

Future Internet Research and Experimentation: The FIRE Initiative

Anastasius Gavras
Eurescom, Germany
gavras@eurescom.eu

Arto Karila
Helsinki Institute of Information
Technology, Finland
arto.karila@hiit.fi

Serge Fdida
University P&M Curie, France
serge.fdida@lip6.fr

Martin May
ETHZ, Switzerland
martin.may@ethz.ch

Martin Potts
Martel, Switzerland
martin.potts@martel-consulting.ch

The opinions expressed in this paper are purely those of the authors and do not necessarily reflect an official position from their companies or the European Commission.

ABSTRACT

The research community worldwide has increasingly drawn its attention to the weaknesses of the current Internet. Many proposals are addressing the perceived problems, ranging from new enhanced protocols to fix specific problems up to the most radical proposal to redesign and deploy a fully new Internet. Most of the problems in the current Internet are rooted in the tremendous pace of increase of its use. As a consequence there was little time to address the deficiencies of the Internet from an architectural point of view.

Within FP7, the European Commission has facilitated the creation of European expert groups around the theme FIRE "Future Internet Research and Experimentation". FIRE has two related dimensions: on one hand, promoting experimentally-driven long-term, visionary research on new paradigms and networking concepts and architectures for the future Internet; on the other hand, building a large-scale experimentation facility supporting both medium- and long-term research on networks and services by gradually federating existing and new testbeds for emerging or future Internet technologies. By addressing future challenges for the Internet such as mobility, scalability, security and privacy, this new experimentally-driven approach is challenging the mainstream perceptions for future Internet development. This new initiative is intended to complement the more industrially-driven approaches which are addressed under the FP7 Objective "The Network of the Future" within the FP7-ICT Workprogramme 2007-08. FIRE is focused on exploring new and radically better technological solutions for the future Internet, while preserving the "good" aspects of the current Internet, in terms of openness, freedom of expression and ubiquitous access. The FIRE activities are being launched in the 2nd ICT call, which closes in October 2007, under the FP7-ICT Objective 1.6 "New Paradigms and Experimental Facilities" (budget €40m). Projects are envisaged to start in early 2008.

Categories and Subject Descriptors

C.2.1 [Computer Communication Networks]: Network Architecture and Design

General Terms: Design, Experimentation

Keywords: Future Internet, Network Architecture, Experimentation, Testbed, Network Neutrality, Situated and Autonomic Communications

1. INTRODUCTION

The architecture of the Internet is approaching 30 years of age. In that time, it has succeeded beyond the wildest dreams of its creators in changing the way people throughout the world live, work, and play. The EU research and industrial communities have played an essential part in the international efforts that have led to the success of networking systems, not least under the successive ICT R&D programmes aimed at building the foundations of the Information Society.

This success has been achieved while the scale of the Internet has increased by 7 orders of magnitude. In order to cope with such growth, the simple, original Internet architecture (16 specification documents) has accreted several hundred additional protocols and extensions. Networks based upon this significantly more complex architecture are increasingly difficult to manage in a way that enables the expected qualities of service for the over 1 billion users to be met.

The increasing, and implicit, reliance on the Internet has stimulated a major debate amongst experts (e.g. [1]) as to whether the current architecture and protocol can continue to be patched, or whether it will collapse under the demands of future applications. Indeed, there are signs that the limitations of the protocol are already appearing in situations such as unusual (but legitimate) traffic load (e.g. flash crowds), delivery of real-time video services, high-mobility of nodes and sub-networks, requirements for multi-homing, weak, asymmetric, and episodic connectivity of wireless channels, unpredictable long delay paths either due to length (e.g. satellite, multi-hop, ad-hoc connections) or as a result of congestion, attacks against the network hardware, software, or protocol infrastructure failures due to mis-configuration or operational errors.

Additionally, the original Internet was designed and built in an era of mutual trust, probably due to the small size of the original community addressed by the DARPA project. Many of the protocol additions/extensions have been to retrofit protection mechanisms that are required in the current Internet environment, which does not merit mutual trust. The volume and types of attempts to subvert the Internet will continue to increase, further stressing the current architecture.

Furthermore, mobile network hosts are rapidly becoming the norm for the devices with which users access the Internet. An increasing number of the protocol additions/extensions have been needed to

retrofit support for mobility into the (initially wireline-focused) Internet architecture. The growing use of mobile sensors will continue to drive the need for solid mobility support in the architecture (and the efficient transfer of small data units).

In Europe, since several years, industrial and academic researchers and the Commission have recognised the need to overcome the intrinsic limitations of the current Internet architecture. Within the framework of FET (Future and Emerging Technologies) under FP6, the SAC (Situated and Autonomic Communications) initiative [2] and other future Internet related projects with an overall budget of ~€30m were started in 2004. These projects address several decentralised (i.e. autonomic or distributed) approaches for new Internet architectures and services, featuring self-organising properties.

In parallel, under the Objective on Research Networking Testbeds under FP6 [3], several testbed projects were started to pave the ground for a future generation of e-infrastructures for research (~€30m funding). These testbeds address (for example) issues like convergence and quality of services. They include two projects, started in 2006, which are envisaged to be instrumental in bootstrapping the experimental facility planned within FIRE: ONELAB [4] which addresses proof-of-concept testbeds for researchers, by extending to Europe the Planetlab approach, and PANLAB [5] which is deriving a framework for the federation of testbeds on a broad scale. There is a multitude of other projects under FP6, which carry out research on networks and services and which address aspects related to the evolution of the Internet; several of them include testbeds, which may be candidates for federation within FIRE. In addition, there are relevant projects being carried out with national or regional funding, including ones under the EUREKA CELTIC initiative.

Many networking researchers around the world have identified the emerging limitations of the current Internet architecture and agree that it is time for research to take a long-term view and to reconsider the basic architecture of the Internet, to see if any better architecture can be identified, even if it does not appear backward-compatible at a first glance.

To be effective and to produce applicable results, this long-term, fundamental research in new communication and networking paradigms has to be tested, at least as a proof-of-concept, in large-scale environments. The objective is to assess the feasibility of these new concepts, verify their large-scale effects (not only at the technological level, but also for their foreseeable implications on users, society and the economy) and derive further requirements, orientations and inputs for long-term research. This kind of experimentally-driven approach avoids that the long-term research will remain at the level of paperwork and will allow exploring dramatic improvements over the current Internet.

FIRE will create an environment promoting strategic, long-term research on new Internet concepts, while at the same time giving the opportunity to Europeans to carry out large-scale experimentation of new paradigms, in order to measure and compare results, give credibility to what would otherwise be only theoretical studies, prove concepts and learn from the experiments. It is possible, for instance, that the current Internet possesses features that we do not fully appreciate. These features would only become noticeable by their absence, once a 'Tabula Rasa' approach is underway. Experimental facilities will also allow for research into the utility of these "hidden" features.

The FIRE Vision

FIRE aims at addressing the emerging expectations which are being put on the Internet by providing a research environment for investigating and experimentally validating highly innovative and revolutionary ideas.

New technological solutions may follow either an incremental approach or a "clean slate" approach. While the first evolves a system from one state to another by implementing incremental patches, the latter demands a radical redesign to offer new abstractions and improved performance, and has no constraint of "backward compatibility". The latter approach might even prove to be disruptive in terms of technologies, services or business models. Whereas the FIRE experimental facilities are meant to be open to both types of approaches, the research carried out under FIRE is addressing the more visionary path; the outcome is more risky, but at the same time more likely to yield significant improvements. The more incremental research approaches are rather part of the research approach currently pursued in national and European programmes and should not be duplicated in FIRE.

It has to be stressed here that incremental and "clean slate" approaches are not competing, but complementary. Where in the short or medium term only incremental solutions can be envisaged, in the long term we have also to consider the possibility of fundamentally changing the Internet architecture or some of the underlying paradigms. Testing them in large-scale experimental facilities will be essential for proving their feasibility, for identifying potential migration paths and for assessing their possible technological and socio-economic impact.

By addressing future challenges for the Internet such as mobility, scalability, security and privacy, this new experimentally-driven approach is challenging the mainstream perceptions for future Internet development.

2. HOW IS FIRE IMPLEMENTED?

Long-term, visionary research on Internet protocols and architectures is expected, as a key and distinguishing core part of this initiative. The main characteristics of such research is its bottom-up, open approach: there should be no boundaries around the research, but rather freedom to address any emerging or radically new but promising concepts that address the fundamental limitations of the current Internet. This can span all layers of the communications protocol - and, of course, explore different paradigms which might not be based on layered models. There should be no backward-compatibility constraints from the outset.

Another key aspect of this research is multi-disciplinarity, as real innovation often comes at the intersection of different disciplines. Moreover, the Internet is a complex system that depends on a delicate equilibrium between technologies, users, services and applications. Evaluating carefully these interrelations will be key to harnessing and exploiting the full potential of the future Internet for the economy and society at large. For instance, bio-inspired principles can be applied to network design, to exploit redundancy, irregularity (random elements) and noise to build a more dynamically robust system. Research can also explore the means by which networks can change, learn and develop spontaneously, thus staying ahead of the curve and perhaps breaking the vicious reactive cycle whereby, for example, security infringements are addressed retrospectively.

In general, this can be labelled as "experimentally-driven long-term research". Large or small research projects in this area should base their research theories and expected results on testbeds and experiments. Building such dedicated testbeds can be in the scope of the research projects themselves, or the projects can envisage exploiting existing or planned testbeds, which will be part of the FIRE experimental facilities under construction. This should be seen as part of a new research methodology in Internet concepts, tightly coupling research and experimentation of long-term, potentially disruptive new ideas. The other dimension of the FIRE initiative is an experimental facility for future Internet technologies which supports both medium-term and long-term research on networks and services and which is gradually built by federating existing and new testbeds, in order that they become an integrated, sustainable, dynamic, large-scale experimental facility for use by both industry and academia.

Large-scale experimental facilities of course cannot be developed anew and repeatedly in every single research project. Therefore there is a strong need to ensure that synergies are achieved, through a federation of testbeds which is sustainable and can dynamically evolve. In a federation, testbeds would normally be geographically dispersed and owned by different organisations. They would however be considered as being part of a single resource, in so far as they are operated within a common management framework under a common management authority.

The federation of testbeds will address issues from early proof-of-concept to validation aspects, thereby allowing industry and academia to collaborate, to exploit synergies, to identify migration paths for new concepts and, in particular, to support the exploitation of research results.

The value of the experimental facilities will significantly increase over time if the data produced by the facility is archived, thereby allowing for better experimenting and comparing at scale or with diversity. To base such an archive on a solid foundation, and to make results comparable, there is a need for additional research on methodologies for testing and benchmarking.

The FIRE facility will be deployed incrementally, as an open recipient responsive to the evolution of the outside stimulations. It will welcome testbeds from different organisations, for instance from SAC, other relevant EU RTD projects (IST/ICT), testbeds from European regional, national, or multinational research initiatives such as EUREKA CELTIC, or the Full Scale European Demonstrator currently under development by the NESSI [6] European Technology Platform. In order to federate geographically dispersed testbeds and to assure the replicability of results, FIRE could also profit from - and build as much as possible upon - GÉANT, the Pan-European Gigabit Research Network. To enable easy replication, interfacing and inter-operation among various testbed components implementation, use of open source software and open hardware design details are considered as key elements.

For FIRE to have maximal economic impact, resulting in new successful European products, services, and companies, it is important that in the long term it can cover the whole chain from basic research to broad pilot projects with real customers. Several of the LivingLabs currently under establishment in Europe [7] have the potential to provide the link to user communities and bring the user into the loop at all stages of research and development, thereby improving the innovation process. The federation of testbeds is challenging in many ways. Managing federated testbeds over

multiple networks crossing multiple administrative domains is difficult and requires professional support. However, federation is necessary to achieve scale, diversity, cost-efficiency and, last but not least, to improve the sustainability and quality of the testbeds.

Sustaining the operation of a testbed comprising advanced technologies beyond the duration of a research project has proven very difficult to achieve in the past. This not only wastes a wealth of experience; but also shows that the business models behind the design of the testbed were fundamentally flawed. Ensuring that the testbeds can be maintained and exploited beyond the lifetime of projects is a key issue for the work in FIRE, as well as finding mechanisms to ensure that the results from past and current projects can be effectively exchanged and compared. Sustainability of the measurement data and its metadata is also an important issue. The huge amounts of information, their heterogeneity and the expected privacy concerns are sensitive issues and solutions for their management (within or outside the FIRE activity) should be found.

Interconnecting different test-beds belonging to different administrative domains means granting some level of access to remote resources that are owned by different stakeholders. This has the consequence that a legal agreement must be in place that governs the rights and obligations of all stakeholders involved, including the handling of intellectual property rights (IPR). In order to establish a long-term sustainable testbed federation, a legal entity might prove necessary to ensure governance.

3. SOCIO-ECONOMIC ISSUES

Nowadays the Internet is not just a technology, but a complex system, and radical technological changes in its architecture could have unexpected consequences at the economic and social level, and even raise ethical concerns. The success of the current Internet is evident by how it has influenced our society. Yet at the same time, the society is continuously raising new challenges for the Internet. A specific and unique strength of FIRE is that it will provide the basis for a scientifically rigorous impact assessment of network architecture proposals, at both technological and social levels.

One core question for the (current and future) Internet is Network Neutrality (NN), a network design principle that states that a maximally useful public information network aspires to treat all content, sites, and platforms equally. NN is perceived by many as being endangered by policies such as differential quality of service, yet by others as a deterrent to innovation and investment in network infrastructure. Opinions are split about how changes to the fundamentals of the network architecture will affect it, and NN is at present debated nearly exclusively in the political arena and in the social sciences, with "hard data" lacking. It is clear that any intelligence in the network - as opposed to today's "dumb pipes" - involves decisions (e.g. what is legitimate traffic and what is an attack?, and what service do packets belong to?) that may affect NN. These decisions should not be the accidental by-products of purely technology- and commercially driven- network innovation. Rather, societal stakeholders should be empowered to voice their requirements for the future Internet in a way that can influence its design, and architectural proposals should be subjected to tests of their likely impacts before being deployed in real life.

The FIRE sustainable, dynamic, large-scale experimental facility on future Internet technologies is an excellent place for doing such rigorous and quantitative research. Specifically, the aim is to treat socio-economic requirements and effects as much as possible in

parallel with the technical requirements and effects: to specify desired outcomes, to define measures and metrics of the achievement goals, to perform analytical and experimental performance measurements, and to derive recommendations.

4. POSITIONING OF FIRE IN THE INTERNATIONAL CONTEXT

Clearly Europe is not doing research in isolation. The FIRE initiative can be seen as strongly related to two US NSF programmes - FIND that is financing research that is similarly addressing new approaches to the Internet, and GENI that is planned to become the US large-scale experimental facility for the Internet; as well as initiatives in Canada, Japan, China and South Korea. A close dialogue on these issues has been started with our North American colleagues and will be continued and expanded to the Far East.

When comparing European research efforts on the future Internet with GENI/FIND, it is worth pointing out that the EU approach is not to build a separate network infrastructure right from the outset. Instead, it is starting from the assumption that such a large-scale platform for experimentation cannot be built or planned before the new architectural concepts have been defined and validated in some way. Whilst European players and efforts are more fragmented, this can be an advantage when looking for synergies, and exploiting them through collaboration. Therefore the EU approach is based on the concept of a progressive federation of testbeds (part of existing or future projects) supporting new, emerging networking concepts.

Another key aspect of European research in this area is a strong social and political perspective on the resulting architectures. If security and accountability have clearly to be enhanced, this has to be done while preserving privacy, and adequately safeguarding the good fundamental principles of the current Internet architecture (e.g. ubiquitous access, the end-to-end principle, the use of open standards and open source, the lack of generalized censorship and its role in ensuring fair use in the circulation of knowledge). Redesigning the Internet architecture cannot be done basing solely on technological and/or commercial constraints, as the Internet has become key for society as well as for the economy. Phenomena such as Wikipedia, YouTube, social networking sites and generally whatever goes under the name "Web 2.0" are having a social and cultural transformational effect on current and future generations whose repercussions have yet to be fully understood.

5. CONCLUSIONS

The FIRE activities [8] are being launched in the 2nd ICT Call [9], closing in October 2007, and projects are envisaged to start in early 2008, under the FP7-ICT Objective 1.6 "New Paradigms and Experimental Facilities" (budget €40m). This first Call relevant to FIRE is focused on advanced networking approaches to architectures and protocols for the future Internet, as well as on interconnected testbeds.

In future Workprogrammes, we hope that the European Commission will continue and extend the approach of experimentally-driven research on future Internet paradigms and architectures and to evolve from gradually federated testbeds towards a sustainable, dynamic, and integrated large-scale experimentation facility. It is envisaged that this will become a major instrument for European and worldwide researchers and industry to reinforce and strengthen their position related to future Internet technologies and services in a globalised world. FIRE is open to any relevant European project within other Objectives of FP7, as well as national, regional or multinational initiatives, to allow usage of the facilities or to federate their testbed within the facility.

In summary, the FIRE activities will be strategic for Europe, given its interest both for industry and academia. The ultimate benefit expected from FIRE is that European industry is better positioned on future Internet technologies and services than it was for the first generation.

6. ACKNOWLEDGMENTS

The authors acknowledge the inputs provided through discussions with many experts from the scientific and industrial community throughout Europe. In particular, Bettina Berendt, Heinz Brüggemann, Costas Courcoubetis, Jon Crowcroft, Christophe Diot, Serge Druais, Anja Feldmann, Silvia Giordano, Dina Kronhaus, Daniele Miorandi, Andrea Nicolai, Yoram Ofek, Radu Popescu-Zeletin, Michel Riguidel, Ioannis Stavrakakis, Joe Svntek, Christian Tschudin, Tanja Zseby, and many others whose names are published on the FIRE website.

Valuable support, ideas and background information have also been provided by members of the new Unit INFSO/F4 of the European Commission. In particular, Per Blixt (Head of Unit), Max Lemke (Deputy Head of Unit) and Fabrizio Sestini (Scientific Project Officer). The full report is available on the FIRE website at <http://cordis.europa.eu/fp7/ict/fire/> [8].

7. REFERENCES

- [1] "The Internet is Broken" in Technology Review, December 2005/January 2006 by David D, Clark, MIT
- [2] <http://cordis.europa.eu/ist/fet/comms.htm>
- [3] cordis.europa.eu/ist/rn/
- [4] www.fp6-ist-onelab.eu
- [5] www.panlab.net
- [6] www.nessi-europe.eu
- [7] <http://www.livinglabs-europe.com>
- [8] cordis.europa.eu/fp7/ict/fire
- [9] http://cordis.europa.eu/fp7/ict/participating/calls_en.htm