Amplifying Future Internet Research and Experimentation for a Sustainable Future

AMPLIFIRE

Support Action, European Commission, FP7-ICT

Grant Agreement: 318550

FUTURE INTERNET RESEARCH AND EXPERIMENTATION: VISION AND SCENARIOS 2020

1ST EDITION

Doc-ID:	D1.1 (M6)
Version:	0.9
Author(s):	Scott Kirkpatrick (HUJI, (lead), Michael Boniface (IT Innovation), Stefan Bouckaert (iMinds), Paul Grace (IT Innovation), José Jimenez (TID), Timo Lahnalampi (Martel), Martin Potts (Martel), Géraldine Quetin (InterInnov), Hans Schaffers (Aalto), Annika Sällström (LTU-CDT)
Date:	07.10.2013
Status:	1st Edition (M6), Final
Availability:	PUBLIC

Document History

Version	Date	Author(s)	Comments
0.1	05.02.2013	Hans Schaffers	Document structure, methodology
0.2	19.02.2013	Scott Kirkpatrick	Introduction
0.3	26.03.2013	Hans Schaffers	Revision of structure, including materials from earlier notes and elaboration of the vision and scenario. Inputs from Géraldine, Michael, José, Annika etc
0.4	22.04.2013	Scott Kirkpatrick, Paul Grace, Geraldine Quetin, Hans Schaffers	Further revision of structure. Incorporate materials reviewed in meetings of 8-9 April.
0.5	05.05.2013	Scott Kirkpatrick, Paul Grace, Geraldine Quetin, Michael Boniface, Hans Schaffers	Completed and integrated chapter 4. Other updates.
0.5a	05.06.2013	Scott Kirkpatrick	Executive summary, updates and editing, next steps
0.5b	07.06.2013	Hans Schaffers	Combining and editing of chapters on sustainability, conclusions in final chapter
0.6	10.6.2013	Hans Schaffers, Scott Kirkpatrick	Update exec summary, reduce no of chapters, trim excess material
0.7	18.6.2013	Scott Kirkpatrick	Version for internal review
0.8	01.7.2013	Scott Kirkpatrick	Version for external review
0.9	07.10.2013	Scott Kirkpatrick, Hans Schaffers	Small edits, final Deliverable D1.1

Executive Summary

The goal of this "FIRE Radar" exercise and document is to sketch the landscape or ecosystem for future internet exploration and experimentation in the years 2015 to 2020, the years in which Horizon 2020 will focus most of its energies. In this landscape we will describe some visions of technologies, services and applications. We explore the ways in which a future FIRE portfolio of experimental facilities can evolve to support the research that can make the visions possible and contribute to their realization in everyday use. AmpliFIRE is an FP7 support action, continuing much of the work previously done by FIRE Station, but the FIRE Radar differs from past efforts in that it provides a range of visions, discusses the gaps that must be bridged to reach them from the present portfolio, and describes the changes in mission and philosophy that will affect the FIRE program in the coming decade.

We provide a historical perspective on how FIRE has been organized and motivated and how those motivations are changing. At present Open Calls and STREPs which are increasingly aligned with the main FIRE experimental facilities are influencing FIRE's evolution from the demand side, by showing customer "pull," supplementing and even replacing technology "push." We conclude that FIRE will increasingly be shaped by demand pull in the radar period, 2015-2020. Its requirements will be based on four pillars of user expectations, consisting of:

- Internet of Things
- Internet of Services
- Internet of Information (media, data, content)
- Internet of People,

all resting on a foundation of adaptable infrastructure, wired and wireless, software-defined, perhaps more open than ever before.

We give examples of where each of these might lead, drawing upon expert advice found in current conferences and white papers. Thus we discuss distributed community clouds, autonomous generation and maintenance of a personal footprint in a real-time wireless "social" network, and explore interoperability issues of the world of smart "things" at several levels, including shared (or "sliced") use of these resources. While federation of access to experimental facilities has been a theme of FIRE to date, the potential for virtualization of processes and services to operate in an extended smart environment seems attractive as a goal for the next generation of FIRE. The business model issues of scaling and extending the potential success of a single smart city to other cities and to whole regions are particularly complex. Finally, we sketch some of the parameters that characterize the space in which our visions of future FIRE prototypes may be constructed. We develop a range of possible scenarios for the future of FIRE, differing along two major axes. Future scenarios might range from highly coherent and coordinated testbeds which link areas of experimentation to fragmented, single-purpose testbeds. And the user community that supports these testbeds can range from individual researchers or developers to communities sharing an interest or an application space. A clear next step for FIRE (and work item for AmpliFIRE) is to identify the most desirable portion or portions of this future space, and develop requirements and instruments that will attract future projects to establish FIRE in that space.

Since the founding of FIRE, many new entrants to the field of experimental networking research and prototyping have appeared, and some of the other existing activities have become more widely exposed. National laboratories in several countries, in North America and in Asia now compete and collaborate with FIRE. On the European scale, the Public-Private Partnership (PPP) and the ICT labs and educational initiatives overlap significantly with FIRE's interests, while addressing different timescales, different customer sets, and different business models.

We address the sustainability of FIRE testbeds. Testbeds must be available for longer than the period of a single research funding cycle, and must evolve as technologies, middleware, and the expectations of both experimenters and end-users evolve. We follow established business methodologies in our analysis, using "CANVAS" charts to identify the resources and customers that must combine to create a successful organization. We observe both academic-oriented and industrially oriented paths that sustainable FIRE facilities might follow. Federation of multiple facilities to permit exploring integrated experiments that exploit multiple new capabilities in different parts of the future internet ecosystem has been a hallmark of FIRE activities to date. We apply our analysis to the question of the appropriate scale on which to federate future activities and the depth to which access through generic experimental control planes might reach.

Table of Contents

1. Introduction

1.1 Overview and initial perspective

The over-all mission of the AmpliFIRE Support Action is to support the FIRE Community to prepare FIRE for 2020, by strengthening the exploitation and impact creation capacities of Future Internet Research and Experimentation (FIRE) facilities. AmpliFIRE develops a vision for year 2020, setting out a transition path from the current situation towards a “FIRE Ecosystem” for 2020.

This document, Deliverable D1.1, contains the first cut at developing a “FIRE Radar,” a sketch of the road along which FIRE can evolve in the coming decade, and the activities in the surrounding environment that make this evolution possible. We first analyze what is changing in the core FIRE projects, ranging from those chartered in Call 5 of FP7 to those starting in the current year. We discuss the influence of peer projects and the changing nature of experimenter demand, as evidenced by the new directions into which the Open Calls are drawing us. We then build and analyze several aggressive but relevant and plausible visions of FIRE’s impact from 2015 to 2020.

The vision of AMPLIFIRE is that in 2020, Internet infrastructures, services and applications form the backbone of connected regional and urban innovation ecosystems**.** People, SMEs and organizations collaborate seamlessly across borders to experiment on novel technologies, services and business models to boost entrepreneurship and new ways of value creation. The Internet comprises a collaborative, community- and market-driven environment where new technologies are rapidly tested in the market, where entrepreneurial activities are fostered and where innovative SMEs are initiating and participating in research, experimentation and product and process innovation in secure and trusted environments.

This “radar” is a forward-looking document in which the FIRE community can tell the world where we may go, what we might accomplish, and how it might happen. It requires broad input from this community and those close to us in order to be successful. This initial version of the radar will be exposed to various audiences before we release it as a first cut, to be updated with a final version at the conclusion of AmpliFIRE.

1.2 Horizon 2020

Horizon 2020 identifies several targets that are of high relevance to AmpliFIRE. Regarding ICT-based e-infrastructures, “the aim is to achieve by 2020 a single and open European space for online research where researchers enjoy leading-edge, ubiquitous and reliable services for networking and computing, and seamless and open access to e-Science environments and global data resources”. Horizon 2020 continues by stressing the need to foster the innovation potential of research infrastructures and their human capital. Innovation should be stimulated, both in the infrastructures themselves and in their supplier and user industry, by developing R&D partnerships with industry, by stimulating the use of research infrastructure by industry e.g. as experimental test facilities or knowledge-based centres and by encouraging the integration of research infrastructures into local, regional and global innovation ecosystems. Also, Horizon 2020 promotes the use of research infrastructures to be leveraged for public services and social innovation. Throughout the Horizon 2020 there is much emphasis on the demand side, the need to engage users to create more innovation friendly markets, and it is clearly stated that the ICT-specific research infrastructures include living labs for large-scale experimentation and infrastructures for underlying key technologies and their integration in advanced products and innovative smart systems.

1.3 Coherence vs. Fragmentation

The gaps between the technologies presently offered in FIRE as testbeds, and the gaps between the layers in which its communities have formed are large. For example, the gaps between wired and wireless networking, between networking researchers and cloud application developers, and between both sorts of developers and end user input all require bridges thatexist today only as research efforts in, e.g. Fed4FIRE. Scenarios and user requirements must shape and drive those bridging activities. This “FIRE Radar” attempts to chart the most direct paths from the present fragmented FIRE portfolio of testbeds, both hardware and user-oriented, to the goals of Horizon 2020. This will require a sustained effort to articulate how the technical goals of the present FIRE activities can be lifted, channelled and amplified to support the societal goals of Horizon 2020. We believe that this places requirements on the FIRE participants which, as engineering teams with an often academic focus, they will need to meet in new ways. Our job is to clarify and justify these requirements and identify new instruments and relationships with business and SMEs that can draw upon the strengths that FIRE has created. We must go beyond the "Portfolio analysis" of the existing projects that FIRE STATION has conducted to date. We must articulate what is not yet being done, what steps are missing, and what communities need to be engaged or created to make these things happen. This helps to create the “pull” that can make FIRE effective as 2020 approaches, and assist the individual projects as they provide the “push.”

1.4 Structure of this document

Section 2 provides the basis for our analysis. It spans from the history of FIRE to the changes that are driving its evolution today. First it introduces the FIRE projects and describes the environment in which this portfolio of testbeds and experimentally-driven research was created and how FIRE and that environment are changing. We consider not only FIRE but also peer activities which make different tradeoffs. Some are happening outside Europe. Next we describe the uses of the testbeds and how the demands upon them are evolving. We pay particular attention to the new directions in which the Open Calls and the latest STREPs are taking FIRE facilities. Finally we survey current research (and the fast moving consumer frontiers) in the areas that FIRE must influence and support.

In Section 3, we develop future scenarios for experimental research and development in the future internet. While our scenarios are rooted in the FP7 FIRE experience, we also shape them to reflect the ambitions of four “pillars” of future networking and its applications: The Internet of Services (IoS); the Internet of Things (IoT), the Internet of Information (IoI), and the Internet of People (IoP). These, we claim, have replaced the more technological objectives that defined ICT in FP7, and will now drive experimental research in the coming decade. They will, however, require an increasingly flexible, integrated and adaptive infrastructure. This foundation will remain a critical part of the FIRE portfolio as we approach 2020.

Section 4 introduces the questions of sustainability. These issues and questions organize our discussion of the future directions of the FIRE activities. Some of these, e.g. business models and legal issues, have been exposed in previous FIRE road mapping activities. We first analyze them here in the context of the individual present FIRE IPs, then extend the analysis to FIRE as a whole, perhaps a federated single entity.

In section 5 we summarize our thoughts on the future scope of federation and prepare to use a "gap analysis" methodology to provide a systematic framework for more extended exploration of the steps and paths needed to move beyond today's activities to the envisioned world of a "FIRE ecosystem" in 2020. This analysis, when done in detail, requires agreement anmong participants and the funding Unit on the envisioned objectives of FIRE. That work will be performed during subsequent months of the AmpliFIRE project and reported when we present a final version of this vision, at the end of AmpliFIRE. We conclude with a summary of "next steps" for AmpliFIRE. An extensive set of references to recent FIRE and related project documents, websites, and relevant articles in the open literature is provided as an appendix.

2. Changes in FIRE: Structure, Environment, and Requirements

This chapter summarizes the current state of the FIRE projects, and the forces that have changed them during FP7 and will continue to have important effects. It is important to understand these if we are to create a valid vision and future scenarios, in particular to address strategies for the future and future sustainability.

2.1 FIRE and related initiatives

The Net Futures initiatives, now including more than 100 projects representing an investment of far more than 800 M€ of which the European Commission funds 570 M€ (The Future Internet, 2011), is addressing this need for early experimentation and testing of Future Internet technologies in multidisciplinary research environments. The FIRE community advances the development and harmonization of experimentally-driven research methods and platforms to ensure the continuous relevance, rigor and robustness of the research and the strategic research agendas. FIRE activities have resulted in many important achievements in terms of federated testbeds, access to testbed facilities, experimental research methods and tools, and collaboration across disciplines and communities and across geographical areas. The FIRE unit extended 40 M€ in funding during Call 2 of FP7, 60 M€ in call 5 (for projects starting in 2010), 20M€for Call 7 (starting 2011), 25 M€ for Call 8 (starting in 2012) and 19 M€ for Call 10, which begins this year.

In addition, several other coordinated initiatives are now operating. The **Future Internet Public-Private Partnership (PPP)** (2011-2014) is a program amounting to 300 M€, focusing on short to mid-term market oriented research. The foundational project of the FI-PPP, entitled FI-WARE, will provide Future Internet developers with a platform and a test bed for rapid application development; another project, called "Infinity", is looking at existing experimental infrastructures available all around Europe; some of the FI-PPP "Use Cases" are strongly oriented towards smart cities, and some are looking beyond infrastructures towards user communities; finally, the coordination action CONCORD involves living labs. Interaction with such projects, especially in Phase 2 of the FI-PPP which will start early 2013, will allow AmpliFIRE to streamline short, medium and long term efforts towards a unified FIRE facility offer. The **CIP ICT-PSP** establishes large-scale trials of existing technology, focusing on service innovation in many areas, including living labs and Internet of Things for smart cities. The European Institute of Innovation and Technology (EIT) has selected the **ICT Labs** to accelerate innovation in Europe in the area of ICT through establishing partnerships between companies, research centres and universities. This initiative uses a co-funding model to catalyze and stimulate R&D – e.g. testbeds - as well as, interestingly, entrepreneurial activities. At the **national level** we also see important initiatives in the domain of future Internet, such as in France (GRIF: Research Group for the Future Internet; FIT: Future Internet of Things), Finland (ICT SHOK), Germany (G-Lab, ICT2020), UK (UK Future Internet).

All these initiatives demonstrate opportunities for transnational collaboration, sharing of knowledge, technologies and facilities, and for synergy creation. The FIRE investment over the past 5 years has created a pool of testbeds/facilities that support networking research leading to future Internet protocols and some of the applications that may drive the continued growth and social impact of the future Internet. The FIRE portfolio has successfully drawn upon various national efforts, some of which have been in place longer and permit more elaborate experimentation or experiments (such as new services) of longer duration. In addition, the most recent additions to the FIRE portfolio have for the first time made application platforms (clouds) and sensor-enriched environments available for broader use. The FIRE STATION coordination action, which ended at the end of May-2013, initiated a "FIRE Architecture Board" as a vehicle to draw the disparate testbeds into a portfolio of federated resources and to bring together the FIRE community (facilities, research projects, and other stakeholders at European and international level). Since October 2012, development of federation technologies is is being carried out in the scope of the FED4FIRE Integrated Project (FIRE FP7 Call 8). Still, at this stage, the users of these FIRE resources are almost exclusively engineers and scientists from academic and research institutes..

Following the fast paced developments in specific research areas in ICT and Internet, the emphasis of the FIRE program must increasingly shift towards integration. The FIRE infrastructure of separate testbeds offers opportunities to develop applications which exploit synergies between the new technologies. Also the testbeds which are in fact application platforms such as clouds and smart cities make FIRE's innovations visible to end users and permit exploring changes in the user experience which may drive growth in computing and communications in the second half of the coming decade. To achieve this requires amplifying the effectiveness, performance and impact of the facilities and working harder to disseminate the know-how which has been aggregated. In the years to come, exploitation, impact creation and sustainability are keys to success.

2.2 What's happening in FIRE?

There have been dramatic changes in FIRE throughout FP7 as a consequence of strategic actions taken by the FIRE Community and the European Commission. FIRE was established from a core of networking testbeds aimed at investigating some of the fundamental issues of the network infrastructure. Driven largely by universities the debate in the early years focused issues such as clean slate vs. evolutionary design, tussles between networking stakeholders, the role of experimental methods in computer science and the relationship with international initiatives in the US, Korea and Japan.

The launch of the European Future Internet initiative and the Bled declaration provided a turning point in FIRE's aspirations. The Future Internet offered a convergence narrative between Networks, Services, Internet of Things, Content and Security and importantly provided a new and important context for FIRE facilities. Suddenly FIRE had to consider not only interesting network research challenges but also how FIRE can deliver general purpose reusable facilities for the Future Internet community avoiding the duplication of effort of testbeds developed within individual research projects. In essence FIRE was now meaningful in the context of the Future Internet and needed to serve the research communities that they represented.

The work programme for ICT-Call 5 was the turning point, when this strategy was implemented. Two facilities were funded targeting Objective 1.2, Software and Services (BonFIRE, TEFIS), one targeting Objective 1.3 Internet of Things (SmartSantander), another building on an important emerging US networking technology OpenFlow (OFELIA), and still another targeting cognitive radio networking (CREW). In ICT-Call 7 the reach was extended to Objective 1.5 social and networked media (EXPERIMEDIA), user centric networking (CONFINE) along with a consolidation of previous networking facilities PanLab and OneLab (OpenLab). The divergence of testbeds after Call 5 raised concerns about duplicate efforts for developing tools to support the experimental lifecycle. This resulted in the addition of a federation IP into the work programme for Call 8 (Fed4FIRE) with the objective of bringing together different efforts through a common high level federation framework. Federation and Virtualization are two long-standing objectives of networking research, both having somewhat loose definitions with some variation from one research project to the next. Federation, as seen in the OneLab series of projects starts with common identity management (logins and access) and proceeds in steps to include common experimental control planes for the dispatch of experiments and collection of the resulting data. Virtualization, allowing the network to be managed in different ways for different types of data in transit or by different applications, is one of the objectives of the OpenFlow software employed in OFELIA.

Without a clear FP7 strategy FIRE has never integrated its different facilities, which provide resources, tools and services to support the needs of different research communities. Fragmentation is not necessarily a bad outcome as the breadth of research challenges faced by experimenters is unlikely to be supported by a single facility. For example, there's a huge difference between what is needed in terms of resources and tooling between cognitive radio and media application experiments. In fact value is created for users when domain specific entry points to facilities are developed by offering technical APIs familiar to those communities. Much of this debate is being had within the FED4FIRE project where the limits of cohesion between facilities are being investigated.

Other work programme constraints during FP7 that have significantly influenced FIRE include:

- **Open Call Constraints IPs**: Call 5 required budgets with 20% reserved for open calls, 20% for coordinated on-demand activity. Call 7 and Call 8 had no constraints and Call 10 reserved 50% for open calls. Following Call 5 all IPs have had some provision for open calls even if the Call text did not ask for it. Open calls have been good for FIRE and projects themselves (e.g. enforces the customer/supplier relationship increasing the quality of results from the core project, keeps facilities relevant/dynamic as new requirements come from experiments, provides a great opportunity for promoting FIRE and the projects, etc). There is a cost and a learning curve for running open calls. FIRE is getting good at this process and the benefits far outweigh the additional management overhead.
- **Facility** Constraints for STREPS: Call 5 had no constraints; Call 8 STREPS had to get letters of support from a FIRE facility. Call 10 again had no constraints. FIRE aspires to fund research projects that use facility projects. This is problematic for various reasons (e.g. contractual issues, inter-project agreements and accountability, etc, etc) but in principle having a STREP identify a facility encourages integration of endeavour and cohesiveness of the programme. At least at proposal time the STREP is thinking about building on FIRE results. There are ways create this relationship, for example, a Memorandum of Understanding (MoU) or include facility partner in STREP (the latter being a recipe for some partners getting big slices of funds). The latter is interesting from an IPR point of view and one option is to treat facilities as BonFIRE treats experimenters (e.g. different contractual terms and access rights to IPR). BonFIRE could then include a facility partner such as iMinds but because they are just offering a testbed to use the experimenters have limited access rights of BonFIRE's or iMinds' IP.
- **Collaboration Constraints for IPs**: Call 5 collaboration was written into the DOWs at 8% of budget, Call 7 was negotiated (EXPERIMEDIA was 2%) and Call 8 had nothing. Collaboration budgets have not been very effective. In Call 5 it was envisaged that FIRESTATION would have authority over how the money was spent towards the vision of federation. In practice, the projects themselves decided what was mutually beneficial

and justified this to their PO in their project reviews. Of course information was provided to FIRESTATION but it was more a case of projects deciding how to account for the 8% rather than FIRESTATION acting with authority.

Figure 1 shows the primary relationships between FIRE projects running in 2013. Nine facility projects exist (green) each serving specific research communities. Some testbeds within facility projects are being federated through the FED4FIRE project. A set of research STREPS are running some of which are directly associated with facility projects as a requirement of Call 8 (yellow) whilst others from Call 7 are independent (white). The diagram shows relationships between projects and not relationships between testbeds within projects. Some testbeds exist within multiple projects, for example the University of Cantabria's testbed is within SmartSantander and FED4FIRE, iMinds' iLab.t is within OFELIA, CREW, OpenLab, FED4FIRE and BonFIRE. The meaning of the relationship needs to be explored for each case. For example, the relationship could mean that the projects share partners or that there's some looser association based on research objectives.

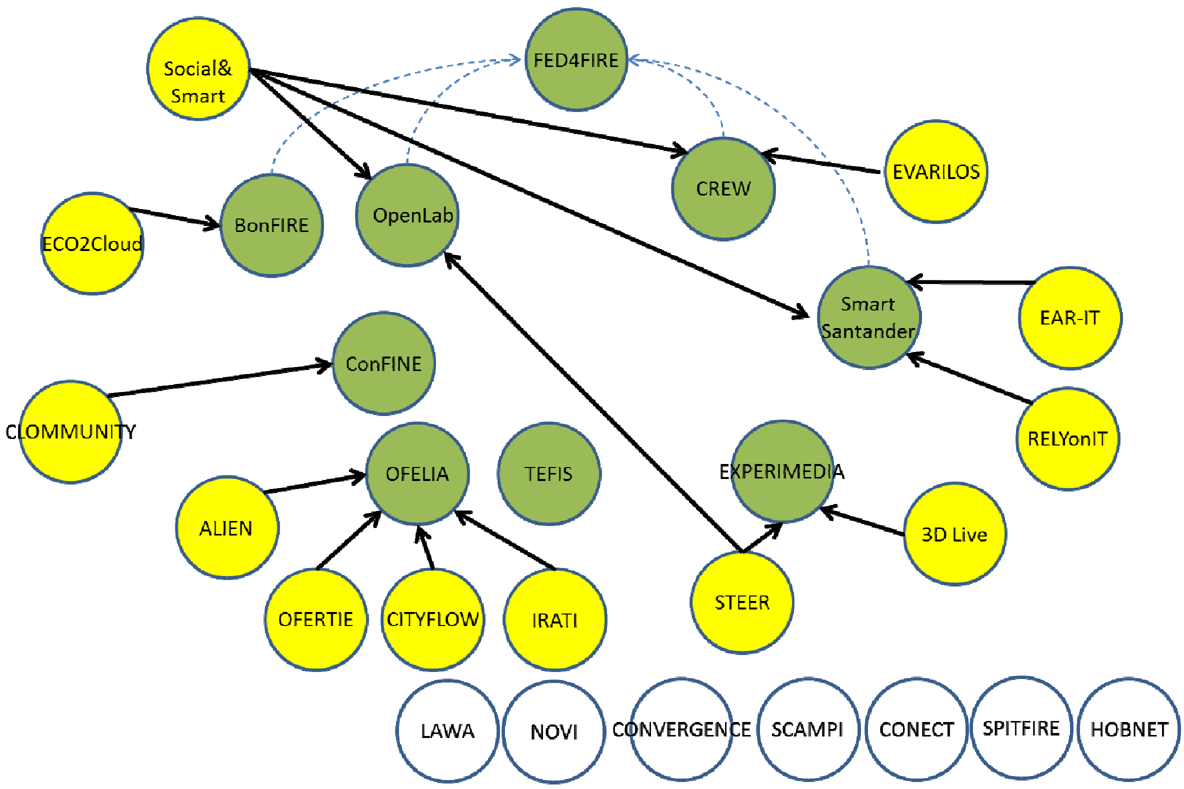

Figure 1 FIRE Projects in 2013

The key point is that the relationships between projects change over time as facility building projects come and go. If we look at what could happen at the start of 2014 (Fig. 3) we have Facility projects finishing in 2013 (orange), two projects starting (blue) and set of Internet Science STREPS funded who are expected to have some association with a facility although the requirements were weaker for Call 10 than for Call 8 STREPs. Therefore, it is likely there will be less coherence and more fragmentation between Internet Science STREPS and facility projects when compared to Call 8.

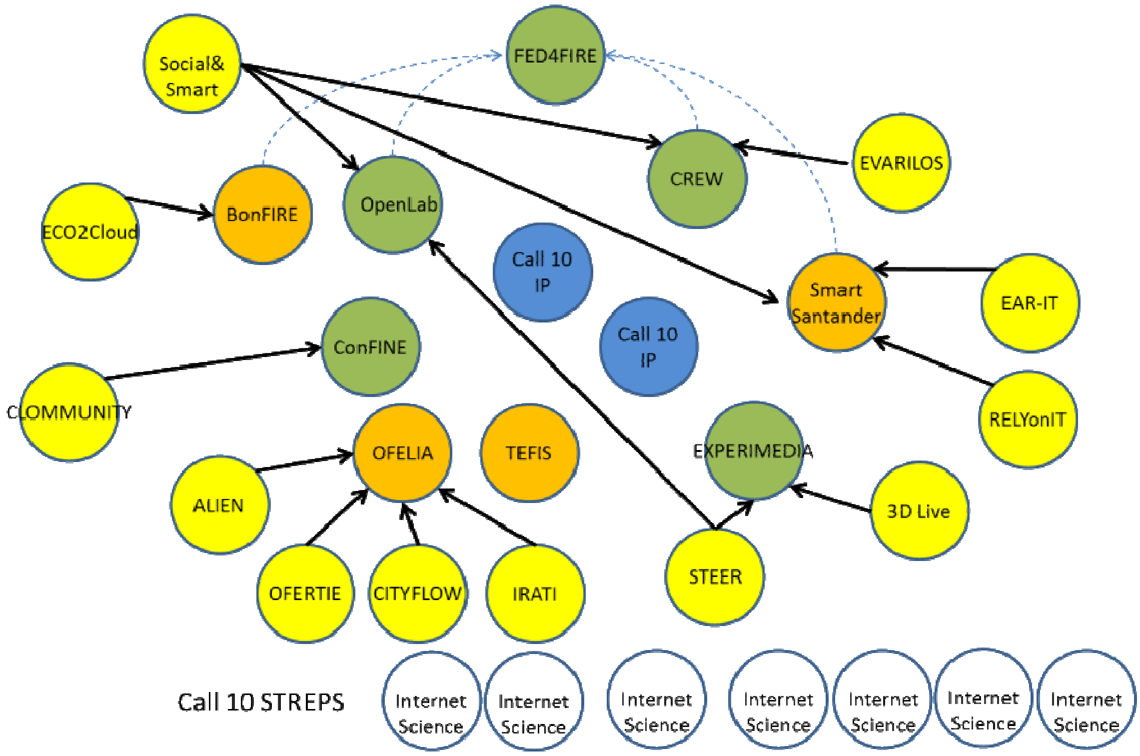

Figure 2: Expected situation in 2014

A key driver for the scope of FIRE is the restructuring of the EC to create DG Connect where FIRE is now part of the Directorate E – Net Futures. The FIRE unit was also renamed "Experimental Platforms" during this process, although we expect the FIRE brand to continue to be used. What this means is that FIRE's target research communities (the primary users of FIRE) are now within the same Directorate where previously they were in different ones. The expectation is that FIRE can increase strategic alignment with potential users both in terms of research goals and potentially revenue streams. We expect the level of funding in FIRE will be similar to that allocated for FP7 (approx. 150 million EUR). This means that the size of FIRE both in terms of number of projects and scale of facility resources is unlikely to change dramatically.

We also note that the European Commission no longer talks about the Future Internet apart from within the context of the Future Internet- PPP. This leads to the question of what does the EC mean by Net Futures in terms of scope of research goals and how can FIRE address the needs of those communities through the facilities it offers? Net Futures consists of "Network Technologies", "Services & Software, Clouds", "Net Innovation" along with FIRE's "Experimental Platforms". We know that Networks and Services will remain key areas but Net Innovation introduces new ideas that may influence the future of facilities. If FIRE's strategy is to provide an offering to Net Futures then understanding what challenges are facing the other units may be critical to FIRE's success.

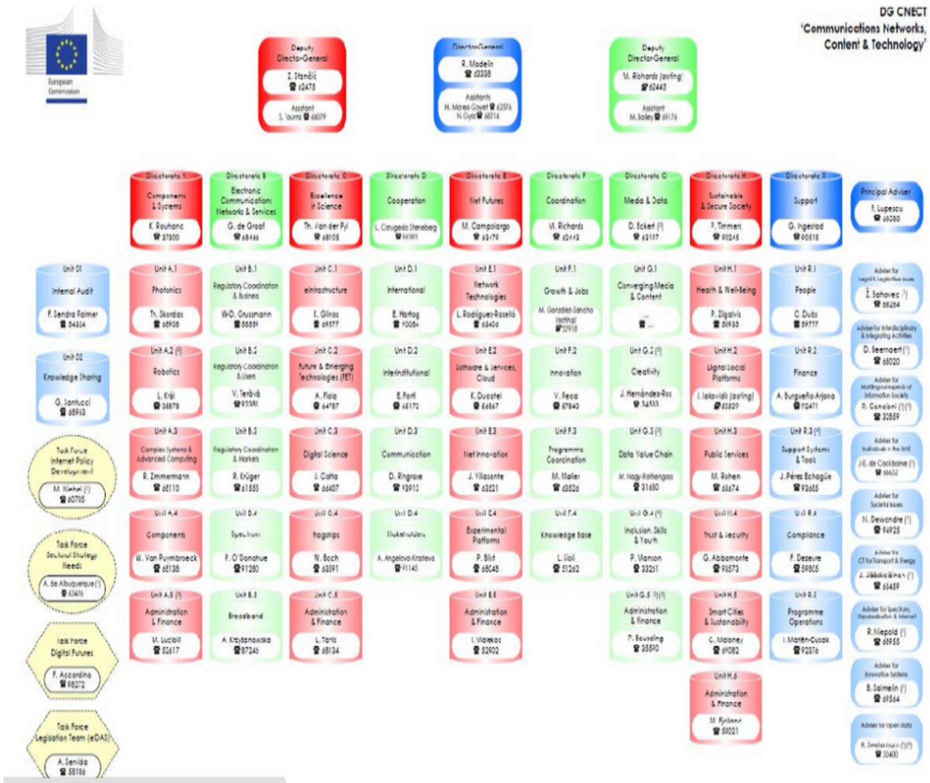

Figure 3: Net Futures in the DG Connect Directorate E

Horizon 2020 (H2020) is expected to launch its 1st call in Spring 2014 with projects starting early 2015. The initial work programme is currently being drafted with an announcement expected in autumn 2014. If we examine the FIRE portfolio at the point when the first H2020 projects will start we know:

- CONFINE, OpenLab and EXPERIMEDIA projects are ending
- BonFIRE open access is ending
- CREW, Fed4FIRE and two call 10 IPs will be running, although each will be over half way through

At the time of writing it is unknown if experimenters will be able to access and run experiments on TEFIS, SmartSantander and OFELIA facilities in 2014. The possible options for these facilities include securing funding associated with a Call 10 IP, establish a business model independent of EC funding or closure of the facility. Call 10 IPs are expected to be continuations of the most successful facility projects so it is likely that some will continue with EC funding.

Sustainability of facilities is an important topic and is dealt with in Section 4 of this report. However, when we consider "what's happening?" it should be noted that FIRE's goal is to provide useful services to users and that the needs of these users change over time. Sustainability does not mean that facilities live forever, just that their costs must be met for the required period of time. Technology comes and goes, and so do the facilities used to research and develop technology. For example, early FIRE testbeds supporting technologies such as IMS may not be needed by researchers today. We are now seeing OpenFlow entering production deployments, so facilities such as OFELIA must evolve or be replaced to support a wider scope of software defined networking studies. Facilities have a lifetime that is related to the adoption of technology and their ability/desire to adapt to new requirements. If a facility provides a context for research over a given period (3-5 years) and then becomes unavailable because it is not relevant to current research should this be seen as a success or a failure?

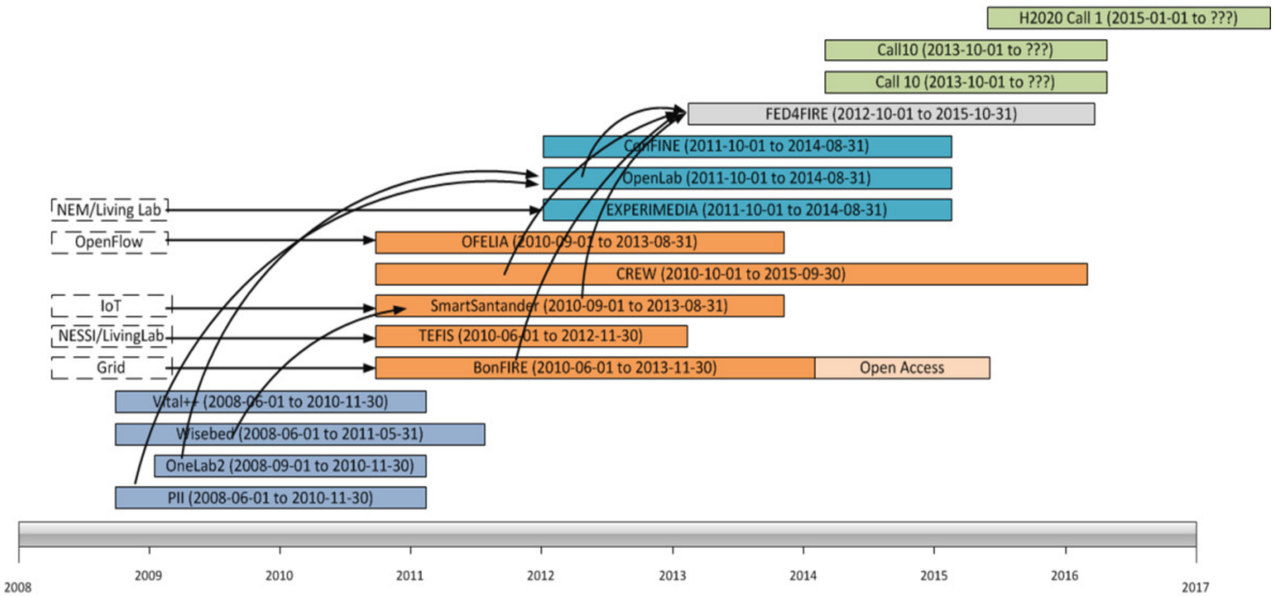

Figure 4: Timescales for FIRE into H2020

The strategy for FIRE following the launch of the Future Internet initiative in 2008 was clear. Yet for H2020 FIRE still needs to establish a Call strategy that builds on the results of FP7 yet meets the demands of researchers and developers from 2015 to 2020. Specific important actions have been taken to secure FIRE's future and to increase its impact (e.g. aligned FIRE and Net Futures Units). The definition of the 1st work programme for H2020 will be key to FIRE's future as the winning projects will establish the foundations. The connections between calls will be important for facilities, as this affects continuity of service. The uncertainty associated with the Net Future's research context needs exploring to understand the scope of FIRE facilities in H2020 whilst inefficiencies in the funding structures must be investigated to ensure that public money is spent efficiently and targeted in the areas required.

2.3 Previous perspectives on the future of FIRE

AmpliFIRE's work continues previous efforts in the FIRE community to understand the common thread in the FIRE portfolio and its evolution. We summarise these discussions, which have contributed to this report.

2.3.1 FIRE STATION's FIRE Roadmap

Within the FIRE STATION Support Action, one of the key activities has been to develop a FIRE Roadmap D3.6 (2012). This roadmap mostly addresses experiment life-cycle management, sustainability (legal, governance, financial aspects), and services (trustworthiness, shared support services). The report presents an overview of trends and developments which is highly useful for AMPLIFIRE as regards the analysis of the current situation and changes. The document does not propose a "roadmap" in terms of targeted activities and solutions and their timeline. AmpliFIRE does plan to develop such a roadmap in later deliverables, factoring in the initial investments made in Horizon 2020.

2.3.2 FI-PPP perspective

Now that the Future Internet PPP initiative is entering Phase 2 in April 2013 it is appropriate to discuss the opportunities of mutual benefit of FI-PPP and FIRE. The FIRE AB document on Horizon 2020 (FIRE STATION 2011b) already mentions that FIRE availability for shorter term trials and experimentation could benefit the FI-PPP especially for Phase 3 trials. Pilot experiments in smart cities and living labs seem to be a good fit. It proposes that FIRE facilities should be used in the FI-PPP where the FIRE facility offer should take into account the requirements from FI-PPP. Also that collaboration with EIT ICT-Labs is to be stimulated. The FIRE AB document considers this as a medium to long-term objective along with ensuring sustainability of FIRE.

The FI-PPP Architecture Board is also looking at interaction between FI-PPP and FIRE. The view expressed in Fig. 5 is to have a common shared pan-European infrastructure made up of a network of federated FI-PPP Datacentre Facilities and complementary sites including FIRE experimentation facilities. The FI-PPP Architecture Board White paper on Phase 2 and 3 (2012) says: "One approach is shown.., in which a single FI-PPP Datacentre Facility and complementary common FI-PPP related sites (e.g., a ID provider site or a site where the FI-WARE Location generic enabler (GE) is hosted) are shared between all the trials. Backend functions can be hosted there using the FI-WARE Cloud Hosting capabilities, as it will be composed of general purpose, commodity storage and processing hardware accessible over the Future Internet. Some of the FI-WARE GEs may also be hosted there, offering their functions "as a Service" (e.g., most of the Data/Context Management GEs, IoT backend GEs, etc). Each trial would also need some other facilities, including all 'in-the-field' elements, which might be provided by trial-specific facilities, or more generic experimental facilities such as those established in the Future Internet Research and Experimentation initiative (FIRE)."

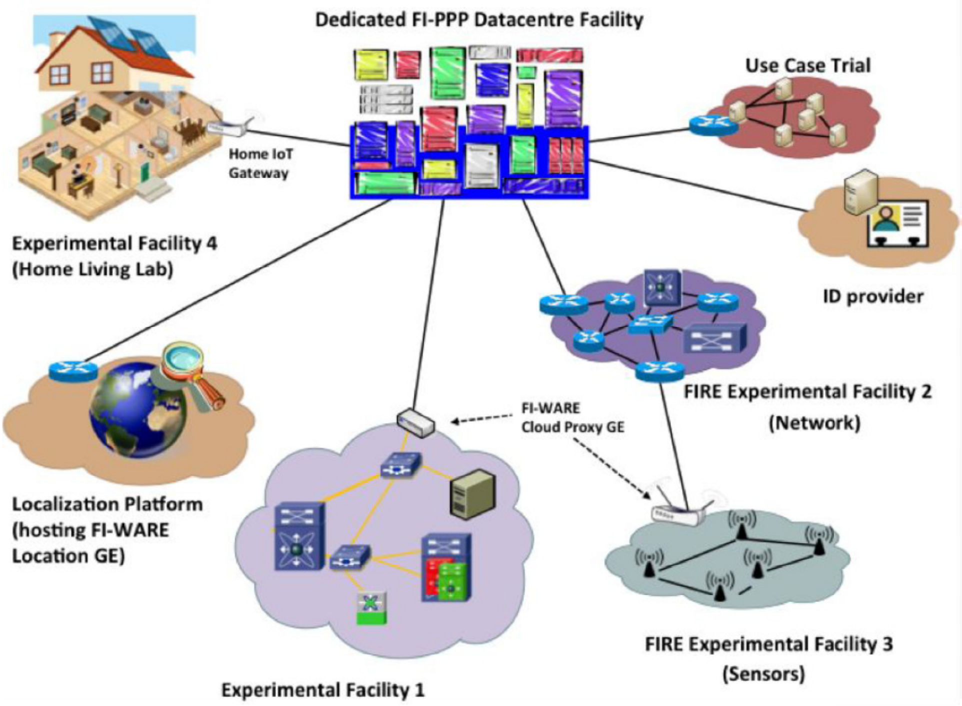

Figure 5 Trials Ecosystem (FI-PPP Architecture Board, 2012)

2.4 Success of the Open Calls

2.4.1 Impact of Open Calls

The Open Calls offered by FIRE integrated projects have had a surprising impact. The response exceeded most facility providers' expectations, and the breadth of organizations responding has changed the demographics of the FIRE community (in a healthy fashion). We study the results of this process here as it has given an early warning of the directions in which the uses of FIRE's facilities are likely to evolve. The FIRE research portfolio has been enriched. It now contains integrated projects, STREPs, the briefer experiments which are absorbed into the integrated projects, and ultimately a wider set of experiments which exploit the FIRE facilities yet are not funded by FIRE (in some cases, formalized as "open access"). The following table summarizes the results that are presently known from the first two rounds of Open Calls in FIRE:

	OpenLab	CREW		OFELIA	TEFIS	BonFIRE
1st Open Call --	**198 proposing organisations – 51 IND & SMEs (26%) 19 proposals selected -- 3 IND / 2 SME selected (18%)**					
Results open call (#proposals received)	14	18	47	21	22	29
Proposing organisations*	19 2 IND-2 SME	24 1IND -1SME	63 9 IND - 19SME	24 1 IND-1 SME	36 2 IND-4 SME	32 9 SMEs
Accepted proposals	4 ⇨ + 6 partners 1 IND	3 0 IND/SME	2 selected ⇨ + 4 partners 0 IND/SME	2 selected ⇨ + 2 partners 0 IND/ 0SME	4 selected ⇨ + 8 partners 2 SME	4 ⇨ + 5 partners 2 IND
2nd Open Call --	**122 proposing organisations – 37 IND & SMEs (30%) 22 proposals selected -- 10 IND/SME selected (33%)**					
Results open call (#proposals received)	15	21	31	17		22
Proposing organisations*	22 2 IND-6 SME	24 3IND-3SME	31 4 IND - 13 SME	23 2 IND		22 1 IND-3 SME
Accepted proposals	4 ⇨ + 5 partners 2 IND	4 +4 Partners 2 SME	4 selected ⇨ + 5 partners 3 IND/ SME	4 selected ⇨ + 6 partners 0IND		6 selected +6 partners 3 IND/SME

*IND = large industries

Table 1: Results of two rounds of FIRE Open Calls

In the first round of Open Calls, announced in 2011, 151 proposals were received, from 174 organizations. The second round has received 106 proposals, from 122 organizations. Open Calls from subsequent integrated projects will soon be offered, and Open Access experiments, which we have not attempted to survey, are starting in BonFIRE, OpenLab, and OFELIA, and possibly others. Existing members of the projects which run the facilities that offer the specific Open Call are not permitted to propose, although they might participate in other FIRE-funded activities. As a result, the process has greatly expanded the number and types of organizations that participate in FIRE.

From the table, we see that nearly one third of the applicants for Open Calls have been industry or SMEs. The proportion of proposals submitted for industry or SMEs increased slightly in the second call, from 26% to 30%. In the first round, seven of these less academic proposals were successful. In the second round, a greater effort was made to tailor some of the offered facilities to the needs of industry and SMEs, with the result that one third (7 out of 21) of the accepted proposals involved industry or an SME. SmartSantander and BonFIRE, which are most visible to end users and present smaller technological hurdles for a potential experimenter to overcome, have received the largest number of proposals. SmartSantander received 78 proposals in these two Calls; BonFIRE received 51. But all projects attracted the attention of more potential new partners than the number of existing partners within the project. The actual number of proposals and new partners accepted, however, has not exceeded four proposals or five new partners in a single call. The finite budgets set aside for Open Calls and the tendency for each proposal to consume at least 50-100K EUR, although less than the suggested maximum of 200K EUR, have kept these numbers down.

2.4.2 Lessons learned from the Open Calls

As the Open Call experience has not been documented and discussed broadly, we take this space to list and discuss some of the lessons learned so far. Open calls for funded experiments are a novelty in the EC's ICT programs, although in other areas of science the process of proposing experiments in order to win access time and support at an experimental or measurement facility is quite common (at particle accelerators or telescopes, for example). The strong requirements that the EC imposes for fairness and openness of the reviewing process and financial oversight of the resulting work while delegating the selection process to an integrated project make FIRE's Open Calls for funded experiments different in some significant ways. The process was initially unfamiliar and considerable extra time was consumed as each project developed procedures and identified appropriate reviewers. The process for each call, from publication to selection, negotiation, and contract signing has required six to nine months. In the course of this some projects have developed consortium agreements that are different for open call partners than for full partners, further complicating the negotiations.

It has proven necessary to get the Open Calls processes started as early as possible in order to learn from a first call and apply that learning in a second call, and to see substantial results before the project ends. No consistent policy has emerged for the extent to which proposals must be kept confidential, when knowledge of some details of the proposed experiments must be shared to help the testbeds get ready to meet the needs of the experimenters. The EC's pre-financing schedules, designed to provide cash flow for partners who are members of the consortium during the entire project, don't always make the funds for an open call available at the needed times. Finally, the rules set by the EC for advertising an open call are an awkward mixture of the procedures for subcontracting and for research calls. The Open Call must be advertised in the national newspapers of three European countries, but this has proved a waste of time and funds, while advertising in technical journals and on the internet is more effective. Projects are still learning the best way to conduct an information day for their open calls.

The FIRE testbeds frequently offer access to a range of hardware and expertise that is simply not available in the commercial market. But applicants to the open calls are also attracted by the opportunity to obtain modest amounts of EC funding for their experiments, in addition to a facility's commitment to offer training and support (as most FIRE IPs have done). Some feel that the relative ease of preparing a proposal from one or two institutions (rather than the 4-6 institutions that participate in most STREPs) is an added incentive. Even when such services are available, as with cloud computing and storage services, the FIRE Open Calls offer free, rather than use-based, service for a limited period. For an industrial or SME proposer, this reduces time to market for an idea that has reached the prototype stage. Finally, participation in FIRE allows extra avenues for publicity, such as participation in public events like the FIA meetings. This should be of value for commercial as well as academic customers.

The testbed owners also benefit from the extra activity. The funded experiments and new partners that they support have increased the number of eyes and hands that contribute to the excellence and capabilities of the testbeds, by a sort of user-driven open innovation. Projects have learned to budget the support effort necessary to make this work, so that user support workload does not come as a surprise or interfere with the research commitments of the project. Users validate the usefulness of each testbed (always a valuable point tested in the projects' reviews), and have contributed to improvements in algorithms, measurement tools, and even equipment. This is especially true when a testbed chooses new partners through Open Calls who bring additional technical resources to the partnership, rather than simply being users of the existing facility. Finally, the larger user community provides a more rapid path to standardization, and aids in assuring interoperability of the hardware or software that is employed on the testbed.

There are some challenges to doing this well that are now recognized and being worked on by several of the FIRE testbeds. The first is the need for better usage metrics and measurements so that the costs of external experimentation can be fairly allocated or anticipated. This is a necessary step in preparing a testbed for future sustainability when usage charges might be a part of its overall financial structure. The experience with initial users may need to be expanded into a package of services which will be attractive for SMEs and industrial users. And the legal implications (especially with regard to intellectual property) of having partners who are users rather than providers of the testbeds have to be understood, so that consortium agreements or subsidiary partnership agreements can be modified appropriately. On the technical side, the FIRE testbeds that see extensive external use will need to develop standard methods (possibly shared within the FIRE portfolio and Net Futures community) for monitoring usage and saving the resulting data to permit best practices comparisons. Managing life cycles of the equipment employed is a second area where the extra work involved in sustaining a testbed beyond a single project to support ongoing experiments can be shared for greater effectiveness.

2.4.3 Validation, extensions and new directions from the Open Calls

Initial guidelines given to the FIRE integrated projects as they prepared to offer Open Calls emphasized verification of the usability of the testbeds and validation of the technical value each testbed provides. These objectives are readily understood by the project managers, since they translate into making each project more successful in terms of its stated objectives. There was discussion within the FIRE architecture board of additional objectives, such as encouraging experiments which enriched the FIRE portfolio, rather than displaying the contributions of a single project in the best light, experiments which are more challenging and risky, and experiments which explicitly link together two or more testbeds in order to understand the needs of an application at a higher level. An initial consensus emerged that FIRE should develop a better understanding of what works and what does not in the open calls process before undertaking additional organizational and legal complexities. Given the apparent success of the open calls in attracting a broader demographic and offering exciting experiments, it is probably time to be more aggressive. And, as we shall see, some of these wider objectives are already being met.

OFELIA, the IP which deployed an OpenFlow routing testbed extending across five "islands" in five European sites, in its first call focused most clearly on extending its reach. At the conclusion of its selection the number of islands had increased to ten. A Brazilian partner gave the testbed an intercontinental reach. OpenLab, in its first round of Open Calls, added partners exploring experimental routing strategies, Internet of Things integration, and wireless content delivery performance. All these could be considered natural applications of specific OpenLab testbeds. Its PlanetLab distributed component is identified as the integrating and deploying core of each of these experiments. Anotherpartner brings in new tools, aimed at locating and characterizing public online storage facilities that can add function to the "middle mile" over which content much pass in global applications. This was probably not an experiment or application that the OpenLab project had initially identified. CREW, a cognitive radio testbed, which is probably the most specialized of the FIRE platforms, selected three experiments. Each deals with particular issues within the scope of cognitive radio. Two experiments explore possible protocols, and the third tests machine learning algorithms adapted to the problem of sensing the presence of transmissions that compete for overall bandwidth but employ different encodings or require different subbands.

In their second Open Calls, each of these technology-intensive projects has branched out, attracting a much wider range of proposals and accepting some which establish some new directions for FIRE. CREW added a project exploring greener solutions which reduce energy consumption, two projects that add benchmarking and field trial capability relevant for the testing of products closer to the marketplace, and one addressing control protocols with an eye to their eventual standardization. At the conclusion of their second call, CREW announced that they would offer public access to the CREW facilities, upon submission and review of proposals, but without additional funding. OpenLab's second call added a partner constructing an experiment dispatch and control plane (a tool), one exploring variants of BGP border routing that offers greater security against hacking, an experiment testing interoperability of telephone company constructed software on the two IMS testbeds, and a software-defined networking approach flexibly linking two wireless facilities across the PlanetLab overlay network. Independently, and perhaps in response to the evident interest in software defined networking, PlanetLab Europe (PLE) has recently offered OpenFlow interfaces across all of its distributed clients. OFELIA in its second call continued to add "islands," but has also bridged into the cloud computing environment as the newest island also is testing support in the network for virtual machines, making the network more application-aware. Other experiments cover a control plane processor for managing very high bandwidth optical pipes, multicast, and video-on-demand caching. All of these explore extensions of the OpenFlow capability that was initially deployed. Coupling the network controls to the applications (e.g. virtual machines) is at the forefront of current SDN research and deployment efforts.

TEFIS, SmartSantander, BonFIRE and Experimedia offer experimental environments in which the path to an eventual product or business opportunity may be quite short. As a result, each of them when thinking about sustainability seems to follow the "commercial model" that the FIRE AB articulated in 2011. They have each used the open call experiments to get initial measurements on the value propositions that they offer experimenters and costs that must be met in supporting them. They have also been able to some extent to identify the alternative cost to such users of finding the same facilities or addressing the same questions without access to FIRE. They are encountering about 1/3 or greater industry and SME presence in their open call proposals.

BonFIRE managed their open call proposers' expectations by publishing several scenarios of possible use cases. Their open call proposals have tended to describe themselves so that they fitted into these use cases. Most of their experiments involve, in one way or another, the ways in which a distributed application, deployed on a multisite cloud will perform when scaled up, how to monitor it and measure it to permit QoS guarantees, Some of the second call experiments are starting to employ large amounts of data, for example, in plagiarism detection. This is not easily tested in a commercial cloud service such as Amazon, because the cost of storage in commercial clouds is often rather high. BonFIRE has by now developed

VPN links to smooth the data transfers needed between their centers so that bandwidth, if not as high as within a single cloud center, is at least predictable. The last characteristic application that has shown up in Santander and probably as well in Experimedia is what might be called the socially-aware wireless network – a network that responds to the needs of its users. This is a theme seen in a wider world outside of FIRE.

2.5 Peer organizations – influence and collaborations

FIRE exists in a world of multiple government and nationally funded activities that address speeding the path to the Future Internet. Nationally funded programs with an emphasis on research similar to that of FIRE exist in North America (FIND and GENI), and several Asian countries. Most of these include testbeds as an essential part of their programs. Additional programs in Europe and in the rest of the world address the migration of future internet technologies and applications into the practical world, incorporate industrial partners and address a shorter term and a faster path to exploitation.

The collaboration between GENI and OpenLab/PlanetLab is quite close. Joint workshops between NSF and EC/ICT people go back some years now. In the experimental plane, Rutgers' WINLAB and OpenLab have collaborated through NICTA, which originally developed the OMF experiment and control software for wireless testbeds and has transferred it to several of the OneLab testbeds, not only wireless ones. Also Japanese, Korean, and Brazilian programs have participated in interactions starting several years ago with discussions, and leading up to joint calls near the end of FP7. GENI racks (purchased by FIRE projects and hosted in some of the OFELIA and Fed4FIRE centers) provide one bridge between the tools developed in FIRE and GENI. For the business-flavored projects, Canada's DAIR is an interesting model.

The OpenFlow open source networking stack (a step towards open software-defined networking) that is employed by OFELIA on its 10 "islands" is also widely employed in advanced work in the US, both in GENI and in some leading edge commercial BigData installations with multiple sites. It has recently been offered by PlanetLab Europe (part of the Fed4FIRE complex) as an overlay network running on the hundreds of distributed clients of PlanetLab.

We discuss FIRE – FI-PPP, FI-WARE, Xi-FI interactions at greater length in the next section. These have a clear focus on the immediate 3 year timeframe. FIRE, as an existing testbed family, is a natural partner for the prototyping efforts in these European FI projects.

A related subject is the flow of ideas and tools from one project to another. Many of the STREPs that were funded early in FP7 (Call 2 and Call 5) have formed key parts of the teams that are now in IPs funded in the later calls, 7, 8 and 10. The directions in which continuing in FIRE has pulled these also contributes to our measurement of future indications. For example, WISEBED has provided key technologies in sensor networking to SmartSantander. Fed4FIRE is picking up many threads of the FIRE portfolio and federating them, not only providing a common login and recording system, but also discovering ways to dispatch experiments that bridge across several of the separate components.

2.6 What's happening in the Future Internet and Net Futures research?

Now let's look ahead into the activities that will occupy FIRE in the coming years. The following sections explore the questions which form the basis for developing future FIRE scenarios (Table 2).

Question	Expected Outcomes
What is happening today in the Future Internet and Net Futures? What is happening today in FIRE?	Set of assumptions on which the scenarios will be based (e.g. H2020 will start Q2 2013, what facilities exist today and in the future, definition of the scope/external environment) Description of the main drivers for change (e.g. federation of facilities, changes to EC structure, etc) Discussion about the trade-offs that allows us to explore uncertainties
What are the possible Future scenarios?	A set of framed scenario description presenting future outcomes based on the drivers we have identified

Table 2 Basis for future FIRE scenarios

The Future Internet is a loosely defined term which describes the migration from current Internet functionality towards infra-structure that is capable of supporting the next generation of applications, services, and networked systems. At present this remains an open-field of innovation, where global initiatives from Europe[1], the USA[2], Japan[3], Korea[4], and others are searching for the prospective breakthroughs. The Future Internet will comprise the expanding network services from a set of identified key domains; [1] [2] commonly specify three constituent domains (or pillars):

- The *Internet of Things* (IoT) is a global, connected network of: tags, sensors, actuators, and mobile devices that interact to form complex pervasive systems that autonomously pursue common goals [3]. The IoT seeks to build connected services upon a highly heterogeneous infrastructure (in terms of network technologies, software platforms, and data) that will scale to trillions of elements communicating vast amounts of data.
- The *Internet of Services* applies the vision of service-oriented computing to Internet-scale systems. A broad-range of service functionality including: software applications, software tools, systems platforms, storage, etc. is globally available to be composed and choreographed using the service-oriented architecture pattern.
- The *Internet of Information* is built upon the ever increasing amounts of data content, where this content may be any type and volume of media, and may be combined, mixed or aggregated to generate new content. It may vary from a few bits (e.g., the temperature that a sensor has measured) to interactive multi-media sessions and immersive complex and multidimensional virtual/real worlds—however, the research problems of large-scale and

[1] http://www.future-internet.eu
[2] http://www.nets-find.net/
[3] http://akari-project.nict.go.jp/eng/overview.htm
[4] http://fif.kr/home.php

heterogeneous media remain common: new technologies will investigate ways to capture, process, store, and output this content [2].

A further pillar is identified by [1]:

- The Internet of People (IoP) connects and unites a growing number of people—including bridges across the digital divide. Everyday users will become the producers as well as consumers of content, and they will form diverse social networks to pro-actively build communities to exchange knowledge; new adaptive networking technologies will be required to support such socially inspired connectivity. Users will become the centre of Internet technology—indeed the boundaries between systems and users will become increasingly blurred by technologies such as Near-Field Communication (NFC), Body Networks, Augmented Reality, and Participatory Sensing.

Hence, the Future Internet offers a broad field for future areas of research that are relevant to FIRE's user communities. To help identify how FIRE will evolve we are exploring these future avenues with the following objectives:

- To identify the requirements for the next generation of experimentation facilities which will allow Future Internet experiment results to be validated, and industry innovations in this domain to be tested.
- To infer future research avenues within the broader research space and identify novel use cases of Future Internet systems to better highlight the significant potential of this field.

Our analysis differs from the traditional approach to projecting the characteristics of the future internet and speculating about the best research avenues to explore. When FIRE was founded, at the start of FP7, a bottom-up approach was most natural, building up from the technical challenges of fully exploiting an infrastructure that permitted converging voice, data, and content distribution, and that extended for the first time to wireless data and voice at the edges of the internet. In the present exercise, we believe it is best to start from the rapid growth of applications that are identified with the four pillars just described, and then ask how the infrastructure will adapt to them, how the development of networking software will evolve, and how some of the fundamental concepts of privacy and security will be defined and provided in the 2020 timeframe.

We are carrying out a two phase analysis of the state-of-the-art, expert opinion and research direction in order to identify the short-term and long-term research challenges in the Future Internet:

- Phase 1: We have analysed the important research areas and topics within the distinct pillars (things, content, and services) for example: Cloud Computing, Big Data, Cyber-Physical Systems, Future Networking, Data and Content-Centric Networking, etc. We have observed the application areas and scenarios where these research challenges will have particular impact. Additionally, we have assessed the requirements of future experimental facilities to validate these research results.
- Phase 2: From the work carried out in Phase 1 we are extrapolating important new relationships and research challenges within Future Internet software and services—particularly those related to complex distributed systems within experimental testbeds.

The Future Internet covers too broad a field of research to perform a thorough state-of-the-art analysis and literature review. Instead, in particular fields existing foresight [4] material is investigated directly. These take the form of literature reviews, expert panels, hot topic workshops, research roadmaps, and expert opinion about future research directions. For example, in cloud computing we consider the following: state-of-the-art literature reviews [5] [6], the HotCloud[5] and LADIS[6] workshop material, interviews with industry experts [7], and research roadmaps. Such material was selected because it is both expert and community driven, i.e. the material is proposed by recognised experts in Cloud Computing (in terms of peer-reviewed articles) and/or forms the collective opinion of the International Cloud Community.

Our analysis of future research in: Future Networking the Internet of Services, Information and People follows; this shows three key outcomes:

- *Integration*: Future Internet systems will integrate a broad range of systems (cloud services, wireless sensor networks, content platforms, mobile users) in large-scale, highly heterogeneous, systems-of-systems.
- *Common research themes*: Scalable solutions, interoperability, new software engineering methods, optimisation, energy-awareness, security, privacy and trust all offer rich fields of research across the combined pillars that form the Future Internet.
- *Federated experimental facilities*: Homogenous experimental facilities are insufficient for the purpose of validating the previously introduced research themes; hence, new federated facilities that are large-scale and highly heterogeneous are required for experimentation.

2.6.1 The Internet of Services (IoS) and Cloud Computing

Two research domains are important for the realisation of the Internet of Services:

- *Cloud Computing*. Built upon a computational model that provides utility computing services, it is naturally suited to underpinning the infrastructure of the Internet of Services. Due to the elastic cost model, new services can be provided with minimal start-up costs, pay predictable costs as the scale increases, and be easily made globally available.
- *Software Engineering of Services*. The success of the Internet of Services hinges on the ability to easily engineer applications and services in order for providers to leverage the benefits of globally built systems with cloud computing costs. Therefore, development tools, autonomic management systems, and new software engineering methods play an important role in this domain's success.

We analyse and present current trends of research in this field and consider how experimental facilities can offer the functionality that such experiments require without significant re-invention of the wheel. Such facilities must support the production of verifiable results, and foster innovation to allow shorter-term to market for new software and services. The results demonstrate that Cloud Computing is indeed a broad field with many taxing challenges to solve. However, initial Cloud experimentation facilities have laid the foundations to address these issues, and it remains the case to add value to these infrastructures.
Cloud Computing consists of a small but well defined and accepted technical taxonomy [5], which we overview here to define the concepts and challenges. Cloud computing software and services exist in one or more of three layers:

[5] 4th International Workshop on Hot Topics in Cloud Computing, June 2012 , Boston, https://www.usenix.org/conference/hotcloud12
[6] ACM SIGOPS/SIGACT Workshop on Large Scale Distributed Systems and Middleware, http://ladisworkshop.org/

- The Cloud Application Layer provides end-users with access to services in the cloud (typically through web portals); this is commonly referred to as Software as a Service (SaaS), and examples include Google Apps[7] and Salesforce Customer Relationship Management (CRM)[8].
- The Cloud Software Environment Layer provides a programming language level environment via a set of service APIs; application developers utilise these services and gain abstractions rich with support for automatic scaling and load balancing. This is commonly referred to as Platform as a Service (PaaS). Examples include Apache Hadoop, Pig, and Google AppEngine.
- The Cloud Software Infrastructure Layer provides fundamental resources to other higher-level layers, which can be used to construct new cloud software environments or applications. Computational resources in the form of Virtual Machines provide the flexibility to deploy defined software stacks. This is typically referred to Infrastructure as a Service (IaaS), and is wholly possible due to advances in virtualisation—Amazon EC2[9] and Eucalyptus[10] are examples of such elastic compute services. Correspondingly storage and access of data hosted on remote disks is referred to as Data as a Service (DaaS)—three main goals of data storage are: availability, scalability, and consistency, however these conflict with one another and different providers will offer different Service Level Agreements (SLA) that satisfy the different dimensions. Amazon S3[11] is an early example of DaaS.

2.6.1.1 Research Challenges for IoS

Programming Models

The growing heterogeneity of application domains being deployed upon Cloud Infrastructure is increasing the demand for new software stacks and programming models that are well-suited to the individual problem requirements. Cloud experts have identified that the prevalent use of map-reduce [8] is not well suited to many problems [7], and in many cases can lead to unnecessary parallelism [9]. The rapid success of systems platforms such as Hadoop[12] (an open source map-reduce execution engine), Pig[13] (a high-level language for describing complex data flows and compiling them to map-reduce programs), and Cassandra[14] (a noSQL database for storing and querying large data sets), have demonstrated the enormous power of Cloud Computing to perform parallel data processing at scale. However, alternative methods must be made equally visible in order to ensure that developers select the most appropriate tool.

This current lack of diversity (particularly within the PaaS space) of established services invites innovation of new programming models and systems software. Customizable workflow execution engines, e.g. CIEL [10]; new novel stream processing middleware[15], and

[7] www.google.com/Apps
[8] http://www.salesforce.com
[9] http://aws.amazon.com/ec2/
[10] http://www.eucalyptus.com/
[11] http://aws.amazon.com/s3/
[12] http://hadoop.apache.org/
[13] http://pig.apache.org/
[14] http://cassandra.apache.org/
[15] http://research.microsoft.com/en-us/projects/grape/

the application of the R programming language [11] are just some of the recent innovations in this area.

In order to support research in this dimension there are key requirements of experiment facilities:

- Availability of cloud test-beds to validate the configuration, scalability, and performance of the new services. In particular, with openness characteristics to allow deep monitoring of the runtime operation of the deployed services.
- Available workload traces from domain applications that can be used to observe the performance of systems in performing tasks against recognised systems benchmarks. For example making a side by side comparison with Hadoop's execution of the Page Rank application.

Optimisations of Cloud Architectures

There is scope to investigate the optimisation of cloud services and infrastructures at each of the cloud levels, i.e., from infrastructure through to SaaS. Therefore, different stakeholders will wish to perform different types of optimisations. For example, data centre management will wish to maximise the utilisation of their resources (i.e. pack in as many applications) in order to reduce their costs without breaking the SLAs in terms of elasticity, performance and reliability; this may be particularly governed by the need to minimise energy costs. Hence they may wish to experiment with local optimisations, e.g. new resource scheduling algorithms [12], or new VM migration techniques.

Beyond scheduling and storage, the optimisation and isolation of network resources available to Cloud applications is often ignored, in turn providing poor network SLA guarantees. For example [13] documents solutions to the "Noisy neighbour problem" where a neighbour in multi-tenancy tries to grab network resources. A key requirement is the provision of predictable network performance and its appearance as a dedicated resource. Software Defined Networking (SDN) and OpenFlow[16] network hardware are key technologies here; for example, CloudTalk [14] researched within the CHANGE STREP[17] examines methods to observe network topology and plan application optimisations using Software Defined Networking solutions.

Alternatively, PaaS and SaaS services may wish to experiment with different configurations and utilisations of the underlying services provided by the IaaS layer. This may be particularly relevant where it is difficult to manage the relationship between cloud applications parameters and IaaS and PaaS guarantees. Experimentation therefore forms part of the software engineering process for innovative software stacks. For example, experiments in Bonfire have examined new approaches to QoS based software engineering of new cloud services[18].

In order to support research in this dimension there are key requirements of experiment facilities:

---

[16] http://www.openflow.org
[17] http://www.change-project.eu
[18] http://www.bonfire-project.eu/innovation/qos-oriented-service-engineering-for-federated-clouds

- The availability of *heterogeneous environments*: different software stack configurations, different data centre hardware in order that new algorithms can be tested in different environmental setups.
- *Scalable experimental facilities*, i.e., the availability of resources to match required workloads; this may be in the 10K-100K range for Virtual Machine instances; and data experiments with Terabytes of data.
- *Rich monitoring* facilities with open access to all elements of the cloud infrastructure, e.g., from information about the resource utilisation of individual hardware machines up to the capacity usage of the data centres themselves. Importantly, energy usage monitoring throughout the infrastructure and carbon footprint breakdown will form additional important models.
- The *availability of SDN* experiment facilities within cloud testbeds in order to experiment with the optimisation of scheduling, storage and network performance.

Privacy and Security

Two of the greatest concerns for the uptake of Cloud Computing beyond its present usage patterns and application towards richer, complex and larger-scale application domains including government and healthcare systems are privacy and security. Is it possible to ensure that data stored in the cloud, and computation that is run in the cloud is secure from unauthorised and unlawful access? This is particularly important given the multi-tenancy architectures of cloud computing resources, whereby processes execute on the same machine and utilise shared resources. Furthermore, the cloud model introduces new points of attack beyond traditional Web vulnerabilities, e.g. during the transfer of VM images, or the use of unpatched software in virtualised settings [7]. To exemplify this significant problem, [15] describes how researchers from Darmstadt were able to expose numerous security flaws in images deployed on Amazon Web Service and then use this to access sensitive data with the potential to cause significant financial damage.

The outsourcing of data also brings privacy concerns, as it becomes increasingly difficult to manage who reads personal or sensitive data (e.g. criminal records, medical records). While data can be encrypted, at some point it will need to be read, and therefore, methods to describe privacy SLAs and enforce them remains an open issue. This is especially important to consider where data is stored across data centres that operate within different legal jurisdictions; or where service providers change the conditions of their SLA.

Therefore, this will necessarily lead to significant research into security attacks, and their patches and novel models of trust in the Internet of Services. The employment of hybrid clouds where systems are deployed across private clouds (where infrastructure is owned and managed by the user) and public clouds (where infrastructure is owned and managed by cloud providers) are an important consideration in the engineering of such trusted services.

In order to support research in this dimension there are key requirements of experiment facilities:

- Security tools to detect exposure of vulnerabilities, attacks and unauthorised access of data, images, and software.
- Security attack simulations, e.g. traces that perform a distributed denial of service attack.

Interoperability and Vendor Lock In

A key requirement of the Internet of Services is the open composition of globally deployed services. However, significant heterogeneity arises where services are developed independently of one another by third parties without knowledge of who will interact with the services in advance. Hence, such composition may face significant heterogeneity problems.

One notable problem associated with the reliance on public cloud providers that do not employ open standards is the vendor lock-in problem. That is, when you choose to utilise a particular provider you tie yourself to their APIs, and essentially close yourself off from the rest of the marketplace. This may be of particular concern where IaaS or PaaS SLAs are changed. Further, alternative cloud providers will offer cheaper pricing plans, and or better service. However, in order to migrate, the cloud applications may need to be re-engineered at a significant cost beyond the potential gains. Open standards are the obvious and well established solution to such an interoperability problem:

> *"the extent by which two implementations of systems or components from different manufacturers can co-exist and work together by merely relying on each other's services as specified by a common standard"* [16].

However, full standardisation is not an achievable or required goal:

- A one size fits all standard cannot cope with the heterogeneity of infrastructure, functionality, and services already available within the different cloud layers.
- New services and applications emerge fast, while standards development is a slow, incremental process. Hence, it is likely that new technologies will appear that will make a pre-existing interoperability standard obsolete; this may be particularly true in the cloud domain where the marketplace drives rapid innovation.
- Legacy services remain useful. However, new standards do not typically embrace this legacy issue; this in turn leads to immediate interoperability problems.

Hence, it is likely that standardisation of parts of the Cloud software stack will gain traction e.g. OpenStack and the Open Clouds standards[19] for key infrastructure services, but Interoperability solutions that broker in the face of heterogeneity and support the migration of Cloud applications will also come to prominence. The EU project Broker@Cloud[20] is investigating mediation in the face of diverse QoS parameters from providers. Alternatively, migration solutions are also being researched: For example, AppMigration [17] proposes new methods to migrate Google App Engine VMs from an OpenStack[21] based architecture to a Eucalyptus based architecture.

In order to support research into interoperability solutions there are two key requirements:

- Availability of heterogeneous configurations that mirror the interoperability problems of complex real world systems. Hence, homogenous, standards-based experiment facilities are typically limited with regards to interoperability research.

[19] http://occi-wg.org/
[20] http://www.broker-cloud.eu
[21] http://www.openstack.org/

- The ability to discover and monitor behaviour at all levels, i.e. beyond the discovery of APIs. It must be possible to observe and monitor network packets, protocol packets, and importantly observe the content and format of the individual data structures employed.

Software Engineering: Autonomic and Self-Managing Systems

The complexity of Internet of Service based systems ensures that their development, deployment, and management will quickly become difficult to oversee and expensive to maintain—and this will limit the future growth of these systems types[22]. Hence, there are particular research challenges in the field of software engineering to develop the next generation of methodologies and tools that will support the engineering side of the Cloud Computing domain. This will typically embrace autonomic computing principles where systems are designed to be self-managing, self-optimising, and self-repairing.

Software engineering research may take the following directions:

- Software engineering methodologies that estimate the cost of migration to the cloud. Taking legacy systems and deploying and managing them in the cloud may be a non-trivial task and it should be possible to estimate and forecast the cost and trade this against any potential gains. Experimental facilities should provide the features and case studies to validate these estimation methodologies.
- Evaluating ease-of-use. From configuration of software stacks in the cloud to the experience of end-users of SaaS platforms it must be possible to evaluate the effectiveness of new solutions and their ease-of-uptake by new users.
- Autonomic solutions. There are numerous research solutions in the field of autonomic computing applied to distributed systems and complex systems in general, and hence these have naturally extended to the cloud, e.g. examining new tuning algorithms, conflict resolution solutions, repair and fault-tolerance strategies. Validation of this research requires the ability to compose large-scale systems, and importantly observe both their behaviour and the behaviour of other competing systems. This is particularly true of cloud systems where isolation may be broken by hidden dependencies—"Icebergs in the Clouds: The Other Risks of Cloud Computing" highlights where a mission critical service uses two providers A and B in a federated cloud but both A and B share a network provider leading to a hidden common point of failure.
- Quality of Service Guarantees and Service Level Agreements. Virtualisation of resources (particularly where VMs share multi-core architectures) have led to a limited range of performance guarantees [5] and weak SLA languages that do not match the requirements of higher-level cloud applications. Hence, there is significant research potential in the definition of new SLA languages, new tools to monitor such agreements, new optimisation methods to maximise resource usage against these SLAs, and new software tools to describe workflows that leverage and broker different service provisions with these capabilities.

2.6.1.2 Experimental Cloud Testbeds

There exist two significant Cloud computing facilities which federate geographically remote facilities, and provide facilities for Cloud Service deployment, testing and experimentation:

[22] A Roadmap for Advanced Cloud Technologies under H2020 http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-expert-group/roadmap-dec2012-vfinal.pdf

BonFIRE:[23] consists of a federation of six facilities providing access to heterogeneous compute, storage and networking resources. OpenNebula is leveraged for the IaaS platform and hence offers a standards-based approach to its usage. A key feature of the test facility is the availability of deep monitoring of the architecture and network behaviour in order to support the validation of experimental research results; this is complemented by a synthetic workload generator that can be used to meet experimental input requirements. Bonfire is a closed consortium of resources, with no model to add future resources and facilities into the infrastructure or federation. Use of the platform is open to users with legitimate requests and requirements which are screened in advance.

OpenCirrus: OpenCirrus[24] is a global, multi-datacentre scale research testbed across 15 sites making available 20K cores. It supports a number of experimental IaaS platforms and Eucalyptus, along with configurable PaaS stacks such as Hadoop and Pig. Wider coverage includes the convergence with software defined networking and networking testbeds e.g. OpenCirrus on GENI— where particular technical challenges include mutual authentication, authorisation and access for OpenCirrus with PlanetLab. A notable feature of the facility is the availability of multiple large-data sets and realistic Cloud workload traces to support the benchmarking of new systems innovation.

The available testbeds present significant functionality to the user. The important element of deep monitoring (not generally available in public service providers) heightens the usability for research experiments. However, there remain questions regarding overall suitability: to what extent are data and traces available? They are not publically available for analysis. Can scalability experiments be realistically evaluated, where the headline number of resource e.g. 10k instances may be reduced to a couple of hundred actually available to a user? How sustainable are the testbeds with respect to the addition and removal of new cloud facilities (cf. Planet lab model)? How easy is software experimentation in terms of configuration and autonomic self-management in largely homogenous software stacks with portal based and API access to services? Is the researcher left to build their own setups atop the bare-bones facilities without re-use potential?

2.6.1.3 Other Important Trends

Green ICT in the Cloud

The EC report: "A Roadmap for Advanced Cloud Technologies under H2020" identifies the need for optimized consumption and efficiency of future platforms because facilities or datacentres currently concentrate 23% of the overall ICT CO2 emissions and are estimated to consume between 2.2 - 3.3% of the UK's total electricity, 2% in the Netherlands, and 1.6% in Germany. Therefore, innovation that reduces energy consumption has the potential to have an enormous impact on the overall ICT energy bill. Potential research scenarios and areas:

- Optimisation experiments within local datacentre facilities. Investigation of the load-balancing of resources versus the trade-off of redundancy. The management and transfer of data. The investigation of autonomic behaviour to optimise energy across a facility. This requires accurate measuring and reporting of energy consumption to identify where gains are made, and or leaks are occurring.
- Placement and utilisation of datacentres in a global federation. Understanding of costs of sites in different domains, e.g. investigating the trade-off between sending computation

[23] http://www.bonfire-project.eu/
[24] https://opencirrus.org/

and data to the site with lower energy costs. This requires predictive costs for applications based upon common and similar traces performed in the experiment facility.

- Energy aware cloud topologies and potential integration with the Smart Grid. Datacentres may be collocated with renewable energy sources that have surplus generation, which instead of being sent to the Grid is used for load-balanced work from the across the federated cloud. Therefore, experiments to predict how, when and where to transfer data and computation can underpin the creation of a smart, energy efficient, green cloud.

Community Clouds

Cloud computing need not be performed in datacentre facilities—computation facilities can be made available from end-users or donated by institutions (e.g. time when a cluster is free). This philosophy builds upon altruistic models of distributed systems as highlighted by the Distributed.NET[25] and SETI@Home[26] projects to solve important scientific problems. Hence, there is significant potential in the building of new cloud services on such spontaneous infrastructures.

- How to provide IaaS services in the face of resources that fluctuate and fail in a completely different pattern to centralised infrastructures. Is computation the only feasible service, or can reliable data storage be created?
- What applications can be employed? Are fully commercial-like public clouds possible, with a corresponding payment and reward model, or will they be constrained to particular, community/charity driven applications only?
- New PaaS implementations will be required in the face of such infrastructure, e.g. Hadoop will not operate directly on a community cloud. The whole of the cloud software stack is open to research and optimisation in this area.

There is significant potential to integrate community cloud testbeds into emerging networking experimental facilities e.g. CONFINE[27] a community-based wireless testbed. Further integration into BigData initiatives will be also fruitful; regarding why users will donate resources: the specific provision to run only specific Big Data experiments e.g. the Virtual Observatory, or analytics of local data important to their community and Smart City.

Network-aware Applications

OpenFlow and Software Defined Networking initiatives are central to future cloud services and applications; application context is a key driver in the movement of data and computation. For example, VMs that can migrate themselves to locations optimised to perform the work, i.e., at the most relevant data source, or where costs are minimised. This raises research challenges with regards to the vertical integration of applications such that they can make the network respond to their needs; how cloud services and infrastructure are federated with software defined networking testbeds; and what is provisioned to support application developers leverage the network services directly.

---

[25] http://www.distributed.net
[26] http://setiathome.berkeley.edu/
[27] http://confine-project.eu/

2.6.2 The Internet of Things (IoT)

The Internet of Things creates a global network of interconnected "things" or objects. The key feature of this environment is heterogeneity. Objects can be one of a number of different hardware devices, including: RFID tags, sensors, actuators, wireless sensors, mobile devices, vehicles, UAVs, workstations, etc. With this heterogeneity, devices will utilise different Operating Systems, different networking technologies (Bluetooth, Zigbee, 802.11abgn, GSM, 3G, 4G, IR, etc.), different software platforms, and communicate data using different protocols and data formats. Therefore, many of the research challenges in IoT are centred upon the taming of such heterogeneity:

- *Identification*. Objects must be identified uniquely and in different ways for different purposes; therefore, an object may have multiple IDs for each functionality c.f. a serial number and an IPv6 address. Objects may number in the trillions and hence easy-to-assign and easy-to-use addressing schemes that are inherently understandable across the architecture are required.
- *Service Oriented Architecture*. Systems are engineered to follow a service oriented vision of objects, i.e., that they are connectable with one another to build co-ordinated and intelligent systems who work together towards common objectives. However, such architectures must ensure that interoperability of objects is meaningful (so called semantic interoperability). The architecture must be extensible and scalable to the extent that architectures of trillions of objects do not affect the performance or require the re-engineering of existing infrastructure. Such architectures are suited to asynchronous, data-centric behaviour (i.e. the flow of data is central to operation) and must tolerate device mobility and periods of disconnection (i.e. delay tolerance).
- *Open middleware* provides the software building blocks to realise the architectural vision. Deployed on objects and in the infrastructure it provides the programming abstractions to help developers realise solutions, and in particular shield them from many of the complexities of building distributed systems (i.e. interoperability, security, fault tolerance, among many others).
- *Security and Privacy*. Systems in the IoT will observe and collect significant amounts of data about users (much of which is personal, and or can be used to infer personal information e.g. location data). Therefore, challenges involve the creation of solutions that are both privacy-aware (i.e. they can determine where potential privacy leaks may occur) and privacy-preserving (i.e. information should be made anonymous).

2.6.2.1 Research Challenges for IoT

Interoperability

The research challenges for achieving interoperability in the highly heterogeneous, complex and pervasive systems that characterise the IoT are well established [18]. A particular novel characteristic is that dynamic interoperation is required to achieve spontaneous interaction between mobile systems. Interoperability cross-cuts many layers of the system architecture and [18] highlights these areas in the software stack:

- *Data Interoperability*. Different systems choose to represent data in different ways, which manifests problems at two levels: i) syntax: where data may employ different formats to represent the data, e.g. XML versus JSON; and ii) semantics. Even with the same format there is no guarantee that systems share the same understanding of data, e.g. the equivalence of data fields labelled "price" and "cost".
- *Middleware Interoperability*. Heterogeneous protocols employ a range of communication protocols for the discovery of other devices, to interact with other devices, and to exchange data. The characteristics of the systems means that common standard protocols

cannot be enforced, i.e. some are delay tolerant, power efficient, etc.—hence, it is not sufficient to say HTTP REST is the solution to the protocol interoperability problem in this domain.

- *Application Interoperability*. Design decisions made by developers introduce interoperability problems, e.g. differences in abstractions: publish-subscribe, streaming, request-response behaviour. The same functionality may be captured in one operation in one system versus multiple operations in another, e.g. (GetData) versus (GetTemperature & GetNoise).
- *Non-functional Interoperability*. Systems may have different security, performance and dependability requirements. Therefore, if functional interoperability is achieved the non-functional requirements may still be violated e.g. a system requiring message latency of 5ms interoperating with a system providing 10ms.
- *Network Interoperability*. Where devices employ different wireless networking technologies (802.11b, Bluetooth, Zigbee), routing protocols, and addressing methods—end-to-end interoperability is required to connect systems to ensure the flow of data.

Solutions to such interoperability problems include global standards and associated middleware platforms c.f. W3C, Grid, Cloud, etc. Enterprise Service Buses also provide bridging solutions to resolve protocol and application differences. Ontologies and the Semantic Web have proposed solutions to semantic interoperability problems. However, these solutions are largely targeted at Enterprise services and their rich pool of resources—the interoperability problem in domains such as wireless sensor networks has been largely ignored.

There is significant scope for research-based solutions to have impact in the area, particularly those that investigate semantic solutions in heterogeneous and resource-constrained environments. Hence, experimental facilities for validating interoperability solutions include:

- Heterogeneous testbeds composed of diverse object devices, utilising different communication protocols and wireless networks. The availability of a broad range of application data exhibiting significant variation in format and semantics.
- The availability of domain ontologies and meta-data describing the behaviour and content of the systems and wireless sensor networks that compose the test-bed.

Privacy

Privacy remains a major research challenge with respect to the uptake and acceptance of IoT systems, which is particularly true of sensor-based systems. This is mainly caused by people not knowing that they are being sensed. [19] argues that systems should be designed with the following properties, and hence research and innovation with regards to technologies providing such properties is required:

- Users must be aware that they are being sensed.
- Users must be able to choose whether they are being sensed and be able to opt-out. This is particularly relevant to localization and tracking technologies: i.e., that there is no way for an attacker to reveal the identity of the person or object, and that localization and tracking is not possible without explicit agreement or knowledge.
- Users must be able to remain anonymous.

While these remain technical challenges—it is not feasible to guarantee such requirements with a purely technical method; legislation must also be created in this domain for the technical solutions to adhere to, and offer preventative punishment for infringements. Privacy requirements (and indeed other non-functional and human centred requirements) require novel experimental facilities beyond simple technical deployments in the wild:

- It is particularly important that the users themselves are able to be observed and form part of the experiments. For example, ethnographic, user evaluation, and quality of experience experiments must be supported to determine the extent to which new privacy technologies are successful, and the level of trust users have in these systems.

Wireless Sensor Networks as a Shared Resource

Initial wireless sensor networks were created as single use resources, i.e., they were deployed to achieve particular application functionality, e.g., a flood monitoring tool at a specific location. However, where sensors are pervasively deployed into multi-purpose facilities such as buildings and cities there is the potential for multiple application types (crisis management, smart energy management, healthcare and transport) to share the computation, storage, networking and data resources of the sensors. There are a number of research challenges with respect to creating multiple virtual sensor networks atop physical sensor networks:

- Virtualization schemes to share the physical sensor network resources. Melete [20] and Federated Secure Sensor Network Laboratory (FRESnel) at Cambridge University[28] investigate solutions to application isolation in shared sensor networks; this is demonstrated with environmental monitoring, office occupancy, and appliance monitoring applications all operating in parallel [21].
- Programming abstractions to shared data resources. Software to virtualise the content provided by the sensor deployment, and in turn trigger different event types to correspond to different application requirements, e.g. using SQL-like statements [22]. This is particularly relevant to sensor deployments where physical resource virtualisation is not possible, or is not necessary for the services required and provided.
- Energy management and optimisation algorithms to maximise the limited resources available and reduce the energy footprint. Indeed, this is particularly relevant in the field of wireless sensors with a limited power supply. While there has been significant work on power saving methods in this domain, the conflicting requirements from resource sharing opens up new dimensions.

A number of experimental facilities and deployments have been developed in this field:

- Wisebed[29] supports resources for dedicated experimentation in the field of WSN, i.e. the testing of new communication protocols across a heterogeneous, federated, Pan European testbed. Experiments gain isolated access to resources, i.e., experiments are scheduled time on the resources where the single VM images are upload to the resources.
- Senslab[30] is a federated wireless sensor network composed of 1K sensor nodes deployed across four geographic location in France. Like Wisebed it supports experimentation of new WSN protocols; however, experiments can run in parallel on the network.

---

[28] http://www.cl.cam.ac.uk/research/srg/netos/fresnel/
[29] http://wisebed.eu/site/
[30] http://www.senslab.info/

- The GENI Sensor/Actuator Network Testbed [23] provides implementation of sensor virtualization using Xen, i.e., it is a full hypervisor approach that is therefore targeted at resource-rich sensor nodes.

Software Engineering

The complexity and challenging requirements of Internet of Things applications and systems will create new engineering and programming methodologies. Abstractions centred upon the programming of individual nodes, known a priori failure models, and a flat network topology have been argued to be not fit for purpose [24]. Instead sensor-based application programming should concentrate on:

- *Goal-oriented abstractions*. High-level definitions of the application requirements should be compiled down onto the infrastructure, only at runtime when the conditions and infrastructure configuration is known.
- *Introspection*. A key functionality of all sensor networks is the provision of information about node, network, and application behaviour in order to inform autonomic management of the deployments and support the evolution of the network behaviour over time.
- *Provenance*. Wireless Sensor networks are inexact and hence fully supported belief systems about the information provided must be built alongside the shared sensor infrastructures. Information about how data was produced (e.g. location and number of sensors) must accompany the data itself.

2.6.2.2 Smart City Testbeds

Smart Cities are a key demonstrator of the Internet of Things philosophy; they are large scale, heterogeneous, and have a diverse range of application domains with the fundamental goal of achieving improved services to the betterment of society. There has already been significant research into Smart Cities in terms of the software and infrastructure, and numerous initial deployments that generally aim to achieve the following common objectives:

- reduce congestion using intelligent transport systems;
- support sustainable waste and energy management through energy usage reduction schemes, and energy & waste recycling schemes;
- provide safer cities using intelligent crime and crisis management systems and reduced emergency response times (e.g. integrated with smarter transport systems);
- offer improved healthcare facilities, especially to an aging population;
- support environment improvement and community initiatives—using monitoring and reporting systems to inspire environment conservation [25] (e.g. reducing CO2 footprint competitions[31]) and using social networking and crowdsourcing technologies to improve local communities e.g. the Open311 project[32] for reporting non-emergency problems in US cities.

Dubuque, Iowa

A public-private initiative between IBM and the city of Dubuque, IOWA provided a demonstrator for the tools and technologies developed by IBM from their Smarter Planet[33] and Smarter Cities projects. In particular, this focused upon the following application areas:

---

[31] http://news.lancs.ac.uk/Web/News/Pages/079B5B068091B9B68025756D0040C3CA.aspx
[32] http://open311.org/
[33] http://www.ibm.com/smarterplanet

- *Water management*. Infrastructure was instrumented with sensors and smart readers to report on water usage in the homes. These reported at either daily or ¼ hour intervals, sending the data over a wireless network to local gateways. This was combined with other data sets: weekly weather data, initial GIS data, housing data, and demographic data to be analysed in a Smart Water Portal (executing in the Cloud). Analytics applied to the data showed local and city-wide patterns and such reporting was used to identify potential leaks and also incentivise households to reduce usage and costs. Additional social networking activities provided further incentives to save e.g. most saved games [26].
- *Electricity management*. A similar infrastructure setup and data analysing tools were applied to the electricity grid to identify leaking/phantom energy usage and support the reduction of energy costs [27].
- *Smart Travel*. Combined measurement data from smart phones and RFID tags was collected about people's movement in the city and combined with transit data, census data, and geo-spatial information. This was analysed to determine how people travel around the city—such that optimisations can be made to encourage public transport usage e.g. bus routes that serve the requirements. The project was part of the City-in-Motion directive from IBM, and the results are being applied to larger cities e.g. Istanbul.

SmartSantander

Dubuque is an example typical of Smart City deployments in that the applications, technologies and usage are pre-designed for a specific purpose, e.g., smart energy. However, SmartSantander[34] which is Europe's flagship smart city deployment seeks to provide general purpose facilities that can be used by a number of different applications and systems; that is, the infrastructure is open to testing and experimentation. For this, the city of Santander will instrument itself with 12K sensor nodes. Three use cases highlight existing use of the infrastructure:

- *Smart parking management*. The use of embedded sensors to detect free parking spaces, with the information disseminated via repeaters on lamp posts. The collected information can be used to inform traffic systems of free spaces such that cars can be directed. The faster time to locate a space reduces fuel consumption and pollution.
- *Environmental monitoring*. Noise, pollution, luminance, and temperature sensors are attached to fixed locations (e.g. lamp posts); and also to city vehicles (buses, taxis, etc.). Information from lighting sensors can help inform the maintenance of the lighting network.
- *Smart Irrigation--Parks and Gardens*. Santander contains over 60 parks which when instrumented with a range of sensors (air pressure, humidity, rainfall, etc.) can be used to inform smart irrigation systems which will reduce water wastage and energy usage of existing irrigation methods.

SmartSantander is notable in that it aims to go beyond an experimental computer science platform for Internet researchers to validate their cutting-edge technologies (protocols, algorithms, radio interfaces, etc.), and instead involves a wider range of stakeholders: industries, communities of users, and experiments to assess new services and applications. However, while the obvious achievements of a large-scale deployment are noteworthy the facility itself does not capture, nor demonstrate the real potential of smart cities and their applications. For example,

[34] http://www.smartsantander.eu/

isolated applications that monitor pollution and parking spaces do not represent integrated systems supporting multiple stakeholders. Nor do they highlight the general purpose nature of the facilities to be used in different ways.

Glasgow

In 2013 Glasgow was chosen as the Future Smart City demonstrator in the Technical Strategy Board (TSB) competition[35]. The 30 submitted proposals for the TSB competition largely focus on the importance of existing city data and the integration of existing systems, with little innovation concerning new application areas—although this can be explained by the local authority planning organisations driving the submissions[36]. The key features of the Glasgow demonstrator will be:

- *Glasgow City Management System*. A Data Repository stores: economic, demographic, employment, CCTV, traffic management, public transport, electricity (supply, demand), gas (supply, demand), water (supply, demand, flow, flood risk management, quality), waste, noise, public realm (footfall), metrological (temperature, air quality) and geological (accessing sub-surface knowledge data) in a cloud-based store together with an ontology to add meaning to the data. An Intelligent Operations Platform provides the tools to collect and analyse data in a meaningful way. A City Dashboard provides tools to visualise the current operation of the city. A City Observatory then provides open access to city data to underpin new innovations.
- *Integrated Social Transport*. Increase multi-tasking sharing of social transport schemes (non-emergency response government vehicles, healthcare transport, etc.) and then integrate these vehicles into Glasgow's bus priority signalling scheme. The goal is to reduce government spending, congestion, and pollution especially in UK cities where these vehicles make up a significant percentage of local traffic.
- *Sustainable, Safe Street Lighting*. Investigating the effect of dimming street lighting, i.e. analysing against crime, and uptake of other activities (walking, cycling, etc.). Also the use of dynamic lighting and CCTV monitoring.
- *Energy Efficiency in Building and Housing*. Investigation of methods to pass on benefits to consumers within smart-managed buildings.

CitySDK

CitySDK[37] is a particularly interesting initiative within the CIP ICT PSP to make available an open, lightweight software development kit to be used to rapidly engineer new services in urban areas. While the SDK is not publically available to evaluate the extent to which it is general purpose—it highlights software engineering at the forefront of experimentation in this domain. There are challenges in how developers will leverage the facilities, and also how they can create tools that real users will build the next generation of smart city systems with. Hence, testbeds with real developers and/or community inspired hacktivists to validate new engineering methodologies are of equal importance to networking infrastructure.

---

[35] http://www.innovateuk.org/content/news/glasgow-selected-to-be-city-of-the-future.ashx
[36] https://connect.innovateuk.org/web/future-cities-special-interest-group/feasibility-studies
[37] http://www.citysdk.eu

FIREBALL – Catalogue of Facilities

FIREBALL[38] has been a Support Action within FP7 which aimed at stimulating Future Internet experimentation in Smart Cities by using the concept of Living Labs as user driven open innovation environments. FIREBALL created a Connected Smart Cities network which is strongly related to the European Network of Living Labs and also worked on building bridges between the FIRE, Smart Cities and Living Labs communities to stimulate collaboration. It envisaged collaboration between different types of facilities and between organisations to create innovation ecosystems for Future Internet research. While there is little evidence of tangible integration and federation of the facilities (in Helsinki, Manchester, Barcelona, etc.) — FIREBALL importantly identified a resource discovery pattern utilising a catalogue of facilities to match the experimenters' requests to what is available [Fireball D1.2, 2012; FIREBALL 1.3, 2012].

2.6.2.3 Analysis of IoT

There are a growing number of demonstrator applications across hundreds of Smart City installations. These have emerged to create self-sustaining cities and concentrate particularly on improvements in transport, energy, safety and healthcare. However, the majority of the applications are self-contained—acting as information silos, and therefore their combination remains an area rich in potential. Do such systems (when combined) identify new optimisations? Do such systems conflict with one another? Many of the demonstrators highlight integration at the data level, e.g. the application of analytics to multiple data sources; however, there are very few examples with the integration and tuning of multiple autonomous sub systems.

Smart Cities are isolated, location-centric systems. That is, at present, they do not consider their roles in wider geographic regions and communities. Are there software and services that support the integration of multiple smart cities in a wider region? Can smart city applications identify improvements based upon shared relationships, e.g., broadening intelligent transport systems (e.g., railway connections, haulage, commuter routes)? Can higher governance improve shared data, computation and infrastructure resource usage by cities? The data from multiple cities can be fused and analysed, and the cloud infrastructure particularly related to security and privacy (e.g. private cloud facilities) can be shared.

Cities have very different characteristics e.g. population size (from tens of thousands to millions), demographics, and facilities. They also place importance on different problem areas, e.g. unemployment versus congestion. Therefore, there may be scope for innovation in the face of different requirements—can the same system be deployed in different cities? Does the application scale in the same way? Is the data provided as accurate, precise, and reliable? Hence, there may be need for multiple facilities to be available providing this heterogeneity such that new research can be validated.

There are a number of potential scenarios in Smart Cities that are worth further exploration and seek to demonstrate combined systems:

- *Pollution management*. Integrating pollution monitoring and noise monitoring WSN, and intelligent transport systems in the city e.g. traffic monitoring WSNs and traffic management actuators. Further, vehicular wireless networks can be integrated, e.g., advising of route changes to individual vehicles. Hence, fine grained monitoring of local

[38] http://www.fireball4smartcities.eu

sensor data can be integrated with traffic observations to determine the level of pollution caused in particular areas. Such information can support autonomic management of traffic flow in the short term (when and where problems occur); and also inform long-term traffic and pollution planning with regards to optimal infrastructure development.

- *Crisis management.* Since 2010 Sao Paolo has been frequently struck by flooding that has caused severe financial costs, as well as the cost of 13 lives in 2011[39]. The integration of environmental monitoring and prediction systems, with city management actuators (e.g. transport and emergency response systems), city wide notification platforms (networks of public displays c.f. the EU public-display initiative www.pd-net.org), and users' mobile devices offer potential city-wide integrated smart crisis management solutions.
- *Commercial activities*. Can real-time data streams from sensed areas of a city be leveraged for commercial opportunities, e.g., for provisioning information centric applications and services like Zoopla, Rightmove and ACORN. Given that traffic management, street lighting, and similar city systems are not particularly innovative. How can infrastructure and data (displays, augmented reality facilities, city media) be leveraged for innovative commercial opportunities: city based games, location-based marketing.

2.6.3 The Internet of People (IoP)

Social computing refers to the use of computational systems to support social behaviours [28]. Web 2.0 systems, social networks (from large scale networks such as Facebook and twitter to smaller, self-created networks), wikis, and blogs in combination with the increase of smartphone usage and embedded sensors have revolutionised the way people interact with one another— hence we are already experiencing the Internet by and for the People [1]. Sensed information builds rich views of users and their interactions, e.g. fine-grained location and activity traces using RFID, Smartcard usage and GPS technologies. User generated and shared content in terms of text, pictures, music, videos and recommendations adds to sensed data to provide a wealth of data that can then be used within software and services that seek to improve the general welfare of users and their communities.

Such services significantly reach beyond M2M systems and traditional interactive systems and place users as the central and most important element. With new abstractions come new challenges, and [28] identifies three grand challenges for research in Social Computing systems and services:

- Environmental sustainability is a key driver of social systems that seek to reduce carbon footprint, minimise energy usage, and seek novel ways to recycle energy and waste. However, social computing technologies must themselves be sustainable and not detract from the overall goal of the systems they implement.
- Promote individual wellbeing. Research should not simply focus on the needs of the community, but also the needs of the user. Novel technologies to improve individuals wellbeing should be investigated, e.g. personal tracking systems (referred to as the quantified self) can help identify how users can improve their health in particular. It is already recognised that the promotion of such systems has the potential for significant savings in global healthcare costs.
- Build fair digital ecosystems. It must be possible for people to regulate their identities to bridge problems such as "The Internet never forgets!"; this includes engineering systems such that hidden inferences are not leveraged. For example, the use of information

[39] http://riotimesonline.com/brazil-news/rio-business/the-costs-of-flooding-in-sao-paulo/

extracted from FourSquare and Twitter data to build the mashup system: "http://PleaseRobMe.com".

2.6.3.1 Research Challenges for IoP

Social computing systems will embed themselves into the realm of the Internet of Things and Services and therefore face many of the same technical challenges, i.e. tackling significant heterogeneity and complexity. Indeed many scalability problems are being driven by user generated content, e.g. 30 billion pieces of content shared every month on Facebook alone [29]. However, the social dimension brings highly tailored research challenges:

Accuracy and Efficiency

Social systems rely on an accurate digital footprint of users that is built up from multiple sensed and user generated sources. Such footprints may in turn require new context measures to be utilised (beyond physical characteristics such as location), e.g. emotional sensing. Hence, new hardware and software techniques to capture this context may emerge, and in turn there must be ways to measure the accuracy of these technologies. While it may be possible to easily verify captured location data, building a psychological portrait is more complex: requiring users and experts as validators.

Context capture is a potentially resource intensive activity, hence research is required into the optimisation of social systems. Continuous context from multiple sources including video and audio capture and analysis may be overkill for the requirements of individual applications. Hence, there is a trade-off between the accuracy of collected information and the efficiency of collecting it; this is especially important for applications with sustainability requirements. Hence, social computing systems must support the experimentation with these parameters in order to fine-tune their operation.

Open Tools

The transformation of social data into knowledge that can inform and power community improving systems relies upon data fusion, analysis, and mining technologies that are often the preserve of experts. This limits technological uptake and indeed is counter intuitive to the spirit of social computing. Hence, there is a need to democratize social software and the corresponding design and development tools such that a broader pool of users can build their own social computing technologies; the previously mentioned CitySDK is one example, and from the Web community—Social Engine[40] demonstrates the idea of providing tools to build and maintain a social networking web site. There remain many interesting software engineering challenges here: can systems in the wild (within towns and cities) be built by motivated communities with limited software project experience? The Web is not the preserve of business and government systems, and neither should systems in the Future Internet be—the availability of diverse Living Lab facilities as part of experimental testbeds will support the validation and development process of such new tools.

Privacy and access to Social Data

Social computing systems and their manipulation and production of personal data magnify the privacy problem, i.e. it zooms in on personal traces as well as broadens its dissemination with community generated data. Hence, there is significant need for advanced privacy preserving

[40] http://www.socialengine.com/

mechanisms with the user in the loop. Privacy mechanisms where different granularities can be specified, middleware platforms that can maintain confidential links between systems etc. are possible directions for future research.

Social media content (and in particular user generated content) is typically locked with individual service providers, e.g. the upload of pictures to Flickr and Picasa. Hence the user loses ownership, there is potential for duplication, user content may be used beyond its original intent, and future access rights may be changed. Hence, mechanisms for separating the content itself from the multiple services that use it [30] better supports a user's ownership rights and allows traces of how the data has been accessed and used to be built.

eHealth and Activity

eHealthcare provides a notable application domain for social computing within the Future Internet. Such systems can range from small-scale individual monitoring applications (c.f. the quantified self) to fully blown computerised national medical systems. While FIRE facilities have limited support for integration with medical facilities (i.e. hospitals) and professionals (doctors)—there is scope for experimentation with individual monitoring and living facilities (particularly within the available living labs). One fruitful area for research is Body area networks which compose activity sensors and medical sensors attached temporarily or permanently to the human body (and in the mobile devices that they carry).

- *Patient Monitoring*. To reduce healthcare costs and improve patient health it is typically preferable for patients to be cared for at home. The remote observation and monitoring of patients (often within assisted living facilities) is a key enabler in the better support of an aging population. Indeed research into quickly and cheaply converting homes is one important direction. Systems reporting the effect of current lifestyle changes to the user; detection/prediction of medical emergency (e.g. early detection of stroke); and integration with emergency response systems all offer challenging research applications that demonstrate the potential of the Future Internet.
- *Elite Sports*: Performance monitoring and training systems that form integrated on-body and external sensors (e.g. cameras and sensor embedded in the environment), e.g. the UK EPSRC SeSaME project[41] employing body sensors to improve sprinter performance.

2.6.4 Future Networks and the Infrastructure of the Future Internet

In reviewing the views of the future technologies that will underpin the next generation of Internet services and applications, we have looked at the keynotes delivered at the major research conferences over the past year or two, considering SIGCOM, INFOCOM, ICC, Mobicom and more specialized meetings, such as HotNets. The views presented in these fora split strongly between projections of far future technologies offering continued increases in bandwidth and computing power over many decades, and a more present-centered emphasis on radically reworking networking software, how it is defined and how it is created. New materials such as graphene may break down the barriers to unlimited scaling that seem to put a horizon on silicon-based electronics allowing computing speeds memory densities and storage capacities to scale to new ultimate limits such as storing one bit in every few atoms of a substrate. An aggressive view is that the era of Silicon will have ended around 2020, to be supplanted by biological or artificial membrane materials. We take a more conservative approach, assuming that steadily improving communications technologies will be phased in alongside the existing backbone of high bandwidth pipes and the enhanced computing power

---

[41] http://www.sesame.ucl.ac.uk

delivered by multicore single CPUs in racks of data center computers and in ever more powerful portable devices will be absorbed into our testbeds without fundamental changes. A bigger technology enhancement occurs at the edge of the internet, as wireless bandwidths are undergoing equally or more rapid increase, though new protocols (e.g. 5G) and the steady progress of cognitive radio, which finds ways to adaptively exploit all the available spectrum. A reasonable expectation is that Gbps bandwidths will be available to wirelessly connected devices by the end of the decade, and that test environments that allow exploring the potential of such capacity with real users and in realistic environments will be of considerable value.

The argument for focusing on radical change in networking software is founded in business as much as in advancing technology. It assumes that the coming decade will see the introduction of cheaper, faster, and commodity hardware for routing, switching and buffering to replace the highly complex and proprietary networking gear in use today. This provides an unparalleled opportunity for standard and open source software to come into widespread use. OpenFlow, which already is exploited in multiple FIRE testbeds, is an open source interface to the routing portions of networking. It has gained rapid acceptance, but addresses only a fraction of the functions that a true general purpose networking operating system must support. Also, the increase in processing power that rapidly evolving standard hardware provides to a router brings the power of BigData into the realm of networking. The databases that routing hardware can make use of are now global in scope, and richer in functional knowledge than was previously thought possible. Algorithms and environments in which these advantages can be realized must be sought, and tested in increasingly real application environments.

An interesting question for the development of future FIRE activities in software defined networking is whether the best approach would be to proceed, as has happened with OpenFlow, with one family of innovative ideas at each step, determining their strengths and weaknesses, and developing a domain of applicability, or to take a clean slate, and create a complete Networking OS in a single broad gauge project or projects. The natural desire to see results adopted by a very conservative and operationally-focused industry has kept the scope of OpenFlow's efforts tightly focused. This may change as rapid adoption of the facilities offered to date (e.g. in Google's internal data communications deployments) makes researchers and developers much more ambitious.

In the past, FIRE's infrastructure efforts, such as those now incorporated into Fed4FIRE, have largely focused on technology development, standardization, and interoperability, leaving little to no overlap with the user-visible testbeds such as SmartSantander and Experimedia. This can change going forward, as a major advantage of flexible networking software is that the new interfaces permit the network to be aware of the applications and the types of flows that it carries. Similarly, the new interfaces permit the applications to become aware of the capacity and contention levels of the network they operate on and exert control over that network to an unprecedented extent. This means that testbeds that bridge the gap between infrastructure and applications are a real possibility in the coming years and should be encouraged. Visionaries in this new field like to say that the best problems that can now be tackled are not just interesting technical challenges but can improve the practices of networking in dramatic ways. Naturally such changes are not easily embraced by industry, although their attraction to start-ups and SMEs should be apparent. Another way of envisioning the power of research in methods of managing networks under open source software is to compare with the software industry or with consumer computing hardware. There is a $10B industry creating tools enables the creating of $100s of billions of hardware or of software and applications. This powerful multiplier makes investment in facilitating the tools and practices of software defined networking a very attractive opportunity.

2.7 What might happen in the Future?

2.7.1 Important Future Research Trends

The analysis of the individual research areas of the Future Internet highlights the growing need for integration, e.g. the integration of multi-purpose, multi-application wireless sensor networks with large scale data-processing, analysis, modelling and visualisation, along with the integration of the next generation of human computer interaction methods: Near Field Communication (NFC), public displays, augmented reality etc. Such scenarios highlight the true potential and requirements for the software and services that will compose systems in the Future Internet.

It is clear that applications will embrace the combined four pillars: things, people, content and services to create truly pervasive, complex systems. This mirrors a significant software engineering initiative in the US[42] focusing on ultra large scale systems.

Crucially, this initiative identifies the similar key themes of research (indeed the biggest challenges in these types of systems are well established by domain experts) and therefore they must underpin the requirements of any future experimental facilities to validate research in these areas:

- *Scalability*. Validating scalability claims requires infrastructure with the potential to increase a system's size, workload, users, or data. Hence, substantial resources must be made available to support realistic research. A claim a system is scalable to 1K nodes is insufficient where comparative industry systems scale to 100K nodes.
- *Interoperability*. The integration of highly heterogeneous technologies across a number of hardware, software, data, and networking dimensions requires new approaches to the interoperability problem. Open standards offer a starting point, but it is unrealistic to expect standards to cover the entire Future Internet.
- *Software Engineering - Taming complexity*. Novel software engineering methodologies and programming abstractions are required (including autonomic ones) to support developers build the next generation of large-scale systems.
- *Security, Privacy, and Trust*. The Future Internet brings new requirements and challenges with respect to the security of systems, and the privacy of users.

There remains a number of interesting research areas in the Future Internet that are not covered by the research challenges and experimental facilities today:

- *Addressing the Digital Divide*. Considering smart solutions for rural areas, technology light society, developing nations. Smart villages, smart agriculture, and water management are among some of the many Future Internet research themes to be addressed to assure that the Future Internet is truly global, and supports the requirements of anyone, anywhere.
- *Crowd Sourcing and Participatory Sensing*. Humans themselves are an important resource for the Future Internet. They perform both intelligent sensing and processing of data, e.g.

[42] L. Northrop, "Ultra-Large-Scale Systems: The Software Challenge of the Future", ISBN 0-9786956-0-7 http://www.sei.cmu.edu/uls/

in the detection of events, capturing information with mobile devices, building communities of users, offering intelligent processing through crowdsourcing methods. Hence, new software and services that realise these capabilities as a central element of systems are required [28].

2.7.2 Driving forces and uncertainties

A first step towards constructing future scenarios of FIRE is to describe the scope. FIRE is a broad topic incorporating diverse methodologies, technologies and communities of practice. Understanding the dimensions of FIRE and how each dimension has the potential to move FIRE towards different futures is a useful tool. A first step is to identify broad areas of change as driver of FIRE's future, and identifying the uncertainties within these areas (Table 3).

Driver Category	Uncertainties
Customer profiles and demand	How will users of FIRE facilities change? (more SMEs, network communities, service communities) What is the expected profile of FIRE users? (Researchers, Engineers, Entrepreneurs, Service Providers, Content providers, Broadcasters, Network Operators) How will users access facilities? (Open access, Open calls, Embedded in projects, linked projects) How will users research, engineer and innovate in the Future? (Open and closed innovation, Scientific methods, Software engineering) How open can facilities be considering the requirements of different customers? What technical research topics are important to users of FIRE?
Facilities infrastructure, tools and services	What facilities will exist at the start of 2015? How will the facilities be structured and organised? Will facilities support a diversity of stakeholder communities or not? Will FIRE facilities fragment or converge? Will FIRE testbed federation further evolve or will islands of facilities result What facilities will continue to remain relevant in 2014 and which ones will not and need to be terminated? Will FIRE testbed capacity and capabilities grow or shrink in H2020? Will there be sufficient expertise/skills to support future needs? What facility resources and support activities are needed to support FIRE users? What's the balance between use of resources and innovation in resources? (e.g. applications using SDN or researchers innovating in SDN) In how far will experiment lifecycle management tools be available widely
Competition	What will competitors do? Are there competitors? Is there really a market beyond academic research? Will researchers continue to build their own testbeds and if so, what are their reasons? What are the barriers to outsourcing?
Financial structure	Will dominance and dependency of EU funding continue or will hybrid models develop? Will adequate instruments evolve at EU level to support efficient distribution of public funding between facilities and their users?

Table 3: Driver categories and uncertainties

An analysis of these drivers and uncertainties within the context of scenarios allows us to explore how each has the potential to move FIRE towards different futures. We want to understand all the uncertain forces and their relationships with each other whilst identifying the most important to the issue of the future of the FIRE programme. As Table 4 demonstrates, often the debate moves between two poles as different stakeholder interests are championed.

Uncertainty	Description
Upstream vs Downstream	New technologies follow a path from research through to maturity in industrial products and services sold in the market. Facilities that support upstream development tend to be lab-based supporting the controlled investigation of technical concepts. As the technology moves downstream towards the real world and integration with wider systems (including economic systems) a greater consideration of socio-economic barriers to adoption is necessary (e.g. legal, operational, scalability, etc). Upstream problems tend to be Deep/Scientific in nature whilst downstream tend to be Fast/Engineering in nature.
Science vs Engineering	The scientific method is a body of techniques to investigate phenomena and acquire new knowledge using empirical and measurable evidence. Engineering is concerned with the application of knowledge to design, build and maintain solutions to life's problems. The primary and related disciplines are Computer Science and Software Engineering. Computer Science is typically hypothesis driven whereas software engineering will aim to verify and validate a system against a set of requirements. Both approaches require "tests" to establish evidence of system characteristics
Academic vs Industry	Academic institutions are driven by the desire to acquire new knowledge and educate populations. Industry is driven by the desire to make a profit. The differences in drivers/performance targets influence organisation behaviour and appetite for risk.
Infrastructure vs User	The Internet is a complex system made up of interacting systems and stakeholders operating in a market of products and services. Facilities exist to support different stakeholder communities and no single facility exists to support everyone (although FED4FIRE is attempting to deliver a broad facility). Infrastructure facilities primarily support networking researchers, Cloud facilities support services researchers and there are even domain specific facilities support areas such as Networked Media. Networking research is a certain type of computer science undertaken in a very specific socio-economic context, operating over slow time scales with large socio-economic impact. Services research by contrast has a broad socio-economic context (e.g. wide range of applications) with the expectation that results can be transferred into new products and services quickly, and where they are developed with participation from users.
Specialisation vs Diversity	The Future Internet is a fairly nebulous concept which is meaningful within the European Research domain. In 2008, the Future Internet meant the convergence of networks, services, media, things and security. In H2020 the definition is still under some discussion but the restructuring of the EC units gives us some indication of the scope. FIRE grew to support the 2008 FI definition by developing testbeds for networks (OpenLab, CREW, OFELIA), services (TEFIS, BonFIRE), things (SmartSantander) and media (EXPERIMEDIA) with FED4FIRE looking to provide coherence between them through federation technologies. In fact, throughout FP7 FIRE was

	always looking for new areas (e.g. incorporation of users in EXPERIMEDIA, FIREBALL, etc) with a certain amount of consolidation (e.g. Panlab+OneLab->OpenLab). However, the level of diversity in the portfolio makes the FIRE offer difficult to communicate and is a challenge for sustainability (e.g. scale, up and down, cannot be achieved by maintaining a set of specialised testbeds).
Best effort vs Guaranteed QoS	Access to resources is dependent on the culture of the communities of practice and operations models of testbeds. Many testbeds and users are happy with best effort either because the usage is so small that access can be negotiated between friends or because the service is free at the point of use and users accept a lower quality of service. In contrast there are some users that demand guaranteed QoS either because they are industrial and expect it or the test to be conducted is only viable if extraneous variables can be minimised. QoS is an emotive word and what this really comes down to is how providers share resources between multiple users. There are different strategies and associated costs.
Large vs Small Scale	FIRE has been built on the promise of large scale facilities but large is a relative term that's rarely quantified. For example, many testbeds in FED4FIRE for wireless networking are specialised small scale facilities in specific buildings and even BonFIRE that offers 5million core hrs and 440000 Storage GB months per year is not considered large scale in the cloud domain. Relative scale must be defined and related to use. For example, BonFIRE does offer larger scale to what EC projects can achieve but is a drop in the ocean in comparison to Amazon.
EC vs Commercial funding	FIRE is funded by EC resources. For the future, and based on new service concepts and demands of experimenter communities and other users, hybrid business models could be foreseen.
Open vs Closed	Research, experimentation and to some degree innovation can be conducted in an open or closed environment. Traditionally, commercial companies, especially SME's, have a preference for operating in a closed environment driven by the desire to protect intellectual property. However, increasingly Internet products and services require community activation to be successful which has driven the need for early user participation through open platforms and open beta programmes. FIRE must consider how different facility architecture, operations and business models support creation, exploitation and protection of intellectual property for different customer groups.
Fragmentation vs Cohesion	Fragmentation is a natural consequence of the EC work programme where the next project needs to differentiate itself from prior work through advances in SOTA and not continuation of an existing service. The instruments added to the work programme (e.g. collaboration funding, open calls, STREP alignment with facilities) are put in place to redress the balance and achieve some cohesion between an independent set of projects.

Table 4: Dimensions of Uncertainty

The next chapter further explores the different scenarios that are spanning the range of possible FIRE futures.

3. Towards Future FIRE Scenarios

3.1 Framing future FIRE scenarios

The uncertainties in Table 3 offer several possible axes to explore the future of FIRE. We have selected two main axes to explore four different logical futures (Fig. 6).

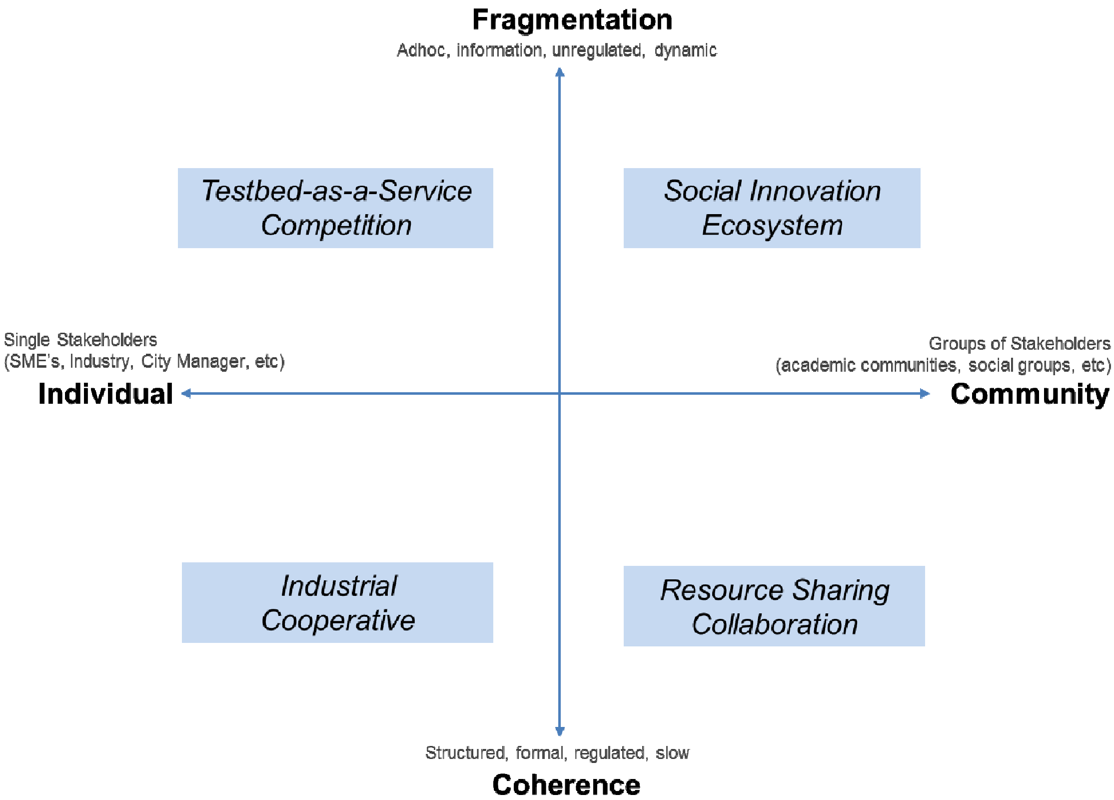

Figure 6: Framing future FIRE scenarios

- **Individual vs Community**: how will researchers collaborate in the research and development of products and services? Individual means people develop alone as a single stakeholder (e.g. SME, industry, city manager). Closed innovation is typical of this space. Dynamic services and market players are typical. Community means that people collaborate opening to achieve a goal where outcomes could be more altruistic rather than financially motivated.
- **Fragmentation vs Cohesion**: how will collaboration be structured and governed? Fragmentation means that structures are adhoc and largely unregulated (e.g. social organisations, informal communication, open markets. Cohesion means that structures are organised and in some cases regulated (e.g. virtual organisations, process-oriented, etc)

3.2 Testbed-as-a-Service Competition

A first possible future scenario is that FIRE will consist of a set of test-beds that provide their facilities as a pay-per-use service. Testbeds may be diverse, highly specialised and isolated, i.e., they offer unique value in the market (e.g. sensing data from a specific location). To a large-extent such a testbed-as-a-service offering is unregulated; new facilities are free to enter the market and offer competing services. However, these testbeds may also form loosely-coupled federations of heterogeneous facilities (or supply chains), where once composed value is gained by the participating members, e.g., to leverage a wider reaching advertising platform. In such a scenario, competition remains a key property; where organisations are in competition with one another they will therefore not trust one another—this will largely drive the type of experimental facilities generally available. Consider two competing organisations, one organisation is unlikely to carry out experimental research utilising the resources and the facilities of another organisation for fear of losing intellectual property and/or competitive advantages.

The business models require testbeds to be operated commercially and for a profit. The strong commercial drivers force testbed operators to focus on financial performance which reduces their appetite for engaging in more risky and less profitable activities. Testbeds become customer focused but mainly customers with the ability to pay for services which excludes some potential users (e.g. academics without a budget). There's a greater emphasis on accountability to customers and delivering Quality of Service, with tradeoffs being made between customers competing for resources. The disclosure of operational decisions is reduced and often hidden from the customer as the principle of information hiding is necessary to retain competitive advantage. The restrictions on observability and control reduce transparency and opportunity for innovation as experiments are now restricted by existing business models. The complexity of accountability and responsibility in supply chains of connected services reduces opportunity for federation. The shift to pay-per-use changes the regulatory environment and requires testbeds to consider relevant law associated with B2C and B2B transactions. As public funding cannot support commercial service operations, the proportion of activities supported by public money reduces significantly to only those required for research and development of advanced features.

3.3 Industrial Cooperative

A second scenario is for FIRE to become a resource where experimental infrastructures (testbeds) and Future Internet services are provided by co-operating commercial and non-commercial stakeholders. The testbed facilities themselves converge to a single large-scale federated facility, based upon well-established standards and common platforms. FIRE is then able to support commercial R&D into new Future Internet technologies that require large-scale trials upon the latest communication infrastructures. Further, the FIRE facility itself also offers commercial opportunities via the provision of new services within the testbed (e.g. data hosting, data processing, communication brokering)—SMEs and start-ups will develop new software that can be installed and deployed across the infrastructure.

The FIRE facility will grow through the addition of new infrastructures (datacentres, sensor networks, and software defined networks) to the federation. Each facility operates through a pay-per-use model in order that value is gained from joining the federation (whether to maximise profit or simply cover operational costs). Infrastructures will be heterogeneous offering a broad range of services —that is, although there is a conformance to standards and software this does not mean that the facilities will be homogenous cloud computing facilities, and FIRE will significantly remain a cutting edge offering.

FIRE will have minimal management overhead—the operation of the federation will require monitoring of continued conformance of stakeholders, there is need to support new joiners and leavers, and support for the day to day running. To ensure fairness, such managements will be carried out by a publicly-funded (or non-profit) organisation.

What are the general objectives?	• Commercially exploitable R&D results in Future Internet applications • Commercial opportunities for infrastructure providers and SME developers of novel Future Internet Services
What does the FIRE programme look like? [Facilities infrastructure, tools and services]	• Towards a convergence of heterogeneous testbeds: Things, Services, Network, Living Labs into a single federation • Flexible federation of co-operating and competing stakeholders • Public and private funded testbed-as-a-service
Who will use FIRE? [Customer profiles and demand]	• SMEs and commercial elements who require minimal costs for elastic testbed resources • Infrastructure owners wishing to add heterogeneous testbed facilities to a global federation for commercial opportunities
What are the research areas?	• Novel business models • Applied domain specific research in the use of Future Internet technologies in specialised application areas such as healthcare, transport • New Internet Technologies, Services and Things

Table 5 Main issues in the Industrial Cooperative scenario

The FI-PPP as an example

The FI-PPP programme offers an illustration of this FIRE vision, albeit at a smaller scale than is envisioned by this FIRE scenario. The FI-PPP seeks "to increase the effectiveness of business processes and of the operation of infrastructures supporting applications in sectors such as transport, health, or energy; to derive possible innovative business models in these sectors, strengthening the competitive position of European industry in domains like telecommunication, mobile devices, software and service industries, content providers and media."

FI-PPP is underpinned by a common software platform called FI-WARE. This consists of a set of Generic Enablers (GEs) which are reusable software services that are deployed in the testbed infrastructures. Generic enablers provide APIs for a wide range of Future Internet services, e.g.: data processing, event handling, and data storage (i.e. those typically provisioned by testbed facilities). GEs must conform to agreed upon API specifications–hence both commercial and open source implementations can sit side-by-side. For example, context broker implementations must conform to the Open Mobile Alliance's NSGI (Next Generation Service Interface) specification.

FI-WARE is installed across infrastructures in the FI-PPP testbed federation. The XIFI project is creating a Pan-European Future Internet facility. Initially, this connects infrastructure hosted at five locations into a federation. XIFI is also developing the federation and business models to allow new facilities to add their infrastructure to the federation both during and after the lifetime of the FI-PPP programme.

The key users of the FI-PPP are an initial set of large-scale, use case trials performing research in diverse application areas: environmental science (ENVIROFI), city safety (SAFE-CITY), smart agriculture (Smart Agri-Food), manufacturing (FITMAN), and smart energy

(FINSENY). These were funded through the first two phases of calls in the FI-PPP. Subsequently, the third and final call will focus on SMEs and entrepreneurs as users of the FI-PPP facilities to add new value to the infrastructure alongside the large-scale trials.

Key points of the scenario
The scenario highlights some key points with respect to the requirements of future infrastructure:

- There will be a convergence of heterogeneous infrastructures through a central pan-European federation.
- Users of FIRE will largely seek commercial goals.

Hence, there is a need for a flexible trust model for the federation in order to allow complimentary commercial organisations (who may have differing levels of trust with one another) to co-operate. Generally such organisations will only participate in a federation if they can get something from the federation where they cannot really compete. For example, telecommunication companies that have their own national networks work together to support international end-to-end connectivity. They are happy to do this as they cannot compete in others spaces: British Telecom cannot create a phone call to Turkey without such peering relationships. However, these companies do compete on other things, and hence have limited trust models that would stop them from participating in a centrally managed, open resource sharing federation (i.e. where anyone else in the federation can leverage their resources).

3.4 Social Innovation Ecosystem

FIRE becomes a collection of heterogeneous, dynamic, and flexible resources; these offer a broad range of facilities, e.g. service-based infrastructures, network infrastructure, smart city testbeds, through to user centred living labs, and highly application-specific sensor-instrumented services. These are open to use and potentially composed with one another. This is unregulated and without uniform APIs but APIs developed to meet the demands of specific user communities. Hence, FIRE itself follows an Internet of Services vision of testbeds within a loosely coupled Service-Oriented Architecture.

FIRE then becomes a driver for innovations for society. Social computing investigates technologies that build communities, often with the goal of achieving beneficial societal impacts. For example, local communities can create initiatives to improve their environment and energy usage. The notable property of experiments in this field is that communities take the central role. In many cases they are the end users, and evaluators of the technologies (i.e. they inform the validation of experiments); they will also be part of the facilities themselves, e.g. in the performance of participatory sensing and crowd sourcing activities. However, and perhaps the most novel involvement, people will be the creators, i.e. creating new applications and services within the FIRE ecosystem in order to validate experiments with new software tools. That is, to engineer new social computing platforms and service that support open, community driven innovation with the technologies and facilities available (often with minimal expertise).

Therefore, Future FIRE ecosystems will make available the open software, technology and tools that will minimise the gap between these users' expertise and the skill sets required to manipulate the current state-of-the-art in Future Internet experimental testbeds. Simply put, make it as easy as possible for people to innovate and validate new technologies.

Openness is a key property here. They must be able to experiment vertically as well as horizontally, i.e. Experimental infrastructure must not offer them a restricted and closed business model. This will allow innovation to be carried out in the infrastructure level (e.g. new protocols, new methods to transfer large amounts of data, new ways to leverage Software Defined Networking) and also at the software and service level, i.e. innovation in new systems built on top of the infrastructure—in the same way new apps are developed for Smart Phones.

What are the general objectives?	• Quest for knowledge advancement – solving tomorrow's grand challenges • Altruistic innovations for the benefit of wider society • Technological innovation in Future Internet {service, things, media, people}
What does the FIRE programme look like? [Facilities infrastructure, tools and services]	• Towards fragmentation of heterogeneous testbeds: Things, Services, Network, Living Labs • Very loose federations and virtual organisation • Open, shareable, heterogeneous data sets
Who will use FIRE? [Customer profiles and demand]	• Everyone and anyone • Open Source Developers • Community Groups • Charity/Non-profit organisations • Academic communities • Schools/Educational Institutions
What are the research areas?	• User centred software engineering • New Services/Things • Social Networking services • User Innovation

Table 6 Main points in the Social Innovation Ecosystem scenario

Illustrative Use Cases in this FIRE Ecosystem

The Raspberry Pi is an example where community-driven innovation has grown around a single idea to be used in unexpected ways. The original purpose of the Pi was to provide a low-cost device that would inspire children to learn to programme in the same way that people were inspired by the BBC Micro in the 1970s and 80s. However, rather than just innovating in the educational arena, the Pi is now the focus of new research and products. As a small, cheap platform it is hardly innovative (there have been many similar boards); however, driven by social media and a growing dedicated community it has become the focus of wider innovation. Hence, it is as important to build a community of innovation. The most advanced experimental testbeds are useless without people motivated to exploit them.

In this context, the already mentioned CitySDK[43] project within the CIP ICT PSP is a particularly interesting initiative. It seeks to provide an open, lightweight software

[43] http://www.citysdk.eu/

development kit. This will then be made available to be used to rapidly engineer new installations within the Smart City testbeds. While the SDK is not currently publically available to evaluate the extent to which it is general purpose, it highlights software engineering at the forefront of experimentation in this domain. There are challenges in how developers will leverage the facilities, and also how they can create tools that real users and developers will build the next generation of smart city systems.

Community driven innovation is not a new concept; hackathons are now a common feature of the technology industry where motivated individuals collaborate intensively over a short period of time to: experiment with new technologies and advance products, drive innovation for pitching new technologies as start-ups, and address societal concerns and problems. Communities with a particular cause are often particularly motivated to solve problems:

- A local community working together to solve a local problem.
- Solve global problems such as water management—the Water Hackathon is a multi-city global event to inspire new technological solutions for water irrigation and management.
- Disaster response e.g. Random Hacks of Kindness and Hackathons that spring out in response to currently occurring disasters.
- Hence, testbeds with real developers and/or community inspired hacktivists to validate new engineering methodologies are of equal importance to networking infrastructure.

Air pollution monitoring measures the quality of air. Typically cities perform these measurements using static monitoring stations. These produce coarse-grained measures, and hence pervasively deployed wireless sensor networks have the potential to offer improved fine-grained results. There is significant motivation to reduce pollution: the direct impact on health (and the associated correlation with increasing healthcare costs), the environmental impact, and the direct financial penalties (e.g. EU levied fines).

Wireless sensors can be deployed across a city using instrumented sensors (e.g. as within lamp posts in Santander), vehicle sensors (e.g. attached to government vehicles in Santander), and participatory sensors where users carry sensors attached to smartphones (as shown in Figure 7).

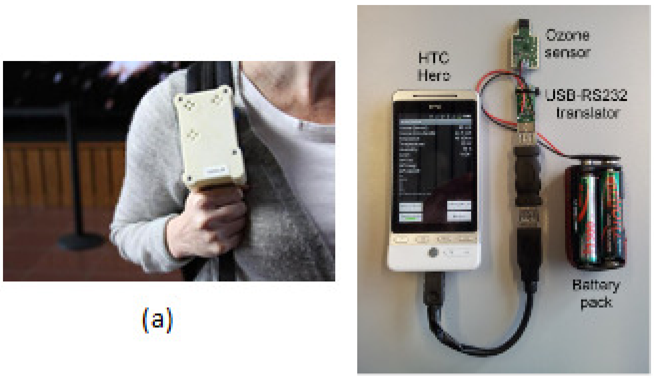

(b)

Figure 7: (a) CitiSense monitor developed at UC. Davis. (b) Smartphone integrated air monitor developed at ETH Zurich with Microsoft

Air pollution is caused by: combustion engine exhausts; factories, offices, homes, and buildings burning fossil fuels; high voltage power lines; pesticides; radioactive fallout; garbage and sewage. Hence, there is a need to integrate further data streams, i.e., traffic monitoring systems, geospatial data, historical pollution data, mapping data etc. User tracking data about travelling and commuting (e.g. using RFID and smart card technologies) may also serve as input.

Smart buildings may be equipped with air monitoring sensors (CO_2 being a particularly important measure)—these have the potential to be federated within a pollution scenario; that is, they may or may not integrate data streams but must support user applications that are reporting local information about air quality (for example, eHealthcare applications for people with respiratory medical conditions).

The analysis of pollution data, execution of modellers and productions of visualisation systems are a resource intensive task. Hence, in the face of rapidly increasing amounts of data from multiple sources, and the requirements for real-time results, there is a significant need to perform these computations using highly scalable resources, e.g., as provided by cloud computing facilities. Measurement and analysis results can be output to different stakeholders:

- Local government to manage levels of air pollution, and observe the impact of pollution saving measures. These results will generally be in (potentially real-time) visualisations and reports, and warning systems of hotspot areas.
- General population users. Smartphone apps, web systems (social networking sites), public information displays (e.g. Augmented Reality systems) all offer potential platforms to both inform the general population about pollution (so they can take precautions), and also inspire and incentivise them to make a difference.

Finally, such systems need not be constrained to local geographic regions; national and international systems monitor broader views of air pollution issues that potentially feed into global environmental and climate monitoring and modelling systems.

Applications

The service and data infrastructure could be leveraged in different ways (beyond the basic monitors and visualisers—although there is scope for different visualisations to be developed):

- Integrated smart city. The most obvious application within a smart city is an integrated traffic system where detected pollution hotspots could trigger changes to traffic systems.
- Healthcare apps. Different smartphone healthcare applications could be developed e.g. a general air pollution monitor. A monitoring system for people with respiratory problems e.g. warning them that their current activity (jogging, cycling, etc.) isn't safe in the current conditions. Building warning system, i.e. the app leverages CO_2 monitoring and warning in smart buildings to be informed of danger.
- Military systems. The infrastructure and data could be used by military applications to monitor for chemical attacks and illegal activities
- Optimised bike route applications are an example of an application that can be built upon real time monitoring data. A bike route can be suggested to the user that does not involve cycling through areas with poor air quality. Potentially new interaction methods to inform real-time changes can be implemented e.g. AR goggles.
- Personal pollution footprint. Applications monitoring personal pollution/environment footprint (could be integrated with other monitoring systems – energy usage). Reductions in footprint could be rewarded with incentives e.g. targeted smart card incentives for use of bicycle, walking, and public transport as opposed to costlier measures.

- Pollution reducing community initiatives. Public displays of areas of pollution e.g. augmented reality view of hotspots. Public displays of neighbourhood reductions. Integration of footprint applications with social networking activities.

Key Points of the scenario
The scenario highlights a number of key points with respect to the requirements of future infrastructure:

- Increasing scale—the scenario may cover a small geographic region with a small number of users. However, it may also scale to a global system with millions of inputs.
- Highly heterogeneous—the technology, infrastructure and software are diverse. Covers a multitude of devices, networking technologies, data, service technologies etc.

The scenario facilities can be leveraged in multiple ways for diverse application types. Within a single application domain there remains room for community and business driven innovation. New services can built upon the infrastructure for local community projects, and or products can be developed and delivered within the facilities.

3.5 Resource Sharing Collaboration

Given the previous requirements for research into large-scale software and service development it is clear that federated infrastructures are required to provide the next generation of testbeds. Single facilities may be able to perform domain specific experiments, but it is their combination that will demonstrate the potential of the Future Internet. For example, within a smart city there may be a city wide test-bed that can be integrated with local testbeds such as: individual smart buildings, WSN instrumented rivers, WSN instrumented road tunnels, instrumented museum, art, or tourist attractions, and public transport testbeds. An interesting research question remains the sustainability and governance of such infrastructures: how do new facilities integrate into the architecture, and how are they then managed.

Fed4FIRE[44] is developing a federation of testbeds, i.e., standardising their integration such that experimenters have access to the entire set of facilities.. The integration of cloud services to model and analyse wireless sensor network data is a clear direction but it remains an open question as to what a truly federated infrastructure should provide.

At the opposite end of this spectrum you have academic collaborations where there is a common goal that is reached through collaboration. The Grid is an important example of resource based sharing; here the members typically trust each other a great deal, and they provide a proportion of resources to the federation, whose use and purpose is then controlled by the federation itself. This is a very open trust model that is open to misuse, however the community driven model weighs this risk against the gains of collaboration.

The OFERTIE[45] project is an example of vertical federation between applications, services and software defined networks. This may seem like a good example of where federation is needed but most of the technical investigations could be done using a single OFELIA island

[44] http://www.fed4fire.eu/
[45] http://www.ofertie.org/

Networking research is done with specific socio-economic constraints. The socio-economic impact of changes to the core network is large. Therefore networks are highly regulated and adoption is a different and lengthy process. This contrasts service innovation where ideas are expected to be monetized quickly.

Illustrative Use Case: Anytime-anywhere device connectivity becoming a true reality
The idea of anytime-anywhere Internet connectivity with any device is far from new. However, with the rising success of 3G and 4G networks and an increasing number of residential internet providers joining the FON initiative (over 7 million Wi-Fi hotspots) or starting similar initiatives of their own the idea is quickly becoming a reality. At the same time there is a renewed interest in personal computer terminals or small form factor computing devices such as smartphones that rely on remote systems and services located in the cloud for computing and storage. Additionally, these personal devices could get information from other neighbouring devices and sensors, or use other devices such as TVs or video walls in their environment, for example to display content. When such scenarios become everyday reality the consequences are:

- More devices (increased scale), more heterogeneity, potential security issues.
- More devices communicating wirelessly means that the wireless medium will be even more loaded than it is today. 3G, 4G systems and their follow-up technologies can offload part of the users to Wi-Fi based systems. Still, this will result in an increased usage of both licensed and unlicensed spectrum – resources that are both very limited. In case co-located devices are interfering with each other, communications disruptions are very likely, leading to end-user user-experience problems or worse.
- Services will consist out of several components, running either locally on a device itself, on a neighbouring device or in the cloud. The distribution of these components will be changing dynamically based on multiple criteria such as the amount of calculation power or storage needed the delay requirements of a service, the available wireless connections and technologies, or characteristics of the device such as remaining battery power.

Dynamic service distribution across many diverse platforms, is clearly an interesting research challenge. To avoid the problem indicated under the first bullet above, new or optimized wireless technologies and protocols based on cognitive radio solutions may be the answer.

Cognitive radios and cognitive networks. The interference problem is a very real risk, and can already be observed today. For example, in office WLANS with many connected devices (typically laptops, smartphones) the unlicensed ISM bands may at times already be fully occupied, resulting in connectivity problems. In the (near) future, these problems are expected to worsen. Cognitive radios are radios that adjust their Tx/Rx characteristics based on the characteristics of the environment in which they are operating. There are two main ways to collect characteristics on the environment:

- Sensing approach: the cognitive devices scan their environment to learn about the use of the spectrum.
- Database approach: based on measurements and/or propagation models, a database is built which holds information on the RF use (and license holders) operating in a particular location.

If one is able to perfectly and instantly detect or get the RF characteristics of an environment, it is possible to use the spectrum in a more efficient way. This can be done in two ways:

- Horizontal resource sharing: the available spectrum is shared optimally between the devices working in the same RF band and having the same priority (horizontal resource sharing).
- Vertical resource sharing: when considering cooperation in licensed bands: "licensed" bands for which one is not the license holder could be used under the condition that the rights of the license holder are not damaged (i.e. if the license holder is using or wants to use the considered band, no device should be operating in that band at the same time).

It is clear that especially vertical resource sharing does not only imply technical challenges but also regulatory challenges. Furthermore, before a license holder ("primary user") even wants to think about accepting "secondary users" on 'his' frequencies, it is clear that he will want to be assured of the fact that "his" frequency will be available whenever he needs it.
If the sensing / adjustment loop can be made a reality, the spectrum efficiency could be considerably improved, leading to more stable wireless connectivity and better user experience. Furthermore, improving the reliability of wireless connectivity also means that it would be possible to use wireless connectivity for more critical applications (e.g. public safety, healthcare, etc.)

Applications
If the anytime/anywhere/every device connectivity and distributed services are realized, several applications are imaginable:

- Widespread reliable vehicle-to-vehicle and vehicle-to-infrastructure communication could help to increase safety and reduce traffic jams
- Faster wireless connectivity –also when on the move- makes it possible for bandwidth-demanding applications to be supported.
- Wireless body area networks could become mainstream and prove their usefulness in health risk prevention
- Reliable wireless cable replacement can be applied in many environments such as in a smart home, in non-critical vehicles or airplane systems (e.g. onboard entertainment). In factory environments this could lead to increased productivity or cheaper installation costs
- Wireless systems are enablers for smart cities
- Device to device communication can be used to learn more about ones' environment and access a wide range of network-accessible services
- End-user devices may become smaller and cheaper if they can reliably connect to services in the cloud or in the environment.

3.6 The FIRE Programme, Uncertainties and Forces

Framing uncertainties along the two axis allows us to explore the impact of current forces on the direction of the FIRE programme. Figure 8 shows the effect of the EC work programme on the positioning of FIRE within our scenarios.

FIRE is currently a fragmented set of facility projects each covering a different technical domain. Periodically the EC launches a new work programme that asks for new facility building projects. Successful proposals must address the call and in doing so must demonstrate that they advance the state-of-the-art. The requirement for uniqueness and advances from what exists today drives fragmentation of the projects because EC ICT research funding cannot be seen to fund follow-ons of existing projects. Of course, facilities and intellectual property does flow between projects in varying degrees but this is achieved by common partners rather than a continuation of a collaborative testbed services.

In addition, the EC's contractual framework considers each project as an independent contract with the EC with minimal programme level governance to maintain coherence between the set of projects funded in a specific area. There are refinements to the instruments (e.g. collaboration budgets, alignment of STREPS, open calls, etc) that the EC can use to increase collaboration and coherence but these tend to be weak in comparison to the drive for fragmentation caused by project uniqueness and independence.

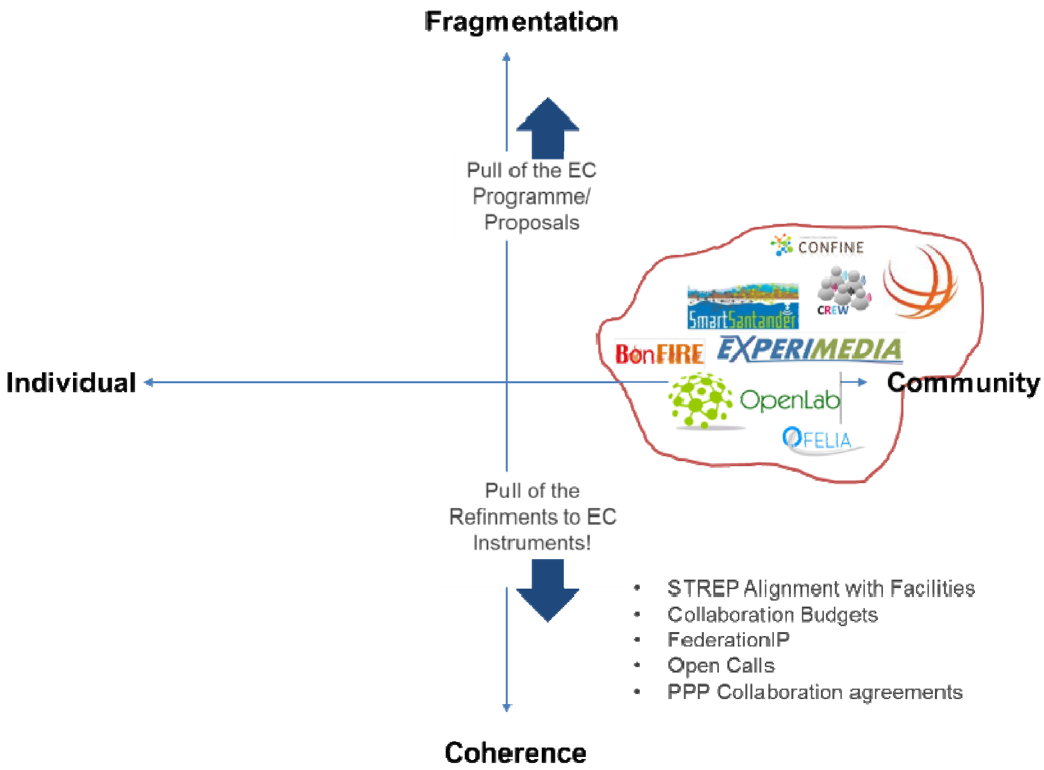

Figure 8: The effect of the EC work programme on FIRE

Figure 9 highlights two important projects (FED4FIRE and XIFI) that are expected to influence the direction of the FIRE programme. FED4FIRE's goal to create a high level federation framework is driving coherence in technology, operations and governance across many of the FIRE facilities. As such FIRE could move from a fragmented set of facility projects to a set of interoperable facilities that can be connected using federation primitives. Early results suggest that FED4FIRE will aim to deliver a federation that's aligned with Collaborative Resource Sharing rather than a loosely coupled set of independently operated testbeds. However, the final structure that is established will depend on the business models of facilities and the demand for federation from customer usage patterns.

There has been significant discussion about the relationship between FIRE and the FI-PPP programme. A key project in the FI-PPP programme related to FIRE is the Capacity Building IP XIFI. As discussed earlier the FI-PPP is considered an industrially-driven cooperative aiming to deliver significant economic growth through the adoption of a core Future Internet middleware FI-WARE in societally important application sectors. XIFI must establish the economic conditions for testbeds and other infrastructures to be used for business experiments as a precursor to production services within the market. This drive towards exploitation and production services is in contrast to what many FIRE users and facilities are focused on.

Although FIRE does support commercial companies in the use of testbeds, the largest proportion of users is from universities and research centres whose primary success measure are publications rather new products and services in the market. The consequence is that most FIRE tests support technical rather than business experiment objectives. It's not black and white just that most users come from that segment and are interested in technical performance rather than commercial viability or legal compliance. In contrast Phase 3 of the FI-PPP expects a large engagement with SME's who will use European infrastructure and FI-Ware to design and deliver new products and services to usage areas.

The difference in objectives and success measures raises an interesting challenge for FIRE facilities within the FI-PPP. A testbed wanting to participate in FI-PPP and FIRE may have to transition to or at least operate in more than one of the scenarios: Collaborative Resource Sharing->Industrial Cooperative->Testbed-as-a-Service Competition. The transformation would require a testbed to adapt operation models, legal context and most likely technical implementations to be useful in each context. This will not be easy as FIRE has not been developed to support FI-PPP business experiments. An alternative approach being explored by XIFI is to build on the strengths of each infrastructure rather than pushing them into different scenarios. As such FIRE continues to support investigations into deeper technical issues and then the FI-PPP provides pathways for users to transition their ideas and technologies towards more production oriented facilities.

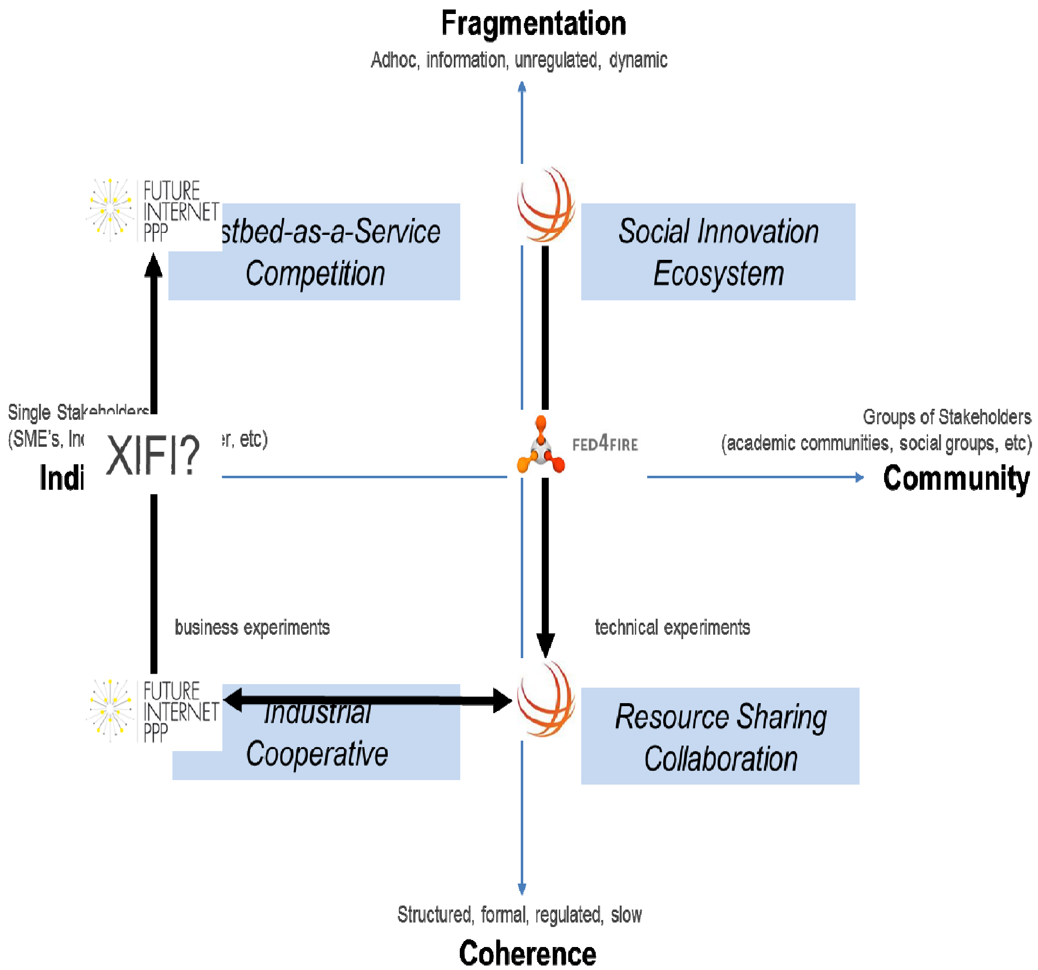

Figure 9: Driving forces in the FI-PPP and FIRE

4. FIRE Sustainability

4.1 Introduction

After presenting our vision of FIRE 20202, this chapter focuses on the issue of sustainability of FIRE and, in follow-up work, we will address the factors underlying the sustainability of the FIRE Vision 2020. Sustainability is defined as “the capacity to endure” (Wikipedia). Like in ecology and biological systems, its application to FIRE emphasises the continuing diversity, evolution and productiveness of the “FIRE System” over time. For that, a healthy ecosystem is required, which is a topic addressed in the D1.2 report. The concept of sustainability as applicable to FIRE works at multiple levels: the individual FIRE projects, but also FIRE as a whole, as a “system”.

The next section 4.2 presents some current views within the FIRE community concerning FIRE sustainability as a starting point. Section 4.3 elaborates how FIRE sustainability assessment is grounded in the **business model** concept proposed by Osterwalder and Piqueur (2010). The idea is that sustainability is enabled by a “business model”, representing the conditions for longer term development. The business model concept is a systematic approach to examine such conditions. Thereafter in section 4.4 we present a concise assessment of several cases of FIRE research and testbed infrastructures in order to analyse the trends and developments that are underlying FIRE future scenarios and future sustainability. Discussed are SmartSantander, OFELIA and TEFIS. The intention is to add other cases, for example EXPERIMEDIA, OpenLab, CREW and others in a later stage of our work.

Building on these sections, the section 4.5 takes a higher level view on the FIRE landscape and ecosystem and covers the question “what are the conditions for long term viability”. Topics that will be covered here include technical issues such as infrastructure and tools, but also non-technical issues such as targeting and engagement of FIRE users, collaboration models etc.

4.2 Approaches to sustainability within FIRE

4.2.1 FIRE Architecture Board on sustainability

Two short White Papers (2011, 2012) developed by the FIRE Architecture Board (AB) define two natural models that the FIRE projects have pursued in their efforts to ensure continuous availability of their facilities from one funding cycle to the next. The more technology-focused testbeds (examples include OpenLab, PlanetLab, and CREW) offer an “academic” platform and services. These meet the needs primarily of university and non-profit research labs. Others, which concentrate on end user-visible services (examples are BonFIRE and TEFIS) offer a more “industrial” environment, addressing the confidentiality and intellectual property concerns of potential industrial customers.

The FIRE AB White Paper on Sustainability (2012) proposes a process approach to tackle sustainability, e.g. define a roadmap, define strategy, attract users, and investigate business model options. Efforts along these lines, appropriate to the natures of the various testbeds, are already taking form in several FIRE integrated projects. In subsequent sections of this “radar” document, we will discuss them and analyze the challenges which they face using standard business modelling analyses.

4.2.2 MyFIRE Routes to Sustainability

The MyFIRE Support Action has developed a short document on Routes to Sustainability in FIRE (2012). This document presents a holistic and high-level analysis of FIRE sustainability. It identifies challenges to sustainability: public funding, commercial funding (which is found not realistic) and discusses a hybrid model e.g. public funding of commercial use: make infrastructure available for SMEs etc (as implemented by CANARIE's DAIR programme in Canada). The document further discusses longer term sustainability of infrastructure from the perspective of efficiently operated federated facilities, enabled by multiple funding streams. It proposes that federation is extended outside FIRE linking to NRENs, Géant and GENI facilities. However it also states that federation should be based on common goals. Over-all the document provides a good perspective however is mostly at the level of principles.

4.2.3 FIREBALL's collaboration frameworks

The FIREBALL Support Action was dedicated to the theme of smart cities as experimentation environments for the Future Internet, introducing the concept of Living Labs in this setting and interacting with three relevant communities: Smart Cities, Living Labs and FIRE. FIREBALL has elaborated several interesting cases of collaboration frameworks where testbeds and living labs interact to provide innovation and experimentation services, e.g. TEFIS, ELLIOT and SmartSantander (Fig. 10). Such collaboration frameworks definitely are important elements of FIRE business models and this sustainability. FIREBALL also attempted to identify the assets of such collaborations such as testbeds, living lab facilities, technologies and know-how, and specified in more detail elements of such business models governing access to these assets or resources e.g. IPR management. In this, FIREBALL also looked into cases at the urban and regional level, for example the ImaginLab in France covering various facilities and research & innovation activities in the Bretagne region, as well as cases in Manchester, Barcelona, Oulu, Helsinki, Thessaloniki and others.

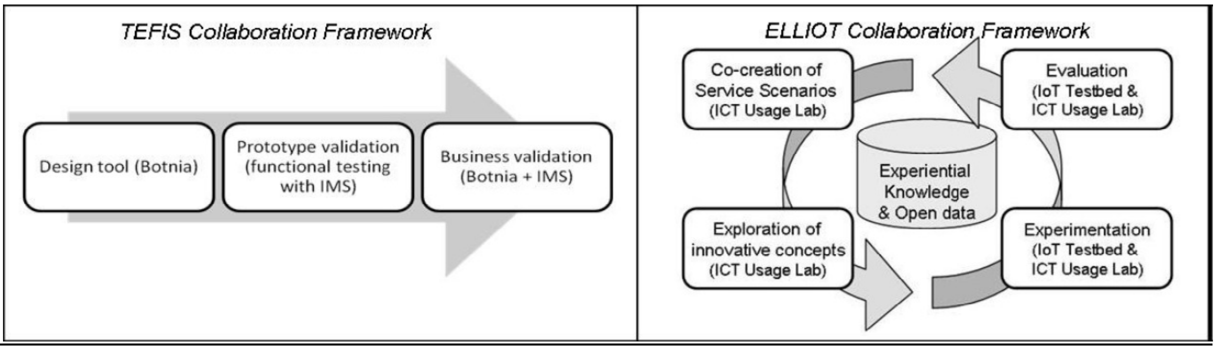

Figure 10: Different collaboration frameworks: TEFIS, ELLIOT [FIREBALL, 2011]

FIREBALL did not undertake a systematic analysis of sustainability, however to some extent issues like access to facilities and managing IP were covered.

4.2.4 OSIRIS and sustainability of Research Infrastructures

The OSIRIS project (Open and Sustainable ICT Research Infrastructure Strategy) in its deliverable D4.1 has identified components of sustainability for ICT research infrastructures. In particular, the project addressed Governance, Sustainability, Access Policy and Operational Principles. OSIRIS has studied different European ICT Research Infrastructures e.g. EGI, DANTE, PRACE, GEANT including the interactions between research – Industry – RI's. Also, OSIRIS analysed the public funding structure and concluded that public funding is 100% mostly but this amount is composed of contributions from national and EC sources of funding. Fig. 11 presents the flow of funds between various actors and entities involved.

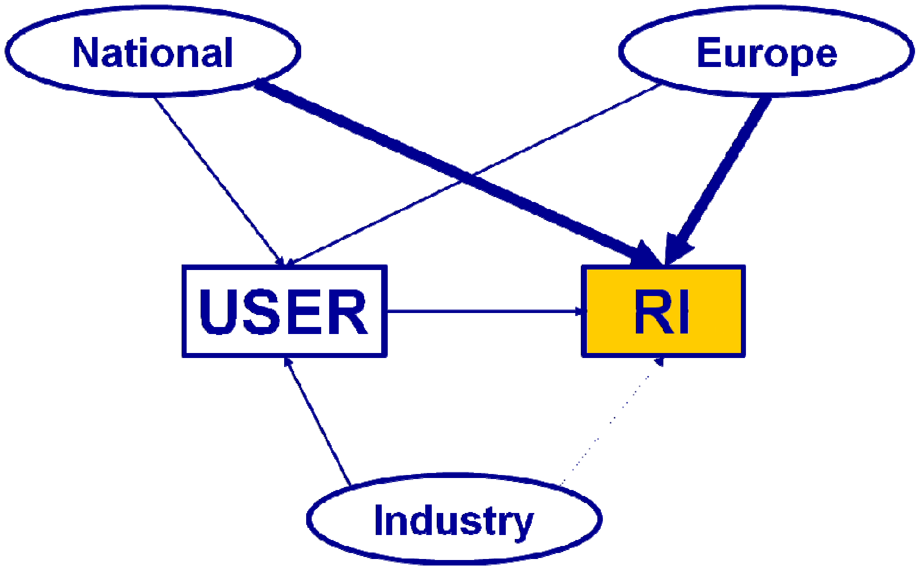

Figure 11 Flow of funds research infrastructures, for Network (Dante) (OSIRIS, 2012)

OSIRIS D4.1 has also elaborated the model for Public Authority – Research Infrastructure interaction in detail, including governing structure, legal forms, partnership arrangements and other elements visualized in Fig. 12. Results of the OSIRIS project will be form an important source for AmpliFIRE's next period activities on sustainability analysis.

Figure 12: OSIRIS analysis of Research Infrastructures business models [OSIRIS, 2012]

4.2.5 Horizon 2020 views on FIRE Sustainability

The document "FIRE in Horizon 2020" (2020) stresses the essential role of experimental platforms in Horizon 2020 as part of "Industrial Leadership" but linking strongly with "Excellent Science" and "Societal Challenges". FIRE aims to cover the whole time to market from advanced innovative Future Internet research and Experimentation to shorter term needs from industry and society, for innovative Future Internet trials infrastructures.

The main message is to continue the developments towards advanced facilities and stimulate their use, while requiring business models and experimental evaluations based on the Future Internet as common enabler to distributed computing services. The document states the importance of FIRE facilities to embrace opportunistically crowd sourced or community and citizen provided experimental resources which would contribute to sustainability. This also would extend FIRE facilities into society and end users domains. This evolution faces a number of challenges including new experimentation support mechanisms and dealing with unreliability and uncertainties of experiment resources but would enable FIRE to work on shorter time trials and experimentations and attract interest of SMEs and industry.

4.3 Business Model Perspective on FIRE Sustainability

In the context of FIRE, several terms and concepts are being used that are related to the FIRE infrastructure, service portfolio and its value proposition. These terms require a better, shared understanding. Below we provide short descriptions of the various concepts.

- **Sustainability**. Sustainability denotes the current and future capacity to remain viable as regards the FIRE facility and FIRE-related experimentation activities. This is based on the attractiveness of its service offer (availability of experimental platforms, support of innovative experimentation) and its funding prospects, and on other conditions like infrastructure and governance. So, future sustainability of FIRE is based on the different elements of its "business model".
- **Infrastructure**. Infrastructure of FIRE includes its experiment facilities, but also its experiment process, and its community network of stakeholders and partners including research communities, businesses and governmental organisations.
- **Service Offer**. This is the portfolio of services offered by FIRE to its users (experimenters), under specific conditions (quality, price, availability …). The user experiences the service offer as a value proposition: does it meet the experimentation demands?
- **Business Model**: this is a consistent description of how FIRE creates value for its stakeholders. It describes the conditions to be set in place in order to provide services and achieve sustainability. These conditions include infrastructure, service offer and value proposition, customers (user base, use conditions), finances etc.
- **Exploitation**: this concept relates to both "business exploitation" of FIRE facilities as testbeds serving experimenter demands, and to the exploitation of FIRE assets such as IP.
- **Governance**: the set of management and decision processes, structures and practices shaping the actual FIRE operations and activities. This also could be seen as part of infrastructure.

<table>
<tr>
<td colspan="2" rowspan="2">Key partners
• FIRE projects partners
• European Commission and national authorities
• Projects and initiatives outside FIRE</td>
<td colspan="2">Key activities
• Experiment practices
• Experimenter core activities
• Exploitation activities: testbeds, service offering</td>
<td colspan="2" rowspan="2">Value proposition
• Service portfolio
• Match with Experimenter demands</td>
<td colspan="2">Customer relationships
• Legal model
• Governance model</td>
<td colspan="2" rowspan="2">Customer segments
• Experiment communities
• Large business
• Advanced SMEs
• National Future Internet initiatives' stakeholders</td>
</tr>
<tr>
<td colspan="2">Key resources
• Testbed facilities and federation
• Experiment tools</td>
<td colspan="2">Channels
• FIRE PR and promotion channels
• Open Calls
• Developer communities</td>
</tr>
<tr>
<td colspan="5">Cost structure
• Testbed facility assets cost structure
• Operations and maintenance
• Marketing and PR, Community support
• Experiment cost</td>
<td colspan="5">Revenue streams
• EC funding
• National funding
• Projects co-funding
• Service pricing revenues</td>
</tr>
</table>

Fig. 13: FIRE business model framework using CANVAS

These concepts find their place in the “business model” concept as represented by the CANVAS model (Osterwalder and Pigneur, 2010). Originally, this “business model” concept was developed for private sector activities. Alternativelyo termed “Operations Model”, it easily can be easily to the (semi-) public sector activities context as well. Specifically applied to the context of FIRE Sustainability, this CANVAS Business Model representation is depicted in Figure 13.

Based on this conceptual framework, the different components of the FIRE business model can be more precisely elaborated with a view towards responding to future challenges. Alternative business models, taking into account options regarding the specific components and how we expect them to evolve over time, can be built as part of FIRE 2020 future scenarios. The D3.6 FIRE Roadmap notes: “There is only a fragmented view of the sustainability of the experimentation facilities, both from a financial as well as technical point of view”. Our challenge is to enhance this state of affairs and propose sustainability models that are viable.

4.4 Sustainability of FIRE Projects

This section will look into the FIRE business model issue at the level of individual existing projects. Applying the CANVAS business model concept to FIRE projects provides us a deeper and more concrete insight in the problems in preserving sustainability at project level.

4.4.1 SmartSantander

SmartSantander (2010 – 2013) is an FP7-ICT Integrated Project, starting with strong institutional support from the Santander municipality. We have selected this case as an example of the exploitation strategy of a FIRE project. To illustrate the sustainability and future exploitation aspects we have included, besides an analysis of exploitation strategy, some materials from the FIREBALL project which studied this case.

4.4.1.1 Introduction to SmartSantander

SmartSantander is focused on providing a Smart City laboratory for testing all types of Smart City solutions, ranging from the use of sensors and their networking technologies to the use of service platforms for collecting sensor information and deploying services. It is based on a number of connected smart city testbeds, including one large Smart City experiment in Santander and smaller experiments in Belgrade, Guilford and Lübeck.

The SmartSantander research facility is sufficiently large, open and flexible to enable horizontal and vertical federation with other experimental facilities and to stimulate the development of new applications by different types of users, including experimental advanced research on IoT technologies, and realistic impact assessment based on users’ acceptability tests. The facility will comprise more than 20,000 sensors and will be based on a real life IoT deployment in an urban setting. The core of the facility will be located in the city of Santander and its surroundings, on the north coast of Spain. SmartSantander embraces the idea of enabling the Future Internet of Things to become a reality applying a living labs approach.

4.4.1.2 SmartSantander testbed infrastructure

Although the main target of SmartSantander is research oriented to create a large-scale testbed allowing open experimentation with key enabling IoT device technologies, it is obvious that such a realistic setting offers the potential of involving real end-users in the experimentation process. There is a long list of applications identified by SmartSantander, in close cooperation with the City Council and the Regional Government of Cantabria, as suitable to be supported by the infrastructure being deployed. Most of them offer a big environmental and social potential: parking spaces and traffic control, environmental management and monitoring (pollution, CO_2, noise, etc.), public installations management (heating, A/C, lighting, etc.), public transportation, parks and gardens control (irrigation, etc), social assistance (elderly, disabled, etc.), etc. Due to time and budget limitations, during the execution of the project just some specific services will be deployed in order to validate the asset deployed. Other interesting and more advanced services are expected to come up later on as a result of parallel initiatives linked to the project at the regional level, as the project is committed to ensure the availability of the infrastructure beyond the end of the project.

The testbed infratructure will be operated and maintained by the consortium during the execution of the project. After that period, several solutions are being considered. Among the choices that are being currently envisaged, and will be further analysed, are the creation of a new legal entity for its exploitation, and/or the transfer of both maintenance obligations and ownership to a third party. In both cases, the use of the asset would have to be bound to legal and financial conditions.

The benefits of the infrastructure addressed by the SmartSantander project are two-fold:

- The deployed facility will enable a wide range of experimentations, supporting different technology aspects and catering for different user groups (researchers, service providers, and end users). Furthermore, through FIREstation, the project collaborates with other FIRE projects to allow the federation with their respective experimental facilities.
- SmartSantander aims at optimizing the societal benefits of investing to build up such a city-scale infrastructure, so it was designed to support real life services, useful to the citizen, at the same time it copes with its primary target of providing an ambitious experimentation platform for the research community. E.g. first cycle deployment consists of a big number of parking sensors able to provide support for experimentation of multi-hop techniques on different topologies, and will also provide the City Council means to control the proper use of the parking spaces reserved to disabled people.

Table 7 specifies the SmartSantander testbed infrastructure assets that can be accessed and shared.

Asset type	Specification of the asset	Shareable asset
Network infrastructure	Heterogeneous Wireless Sensor Network, with specific experimentation capabilities allowing remote configuration of the different types of nodes (sensors, repeaters, and gateways).	It will be available under specific conditions: experiments to be carried out on top of it should pass a ‘sanity test’ to ensure they do not compromise the infrastructure itself. Deep technological knowledge would be required.
Software applications	Basic applications for node configuration and management in order to be able to validate the operation of the system. Initial approach of first set of service oriented applications related to the management of the parking spaces.	Access to basic applications would be granted for experimentation purposes in case it is required. Applications for specific services being competence of the municipality not within the scope.
Innovation environments user communities	Currently not available. They will be addressed during the execution of the project, once the infrastructure is available, to involve third parties and end-users in the creation of services based on the sensors’ data.	Will be available in the future, based on a Living Labs approach. The access will be limited to non-sensitive information to guarantee personal data protection, and prevent misuse of the information provided.

Sustainability and exploitation plan	Information of the outmost importance to describe the models being considered within the project, with emphasis on those more suitable to guarantee the sustainability of the infrastructure.	The report analyzing potential exploitation models will be also publicly available through the project web-site at: http://www.smartsantander.eu.
Public data / information	A number of different information categories will be opened up to the public, to enable the use of applications, and the development of new ones.	Open APIs for accessing data will be made available at three different levels: research and developers' community, Service Providers (ISPs), and end-user.

Table 7: SmartSantander project most important common assets [Source: FIREBALL]

Apart from this, SmartSantander is aware of its potential to reduce time to market for new services, by shortening required R&D cycles, providing a fast end-user feedback for the assessment on socio-economic impact to the European researchers and service developers, and helping to make technology benefits more visible to the EU citizens. This will be facilitated by the deployment of novel IoT solutions and application pilots on a realistic target environment involving real end-users. Besides, an early end-user exposure to the first applications and services based on IoT technologies can encourage its adoption and lower the boundaries of social acceptance by the public, which often acts as an inhibitor of technological advance.

The first deployment phase was carried out in Santander. By June 2011, most of the first 2,000 sensors corresponding to the first phase of the project had been deployed across the city. Using this preliminary approach to the final testbed, the project will issue the first Open Call to select proposal to be funded in order to run experimental research on top of it. At the same time, end-user perception with regard to the first services was analyzed by means of surveys among the citizenship, and some services related to specific urban mobility use-cases will be further improved under a Customer Driven Innovation approach (CDI). These methodologies are also common to most Living Labs experiences. In the future stages of the project, and once the assets become progressively and publicly available, it is expected to involve wider communities in the usage of the infrastructure.

4.4.1.3 Exploitation plan SmartSantander

SmartSantander has an exploitation WP where detailed plans on the exploitation possibilities of the infrastructure are being drawn in Figure 14. Even if those plans are not yet completed, the main idea would be the creation of a number of organizations, around the Santander Council, to exploit the sensors and the platform. The Council, who owns the sensors and locations, will offer to industrial companies the use of those sensors for the provision of a number of city services, which will be paid by the Municipality using the conventional public offering model. Sensors should allow easier and more efficient service provision and therefore, the company interested in offering the service to the Santander Municipality will find it useful to pay for the information extracted from the sensors. The service platform is also offered to the Municipality by the Operating Company, Telefónica in this case, provided that infrastructure support and maintenance services are contracted to them.

The exploitation model of SmartSantander is self-sustained since it is based on existing contracts to the Santander Municipality. Also, it is expected it will represent significant savings to the Municipality, since the prices charge will be reduced. Companies offering services will also benefit, because they can offer a more competitive service and be more efficient. Another important aspect will be the change in service philosophy: instead of paying for a service to be performed, payment will be made based on the results. For example, in the case of garbage collection, the company will not be paid by number of trucks used or amount of material taken, but rather by the final results on the street, which can be measured by the sensors in real time.

In principle, the benefits of the project are then transferred to Santander Municipality and Santander citizens, who have more efficient and cheaper municipal services. Utilities offering the services also benefit from a more effective situation. The use of the same set of sensors for different services should bring in economies of scale.

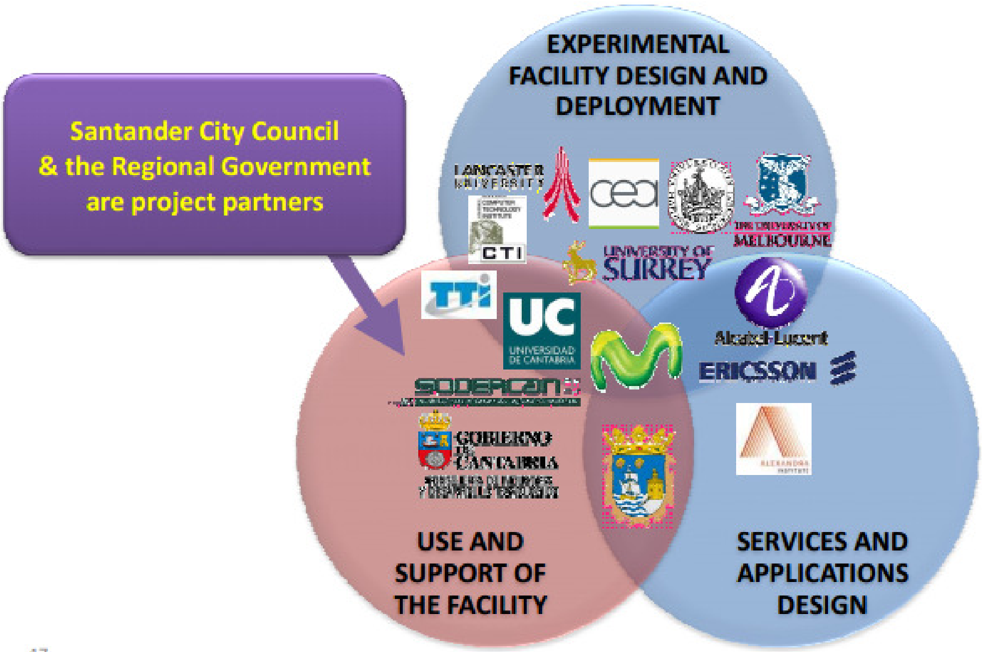

Figure 14. SmartSantander exploitation ecosystem

Besides those exploitation plans, SmartSantander experiments have already been very beneficial to the region, since a number of Start-ups, many of them based on the university staff, have been created to develop technology solutions in this industrial activity area.

4.4.1.4 Limitations to the present exploitation plans

Even if the SmartSantander exploitation model is attractive, it requires significant involvement from the companies providing local services which are in the end, paying for the sensor deployment. In the case of Santander, there has been a very important effort to involve the local companies from the beginning, since an important objective of the project has been the promotion of the Cantabria region where Santander is located.

As regards sensor deployment and their usage, even if the process has been carefully planned, in some case, there are some complaints that the sensors are not ad-hoc designed for the specific service objective. Day to day operation may require additional infrastructure deployments. There could also be difficulties in actual operation, since a large proportion of the local services companies are not trained in the new equipment. However, the main limitation for replicating this model is that SmartSantander required the decided impulse Santander Municipality and an enthusiastic Major. In the case of SmartSantander, the Municipality considered the project as their own from the start. It would be very difficult to really replicate the success in a different environment if their collaboration cannot be assured. Also very important is the implication of the University of Santander, who is working in close collaboration with Santander's Mayor and local industry.

The support of the University in the initial stages of the project has proved to be essential. It provided, partially thanks to the EC support, the specialized workforce to analyse protocols, study alternatives and also to publicise the project in order to get extra funding without which the endeavour would have been very difficult to implement.

Besides that, even counting with the support of the Municipality, it is not simple to replicate the model. SmartSantander has been a very successful model to start the Smart city concept, adjust ideas, and develop a new platform and get significant funding. Unfortunately, it is difficult to imitate because, even if many cities are interested, they are actually looking for funding at a narrow local level, where the innovation dimension is no longer present. Also, for the operating companies (such as national or global service providers) it is not useful to repeat the same experiments again and would like rather to re-use the results of Santander or from experiments made elsewhere. All this makes it much more difficult to get the funding, and thus the enthusiastic support of the local authorities which was essential in the project

4.4.1.5 Possible alternative approaches

There are two alternatives for pursuing the SmartSantander results: the reuse of the facilities in Santander and the use of the experience related to the platform and protocols. As for the use of the sensors and facilities, SmartSantander exploitation model allows the use of the facilities by large companies, i.e. sensors and platform, to perform experiments of their own in the city of Santander, or, alternatively in other cities of the consortium. Those experiments can consist of testing new protocols and sensor models, and even include some preliminary services design. These experiments could receive some feedback from the local population. This way, Santander could become a living "Smart city testing environment". The attractiveness of this approach has been already demonstrated, and some companies are approaching the centre for research on smart cities (CiCiS) recently created at a local level in parallel to the project.

An alternative possibility would be to use the experience gained in the platform and protocol development in other, larger projects to continue building on the Smart City concept, which should go beyond just offering services in a more effective manner. The PPP Future of internet could be a good area for innovation, which could be extended in H2020. This is also an interesting opportunity and there are a number of actions already under way to include Santander in the test-beds. However, new experiments and ideas should be designed to make it attractive beyond existing projects.

4.4.1.6 SmartSantander sustainability, a summary

Based on preceding materials, an initial analysis and assessment of underlying forces affecting future sustainability of SmartSantander can be performed. To this end, the earlier introduced CANVAS framework for business model analysis is used in Fig. 15.

<table>
<tr>
<td rowspan="2">Key partners
Institutional support from Santander municipality
Consortium partners e.g. Telefónica, Santander Municipality and several companies
Santander Council as key actor for exploitation</td>
<td>Key activities
Building a Smart City laboratory
Testing smart city solutions</td>
<td rowspan="2">Value proposition
Use of sensors for provision of city services
Paid by the Public Offering model
Service platform offered to the Municipality by the Operating Company (Telefónica)
Infrastructure support and maintenance to be contracted to Telefónica</td>
<td>Customer relationships
Santander Council and connected organizations to exploit the sensors and platform</td>
<td rowspan="2">Customer segments
Organisations around the Santander Council are exploiting the sensors and platform.
Industrial companies will use the sensors to provide city services</td>
</tr>
<tr>
<td>Key resources
Smart Santander facilities
How is maintenance of the facility ensured after 2013?</td>
<td>Channels
Santander Council</td>
</tr>
<tr>
<td colspan="3">Cost structure
Platform facility cost structure
Operations and maintenance
Marketing and PR</td>
<td colspan="2">Revenue streams
EC funding will be finalized in 2013
Will larger companies and local users invest in continuation?
Savings to the Municipality because of prices reduction: prices are result-based</td>
</tr>
</table>

Figure 15: SmartSantander business model framework using CANVAS

This presentation identifies the key factors, uncertainties and questions that determine current and future sustainability, and is a starting point for further analysis.

4.4.1.7 Conclusion

SmartSantander has evolved into a platform which seems to be capable in principle to provide commercial services in the public domain (e.g. garbage collection). It has been noted that it is not simple to replicate the model as funding is a problem. The SmartSantander case is very different from other FIRE projects as SmartSantander had always the orientation towards practical exploitation (whether this is viable remains to be seen). It might be considered as a kind of "spin-off" of FIRE, and FIRE as facility and experimentation infrastructure might have the challenge to generate more of such spin-off projects especially as regards its research and experimentation projects.

4.4.2 OFELIA

4.4.2.1 Introduction

OFELIA is a large FIRE Project whose main objective is to provide an experimental facility to test network architectures and solutions. It is based on the use of Open Flow over a number of programmable switches. Open Flow is a novel approach to provide network virtualization through secure (and standard) interfaces. Total project budget is 6.3 M€ and has a duration of 3 years, starting in September 2010. The project is led by EiCT (which is a Public Private Research Center) and has participation from major operators, such as Deutsche Telecom, Instituto de telecomunicacoes Aveiro –representing Portugal Telecom, I2CAT, as well as other universities and research centers. OFELIA is offering up to 10 Open flow islands, located in several locations in Europe. OFELIA has become the pivotal project related to experiments and deployments using Open Flow and new protocols that could help to increase the efficiency of future networks. It has become a focal point for its research community and is producing very useful results.

The main objective of OFELIA is to provide experimental facilities in an open source environment in which to explore software-defined networking (SDN) using OpenFlow as a critical enabler. The project is very much focused into R&D research as indicated by the relatively large number of universities. As in the case of SmartSantander, there have been a number of use cases, using OFELIA resources, such as Vertigo, which is analyzing specific features of the Open flow protocol.

4.4.2.2 Infrastructure

OFELIA is offering up to 10 Open flow islands, located in several locations in Europe.

- Berlin, Germany (TUB) – partial replacement of existing campus network with OF-switches
- Ghent, Belgium (iMinds) – central hub, large-scale emulation
- Zurich, Switzerland (ETH) – L2 (NEC) switches mesh, connection to OneLab and GENI
- Barcelona, Spain (i2CAT) – L2 (NEC) switches and optical equipment (ROADM ring)
- Bristol, UK (UNIVBRIS) – national hub for UK optical community; optical (ADVA, Calient), L2 (NEC, Extreme) switches, FPGA testbed
- Catania, Italy (CNIT) – two islands, based on NetFPGA and OpenSwitch technologies, with focus on ICN (Infomation Centric Networking)
- Rome, Italy (CNIT) – based on NetFPGA and OpenSwitch technologies, with focus on ICN - under deployment
- Trento, Italy (CREATE-NET) – a city-wide distributed island based on L2 (NEC) switches and NetFPGA; opt-in users via heterogeneous access technologies
- Pisa, Italy (CNIT, 2 locations) - based on NetFPGA and OpenSwitch technologies, with focus on Cloud Data Center management - under deployment
- Uberlândia, Brazil (UFU) - under deployment.

4.4.2.3 Exploitation plans

Present exploitation plans in OFELIA can be considered only preliminary. OFELIA partners consist essentially of Universities and research centres and the main objective is to provide means for research and continued academic activity. The main plan, therefore, is to continue going to FIRE projects and making experiments. At individual level, some of the companies in OFELIA have some plans to reuse the infrastructure at local level, through local or regional funding. Also, in the past, before the introduction of Open Flow, there were some initial ideas about a model based on pay-per use, but they did not materialize. Those plans are being considered by the consortium (or parts of it) and they may be continued, however, there is not a really effective action in that direction. Lastly, some of the nodes in OFELIA are being replicated by GEANT and incorporated to their network. This could also be a very interesting possibility for continuing OFELIA activities.

4.4.2.4 Limitations to present exploitation plans

The main problem of the present exploitation model is the fact that it is very limited to the project partners. Even if companies and universities in OFELIA would be most interested to find a suitable exploitation model, getting funding from external, industrial, companies, this this not straight forward due both to internal and external reasons. As regards internal reasons, there are a significant number of partners who belong to universities and public institutions. This makes it very difficult to react with enough flexibility and rapidity. Certainly this is a minor problem, since most public organizations would be willing to sell the equipment.

However, the relatively low cost of the equipment makes it very difficult to justify the expenses to the buying company.

Concerning external reasons, if the external commercial companies would need to make a significant investment in the facilities, they would need to make a careful business analysis of the advantages of using OFELIA facilities as compared to owning the equipment. In the business case, important aspects would be the service level of the facilities.

4.4.2.5 Possible alternative approaches

OFELIA facilities are one of the most interesting in Europe for testing Open Flow, and they are useful to many operating companies which would need to deploy similar equipment for their own experiments. The easiest way forward would be to continue offering those facilities to research groups, perhaps in the context of EU funded projects. Alternatively, it would be easier if those facilities could be transferred to a pan-European R&D centre which could rent or lend those facilities under the adequate contract, such as GEANT. As indicated before, this approach has difficulties, even if it could be applicable in some cases. As a third alternative, OFELIA could be offered to local or national operating companies, based on prices paid by experiment or by project. This approach however, has the disadvantage that it will be difficult to maintain the ten node network connected since it will be difficult to get support for all partners.

The possibilities of offering the services of OFELIA to commercial companies have to be analysed and implemented quickly. Many vendors could offer test beds to operating companies at a very good price, and the advantages of OFELIA could become less significant in the coming months. Also, the equipment and results would become obsolete relatively fast, so swift action would be needed.

4.4.2.6 OFELIA sustainability

Based on preceding materials, an initial analysis and assessment of underlying forces affecting future sustainability of OFELIA can be performed. To this end, the CANVAS framework for business model analysis is used (Fig. 16).

<table>
<tr>
<td rowspan="2">Key partners
• Uncertain prospects for partnering after project completion
• Should partner within FIRE</td>
<td>Key activities
• Provide experimental facilities to perform Open Flow experiments</td>
<td rowspan="2">Value proposition
• Offer experimentation services to FIRE projects
• What is the potential to reuse the infrastructure at local level?
• What is the potential to reuse parts of the infrastructure within GÉANT?</td>
<td>Customer relationships
• How are customer relationships built up and exploited?</td>
<td rowspan="2">Customer segments
• GÉANT
• Other FIRE projects
• Companies demanding testbeds</td>
</tr>
<tr>
<td>Key resources
• OPHELIA facilities, experiment procedures etc</td>
<td>Channels
• Which channels are being used?</td>
</tr>
<tr>
<td colspan="2">Cost structure
• OPHELIA facility maintenance costs
• Experiment costs</td>
<td colspan="3">Revenue streams
• Dependency on FIRE funding of experiments
• Is local or regional funding a possibility, in order to reuse the infrastructure ate local level?
• Is pay-per-use a realistic possibility?</td>
</tr>
</table>

Figure 16: OFELIA business model framework using CANVAS

Like SmartSantander, the picture forms a starting point for further analysis, identifying the key factors, questions and uncertainties determining current and future sustainability.

4.4.2.7 Conclusions

OFELIA has the mission to provide experimental facilities (to perform Open Flow experiments). It is much focused on R&D. As noted in such circumstances it is difficult to develop a sound exploitation plan involving business and local governments. Exploitation would lie more in integrating the project results and the created assets and technologies into other facility and experimentation activities within FIRE or GÉANT.

4.4.3 TEFIS

4.4.3.1 Introduction

TEFIS supports Future Internet of Services research by offering a single access point to different testing and experimental facilities for communities of software and business developers to test, experiment, and collaboratively elaborate knowledge. It offers an open platform to access heterogeneous and complementary experimental facilities, including living lab facility, and testing tools to be used by service developers supporting the service development life-cycle. The platform provides the necessary services that will allow the management of underlying testbeds resources throughout the entire service-development lifecycle.

4.4.3.2 Testbed Infrastructure

TEFIS is selected as example of bringing together Future Internet / IoT and living labs resources for the purpose of smart city innovations or other desired outcomes of the project because of the following:

- An experimental platform for Smart Cities development empowered by Future Internet technologies
- An open framework that will allow efficient combination of various experimental facilities to support the heterogeneity aspects of Future internet experiments including the end-user involvement
- A platform to share expertise and best practices for higher “smartness” by shared intelligence and experiences

Two main types of assets are available via TEFIS for future Smart Cities experimentations: the platform and the testbed facilities provided by partners of TEFIS (Table 8). The TEFIS platform is organised into four main functional blocks: the portal, core services (middleware), testbed connectors and user tools. It offers different types of support for Future Internet experiments such as designing, planning, management of experimental workflow, configuration assistance, experimental data management, reporting, knowledge sharing with other experimenters and access to different testbed facilities and service offers independent of geographical location. The testbed facilities provided by testbed partners of TEFIS include a wide spectre of testing and living lab opportunities.

The following project case illustrates how in TEFIS resources are combined and shared. This specific Future Internet experiment is combining experimental resources from two different testbeds; the SQS IMS testbed in Spain and the Botnia Living Lab in Sweden. The experiment is focused on a mobile application over IMS, and is divided into three different phases of the service development life-cycle: concept development, prototype development and business model definition. First, this experiment will explore end-user feedback to check if the application is suitable and would be useful for users by access to Botnia Living Lab assets. In the second step, they will use the IMS-testbed facilities as a validation tool to perform system acceptance testing (including functional and non-functional), and Botnia Living Lab for usability evaluation with end-users. In the third step, they want to identify the correct business model for long-term sustainability. In this third phase both end-users feedback and network usage is monitored and analysed, and for that purpose the IMS testbed and Botnia Living Lab are combined.

Asset type	Specification of the asset	Shareable asset
Network infrastructure	Planetlab: powerful infrastructure consisting of 1018 nodes for testing and evaluation of network protocols and distributed systems on a large scale. PACA Grid: a computing infrastructure for large-scale computations and a number of tools to automatically deploy and execute distributed applications and to monitor the progress of the computation and retrieve the results. ETICS: a build and test job execution system based on the Metronome software and an integrated set of web services and software engineering tools to design, maintain and control build and test scenarios. SQS IMS: Assets: The emulated IMS platform with IMS Core services, Presence and Group management, Push-to-talk, IMS Messaging, Instant messaging and Instant Multimedia Messaging, GSMA video/image share and enhanced VoIP and IMS Core Network emulator. Wizards and templates included in the tools are used for testing purposes. KyaTera: A high speed network of over 266 km of optical cables with 8 to 144 fibres and a network measurement tool to measure network status as bandwidth, jitter, delay, ping between two nodes, packet loss etc.	For sharing outside the TEFIS CA of these assets each Testbed facility provider has its own regulation for sharing and access to their assets.
Software applications	The TEFIS platform is organized into four main functional blocks: TEFIS Portal, TEFIS Middleware, TEFIS testbed connectors and TEFIS User tools.. The User tools will be external tools, which could not be free, that the TEFIS platform can embed in a future next step	The TEFIS platform it is being developed under the conditions of the Open License Terms.
Innovation environments user communities	Botnia Living Lab: Research expertise in end-user evaluation and testing, the FormIT methodology for end-user involvement, a database of 6000 creative end-users in Sweden and access to end-users around the world via 3rd parties.	These assets are available to any user and access is regulated depending on what kind of resources, Handbooks are available
Sustainability and exploitation plan	Assets above provided via the different actors of TEFIS are in use today in internal cases and with external actors. Exploitation work is in progress on the networked offers for users of the facilities and for the Tefis facility itself. A specific framework is used for the exploitation and sustainability processes.	Framework for business model creation, development and evaluation.
Public data / information	Depending on the users and each experiment data can be made public. At the minimum general information about each experiment is to be public available for knowledge sharing and visibility.	General information about each experiment using the Tefis portal for their performance.

Table 8: TEFIS Project Most Important Common assets (Source: FIREBALL)

4.4.3.3 Exploitation plans

The TEFIS single-access point has been used in different experiments and they serve as the baseline for exploitation showing the diversity of potential users of TEFIS. Other assets for exploitation include the TEFIS toolkit for testbed management. The exploitation strategy is based on individual partner exploitation plans (exploitation of TEFIS components, know-how and individual testbed services) and joint exploitation of TEFIS by an informal partner network established early 2013. This network aims to further develop the concept of Testbed-as-a-Service, and interact with other initiatives within and outside FIRE) for long term synergies and strategic development. The network will act as the TEFIS operator when the project is finished by mid-2013. Limitations of the exploitation plans, which are relevant for the business model discussion, are seen in three aspects: 1) the TEFIS revenue model, 2) Testbed usage policies, 3) User expansion.

The TEFIS revenue model can shortly be described as a prosumer-model with shared interest and responsibilities with R&D projects to fund usage, development and maintenance. To this end it means that new projects where TEFIS usage is included needs to be approved for external funding. TEFIS has been included in 5 FP7 proposals for Call 10 and they serve as the initial potential continuity for TEFIS. To this end TEFIS parties they need to be active in new proposals as well as to include TEFIS exploitation into their daily operations for TEFIS to be sustainable. TEFIS must offer key-benefits for them in different dimensions: Testbed operation, partner collaboration and R&D.

One additional limitation regarding exploitation is the diversity in usage-policies. During the project this has not been an issue as the users have been partners of a consortium and therefor usage have been easily handled. In further exploitation the differences in usage-policies among testbed providers has to be considered and handled when new users access the testbed services via TEFIS. This is both a technical issue as well as part of the TEFIS revenue model.

Regarding users expansion the usage have so far been funded by EC funding via Open call and this has been a valuable instrument for TEFIS development. For TEFIS use similar instruments are necessary to foster usage as well as to sustain and develop the infrastructures. In the TEFIS development the users of TEFIS have been the key-drivers and they are the core-heart of the solution. TEFIS must continue to be their environment and to get more users on-board some more investments from EC will be necessary – not only for development but also for maintenance and use. The concept of testbeds is still very immature and unknown and to develop the concept as well as to invest in usage some public grants are necessary.

An alternative approach for TEFIS exploitation could be to formalize a commercial actor as key-exploitation body. This alternative has not been chosen due to market immaturity as well as the interest for partners were more according to form an informal association/network for inter-organizational R&D.

Considering users TEFIS has so far been approaching a very diverse user-base and from the Open call users we reached our target. Another approach could have been to target a specific sector of Future Internet R&D that is in need of fore-front testbed infrastructures meaning those not having this infrastructure themselves. From the Open Call we saw an interesting phenomena where users also could become resource-providers and therefor an alternative approach could be to target users with own infrastructure as both users as providers of TEFIS – to implement a shared interest and equality of every party. Platforms for Future Internet experimentation for the future should in this sense not only focus on technical assets but also on social capital and "soft assets" and this could be an alternative approach for TEFIS exploitation.

4.4.3.4 TEFIS sustainability

Based on preceding materials, an initial analysis and assessment of underlying forces affecting future sustainability of TEFIS can be performed. To this end, the CANVAS framework for business model analysis is used (Fig. 17).

<table>
<tr>
<td rowspan="2">Key partners
• What are the partners that could exploit TEFIS?</td>
<td>Key activities
• How to maintain the testbed infrastructure after project end?</td>
<td rowspan="2">Value proposition
• What is the service that can be marketed by TEFIS?</td>
<td>Customer relationships
• How to attract public funding to support attraction of users?</td>
<td rowspan="2">Customer segments
• Who are the TEFIS customers?
• How to manage differences in usage policies?</td>
</tr>
<tr>
<td>Key resources</td>
<td>Channels</td>
</tr>
<tr>
<td colspan="2">Cost structure
• Testbed facility assets cost structure
• Operations and maintenance
• Marketing and PR, Community support
• Experiment cost</td>
<td colspan="3">Revenue streams
• EC funding will end soon
• What are prospects for continuing based on national funding
• Is it realistic to expect service pricing revenues</td>
</tr>
</table>

Fig. 17: TEFIS business model framework using CANVAS

4.4.3.5 Conclusion

The TEFIS project provides some interesting opportunities for exploitation. Its business model could be based on creating a close collaboration between domain actors, service providers and technology offers.

4.4.4 SWOT analysis of FIRE Facility sustainability

Based on materials presented here, only an initial SWOT analysis of the sustainability of FIRE facilities can be performed, as a basis for testing and validating the current business models, finding weak and strong spots, and preparing for enhancing the business models. The SWOT analysis in Fig. 18 presents a very initial impression gained so far and forms a starting point for more detailed analysis.

Opportunities	**Threats, vulnerabilities**
• Federated FIRE facility as core backbone of pan-European research and innovation ecosystems	• Uncertainty regarding EU funding • Dependency on EU funding
Strengths	**Weaknesses**
• Facility infrastructure • FIRE experimenter community	• Lack of business and SMEs interest • Service offering • Governance model

Fig. 18: SWOT analysis of FIRE sustainability

4.5 FIRE System Sustainability

4.5.1 Trends and developments affecting FIRE sustainability

Using the CANVAS framework introduced above, and having examined the project-level sustainability in FIRE projects, we will study the current and future sustainability of the FIRE landscape as a "system" first by looking at the current elements and the trends and changes that can be observed. Within the FIRE program, several studies have been performed which provide some insight in elements of FIRE sustainability and their trends and changes. The main source for the below Table 9 which summarizes these elements and trends is FIRE STATION's D3.6 FIRE Roadmap.

	Current Situation 2012-2013	**Trends and changes expected 2013-2015**
Infrastructure	Set of FIRE testbeds in place Ongoing work towards federation FIRE Community partner network emerging	Towards increasing federation support, easy access by user communities according to their needs, availability of experiment lifecycle management tools etc.
Research / experiment practices	From traditional design orientation to experimentation, towards experimentally driven RTD	Experimentally driven RTD established Stakeholder engagement and user acceptance
User / customer base (experimenters communities)	Experimenters from Internet communities: academic, industry, and academic-industry partnerships Layered organization of experimenter communities with own facilities	More emphasis to experimenters from the Web community More emphasis to the Future Internet ecosystem as experimenter environment Interlinked experimenter communities Large ICT-industrial users
Experimenter core activities	Development of Internet IT and networks products and services	Development of Web-based applications, services and content Federation of resources and data
Experimenter critical demands	Project-based demands mainly	Experiment life cycle management resource provision and tools Involvement of different communities Contractual and legal aspects Responding to service orientation (PaaS, IaaS) Open access to facilities
Service Offer / Value proposition	See FIRE STATION	Towards federated testing service offering Experiment life-cycle management in federated environments Trustworthiness enhancing services (D3.6 – 4) e.g. federated identity management and access control, SLA management Shared support services (D3.6 – 5) Support new, nomadic, large scale experiments on demand – See also: slides of Piet (FIRE 2020)
Legal model	Without substantial legal arrangements defining the collaboration and its risks	Legal base of operation e.g. access, property rights gets more attention. Emerging issues are federation contracting, multi-party collaboration agreements (ownership, access rights, property rights, legal protection etc)
Governance model	Ad-hoc, project based	Governance principles to be aligned with the requirements of federated testbeds. New trend of Joint Operations Centre for managing testbed resources. Wider range of governance models e.g. self-organisation, special organizational model, non-commercial organization.

	Current Situation 2012-2013	Trends and changes expected 2013-2015
		Joint operations centre model for managing Testbed resources. Federation Authority managing the interconnection of testbeds. Federation Mechanisms (D3.6 – 3.3.3)
Cost structure	Cost structure based on project facilities and project services. Hardware, software, personnel costs. Cost drivers …	Cost structure based on federated facilities and services including governance. Opportunities of sharing and exploiting economies of scale and scope. Cost accounting approaches such as activity-based costing.
Financing model	National and EC funding mechanisms targeting the financing of project-based facilities. Project financing. Use of Open Call mechanism.	Hybrid financing models combining different sources of funding (public at EC, national, regional + usage related) National and EC funding mechanisms targeting Federation strategies. Research Infrastructures funding. Usage-based fixed-variable pricing strategies (compare telecoms, energy pricing in network markets)
Exploitation		Exploitation of FIRE facilities for smart cities, living labs, science parc activities etc Exploitation of FIRE assets for large business users

Table 9: FIRE initial Sustainability analysis based on analysis of trends and changes

In follow-up activities we will undertake a more systematic analysis of the "FIRE System" business model and its trends and changes. Still, the Table highlights a number of highly important developments which affect FIRE's future sustainability, such as the trends towards federation, more easy access by user communities, more emphasis to testbed facility's services based on user (experimenter) needs in terms of experiment life cycle support, attempts to widen the user base and establish collaboration with related types of stakeholders, more attention paid to the legal model of accessing facilities, and also increasing emphasis to governance of (federated) testbeds.

4.5.2 CANVAS analysis of FIRE sustainability

Using the CANVAS framework at the level of "FIRE System" we use the materials identified in the previous section to get a glance of what the driving forces are that are acting on the current FIRE facility business model and what kind of vision could be thought of for the future as regards the evolution of this business model. The following Figure 19 presents a view on these forces based on identification of the main certainties as well as uncertainties. This could be seen as applicable to the FIRE facilities individually (projects) and to the European-level FIRE facility as a whole.

The main **certainties** acting as drivers affecting all FIRE future scenarios seem to be the following: Increasing pressure to federation, openness and cooperation of research infrastructures at European level; Increasing importance of FIRE service offer and meeting experimenter demands; European funding will continue to be the main source of funding. The main **uncertainties** affecting different aspects of FIRE sustainability include the level of EC funding of FIRE facilities on the longer term; FIRE's business attractiveness and prospects for business funding; and concrete prospects of collaboration with other initiatives like FI-PPP, EIT ICT-Labs, GENI, Géant.

<table>
<tr>
<td rowspan="2">Key partners
• How to attract the business community and national initiatives
• Economic downward trend</td>
<td>Key activities
• How to continuously improve experiment practice
• How to keep up with new technologies
• How to introduce and adopt new models for service offering</td>
<td rowspan="2">Value proposition
• Increasing pressure to openness, federation, cooperation across RI’s
• Service offer and responding to user demands is to become more critical
• How to keep up with the changing experimenter demands
• What are the new service models and how to implement them (TaaS etc)</td>
<td>Customer relationships
• How to implement demanding legal models
• How to introduce and implement professional governance model</td>
<td rowspan="2">Customer segments
• How to attract business communities and collaborate with national initiatives’ stakeholders</td>
</tr>
<tr>
<td>Key resources
• How to maintain and extend the main resources: the facilities, services, human capital, community networking</td>
<td>Channels
• How to modernize PR and promotion channels
• How to extend the Open Calls
• How to build Developer communities</td>
</tr>
<tr>
<td colspan="3">Cost structure
• How to decrease cost of experimentation and increase the scope of experiments
• How to manage cost by collaboration and synergies (sharing other research infrastructures and resources, joint research, development, experimentation)</td>
<td colspan="2">Revenue streams
• European funding remains the principal source now and in future, however the level of continuation is uncertain
• How to attract new sources of funding (industry, national funds), and how to introduce service-based pricing</td>
</tr>
</table>

Fig. 19 Driving Forces acting on FIRE Business Model and Sustainability

4.6 Concluding remarks

This chapter has performed an initial analysis of FIRE sustainability at two levels: 1) project level, 2) system level. The CANVAS framework provided a useful tool to systematically analyse the conditions for future sustainability at the levels. Sustainability is much more than the funding model only. It comprises key elements such as FIRE’s value proposition and service offering, and focuses attention to its future customer base – the “users” – and how to identify their needs.

5. Conclusions and Next Steps

5.1 Summary of conclusions

The main changes in the FIRE landscape and its environment that must be taken into account in developing the FIRE 2020 vision are 1) changes in FIRE's structure, 2) changes in FIRE's role within the wider Future Internet ecosystem, 3) changes related to the FIRE demand side, and 4) changes in Future Internet technologies, applications and infrastructures.

Changes in **FIRE's role within the Future Internet ecosystem** are the result of the emergence of a number of related initiatives such as EIT ICT Labs, Future Internet PPP, and initiatives regarding research infrastructures at both the national and European level (Eureka, Géant etc). Several platforms or communities creating research agendas have emerged around different thematic areas related to the Future Internet, such as Net!Works, NEM, NESSI, Photonics and others. As the Future Internet landscape in terms of strategic priorities and stakeholder platforms is changing, FIRE's role must reflect on new opportunities as regards collaboration and synergy creation. FIRE as a program needs a careful positioning in this emerging landscape.

As regards the **structure of FIRE**, our analysis demonstrates dramatic changes in FIRE throughout FP7 as a consequence of actions taken by the FIRE community and the EC. Increasingly FIRE's mission has become to deliver reusable facilities for the Future Internet community avoiding the propagation of testbeds within individual projects, resulting in the current emphasis on federation. In this context, FIRE will need to explore how to find a balance between coherence and fragmentation. On another level, FIRE is also giving rise to new instruments such as Open Calls, and interactions with other communities such as living labs and smart cities.

The **FIRE demand side** shows developments and changes as well, as is represented by the changes in experimenter demands and requirements, and the emergence of new types of service concepts (e.g. Testbed as a Service). This also affects the methods and tools, the channels to offering services to new categories of users, and the collaborations that must be established to deliver the services.

At the level of **Future Internet technologies and infrastructures and applications**, a number of key domains have emerged around Future Networks, Internet of Things, Internet of Services and other concepts which give rise to new research and innovation challenges. Our analysis shows several key trends, such as the integration of a broad range of systems (cloud services, wireless sensor networks, content platforms, mobile users) within Future Internet systems in large-scale, highly heterogeneous systems-of-systems.

These changes and developments bring with them lot of uncertainties in their outcomes and implications, which we have identified in some detail. Our analysis of drivers and uncertainties allowed us to explore how each has the potential to move FIRE towards different futures. The **key uncertainties** that we have identified shaping FIRE's future can be summarized along two dimensions: 1) the dimension of individual vs community: how will researchers and users of facilities collaborate in research, development and innovation of products and services? 2) Fragmentation vs cohesion: how will collaboration be structured and governed; will it be ad-hoc and largely unregulated or will it be organised and regulated? As a result, at least **four distinct scenarios of FIRE's future** emerged from our analysis:

1) **Testbed as a Service competition**: FIRE as a set of testbeds providing their facilities as a pay-per-use service.
2) **Industrial cooperative**: FIRE becomes a resource where experimental infrastructures (testbeds) and Future Internet services are provided by co-operating commercial and non-commercial stakeholders.
3) **Social Innovation ecosystem**: FIRE as a collection of heterogeneous, dynamic and flexible resources offering a broad range of facilities e.g. service-based infrastructures, network infrastructure, smart city testbeds, support to user centred living labs, and other.
4) **Resource sharing collaboration**: federated infrastructures provide the next generation of testbeds, integrating different types of infrastructures within a common architecture.

These four scenarios are different in their openness characteristics, business models, services and many more. They should be interpreted as inspirational images of possible FIRE Futures, not as predicted outcomes. Several elements of the scenarios, such as the services they deliver, may co-exist. Scenario analysis serves as a tool for creativity and thinking about the future. Still, these scenarios can be analysed in terms of their desirability, in terms of their business model needs, and in terms of the forces they represent and their impact on FIRE future. Given the scenarios, AmpliFIRE's over-arching vision of FIRE for the future is, in a summarized form:

"In 2020, Internet infrastructures, services and applications form the backbone of connected regional and urban innovation ecosystems. People, SMEs and organisations collaborate seamlessly across borders to experiment on novel technologies, services and business models to boost entrepreneurship and new ways of value creation."

We concluded this report by analysing the conditions for FIRE's sustainability using the CANVAS business model framework. The analysis of FIRE projects' sustainability led us to recognizing the forces and problems associated with realizing sustainability and exploitation strategies in some of FIRE's current projects. Evidently, at the level of FIRE as a system we need to take into account different business model conditions than at the project level. Key aspects of longer term sustainability include the structure of funding (will EC funding prevail as main source or will other sources of funding become necessary); the user base (will it be possible to target stakeholder groups such as smart cities, large companies involved in ICT, or – in collaboration with other initiatives – scientific organizations); the value proposition (what kind of services will be offered to what kind of user groups). This already shows that the sustainability discussion is strongly related to the FIRE strategy towards 2020.

5.2 Follow-up work

Following consultations with other members of the FIRE community and its EC directorate at FIA 2013 in Dublin (May 2013), we shall wrap up this first stage of developing our FIRE radar – a vision of FIRE potentials and the road ahead from 2015 to 2020. This will leave much to be done in AmpliFIRE – detailed roadmaps and assessments of the full set of FP7 investments and the first Horizon 2020 activities as well as the gap analysis that at this point we can only describe as a process to be followed. The second iteration of this Radar document will be created at the end of the AmpliFIRE project, mid 2015, incorporating the results of these further analyses, and benefiting from another year of discussion with members of the FIRE community. Several questions are already evident, and should be grounded in the material presented in this document. This chapter summarizes the main findings of this report and identifies some of the questions to be addressed in follow-up work.

5.3 Issues for further study

5.3.1 Understanding federation

One question that has already emerged from the ongoing discussions that we have had and the analytical framework for sustainability that we present here. Is federation of all the FIRE facilities under a common set of user identities, common logon, a common ontology that permits description of user privileges and requirements across all the FIRE testbeds, and a common set of experimental plane dispatch, control and analysis software possible. Is it desireable? To oversimplify the question, does one size fit all FIRE's customers?

The initial efforts towards growing sustainable testbeds within the FIRE community has taken different forms in the "academic" and quasi-"industrial" extremes, although the level of organization that we find in each already extends beyond the boundaries of individual institutions. Typically a small group of members of a present FIRE integrated project settles upon a "brand" under which they can continue to support their existing facilities, and present them both to customers and to sources of external funding so that they can evolve and grow in function. Will these nascent "brands" survive and grow separately, or will they merge, perhaps into an ultimate FIRE single federated testbed? We hope to know more by the end of AmpliFIRE.

5.3.2 Gap analysis

Gap analysis is a well-established methodology for building a plan to bridge from the realities of today to agreed-upon, potentially feasible goals for tomorrow. We will employ this approach in the remainder of the AmpliFIRE project. However, we must remember that the requirements for gap analysis are detailed reality-based assessments of the present capabilities of an organization and agreed-upon objectives to be reached within a specified time frame. The portfolio analysis that AmpliFIRE performs in WP2 is a requirement still to be done. This Radar is only a sketch of both today and the possible futures that we should strive for. There will be much work to be done.

5.3.3 Business model and FIRE sustainability

The FIRE vision is built upon scenarios that embody uncertainties. A better insight into the "business model" of these scenarios will contribute to a better assessment of the pros and cons of each scenario and will contribute to the understanding of the forces underlying these scenarios. Our analysis has started with assessing the sustainability and business model factors of individual currently running FIRE projects, then we took a wider view of FIRE as a system. Follow-up work in AmpliFIRE will be to elaborate our findings regarding business model analysis and sustainability in more detail, focusing on particular scenarios in validating the business model assumptions, and generalizing the findings into lessons learned and recommendations for future projects as well as for the FIRE program level.

References

FIRE project documents

FIRE STATION Support Action : Working Paper on FIRE Sustainability – FIRE Architecture Board. Dated 27 April 2011a.

FIRE STATION Support Action: FIRE in Horizon 2020. Position Paper and Annex from the FIRE Architecture Board, 2011b.

FIRE STATION Support Action, Deliverable D3.4 FIRE Roadmap. August 2011 (V1.2a).

FIRE STATION Support Action Deliverable D2.1: 1st FIRE Portfolio Update. 17th October 2011.

FIRE STATION Support Action Deliverable D2.2: 2nd FIRE Portfolio Update. 30 November 2012.

FIRE STATION Support Action Deliverable D3.6: FIRE Roadmap. 3rd May 2012, v1.0.

MyFIRE Support Action, Deliverable D4.1 Landscape Analysis Report, December 2010.

MyFIRE Support Action, Routes to Sustainability in FIRE. June 2012.

MyFIRE Support Action, Contribution to FIRE Roadmap, September 2012.

FIREWORKS Support Action, White Paper on Experimentally Driven Research. Version 1, 2010.

Other relevant project documents

FIREBALL Support Action, Deliverable D1.2: Common Assets Identification and Characterisation. April 2012.

FIREBALL Support Action, Deliverable D2.1: Landscape and Roadmap of Future Internet and Smart Cities. April 2012.

FIREBALL Support Action, Deliverable D1.3: Operational Infrastructure of IPR Arrangements for Access, Sharing and Reuse of Common Assets, 2012.

FIREBALL Support Action, White Paper on Smart Cities and the Future Internet, 2012.

OSIRIS Deliverable D4.1 Model for Public Authority — Research Infrastructure Interaction. Business models for ICT Research infrastructures (2012). http://www.ami-communities.eu/bscw/bscw.cgi/d905901/248295_OSIRIS_Deliverable_D4-1%20vs%2012-06-04.pdf

Workshop reports

The Role of ICT Infrastructures in Horizon 2020. Workshop, 20 June 2012. Summary of presentations and discussion. http://www.ictresearchinfrastructures.eu/communicraft-cms-system/uploads/Report_on_the_Event_20_june_2012.pdf

Websites

Website on European Web Entrepreneurship Strategy.
http://daa.ec.europa.eu/content/special/towards-european-strategy-web-entrepreneurs

Digital Agenda, Action 54: Develop a new generation of web-based applications and services
https://ec.europa.eu/digital-agenda/en/pillar-v-research-and-innovation/action-54-develop-new-generation-web-based-applications-and

CANVAS Business Model: http://en.wikipedia.org/wiki/Business_Model_Canvas

Relevant Documents beyond FIRE

NSF/EU Workshop on Future Directions in Pervasive computing and Social Networking for Emerging Applications.
http://www.ami-communities.eu/bscw/bscw.cgi/d905905/NSF_EUWSREPORT_Feb25.pdf

EGI Role towards Europe 2020.
http://www.ami-communities.eu/bscw/bscw.cgi/d905895/EGI%20role%20towards%20europe2020.pdf

FIA Research Roadmap towards 2020:
http://fisa.future-internet.eu/index.php/FIA_Research_Roadmap

Future Internet 2020. Call for action by High-Level Visionary Panel, 2009.

OECD, Building Blocks for Smart Networks. OECD Digital Economy Papers, 215. January, 2013.

Future Networks, The Way Ahead. European Commission, DG INFSO & Media, 2009.

Oxford Internet: Towards a Future Internet. Interrelations between technological, Social and Economic Trends. Report for DG Information Society and Media, 2010.

Rand Europe: Towards a competitive European Internet industry. A socio-economic analysis of the European Internet industry and the Future Internet Public-Private Partnership. FI3P Final Report, 2012.

Consultation on Future Network Technologies Research and Innovation in HORIZON 2020, 2012. http://cordis.europa.eu/fp7/ict/future-networks/documents/consultation-fn-h2020-report-final.pdf

The Future of Cloud Computing. Opportunities for European Cloud Computing Beyond 2010, 2010. http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-report-final.pdf

A Roadmap for Advanced Cloud Technologies under H2020, December 2012.
http://cordis.europa.eu/fp7/ict/ssai/docs/cloud-expert-group/roadmap-dec2012-vfinal.pdf

Workshop report: H2020: The challenge of providing cybersecurity.
http://cordis.europa.eu/fp7/ict/security/ws-report-cybersec-v2_en.pdf

Present and Future of Networked Media collaborative R&D in Europe.

http://nem-summit.eu/wp-content/plugins/alcyonis-event-agenda//files/Present-and-Future-of-Networked-Media-collaborative-R&D-in-Europe.pdf

Future Media Internet Research Challenges and the Road Ahead. April, 2010. http://cordis.europa.eu/fp7/ict/netmedia/fmi-research-challenges.pdf

SESERV Deliverable D2.2Final Report on Economic Future InternetCoordination Activities, August 2012.
http://www.scribd.com/doc/105908036/D2-2-v2-0-public

SESERV Deliverable Deliverable D3.2. Final Report on Social Future Internet Coordination Activities, August 2012.
http://www.scribd.com/doc/105908010/D3-1-2-v2-pdf

FInES 12, Research Roadmap 2025, March 2012. http://www.fines-cluster.eu/fines/wp/finesroadmap/

Report on the Consultation Workshop on the FP7 ICT Work Programme 2013: Networked Media Systems, January 2012.
http://ec.europa.eu/information_society/activities/netmedia/events/consultations/wp2013consult/docs/consultation-workshop-270112-final.pdf

Future Networks Research in FP7 and H2020, 2012. https://ripe64.ripe.net/presentations/225-RIPE64.pdf

Internet of Things in 2020. Roadmap for the Future. http://www.iot-visitthefuture.eu/fileadmin/documents/researchforeurope/270808_IoT_in_2020_Workshop_Report_V1-1.pdf

Internet of Things Strategic Research Roadmap, 2011. http://www.internet-of-things-research.eu/pdf/IoT_Cluster_Strategic_Research_Agenda_2011.pdf

IoT governance roadmap, June 2011. http://ec.europa.eu/transparency/regexpert/index.cfm?do=groupDetail.groupDetailDoc&id=7573&no=3

Tech transformation in the age of uncertainty, June 2011. http://www.slideshare.net/ictseserv/stephen-minton-tech-transformation-in-the-age-of-uncertainty-seserv-se-workshop-june-2012

Emerging trends in IT technology, 2012. http://www.slideshare.net/ChetanSagar1/emerging-trends-in-it-technology

Future and Emerging Technologies in H2020, 2012. http://www.fp7.cz/dokums_raw/microsoft-powerpoint-fiala-re236im-kompatibility_1351504070.pdf

Hola: Software and Services Interactive Workshop. https://docs.google.com/document/d/16s8P-HfyYAUIa9PiXYkmX40PJUSXqAzS7V4e9oyoCy4/edit?login=1&pli=1&overridemobile=true

Big Data, A New World of Opportunities, 2012. http://www.nessi-europe.com/Files/Private/NESSI_WhitePaper_BigData.pdf

NESSI Response to the European Cloud Strategy, 2012. http://www.nessi-europe.com/Files/Private/NESSI_Position_EuropeanCloudStrategy.pdf

Net!Works Horizon 2020 – The Framework Programme for Research and Innovation, 2012. http://www.networks-etp.eu/fileadmin/user_upload/Publications/Position_White_Papers/Horizon_2020_-_Detailed_comments_from_Net_Works_ETP_Final.pdf

Net!Works ICT - National Critical Infrastructure, 2012. http://www.networks-etp.eu/fileadmin/user_upload/Publications/Position_White_Papers/NetWorks_White_Paper_ICT_CritNatInfra_Final_12-07-02.pdf

Articles

[1] D. Papadimitriou, "Future internet: The cross-etp vision document," European Technology Platforms, 2009.

[2] V. Issarny, N. Georgantas, S. Hachem, A. Zarras, P. Vassiliadist, M. Autili, M. Gerosa and A. Hamida, "Service-oriented middleware for the future internet: state of the art and research directions," *Journal of Internet Services and Applications,* vol. 2, no. 2, p. 23–45, 2011.

[3] L. Atzori, A. Iera and G. Morabito, "The Internet of Things: A survey," *Computer Networks,* vol. 54, no. 15, pp. 1286-1389, 2010.

[4] R. Popper, "How are foresight methods selected?," *Foresight,* vol. 10, no. 6, p. 62–89, 2008.

[5] L. Youseff, M. Butrico and D. Da Silva, "Toward a unified ontology of cloud computing," in *Grid Computing Environments Workshop GCE '08*, 2008.

[6] Q. Zhang, L. Cheng and R. Boutaba, "Cloud computing: state-of-the-art and research challenges," *Journal of Internet Services and Applications,* vol. 1, no. 1, pp. 7-18, 2010.

[7] G. Blair, F. Kon, W. Cirne, D. Milojicic, R. Ramakrishnan, D. Reed and D. Da Silva, "Perspectives on cloud computing: interviews with five leading scientists from the cloud community," *Journal of Internet Services and Applications,* vol. 2, no. 1, pp. 3-9, 2011.

[8] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," *Communications of the ACM,* vol. 51, no. 1, p. 107–113, 2008.

[9] M. Schwarzkopf, D. Murray and S. Hand, "The seven deadly sins of cloud computing research," in *4th USENIX conference on Hot Topics in Cloud Ccomputing, HotCloud'12*, Boston, MA, 2012.

[10] D. Murray, M. Schwarzkopf, C. Smowton, S. Smith and S. Hand, "Ciel: a universal execution engine for distributed data-flow computing," in *8th USENIX conference on Networked systems design and implementation, NSDI'11*, 2011.

[11] S. Venkataraman, I. Roy, A. Young and R. Scheiber, "Using r for iterative and incremental processing," in *4th USENIX conference on Hot Topics in Cloud Ccomputing, HotCloud'12*, Boston, MA, 2012.

[12] J. Ahn, C. Kim, J. Han, Y. Choi and J. Huh, "Dynamic virtual machine scheduling in clouds for architectural shared resources," in *4th USENIX conference on Hot Topics in Cloud Computing, HotCloud'12*, Boston, MA, 2012.

[13] V. Jeyakumar, M. Alizadeh, D. Mazieres, B. Prabhaker and C. Kim, "EyeQ: Practical network performance isolation for the multi-tenant cloud," in *4th USENIX conference on Hot Topics in Cloud Computing - HotCloud'12*, Boston, MA, 2012.

[14] C. Raiciu, M. Ionescu and D. Niculsescu, "Opening up black box networks with

cloudtalk," in *4th USENIX conference on Hot Topics in Cloud Ccomputing, HotCloud'12*, Boston, MA, 2012.

[15] S. Buigel, S. Nurnberger, T. Poppelmann, A. Sadeghi and T. Schneider, "Amazonia: when elasticity snaps back," in *18th Conference on Computer and communications security, CCS '11*, 2011.

[16] A. Tanenbaum and M. Van Steen, Distributed systems - principles and paradigms, Pearson Education, 2007.

[17] N. Chohan, A. Gupta, C. Bunch, S. Sundaram and C. Krintz, "North by northwest: infrastructure agnostic and datastore agnostic live migration of private cloud platforms," in *4th USENIX conference on Hot Topics in Cloud Computing, HotCloud'12*, Boston, MA, 2012.

[18] G. Blair, M. Paolucci, P. Grace and N. Georgantas, "Interoperability in complex distributed systems," in *M. Bernardo and V. Issarny (Eds.)*, Springer, 2011, pp. 1-26.

[19] M. Langheinrich, "Privacy by design - principles of privacy-aware ubiquitous systems," in *Ubicomp, volume 2201 of Lecture Notes in Computer Science*, Springer, 2001, p. 273–291.

[20] Y. Yu, L. Rittle, V. Bhandari and J. LeBrun, "Melete: Supporting concurrent applications in wireless sensor networks," in *ACM SenSys*, 2006.

[21] I. Leontiadis, C. Efstratiou, C. Mascolo and J. Crowcroft, "Senshare: Transforming sensor networks into multi-application sensing infrastructures," in *EWSN, volume 7158 of Lecture Notes in Computer Science*, Springer, 2012, pp. 65-81.

[22] S. Madden, M. Franklin, J. Hellerstein and W. Hong, "Tinydb: an acquisitional query processing system for sensor networks," *ACM Transactions on Database Systems,* vol. 30, no. 1, p. 122–173, 2005.

[23] N. Sharma, D. Irwin, P. Shenoy and M. Zink, "Multisense: fine-grained multiplexing for steerable camera sensor networks," in *2nd annual ACM conference on Multimedia systems, MMSys '11*, 2011.

[24] A. Dearle and S. Dobson, "Mission-oriented middleware for sensor-driven scientific systems," *Journal of Internet Services and Applications,* p. 133–139, 2012.

[25] A. Boucher, D. Cameron and N. Jarvis, "Power to the people: dynamic energy management through communal cooperation," in *Designing Interactive Systems Conference, DIS '12*, 2012.

[26] M. Naphade, "Smart water pilot study report," IBM Research, 2011.

[27] M. Naphade, "Smarter electricity portal pilot study report," IBM Research, 2012.

[28] L. Capra and D. Quercia, "Middleware for social computing: a roadmap," *Journal of Internet Services and Applications,* vol. 3, no. 1, pp. 117-125, 2012.

[29] J. Manyika, M. Chui, B. Brown, J. Bughin, R. Dobbs, C. Roxburgh and A. Byers, "Big data: The next frontier for innovation, competition, and productivity," McKinsey Global Institute, 2011.

[30] M. Krohn, A. Yip, M. Brodsky, R. Morris and M. Walfish, "A World Wide Web Without Walls," in *6th ACM Workshop on Hot Topics in Networking (Hotnets)*, Atlanta, GA, 2007.

ERCIM News, 77, Future Internet Technology. April 2009.

A-M. Oostveen et al. Cross Disciplinary Lessons for the Future Internet. The Future Internet, Springer, 2012.

H. Schaffers, A. Sällström, M. Pallot, J. Hernandez-Munoz, R. Santoro, B. Trousse, "Integrating Living Labs with Future Internet Experimental Platforms for Co-creating Services within Smart Cities". Proceedings of the 17th International Conference on Concurrent Enterprising, ICE'2011, Aachen, Germany, June 2011.