societal goals and possessing the flexibility to evolve towards unforeseen future requirements and uses. For this reason, there is a need for prior work to set the reference framework within which the Future Internet can be designed.

This framework is the **Reference Architecture** of the Future Internet**.** A reference architecture provides a long-term technology independent perspective, and a conceptual framework for understanding the functional components and interfaces that will exist in the Future Internet. This

provides an intellectual context against which requirements and technical solutions can be discussed. The reference model also helps to shape and direct ongoing R&D.

Papadimitriou explained that a reference architecture contains common design objectives, principles, invariants and interfaces; these are identified, defined and accepted by the research community which can then use the architecture as a common baseline on which to build design and implementation models.

Design principles are drawn from requirements covering both functional and performance objectives. The architecture will comprise a set of components (functional entities, procedures, data structures etc.) and the characterisation of their interactions (messages, calls etc.). The architecture must address communication, service, management, content, objects and resource aspects and the distributed and dynamic interactions between entities. Our understanding of current systems will be fundamental to assess both quantitative and qualitative Future Internet challenges.

Achieving a useful and enduring model for the reference architecture is an extremely complex, challenging and conceptually immense task, and must draw on a wide variety of expertise and experience. “We have to identify and understand every possible Future Internet requirement, and balance these requirements against technical, societal and commercial constraints,” said Papadimitriou, “all the while achieving the necessary flexibility for new applications and services to reach a wide range of users and customers.”

Requirements that are either neglected, or addressed with ad hoc solutions of limited effectiveness today, include:

- quality-of-service;
- scalability;
- mobility;
- the balance of wired and wireless connection;
- multi-layer and multi-domain integration;
- inter-domain control and management;
- unification of heterogeneous technologies;
- reliability;
- availability;
- identity management;
- trust;
- privacy;
- openness;
- flexibility;
- evolutionary potential;
- low entry cost for inexpert stakeholders;

- power consumption;
- environmental friendliness.

"This is a large problem to solve, but before we can solve anything, we have to discover the problem – and a way to discover the problem in the first place," Papadimitriou asserted.

Various approaches are possible. A consensual top-down approach may be followed. Alternatively, different groups might propose different solutions; they might then come together searching for consensus, or they may exploit more or less competitive approaches to find the best solution. There is a danger that individual solutions may concentrate too much on particular points, perhaps of known failures in current systems, to the detriment of the overall requirement. Another danger is Conway's Law from 1968, that "organisations which design systems are constrained to produce systems which are copies of the structures of these organisations" *(quoted by Arian Zwegers)*. However, success will be easy to measure: a globally agreed architectural solution.

Virtualisation

Existing European R&D projects are making tangible progress towards a model for the Future Internet Reference Architecture. These contributions show that the effort is productive and worthwhile. The presentations of different architectures in this session revealed some common themes and gave insights into the features that will emerge in the consensus Reference Architecture.

A common theme is "**virtualisation**". This is being proposed for different Future Internet layers, for example in the application, service, network and communication layers. Virtualisation involves the interposition of functional elements between other entities which makes it possible to add useful services to their interaction.

To give a simple example, an application in today's internet may request and use a communication link, interacting with it via an established interface. If that communication service is "virtualised", then the application requests a virtual (that is, abstract) communication service from an intermediary. It is actually the intermediary that has to find and control the communication service. We might call this new element a "broker", "resource manager" or some such name.

Does virtualisation add value? It might add value by locating several suitable links, thus allowing the application to use different communications media at different times (for example wired or wireless). It might permit the user application to specify more tightly the parameters of the communications path than currently possible. It might transparently re-route communications in the face of problems. In this way, virtualisation adds, among other things, to the flexibility, technology independence, evolution potential, openness, inter-domain unification, reliability and availability of the internet.

The following figures give a flavour of some of the architectural features presented in this session (see presentation slides at http://bit.ly/fia-valencia-architectures for more detail). The diagrams show possible views of a reference architecture with virtualised service and network support and elaborate additional functional elements that may exist.

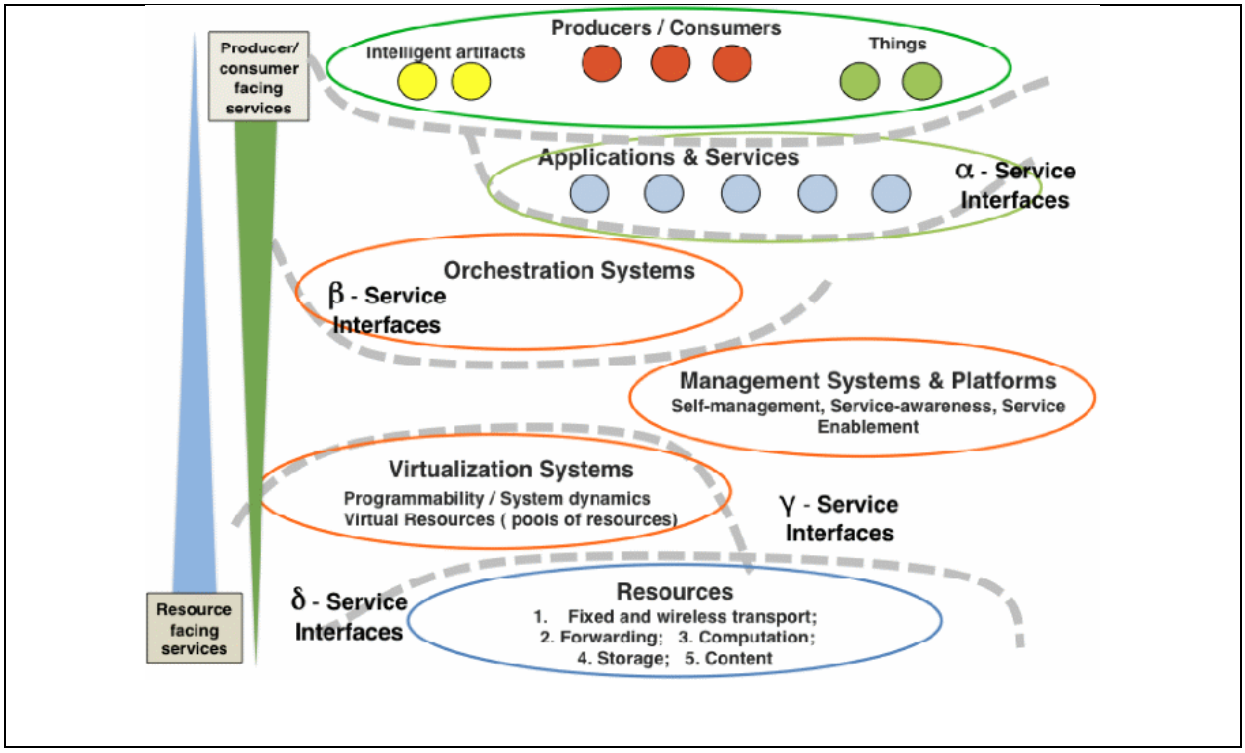

Future Internet Network, Service and Management Architecture (MANA) consolidated approach, picture from MANA Reference Model presentation

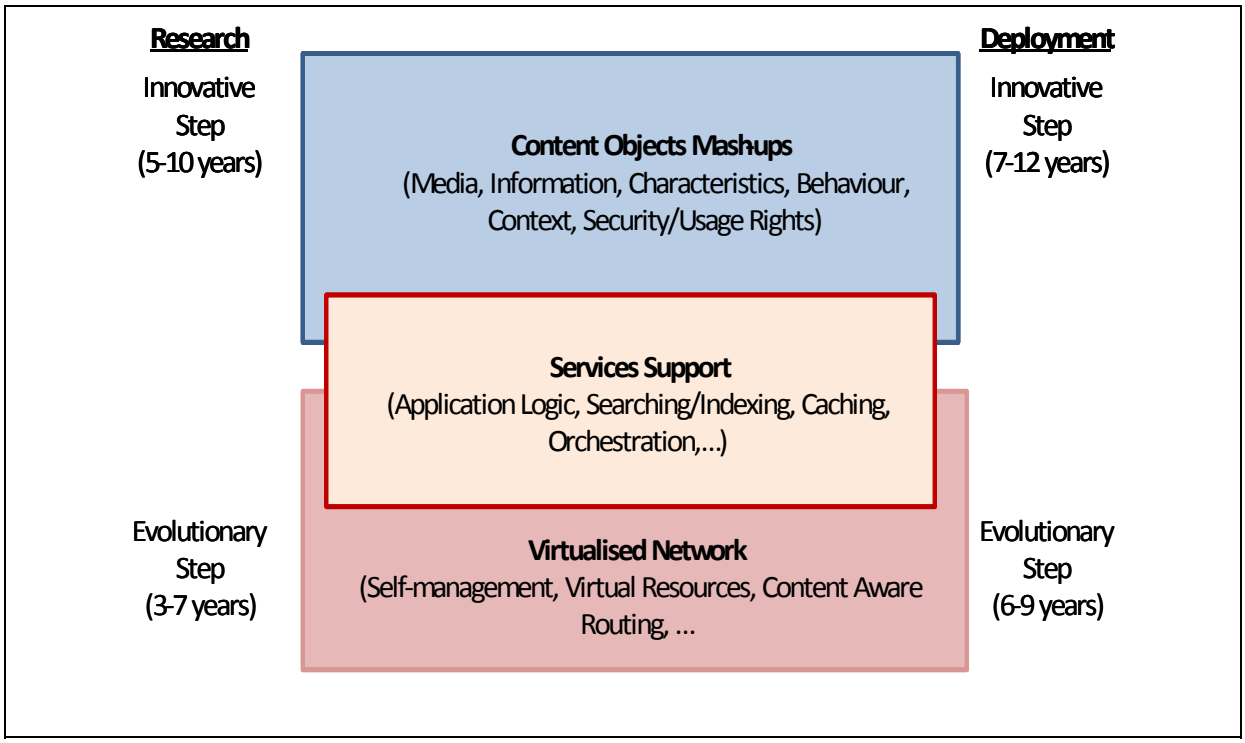

Future Internet Content Architecture Consolidated Approach, picture from FCN Reference Model presentation

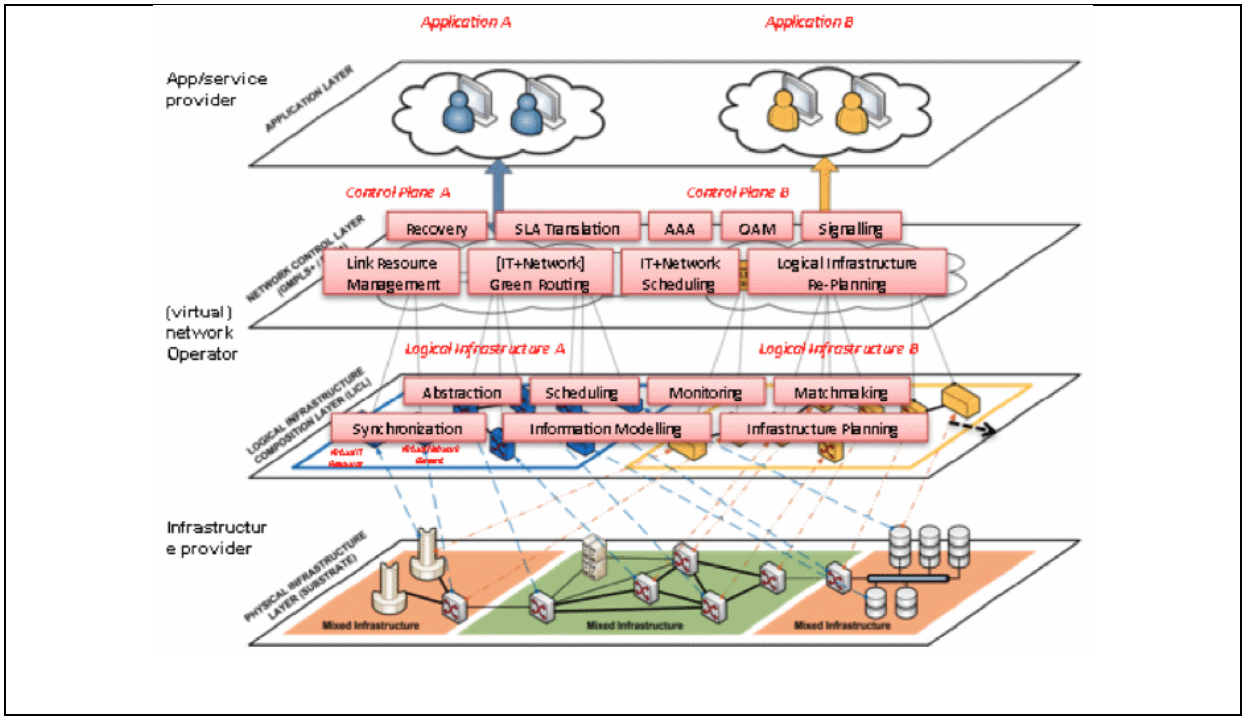

Application-Network Interaction: Logical Infrastructure Composition Layer, picture from Converged Network Reference Model, GEYSERS project

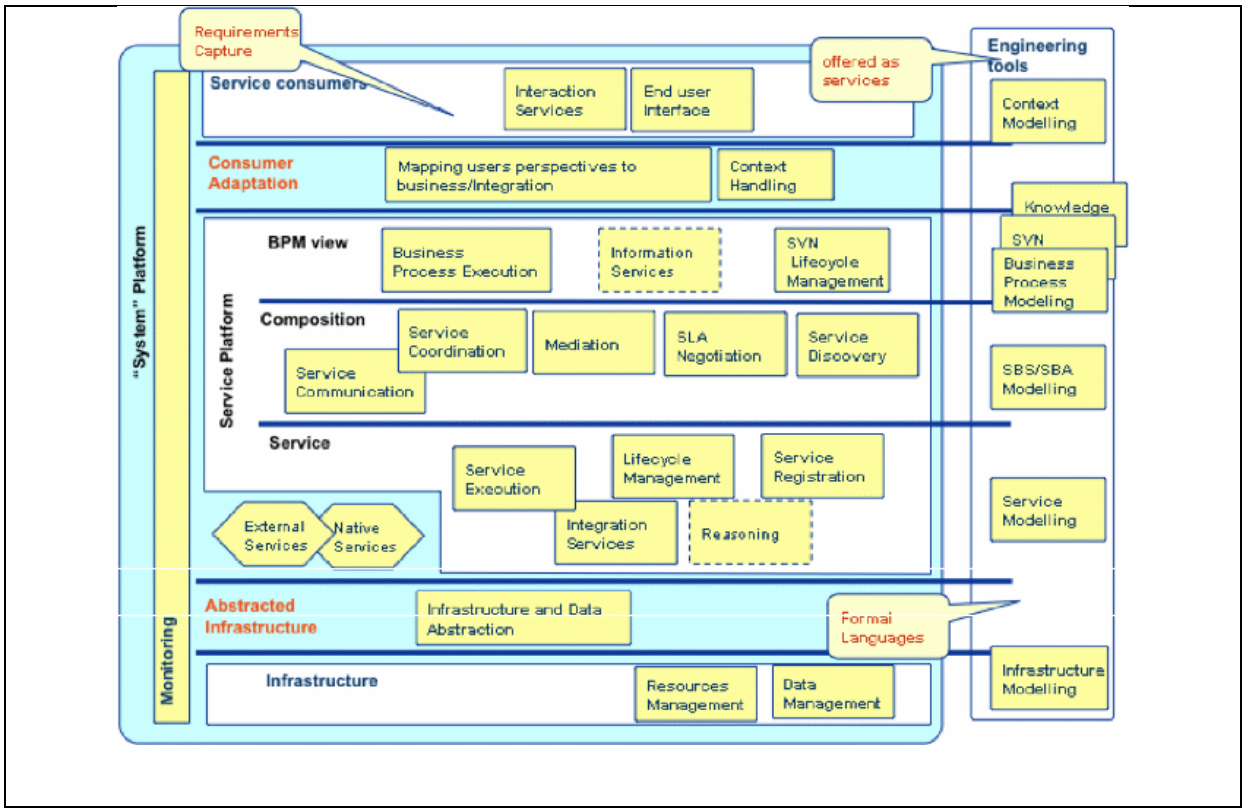

Service Platform Component Interaction, picture from FISO Reference Model presentation, NEXOF project

The current internet has a view of the world where the principal resources are communication links and the main objects are end-points to be linked together. But the Future Internet architecture needs to be much more comprehensive. The principal resources should also include processing power, items of data (or content) and storage, as shown in the MANA reference model. The ability of

an application or service to request data content simply by knowing its name will be a major simplification, while the explicit provision of storage in the network will facilitate content caching and distribution (see also keynote address by Van Jacobson). Naturally, the Future Internet will have to contain functional infrastructure elements that translate a content object's (virtual) name into a (physical) way of reaching it.

Content-centric networking (CCN, see also keynote address by Van Jacobson) is the principle that data (and by extension anything that a human can perceive with the senses) in the internet will be found using its unique name and *not* its location within the network topology. This will simplify access to data, enabling application developers to have less expertise in communications and networking. CCN will effectively virtualise the location of data and content. The data can migrate to form as many copies as are needed, which can be located at the most convenient points on the network. New service elements will be needed for data and content to make known (or publish) their existence in the internet, and for user services to find, subscribe to and access that data, as proposed, for example, in the Publish-Subscribe Internet Routing Paradigm (PSIRP, see paper presentation during the opening plenary session).

Virtualisation is potentially valuable in the service and application domains as well as in communications and localisation services, as described above. It permits composite and mash-up services to be composed by the orchestration (that is, the bringing together) of other and presumably simpler component services. This supports the richness of services offered, the openness of the internet to service providers who are not experts in internet technologies, and reduces the time and effort of service realisation.

Network governance and management

The management of the internet will be a major topic in the future Reference Architecture. Many types of management are needed. The examples above illustrate the need for management systems underlying object virtualisation. It will also be required for reliability, security, privacy and quality of service management. However, the Future Internet would have a serious weakness if it had to rely on central management; research is required to develop suitable alternative management paradigms. It may be possible to have conceptually centralised but actually distributed management functions. It may be possible, for example in the case of data security and trustworthiness, to embed the security in data and other objects rather than relying on a central, unitary management mechanism.

Autonomic management, or self-management, is a most desirable goal within the Future Internet for two reasons. The first is that it avoids or reduces the high cost inherent in using highly skilled people to keep networks operating. The second relates to internet openness, as it assists non-expert users to connect with and use the services of the internet. Autonomic management combined with cognitive (i.e. artificially intelligent) functional elements will be especially valuable. Cognitive components and systems discover for themselves the network and service environments around them and so do not require difficult, expert configuration.

Finally, the Future Internet has to recognise and cope with **polymorphism**. Different network configurations, virtual network interconnection structures, applications and service compositions will develop and appear at different times and in different places. The Future Internet may well federate into separated clouds of computers, services and users. Nonetheless, these clouds will need to be able to create associations with one another and achieve connections at any level via service-based interfaces when required. While one might instinctively yearn for the seeming simplicity of the physically orientated TCP/IP model of the current internet, it is well, remarked Frederic Gittler, to remember Albert Einstein's insight: "Make things as simple as possible, but not simpler".

Evaluation and comparison of model architectures

Evaluation of reference architecture models is a necessary and difficult part of the process of finding a suitable Future Internet Reference Architecture, remarked Kurt Tutschku who offered a comprehensive view of how proposed architectures could be compared. For some aspects, absolute evaluation against definite requirements will be essential; however, for others a relative evaluation may be more appropriate, he argued. Some aspects are measurable, some are harder to quantify but can be made measurable, while others (e.g. maintainability) are very hard to measure. A point-scoring approach using normalised scores for a goal-derived metric set could be used. But Tutschku ended with a note of caution: context is important. The idea of 'quality' depends greatly on different intended uses and purposes.

Discussion

The discussion in the first of the architecture sessions concentrated on the methodologies being followed to develop reference architecture models. There was the issue of many proposals versus one: should everyone develop his or her own architecture then compare and compete to find the best, or should everyone work together and contribute to a single project from day one? Was it indeed possible to proceed without first agreeing a top-level abstract definition of the internet? Because we are still discovering the problem to be solved, there are dangers in fixing too much on a high-level view, so some people felt it was safer to encourage voluntary input and permit 'bottom-up' work where various groups tackle known problems. The example of networks based on Asynchronous Transfer Mode (ATM) shows clearly the danger of commitment to a top-down approach. On the other hand, it would be seriously unhelpful to have too many proposals, as it would become very difficult to compare them and reach a balanced conclusion.

"Might a Euro-centric approach be globally impoverished?" Of course a global approach is essential, hence the work with standards bodies, although current European initiatives are necessary to gain experience and give direction to European research. Although this work will take time, it should not be overly delayed, or else visions will not be visions any more.

A speaker made a plea that flexibility should not be sacrificed in the pursuit of performance. This means that the model for the Reference Architecture should not sacrifice the openness of the internet and its ability to evolve in order to get short-term speed and performance.

Session 2: Concrete Results from Existing Projects

Chair: Marcus Brunner, NEC Labs Europe

Objectives

The second session showed concrete results from various European architecture projects. Both evolutionary approaches as well as "cleaner slate" lines of attack were presented. Engaging with various themes, such as networking, virtualisation, identity and content delivery, the presentations offered reference models that showed conference participants what the Future Internet architecture might be like. This session finished with an open discussion on how to build on these results to achieve a European Reference Architecture for the Future Internet.

Presentations

Polymorphic internet

• ANA project	Martin May (for Christian Tschudin)	University of Basel
• The 4WARD system model – defining network architectures for future networks	Martin Johnsson	Ericsson

Virtual Infrastructures (Network, Service, Storage)

What is the role of abstraction? How far are we from the commercial world? Are we achieving transparency? Validated promises and results? Management and operation of virtual infrastructure.

• Service computing clouds: RESERVOIR project	Yaron Wolfstal	IBM Haifa Research Laboratory
• Virtual networks: 4WARD project	Djamal Zeglache	Telecom SudParis, Institut Télécom
• Management of virtual infrastructure: AutoI project	Alex Galis	University College London
• Sustainable identity framework for the Future Internet: SWIFT project	Amardeo Sarma	NEC

Content-based networking

• Publish-subscribe networking: PSIRP project	Arto Karila	Helsinki Institute for Information Technology
• Content delivery: SEA project	Giovani Pau	UCLA
• Media centric networking	Adolfo M Rosas	Telefónica Investigación y Desarrollo

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-architectures).

Summary of Presentations

Architecture is by nature abstract, so presentations and discussions about architecture use a lot of abstract terms and concepts. On the one hand, there may lie behind these terms much profound thought, elaboration, development and solved problems. On the other hand, an abstract function on a diagram may express a desire that a problem will be solved. Accordingly, it is most important to have concrete work, ideally verified by experimental application.

The **ANA project has** worked on a meta-architecture for various addressing and naming styles and core networking concepts. It supports polymorphic network structures and aims to identify fundamental scaleable autonomic networking principles and to demonstrate autonomic steering. The project has developed tangible open source software for Linux and smart phones.

4WARD is developing a system for flexible and modular network components and architectures ("Let a thousand networks bloom" is its catchphrase). It also provides an extensive framework for a network of information (NetInfo), virtualisation of networks and resources, in-network management and self-management. These models, while future-proofed, can be applied to today's systems also. The 4WARD system models have the potential to replace the ISO seven-layer model commonly employed in the ICT world. It is implemented on the Heterogeneous Experimental Network (HEN), a cluster of over 110 computers.

The **Reservoir** project is creating a revolutionary service cloud infrastructure where resources and services can be transparently and dynamically managed, provisioned and relocated like utilities, virtually and "without borders". Four application scenarios are driving the design, taken from the real worlds of business, communications, computing and government. The system should enable utility-like deployment of services, creating the basis for service products in the Future Internet (see also the paper presentation by Stuart Clayman from the opening plenary session).

AutoI is developing a self-management infrastructure for virtual environments (virtual networks and services) based on virtual machines. The new infrastructure is based on service enablers, orchestration, knowledge, management and the virtualisation planes for physical resources. The project has developed a number of tangible open source platforms for managing virtual environments. There are two test-beds of many processors, including the G5000 test-bed, which has 5000 nodes, supporting experiments in self-configuration and self-performance management.

Amardeo Sarma focused his presentation **(Sustainable identity framework for the Future Internet)** on concrete results coming from several identity and privacy projects since FIA Stockholm. He also looked at how consensus is building within the research community, based on R&D and the experience of experimental privacy-enabled identity systems and infrastructures.

This work has been coordinated by the FIA Trust and Identity and MANA caretakers and has involved a number of projects including Think-Trust, PrimeLife, TAS3 and SWIFT. **SWIFT**, for example, is examining the concept of partial identities provided across the layers, whereby user IDs, service IDs and application IDs are made available on a 'need to know' basis. The premise behind partial identities is that applications do not need to request the whole of a user's identity data; typically only a subset of ID information is actually required. For example, dating services may need to know that you are an adult, but do not need to know when you were born. The overall aim of the projects working in this area is to improve on the privacy and security available today and to make privacy a "European trademark of the 21st century" based on R&D funded through the FI PPP and, in the longer-term, through FP7 and its successor.

The **PSIRP** project concentrates on the 'publish-subscribe' paradigms necessary for content-centric networking (see paper presentation by Arto Karila from the opening plenary session). The project has achieved an open-source implementation of the core architecture and is running some test applications. An integrated system with some applications will be demonstrated at ICT 2010 in Brussels in September 2010.

The FP7 **SEA** project is aiming to provide an efficient content delivery and adaptation system across different network technologies and architectures. Its PlanetLab infrastructure, consisting of over 954 nodes at 479 sites across the globe, provides an experimental test-bed to evaluate the project partners' different technologies.

Content delivery is a topical problem for telecommunications operators. As they see data flows rising exponentially, they find that traditional network architectures, with data centres in their transit layer, are inadequate. Adolfo Rosas presented Telefónica's development of **ENVISION**, which provides a new content distribution network with its own control layer. Incremental overlays, such

as this, can be quick and efficient to deploy, provided there is cooperation with the telcos to expose the best features of the underlying networks.

Discussion

The discussion at the second session addressed the issue of c**ommercialisation of results**. Some of the projects had work that was deployable right away, having commercial partners who were keen and interested in getting solutions to market. Examples of exploitable results were in the areas of autonomic networking and service clouds. However, forward-looking research is not always a fast win; it can take time to promote understanding of new concepts.

One questioner was most disappointed that so little of the current work appeared to have addressed power consumption and "green issues". The panel agreed that this was important and must be taken into account, although they felt it would not radically change the architectural thinking so far achieved. Cloud computing, for example, had been demonstrated to reduce energy consumption. Of course there were basic hardware as well as architecture issues to consider.

*More details about these projects can be found in the specific presentation available on the FIA Valencia website (*http://bit.ly/fia-valencia-architectures*).*

Session 3: Standardisation and Next Steps

Chair: Henrik Abramowicz, Ericsson

Objectives

The third session focused on how to move research results forward to standardisation bodies and so stimulate take-up by the international community. Several standardisation initiatives regarding the Future Internet were shown. The session also looked at alternative ways of disseminating results through open source solutions and through explicit competition. Finally, a European Commission initiative to support the development of a common model of the Reference Architecture ready for presentation at FIA Ghent was unveiled. It was hoped that this session (and the previous sessions) would lead to a clearer understanding and common view of an action plan and a roadmap to a European Future Internet Reference Architecture.

Presentations

• Move Future Internet research into practice - closing the standardisation gap	Didier Bourse	Alcatel-Lucent
Standards Organisations		
• ITU-T SG 11 focus group on future networks	Takashi Egawa	ITU
• ETSI ways of standardising Future Networks	Ultan Mulligan	ETSI
• ETSI Future Internet architecture ISG – industrial specifications	Guillermo Cisneros	Technical University of Madrid
Next steps		
• How to make use of open source (as alternative way of creating "standards") / What facilitators are required (directory, IPR problems)?	Rui Aguiar	University of Aveiro

• The Future Internet Tournament: competition & cooperation - how to steer internet evolution	Tanja Zseby	Fraunhofer Fokus
• EU initiative: towards architectural design principles and a reference architecture for the Future Internet"	Arian Zwegers	European Commission, DG-INFSO

*The presentations from this session are available on the FIA Valencia website (*http://bit.ly/fia-valencia-architectures*).*

Summary of Presentations

The path ahead

To be effective, the Reference Architecture must be widely accepted and used, and so must be promoted globally by way of standards. All this effort must be strategically directed. To this end, a number of standardisation initiatives have begun or are about to begin and several of these were introduced to the delegates.

Innovation via the standards process implies the recognition of a target eco-system, and an alignment of the research and standardisation cycles. The traditional standards process is not normally well adapted for research, which requires more pre-standardisation work, a more lightweight approach and an iterative process that is open to academic participants. In contrast, the standardisation process is a mature cycle, while the research cycle needs a well-defined methodology and regular evaluation of results.

A coherent strategy with a roadmap and clear plans will make a clear impact. European institutions have plenty of experience of standardisation processes. We will have to standardise the functional and network architecture components of the proposed solution, and the communication protocols between them.

Nevertheless, the standards bodies are open to the research process. They invite and encourage industry participation in pre-standardisation discussions, such as the ETSI Industry Specification Groups (ISGs) and the ITU-T focus groups. Standards are pervasive in the ICT sector; all of them having been established by positive, voluntary contribution rather than governmental dictate.

But simply submitting some input to standards bodies is not enough: there must be ongoing participation in accordance with a coherent standards plan. It is not yet conclusively clear, however, whether current work on the model for the Future Internet Reference Architecture is sufficiently advanced to proceed to the standardisation process.

There is a delicate balance between the roles of competition and cooperation in the innovation and standardisation processes, noted Tanja Zseby. On the one hand, competition triggers evolution; on the other hand, cooperation enables advanced solutions. Competition must therefore be encouraged, though it is necessary to control the selection processes, taking care to avoid weak compromises. Heterogeneity will be a feature of the Future Internet, allowing systems to specialise for different situations, though diversity can also spread vulnerabilities. “And disruptive ideas are helpful,” Zseby asserted.

Zseby went on to announce the Future Internet Tournament, where Future Internet solutions will be subjected to a programme of challenges and compete against one another to solve network

problems. However, cooperation to develop entries for the tournament is encouraged. Cooperation requires incentives, for example information and resource sharing, and these incentives need to be explicit in the form of money savings or protection against threats. Ideally, it should be easy to find partners and set up dynamic coalitions, for example through the Node Collaboration System proposed by Fraunhofer Fokus.

Rui Aguiar strongly advocated the role of open source software (OSS) in the Future Internet. This is more than an implementation issue, he argued, suggesting that it is a development method that harnesses the power of transparency and distributed peer review. “OSS promises better quality, higher reliability, more flexibility, lower cost and an end to predatory vendor lock-ins,” Aguiar remarked. In a world of increasing complexity, open source may facilitate the development of an effective Future Internet by sharing solutions that can be refined and verified by everyone. Arguably OSS is an unavoidable trend that should be promoted as a way of impacting the Future Internet by effective and accepted standardisation.

Work in the standards bodies

The European Telecommunications Standards Institute (ETSI) is open to research results, and encourages the formation of Industry Specification Groups (ISGs), open to non-members of ETSI. Three relevant ISGs, deriving from FP7 R&D projects related to future networks, include:

- AFI (Autonomic Network Engineering for a Self-managing Future Internet)
- INS (Identity and Access Management for Networks and Services)
- MOI (Measurement Ontology for IP Traffic)

A proposal currently exists for the setting up of a new ISG, ISG FIA on Future Internet Architecture.

In ITU-T, various study groups (SGs) develop standard recommendations. SG 13 (future networks including mobile and next-generation networks) is clearly very relevant to the Future Internet architecture. It brings together experts in communication network architecture. A focus group on future networks (FG-FN) has been established and is open to anybody, including non-ITU members such as academics and independent experts. Its work is proving helpful for developing future network virtualisation, energy-efficiency and identity recommendations. “Making standards is easier than making useful standards,” Takashi Egawa quipped, “and ITU-T does not want to produce only paper!”

Towards a draft model for the Reference Architecture

The European Commission has enabled a team of about five editors backed by a reference group of 10 to 15 experts, to actively invite contributions and comments from the community at large and draft a Reference Architecture model. This action will support the process of creating a ‘common’ model for the Reference Architecture, aiming for a first version in September 2010 and a second version to be presented at FIA Ghent in December 2010.

Discussions

The discussion in the last of the three architecture sessions tackled the issue of whether there was enough solidity and consensus among the current proposals to start the standardisation processes. If it is premature to start now, as some people think it is, then is there a danger that the process will stall? The resulting loss of interest would set progress back relative to what could in due time have been achieved.

The general view of the panel was cautiously optimistic. It will be a difficult, though not impossible task, and much discussion will be needed to reach a shared consensus. It is probably right to be starting now, the panellists thought, to take in continuous input and build up a critical mass of shared understanding over the next six to 12 months.

'Search' in the Future Internet: SME and Citizen Perspectives

Caretakers: Petros Daras (CERTH/ITI), John Domingue (The Open University) and Nozha Boujemaa (INRIA)
Chair: Petros Daras

'Search' means making the best use of available (human or machine-generated) knowledge to find a desired information artefact, even when the user request might be poorly formulated or unanticipated by the system.

The value of a search engine depends on how efficiently it manages all the knowledge at its disposal (i.e. how it is acquired automatically, enriched, structured, retrieved, filtered, interpreted) and how easily the information is accessed and understood by the end-user.

At FIA Stockholm we had a technology-based discussion, looking at how search might function within future content networks, the Internet of Things and the Internet of Services. This previous meeting raised several questions:

- **Given the polymorphic facets of the internet, what should a search environment look like**?
  - Are the open, federated search services a/the solution?
- **How can we implement/design search and discovery over a range of information artefacts?**
  - How can you integrate services, sensor networks and rich multimedia content?
  - What types of indexing/caching could support Future Internet search and discovery?
- **What types of representations (metadata) could we use to support Future Internet search across the heterogeneous resources?**
  - How would these be created and maintained?

Objectives

The session decided to discuss Future Internet searching from user and SME perspectives. The overall aim was to consider the requirements and research challenges raised by the application of search within different vertical application areas. The session also aimed to reflect upon the role that search technologies could play within Future Internet applications. The main areas of Future Internet application investigated included health, energy grids, city management, information and content/media. The aims of the session were to:

- identify the most important user needs and unsolved challenges that will require the development of new search technologies and techniques;
- identify potential Future Internet related applications that require interactive forms of ICT collaboration among users and providers.

Presentations

• Search in content/media: state-of-the-art and future challenges in multimedia search	Francois Bourdoncle	Exalead
• Search in multimedia health records	Martin Huber	Siemens
• Search and mine, with hindsight	Julien Masanes	European Archive Foundation

• Green IT challenges: uses-cases and related search technologies	Daniel Barthelemy	AMAP
• Smart cities and search applications	Cedric Ulmer	SAP

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-search).

Summary of Presentations

State-of-the-art and research challenges

François Bourdoncle covered the state of the art in multimedia search from the perspective of the QUAERO project and from an Exalead perspective. New functionalities described included the ability to:

- search for an image containing a face;
- find images with a particular colour hue;
- finding specific segments of a video which match a given phrase.

These types of searches are finding their way into today's applications and services. However, Bourdoncle went on to present a long list of research challenges. First, there are a number of problems associated with searching for and within specific types of media, requiring further research into:

- high-quality speech-to-text technologies;
- high-quality, general-purpose machine translation.

Other research challenges include the development of:

- image classification technologies;
- integration of semantic web technologies;
- real-time (i.e. online) algorithms;
- applications with good usability features
- solutions to search and navigate pervasive video;
- methods to index all content types (e.g, Flash, iPad apps, etc.).

Bourdoncle also covered the (partially historical) relationship between search and cloud computing (see below).

Search in multimedia health records

Martin Huber described methods currently being investigated to search for information in electronic health care records. Three different search approaches are being explored: text search, knowledge-driven semantic image searching and data-driven, similarity-based searching.

Huber said that text-based searching in medical records is already well progressed, using natural language processing techniques combined with semantic technologies. However, before useful applications can be developed the searches must reliably retrieve high-quality (i.e. relevant, up to

date, validated/trusted) and personalised information that is relevant to both professionals and self-informed patients.

“Search in heterogeneous health records requires information extraction from unstructured data,” Huber explained. “This can be supported by semantic and machine-learning technologies.” But he warned that this was no easy area with a quick fix, 'one-size-fits-all' approach. With respect to multimedia in medical searches, Huber advocated more targeted solutions.

Search and mine, with hindsight

Forget simple text searches with Google – search in the Future Internet will go up a gear, or several gears even. When it comes to mining the web you get some very important scale effects. For example, the accuracy of comparison tables (e.g. the height of the highest mountain in each continent) generated by mining over 14 billion pages is extremely high.

Julien Masanes suggested the following research challenges in this area:

- The building of an open, neutral and sustainable virtual observatory of the web for European research. This would require large scale crawling, storage and indexing of web data (10+ petabytes) and not just limited to text.
- The creation of a baseline distributed analytics services (large-scale information extraction, natural language processing, distributed and efficient processing and storage). We also need to standardise and define a baseline in this domain to create a platform for multimedia search engines, social media research etc. The analytics service could be used to research issues such as:
    - Hadoop-style abstractions over internet-wide repository/processing clouds;
    - the optimisation of data placement (partitioning and replication) for analytics;
    - distributed indices.
- Temporal indexing of the significant characteristics of networked content (from distribution to semantics). There needs to be a wide spectrum of research in information extraction, information retrieval, network topology etc.
- Make this infrastructure acceptable by society (respect privacy, transparency, IP rights, etc.).

Green IT

In this presentation Daniel Barthelemy highlighted the need for sustainable agricultural development and the conservation of biodiversity. He argued that conservation work and sustainability initiatives will be based on accurate knowledge of the identity, geographic distribution, production and uses of plants. He identified many research challenges such as:

- geometric/topological models for image-based visual plant identification;
- users' perception of interactive identification tools;
- active multi-class machine learning for multiple identification criteria;
- interactive similarity based visualisation and navigation methods;
- scalability issues related to image-based indexing;

- modelling the geographic distribution of plant species and communities.

Search in smart cities

Cedric Ulmer outlined SAP's vision for smart cities. “Smart cities are about smart citizens, embedded devices and prosumers,” he said. He suggested that the smart city is made up of smart city managers, smart partners and smart platforms (Internet of Things, Internet of Services). But where does search come into it?

Search is at the very heart of smart cities. With pervasive networks, including large sensor networks, CCTV, traffic monitoring, etc., the volume of data – in many different multimedia formats – will need to be stored and easily retrieved, possibly in real time. Ulmer presented numerous challenges for search in smart cities:

- **Security and trust:** search engines will need to aggregate vast quantities of data but retain individual privacy. There also need to be methods to measure trust in search results.
- **Speed:** citizens will adopt the tools if they are fast.
- **Simplicity:** smart city managers and citizens need search tools that can be mastered quickly and with little effort.
- **Relevance:** searching becomes inefficient if results are not relevant.
- **Interpreting results:** how do you visualise and display the results of a query and how to select the best/most appropriate visualisation mechanisms.

Discussion

The moderators Nozha Boujemaa (INRIA) and John Domingue (OU) centred the discussion around the issue of how to include users in the loop. “What search technologies can help citizens and answer urgent needs?” they asked. “Why is search so important to everyday internet access in the context of citizens' everyday life and what needs fixing within the next five to 10 years?”

The questions, comments and responses from the panellists and other session participants identified the following issues that need to be addressed in future research programmes:

- **Usability:** This is a key factor because it determines the extent to which search queries can be formulated and reformulated. Visualisation techniques also influence the usability of a search because they determine how much the search results can be accessed and understood.
- **Bespoke solutions:** There is no 'one size fits all' to searching, so there is a considerable opportunity for SMEs to exploit markets for niche search solutions. Even within a single application area it is clear that there are different search niches which can be filled by quite different tools. Within the medical sector this landscape has given rise to the emergence of SMEs, each one of which fills a specific search segment.
- **Cloud computing:** We can see that all the premier league cloud computing players – Google, Yahoo, Microsoft and Amazon – have strong ties to search. The key technical requirements for cloud computing have a significant overlap with the technical requirements associated with internet-scale search, namely: distributed storage (e.g. Google’s GFS), distributed computing (e.g. Hadoop, MapReduce) and scalable semi-structured databases (e.g. Google’s BigTable).

Bringing all the discussion together, the moderators concluded that search should not be seen as a service or a necessary function of the Future Internet, but rather as a key platform layer on which Future Internet applications can be built. The primary requirements for this platform are:

- trust in the search engine;
- speed and scalability;
- simplicity;
- high relevance of the retrieved results.

'Deploying' on Future Internet Research and Experimentation (FIRE)

Chair: Susanna Avessta
Session rapporteur: Anri Kivimäki, Dimes Association

The Future Internet is not going to be built overnight, nor are we going to ever experience an official or formal switch to Internet 2.0. The transition will almost certainly happen step by step. But that does not mean that the new technologies and infrastructure will need rigorous testing before they are deployed.

But how do you test something as big and complex as the Internet? There is no point experimenting at small scales; you have to prove that new architectures and technologies will work on the massive scale of the Internet itself. But you cannot wire up some kind of prototype Future Internet either.

This is where FIRE steps in – the name given to the global effort to develop test-beds and put prototype Future Internet technologies and architectures through their paces.

Objectives

FIRE is certainly evolving. From its initial phase, to build Future Interest testing facilities and run experiments, FIRE is now moving into a second phase, with more extended and diverse research and support facilities on offer. This session was designed to reflect this transition: the first half of the session aimed to update delegates on the current state and results of FIRE; the second half looked forward, with sneak previews of FIRE 2.0 and plans for phase 2 projects.

The current set of FIRE facility projects covers a major part of the basic structure. The second wave projects will bridge some gaps, moving the experimental facilities focus to higher layers, i.e. to services (BonFIRE, TEFIS), and also to the converging technologies in terms of transport and taxonomy (Openflow). The second wave of projects also strengthen the connections to the physical world, with e.g. test-beds dedicated to cognitive radio (CREW). All facility providers are preconditioned to work for a common interface and a portal of FIRE experimental services, and are open for all research customers in a transparent way. FIRE-wide federation and sustainability concerns remain to be tackled by the new projects and customers. Aspects needing further contribution from the community are the common high-level federation model (in practice, operational terms), standardisation, business cases for the addition of commercial test-beds as well as brokering and common management of the FIRE test-beds.

Presentations

FIRE mid-2010		
• White paper on experimentally-driven research	Anastasius Gavras	Euresom
• FIRE portfolio analysis	Scott Kirkpatrick	Hebrew University of Jerusalem
	Dirk Trossen	University of Cambridge
	Jerker Wilander	Dimes
• FIRE internationally	Serge Fdida	UPMC-LIP6
FIRE 2.0		
• FIRE plans	Jerker Wilander	Dimes

• Test brokering with Teagle	Florian Schreiner	Fraunhofer Fokus
• Creating a European market for testing and experimentation facilities	Esteve Admiral	Universitat Pompeu Fabra
• Including end-users	Michael Nilsson	CDT

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-fire).

Summary of Presentations

FIRE burning bright

The concept of multi-disciplinary, experimentally-driven research has been embraced by the FIRE community. There are 12 research and facility projects that have been running for nearly two years and another set of FIRE projects has been selected and will be launched in autumn 2010.

Anastasius Gavras presented an analysis of all the European activity in this domain. "FIRE covers systematic experimentation for which theoretical experimental models do not exist," he stated. "FIRE is/can be seen as a crystallising experiment – an experimental environment, but one in which experiments will run."

Gavras said that current FIRE initiatives were focusing on four priority areas: verifiability, reliability, repeatability and reproducibility. Progress has been encouraging and the first milestones of the first batch of projects are now being achieved.

But it is important to maintain a distinction between testing and experimentations. ETSI has published its definitions, but there needs to be some agreement on where to draw the line between these two activities.

Scott Kirkpatrick has conducted an analysis of the portfolio of FIRE projects and activities. He outlined how the earlier calls of FP7 are being met; he also identified the research gaps that still remain and what needs to be addressed in future calls. His analysis of the FIRE landscape highlighted several points:

- **The "Wise men's report":** This report points out what is missing in FIRE research and presents a suggested collaboration structure (see figure below). The report highlights the big differences between FIRE researchers and FIRE end-users. It also identified several technical issues related to FIRE.

- A total of 15 FIRE projects were funded through the second call of FP7, but together these projects still do not produce a "federated FIRE facility" – some of the projects will only produce architecture that will not be sustained after the project ends. Nevertheless, the good news is that there are now FIRE users and some good use cases are now available too (e.g. FEDERICA FP7 project; PlanetLab Europe).

- The FIRE community is actively discussing these issues and working to improve its understanding of federation and topics such as peer-to-peer access, heterogeneous federation; and hierarchy creation. The concept of federation is much debated (trying to establish the drivers and main benefits for federation, including cost savings and prototype federation in the future). Legal and operational difficulties also need to be overcome.

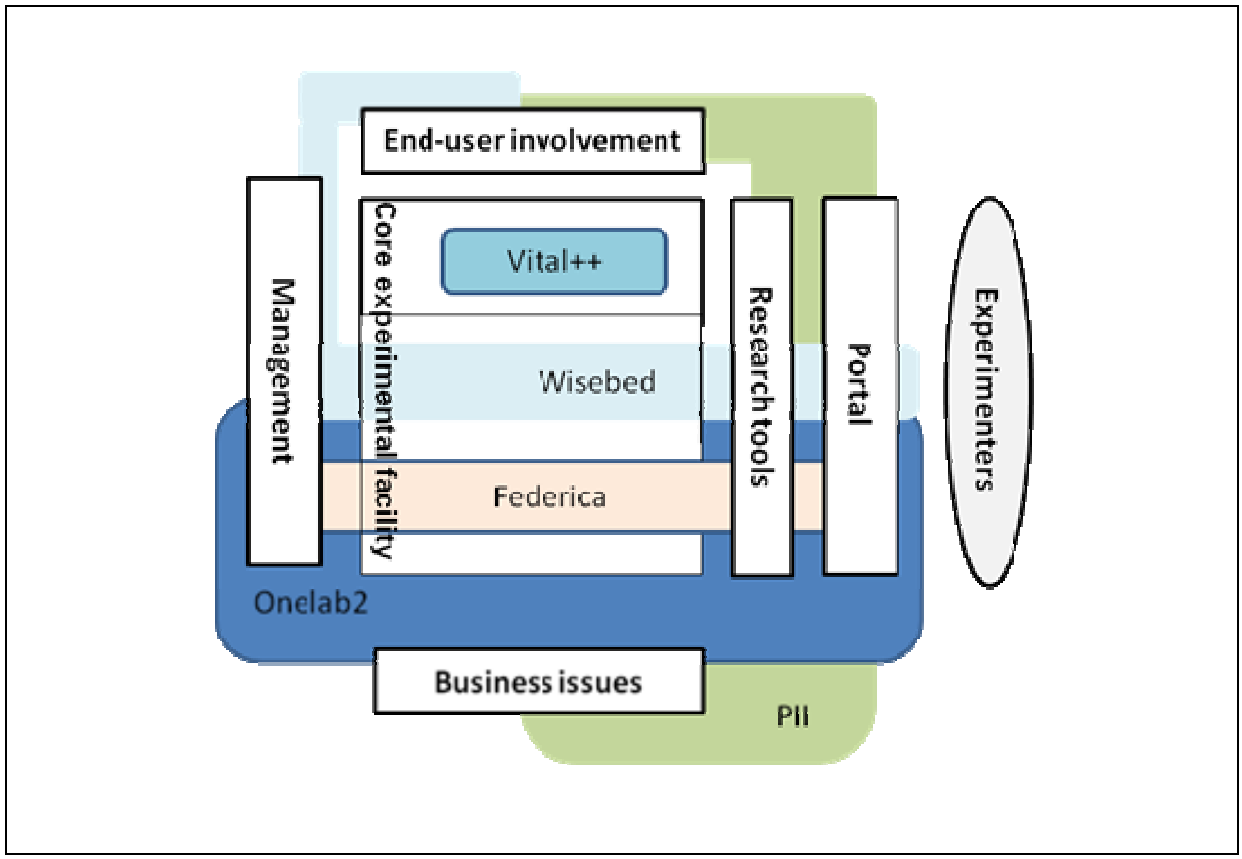

FIRE Research Collaboration

Serge Fdida reminded the participants that FIRE is not just a European initiative. The nature of the Future Internet means that test-beds have international significance. Fdida remarked that:

- the development of test-beds is extremely important;
- the use of test-beds is currently quite limited;
- Europe is unable to "think globally but act locally" – there is global competition, but best practices have to be shared beyond the EU;
- In the US the GENI initiative is a continuous process of development and deployment of different technologies;
- As its international role continues to extend, China has expressed its interest in cooperating with projects (e.g. CERNET2, 3TNET);
- ASIA FI is a Future Internet R&D coordination initiative including education and events (see http://www.asiafi.net).

FIRE has strong links beyond Europe. The everyday cooperation with American and Asian counterparts gives Europe's activity an important and influential position and connects Europe to the global scene. Indeed, the global development of experimental facilities and their federation is strongly anchored in Europe.

Fdida highlighted some important international events:

- Joint Workshop on Federation (Princeton, May 2010);
- annual joint EU/Japan seminar in October 2010
- joint Brazil /EU/US workshop on "Test-beds for Future Internet".

Workshops are being organised on a regular basis by interested parties.

Fdida finished with some remarks about federation: "Federation is also about managing complexity. The reasons for federation revolve around mutual benefits, but in practice they are quite different. There are many forms and standards for federation. When building a facility, you have to think about continuity!

"Creating something sustainable depends in part on funding mechanisms. There are different funding mechanisms in different places and intensive international collaborative effort is required. And we must not lose sight of the main driver for federation: users."

FIRE ahead

Four speakers gave the session participants a glimpse of how FIRE will develop in the near future as it evolves into a more mature "FIRE 2.0". A second wave of FIRE projects aims to bridge some gaps between the first round of projects, moving the focus of experimental facilities to higher layers, i.e. to services (BonFIRE, TEFIS), and also to converging technologies in terms of transport and taxonomy (Openflow).

The second-wave projects also have strong connections to the physical world (e.g. with test-beds dedicated to cognitive radio (CREW). Among these projects, all the facility providers are expected to work towards a common interface and a portal of FIRE experimental services which makes them open and accessible to the Future Internet research community.

FIRE-wide federation and sustainability still face challenges, which the newer projects will continue to tackle. The FIRE community still needs to work out how federation can be achieved in practical and operational terms, looking at issues such as standardisation, the business case for the development of commercial test-beds, brokering and common management schemes.

Jerker Wilander offered the session participants a top-level analysis of the current FIRE R&D portfolio, focused on the five integrated projects that were funded through FP7 Call 5. He remarked that these projects overlapped in certain areas and called for even closer collaboration and coordination between them.

Wilander said that these projects need to get together to discuss which items and resources they could share and joint strategies for running experiments. For example, all the projects have requirements for experimental data storage; they could perhaps also develop a joint portal. These projects also have a lot to do in the areas of business and policies. "If we are going down the path of federation, then we need to be sure that there is only a single agreement between us," Wilander remarked.

He also suggested that the Future Internet community should analyse when it is most appropriate/optimal to use experimental research because it is still so difficult to actually build a functional experimental facility – many FIRE projects are still not ready to open up to users. The construction of test-beds is long-term exercise and costly.

Wilander noticed that researchers often use PlanetLab's connections and facilities, but this is not a European facility. Why is this happening? "Ultimately, the whole point of FIRE is to build test-beds for testing novel architectures. Fortunately, that is exactly what most projects are working towards, but it does take a long time."

Wilander was disappointed to discover that most projects do not spend enough time in user-facing tests. "It seems that it is going to be a long process to get users involved," he concluded.

Florian Schreiner introduced the novel concept of a brokering service for test-bed facilities. The Teagle project is analysing how to federate facilities, for example Federica with other test-beds in development through different projects. Teagle is developing a service domain system targeted at the research community. The Teagle service provides a facility for large-scale testing and experimentation by federating test-bed resources (see http://www-fire-teagle.org).

Europe's Living Labs (LL) concept has several lessons for the FIRE community, said Esteve Almirall. In a lively presentation on Living Labs, Almirall suggested that FIRE researchers should think of the Living Labs as open innovation intermediaries. Users are central in the LL approach because the users are the ones who generate the innovations; innovation is a societal process. Sophisticated users, such as those requiring Future Internet test-bed facilities, are a key element in the LL environment.

So how do you get end-users involved in FIRE? Michael Nilsson provided delegates with a few tips. First, he suggested projects should decide on what kind of users they want. However, there is no rigorous method available for making this decision. Therefore, more discussion is necessary within the FIRE community about users:

- What are the users' roles?
- What level of control could they have?
- How much responsibility do they want/should they be given?
- What are the most suitable FIRE structures for users?

"The development of test-beds will undoubtedly benefit from user involvement," Nilsson concluded. "Involving users does create new challenges, but we must learn to listen."


The European RTD Framework Prog
From Economi
to Sustainabili
IV Conference on the FP7 in Spain
The European RTD Framework Programmes
From Economic Recovery
to Sustainability
IV Conference on the FP7 in Spain
Auditorio
From Economic Recovery
to Sustainability
From Economic Recovery
to Sustainability

Applications of the Future Internet

What Can the Future Internet Mean for Smart Energy?

Chair: Pierre-Yves DANET, France Telecom/Orange Labs

Scientists, economists and policy-makers are calling for CO_2 emissions targets of at least 20% below 1990 levels in 2020. It is widely accepted that ICT could help to reduce our carbon footprint in three different ways, specifically by:

- using smart ICT applications to reduce the carbon-footprint of many everyday activities and technologies (it has been estimated that ICT may help other sectors to reduce their CO_2 emissions by 15% in 2020 which is a significant step towards reaching the 20% target);
- enabling green energy (i.e. the generation and distribution of energy without CO_2 emissions);
- reducing its own energy consumption and carbon footprint.

In the specific area of energy efficiency, ICT has a role to play in two primary areas:

- **Network efficiency:** designing networks and the telecom infrastructure to minimise energy consumption.
- **Efficiency applications:** helping people reduce their energy consumption through innovative internet-based services.

Objectives

This session on smart energy aimed to develop a common vision in the Future Internet community regarding the ways in which the Future Internet could help society to save energy, either through reducing consumption within the telecoms/internet infrastructure itself, or by helping individuals and businesses to reduce their consumption through internet-enabled tools and applications (e.g. smart metering, home automation, e-applications, tele-services, etc.). The session was designed not just to identify research areas where ICT could be deployed to decrease energy consumption, but also to see how existing research could also contribute to the architectural design of the Future Internet. There was also scope for identifying research gaps and proposing new project ideas.

Presentations

Smart energy panorama

• Energy utility vision	Asier Moltó Llovet	REE
• Smart grid vision	Duncan Botting	Smartgrids ETP
• Open challenges enabled by the Future Internet	Mikhail Simonov	ISMB
• M2M standardisation activities	David Boswarthick	ETSI M2M

Socio-economic analysis

• Socio-economic perspectives on smart energy	Gabriella Cattaneo	IDC

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-energy).

Summary of Presentations

It's a grid, sir, but not as we know it

Asier Moltó Llovet from Red Electrica De Espana outlined a compelling case for smart management of the electricity grid.

European Energy Policy is driven by the 20/20/20 objectives: a 20% reduction of greenhouse gas emissions, 20% of energy consumption from renewable sources and a 20% reduction in energy consumption by 2020. National policies in Europe are therefore oriented to comply with these targets.

During the past decade, renewable production has exploded in European countries. But for systems operators, this has created a major challenge: data in control centres has multiplied and new tools for forecasting and production management need to be developed. "Fortunately today we can talk about this issue in terms of success," Moltó Llovet asserted. "ICT has been the enabling technology."

But the next decade will present challenges in the area of demand-side management, where the grid must account for customised management of millions of loads in different consumption sectors. One response to this challenge is for the internet to provide consumers with access to information on their own energy consumption and production, for example with hourly energy prices and best practices for sustainable energy consumption.

Duncan Botting, vice chair of the SmartGrid European Technology Platform, gave a slightly different view on smart grids. Indeed, he suggested that smart grids could only become reality with the Future Internet.

Many stakeholders exist in the current market supply chain and new ones are being added all the time – the grid is extremely dynamic and complex. Customers are now becoming suppliers; aggregators (virtual power plants) and demand-side ancillary services are some of the possible new players that will emerge. Utilities are no longer the "command and control" entities they used to be. "The customer is now becoming the main controller," Botting remarked. "And it is not only consumer behaviour that will have to change; much of the market will need to change including legislators, regulators and other market based stakeholders."

Botting outlined some important features of smart grids, based in part on some of the preliminary solutions studied by the SmartGrid ETP:

- Utilities are not the custodian of the smart grid design, they are a component as much as the ICT is.
- R&D is required to develop a cost effective system for communicating with the millions of sensors that are necessary for a true smart grid.
- Innovative ICT solutions will be needed to provide affordable and reliable application solutions.
- Providing data/information to end-users (i.e. smart metering) is not the same as a smart grid. If price signals are to be the only innovation then a smart grid will not develop. Fully automated and real-time demand-side participation will also need to be integrated into the system if it is to become a true smart grid.
- Generation is only any use if it is connected to end-users – often generation is considered as the only element to secure supply, and distribution networks tend to be ignored. A smart

grid must deliver flexibility and options for future architectures, otherwise security of supply will be compromised.

- Smart grids rely on four important and parallel activities: technology, economics, and environmental and cultural solutions. ICT will provide a contribution in each of these areas and be a key enabler of the future smart energy grid.

Open challenges

The Future Internet plays a key enabling role in the development of smart grids and dynamic energy management. But there is plenty of work to do before the prototype grids in development today can be rolled out en masse. The main challenges are:

- the integration of interoperable grids into a system-of-systems;
- the integration of hybrid electric vehicles as dynamic nodes into the grid;
- the development of e-energy businesses (for example involved in online, real-time energy trading);
- the development of new control paradigms to optimise and balance individual and collective interests;
- the refinement of forecasting models to optimise supply and demand;
- the scale of new smart grids and the volume of data and nodes they will have to manage;
- the need for grids (including national grids) to work independently of one another, but also in collaboration with others.

Mikhail Simonov gave four concrete examples of where the Future Internet could help in the field of smart energy:

- A wider use of photovoltaics could increase generation from renewable sources beyond 10-15%. But forecasting using real-time data must improve, especially at a high resolution, so that the grid can be managed actively by anticipating peaks and troughs in generation at a local level (e.g. because the sun goes behind a cloud).
- By making consumers aware of their consumption, but also using human-machine interfacing to complement smart metering so that users can see how their consumption is also being automatically managed.
- The volume of real-time data can help loads be managed in a proactive manner – variations in consumption can be anticipated and supply managed actively.
- Energy could be stored in grids by integrating electric vehicles into the grid (vehicle batteries could provide energy in peak demand and be charged during periods of lower demand). This integration would add a whole new set of energy services into the mix, and offer a very different way to balance the loads.

Machine-to-machine standardisation

The smart energy grid will require a high degree of standardisation so that a wide range of different devices and technologies can be integrated and controlled automatically, said David Boswarthick from ETSI.

ETSI is highly active in domains such as machine-to-machine (M2M) standards, smart metering, wireless, fixed and power line data transmission, as well as many radio frequency technologies. Such areas are essential building blocks for the realisation of the smart grid and eventually, the Internet of Things.

Boswarthick provided an overview of the initial work already completed by ETSI TC M2M, and how it may fit into the standardisation required for the smart grid. He also provided an update on the ongoing work for the European Commission's mandate on smart metering inter-working (M441).

A socio-economic perspective

Technologies do not enter a vacuum – their success or failure is ultimately determined by the behaviours and attitudes of the people who use or stand to benefit (or suffer) from them. Gabriella Cattaneo from IDC took quite a different slant from the preceding presentations.

"The Future Internet promises a new wave of radical innovation," she noted, "leading to new thresholds of pervasiveness of information infrastructures in the economy and society. It will deeply transform production, distribution and usage patterns of ICT. Technical innovation is an important driver of this transformation, but by no means the only one."

Cattaneo explained that the pace and nature of the development of the Future Internet in Europe will be influenced by general economic growth, the availability of investments for ICT innovation, the effectiveness of ICT investment and the research policies instigated by governments. Success or failure will also depend on the ability of businesses to exploit and embed ICT innovation in their business models, and the willingness of consumers to adopt new services.

Cattaneo introduced a newly launched study into the potential socio-economic impacts of the Future Internet. The study is commissioned by DG INFSO to prepare the ground for the Future Internet PPP. It will explore these main socio-economic trends and assess Future Internet development under a number of different scenarios.

Preliminary work suggests that smart energy is certainly one of the most promising areas for Future Internet applications. According to IDC research, ICT is an important tool for the low carbon economy. Intensive use of ICT technology in the G20 countries could reduce emissions by over 25% annually by 2020 compared with 2006 levels. These potential emissions savings would be highest in the energy creation and distribution sectors, particularly through the integration of significant renewable energy resources into energy distribution using smart grids. There are already positive signs of dynamism in the use of ICT in the energy sector in Europe: according to IDC Energy Insights, spending on intelligent energy grids in EMEA (Europe, Middle East and Africa) will reach $8 billion (€6.4 billion) in 2010.

Discussion

Moderator: Roger Torrenti, Sigma Orionis

The speakers outlined what they thought were the R&D priorities in the area of smart energy. They suggested more work was needed to:

- develop ICT skills so that end-users could make the most of ICT tools to help them improve their energy consumption behaviour;
- integrate power grids with ICT networks;

- ensure interoperability between technologies, systems and components through standardisation;
- manage large and real-time data sets;
- improve the energy efficiency of ICT networks;
- deploy sensors and smart energy grid communication in a cost effective manner.

Next Steps

Pierre-Yves Danet outlined the next steps for the FIA Smart Energy working group:

- Collect the names of people interested in the subject and set up a mailing list of these people.
- Launch a survey to collect ideas from participants about FIA involvement and activities around smart energy (end May).
- Analyse/classify ideas (end June) with a view to establishing an editing group in this field.
- Identify which ideas are already covered by existing projects.
- Communicate to the FI PPP and Future Internet ETPs the smart energy topics that are not covered by existing projects or calls so that these ideas can be considered for future calls. (end September).
- Publish a white paper covering 'the story so far' (in preparation for FIA Ghent). The purpose of the white paper will be to help existing smart energy research projects also contribute to aspects of Future Internet design. The publication will also recommend ways in which projects involved in Future Internet design and architectures could also take into account the specific requirements of the smart energy sector.

What Can the Future Internet Mean for Smart Health?

Chair: Paul Moore, Atos Research & Innovation

The challenges that Europe's health care systems have to face over the coming years are massive, but the quality of health care that we receive depends on us overcoming numerous hurdles. The increasing demands of people for good quality health care, the expanding ageing population, the rise of patients with one or several chronic conditions and the lack of sufficiently qualified health care workers are just some of the very real issues and challenges that lie ahead.

At the same time, the Future Internet holds the promise of revolutionising health systems and the way health care is delivered. The Future Internet will offer new approaches and opportunities that could help to overcome most of these challenges. Future Internet-based health services will provide better assistance to patients, reduce inequalities, improve health care delivery to patients, and minimise human errors and unnecessary duplication of work. In the longer term, Future Internet applications will help to reduce the costs of health care, reduce hospitalisations, improve patient outcomes and facilitate disease surveillance.

Last but not least, the Future Internet will facilitate the delivery and exchange of health information and training to both patients and caregivers. This will help to reduce health inequalities not only across countries, but also between rural and urban populations. Through better access to information citizens will begin to change their behaviours and attitudes towards health and adopt more healthy lifestyles, which will help to prevent disease.

The convergence of various scientific advances in different fields has lead to the development of a new concept called 'smart health'. This ICT-based approach aims to overcome not only financial and population pressures (relating to cost, access and quality), but also to set up the basis for clinical and biomedical data integration that will boost the translation of research from laboratories into clinical practice.

These huge challenges are significant catalysts in the development of Future Internet-related technologies. Indeed the Future Internet will not just help to overcome these challenges, but also radically alter the framework in which we conceive health care today. Therefore, the area of smart health is a one of the most advanced sectors for Future Internet development – scenarios embrace all five pillars of the Future Internet and will serve as a driver for the deployment of Future Internet-based services and applications.

Nevertheless, the promise of these services depends on whether R&D can deliver solutions that mitigate or resolve the many challenges and help to advance a health "infostructure".

Future Internet-based solutions, while exciting and promising, also present new ethical, legal, technological and even financial challenges, particularly related to:

- the adoption of widely acceptable standards for interoperability;
- the provision of a regulatory framework that will overcome jurisdictional boundaries;
- the allocation and stimulus of public and private investments;
- the development of safe and secure data infrastructure that takes into account the special requirements related to privacy, security, confidentiality and data integrity.

Objectives

The aim of this session was to explore the potential of Future Internet technologies to deliver improvements in health care services so that European citizens can continue to enjoy world-class medical care in the context of a dynamic and changing society. The session was designed so that delegates could discuss the impact of demographic changes, the increasing mobility of citizens (which requires greater interoperability among national European systems and better ways to share health data) and the growing requirement for patient empowerment and patient-centric health care management (i.e. personalised treatments and prevention).

Presentations

• The e-health pillar in the FI PPP	Blanca Jordan	ATOS
• The Virtual Physiological Human: a metaphor for the future internet?	Marco Vicceconti	Instituto Ortopedico Rizzoli
• E-health, the business case of prevention : lifestyles and wellbeing	Dr Alberto Sanna	Scientific Institute San Raffaele -
• The end-user perspective for the Future Internet and e-health	Dr Antonio Campos	CTIC Foundation

Summary of Presentations

State of the art

The traditional e-health care model is depicted in the figure below.

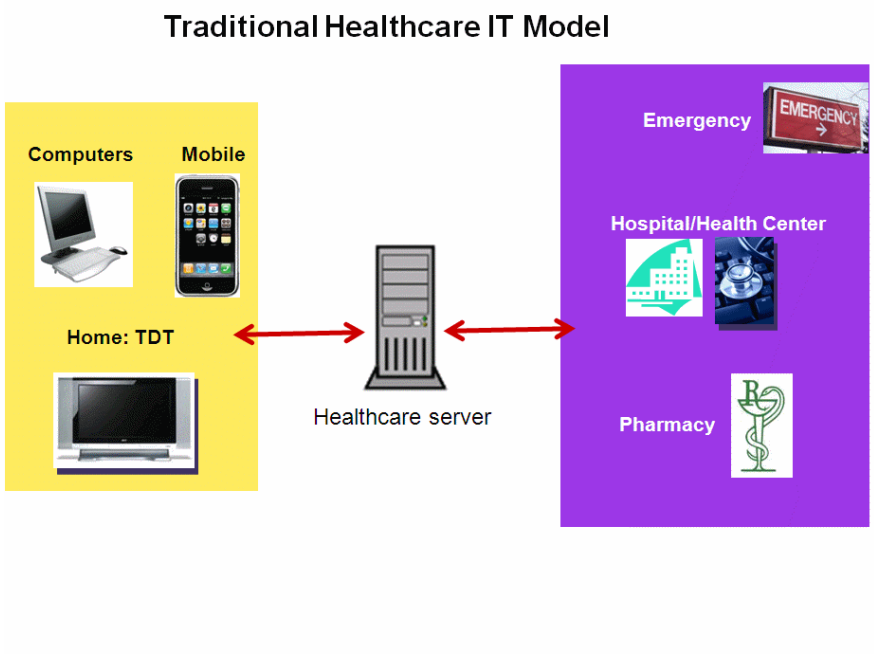

TDT: terrestrial digital television

Traditional Health Care IT Model

From its start, the main goal of e-health services has been to overcome the time and distance barriers that separate the health care provider from the patient, but in a way that does not affect clinical practice, which focuses on pathology rather than the patient.

Numerous innovations are already supported: computer-based patient records, remote consultations, isolated clinical information systems, computer-based decision-support tools, mobile and wireless terminals, and new ways of distributing health information to physicians and patients. But the wider adoption of e-health services has been hampered by technological, regulatory, ethical and other barriers, including the suspicion of patients, poor uptake from the medical community and a lack of financial investment.

The new care model

One of the key challenges for modern medicine is to advance our understanding of the complexity of biological systems and how disturbances lead to the progression of illnesses. Integrated health approaches that will tackle a holistic vision of a patient are vital, first to understand the risks and specific interrelationships between illness causes and individual genotypes and, second, to provide individual therapy. More focus on individual patients is gradually leading to what is called **personalised medicine** which takes account of numerous individual factors, but requires much more health monitoring and risk management.

Today, with increasingly powerful technologies and ICT infrastructures that support the Future Internet paradigm, we are able to provide:

- the needed technology for ubiquitous communication between biosensors and remote monitoring devices (Internet of Things);
- the safe and real-time transmission of huge amounts of data;
- data storage and computer capacity for processing large amounts of data;
- capacity to develop multi-scale models able to provide knowledge rather that a mere flood of information (thanks in part to grid/cloud infrastructures supported by high-speed connections);
- services (using Web 2.0, or the Internet of Services) that can actually deliver useful functions.

Nevertheless, people – patients and health care professionals – must choose to adopt this technology. Citizens will not sit back and simply see themselves as passive beneficiaries of technological innovation. So it is important that they can also participate in their own health management, using the Future Internet to share their experiences, advice and feelings with peers and health care workers.

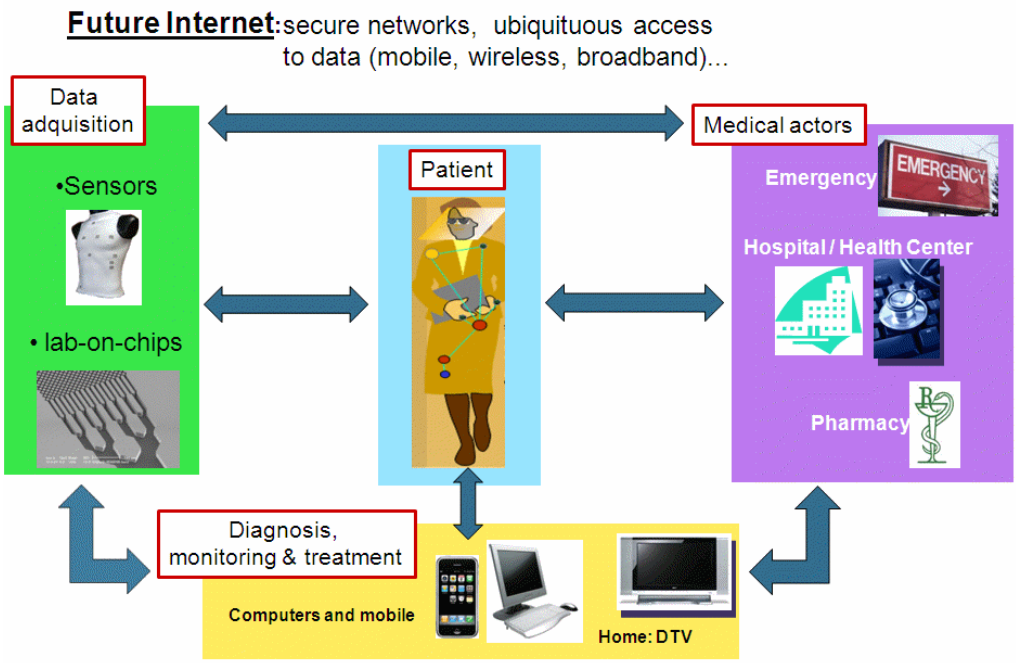

The Place of the Patient in Smart Health

Applying the Future Internet to health opens a wide range of possibilities for applying the most innovative technology in ICT to resolve multiple health challenges as well as to improve the quality of life for European citizens. This new paradigm will be built based on the new biomedical scientific advances wrapped by the Future Internet framework.

This new scenario means a re-engineering of business processes and models in the health care sector. This will change not only the way that health care services are run and provided, but also our understanding, as citizens, of how the health system will fulfil our expectations of having improved care.

A call for research

In the context of the discussions about the Future Internet and the FI PPP in this area, a working group has been created to discuss smart health and the Future Internet. A number of different papers and internal working documents have been published and a position paper is currently in production. The working group is discussing how to proceed with the different proposals that are on the table, but there is a concerted effort to forge collaborations between the most relevant actors and build on Europe's excellence and experience in this area as the research community prepares for the first FI PPP call for proposals.

A competition scheme is being considered to provide a level playing field for the research community (including the members of the working group) across the EU. If this idea seems viable, a public website will be launched with details about how this scheme could work.

E-health, with people at the centre

The session chair, Paul Moore from Atos Origin, presented a brief introduction about the Future Internet not only from a technological point of view but also looking at potential applications, with

special emphasis in the area of health. He mentioned that the Future Internet is not only a technological challenge, but also a political and social challenge; a future paradigm that aims to be the solver of existing problems is bound to be a political "hot potato" so there will be considerable support and interest from politicians at a very high level.

Moore stated that, technically, the current situation is that the internet continues to expand in all directions and at all levels. But the pace of growth may mean a decrease in the quality of the final service. At the same time, numerous problems are coming to light: issues with copyright, intellectual property and identity theft, for example.

Moore thought there was a good opportunity to boost the European ICT industry, making it a leader rather that a follower. "We have the momentum, we have the opportunity, we have the technology and the political support, so what should we do?" he asked. "We have to take advantage of this situation and not be left in second place. Europe should be able to [put] itself this time in a leading position and seize the advantage to create new opportunities for our economies."

Moore pointed out that, if we combine the Future Internet paradigm with the new model of personalised medicine, we end up with a model that places citizens at the heart of a technology-powered health system. Within this new framework, clinical practice and biomedical research will be improved; data acquisition, information management and treatment will be easier and much more complete than currently, and this will have a direct impact at every stage of medical intervention: prevention, prognosis, diagnosis and treatment. He suggested a new way to look at data (for example, the data from biosensors), not as merely 'patient data' but "data for this specific person with a specific clinical history and with a specific set of measured parameters".

Collecting physiological data in combination with so-called 'omic' data (genomic, phenomic and proteomic) will contribute to appropriate models that could effectively simulate the biology and/or pathology for individual patients. This information will have to be transmitted and exchanged (from the sources to the treatment and backwards), but the Future Internet will ensure that this travels in a safer way and is retrieved in an enriched manner. The analysis and the visualisation of the information will also be improved because of new techniques to produce and display content.

Moore concluded that the area of health could provide a major fillip to the development of the Future Internet and stimulate advances in other application domains.

Health: a pillar of the FI PPP

Blanca Jordan continued to establish the case for smart health as a pillar within the FI PPP. She noted the current momentum in Future Internet R&D in this domain and she agreed with Paul Moore that now would be the time for action for positioning Europe ahead of others in the global economy. She expressed her idea that the development of the Future Internet should be targeted for maximum impact, focusing on areas that would deliver real added value and efficiency and all within a trusted framework.

Jordan proposed that applications should work to create a "virtual intermate" which she described as an internet-linked friend and health expert. The virtual intermate would be much more than an aggregator of existing or improved tools and services; instead it would actually create a system that could replace existing processes, not just in health but also in transport, energy, environment, leisure and tourism.

"Health is the perfect domain for applying the Future Internet paradigm," Jordan offered, "since health is a vertical domain with a high impact on society. The inclusion of ICT in the health sector is really slow at present, even though the Future

Internet R&D activity in this domain is relatively mature." She argued that research initiatives lack clear action related to innovation, application and how the Future Internet could radically transform the very way we 'do' and manage health care in Europe.

Virtual biology

Marco Vicceconti took the discussion to a much more personal level, introducing participants to the European research effort on the development of the Virtual Physiological Human (VPH). Vicceconti defined the VPH as "the vision of integrative biomedical science – shared by a number of distinctly innovative new approaches including systems biology, multi-scale modelling and the physiome." However, this vision will become practically possible only when an entirely **new framework of methods and technologies** has been developed for investigating organisms as single systems.

The VPH "infostructure" would include the following elements:

- **A secure health data cloud:** allowing data to be exploited for information.
- **Personalised models:** taking information and turning it into knowledge.
- **A web of predictive models:** the exciting step that turns knowledge into wisdom.

The development of predictive, personalised models of patients adds weight to the 'business case' for the importance of the Future Internet in preventative medicine. In his presentation, Alberto Sanna argued that lifestyles and wellbeing are an important aspect of smart health. He gave many examples of how internet-based technologies could be exploited to identify trends or inform individuals about how their lifestyle choices could have an impact on their health.

The presentations were rounded off by Antonio Campos who gave a compelling account of the expectations and desires of end-users for the Future Internet and smart health. "The basic goal of smart health from the end-user perspective," he stated, "should combine health and technology to improve both the number of citizens covered and the services on offer. Moreover, the Future Internet promises to improve the quality of these services whilst making the actual technology transparent."

What Does the Future Internet Mean for Enterprise?

Caretakers: Man-Sze Li, IC Focus (FISO, FISE), Stefano De Panfilis, Engineering (EFII), Sergio Gusmeroli, TXT (FISO), John Kennedy, Intel (FISO, RWI), Jean-Dominique Meunier, Technicolor (NEM, FCN), Michele Missikoff, CNR (FISO)
Chair: Man-Sze Li

This session was motivated by the outcomes and follow-up activities of the FIA Stockholm session on enterprise. The FIA Stockholm session debated the question "What will the Future Internet deliver for enterprises?" The debate continued via open consultation on the FIA Enterprise public wiki for three months after the event, leading to a set of tentative priority topics for further investigation.

The enterprise session at FIA Valencia focused on the following themes, emphasising the central role of innovation in catalysing Europe's economic recovery and pursuing future growth and prosperity:

- **Vision:** innovation as part of routine business and smart enterprise.
- **Business models:** business models to support new value propositions and drive new business values.
- **Future Internet systems:** next-generation systems that will support enterprises to innovate and thrive in the post-crisis landscape.

Objectives

The objective of the FIA Valencia Enterprise session was to further advance the good work of FIA Stockholm by:

- stimulating additional debate on several of the priority topics;
- attracting new input from a broad spectrum of FIA stakeholders;
- reaching agreement (where possible) on the direction of Future Internet research that would motivate, enable and support enterprises, including SMEs, to achieve their business aspirations and objectives, thereby creating a positive impact on the economy and society.

The multi-disciplinary nature of the enterprise domain was emphasised in the FIA Stockholm discussion and subsequent consultations. Accordingly, the FIA Valencia session was targeted at all interested parties in Future Internet R&D. It was specifically organised to encourage energetic debates and active participation in advance of, during and after FIA Valencia.

Presentations

• What can Future Internet technologies deliver for enterprises and what can enterprise as a research domain contribute to Future Internet research and the FI PPP?	Stefano De Panfilis	Engineering / EFII
• The Future Internet – SAP's vision in an enterprise context	Thomas Michael Bohnert	SAP Research CEC Zurich
• What does Future Internet mean for enterprise?	Miguel Borras	Antara Information Technology

*The presentations from this session are available on the FIA Valencia website (*http://bit.ly/fia-valencia-enterprise*).*

Summary of Presentations

Background

Recent research and ongoing consultations suggest that the new drivers for prosperity and growth will come from innovation and from using resources better, where the key input will be knowledge. Future competitiveness will be driven by factors far beyond conventional economic dynamics. Instead, the focus is increasingly on conserving and making more effective use of energy, natural resources and raw materials; it is also on social cohesion, tackling unemployment and fostering social inclusion. European enterprises have the opportunity to thrive in this post-crisis landscape by means of environmentally and socially responsible business innovation and creativity.

Enterprises of the future are envisioned to be ever more open, creative and sustainable – and they will become smart. Smart enterprises will reap competitive advantage through innovation. Innovation occurs at many different levels. It includes not only product, services and processes, but also the organisational model and the full set of relationships that comprise the enterprise's value chain.

The unifying glue is the enterprise's business model. New value propositions and business models will arise, generating new demands for and from ICT. The Future Internet will best support and enable enterprises by directly meeting the requirements that are determined by the business models. It will give enterprises a new set of capabilities not possible today. Specifically, the Future Internet will enable enterprises to innovate through flexibility and diversity in experimentation.

The time has come to consider what the Future Internet will be able to deliver to and deliver for future enterprises. Whatever they might be, one thing is certain: future systems will not be based on technologies in silos. Instead, they will almost certainly reflect the "DNA of the Future Internet" – simple to use, adaptable to dynamic needs, customisable to highly specialised markets, affordable to small budget holders. These systems will also have the required technical attributes of accessibility, reliability and interoperability. They will be enterprise-centric rather than technology-centric.

The availability of such systems should lead to an explosion of adoption, particularly by SMEs. The DNA of the Future Internet would become the building blocks for potentially an unlimited array of value-added enterprise applications.

Not business as usual

In her welcome, Man-Sze Li described the context of enterprise research within the ICT research landscape. She also presented the result of the open consultation on the research priorities for the enterprise domain which was launched at the FIA Stockholm enterprise session. While a range of opinions had been expressed, there was unanimity among all contributors that the top priority was more research on "business models and relationships". The majority of the contributors also thought that it was important to define more clearly the different Future Internet research streams, and especially where enterprise research should be positioned within these strands. Contributors also called for studies on the impact of the Future Internet on enterprises particularly in relation to SMEs.

The overall message was:

- for enterprises, "business as usual" is over;

- Future Internet research, including the work of the forthcoming FI PPP, needs to produce a positive, high impact on society at large, and bring concrete benefits to European enterprises;

- there is a need to understand, define and specify the "DNA of the Future Internet" (not just its technical architecture, but also its values, properties, attributes, and ultimately standards).

A mutual relationship

Stefano De Panfilis recalled the starting point of Future Internet research and described three main groups of beneficiaries: citizens, (dynamic) communities and enterprises. For the enterprise, the Future Internet promises significant opportunities, at least indirectly. Who would have believed that an app for 'painting ' with your fingers on an iPhone could make a developer a millionaire!

De Panfilis discussed also the potential direct impact of the Future Internet on enterprises – so long as enterprises evolve, and are willing to change and embrace the new rules that the Future Internet will inevitably create. To help enterprises make the most of the Future Internet, supporting technologies are needed, which De Panfilis called "XaaS". He then described the vision of the European Future Internet Initiative (EFII) regarding the Future Internet Reference Architecture (see preceding chapter on "Architectures of the Future Internet"). Such an architecture is itself based on generic enablers, requiring high flexibility and standardisation, based on common needs, and requirements for easy adoption and future evolution. He concluded by throwing out a suggestion to the delegates: why not have a use case in the FI PPP from the Future Internet Enterprise Systems (FInES) community?

Where does the Future Internet fit?

In attempting to establish what may be the emerging Future Internet requirements for enterprises, Thomas Michael Bohnert presented evidence to demonstrate that services are driving economies. Global business networks are emerging, global competition is increasing, and risks and risk potentials are also increasing. A SAP survey shows that the main business applications to drive ICT innovation are web-based services, business intelligence, modelling and design. Services that help businesses acquire, create and exploit knowledge will add tremendous value to the future enterprise.

The concept for the Future Internet requires a holistic framework, Bohnert asserted, which includes the Internet of Things (IoT) and the Internet of Services (IoS). He discussed the value proposition of the IoT and showed a video of integrated car communications with advanced features, as an example of the business opportunities that IoT can offer.

The Internet of Services (IoS) is the foundation of web-based service economies – it is an open service platform that enables the activities of various communities of networked participants. The IoS is itself supported by network enablers and a foundation architecture (secure network infrastructure, IoT, cloud computing, etc.). Bohnert discussed the value chain of the IoS and SAP's proposal for a Universal Service Description Language (USDL) in addressing the needs arising from diversified partnerships. Finally, he presented the architectural recommendation from the Future Internet Research Alliance (FIRA) being spearheaded by SAP (Note: the FIRA recommendation is not to be confused with the G15 proposal).

How will enterprise change?

SMEs are the bedrock of the EU's economy, so Miguel Borras explored what the Future Internet might mean for an entrepreneurial SME. He highlighted the difficulties and issues of technology and research for SMEs. Security, trust and usability are still big problems for SMEs, he said, and there are

major risks associated with cloud computing (e.g. who has control of my data?). Cloud computing is certainly a buzzword, but is it just another term for Software as a Service (SaaS), a model that never really took off, even for SMEs? And if they are different, which will prevail?

For enterprises, Future Internet means software and services become a commodity. Web 2.0 has dealt a blow to the "semantic faith". "What will become of semantics in Web 3.0?" Borras asked. "Indeed, will there even be a Web 3.0?"

There is certainly a need for a semantic grid on the consumer side, in other words semantic-based competitive intelligence. However, Web 2.0 has not been without its problems as a "leaky tap for enterprises" – the growth of collaborative and Web 2.0 environments rapidly increases the leaking of commercially sensitive information. Therefore, the Future Internet also means companies will have to manage their intellectual rights – an issue that Borras proposed should be added to the FinES research roadmap.

Discussion

The panel discussion section in this session followed a structured format, divided into three topics to explore the Future Internet vision for enterprises, future business models and future systems for business.

The Future Internet vision: smart enterprises of the future and routine innovation

The discussion moderators Man-Sze Li and Stefano De Panfilis asked the panel of speakers and FInES caretakers four questions:

- **Exit from the crisis:** What does it mean for enterprises?
- **Enterprises of the future:** Is it time to re-think and even to re-invent the nature of business, the characteristics of enterprises and the role of firms in a potential "new global order"?
- **Innovation union:** What is the role of the Future Internet for enterprise (and especially SME) innovation?
- **European Future Internet research:** Does it make sense to consider "enterprises" as a generic research domain?

Delegates also put forward their own questions about the Future Internet vision for enterprises:

- **Future Internet robustness:** What happens if catastrophic failure occurs? If everything is online, what is "Plan B" to deal with emergencies?
- **Future versus present:** Are we discarding the old? Or can the old be "upgraded"? How do we get from the 'now' to the future?
- **Costs:** What about SMEs that aren't – or can't afford to be – "ICT savvy"?
- **Reinventing products:** What does this really mean? Can generic strategies be applied?

The panel offered several opinions on getting from the *now* to the *future*. On the one hand, backward compatibility needs to be maintained in order to enable enterprise growth; IT support tools could help with this (e.g. the maturity models investigated as part of the COIN project). On the other hand, a revolutionary path to the Future Internet could also be expected, because the context for doing business has changed and will continue to change.

It was asked whether there is such a thing as a Future Internet that can be "switched on"? Potentially, the Future Internet is an evolution that also supports revolution, the panel suggested. Importantly, enterprise mission, organisation and intangibles need to be looked into, not just technology. There was a view that the (re)invention facilitated by the internet is not about the business *per se*, but the enterprise's market position and its interactions. In sum, there is a need to reconcile evolution and revolution, which is potentially a major challenge.

Some panellists also expressed the view that Web 2.0 tools should be more extensively used in the enterprise environment, especially by SMEs. Greater adoption of existing tools will place SMEs in a better position to embrace Future Internet applications as they become available.

There was unanimity on the panel that the Future Internet needs to be based on open standards and should be relevant to enterprises of all sizes. The internet must be a community, and an inclusive one, they said. Once again, the need to increase the awareness of SMEs of the Future Internet opportunities was underlined.

One panellist expressed the view that "If you go into business, you want to win". Future enterprises need to think carefully about the value they offer and how they do it. These include both "hard" (financial) and "soft" (non-monetary) values. Especially in light of the global crises in financial systems and other environmental challenges, the nature of "what business is" is potentially changing. More so than before, those willing to take risks are more likely to succeed. That said, it was felt that Europe "missed the boat" on Web 2.0 and many of the new opportunities enabled by the web. The question was asked: how can this be avoided when the Future Internet comes on line?

Business models

The discussion moderators Michele Missikoff and Jean-Dominique Meunier raised four issues for the panel of speakers and caretakers to tackle:

- **What will be the key drivers for business models in the internet economy of the future?**
  Technology? Customers and end-users? Organisations and staff? Public Sector, laws and regulations? Others?

- **What will be the characterising elements of those business models?**
  Value proposition (goods and services)? Business processes and enterprise architectures? New value-creation paradigms? New forms of innovation and intellectual property rights? Others?

- **Are there emerging examples of those business models?**
  Apple/Google Apps factory, emerging idea of iAd, SaaS-based business models, Amazon Mechanical Turk, Daimler Car2go, etc.

- **What are the lessons that can be learnt so far?**
  Are there already major barriers to business model experimentations (e.g. cultural resistance, lack of innovation, inertia of enterprise organisations, costs of re-engineering, lack of appropriate or new skills)?

Several views were expressed by panellists and delegates on the implications and impact of the trend towards ICT commoditisation. It was observed that ICT companies might need to focus on value-added services (possibly "giving the rest for free"). A statement from the SAP speaker at the FI PPP launch of the previous day was recalled: "Greatest innovation comes from commoditised businesses". Thomas Michael Bohnert mentioned his company's Business By Design product as an

example of addressing customer needs, on demand – this product has features set ten times over that of rival Salesforce. The general view was that, with commoditisation, the customer's voice is getting stronger. ICT providers need to transform themselves and continuously adapt their offerings.

Several views were also expressed on the patterns of emerging business models. But many people said that the new models were dependent on traditional models of enterprises – companies want to improve their business, irrespective of whether the models are about the Future Internet.

Others stated that there are huge changes even in traditional enterprises such as automotive manufacturers, leading to the observation that there may be "no such thing as traditional enterprises". Yet others drew a distinction between US enterprises and European enterprises, claiming that the former are (already) largely internet-based, whereas European enterprises tend to use the internet as a means to expand on the existing business. There are, however, major exceptions to this, notably European mobile phone players which have emerged from completely different industries and provide fine examples of business transformation. The question then becomes: do we in Europe tend to think about "evolving" existing businesses rather than thinking outside the box?

Some people expressed their concern that the timeframe from research to innovation, typically five years, was a long time in business. They pointed out that this is a very long time for SMEs, so it does not encourage them to get more involved in research.

It was observed that many world leaders in ICT are not European based. There was a view that such leadership is also "not telco based". A delegate asked where an entirely new industry could be launched in Europe based on services? The panellists stated that, whenever you follow others, you will always be behind! But, again, the mobile industry was mentioned as an example of new businesses "starting from nothing". Was this achievement a one-off, or could it be replicated? The fragmentation of Europe as a market (due to language and cultural differences) is certainly a weakness, but it could potentially be turned into unique opportunities. There was a general view that the FI PPP is a key opportunity for forming a new platform to bolster the strength of Europe.

Next-generation systems to support enterprise innovation in the post-crisis landscape

To launch the debate, the discussion moderators Sergio Gusmeroli and John Kennedy asked the panel of speakers numerous questions:

- **Which Future Internet technologies will specifically help revolutionise enterprise systems**?
  - Future networks (universal business infrastructure / network convergence)?
  - Internet of Services (cloud computing / service web / public data access)?
  - Internet of Things (sensor networks / smart objects / distributed intelligence)?
  - Internet by/for People (social networks / empowerment / Enterprise 2.0)?
  - Internet of Contents/Knowledge (3D media / Fifth Freedom / openness)?

- **In which phase(s) of a product's lifecycle does the Future Internet offer the most promising post-crisis exit strategy for European industrial SMEs?**
  - New product development (3D models / virtualisation / Open Innovation Living Labs)?
  - Sustainable manufacturing (smart / virtual / Digital Factory of the Future)?
  - EU single marketplace (virtual-physical points of sale / collective intelligence)?
  - After sales services (extended products / intangibles / dismantling / recycling)?
  - Where to start from for ICT as a service in industry? (privacy / legal / trust / security)?

- **Is the Future Internet an innovation opportunity for European ICT SMEs?**
  - What is the role played by EU ICT SMEs in the internet of the future?
  - Infrastructure, platform, software, consultancy as a service?

- Smart niche applications on top of an open core platform?
- Is what sense is 'openness' intended / understood (open source / common specifications / standards)?
- Closed innovation / collaborative innovation / open innovation (Living Labs)?
- Is partnership with ICT big names mandatory? Just EU or also US big names?

The view was expressed that combinations of Future Internet technologies are needed to deliver value. But what combinations deliver the most value? The FI PPP will address this question. The FI PPP is not about developing new technologies, but about finding ways to 'federate' existing generic enablers and integrate them so that specific combinations can be built together. The eight use cases to be selected for the FI PPP should be able to provide broad understanding of the application of generic enablers and how they may be shared across application areas. The five most promising use cases will be pushed forward into large-scale pilots.

There was a general view that the Future Internet gives Europe a significant opportunity to maximise value for businesses, individuals and society. The concept of 'value' must be clearly defined. For instance, is there a valid comparison between the US and European approach to value? What about the rest of the world? What if China buys companies based in Europe? Indeed, is there a difference between the perspectives of eastern Europe and western Europe on the Future Internet? A delegate pointed out that Living Labs could offer one way to help researchers identify how value is constructed in different contexts.

A Slovenian participant informed the session that a Danube River region is now forming. The river used to be a border and barrier, but now it is bringing companies from neighbouring countries together. New ways are needed for companies to work together – in cross-border collaboration and experimental approaches; ICT is helping to create new and better businesses. Stakeholders in the Danube region are keen to engage with the broader FIA Enterprise community. If the Future Internet could add value to regional developments, even more opportunities would open up on a global scale.

A delegate asked the panel whether they knew of any mechanisms for extracting the lessons learned from individual R&D projects so that they could be applied to other sectors. For example:

- How can findings about smart energy grids be applied in other sectors?
- What can other projects involving SMEs teach Future Internet projects focusing on SMEs?
- What does the transition to services mean for product-based companies?
- How will businesses get value from data in the future?

There was a suggestion that the EU's research Framework Programme needs to be updated so that these kinds of lessons can be shared more effectively and widely. A participant asked whether the Knowledge and Innovation Communities (KICs) being established by the European Institute of Innovation and Technology (EIT) would address this particular issue.

Several participants highlighted their concern that businesses would need education to see the opportunities presented by the Future Internet. They pointed out that there is still a considerable gulf between the research community and the business community which will ultimately be the user of Future Internet technologies.

Conclusions and Next Steps

The following messages emerged from the presentations and discussions:

- Everyone acknowledged that the opportunities of the Future Internet for enterprise are immense, both indirect and direct.
- There were concerns raised regarding trust, "failover" alternatives, ease of use for SMEs, leaking of intellectual property and competitive information, and cost.
- There was consensus that the Future Internet is relevant to existing "traditional" enterprises, as well as to entirely new industries. Indeed, as mentioned by the SAP speaker at the formal launch of the FI PPP, commoditised businesses often generate the greatest innovation. Examples were given of the traditional automotive industry embracing the Internet of Things, and of the SME dedicated to a finger painting iPhone app.
- The customer's voice is getting stronger – businesses that listen will be more successful.
- Evolutionary improvements of technology will enable revolutionary business models – but evolution and revolution will have to be reconciled.
- Companies need to be prepared to learn, to evolve, to change, maybe even to transform. And not necessarily just in relation to their products or their services; but also their vision, their values and their business processes.
- There were many references to successes in the US and elsewhere. But if we follow others all the time, then we will never be in the lead. Enterprises need to think out of the box: the Future Internet will enable a whole new world of business.
- Interesting insights were given into the perspectives of both EFII, 'G15' and the Future Internet Research Alliance.
- There was broad support for a proposal from FInES to contribute a use case to the FI PPP.
- There was a wide spectrum of views on the kinds of business models which would be promising for European businesses and how Future Internet technologies could support and enable them. There was a general view that this topic could be pursued at FIA Ghent. The participants agreed that a session could be used to develop scenarios for the future of businesses on the internet.

What can Future Internet mean for Smart Cities?

Caretakers: Nick Wainwright (HP), Alex Gluhak (University of Surrey), Mirko Presser (Alexandra Institute)

The Future Internet offers solutions to many of the operational activities and challenges that occur in cities: community building, mobility, the deployment of efficient services, new applications and services, rethinking utilities, culture and the built environment. So cities provide a unique opportunity to Future Internet research; they offer real challenges, real users at a high density, realistic societal, organisational and operational structures, self-sufficient governance and decision making.

But there is a noticeable gap between researchers grappling with city issues and the Future Internet research community. Smart city research projects tend to explore the problem space by piloting novel applications and conducting experimental research in urban settings. They look at design issues and consider ICT to be just one (albeit important) component to the 'smart city solution'. Future Internet researchers on the other hand tend to explore the possibilities of open, internet-scale infrastructure and platforms, and see how it can be applied within urban contexts.

Objectives

The session brought together these two research groups in an attempt to foster an understanding between them and provide a stimulus for collaboration. The session was organised around three different presentations, each highlighting a particular 'problem area' that smart cities will have to address. The presentations and panel discussions aimed to identify gaps and/or overlaps in the research efforts of the two research communities and to elaborate on any challenges facing Future Internet research is this domain.

Presentations

• Living in smart cities	Martin Brynskov	Centre for Digital Urban Living, Aarhus University
• Running smart cities	Cedric Ulmer	SAP Research
• Smart urban transport	Antonio Marques	ETRA

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-cities).

Summary of Presentations

The tech-savvy city life

Martin Brynskov talked about what it might be like for citizens to live in a smart city. He highlighted the need for researchers to understand that ‘smartness’ is more than just clever coordination and control of city services (plumbing). It also has to take in some poetry – engagement with the citizens. Brynskov's key message was that technology development often focuses on “serious stuff” followed by “less serious stuff”. But the citizen/consumer is increasingly driving the agenda and they are interested in technology for fun and leisure. “We need to set up processes where researchers can take both perspectives into account from the outset,” Brynskov suggested.

He highlighted several challenges that we need to address in this context:

- Create desirable but realistic visions (not just promotional material like the IFEZ u-City in South Korea), for example the 'Aarhus by light' interactive light 'sculpture', which will show how the Future Internet can mediate between communities and social interactions, provide new forms of social interaction and engaging experiences for citizens.
- Build “openness” into designs – conceptual and physical openness and 24/7, always-on accessibility.
- Identify new genres, build new platforms.
- Find the right people (both poets and plumbers).
- Organic growth is harder, but also more important (permanent solutions can be life changing whilst one off initiatives may have no lasting impact).

Who's in charge?

Smart cities require 'smartness' from different actors: citizens, associations, companies, public organisations, and city organisations. Following Brynskov's call for more citizen-focused applications of the Future Internet in cities, Cedric Ulmer looked at how the Future Internet could prove useful for city managers. “Managing cities is challenging,” he remarked, “and city managers need tools and technologies to cope. And it gets even more challenging when managers are asked to report on sustainability, economics, and satisfaction. Solving such challenges requires an end-to-end view of the city.”

Ulmer outlined the specific needs of city managers and showed how the Future Internet might be able to provide them with the level of data they require to ensure that city services are run optimally to meet financial, environmental, quality and safety targets and standards. Real-time active management would be supported by a vast network of sensors and other data sources, ranging from flow rate detectors in sewers, CCTV and automated air pollution monitoring.

Ulmer suggested that managers would need data to be presented at different levels of granularity, from single rooms, to buildings, streets and entire districts. This data would be presented using dashboard-style user interfaces, but behind these control panels would be a vast wealth of interconnected data sources and analytical tools.

From A to B

Antonio Marques presented his view of how the Future Internet could help to solve urban mobility problems. “Mobility is about people and we should think about *movin*g people as the end goal,” he pointed out. “Key measures of urban mobility are the economy, efficiency, convenience, quality, prestige, safety, and security. Poor mobility can have a big impact – congestion, for example, is estimated to account for 1% of GDP.”

Cities have tried just about every trick in the book to try and deal with congestion and improve urban transport networks. But these solutions – better access into cities, public transport, park and ride, seamless multimodality, traffic calming, cleaner fuels and vehicles, and even soft measures such as car-sharing and awareness-raising – have not solved the problems. Can the Future Internet get us out of this conundrum?

Marques provided some context to the problem of urban transport. Key drivers for change are climate change, cost, safety and ‘Generation Y’ (the digital natives). Indeed, this generation's desire

for choice (citizen empowerment) is one of the key drivers that will promote structural change in transportation.

The overriding aim for cities is to transform citizens into **consumers of mobility**, but not consumers of cars, fuel, train or bus journeys *per se*. "We must enable transport users to take well-founded travel decisions based on current, relevant, global, integrated information. People must be empowered to choose," Marques argued. "Given the right information, we may be surprised how 'wise' people will be. So we need a platform which can be a utility for road users, public transport operators and citizens alike."

It is clear that the Future Internet is necessary for this purpose because such a utility would require access to a massive range of real-time data. But what kind of useful services could be on offer?

Marques described a tool that would let people see where the congestion was on the roads, letting traffic lights respond to where travellers are actually going, telling people what their alternative options might be in real time and integrating all of this with cost information and billing services. As a real-time service, it would be possible to introduce variable pricing based on numerous parameters (e.g. pollution levels, traffic flows, capacity in public transport systems, comparative fuel consumption, etc.).

"These are the kinds of things that would let transport users 'consume mobility' and not have to manage every journey on an individual basis," Marques concluded.

Discussion

Following Martin Brynskov's inspirational presentation on citizen-focused Future Internet applications, the panellists (Martin Brynskov, Malte Behrmann and Peter Ljundstrand) observed that today there is a much better mix between the 'softer' and 'harder' sciences and this can lead to more effective results and applications. However, there is a chicken-and-egg problem: how can you perform integrated research without necessarily being based on technology push? The city is about citizens, the panel agreed. But how do you get them on board when it comes to technology development and design?

Discussion highlighted different methodologies for citizen participation and delegates mentioned some examples of participatory design within existing projects. Games could be also used as tools to capture user feedback at early stages someone suggested. It is, however, important not only to focus on citizens as the only users, but ensure that all stakeholders are engaged, including architects, builders, infrastructure companies, etc.

Brynskov mentioned that the Future Internet would allow every surface in a city to become an interactive input/output (I/O) device. But what would be the implications of such a scenario for humans, a session participant wanted to know. What would be the impact on our daily lives? If the Future Internet is to be so intrusive, can we get away from it?

The panel observed that the home offers barriers and dividers of physical space; walls offer protection and organisation in a natural way. The social context of a user and their relationship with an I/O surface could also affect how people respond to them. Nevertheless, as the example in SongDo illustrates, South Korea takes the smart city topic very seriously (perhaps more so than Europe) and it is certain the technology will be pervasive.

A delegate asked Brynskov about his vision of how people and their homes fit into the smart city. Brynskov accepted that most of a city is where people live, so home environments would also be included in smart city networks. The main differences lie in the scale that the solutions have to

support and the social aspects. Some reference was made to games such as the Mirror's Edge, which take place in a city context.

Some comments were made about the complexity of cities as systems. Could complexity studies and tools be useful for people to understand how to design smarter cities? Certainly designing smart cities from scratch would probably not be best; evolutionary design seems more promising as it would incorporate our experience and knowledge of 'typical' or emergent behaviours and responses.

After listening to Cedric Ulmer's portrayal of smart city managers, a fresh panel (Cedric Ulmer, Barbara Daskala and Nigel Baker) wondered how citizens might respond to city managers having an end-to-end view. They discussed how privacy within the smart city context could be addressed. It was observed that the 'smart' in cities can be dangerous; trust, governance and privacy would have to be addressed and citizens engaged before they would be happy to live in such a pervasive monitoring environment.

At what level do we seek to provide an end-to-end view of city activities to the city managers (the city authorities)? Will smart cities be transparent about the data and how they (and whoever else) use it?

Ulmer suggested that data would have to be kept at a high level (i.e. 'anonymised' and aggregated) and that it was vital to have complete transparency. The panel suggested that citizens living in the city should be involved in debating what information should be made available to the city authorities.

It was observed that if there are a lot of sensors embedded in cities this can provide a lot of fine-grain information and that good governance of that data is essential. Ownership of that data gives power to the owner of the data, and trust will play a major role in acceptance of real "smarts" in cities. Research in trust and identity must be applied to this problem. Perhaps there is also a need to develop algorithms to manage fine grain sensor data in a transparent way that respects privacy.

It was noted that (like in the earlier talk) when you deal with ICT and people, you have to work on both angles simultaneously. Responsibility for transparency lies with those managing the city and the information.

The issue of privacy led the panellists to think about how it would be possible to create open, innovation-friendly systems that will foster new applications and services? It was observed that today, lots of city data is in sector silos (fire, police, water, power, etc.) and locked into proprietary formats. Making this data interoperable would be an important step to demonstrate how access to more data could have a big impact. However, business models may have to be rethought: there are many private companies in the end-to-end city system. How can they open up their data without losing their business?

The question of security came up, and someone pointed out that the notion that closed systems are the most secure is being challenged. Again, engaging the user to establish trust is a good technique.

Brynskov pointed out that people are not, by definition, averse to being monitored – individuals were quite happy to being 'sensed' in the Aarhus by light installation. So context is important – being sensed by city officials for management purposes may not be quite so popular. It is a question of "control versus poetry".

Many further questions were raised by the audience:

- Is the city manager more important than the citizen?

- How can we take a bottom-up approach (as advocated by the Living Labs initiative) to city management and the application of the Future Internet in this context?
- What is a good metric for the level of digitisation or “smart” in a city?
- The end-to-end view makes the city manager happy, but what about the citizens, what makes them happy?

In conclusion, the panel suggested that the key message was that citizens want – and expect – well managed city services including emergency services. 'Smarts' are vital to this, but during the building process, we must engage the citizens too and take care that, in solving one problem, we don’t create another.

Lively discussion also followed Antonio Marques' presentation on the application of the Future Internet in solving urban transportation problems. Delegates felt that this opportunity to inform end-users (i.e. travellers) and give information and advice is a key opportunity for the Future Internet. At present, you can only do this on a small scale, but it would be so much more powerful to have applications that could manage journeys on a city, national or even international scale. The panellists (Antonio Marques, Fiona Williams and Jonathan Cave) agreed that there was a great case for integrating data, but the tools would have to be easy and simple for the user. “If it is difficult to understand, then I will just take my car,” one panellist said. At present, there are the 'haves' (who already know how to access the right websites and obtain this kind of information) and the ‘have-nots’ (who are disadvantaged and just take the default option of travelling by car).

The first transport revolution (trains) shrunk distances and rewired the world. The second revolution (cars), let transport be individualised, but also created new problems and externalities (e.g. congestion). The third revolution may be one where we internalise those externalities.

But will people become ‘mobility consumers’ and give up ‘inessential freedoms’ (e.g. control of the car)? We must consider how the demand side interacts with the supply side, the panellists agreed. It is not just about giving information to travellers, but also empowering the supply side with real-time data so that the complex urban transport network functions like you would expect. There is no point informing a citizen that a bus will be along in three minutes if there is no room on the bus!

A profound change in mobility will occur when people begin to think differently i.e. as consumers of mobility. Experiments in the direction of smart urban mobility are good, the panellists agreed, and they looked forward to seeing an impact over time as our views and behaviours change. The issue of internalising costs is being provoked by the issue of electric cars – one cannot compare just the selling price of electric cars versus ordinary cars. We must take into account all the other costs and let people be fully informed.

Noting the international travel chaos caused by Eyjafjallajökull's ash clouds, it was observed that the problems of inter-city transport are not the same as urban transport, and the problems created by the ash cloud are not the everyday problems of urban transport. Urban transport addresses what it is like every day when it takes one or two hours to go to work. Urban transport must focus on daily life and the quality of life. Interlinking international and inter-urban transport is complex and also involves many non-technical challenges.

A session participant wondered whether and how much transportation could be self-regulated. Marques thought that urban mobility could never be an entirely self-regulated system. He pointed out that if the network manager, who has a legal mandate to manage the system, has to set the traffic light to red then you have to stop! Empowerment is not the only thing, safety must also be considered so complete self-regulation would be practically impossible.

Someone asked whether the Future Internet could not just optimise urban transport, but eliminate it entirely through remote working, collaboration, teleconferencing, etc. It was observed that there are trends in urban planning to try to get the demand for transport down by locating work, shops and transport in clusters. However in many cases virtual reality and remote working it not a good substitute for physical presence – you only have to look at the interest in FIA Valencia to realise the benefits of meeting people face-to-face. However, when normality is disrupted, people develop alternatives to transport (e.g. when people couldn't get to work because of earthquake damage in LA they managed to work at home). But these new patterns usually return to normal once circumstances allow. And even if we insisted on teleworking, cities would then be divided by geography; the rich and powerful would live in the city centres where they could also meet face-to-face, and the teleworkers would be disenfranchised away from the power hubs.

In conclusion, it was noted that (despite many believing the contrary), mobility is a human right, driving is not! Although this wasn't stopping several delegates frantically searching the web for hire cars to get them home under the immovable ash cloud.


The European RTD Framework Programmes
From
Economic Recovery
to Sustainability
IV Conference on the FP7 in Spain
12-16 April 2010 / Valencia - Spain

Socio-economics of the Future Internet

Foundations of Trust

Caretakers: Nick Wainwright (HP), Volkmar Lotz (SAP), Jim Clarke (WIT), Michel Riguidel (ENST)
Chair: Volkmar Lotz, SAP

Europe's Future Internet initiatives will extend the reach of internet services into many more aspects of business and our personal lives than we have hitherto experienced. We will encounter new business models, new platforms and new services, more open and global online markets, increasingly sophisticated data mining technology and deeper integration between the physical and virtual worlds. These new technologies all create new challenges for the creation of a trustworthy Future Internet.

Objectives

The purpose of the session was to explore the human and technical foundations of trust in the Future Internet with a view to drawing up a roadmap for future cross-domain research. The aim was to address the very real question of how to ensure that end-users can have sufficient confidence and trust in the infrastructure and services of the Future Internet so they participate readily in a digital life and digital society.

Presentations

• Security challenges for the Future Internet	Prof. Evangelos Markatos,	FORTH and University of Crete
• Provenance in the Future Internet	Dr Jose Manuel Gomez-Perez	iSOCO
• Economics of trust and security	Dr Simon Shiu	System Security Lab, HP
• Legal frameworks for trust in the Future Internet	Dr Mireille Hildebrandt	Vrije Universiteit Brussel and Erasmus University Rotterdam

*The presentations from this session are available on the FIA Valencia website (*http://bit.ly/fia-valencia-trust*).*

Summary of Presentations

In his introductory remarks, session chair Volkmar Lotz stated that there was a need to come back to the foundations of trust in the Future Internet because security and trust is a really "cross-cutting" topic, and must be addressed from many different angles to achieve confidence and trust in the services and infrastructure of the Future Internet.

He noted that one of the Trust and Identity sessions during FIA Stockholm, introduced the important topic of measuring trust; it was agreed that this area would clearly benefit from interdisciplinary work amongst the work streams of FIA. For FIA Valencia, it was decided to focus on four topics in the session laying out the foundations on which trust in Future Internet can be built.

Security challenges

Evangelos Markatos of FORTH in Greece spoke about the work of the FP7 ICT Coordination Action FORWARD which has explored internet security challenges. In particular, it has found that hackers

are getting increasingly more sophisticated. Large-scale viruses that compromised thousands or even millions of computers are a thing of the past. Today's hackers use social networks, Twitter, corrupt files, etc. to catch users. One example is 'Koobface', a worm on Facebook, which tricks friends into downloading malware. Hackers use Facebook to launch phishing attacks, exploiting the trust between users and friends to trick them. Another example is that people are directed towards bogus websites via Google and search engines; hackers use topical websites, spurious charitable websites and similar deceptions to trick users. Twitter is used to distribute URLs that point to malware.

The impact of cyber attacks is getting larger. Quoting from Viviane Reding, Markatos pointed out that attacks may have a widespread impact on real lives. For example, computers at the Houses of Parliament in the UK were infected. In other examples hackers brought down train signals, and in one case a disgruntled employee remotely disabled more than 100 company cars. Given the widespread use of networked computers in all aspects of our critical infrastructures and everyday life the possible consequences of serious cyber attacks are extremely worrying.

Markatos described the main areas of focus for FORWARD's three working groups: malware and fraud, smart environments and critical systems. The project has also used expert working groups to identify and rank emerging threats in ICT infrastructures. High-priority topics that emerged from these think tanks included:

- **The underground economy**: There is a dramatic change in the goals and models of hackers, shifting from hacking for fun to hacking for profit. This underground economy is flourishing. There are even support structures for this underground economy – markets for information, bullet proof hosting and 'rogue' networks. Possible solutions to this problem include methods to attack fraudulent transactions (by flooding the bogus systems with useless data), large-scale tracking and data correlation to identify the places where these markets are lurking.

- **Social networks**: Social networks attract hackers due to the high number of users and the large degree of trust held between users and their trust of the social networking service. Hackers capitalise on this trust network. Most social networks have third-party applications, many (most) are games, which may give hackers access to the private information of the player of the game and/or access to the user's hardware (e.g. by uploading files). Possible solutions (e.g. fine grain mechanisms) will require a united effort, with collaboration from social network providers.

- **Threats due to parallelism**: Parallelism is hard to implement correctly, which opens the door to hackers who can search for bugs and exploit race conditions, etc. Most people are not very good at programming for parallelism, so we need to invest in the development of new, more secure programming languages, apps, libraries and operating systems. They must be designed with parallelism in mind. Virtualisation and hardware security innovations may help with this.

- **Threats due to scale**: We are vulnerable to attacks that leverage and amplify minor vulnerabilities in millions of devices.

- **Mobile device malware:** Mobile devices are a soft spot in our defensive armour because they are mobile and highly connective (e.g. automatically logging on to Bluetooth and WiFi hotspots). Moreover, physical security is a real issue – they are too easy to lose. Possible solutions include apps in sandbox, intrusion detection and server replication of every phone's state.

Markatos concluded by introducing delegates to SysSec, a new European Network of Excellence that will explore how to manage threats and vulnerabilities. At its heart is a new game-changing approach to cyber security. Currently, researchers are mostly reactive; they track attacks after the attack has been launched and work out protection once the attack is understood and fully characterised. But this always puts researchers one step behind attackers. SysSec advocates that researchers should be more proactive and should anticipate attacks, predict and prepare, and warn the research and security community before attacks materialise. SysSec creates a distributed centre of excellence in the area of emerging threats and threat analysis.

Provenance in the Future Internet

So how do you know you can trust the information on your computer? This is a good question to ask – and you can only really trust data and information if you can be sure of its provenance. In terms of building trust, provenance is an enabling technology, and it urgently needs some serious attention.

The notion of information quality must find a good balance between the security of a system and the functionalities that the system provides. For example, locking all your cash into a deep bank vault is extremely secure, but you can't spend it. On the other hand, carrying wads of cash around the shops makes it highly 'functional' – you have enough to buy whatever you want – but there's a risk you might lose it or even be robbed.

It is therefore important to have an intelligent and automated method to evaluate the extent to which any item of data can be trusted. It must build on knowledge about who produced the resource being accessed, what process produced such results and how the information was transformed during these processes.

Provenance information is a record of the sources of information, with evidence of its authenticity, explained Jose Manuel Gomez-Perez. Provenance information is valuable, but it is hard to collect and verify. But armed with this information it is possible to assign 'credit' when data is good, and in some cases, 'blame' when data is bad. The use of provenance information would be valuable, for example, in situations where false information is published about events using fake websites, fake news items, fake Wikipedia entries, etc.

It is also important to know the provenance of linked data, so we know where the data came from, who produced it and why. Open government initiatives are good examples, for example the open government initiative in the US, and the data.gov.uk initiative in UK to publish government data in RDF format.

Provenance allows people to get a much clearer picture of the quality of information – its timeliness and consistency, along with stable and meaningful data links.

So what does provenance information look like? It is usually represented as a graph; provenance models define types of provenance, elements, and relationships between them and allow us to answer the following questions:

- Who created the content (attribution)?
- Has it been manipulated or processed, and by whom?
- Who is responsible for providing it, and where is stored?
- How can I believe this provenance information?

The W3C provenance group – started in 2009 – has developed a set of key parameters for provenance information and published a number of user cases. It has developed user and technical requirements; a state-of-the-art report is due for publication in June 2010.

Key technical challenges for the future include:

- more work on developing vocabularies for provenance information;
- assessing granularity – how much provenance information is useful, manageable and scalable?
- defining information quality and how this relates to trust;
- measuring the evolution and updating of data – tracking timeliness and versioning;
- integrating provenance into data consumption, visualisation and navigation.

But an arsenal of provenance-related security functionality is just the start, warned Gomez-Perez. People and organisations must learn how to implement policies based on a sound knowledge of the provenance of information. How provenance translates into trust depends greatly on context, but there may be general principles on the associations between sources and their level of trustworthiness (e.g. if a source is an oil company vs a blog). There appears to be a need for authoritative agencies to underpin the quality of provenance data.

Perhaps the big question is how to 'incentivise' the web to make use of provenance so that content and service providers move to a provenance-aware paradigm. Service providers will have to generate provenance metadata, while search engines must consume and exploit it (e.g. increase rankings in search engines, attract internet traffic, increase automation according to the provenance information of content).

Economics of trust and security

Organisations typically make poor quality decisions about where and how to make investments in security management systems, technologies, and processes, asserted Simon Shiu at the start of his presentation. But help is at hand: economic methods can be applied to the area of trust and security to help large organisations make better security management decisions.

In the security management lifecycle, the CIO has to think about risks to the organisation. This guides the organisation's policy concerning which risks to address, and how. The guidance governs the implementation of systems and processes to address those risks. Finally, the systems and processes are analysed to determine how effective they have been, and the results of this analysis is fed back to the CIO who uses the data to modify the policies.

So where do you start in this repetitive loop? You have to understand the risks because the risks ultimately determine policy, investment and outcomes. Security investments affect multiple outcomes – budget, confidentiality, integrity and availability; predictions are made with high degrees of uncertainty, outcomes are inter-related – but the link to investments is poorly understood.

The classical business justification for an investment – the return on that investment – is poorly evaluated, Shiu argued. Indeed, many of these points are typically glossed over and the link between investment and business outcomes is weak or even absent.

Economics has many techniques and tools to frame and analyse these types of problems. A good analogy of the type of decision framework we seek to create is with the 'central bank problem' which uses a 'utility function' to determine its interest rate policy. In security, we need a utility

function that will tell us where to invest in security measures to achieve satisfactory levels of (for example) cost, confidentiality and availability.

So once you understand your risks, you then need to know what you want your security measures to actually achieve. However, most security stakeholders do not seem to take enough account of the multiple business relevant outcomes of security breaches. Shiu described HP's work using a process of preference elicitation to move from key issues (confidentiality, availability, cost) to consequences (impact of breaches, SLA violations, etc.) which can help managers compare the potential impact of security investments (or lack of investments) on the business.

HP is leading two research projects – Trust Economics and Cloud Stewardship Economics – funded by the UK Technology Strategy Board. The Trust Economics project is working to integrate many scientific disciplines into enterprise security management.

The Cloud Stewardship Economics project meanwhile is looking specifically at cloud computing. Here, the enterprise is one type of service consumer; security properties and decisions are contingent on complex incentives (obligations, preferences, requirements, expectations) and interactions within the cloud services ecosystem. There are many intuitive analogies that suggest (micro) economics will be a good tool for exploring these incentives. For example:

- the service provider knows more (or at least different things) about costs and risk than users or regulators (i.e. there is an information asymmetry);
- there are many situations where being secure costs me more than I gain, even though others in the community gain too (i.e. what are the external costs and issues associated with public and/or club goods?);
- the challenge that providers and consumers currently have trying to assess the value of bundled security characteristics and developing competitive pricing strategies (i.e. there is heterogeneity between services and users).

Left unchecked, it seems likely that the ICT services market will prioritise low cost and flexibility, Shiu said, ignoring the negative 'security externality' effects. It is therefore essential that organisations become more explicit about their current and future information security lifecycle and needs.

A new legal framework

Mireille Hildebrandt presented the delegates with a number of topics that she believed needed to be addressed in parallel with the technical developments and architecture design of the Future Internet. Today's internet is already challenging existing legal frameworks and the Future Internet will almost certainly take this assault to new levels, so it is essential that numerous legal issues are addressed now.

- **The concept of trust**: It is important to acknowledge that trust is a way to reduce complexity, thus enabling citizens, consumers, governments and companies to interact despite uncertainties due to the complexity of the socio-technical infrastructure (cf. Luhmann). Without trust, people may refrain from taking risk and the creativity and added value that is generated by trust will be lost. Trust and security may be at odds with each other: too much focus on security stifles creativity (cf. Nissenbaum).
- **Legal certainty:** Legal certainty is an important instrument to sustain trust because it stabilises legitimate expectations. Within a constitutional democracy core values like privacy, freedom from discrimination and due process must be sustained to promote such trust. These legal protections must be rethought and re-articulated in the era of proactive

environments that provide personalised responses to people on the basis of inferred computational knowledge (i.e. knowledge that is inferred from possibly even 'anonymised' aggregated data).

- **Ambient law:** Legal protection needs to be articulated at the deepest level of the socio-technical infrastructure to be effective. The democratic legislator must inscribe privacy rights and especially transparency rights into the smart infrastructures that form the core of the Future Internet. This is the only way to give substance to the transparency right of Article 12 of the European Data Protection Directive D 95/46/EC, which stipulates that people should be informed about the logic of processing that determines how they are profiled. Users must have adequate feedback on how they match the computational knowledge that is 'out there', so they can practice 'smart data minimisation'. Control over what data you release and to whom is much better than trying to hide all of your data indiscriminately.

Discussion

The panel discussion explored the issue of which threats to focus on. Someone asked about the well known Kaminsky attack on domain name systems (DNS), but it was not discussed because this vulnerability emerged a couple of years ago and is now well known and understood. The panel agreed that the focus for current work is to understand what new threats may emerge.

One questioner asked how we can take into account the human being in the system. Humans affect security outcomes quite strongly, whether we are dealing with behaviours, interfaces or whether people make reasoned decisions or quick ad hoc decisions. All these factors significantly affect outcomes and trust in the Future Internet.

During the presentation on provenance, and afterwards in the panel debate, various delegates and panellists debated how provenance-based approaches could complement social-based approaches to define a notion of trust.

The need to establish a notion of system accountability in the current and Future Internet was also mentioned in the presentation about legal frameworks, where consumers of internet services have the right to know and approve the use of their personal data. Provenance has been applied to address this problem in the past and will contribute to this respect as well.

Delegates voiced their concern about the escalating quality of data and metadata that needs to be generated in order to support high-performance networks like content-centric networking. In the provenance area, this implies that provenance vocabularies need to improve their level of expressiveness, including a high-level of abstraction. Data compression techniques will be needed to minimise the size of provenance records and reduce the effects of 'metadata overload'.

Finally, the panel discussed ways to move beyond 'the system with the technology of the script' (i.e. sequential programming) to address Future Internet issues. If we want to sustain security protection in the Future Internet we have to think about how to implement systems that operate in parallel. Machines will begin to make use of emergent behaviour and predict agent behaviours and actions; we are going to be anticipated in unpredictable ways by machines so we must also think about the field of liability – if ambient technology tries to predict our actions and behaviour based on complex algorithms and emergent behaviour, it would become impossible even for us to be able to say for sure what predictions the technology might come up with. This will make it very difficult to say who is actually liable if something goes wrong.

The Economics of Information for Citizens, Communities and Commerce

Caretakers: Michael Boniface, Man-Sze Li, Tuan Trinh, Jonathan Cave
Chair: Michael Boniface

Digital information is the principal asset of the internet and systems are increasingly focusing on evolving networks of autonomous applications and people who interact to produce, publish and retrieve information. The growth in internet usage and system-to-system interactions will require infrastructures that support billions of information exchanges.

Digital information is now the enabler for creativity, innovation, decision making, economic output and enjoyment, but also of retrograde changes. For example, the phenomenon of 'internet addiction' shows that what ought to be a means to better transactions, interactions, decisions, etc. can also negatively consume vast amounts of attention. 'The internet' becomes an end in itself and information overload renders many decisions so complicated as to threaten the effectiveness with which people make them. Understanding the nature of digital information, how it can be used for societal and economic benefit, and how it is governed will be essential for the success of the Future Internet.

In the recent European Commission communication, 'A public-private partnership on the Future Internet', the Future Internet is described as a tool for a smarter world. Smart infrastructures are cited in energy, environment, transport and health care sectors all promising to make extensive use of connectivity and distributed information processing to redesign their business and operational processes and make them 'smart'. But how do infrastructures become smart? Or to put it more explicitly: how can infrastructures determine what is important among the people, applications, sensors, actuators, etc. with which they interact? How can they make clear decisions at a huge scale and in real time, bypassing irrelevant information and achieving solutions smoothly, effectively and efficiently?

Linking to content and linking through content will be increasingly important. We expect to get beyond current textual communities to non-textual material where presumptions of a common language and linear, sequential access are no longer assumed – for example in the use of video and music (where at least the sequence may hold until mashing takes over) and images (where even the 'language' of expression may differ by individual or group). In this case, we may hope for a more 'generative' discourse. So 'information' develops much richer societal meanings, property rights become fluid and collective, and defined as much by users' attributes as by any payment or technological channel of access.

The key ingredient that will give value to the Future Internet will be the fundamental ability to share information (e.g. network, application, users, location, time, etc.) between citizens, communities and commerce (e.g. smart energy systems that deliver efficiency savings will require accurate real-time information from consumers (smart meters), the environment (for meteorology forecasts), and transport infrastructures to predict future demands).

Similarly, consumers want more accurate and timely 'quality of experience' (system behaviour as experienced and measured by consumers) and quality of service (system behaviour as measured by providers) to intelligently select the best supplier for their requirements.

Network operators, meanwhile, want to know the characteristics of application packets to optimise delivery paths, business performance and to preserve levels of investment. E-commerce retailers, search engines and social networks want personal information so they can derive information about people's behaviours and create applications that deliver highly personalised advertising.

In the business sphere, enterprise systems will increasingly rely on knowledge sharing as well as digital assets trading to become far more flexible, adaptable and open than today, thereby enabling enterprises to tap into the latest business opportunity and form dynamic value networks.

There is a strong assumption that increasing the quantity and availability of information (regardless of its alignment with cognitive capacity, power to act or objectives) is a good thing. But our current systems of markets, laws, etc. do not assume this; indeed, they take great care to align information with other characteristics. So the development of the Future Internet calls for an enormous leap of faith and should not be made without very careful reflection.

In all these cases, individuals or businesses assess the relative benefits of protecting or disclosing information, although sometimes governments mandate disclosure. For citizens, this decision is often not a conscious process, but for most businesses and governments risks are assessed, even if imperfectly. Non-disclosure may produce information asymmetry in markets (P2P vs network operators) and significantly affects the balance of power and its overall performance. In such cases, where greater balance is required, regulatory bodies can mandate information disclosure, but only by assessing the rights of one side over another considering specific relative costs and benefits of those involved. Disclosure can also create asymmetry, unless it takes into account differences in:

- reasoning power;
- prior knowledge;
- common knowledge (what each person knows that others know, etc.. which in turn lets them draw appropriate inferences from what they hear from others or see them doing);
- powers to act, etc.

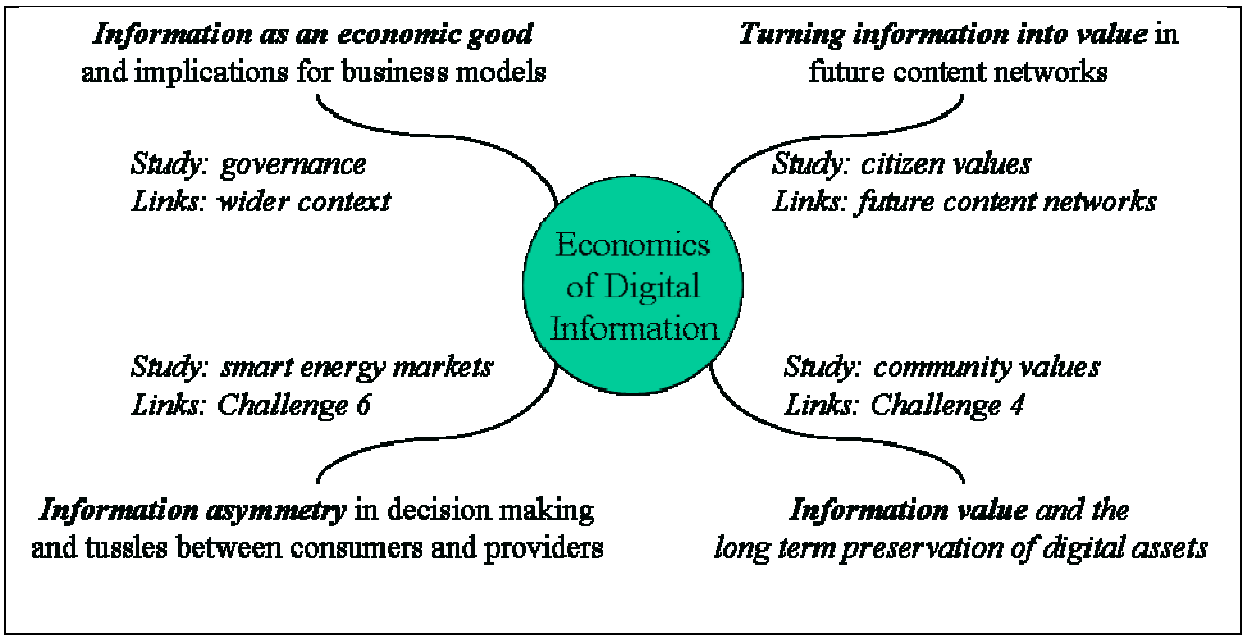

Economic Aspects of Digital Information

For enterprises operating in the digital economy, determining and maintaining the value of information and risks in respect to making decisions is increasingly important, and decision makers are increasingly reliant upon information. However, digital information is extremely heterogeneous (e.g. media content, sensor data, software, etc.) and business must adopt corporate structures, processes and agreements to govern information, but in a way that maximises innovation for themselves, their partners and customers.

Of course, what constitutes 'business innovation' is still not precisely defined, which makes it even harder to understand the economic properties of digital information and how it can practically be governed.

This lack of understanding has led to tussles between those that want to close and control information and those that want openness and freedom. Technical platforms, business models and laws have been developed that try to assert or apportion control over digital information (e.g. Amazon, Apple, and – according to some – Google); these parties argue that they need to protect investment. In contrast, open communities continue to promote the use of the "commons" as the basis for greater innovation and societal good. The challenge is not whether one is right or wrong, but how both situations can coexist.

Objectives

Information sharing is a complex issue with many deep socio-economic concerns, phenomena and tussles. It is related to aspects such as open versus closed cultures, intellectual property, privacy, information value, risks and rewards, incentives and even societal freedoms and values. The objective of this session was to examine information sharing from an economic perspective as the basis for providing insights into how 'smartness' can be valued as well as achieved in the Future Internet.

Presentations

• Information as an economic good and implications for business models	Claudia Keser	Georg-August Universität Göttingen
• Turning information into value in future content networks	Doug Williams (BT), Peter Stollenmayer (Eurescom), Adolfo M. Rosas (Telefonica)	
• Information value and the long-term preservation of digital assets	Roeland Ordelman	Netherlands Institute for Sound and Vision
• Information asymmetry and tussles between consumers, providers and operators	Dr Tuan Anh Trinh	Budapest University of Technology and Economics
Projects		
• FI3P socio-economic study	Jonathan Cave	Rand Europe
• SESERV Support Action	Michael Boniface	IT Innovation

The presentations from this session are available on the FIA Valencia website (http://bit.ly/fia-valencia-economics).

Summary of Presentations

Embracing common interests

This was the session not to miss – at least if you wanted the chance to win some sweets! Claudia Keser used a simple game to make a point. Each session participant was given a raffle ticket and they had to decide (without telling anyone) whether they wanted to invest their ticket into an individual account (with a guaranteed, fixed return) or a common account. Keser's game – a generalised

version of 'the prisoner's dilemma' – demonstrates that you can exploit a common pool for societal benefits and that the existence of public goods can increase contribution.

Putting the sweets aside, Keser outlined Ostrom's principles for successful self-organisation of the commons and showed how these are already manifested in the internet today (e.g. open source software, Wikipedia, etc.). Keser concluded with a list of challenges that would need to be addressed to maintain the coexistence of the commons and property rights (economic value of the commons, negative market effects, governance models, reputation, trust and quality assurance).

Hats off to the future

From sweets to Thinking Hats. Inspired by Edward de Bono, Doug Williams, Peter Stollenmayer and Adolfo Rosas offered their thoughts on the Future Internet from three different perspectives: the pessimistic conservative telco, the ambivalent "nothing new" stance, and the optimist who embraces future opportunities. The purpose of their role play was to stimulate discussion on whether the Future Internet was merely a pipe dream or how it really could bring added value to business (see the following section covering the discussion of this presentation).

Roeland Ordelman's presentation focused on practical applications, specifically investigating how to identify and deliver value to communities of interest which may have competing values. Ordelman described work from the Netherlands Institute for Sound and Vision, showing some behavioural experiments where the institute had targeted specific communities using social networking technologies and online gaming to crowd source and automatically collect/share annotations. He was able to demonstrate that the commons can be exploited to create value, if the right people or audiences are targeted.

Tuan Trinh's final presentation discussed the applications of smart information management to deliver energy efficiency in networked systems, including energy grids and next-generation mobile networks. Trinh presented results from the **EARTH** project which is analysing energy efficiency in mobile networks, looking at how to deliver maximum energy savings. He outlined the challenges that smart grids must overcome, including metrics/measurements, efficient communications and management, and the development of appropriate incentive mechanisms for engaging all market stakeholders. For example, he highlighted initiatives to provide energy users with more information about their usage to try and influence their behaviour (e.g. energy monitors and displays in the home). "With efficient management of information," Trinh concluded, "it is possible for the stakeholders of the Future Internet to run profitable businesses, but also make them energy efficient. I think that is a win-win situation."

Discussion

Following Keser's presentation, a session participant asked about the roles and limits that economic games could have in the development of the internet. There was some discussion about the difficulty in practical application of such approaches in controlled technical experiments due to the complexities in human behaviour and decision making.

The 'Thinking Hats' role play certainly got the delegates thinking and stimulated a lively discussion. Much discussion centred on whether it would be possible to deliver the expected value when the complexity of information is increasing beyond the limits of cognitive ability.

The crowd sourcing experiments to tag video clips prompted one delegate to ask about how the interest groups were identified in the first place. They also wanted to know about techniques to capture, evolve and enforce rules of engagement in relation to common pool resources. The development of acceptable participation technology (e.g. social metadata, etc.) that links people through content and goes beyond textual links is absolutely necessary, Ordelman argued. Tracking

social relationships through their interests in library resources would be the next step, although the availability and legality of monitoring and analytical tools (e.g. for tracking the evolution of relationships and their meaning through content) would need to be investigated.

Closing Plenary

16.04.10 1500: Ash cloud still covers N Europe. STOP. All major airports closed. STOP. Situation not easing. STOP.

Feedback from Parallel Sessions: Conclusions, Discussions and Next Steps

Chair: Theodore Zahariadis , Synelixis

Panel: Van Jacobson, Markus Brunner, Susana Bañares Hernandez, Mireille Hildebrandt, Martin Brynskov, Miguel Borras, Piet de Meester, John Domingue, Paul Moore, Tuan Anh Trinh

That Friday afternoon feeling was settling in. Satisfied with excellent paella and wine, full of new ideas, pockets stuffed with business cards, the delegates were ready to round off this excellent event. Few seemed worried that most of European airspace was closed, that news sites were predicting flights could be grounded for several more days. After all, now that the sun had burned off the haze, being stuck in Spain couldn't be so bad, could it?

Theodore Zahariadis explained that he had spoken to the caretakers of the parallel sessions. His challenge was to highlight the main issues that had come to light, then open the floor to questions and discussion points for the panel.

The main points of the parallel sessions are found in the detailed coverage of the sessions in this report.

Before opening the discussion Zahariadis offered his own big questions:

- What should an architecture of the Future Internet contain to offer the applications of the future?
- Will application domains or socio-economic considerations deliver the right requirements for the Future Internet, or can it be designed without any application in mind?
- How can we test and validate the Future Internet core platform/components/architecture?
- What are the next steps? How can we move towards a consolidated European view?

Each of the session caretakers then had 1 minute to summarise what they thought were the major issues that needed to be addressed within their particular areas of focus.

- **Smart energy (Susana Bañares Hernandez):** "The next decade will focus on the demand side, where Future Internet will play a very important role. We need to take a massive volume of data and integrate it safely into the control of the grid. Plus, we need more informed consumers."
- **Smart cities (Martin Brynskov):** "Future Internet will have to deal with the crystal palaces and mud homes. That's applications for urban planners and municipal managers and citizens. The problem is the management systems need data that's currently in the hands of citizens. They need participation, so will have to prove future systems are safe, efficient and trustworthy."

"Future Internet will have to deal with the crystal palaces and mud homes"

- **Smart health (Paul Moore):** "Health is one of the biggest drivers of Future Internet applications and the domain where Future Internet is most evolved. Short cuts are not allowed, so health is a vital proving ground for the future."

- **Enterprise (Miguel Borras):** "Today it seems to be the consumer who gets all the killer apps. We need to focus on value-added applications, built on semantics which is essential for business intelligence. Trust in business must also be built back up and we also need to detect 'leaky hole' through which competitive information escapes."

- **Search (John Domingue):** "One size does not fit all, searches will be specialised for niche applications. This is a great opportunity for SMEs to develop specialist, efficient search solutions. Challenges included improving user interfaces and increasing a user's trust in a search engine to find the right answer. We shouldn't think of search as a service anymore, but a platform on which to build future applications."

- **Socio-economics (Tuan Anh Trinh):** "There is plenty of information out there in different forms. The challenge for the Future Internet is to turn information into something of value. The socio-economic requirements of a Future Internet must also be input into its development at this stage."

- **Architecture (Markus Brunner):** "The Future Internet is still conceived as the 'middle bit' between applications and telecommunications technologies. This perception will be very hard to change at a global scale. We need to analyse what we can't change once we've started on a Future Internet. We need to agree on where we need standardisation and where we need technological competition. We need to make a smart choice about a set of core components that everyone needs to agree on."

- **Trust (Mireille Hildebrandt):** "Big issues are democracy and privacy. The current legal framework is not geared up for the Future Internet. Lawyers are used to expressing law in text and think things can't be regulated… they can't be written down on paper. We need to find ways to articulate law within the social-technological infrastructure. We need ambient law – law embedded into the environment and the technology around us. Designers and engineers need to sit down with lawyers and legislators."

- **FIRE (Piet de Meester):** "It is clear we need to provide experimental facilities, but we need to think globally. We need to think about FIRE 2.0 which involves international collaboration, with more efficient use of resources, heterogeneous test-beds. And there needs to be more interaction between the experimenters and those providing the experimental facilities."

Panel discussion

A delegate raised the issue of the scale involve in data collection, especially for smart energy and smart cities. "Are you investigating how this **data could be misused?**" they asked. Susana Bañares Hernandez acknowledged that misuse was certainly an issue, but suggested that it would also be unethical not to collect data to try and meet climate change objectives. She accepted that it would be possible to infer patterns of consumer behaviour from their energy consumption data, but pilot trials were too small to tell if this would be a big problem.

"Are you investigating how this data could be misused"

Mireille Hildebrandt said that inference of knowledge from data was not covered by data protection law.

Van Jacobson pointed out that data protection legislation was very different in the United States where consumers have very little protection. "As soon as it has left your house it is considered a matter of public record," he said.

A comment from the floor noted that information means market power. He argued that competition law would have to address this issue because any company in possession of personalised data

effectively owns a relationship with the consumer and would make it much harder for the consumer to search for alternative options. Hildebrandt accepted that consumer protection would be an issue. "Privacy is not about hiding data, but having control over how you give what data to whom," she said.

A participant pointed out that the internet actually started as a test-bed some 40 years ago, but has grown far beyond anything imaginable back then. They wanted to know what policies will ensure that such growth is possible in the Future Internet, but at the same time respect out privacy? Martin Brynskov responded by saying that the Future Internet would establish its own barriers of privacy and acceptability. "Today, people are happy to reveal their darkest secrets on Facebook. That would be unimaginable before. I think new barriers will fall into place."

"[The] internet actually started as a test-bed some 40 years ago, but has grown far beyond anything imaginable"

A delegate involved in the European Network of Living Labs observed that he had heard a lot from the conference that seems to be a prescriptive approach to the Future Internet, imposing solutions from the top down. He asked about opportunities for input from the bottom up. Hildebrandt backed up this view with an analogy to legislators either imposing new rules or actually listening to citizens. "We need to create an incentive structure that allows citizens to get involved in organising themselves. Today, and more so in the future, the very infrastructures in society determine how we live our lives."

From Bled to Valencia: Future Internet support actions

Michael Boniface, University of Southampton

"We have made significant progress since those early days of Bled," enthused Michael Boniface. "We have a community which is broad and cross-disciplinary. We have a strategic map, discussions, position statements, white papers and more. But we want more."

Boniface revealed how the FIA and Europe's Future Internet research community had access to significant resources to support their work. There are now several EU-funded specific support actions (SSAs) which will be helping the Future Internet community to clarify the European Future Internet vision, identify future research challenges and build closer links with FP7 projects.

Boniface mentioned some of the support on offer to FIA, for example:

- revitalising the Future Internet portal, especially looking to create quality content that will showcase the results of the core of the FIA – the more than 100 projects;
- publishing a vision for a Future Internet architecture (the European Reference Model);
- road-mapping activities;
- facilitating international collaboration and standardisation activities.

Boniface also asked delegates to play their part in helping to shape FIA Ghent. They could vote on workshops and discussion topics – or submit their own – over the web (www.bit.ly/fia-ghent-topics). Voting has now closed; 33 ideas were submitted and 206 people cast a total of 824 votes.

Towards FIA Ghent

Wim De Waele, IBBT

With FIA Valencia riding so high, FIA will once again piggyback on another related conference The Future Internet Conference Week, to be held in Ghent, Belgium in December 2010. Wim De Waele of the IBBT in Belgium said that the week of activities and events promised to be the biggest Future Internet event yet to be held in Europe with more than 2000 delegates expected to attend.

FIA would share its plenary sessions with the iMinds conference, including a keynote speech from Ben Verwaayen, CEO of Alcatel-Lucent. De Waele promised stimulating presentations, discussion and debate. Or following the sentiments of French singer songwriter Jacques Brel: "it won't be business as usual".

FIA Ghent conference website: http://www.fi-ghent.eu.

Closing Messages

Luis Rodríguez-Roselló from the European Commission took the podium to round off the event. He thought it was the best FIA yet and had greatly benefited from broadening its audience and introducing researchers involved at the application layer of the Future Internet.

"One of the recurrent themes that seems to have come through much of what I have heard these past two days is the importance of users," he said. "We need to find ways to get them involved."

He also commented that Future Internet research was maturing with the launch of the FI PPP. This initiative, along with the support actions, will play in important role in the future.

"There has been some criticism in the past that our work is not tangible enough," Rodríguez-Roselló admitted. "I'm now glad to see that the work of the FIA is becoming more concrete through the work of the specific support actions. The idea is to prevent the formation of new closed circles, but open out the debate and discussions though cross involvement with other activities."

Rodríguez-Roselló thanked the Spanish Presidency for hosting the event, his colleagues in the European Commission and the session organisers and caretakers.

"I was also going to wish you a safe return home," he ended, "but it looks like some of us will be going nowhere for a while!"

"The plane is there, we just now need to get it flying"

Mário Campolargo closed FIA Valencia with some remarks about Future Internet research in Europe coming of age. "I sense that this assembly has been a turning point, and not just because FI PPP is now consolidated at the political level. I think from here we have a collective mandate to make FI PPP come alive and complement the work of Challenge 1. We must not lose the goal, to make European society more dynamic but with a sustainable future."

"The plane is there, we just now need to get it flying!"