## Minutes of the Workshop on Measurements and Measurements Tools

## Aalborg (Denmark) - May 9, 2012

**Number of attendees**

Up to 25 people attended the workshop, most of them to the majority of the presentations

**Changes in the workshop program**

The initial agenda was the following:

09:20 Welcome

09:30 Introduction and motivation of the workshop topic, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project

10:00 Session 1

Session chair: Davide Careglio (Universitat Politècnica de Catalunya - UPC, Spain), EULER project

*Experimental performance evaluation of sensor-based networking for energy efficiency in smart buildings*, Sotiris Nikoletseas (Research Academic Computer Technology Institute - RACTI, Greece), HOBNET project

10:30 Coffee break

11:00 Session 1 (continuation)

*Metrics and measurement tools for assessing the channel condition in a wireless experimentation environment*, Ingrid Moerman (Interdisciplinary Institute for Broadband Technology - IBBT, Belgium), CREW project

*Methods and tools for temporal Web analytics*, Marc Spaniol (Max-Planck-Institut für Informatik - MPI-INF, Germany), LAWA project

*Metrics and measurement tools needs in Information-Centric Networking and CONVERGENCE in particular*, Nicola Blefari Melazzi (University of Rome "Tor Vergata", Italy), CONVERGENCE project

12:30 Lunch break

13:30 Session 2

Session chair: Josep Lluís Marzo (Universitat de Girona - UdG, Spain), EULER project

*Metrics and measurement tools needs in OpenFlow and OFELIA in particular*, Hagen Woesner (European Center for Inf. & Communication Technologies - EICT, Germany), OFELIA project

*Delivers control and experimental plane middleware to facilitate early use of its large scale shared experimental facility by researchers*, Anastasius Gavras (Eurescom, Germany), OpenLab project

*Metrics and measurement tools for distributed and adaptive routing algorithms*, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project

14:50 Coffee break

15:20 Session 3

*Measurements and measurement tools in OpenLab*, Javier Aracil (Universidad Autónoma de Madrid - UAM, España), OpenLab project

*NOVI's experience in monitoring tools and measurements*, József Stéger (Eötvös Loránd University - ELTE, Hungary), NOVI project

*NITOS: Methods and measurement tools for experimentation on wireless testbeds*, Thanasis Korakis (Centre for Research and Technology Hellas - CERTH, Greece), CONECT project

16:00 Round table

17:00 Concluding session

17:15 End of the workshop

The final agenda of the workshop was different than the initial one, because unexpectedly Hagen Woesner couldn't come to the workshop. His talk was cancelled and several talks were moved backward and reordered. Moreover, some of the presentation's titles were changed. The final agenda was then the following:

09:20 Welcome

09:30 Introduction and motivation of the workshop topic, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project

10:00 Session 1

Session chair: Davide Careglio (Universitat Politècnica de Catalunya - UPC, Spain), EULER project

*Experimental performance evaluation of sensor-based networking for energy efficiency in smart buildings*, Sotiris Nikoletseas (Research Academic Computer Technology Institute - RACTI, Greece), HOBNET project

10:30 Coffee break

11:00 Session 1 (continuation)

*Metrics and measurement tools for estimating wireless channel condition*, Ingrid Moerman (Interdisciplinary Institute for Broadband Technology - IBBT, Belgium), CREW project

*Methods and tools for temporal Web analytics*, Marc Spaniol (Max-Planck-Institut für Informatik - MPI-INF, Germany), LAWA project

*Enhancing the Internet with a content-centric, publish-subscribe service model, based on a common container for any kind of digital data, including representations of people and Real World Objects*, Nicola Blefari Melazzi (University of Rome "Tor Vergata", Italy), CONVERGENCE project

12:30 Lunch break

13:30 Session 2

Session chair: Josep Lluís Marzo (Universitat de Girona - UdG, Spain), EULER project

*Heterogeneous testbeds, tools and experiments - Measurement requirements perspective*, Anastasius Gavras (Eurescom, Germany), OpenLab project

*Measurements and measurement tools in OpenLab*, Javier Aracil (Universidad Autónoma de Madrid - UAM, España), OpenLab project

*NOVI's experience in monitoring tools and measurements*, József Stéger (Eötvös Loránd University - ELTE, Hungary), NOVI project

14:50 Coffee break

15:20 Session 3

*Metrics and measurement tools for distributed and adaptive routing algorithms*, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project. Additional talks/demos: *Simulating routing models on large-scale topologies*, Aurélien Lancin (Institut National de Recherche en Informatique et en Automatique - INRIA, France); *Impact of power-law topology on IP-level routing dynamics: simulation results*, Amélie Medem Kuatse (Université Pierre Marie Curie - UPMC, France); *Distributed UDP Ping*, Elie Rotenberg (Université Pierre Marie Curie - UPMC, France).

*CONECT-NITOS: Methods and measurement tools for experimentation on wireless testbeds*, Thanasis Korakis (Centre for Research and Technology Hellas - CERTH, Greece), CONECT project

16:00 Round table, Javier Aracil, Thanasis Korakis, Davide Careglio, Nicola Blefari Melazzi, Dimitri Papadimitriou (moderator)

17:00 Concluding session

17:15 End of the workshop

***Introduction and motivation of the workshop topic*, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project**

Measurements and measurement tools will be key elements in the operation and management of future network infrastructures, at the equipment and network performance monitoring level but also in support of higher-level control functionality such as on-line analysis and diagnostic. These tools also play a fundamental role in measurement-based experimental research relying on the experimental evaluation and benchmarking of project outcomes including protocols, systems, etc., by means of reliable and verifiable tools.

The workshop aims at presenting current developments on measurements and measurement tools in research projects within FIRE (Future Internet Research and Experimentation) initiative of the EU 7th Framework Programme (FP7). Speakers from nine different FIRE projects will present their current needs and developments on measurements and associated tools in the context of experimental research in the areas of wireless/sensor networks, informatic-centric networking, large scale experimental facilities, OpenFlow-enabled networks and Internet. The expected outcomes of this workshop are: i) to identify what can be performed/reached by means of cooperation between projects from a directory of tools accessible to the FIRE community at large up to the joint development of tools, under which conditions, etc., and ii) to determine needs and document best practices in tools development for measurement-based experimental research.

An experimental methodology is an iterative process that includes the following steps:

- specification of the performance objectives, (technical and non-technical) constraints, and description of expected results, and the definition of relevant performance criteria and metrics.
- description of the modus operandi including configuration, initialization, and running conditions and (iterative) procedure(s) to be executed.
- reporting on observations and the resulting analysis, and the feedback on each iteration before reaching (partial) conclusion, and/or another iteration of the process.

In experimental ICT research measurements are distributed, which leads to the following challenges: tool properties and calibration; timing and synchronization; granularity, sampling and representativeness; data properties (no meta---info, timestamps); data analysis, reproducibility and referential. Moreover experimental measures should have the following properties: be reliable, repeatable, reproducible and verifiable.

***Experimental performance evaluation of sensor-based networking for energy efficiency in smart buildings*, Sotiris Nikoletseas (Research Academic Computer Technology Institute - RACTI, Greece), HOBNET project**

The main objective of the HOBNET (HOlistic Platform Design for Smart Buildings of the Future InterNET, http://www.hobnet-project.eu) project is to ease and maximize the use of experimental platforms by multidisciplinary developers of future Internet applications focused on automation and energy efficiency for smart/green buildings. The project's research addresses algorithmic, networking and application development aspects of future Internet systems of tiny embedded devices. HOBNET take a holistic approach addressing critical aspects at different layers (networks, algorithms, applications/tools) in an integrated way, including the following hierarchy:

- At the low level, network protocols and architectures, mainly based on IPv6, are studied, with an emphasis on heterogeneity and interoperability.
- At a second layer, we provide algorithmic models and solutions for smart buildings, with a special care for scalability.
- An interface layer for the rapid development and the evaluation of building management applications is provided at a third level.
- Finally, proposed research solutions and key innovations are organically evaluated in the context of the platform integration.

The main HOBNET testbeds are three: the CTI testbed (University of Patras, Greek), the UNIGE testbed (University of Geneva, Italy) and the MI building testbed (Mandat International Building in Bellevue, Geneva, Italy). These testbeds have a high diversity of devices (sensors actuators, etc.), different thematic emphasis (energy, tracking, visualization, etc.), and have been used to implement many scenarios (local adaptation to presence, $CO_2$ monitoring, garden watering, maintenance control, electric device monitoring, etc.).

Concerning algorithmic models and solutions for smart buildings, HOBNET has proposed two routing protocols, the distance-based $P_i$ protocol and the energy-based $E_i$ protocol, which have been implemented and experimentally evaluated in testbeds. A control application in a PC is responsible of detecting the sensor network, allowing the administrator to interact with them, and providing an interface for computing the required performance metrics. They also have implemented the so-called MeasurementsLogger component, whose role was to setup the parameters of the experiments (event generation rate, energy sampling rate, duration, etc.) and to monitor the evolution of the protocol by enabling the logging of the performance measurements. Performance measures were the delivery latency, the average energy consumption, and the success ratio.

Several critical options for the experimental scenario were detected: structured topologies vs randomized deployments; homogeneous sensor deployments vs heterogeneous deployments; all sensors at the start vs sensors added during network evolution; uniform node density vs high density diversity; static deployments vs mobile deployments (and hybrid combinations). Performance evaluation objectives were the following: scalability with respect network size (this implies the need for very large input sizes), fault-tolerance (this implies the need for diverse fault models), and inherent trade-offs such as energy vs time, and other.

As an example, a smart irrigation system was presented, able to adapt to specific watering needs (of different plants and of different areas) and diverse weather conditions. Implementation details and performance evaluation results were shown.

Finally the REST architecture was presented. It is an architecture for building network-based systems that allows for a better convergence and integration of collected measurements. Every distinguishable entity is a resource (anything from a physical device, like a sensor/actuator, to a Web site, XML file, etc.) uniquely identified by an URI. Below there is an all IPv6/6LoWPAN infrastructure, with the use of IPv6 to integrate heterogeneous technology (sensors, actuators, mobile devices etc).

Measurements based experimental research nicely complements rigorous performance analysis and simulation based evaluation. The results are more realistic and can contribute to validate and fine tuning the abstract algorithms. A large variety of realistic topologies, mobility profiles, traffic patterns is required. Novel network parameters as well as performance measures (and their trade-offs) arise. Ad hoc approaches are useful but there is a need to converge to widely accepted, common integrated approaches, systems and tools

***Metrics and measurement tools for estimating wireless channel condition*, Ingrid Moerman (Interdisciplinary Institute for Broadband Technology - IBBT, Belgium), CREW project**

The main objective of the CREW (Cognitive Radio Experimentation World, http://www.crew-project.eu) project is to establish an open federated test platform, facilitating experimentally-driven research on advanced spectrum sensing, cognitive radio, cognitive networking, and spectrum sharing in licensed and unlicensed bands.

Spectrum sensing is an important part of any cognitive radio solution and any wireless network experimentation. Some of the main issues here are the lack of (external) interference control and the lack of repeatability (due to unwanted external interference). Another challenge is the existence of heterogeneous wireless devices that provide different measurements, and therefore there is a need to find a way to characterize, compare and combine measurements from heterogeneous devices.

The CREW project proposes a benchmarking framework for spectrum sensing through heterogeneous devices composed of the following steps:

- Pre-calibration of the heterogeneous hardware, using a common metrics. The main metric is the Power Spectrum Density (PSD), with specific values of the span, resolution bandwidth, sweep time, time resolution, and energy detection threshold. Device heterogeneity induces different signal strength attenuation in the receiver chain, called power offsets, which have to be measured for later processing of the measured data (in the post processing stage)

- Set up the experiment, using metadata to unambiguously describe the experiment. Metadata would be the following: transmission signal pattern, transmission power level, device name, device location, device power offset, start time, frequency bins, resolution bandwidth, etc.
- Measure, using a clear methodology. The target would be the characterization and comparison of heterogeneous spectrum sensing devices. The approach would be to select a frequency band, configuring the transmission signal and to measure the PSD.
- Comparison, using a common data format, post processing and scores. Measurements would be in the form of a matrix containing PSD and relative time stamps. The post-processing would compensate for hardware heterogeneity, through pre-calibrated power offsets, and for software heterogeneity, through averaging and resampling the PSD matrix so that all devices have a common reporting rate in time domain. Then, through a given energy detection threshold, signal detection is determined. This threshold is obtained through the receiver operation characteristic plot that relates the Probability of False Alarm (PFA) and the Probability of Missed Detection (PMD). Finally, a “device performance score”, which indicates its performance, is derived.

Note that a “metric” here means a quantitative measure of a certain quality of the System Under Test (SUT), while “score” means an abstraction of a set of metrics (e.g., device performance, reliability of the experiment, cost, repeatability). Score hides performance evaluation details and is useful to compare a large number of experiments, to evaluate a solution by non-experts, and to automate the performance evaluation.

The experiment cycle in spectrum sensing through heterogeneous devices consists of a set of phases: warming-up, pre-assessment of wireless environment (create PSD map), in situ assessment of wireless environment (create PSD map), and post-assessment of wireless environment. The assessment of presence of external interference is made through analysis of pre/post/in situ PSD maps. The correlation between subsequent PSD maps from the repetition of the experiment cycle would provide a measure of the validity of the experiment (i.e., the impact of unwanted interference). Finally a “quality of experiment score” is derived.

***Methods and tools for temporal Web analytics*, Marc Spaniol (Max-Planck-Institut für Informatik - MPI-INF, Germany), LAWA project**

Web archives contain timestamped versions of Web sites over a long-term time horizon. This longitudinal dimension opens up great opportunities for sociological, political, business, and media analysts. For example, one could compare the notions of “online friends” and “social networks" as of today versus five or ten years back. Similar examples relevant for a business analyst or technology journalist could be about “tablet PC” or “online music”. This requires finding all Web pages from certain eras that contain these and/or other related phrases, disambiguate the entities contained and, finally, group and/or filter them accordingly. The main objective of the LAWA project is to build an Internet-based experimental testbed for large-scale data analytics, with a focus on developing a sustainable infrastructure, scalable methods, and easily usable software tools for aggregating, querying, and analyzing heterogeneous data at Internet scale. Particular emphasis is given to longitudinal data analysis along the time dimension for Web data that has been crawled over extended time periods. A Virtual Web Observatory will be created, to support data-intensive experimentation with Web content analytics. A demonstrator is planned which will allow to interactively browse, search, and explore born-digital content along the time dimension.

The LAWA methods and tools for temporal web analytics are the following:

- Distributed access to large scale data sets, based on wide area operations and heterogeneous distributed indices. The LAWA Reference Collection contains the UK web content dating back until 1998, including government sites, health portals, education, science, finance, and other. Concerning the collection management, the LAWA project uses the Hadoop framework, which provides an easy way to store the large amounts of data required for the project in a reliable fashion, and specifically the Hadoop Kelvin, a network monitoring system that monitors data (not control) traffic between Hadoop nodes and provides multiple ways to store, visualize and access the stored monitoring data.

- Distributed temporal Web analytics, including distributed web data aggregation, querying and ranking; entity detection, disambiguation and tracking; interesting phrase mining; text-entity-time analytics; community detection and analytics. AIDA (Accurate online Disambiguation of Named Entities) is a framework and online tool for entity detection and disambiguation, where given a natural-language text or a Web table, it maps mentions of ambiguous names onto canonical entities like people or places, registered in a knowledge base. YAGO is a huge semantic temporal knowledge base, derived from Wikipedia, WordNet, and GeoNames, which knows almost 10 million entities (e.g., persons, organizations, cities), and 120 million facts about these entities.
- Studies on Web-scale data: measurement, mining, and classification services; experiment-driven research on large-scale. The LAWA experimental testbed is a platform for analytic tools that will showcase a Virtual Web Observatory (VWO). The VWO implements a collection of core analysis solutions for the application stakeholders by developing methods and providing data that contribute to the FIRE experimental facilities. The VWO will support experimentally driven research by enabling cycle of experiments on large scale Internet data from the continuously growing reference collection. The initial configuration will be based on the following data, measurement facilities and tools: LAWA Reference Collection, Wide-Area Hadoop Operations (Hadoop Kelvin), temporal knowledge base (YAGO), named entity disambiguation services based on AIDA, service(s) for Web content classification.

***Enhancing the Internet with a content-centric, publish-subscribe service model, based on a common container for any kind of digital data, including representations of people and Real World Objects*, Nicola Blefari Melazzi (University of Rome "Tor Vergata", Italy), CONVERGENCE project**

Originally conceived as a "network of hosts", the Internet is evolving into an Internet of services, an Internet of media, an Internet of people and an Internet of "things". This implies a strategic shift from "host-centric" to "content-centric" and "data-centric" networking. Against this background, the CONVERGENCE (http://www.ict-convergence.eu) project proposes to enhance the Internet with a novel, content-centric, publish-subscribe service model, based on a common container for any kind of digital data:

- The "common container" of resources is the Versatile Digital Item (VDI), a standard–based (ISO-MPEG), self-contained, "all-inclusive" data unit. VDI is a container for any kind of digital data, including media, representations of people and virtual or physical objects (Real World Objects –RWOs-). VDI contains data (other VDIs, audio, images, video, text, descriptors of people, descriptors of RWOs, etc) and meta-data (information describing the content of the item; authentication and protection; rights to use the item; expiration date -supporting "digital forgetting"-). VDIs have a unique identifier.
- The main players in the CONVERGENCE framework are two: the publishers, which advertise their resources (both data and service-access-points) on the CONVERGENCE system; the subscribers, which express their interest in specific resources on the CONVERGENCE system. The CONVERGENCE system notifies subscribers when the resources they are interested in are available. The CONVERGENCE framework provides functionalities to create a VDI (defining related licenses and rights), publish a VDI, search and retrieve a VDI, subscribe to a VDI, verify the authenticity of a VDI, monitor the use of published VDIs, signs and/or encrypt a VDI, communicate with owners of VDIs, update VDIs, delete VDIs, etc.

The main functional levels of the CONVERGENCE system are:

- Applications and tools, e.g. the VDI Creator, VDI manager/browser
- Middleware, for powerful content identification and handling (publish VDIs, subscribe to VDIs, search, security, etc.)
- Information Centric Network (ICN), i.e., the networking functionality that allows storage and routing of VDIs through the network.

In ICN the network layer provides users with contents, instead of providing communication channels between hosts, and is aware of content identifiers. Its basic functions are: address contents, adopting an addressing scheme based on names (identifiers), which do not include references to their location; route a user request toward the closest copy of the content with such a name (name-based anycast routing); deliver back the content to the requesting host

Some of the advantages of the ICN approach are: an efficient content-routing, in-network caching, support for peer-to-peer like communications, per-content quality of service differentiation, handling of mobile and point-to-point multipoint communications, content-oriented security model, support for time/space-decoupled communications. Some of the disadvantages are: changes in the basic network operation, scalability concerns, cumbersome support for “conversational” communication patterns.

Convergence Network (CONET) protocols may be implemented in different places leading to different number of CONET subsystems (CSSs) in the network: only in user equipments (one single CSS, i.e., the current Internet); in current border gateways, i.e., where BGP runs (one CSS per AS); in all current routers (one CSS per IP subnet); in layer 2 devices (one CSS per layer 2 networks).

The chosen option in CONVERGENCE is to implement the network protocols in some of the current routers. Once an end node sends a contents request, border nodes would use a routing table based on content names to forward the request to the next node, till the serving node, which then would send the content back. Besides routing by name, border nodes would also do caching, and in the name-based routing table would also have entries related to the cached content. Therefore, routing tables would include IP routes, name-based routes and cached content. In case of a missing routing entry, a Name System provides the entry, which is temporary stored in the table. The Name system also handles best replica selection.

The CONVERGENCE project has performed measurements and developed measurements tools in the levels of applications, middleware, and network:

- Applications: professional photography, audiovisual archives, podcasts with synchronized slides, and retailing supply chain for electronic products.
- Middleware, to publish/subscribe functionalities, search and provide security mechanisms.
- Network, evaluation of scalability in routing and naming, the caching performance, native peer-to-peer support, native mobility support, quality of service differentiation, quality and network usage of video streaming.

***Heterogeneous testbeds, tools and experiments - Measurement requirements perspective*, Anastasius Gavras (Eurescom, Germany), OpenLab project**

The OpenLab Project (http://www.ict-openlab.eu/) aims at providing a large-scale network testbed enabling advanced experimentation of innovative applications, protocols and architectures in the area of the future Internet. Several experimental testbeds across Europe, showing heterogeneous transport and control technologies, are available under the OpenLab Project umbrella (e.g., NITOS, PlanetLab Europe, w-iLab.t, etc.). In this scenario, it is a key goal of the project to deliver a control and experimental plane middleware facilitating the combination of resources from these testbeds to allocate large-scale future Internet experiments.

The following are some examples of experiments being conducted in OpenLab:

- A mechanism to find the location map of the content servers in a Content Delivery Networks (CDNs), which provides a particular service in the actual Internet.
- Evaluation of a novel replica placement and virtual resource mapping framework for realizing a large-scale scenario of wireless CDN (i.e., CDNs where some clients use wireless access links).
- Evaluation of a novel architecture based on locator-identifier separation for supporting mobility and multihoming in the Internet.
- Evaluation of a new greedy geometric routing scheme for the Internet, where nodes are embedded into the hyperbolic metric space, which aims to solve the known problems of the Internet routing system (the

growth of the routing table size and the number of routing information exchanges produced by topological or policy changes).

OpenLab testbed allow researchers to evaluate and measure the performance of their proposals on a continental-wide experimental facility. In this presentation particular emphasis is put on the calls for experiments that the OpenLab Project opens, seeking for new partners (either industry or academics) responsible for conducting innovative experiments in the context of the future Internet, pushing the limits of the OpenLab infrastructure and exploiting its heterogeneity. While the 1st call for experiments was closed in Nov. 2011, a 2nd call is expected by the end of 2012- beginning 2013 (more info will be available at: http://www.ict-openlab.eu/open-calls.html). This may be a good opportunity for EULER project to test their Internet routing protocol prototypes on OpenLab, taking advantage of its technological features and creating synergies between both projects.

***Measurements and measurement tools in OpenLab*, Javier Aracil (Universidad Autónoma de Madrid - UAM, España), OpenLab project**

The work and ideas disclosed in this presentation have been also generated within the OpenLab Project framework (http://www.ict-openlab.eu/).

To start, the presentation highlights the problematic of obtaining measurements in a heterogeneous large-scale network testbeds like OpenLab, where multiple measurement tools, each one showing specific formats, may be involved in end-to-end performance evaluation. It is thus of high interest in such scenarios to come up with integrated and comprehensive measurements, so as to be reported to researchers using the experimental facility. The OpenLab approach for integrating information resulting from multiple management tools into a common user interface is then presented. The central entity in this procedure is the TopHat system entity, which collects the system (i.e., node) and network measurements reported by the different monitoring tools, and provides testbed users with a unified measurement view along the slice of resources that they have booked. Particularly, network measurements (obtained by measurement tools like ETOMIC/SONoMA, DIMES or Gulliver) are requested by TopHat to the TDMI agents (a measurement platform running in OpenLab network nodes). System and network statistics are completed with other information sources like geo-location data, IP-ASN or AS taxonomy. Finally, the presentation ends giving some information about standardization bodies efforts towards a common ontology for network measurements.

***NOVI's experience in monitoring tools and measurements*, József Stéger (Eötvös Loránd University - ELTE, Hungary), NOVI project**

The NOVI project (http://www.fp7-novi.eu/) investigates efficient approaches to compose virtualized resource slices spanning multiple heterogeneous infrastructures (cloud and network) in the context of the future Internet. Special attention will be put into the formal description and brokerage of virtualized resources within a federation of future Internet platforms, as well as on the performance monitoring of the provisioned virtual slices. Prototypes resulting from NOVI are expected to be deployed and validated on PlanetLab Europe and FEDERICA testbeds.

This presentation reports the functionalities and architecture of the monitoring service of NOVI. As seen in NOVI, the monitoring service has a twofold role, namely, helping in the network resource maintenance and provisioning procedures (e.g., assessing the available resources for a certain slice according to the network measurements and the constraints imposed by the end-user), as well as offering measurement and monitoring features to end-users so as to supervise the appropriate performance of the slices they book (NOVI end-users can book monitoring tools jointly with their slices, which can trigger alarms in case that certain performance metrics fall below a certain threshold). For metric measurements, both active and passive tools are employed (i.e., depending on whether they insert extra-traffic in the network or not), which report both system (CPU load, disk usage, memory usage etc.) and network (delay, bandwidth, loss etc…) performance metrics. The presentation then continues with the current status of the NOVI monitoring service implementation, as well as the information model proposed and adopted within NOVI.

Given the rich features of the NOVI monitoring service, it could be analyzed within EULER if an experimental evaluation of the protocol prototypes resulting from the project could lastly be performed in PlanetLab/Openlab, and thus take advantage of this service to validate the prototypes behavior. Going towards this possibility, it should be investigated if the measurement tools/metrics that the NOVI monitoring service controls/collects are really appropriate to effectively catch the behavior of the protocol prototypes developed within EULER.

***Metrics and measurement tools for distributed and adaptive routing algorithms*, Dimitri Papadimitriou (Alcatel-Lucent Bell Labs - ALBL, Belgium), EULER project. Additional talks/demos: *Simulating routing models on large-scale topologies, Aurélien Lancin (Institut National de Recherche en Informatique et en* Automatique - INRIA, France); *Impact of power-law topology on IP-level routing dynamics: simulation results*, Amélie Medem Kuatse (Université Pierre Marie Curie - UPMC, France); *Distributed UDP Ping*, Elie Rotenberg (Université Pierre Marie Curie - UPMC, France).**

The main objective of the EULER (Experimental UpdateLess Evolutive Routing, http://www.euler-fire-project.eu) project is to investigate new routing paradigms so as to design, develop, and validate experimentally a distributed and dynamic routing scheme suitable for the future Internet and its evolution.

The EULER project is confronted to the situation that there is no Internet scale facility where routing schemes can be experimented before being deployed. When the scale of the real environment becomes difficult to reproduce (such as the Internet) and the number of states and their dynamics are themselves difficult to model, a new approach shall be considered: simplify the experimental corpus using abstraction while maintaining its functionality, and reproduce significant environmental phenomena derived from measurements of the actual environment. Therefore the proposed methodology is the following:

- Develop measurement tools in order to derive representative conditions of the actual environment, i.e., the Internet topology and its dynamics
- Derive patterns for modeling topology dynamics
- Use these (dynamic) models for generating experimental scenarios in large scale routing model simulation and emulation.

Concerning the measurement tools, the EULER project has developed tools to measure the degree distribution, the dynamics of topology and events, and the stability of routing paths:

- Distributed UDP Ping, a tool for measuring the node degree. UDP Ping works as follows: the application sends a UDP packet towards an unallocated UDP port of a target router; the router sends back an ICMP (destination unreachable) packet through one of its interfaces; from the received packet the application finds out the IP address of the target router. Distributed UDP ping is made up of a large set of monitors located in different places that together hopefully obtain the whole set of IP addresses of the target router, and then estimates its number of interfaces. From the data obtained the goal is to estimate the degree distribution. It has been deployed and experimented on PlanetLab.
- Tracetree – radar, a tool for measuring topology changes, routing path dynamics, and load balancing. Tracetree is similar to traceroute, but it obtains a routing tree of IP paths from 1 source to many destinations. Radar performs periodic measures with Tracetree to observe the dynamics. Dynamics in routing paths come from physical topology changes, routing path dynamics, and load balancing. From the data obtained the goal is to provide a dynamics topology model that fits the observations. It has been also deployed and experimented on PlanetLab
- A tool to measure the stability of BGP routing paths. BGP routes are affected by two types of instability, policy-induced (conflicting policy interactions) and protocol-induced (path exploration), which may lead to non-deterministic unstable states and delayed BGP convergence. A set of metrics that compares the BGP routing table at periodic times measure the stability of its routes. A tool has been developed that parses

significant volumes of real BGP datasets from routers under the control of RouteViews and computes the stability metrics.

***CONECT-NITOS: Methods and measurement tools for experimentation on wireless testbeds*, Thanasis Korakis (Centre for Research and Technology Hellas - CERTH, Greece), CONECT project**

This presentation focuses on the measurement methods used in a wide range of experiments involving different wireless technologies. These experiments were carried out in the NITOS testbed. Nitos testbed includes WiFi nodes, USRP boards, sensors (such as cameras, temperature, humidity), WiMAX nodes, LTE and 3G femtocell components. The main objective of this testbed is to allow work at packet level at MAC layer, so you can analyze and implement new schemes at the MAC and application layer. The use of Linux and open source drivers in some nodes allows also working with the physical layer on subcarrier allocation and power control.

Access to resources takes place through the slice abstraction, isolated resource containers accessible by one or multiple users. Apart from the wireless network itself, the testbed is comprised of three discrete wired local networks: the control network, to log into the nodes via a server; the chassis network, to power nodes on or off; an OpenFlow-capable experimental network. The testbed also provides tools to assist experimenters in assessing the testbed's wireless environment properties and selecting an appropriate topology for their experiment.

To cope with the heterogeneous set of measurements in these experimentation efforts, NITOS use the OMF measurement framework adopted from the EU project Openlab. cOntrol and Management Framework (OMF) is focused to enable an efficient management of testbed resources providing a clear and easy way to define an experiment, execute it and collect its results (http://www.nicta.com.au/research/projects/tempo).

OMF Measurement Library (OML) is based on customizable measurement points inside applications running on the resources and provides a well-structured solution for capturing, processing, filtering, storing and visualizing measurements. It usually sets up an OML server responsible of gathering the measurements and storing them in a database, and OML clients capable of injecting measurements generated at measurement points (MP) inside applications into streams towards the OML server.

Extra features include OMF graph generator (displaying results) and availability of using proxy OML servers to cover disconnected parts of the same experiment (mobility suport).

The following are some examples of experiments conducted in this testbed:

- Design and evaluation of cooperative networks, which exploit different paths through the aid of possible relays that carry out the traffic. The goal is to increase throughput and minimize power consumption.
- Demonstration of an scenario where a vehicle equipped with sensors gathers measurements from its environment and communicates opportunistically with road-side units to forward the measurements to a centralized framework for their storage and analysis.

***Round table*, Javier Aracil, Thanasis Korakis, Davide Careglio, Nicola Blefari Melazzi, and Dimitri Papadimitriou (moderator)**

The moderator started the round table introducing the following issues/questions to the all the panelists:

- Determine needs, new domains of measurements (not currently addressed), and document best practices in tools development for measurement-based experimental research
- Which tools have been developed and applied outside the scope of FIRE experimental facilities: determine the missing elements for large-scale experiments on these facilities
- Which tools that can be combined (with potential extension) for larger experiments, e.g., multiple STREPs joining efforts and what would be missing to realize larger experiments
- Which tools that are mature enough to start a basis for re-use by other projects

- Identify what can be performed and reached by means of cooperation between projects from a directory of tools accessible to the FIRE community at large up to the joint development of tools, under which conditions, etc.

Everybody agreed that something has to be done with respect to the measurement tools, testbeds and experimental research in network technologies in order to join efforts. These are the main challenges:

- Bring people from different projects to put together their experience and set a common base or framework in order to start sharing tools.
- Testbeds are different (wired, wireless, different hardware, different set-ups, etc) and tools designed for one are likely to be incompatible with others. There is a need for well-defined interfaces of different tools with implementations adapted to different equipment, and some standardization efforts in the formats of the collected results (data) and, if possible, also in the control of the experiment (control).
- Need of code factorization. Well defined and well documented modules so everybody could use them independently (reuse).
- How can the huge number of tools that already exist be shared/reused? Scientists should create joint working groups involving different projects and identify common things that can be shared with other projects.

Some open questions:

- In order to start to put things in common, is it better to bring together scientists from many different projects altogether? Or is it better to start with small groups in the same areas (e.g. people from wireless projects, people from optical networks projects, etc), and after that, to start working inter-group searching for larger common things?
- Is it necessary to ask first to the FIRE community about their interests, the tools they need, and not to waste time working with tools that are not going to be used? Or is it better to start working with the tools that already exist and making them available in a repository?
- Should be forced a policy consisting of “a group or a project can make use of the repository of tools if and only if they contribute to that repository”?

Everybody agreed on the following:

- There is a need for advance in this topic now.
- The researchers’ community should be in charge, not the industry (otherwise it ends up with a proprietary thing).
- There should be some leader (person or group) with long term support and deep involved over time. Otherwise there is a danger of lack of commitment. There should also be some sort of gain / benefit for the leader (licensing?).
- All these tools and standards should be open-source.
- Concerning the implementation of the repository, it is better to start with a simple repository and then add improvements (bottom-up) than designing a powerful but empty thing (top-down). Even it could start with just providing data from experiments (data sets), so that other people could analyze other aspects and obtain more results.
- It is important that all these tools could also work / be adapted from the testbed scenario to the emulation scenario and to the simulation scenario. This will facilitate the comparison of results.
- It is not desirable to end up having ten mini-catalogues. Some effort should be done by the FIRE Station support action in order to give some umbrella and help with the establishment of a repository of tools at a European level that groups most of the projects.

These are some examples and experiences that could help:

- Netflow. It is an example of a bottom-up tool, which initially works, but with the lack of a common framework regarding syntax, semantics and metrics. Therefore there is a real danger that it ends up with a Cisco NetFlow, a Juniper NetFlow, etc., all of them incompatible with the others.

- Free T-Rex (www.free-t-rex.net). It is an example of a repository of tools where, after an initial effort, the leader left and this repository became orphan and not well maintained.
- Caida.org. It is an example of a successful repository of tools with available measurement data sets and tools. Its main characteristic is simplicity. Why a similar thing cannot be done in Europe?