Notes from the FIRE Panel Discussion at SAC Workshop in Turin

5th March, 2008

Introduction

Fabrizio Sestini, EC, summarised the philosophy of FIRE (experimentally-driven, long-term research), the projects and themes accepted in the 1st Call of the "Network of the Future" (Unit D) and the topics covered by projects under negotiation from the 1st FIRE Call. The latter include:

- Cross-layer /cooperative design
- Multi-hop communications
- Re-thinking the distribution of the intelligence within the network
- Content- and human-centric networking
- Autonomic communications
- Achieving energy-efficiencies
- Low-cost communications (social aspects)
- Virtualisation

Questions for the panel:

1. Do we need FIRE?
   „FIRE is a research environment for investigating and experimentally validating highly innovative and revolutionary ideas"

2. Span of the experimentally-driven long-term research:
   - Concrete examples – at which level?
   - How to bring in a real multi-disciplinary perspective?

3. What kind of testing is possible on the FIRE facility?:
   - Concrete examples (ie. Testing of Internet protocols/architectures, of service architectures, of middleware/applications, etc.
   - Openness of the facility

4. How to operate FIRE?
   - How to get industrial commitment and academic involvement

5. The BLED (Future Internet) conference

Q1: Do we need FIRE?

There was a clear consensus that we need the sort of activities that FIRE is helping to finance.

It is an innovative concept and is the only place in the EC Framework Programme 7 for longer-term, potentially disruptive research - with a focus on practical experimentation. The association between the long-term research and the "experimental facilities" (this term is preferred to "testbeds") is unique and needed.

There were many interesting proposals, though no clear "big picture" of the future Internet can be extrapolated from the set of projects under negotiation.

It is also possible that some of the proposals that failed the evaluation were more innovative than some of those that were accepted, but funding was limited (only 1 out of 6 proposals could be accepted) and proposals had to meet *all* the evaluation thresholds (not only S&T).

Q2: Span of the experimentally-driven long-term research:

- **Concrete examples – at which level?**
- **How to bring in a real multi-disciplinary perspective?**

Span:
Even with innovative paradigms, it is the case that many aspects (eg. interoperability), could be tested "in a room". FIRE should focus on experimenting with those features that need large-scale facilities.

However, large-scale facilities are expensive, difficult to put together and maintain, and those at the leading edge of modern technology become outdated very quickly (GENI will issue new compliance specifications every 2 years for nodes that want to participate). Due to the costs for construction and maintenance, the FIRE facilities must therefore be shared (also with projects from other EC Units).

Scalability may be an issue: a content distribution test network has been proven to be feasible with 500 people, but not 50'000. Often the limit will only be discovered during the course of the project.

Examples of technologies for the future Internet
The general consensus was that it is too early to identify the next *breakthrough* technology. Time is needed to see what develops from the accepted projects (which are, by definition, *long-term*) …. that have not even started yet. Usage trends are key. The future Internet will be driven by the wishes of users for (e.g.) new social networking applications, or for making life easier.

Although other research programmes may take a decision to focus on a particular topic, or area of inter-disciplinary research (eg. in Japan, the government has decided to finance bio-inspired communication networks), this approach was not favoured by the participants. It

was preferred to keep the scope broad ("let a thousand flowers grow") in order to have more chance of catching the case that will change the Internet.

The Internet is a "built thing", but has evolved to do many things we didn't know it would be used for. This confirms it is impossible to plan everything in advance. The recent evolution has been driven by the Web, P2P, Facebook.

Will there be more than one Internet? Are we already in this situation? – the GÉANT network (for researchers) has different characteristics (because it is massively over-provisioned) than the "regular Internet". Also, time-critical services (IPTV, VoIP) on the "regular Internet" are already prioritised over (for example) Web browsing in the access network and in the home.

Work is needed on inter-domain agreements (especially a common terminology relating how to interconnect).

More meshing, more distributed routing, smaller routers in more places, research on routing protocols.

Research on data-driven architectures.

Consider re-designing the Internet with P2P communication as the primary service. What would it look like?

Multidisciplinarity
FIRE should be open to test new applications, particularly those which impact on the network (eg. mobile, wireless).

A project in Germany is building a "100Gbit/s Internet". This will also impact on the higher layers, as it makes new applications possible.

Methodology for building the experimental facility:
The sequence is important ….. the research should come first, then design the experimental facilities ("bottom up").
Aim for small, cheap, fast deployments, with a time horizon of 2-3 years, or else reality will by-pass the research.

Start with a small set of existing testbeds and then expand them step-by-step with new functionalities (ie. federate). Make them useful quickly for users. Wireless, satellites, SAC, … extensions to the testbeds can be added later).

Build experimental facilities per technology, then:
- *federate* them as they become mature, and
- *delete* them as they become obsolete

The federation is challenging, since the testbed technologies will be homogeneous only at a certain level (physical, network, application, …). These levels will be different between various experimental facilities. However, the alternative (trying to define the future Internet architecture and technologies first) will result in it being obsolete before it is built.

NSF/GENI has also recognised this, and will also now exploit existing testbeds. Calls for the first new pieces of the GENI infrastructure have been issued.

The main added value of the different testbeds is through interconnecting them. Further value will be gained if the testbeds include both evolutionary and disruptive research.

The *management* of the federation is a big issue (legal aspects (eg. privacy), common procedures on all testbeds, managing the user access / time-slicing, archiving the measurements, reproducibility of the tests, …)

Q3: What kind of testing is possible on the FIRE facility?:

- **Concrete examples (ie. testing of Internet protocols/architectures, of service architectures, of middleware/applications, etc.**
- **Openness of the facility**

Is it a goal of FIRE to promote virtualisation? Virtualisation is a "hot topic" at the moment. It can be a way to "squeeze" multiple virtual routers into one physical router, thereby saving energy and operational/financial costs for organisations by sharing computing resources which are not physically on their own premises. What are the research topics for virtualisation? In order to ensure there is no interference between users of the same physical link bandwidth, QoS mechanisms are needed (VPNs).

Openness:
The facility should be as open as possible for projects to test upon (openness and commitment to interoperate should even be mandatory), but procedures have to be in place to prevent confidential data being released. One solution would be for projects to send scripts to be run on archived traces and to be sent back the results.

Q4: How to operate FIRE?

- **How to get industrial commitment and academic involvement**

There are 2 fundamental business approaches (FIRE follows the latter one):

- Improve what exists (more profitable in the short term - was likened to the approach of Huawei, Nokia, Ericsson, …)
- Develop something new (longer-term view - was likened to the approach of Cisco, Sun)

Operation:
Support and maintenance is important. For example, whilst PlanetLab claims to have over 800 nodes, in practice only 200 may be available at any one time.

Measurement and monitoring:
Access to archived Internet traffic measurement data is as valuable for predicting the future of the Internet as is (for example) astronomy and weather data. However, care has to be taken to preserve the privacy/anonymity of communications data. It is an interesting research direction to determine how to achieve this.

Measurement and monitoring allows to visualise the dynamics of the Internet in real time. The issues relate to: what to measure? where to measure? how to ensure that the measurement data does not influence the genuine Internet traffic? how to bring all the measurement data back to the place where it will be analysed?

Q5: The BLED (Future Internet) Conference

March 31st - April 2nd
The 2nd and 3rd days of this conference are planned for the discussions of issues such as FIRE.

Similar questions to the ones used for this panel session could be asked again, in order to try and collect further needs for experimentally-driven research.

More-radical ideas:
The EC could be more aggressive (eg. obtain a 3G licence exclusively for research, reserve DVB channels, satellite channels).

"Squeezy-phone".

Create a network of virtual routers then "move them around" according to traffic demand. Similarly, move content around dynamically.

Are we thinking of transferring anything different than "bits"?

Other comments:
"Real users" are difficult to incorporate into potentially disruptive networks – as they need reliable, available networks.

It is difficult to get good researchers to do experimental research, since they have to produce papers. Work on new paradigms and the subsequent measurement results may typically generate 2 or 3 papers over a 3-year timeframe; this is less than would be possible on more-theoretical topics.

EC projects are expected to end in success, but it can be argued that this means they are not doing leading-edge research (where it would be expected that 90% of projects might lead nowhere). NSF, on the other hand (and also recently the German National Science Council) allows for failure – this culture needs to be spread further across Europe.

The recognition of success or failure of a project might take many years to materialise (IP, PCs, WLAN, Multicast are all examples of technologies that took many years to be recognised as successful).

It is also reasonable to consider revisiting older research work that was abandoned, because new technology might now make the original ideas feasible.

GÉANT can not result interesting as a realistic test facility as it is always over-provisioned.

Technology can lead to changes in laws (ie. what may appear to be regulatory barriers today, may not be tomorrow … and vice-versa).

Notes kindly drafted by Martin Potts