EUROPEAN RESEARCH ON FUTURE INTERNET DESIGN

PETER STUCKMANN AND RAINER ZIMMERMANN, EUROPEAN COMMISSION

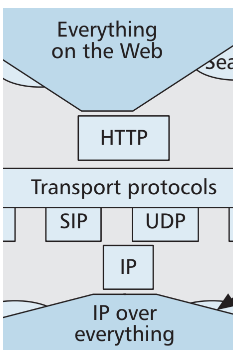

Future Internet is the federating theme for European research on communication networks and services. In the core lies research on communication networks towards an efficient, scalable, and reliable Future Internet coupled with research on the underlying technologies.

ABSTRACT

Future Internet has become the federating theme for European research on communication networks and services. At the core lies research on communication networks toward an efficient, scalable, and reliable future Internet coupled with research on the underlying technologies, in particular mobile and wireless access and optical networks. This article first presents the motivation for a bold initiative for future Internet research in Europe. In this context a changing business environment for telecommunications and the Internet, and the opportunities to provide future Internet services are discussed. From a technical perspective the limitations of the current Internet technology are outlined. The research activities that address the challenge of future Internet research are introduced under three main lines: future Internet architecture and network technologies, spectrum-efficient access to future networks, and converged infrastructures in support of future networks. Examples of the first promising approaches to significantly change the principles of the Internet architecture and protocols are presented.

INTRODUCTION

The European Union is presently funding collaborative research and development activities in the 7th Framework Programme (FP 7), which covers the period of 2007–2013 [1]. In the ICT Work Programme 2009–2010 [2], Future Internet has become a federating theme for European research on communication networks and services. This focus on future Internet is motivated by the changing requirements to the current Internet that was designed in the 1970s to support communication between computing systems for communities of expert users. It was not designed to cope with the wide variety and ever growing number of networked and mobile users and applications, business models, edge devices, networks, and environments that it now has to support. Its structural limitations are increasingly being recognized worldwide.

Clean slate or evolutionary approaches or a mix of those can be equally considered. The evolutionary approach builds on the evolution of the current existing Internet to conceive pragmatic and viable solutions for commercial rollout [3]. A clean slate approach eliminates legacy Internet design constraints. Both approaches target the same usage vision and will have to be synchronized on phased agendas. Once promising clean slate architectural concepts have been identified and evaluated as fulfilling the design goals, evolutionary paths for deployment have to be identified. Since the current Internet has grown to become so large, it will be commercially and operationally very challenging to introduce new architectural principles [4]. One deployment possibility is virtualization, which would enable logically independent networks built on a common physical infrastructure. In this case new network functionalities and protocols could be deployed, but also specialized networks could be provided by building overlay or underlay techniques running new protocols on top of or below the network and transport protocols. In this transition scenario a parallel network can be run for applications that truly need the improved functions. Users or providers would migrate to the new system over time, similar to the way some are now abandoning the traditional telephone system for Internet-based phones, even as the two networks run side by side.

A big challenge in clean slate research on future Internet is the need to evaluate the designed architecture. Whereas analysis and simulation can allow a first performance evaluation, the use of prototypes is crucial, as a system has to be built in order to evaluate it and to convince others that it is the appropriate solution. For global systems such as the Internet, it is almost impossible to get a new idea adopted that has not yet been tried at scale and under realistic conditions. In addition, an experimental facility would be useful intellectually as it enables researchers to uncover things that would otherwise have been omitted [4].

The views expressed in this article are those of the authors and do not necessarily reflect the official view of the European Commission on the subject

Besides the necessities and opportunities that lie in network research, key innovations in underlying technologies such as mobile and wireless communications and optical networks are expected to drive the development of the future Internet. Therefore, these aspects have to be studied jointly with network architecture research as they build together the future network infrastructure foundation. Modern communication infrastructures will be characterized by mobile and wireless broadband access on the last mile or meter and the interconnection to ultra-high-capacity all-optical networks.

Changing Application Requirements for the Future Internet

The Internet has emerged as a critical infrastructure for society and the economy as a whole, similar to any other utility (e.g., infrastructure for electricity and water supply). Society is undergoing a paradigm shift, the evolution of the society and the Internet now being tightly interconnected. Daily life factors including health, transport, knowledge, and culture rely increasingly on the Internet in the developed world, and it is bringing economic development to emerging economies.

In addition to 1 billion fixed Internet hosts expected to be connected by 2011, it is expected that mobile and other types of handheld devices will be directly connected to the Internet, leading to about 3 billion connected hosts by 2011 [3].

Since the Internet was designed for fixed terminals, it shows inefficient behavior for mobile and nomadic terminals. Therefore, generic and efficient support of mobile terminals and mobile applications is one of the major design goals of future Internet architectures and technologies.

Besides the increase in number of users and connected devices expected, new application requirements for the future Internet are emerging.

High-Quality and Shared Content Dissemination

As digitalization of data progresses, it is now expected that the majority of new media will arrive in digital form, with the analog form being the exception. For instance, digital videos will not only increase in number, but also in size, due to increases in resolution and the ease of creation and manipulation. The increase in number of digital videos and their distribution over an increasing number of locations creates the need for specific multimedia search engines. Progress in network multimedia communication is also leading to 3D videos, virtual reality, and gaming. Digital TV channels are also progressively penetrating the Internet space.

Not only high-quality content is made available over the Internet by large content providers. Users are also enabled to easily produce, offer, share, and consume content on the Internet, and are becoming *prosumers*. Whereas communication will remain an important service to be supported by the global network, dissemination of content, either distributed by content providers or made available by prosumers, is expected to be one of the main functions of the future Internet.

Connecting Objects and Things

While the current Internet is a collection of rather uniform devices, it is expected that the *Internet of Things* will be characterized by a much higher level of heterogeneity, as objects totally different in terms of functionality, technology, and application fields are expected to belong to the same communication environment. The Internet of Things can be defined as "a worldwide network of uniquely addressable and interconnected objects, based on standard communication protocols." This enables applications involving real-world objects, but also business applications based on network-assisted machine-to-machine interaction.

Service-Oriented Internet

Whereas a lot of computing and storage applications are still executed locally on end-user devices such as PCs, a service-oriented Internet would allow access to complex physical computing resources, data, or software functionality in the form of services. One example is the *cloud computing* approach to infrastructure services, where large-scale data centers provide virtual execution and storage environments as Internet services with the same functionalities as physical machines but far greater flexibility and scalability.

These sets of expected applications for the future Internet are only examples; the target architecture should generically and flexibly support different requirements and traffic patterns.

Future Internet Service Provisionioning

Today the telecom and Internet industries are structured differently with regard to applications. On one hand, in the telecom industry applications are network-infrastructure-centric, including many standards. Applications may look similar but are expected to be fully interoperable. On the other hand, in the Internet industry service providers are mostly answering a customer need without paying attention to interoperability and relying on lightweight infrastructure. Network infrastructure providers (Internet service providers [ISPs]) get their main revenue from end users who pay for network connectivity. Service providers (e.g., Google) get their main revenue from advertisers. In the past user micro-payments have proven to be too unpredictable and too much of a burden to be acceptable to users. History has shown that users usually prefer flat rates or subscriptions.

The future Internet shall provide a twofold path: enabling focus on a specific user need and developing a solution without paying attention to the network infrastructure. Once a leader on a specific application emerges, the future Internet should also enable the growth of an ecosystem of players able to build their own applications on top of its system.

While the current Internet is a collection of rather uniform devices, it is expected that the Internet of Things will be characterized by a much higher level of heterogeneity.

It is important that the Future Internet is designed to accommodate conflicting interests, so called tussle networking such as conflicting policies, traffic patterns, and compensation modes.

In the future Internet, access to the network will be made available ubiquitously, and connectivity will become a fundamental service that communities use and rely on. The transmission technologies and access networks will increasingly be designed and deployed for horizontal integration and service-agnostic platforms, as addressed under the *technology neutrality* paradigm where services can be provided through different networks and different technologies.

Bit Pipes and Advanced Network Services

The provision of a *bit pipe* is becoming a commodity business as end-user services are separated from the network infrastructure. However, based on the pure data forwarding service provided by the communication infrastructure, various kinds of advanced network services can be established. While the focus of today's Internet is mainly on non-real-time messaging, Web browsing, and multimedia, it is expected that new applications will demand new capabilities from the networks. Bulk data transfers, real-time data transmission, and information delivery are examples of high interest, especially for industrial applications. Interconnecting more than computing machines (e.g., sensor networks) will push connectivity needs that have nowadays reached the limit of the original design. The network operator is also well positioned to provide context information to service providers for future innovative applications (e.g., based on location information). For content dissemination, already today video download service providers are buying content distribution network (CDN) services from infrastructure providers in order to be able to provide their content at acceptable quality. As user numbers and the amount of content are expected to rapidly increase, this kind of support service for efficient content dissemination is promising to have further potential.

Controlling the Internet

Today's Internet control architecture implicitly restricts business to a set of primitive models such as settlement-free peering or flat-rate subscription. Any more sophisticated approach has to be either implemented as a *walled garden*, losing the benefits of global interconnectivity, or built by abusing the current architectural components, using ad hoc solutions for provider independence [5].

It is important that the future Internet is designed to accommodate conflicting interests, so called *tussle networking* [6], such as conflicting policies, traffic patterns, and compensation modes. It is fundamental to recognize the powerful capability of the current Internet to accommodate new applications developed by a growing user community. It is thus essential to keep the entry barrier as low as possible and design the future Internet so as to allow and steer open and innovative application development without impeding the Internet's genericity, evolvability, openness, and accessibility. The future Internet shall thus cultivate the opportunity for new players to benefit from not only the infrastructure foundation but also the pillars of the future Internet without sacrificing its global architecture objectives and principles. A future Internet should be able to adapt to the changes in society's demands on the Internet as they occur without requiring permanent redesign. Additions and extensions to the network architecture should be facilitated without rigorous standardization processes but also without replacement of the infrastructure equipment.

This approach of *design for tussle* is also motivated by today's tussle over resource control by users and network providers. Network providers offered several pricing schemes, but flat rate prevailed. Users of peer-to-peer applications took advantage of that and, by increasing the number of their parallel transport connections, absorbed all the bandwidth of the network, making performance dismal for traditional interactive Web browsing users. To prevent congestion in their networks, operators introduced techniques such as port-selective packet dropping and deep packet inspection, while users reacted by encrypting their traffic.

Besides filtering of traffic, operators are also starting to use differentiated charging for transporting content. Operators want to charge higher fees for traffic belonging to premium customers. Advocates of net neutrality see this to be against the principle that all Internet traffic should be treated equally. The principle is seen as the basis for the preservation of current freedoms, and the Internet as an engine for innovation and creativity where business services and user-generated content can easily be offered and accessed with global reachability.

This discussion is closely linked to the governance aspect that will become an important part of the future Internet, particularly if it comes to issues with a public policy component addressing aspects like privacy, security, freedom of expression, intellectual property rights, and data protection. The term *Internet governance* emerged in the early 1990s to describe elements of the management of critical Internet resources, in particular the domain name system (DNS), but also IP addresses, root servers, and Internet protocols. It was based on a conceptual understanding of self-regulation and private sector leadership, of *governance without government*. One concrete result of this discussion was the establishment of the Internet Corporation for Assigned Names and Numbers (ICANN) in 1998 and, as the next step, the United Nations Working Group on Internet Governance (WGIG) in 2003 and the Internet Governance Forum (IGF) in 2006 [7].

Technical Limitations of the Current Internet

Internet technology can be characterized by its design principles: layering, packet switching, a network of collaborating networks, intelligent end systems, as well as the end-to-end argument [4, 8].

It has been simple to link any new network to the Internet, providing instant benefits resulting from the interconnectivity with a huge range of communicating peers. The transparency of the Internet has facilitated the deployment of successively more complex network-agnostic applications and services. Together, these two attributes characterize the *hourglass* approach to protocol architecture — the network layer, the waist of the hourglass, is simple enough to operate on top of and integrate with any link technology, and transparent enough that anything can run over it (Fig. 1). The hourglass approach has led directly to a virtuous circle of increased network reach enabling new styles of usage and vice versa.

However, with its increasing success, the Internet architecture is progressively losing its original simplicity and transparency. Some of the main causes are the rise of new classes of applications, additional operational and management requirements, variety of business models, security mechanisms, and scalability enablers that give rise to ad hoc solutions that extend the architecture without regard to the original key design principles.

Examples are firewalls to support end-user and site security, and Network Address Translation (NAT) to cope with the exhaustion of IPv4 address space. There is, however, a growing consensus among the scientific and technical community that the methodology of continuously *patching* the Internet technology will not be able to sustain its continuing growth at an acceptable cost and speed.

This loss of flexibility is already being felt as the number of Internet nodes grows another order of magnitude. The size and scope of today's Internet make the deployment of new network technologies very difficult while experiencing increasing demand in terms of connectivity and capacity. This situation, where technological innovation meets natural resistance, is called *ossification*. Examples are the slow deployment of technologies such as multicast or Internet Protocol version 6 (IPv6). Innovation has happened mainly in the applications and underlying transmission technologies, rather than in the core technology, the network and transport layers TCP/IP (Fig. 1). The following technological limitations have to be overcome to meet the future challenges for the global communication network.

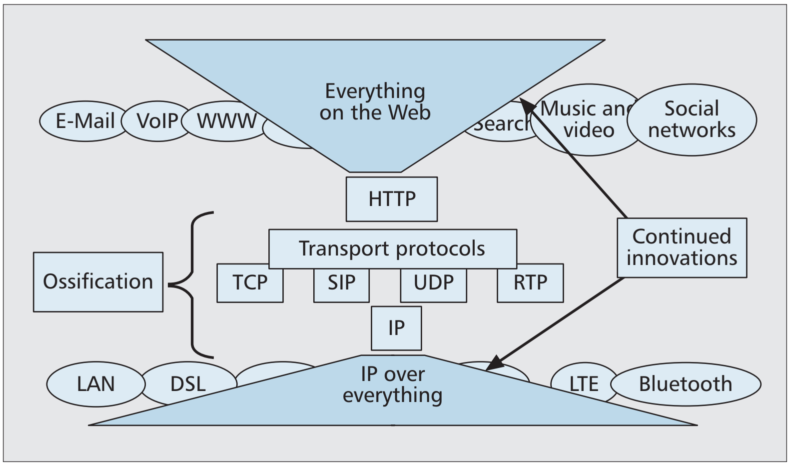

Figure 1. *Innovation and ossification in the Internet.*

Limited Support of Mobile and Wireless Terminals

While the Internet was designed for stationary computers, today laptops and smart phones are constantly on the move. With today's technology, a laptop changes its address and reconnects as it moves from one wireless network or access point to another, disrupting the data flow. Alternatively, the Internet standard Mobile IP allows routing all traffic back to the first access point as a laptop moves to a second or third location, but delays and inefficiencies may result. As a clean slate solution, the address system would have to be restructured so that addresses are based more on the device and less on the location. This way, a laptop could retain its address even if it moves from one wireless network or access point to the other.

Lack of Built-in Security

With the evolution toward Internet-based services, traditional telecom networks as backup for mission-critical services are expected to gradually disappear. Therefore, built-in security mechanisms are one of the main goals of future Internet design.

The Internet was designed to be open and flexible, and all users were assumed to be trustworthy. Thus, the Internet protocols were not designed to authenticate users and their data, allowing spammers and hackers to easily cover their tracks by attaching fake return addresses onto data packets. Internet applications such as firewalls and spam filters attempt to control security threats. But because such techniques do not penetrate deep into the network, bad data still get passed along, clogging systems and possibly fooling the filtering technology.

The network would have to be redesigned to be skeptical of all users and data packets from the start. Data would not be passed along unless the packets are authenticated. Faster computers today should be able to handle the additional processing required within the network.

Scalability Issues

In addition, the vision is also under threat from basic engineering problems. The routing system, which is the single most critical part of the Internet infrastructure, is facing significant scalability issues [4]. Just when the Internet is becoming critical infrastructure, the core protocols may become increasingly fragile as more manual configuration is needed to avoid cascading problems due to overload, accidental misconfiguration, or attack. The IPv6 standard allows expanding the address pool, but nearly a decade after most of the standard was completed, the vast majority of software and hardware still use the older IPv4 technology. Even if more migrate to IPv6, not all addressing issues would be solved. Researchers are questioning whether all devices truly need addresses. Sensors in a home could communicate locally and relay the most important data through a gateway bearing an address. As routing and addressing are becoming the main challenge of a global network with billions of users and objects, scalability is one of the most critical design criteria for future Internet architectures.

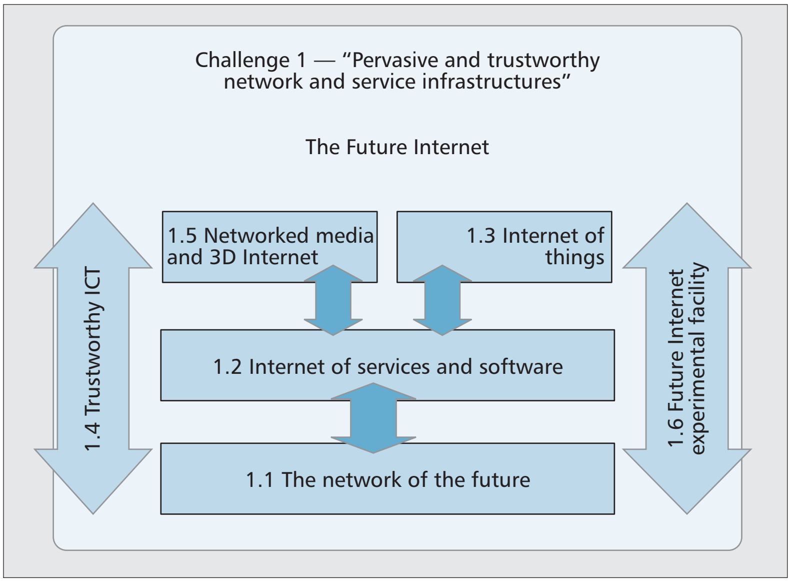

Figure 2. *Future Internet as a federating theme.*

Performance and Quality of Service Challenges

While mechanisms for providing quality of service (QoS) within the Internet as well as asynchronous transfer mode (ATM) networks have been very well studied, the interaction problems between the network layers are still unresolved, and the management of such services, including configuration, policy setup, charging, and interprovider setups, is still a challenge [4].

Research Program on Future Networks

In the ICT Work Programme 2009–2010 [2], future Internet has become the federating theme for European research on communication networks and services. The planned activities under the heading "The Network of the Future" can be seen as the basis for future network infrastructure foundations enabling future Internet services and applications that are complementary, funded under the headings "Internet of Services," "Internet of Things," and "Media Internet." In addition horizontal activities like "Trustworthy ICT" and "Future Internet Experimental Facility" are complementary parts of the program (Fig. 2). Together with the ongoing activities under work program 2007–2008, the total funding for Future Internet research lies in the order of €800 million, roughly half of which is allocated to research on future networks. The ongoing research activities and the target outcome of future projects are structured under the following three major lines.

Future Internet Architectures and Network Technologies

The main goal of this activity is overcoming structural limitations of the current Internet architecture resulting from an increasingly larger set of applications, devices, and edge networks to be supported. The first target outcomes are novel Internet architectures and technologies enabling dynamic, efficient, and scalable support of a multiplicity of user requirements and applications with various traffic patterns, variable end-to-end QoS, point-to-point or point-to-multipoint distribution modes, and supporting legacy and future service architectures. The target architecture should support personalized rich media networking, machine-to-machine communication, wireless sensor networks, ad hoc connectivity networks as well as personal and body area networks. It should also be wireless-friendly, natively support mobility, be spectrum- and energy-efficient, and support future very-high-data-rate all-optical connections as well as heterogeneous wired/wireless access domains. Routing and location-independent addressing or naming, dynamic peering, signaling, resource virtualization, and end-to-end content delivery techniques are related research issues.

Besides the architectural concept, flexible and cognitive network management and operation frameworks are needed. They are expected to enable dynamic, ad hoc, and optimized resource allocation, control, and deployment, and administration with accounting that ensures both a fair return on investment and expansion of usage, differentiated performance levels that can be accurately monitored, fault tolerance, and robustness associated with real-time troubleshooting capabilities. The management architecture should target self-organized and self-healing operations, cooperative network composition, service support, and seamless portability across multiple operator and business domains.

Migration paths and coexistence through overlay, federation, virtualization, and other techniques should be investigated to support several network and management architectures including legacy systems. Benchmarking capability of the proposed architectures is to be considered from the onset. Clean slate or evolutionary approaches or a mix of those can be equally considered.

Spectrum-Efficient Radio Access to Future Networks

One of the important infrastructure foundations of the future Internet will be next-generation mobile radio technologies. They should be designed to be cost-, spectrum-, and energy-efficient, and adapted for implementation in future high-capacity mobile radio systems. Key technology building blocks expected to be addressed are adaptive modulation and coding schemes, multiple antenna and user detection schemes, cross-layer design and low-latency transmission schemes. They are expected to be complemented by cooperative technologies at the base station and/or terminal level, novel network topologies, and related dynamic channel modeling and estimation. Integrated projects are expected to take a comprehensive approach to the key technology building blocks and develop system evolution paths by jointly designing radio transmission techniques and radio interface protocol stacks, and considering spectrum coexistence and sharing.

Cognitive radio and network technologies are to be developed in order to reduce the management complexity and enable seamless service provisioning in a radio environment with a large number of heterogeneous radio access technologies. These should support environment-aware, self-reasoning- and learning-capable mobile devices that can change any parameter or protocol based on interaction with the environment, with or without network assistance.

As a complement to cellular systems, novel radio network architectures enabling the innovative usage of licensed, unlicensed, or unused radio spectrum with the aim of radical cost and energy reduction are to be researched. Target environments range from short to medium distance including systems based on femtocells, ad hoc networks, and vehicular networks, up to wide-area terrestrial and satellite-based radio access networks.

Converged Infrastructures in Support of Future Networks

The second important infrastructure foundation that comes before the last mile or meter encompasses ultra-high-capacity optical transport and access networks. They are expected to be based on state-of-the-art photonics with transparent core-access integration, optical flow and packet transport, dynamic wavelength allocation, and end-to-end service delivery capability. They should overcome the limitations of segmentation between access, metro, and core networks and domains, enable lower-cost optical access, and address the need for energy efficiency. Integrated projects are expected to also address a network control plane supporting flexible management capability of multidomain and multi-operator contexts with end-to-end carrier grade performance.

Another objective is converged service capability across heterogeneous access. Concepts should be developed that go beyond incremental steps in service platforms. Breakthrough technologies and architectures are needed for seamless ubiquitous broadband services, integrating wired and wireless, fixed, and mobile technologies in hybrid access networks, including hybrid satellite networks. These enable generic support for service portability and continuity across composite networks through the service-network interface, with ubiquitous access from any network, from any technological or administrative domain, from any location and with a variety of access devices.

Future Internet Research Coordination

While in future Internet research a bottom-up approach is needed to generate ambitious and long-term ideas, concepts, and solutions, coherence is needed to generate critical mass. The coordination structure for European future Internet research is characterized by three pillars (Fig. 3). Academia is expected to play an important role in clean slate thinking, and major contributions are expected toward new architectural concepts. After an initial expert consultation in March 2006 [9], the EIFFEL think tank was established in October 2006 for discussing and exchanging ideas and research trajectories on the future of the Internet architecture and governance building as a foundation of the future networked society [10]. To have a significant impact, industry leadership is needed, and the program was already successful in the first call in 2007, attracting the major industry players to take up the challenge of future Internet research. The European technology platforms in the area of networks, satellite communications, electronic media, software and services, and smart systems have joined forces and drawn up a common vision toward a public-private partnership on future Internet technologies [3]. To provide a forum for the ongoing projects to exchange results and develop synergies, the Future Internet Assembly (FIA) was established, especially focused on cross-domain discussions to bridge the gap between the different research communities of networks, services, content, and network security. The Bled conference in March 2008 was the kick-off of the FIA, where the Bled declaration [11] was signed by 63 European research projects and European technology platforms to express their commitment and call for European action toward the future Internet. A first milestone has been the publication of the book *Towards the Future Internet — A European Research Perspective* presenting the first project results [12].

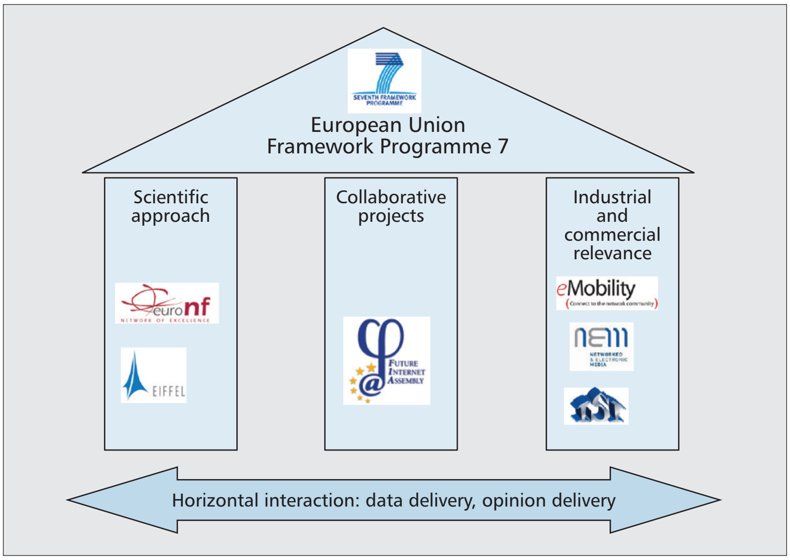

Figure 3. *Structure of future Internet research coordination.*

First Approaches to Future Internet Design

Research on future Internet architectures and technologies has already started. Both industry-driven integrated projects generating critical mass and focused projects concentrating on specific ideas have been launched in the first phase of FP7. The project portfolio is structured in three clusters: Future Internet, Radio Access and Spectrum, and Converged and Optical Networks (Fig. 4). In the following, examples of the first promising approaches to significantly change the principles of the Internet architecture and protocols are presented.

The current semantic overload of the IP address as both node identifier and locator, indicating the current point of attachment in the network topology, is replaced by a clear separation of information self-certifying object identifiers and locators.

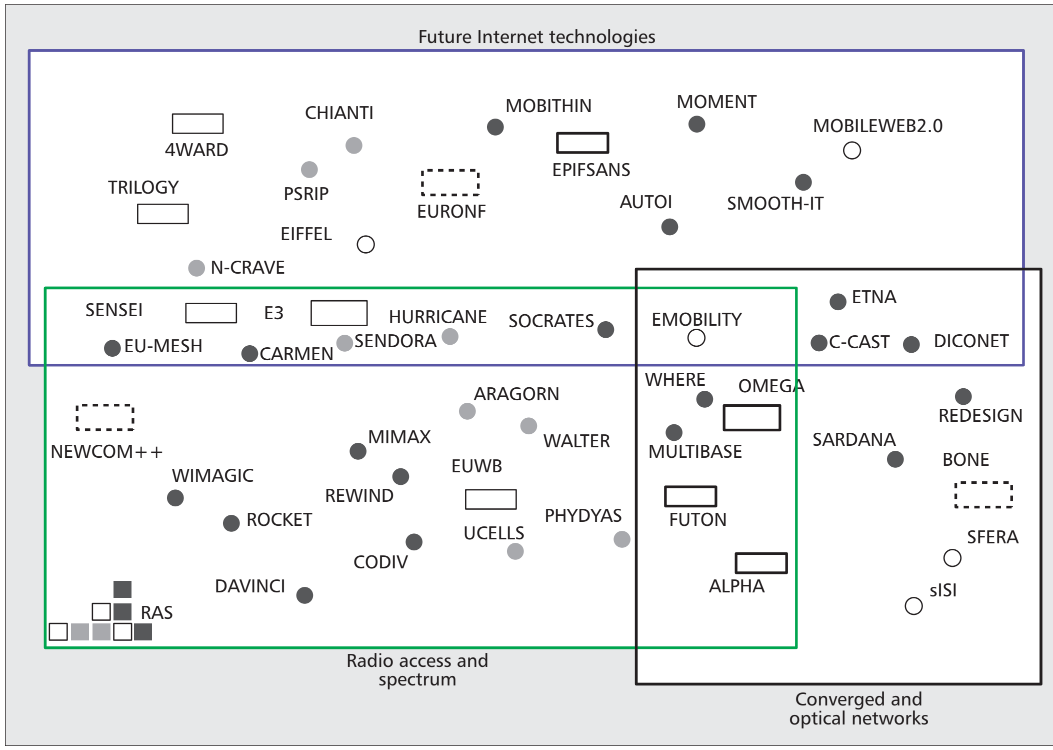

Figure 4. *Future networks project portfolio and clustering.*

Information-Centric Paradigm and New Addressing Concepts

While today's Internet is based on a node-centric paradigm, several projects are adopting an information-centric paradigm. In this paradigm, the communication abstraction presented to applications is based on transfer of application data objects instead of the end-to-end reliable byte stream used by the majority of applications today.

Networking across Multiple Paths and Layers with the Generic Path — The project 4WARD [13] is integrating much of the functionality of peer-to-peer overlays by including caching functions where the "copies" are treated as the originals. This is done in a common and open information networking service generalized for use by applications.

The current semantic overload of the IP address as both node identifier and locator, indicating the current point of attachment in the network topology, is replaced by a clear separation of information self-certifying object identifiers and locators. While previously proposed models for abstracting the location and focusing on networking between (mobile) hosts are considered, the project is designing a network architecture where mobility, multihoming, and security are intrinsic rather than add-on solutions.

The notion of a generic path (GP) is defined as a framework able to efficiently realize *networking of information* by exploiting cross-layer optimization and multiple network paths. A GP is defined as a means to organize the accessibility of a sufficient number of parts or copies of information objects stored in a group of hosts. Incorporating the paradigm of information-centric networks means that a GP is actually hiding the physical location of information objects. It does not matter where chunks or copies of information are stored; the GP takes care of delivering it to the destination. Because cross-layer information is available, new transmission techniques can be used inside a GP. This is especially interesting for the introduction of network coding into fixed and wireless networks. Here, multipath routing needs to be combined with specific capabilities of nodes. An important advantage of this concept is that mobility of information objects and hosts becomes conceptually equivalent and is dealt with by the GP internally.

The Publish/Subscribe Routing Paradigm — The PSIRP project is redesigning the entire Internet architecture based on the publish/subscribe routing paradigm [14]. In this already existing concept, senders *publish* what they want to send and receivers *subscribe* to the publications they want to receive. In principle, no one receives any material in which they have not explicitly expressed an interest by way of subscription. By explicitly subscribing to information, subscribers offer their resources to be *used by the sender*, which means that resources are added to the communication in a scalable fashion as the number of subscribers grows. In such a new Internet, multicast and caching will be the norm, and security and mobility will be designed directly into the architecture, rather than added as afterthoughts.

PSIRP's work is focused on the intersection of security, routing, wireless access, architecture design, and network economics, in order to design and develop efficient and effective solutions. The architecture also addresses multicast and caching functionalities to achieve optimal operational efficiency, and embody security and mobility directly into the foundations of its design. Moreover, a system of identifiers and labels, combined with methods for scoping information, allows information to be addressed as opposed to end hosts, effectively providing a locator-identifier split. The result is a powerful yet flexible infrastructure with a high degree of resilience. The new pub/sub-based internetworking architecture will restore the balance of network economics incentives between the sender and the receiver, and is well suited to meet the challenges of future information-centric applications and use modes.

The conceptual architecture is based on a modular and extensible core, called the PSIRP component wheel. The architecture does not have the traditional stack or layering of telecommunications systems, but rather components that may be decoupled in space, time, and context.

Architecture Based on Reachability and Forwarding Planes

The Trilogy project is developing a baseline architecture [6], which is comparable in scope to the current Internet network and transport layers, but with a subtly different internal structure. A fundamental assumption of the architecture is that it is based on a minimal packet delivery service. The ideal case is that packets are entirely self-describing, meaning that other concepts such as connections, flows, and sessions are higher-level constructs, invisible to the delivery service, and the network delivers each packet independent of every other. In particular, the functions involving the networking infrastructure are divided into two planes, for reachability and forwarding, and are distinguished from the transport services to which the network is totally transparent.

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

Distinct from the packet delivery service, the functions that are implemented are identified in a pure end-to-end fashion. Transport services include functions, such as reliability, flow control, and message framing, that are totally invisible to the packet delivery service. The reachability and forwarding planes are separate and together achieve the packet delivery service. The key identifier space is the destination locator, which is handled in detail only by the reachability plane. The forwarding plane treats locators as opaque tags that only have to be tested for equality in order to test for path consistency. The transport services use endpoint identifiers that label the communicating parties, but these are totally independent of the locators used in the reachability plane.

For resource control, an accountability framework is proposed that is based on congestion volume. Accountability enables a rational basis for sharing resources among users on a future Internet that is a playground for competing users, applications, and businesses. In addition, the concept of resource pooling is introduced, which enables separate network resources to behave like a single large pooled resource. Resource pooling enables better resilience and efficiency.

Conclusions

Future Internet has emerged as the federating scheme for European research on communication networks and services. Besides the focus on network science, underlying technologies such as mobile and wireless access and optical networks remain the key drivers for future Internet developments. The main challenges for future Internet design are support of mobile broadband applications, manageability, and scalability including QoS support, security, and trustworthiness as well as support of advanced high-quality content, including 3D applications.

A first set of projects has been launched, and the first promising concepts are under development. Coordination is ensured by providing platforms for interproject cooperation to ensure coherence and the generation of critical mass. Increased coordination effort is needed to further build the Future Internet Assembly. Cooperation with EU national initiatives should be reinforced, and international cooperation with regions having future Internet initiatives is desired.

In the next phase, projects should continue developing the technological and architectural foundations of the future Internet. In addition, an integrated approach to the individual innovations is needed. Conceptual results should be evaluated and implemented in prototypes. The evaluation should show to what degree design criteria and design goals are met. Benchmarking of architectures is needed to compare different approaches. Once architectural concepts are evaluated as promising, evolution paths need to be studied, including assessment of their risks.

In the next phase, projects should continue developing the technological and architectural foundations of the Future Internet. In addition, an integrated approach to the individual innovations is needed. Conceptual results should be evaluated and implemented in prototypes.

References

[1] P. Stuckmann and R. Zimmermann, "Toward Ubiquitous and Unlimited-Capacity Communication Networks: European Research in Framework Programme 7," *IEEE Commun. Mag.*, vol. 45, no. 5, May 2007, pp. 148–57.

[2] European Commission, "ICT Work Programme 2009–2010," version agreed by program committee, Nov. 2008; http://cordis.europa.eu/fp7/ict

[3] D. Papadimitriou, "Future Internet — The Cross-ETP Vision Document," v. 1.0, Jan. 2009; http://www.future-internet.eu

[4] A. Feldmann, "Internet Clean-Slate Design: What and Why?," editorial note submitted to CCR, June 2007; http://ccr.sigcomm.org/online/?q=node/229

[5] L. Burness, P. Eardley, and R. Hancock, "The Trilogy Architecture for the Future Internet," in *Towards the Future Internet — A European Research Perspective*, G. Tselentis *et al.*, Eds., IOS Press, 2009, pp. 79–90.

[6] D. D. Clark *et al.*, "Tussle in Cyberspace: Defining Tomorrow's Internet," *Proc. ACM SIGCOMM*, 2002.

[7] D. Hausheer *et al.*, "Future Internet Socio-Economics — Challenges and Perspectives," in *Towards the Future Internet — A European Research Perspective*, G. Tselentis *et al.*, Eds., IOS Press, 2009, pp. 1–11.

[8] D. D. Clark, "The Design Philosophy of the DARPA Internet Protocols," *Proc. ACM SIGCOMM*, 1998.

[9] P. Stuckmann *et al.*, "Communication Networks of The Future," consultation mtg. rep., CORDIS, Mar. 2006; ftp://ftp.cordis.europa.eu/pub/ist/docs/ct/report-communication-networks_en.pdf

[10] P. Mähönen *et al.*, "EIFFEL: Evolved Internet Future for European Leadership," white paper, EIFFEL Think Tank, Dec. 2006; http://www.fp7-eiffel.eu/fileadmin/docs/EIFFEL-FINAL.pdf
[11] Bled Declaration; http://www.fi-bled.eu/Bled_declaration.pdf
[12] G. Tselentis *et al.*, Eds., *Towards the Future Internet — A European Research Perspective*, IOS Press, 2009.
[13] H. Abramowicz *et al.*, "A Future Internet Embracing the Wireless World," in *Towards the Future Internet — A European Research Perspective*, G. Tselentis *et al.*, Eds., IOS Press, 2009, pp. 91–101.
[14] S. Tarkoma *et al.*, "The Publish/Subscribe Internet Routing Paradigm (PSIRP): Designing the Future Internet Architecture," in *Towards the Future Internet — A European Research Perspective*, G. Tselentis *et al.*, Eds., IOS Press, 2009, pp. 102–11.

Biographies

PETER STUCKMANN [M] (peter.stuckmann@ec.europa.eu) received his engineering degree in 1999 and his doctorate degree in 2003, both from the Faculty of Electrical Engineering and Information Technology of Aachen University, Germany. He is currently project officer with the European Commission, Directorate-General Information Society and Media, where he is managing the research cluster on mobile and wireless communications as part of the European research programme FP7. Before joining the European Commission in 2004, he occupied several engineering and managerial positions in both industry and academia in the area of radio interface engineering. He is the author of the text ook *The GSM Evolution* (Wiley, 2002), two book contributions, more than 20 journal and conference publications, and one patent. He is a member of the German VDE/ITG.

RAINER ZIMMERMANN (rainer.zimmermann@ec.europa.eu) graduated in engineering from the Technical University of Berlin in 1977 and obtained his Dr.-Ing. degree in 1985 from the same university. He worked from 1977 to 1985 as a researcher for the Technical University of Berlin and the Fraunhofer Society. He then joined the European Commission as a project officer for projects in the fields of production engineering, high-performance computing (1990), and software (1992). Since 1995 he has held positions as head of unit for telematics between Administrations-IDA, Software Systems and Services, Nanoelectronics, and Photonics. He is currently head of unit for Future Networks of the Directorate-General for Information Society and Media.