Enabling Future Internet Research: The FEDERICA Case

Peter Szegedi, Trans-European Research and Education Networking Association

Jordi Ferrer Riera and Joan A. García-Espín, Fundació i2CAT

Markus Hidell, Peter Sjödin, and Pehr Söderman, KTH Royal Institute of Technology

Marco Ruffini and Donal O'Mahony, Trinity College Dublin

Andrea Bianco and Luca Giraudo, Politecnico di Torino

Miguel Ponce de Leon and Gemma Power, Waterford Institute of Technology

Cristina Cervelló-Pastor, Universitat Politècnica de Catalunya

Víctor López, Universidad Autónoma de Madrid

Susanne Naegele-Jackson, Friedrich-Alexander University of Erlangen-Nuremberg

ABSTRACT

The Internet, undoubtedly, is the most influential technical invention of the 20th century that affects and constantly changes all aspects of our day-to-day lives nowadays. Although it is hard to predict its long-term consequences, the potential future of the Internet definitely relies on future Internet research. Prior to every development and deployment project, an extensive and comprehensive research study must be performed in order to design, model, analyze, and evaluate all impacts of the new initiative on the existing environment. Taking the ever-growing size of the Internet and the increasing complexity of novel Internet-based applications and services into account, the evaluation and validation of new ideas cannot be effectively carried out over local test beds and small experimental networks. The gap which exists between the small-scale pilots in academic and research test beds and the real-size validations and actual deployments in production networks can be bridged by using virtual infrastructures. FEDERICA is one of the facilities, based on virtualization capabilities in both network and computing resources, which creates custom-made virtual environments and makes them available for Future Internet Researchers. This article provides a comprehensive overview of the state-of-the-art research projects that have been using the virtual infrastructure slices of FEDERICA in order to validate their research concepts, even when they are disruptive to the test bed's infrastructure, to obtain results in realistic network environments.

INTRODUCTION

The European Community co-funded project FEDERICA (Federated E-Infrastructure Dedicated to European Researchers Innovating in Computing Network Architectures) [1] supports the development of the future Internet by definition. The project consortium is an optimal mixture of research institutes, universities, and industrial partners in order to foster cross-sector collaboration. Although, the two-and-a-half-year project officially ended on 30 October 2010, after a four-month extension, the infrastructure and its services are still up and running thanks to the voluntary efforts of the European Research and Education Networking (NREN) organizations, which are actively participating in the operation of the infrastructure.

FEDERICA's architecture has been designed in such a way that allows users to request "virtual slices" of the infrastructure that are completely isolated from each other and behave exactly as if they were physical infrastructure from the point of view of the users' experiments. The major objective of the deployment of such an infrastructure is twofold. On one hand, FEDERICA allows extensive research on virtualization based architecture itself. Experiments on various resource virtualization platforms, virtual resource descriptions, monitoring and management procedures, virtual network slice compositions, isolations or federations can be performed. On the other hand, the virtual slices of FEDERICA enable multidisciplinary research on future Internet within the slices.

This article is organized as follows. After a brief summary of the FEDERICA virtualization

capable infrastructure and its services, we provide an overview of similar initiatives around the world emphasizing the importance of federated facilities. Then, we discuss the special features of FEDERICA that make the infrastructure unique (as of today) by way of supporting a wide variety of future Internet research activities, maybe leading to new paradigms. As FEDERICA is highly user-centric, we introduce the broad spectrum of its user community by looking at the actual users and use cases. The contribution of FEDERICA user experiments to the whole Internet research space is also illustrated in this article that finally concludes with the project's perspective for the future.

FEDERICA AS A SERVICE

FEDERICA project participants have designed and deployed a virtualization capable infrastructure substrate, including programmable high-end routers, multi-protocol switches, and PC-based virtualization capable nodes, on a pan-European footprint (Fig. 1). The physical network topology is composed of 13 sites, 4 core and 9 non-core nodes, connected by 19 point-to-point links. On top of this substrate, virtual infrastructure slices can be created, which realize whatever topologies are requested by the users. The facility is also connected to the public Internet, thereby enabling easy access from any location and type of connectivity (wireless or fixed line) [1].

Virtualization is the key having a profound influence in technology. In its network-oriented meaning, it corresponds to the creation of multiple virtual networks on top of a physical infrastructure. The network virtualization can be performed at most layers of the ISO/OSI stack, from the data link to the network layer. In its system-oriented sense the process is even faster and more intriguing, leading to multiple operating systems with varying tasks running on the same hardware platform.

The service architecture of FEDERICA follows the Infrastructure as a Service (IaaS) paradigm. IaaS, in principle, is the common delivery of hardware (e.g., server, storage, network), and associated software (e.g., operating systems virtualization technology, file systems) as a service. It is an evolution of traditional hosting that does not require any long term commitment and allows users to provision resources on demand. Amazon Web Services Elastic Compute Cloud (EC2) and Secure Storage Service (S3) are examples of commercial IaaS offerings [2]. FEDERICA has two major services:

- The provisioning of best effort or QoS IP slices with the option of both preconfigured and unconfigured resources (i.e., routing protocols, operation systems, QoS parameters)
- The provisioning of raw resources in terms of data pipes (i.e., native Ethernet links) and unconfigured resources (i.e., empty virtual machines, clear routing tables, no network protocols)

The current FEDERICA services may naturally evolve toward Platform as a Service (where various preconfigured images and network scenarios can be provided from a repository), and Application as a Service (where a set of applica-

Figure 1. FEDERICA infrastructure, the substrate topology.

tions can be preselected by users). This may eventually lead to a complete portfolio of cloud services, where all of these services share simplicity in access, typically through a web interface, simple scalability, and ease of use.

FEDERATED FACILITIES WORLDWIDE

FEDERICA does not operate in isolation. There are initiatives all over the world, which are of particular relevance and represent the natural set of peers for liaison and collaboration with FEDERICA.

United States — The Global Environment for Network Innovation (GENI) initiative [3] in the United States is an activity that aims at creating a unique virtual laboratory for at-scale networking experimentation. GENI is in its third phase, called spiral 3, of exploratory rapid prototyping with about 50 experiments and facilities operational. Work in spiral 3 should improve the integration, operations, tools, and documentation of GENI prototypes to lower the barriers of use. User projects have to collectively agree on a federation architecture to create the GENI environment spanning multiple facilities. The FEDERICA project has obtained excellent interactions with the GENI project office and some GENI experiments (e.g., ProtoGENI, GpENI) through regular meetings.

Japan — The AKARI Architecture Design Project in Japan [4] aims to implement a new generation network by 2015, developing a network architecture and creating a network design based on that architecture. The philosophy is to pursue an ideal solution by researching new net-

Figure 2. User segmentation results on the exploited unique features of FEDERICA.

work architectures from a clean slate without being impeded by existing constraints. Some prototyping projects have been started at the University of Tokyo in collaboration with the National Institute of Information and Communications Technology (NICT).

Europe — In Europe, the FIRE initiative [5] from the European Commission funds various infrastructure projects (e.g., Panlab, Wisebed, OneLab) with which FEDERICA already has strong collaboration. The federation of OneLab and FEDERICA facilities is an excellent example in terms of combining guaranteed and non-guaranteed resources. FEDERICA can offer resources to OneLab that have full access to the network layers and OneLab can provide access to a larger set of distributed computing resources. During the last years, OneLab has developed an architecture to federate different domains based on PlanetLab that is called Slice Federation Architecture (SFA) [6]. SFA enables FEDERICA and similar infrastructures to become a worldwide federated facility supporting collective on future Internet research.

UNIQUE FEATURES OF FEDERICA

FEDERICA can be easily compared with PlanetLab [7]. PlanetLab started in Princeton, New Jersey, and is a well-known global research facility based on a large set of end hosts on top of the public Internet. The hosts are running a common software package that includes a Linux-based operating system. The key objective of this software package is to support distributed virtualization (i.e., the ability to allocate a slice of PlanetLab network-wide hardware resources to a user application). In contrary, a FEDERICA node is not only considered as an end host as in PlanetLab, but a combined network and computing element that itself routes traffic and affects network reachability. In addition, the user is not forced to use a specific operating system or

application but it is free to install any software component within the slice. While PlanetLab connections are based on the public Internet, FEDERICA allows the user to perform experiments on the lower layers as well.

FEDERICA architecture is more than an evolution of the concept of a traditional test bed. Test beds are usually devoted to a few technologies and oriented to production service or to pure research, and are not flexible enough to accommodate both types of use. The FEDERICA environment is also different from a commercial cloud environment, where the delay between nodes is typically negligible. In FEDERICA the delay corresponds to physical distance of the substrate nodes. Such naturally distributed architecture provides an ideal environment for testing medium size experiments and migration paths in a realistic environment.

The unique features of FEDERICA (as of today) can be summarized as follows:

- FEDERICA has its own physical substrate fully controlled and managed by the FEDERICA Network Operation Center. This allows the user to take over control of the lower layer resources within their slices. At the moment, only raw Ethernet resources are available for the users but in the near future native optical resources (i.e., lambdas) may also be provided.

- Thanks to having full control over the substrate links, specific QoS parameters can be assured for the virtual connection of the user slices and real transmission delays can be experienced. Currently the links run up to 1 Gb/s but this may be upgraded radically in the near future.

- FEDERICA users have full flexibility and freedom to choose any networking protocol or operating system to be installed on their virtual nodes. To the virtualization platform end, JUNOS-based programmable high-end Juniper platforms and VMware-based SUN servers are both currently deployed in the FEDERICA substrate.

- FEDERICA can ensure the reproducibility of the complete testing environment and conditions of the user experiments at a different location or time. Repeatability of the experiments can also be ensured in the sense of obtaining the same results given the same initial conditions at any time.

- The overall architecture is federation-ready in line with the Slice Federation Architecture concept. Moreover, FEDERICA provides, for example, non-web-based federated SSH access to all of its resources supported by the Single-Sign-On infrastructure of the research and education community.

However, it is important to mention some of the limitations of FEDERICA, too. Typically, compared to PlanetLab or commercial clouds, the scalability of the FEDERICA virtual infrastructure is limited. The scalability is the function of the given size of the physical substrate (that may be extended in the future) and the number of active user experiments requiring QoS assurances on links and/or computing nodes. As a consequence of this, user access to the FEDERICA slices must be governed with care. This role is undertaken by the User Policy Board of the project, which collects and approves (or rejects) user slice applications.

The exploitation of the aforementioned unique features of FEDERICA is illustrated in this article by a wide range of user projects.

USERS AND USE CASES

During the active lifetime of the FEDERICA project (2008–2010) the consortium partners approached and consulted more than 40 potential user groups all over Europe that resulted in 15 slice requests for specific user experiments, as of 30 October, 2010. The segmentation and analysis of the user community have been considered to be important to FEDERICA. The aim of this segmentation was to understand the user community better and the type of experiments that can exploit the unique features of FEDERICA. This should help to steer the future infrastructure development and user consultation activities in the right direction so as to support the needs of the community better.

USER SEGMENTATION

The detailed results of the FEDERICA user segmentation are available in the public project deliverables [8]. In the following section, we highlight the main characteristics of the user community.

Analyzing both the slice requests already submitted and the preliminary requirements collected from the possible user groups, Fig. 2 shows that 40 percent of the performed experiments in FEDERICA use a slice as a fully configured IP network with best effort connections. There is nothing unique in that sense, so other motivations (e.g., cost, economy of scale, simplicity of use) must exist in the case of that 40 percent. In contrast, 60 percent of the experiments performed use FEDERICA because of its unique features (e.g., IP quality of service assurance or unconfigured resource provisioning). In particular, 26.6 percent of users have requested raw resources and data pipes (i.e., IaaS) in order to fully configure and manage their slices. Analyzing the potential broader user base, we expect a few more user experiments in the future exploiting the above mentioned unique features of FEDERICA.

MAJOR EXPERIMENTS

FEDERICA's mission is to support the development of the future Internet via Future Internet Research projects. To measure its success, FEDERICA has collected feedback from users in order to understand, say, the contribution of the experiments' results to the ICT research community as a whole.

Figure 3 shows the results of the users' feedback on this particular aspect. Sixty percent of the experiments performed in FEDERICA contribute to international research projects partly funded by the European Commission. More than 10 percent of experiments are aimed at directly or indirectly contributing to standardization activities in the field of networking. Requests from the commercial or private sector has not yet been received during the active lifetime of the project, but some discussions within the community suggest that it could be possible in the future.

Figure 3. Contribution of FEDERICA experiments to ICT research community.

CONTRIBUTION TO FUTURE INTERNET RESEARCH

In the following section, we illustrate the usage of FEDERICA via a wide spectrum of its user experiments. The examples of user experiments are grouped in three categories according to their main objectives; validation of virtual infrastructure features, evaluation of multilayer network architectures, and design of novel data and control plane protocols.

VALIDATION OF VIRTUAL INFRASTRUCTURE FEATURES

This group of experiments aims at validating the basic principles of virtualization capable infrastructures in general and particularly the unique features of FEDERICA.

The Universitat Politècnica de Catalunya (UPC), Spain, requested a FEDERICA slice, called ISOLDA, in order to perform basic network performance and slice isolation tests. Parameters such as bandwidth, latency, jitter, and packet loss were measured in parallel slices with the aim to prove sufficient isolation between them. Figure 4 depicts the test scenario where two virtual connections (red and dashed green slices) share a physical port on the FEDERICA substrate node in the middle. This shared physical link contains two different VLANs, one for each virtual connection. VM1 and VM3 are located on the same virtual machine server, called VMServer 1, while VM2 and VM4 are located on another server, called VMServer 2. The other virtual connection (dotted blue line) does not share any resource.

The isolation test at network level was done by using various control and management protocols to identify whether the isolated virtual machines can accidentally connect to each another. Isolation of virtual machines was also proven by ping tests. The results from these tests were affirmative, as the virtual machines on different slices were not visible to one another. In conclusion, information was obtained about how the FEDERICA substrate nodes should be configured to assure the isolation feature between

Figure 4. *ISOLDA: Network performance and slice isolation tests.*

slices that share physical resources. Following these defined procedures in FEDERICA, the isolation between slices is completely ensured.

KTH Royal Institute of Technology, Sweden, requested a slice, called METER, in order to study another important feature of FEDERICA: the repeatability of its experiments. This user project dealt with problems related to repeated experiments on a shared network, where other external activities may influence the set of results. KTH investigated a method to identify time periods of comparable network conditions based on metadata-contextual information about the environment where the experiment was executed.

During the timeframe of an experiment, active background measurements were run to collect metadata parameters. Experiment data was time-stamped and KTH used statistical analysis on the metadata to determine if experiment data was gathered during comparable network conditions. An important goal was to be able to do this without interrogating the experiment data itself. Within the slice, service degradation in terms of resource competition with experimental activities in other slices was rarely detected. An important part of the experimental work was to detect differences in network conditions e.g., in terms of latency. A set of background measurements was used to identify periods of comparable network conditions. Based on the measurement results, it was possible to reduce the confidence intervals of the outcome of the repeated experiment, so the remaining values were suitable for comparison.

The Friedrich-Alexander University of Erlangen-Nuremberg (FAU), Germany, in cooperation with FEDERICA partner Deutsche Forschungsnetz (DFN), requested a slice in order to perform Hades Active Delay Evaluation System (HADES) measurements over the physical infrastructure of FEDERICA. The nodes of the FEDERICA substrate were measured in a full mesh topology; with bidirectional measurements, this yielded 42 measured links.

HADES was developed and deployed as a tool that provides performance measurements and offers IP performance metrics such as one-way delay, one-way delay variation, and packet loss [9]. The main purpose of the experiment was to collect and archive these IP performance metrics over an extended period of time and make the data available to other FEDERICA project partners and experiments that needed reference data as part of their investigations on the behavior of virtualized networks and slice processing.

The data was also used to show that FEDERICA users had a stable network environment and repeatable conditions for their experiments.

EVALUATION OF MULTI-LAYER NETWORK ARCHITECTURES

This group of experiments illustrates the capability of connecting external test beds or virtual slices provided by other facilities to the FEDERICA user slice. In this way, multi-layer and multi-domain network architectures and federation principles can be evaluated.

Trinity College Dublin (TCD), Ireland, requested a slice, called OIS, in order to carry out experiments on the Optical IP Switching (OIS) concept, a network architecture [10] that was developed by the CTVR Telecommunications Research Centre based at the University of Dublin. The idea behind OIS is to build a packet router that can analyze traffic flows and use an underlying optical cross-connect to bypass the router electronics for suitable IP flows. The FEDERICA slice was used to analyze how dynamic creation, modification or cancellation of optical paths can degrade the quality of applications transported over TCP protocols, on a large-scale testbed implementation. The results presented substantial differences to those previously obtained through laboratory-based experiments.

The OIS network architecture was setup in TCD's laboratory as a core network and interconnected with vanilla IP domains, which were realized within the FEDERICA slice. In the experiment the FEDERICA slice was organized in such a way that emulated a simple access network, where user data were aggregated by an IP router and sent toward the TCD testbed. Here packets were routed back to FEDERICA towards the destination network. TCD could then verify how their testbed would work when connected to external legacy IP networks. The results showed, for example, that UDP and TCP transmissions are impaired when the signals are dynamically switched from an electronic to an optical connection if both upstream and downstream links are congested. The experiment also gave TCD the opportunity to evaluate feasibility and efforts required to use virtualized networks in combination with physical testbeds.

The Telecommunications Software & Systems Group (TSSG), based in Waterford, Ireland, was one of the 10 partners of the European Commission funded project PERIMETER [11] which is finished in May 2011. The project addressed challenges such as quality of experience (QoE), security, and end user impact in order to establish a new paradigm for seamless mobility. TSSG requested a FEDERICA slice (incorporating five virtual nodes and routers) to extend the existing physical test bed and carry out more robust testing of the PERIMETER application particularly in the areas of routing, scalability, and authentication. This was a relatively inexpensive way to dynamically pool and use resources for experimentation of new networking concepts. The federated PERIMETER testbed and FEDERICA slice were used to conduct a number of testing processes including scenario conformance tests, interoperability tests, application tests, Living

Labs tests, and performance and scalability tests in the project. These testing processes could not be performed while achieving such a level of realism without the use of FEDERICA. In conclusion, it was validated that the federation of a physical test bed and a virtual slice allows experimental testing to proceed, using real-size platforms, real operating systems, and real applications in a realistic environment, which is similar to the actual target system.

The Universidad Autónoma de Madrid (UAM), in Spain also requested a FEDERICA slice, called ANFORA, in order to interconnect that with two other infrastructures. This complex networking experiment used resources of three facilities: FEDERICA, OneLab, and PASITO. Beside the FEDERICA slice, UAM already had OneLab nodes with specific monitoring cards in order to perform quality of service measurement and was also connected to PASITO, the Spanish national layer 2 networking facility [12], to provide lower-layer connections at the FEDERICA points of presence. The multilayer test facility that was created (Fig. 5) was used to evaluate the Path Computation Element (PCE) protocol, and to validate multilayer traffic engineering (MTE) algorithms.

The FEDERICA slice allowed the creation of a virtual network topology, which played the role of an IP service provider's backbone network. OneLab measurement tools, implemented on top of FEDERICA virtual nodes, ensured precise monitoring of traffic engineering (TE) parameters, like bandwidth or delay. The end-to-end monitoring information was sent to the PCE database so, PCE nodes knew the actual performance of each layer.

During the experiment, end-user traffic was routed using the IP layer (i.e., FEDERICA slice) by default, while the network performance was adequate. However, when there was IP congestion, the lower layer connections of PASITO infrastructure were used to provide the desired quality of service for the end users. PCE computed the end-to-end routes and decided on which network layer the traffic should be transmitted. As a result of the experiment it was concluded that, PCE is suitable for multilayer algorithms since their computation is more complex than traditional routing algorithms. Moreover, the PCE database was filled with information of both layers providing a global view of the network to the operator. The interconnection of FEDERICA, OneLab, and PASITO provided a unique test environment in order to successfully validate various MTE algorithms.

DESIGN OF NOVEL DATA AND CONTROL PLANE PROTOCOLS

This group of experiments aims at designing and validating novel data and control plane protocols as well as architectures.

The Politecnico di Torino (PoliTo), Italy, requested a slice, called VMSR, in order to experiment with the Virtual Multistage Software Router (VMSR) distributed architecture. Distributed routers are logical devices whose functionalities are distributed on multiple internal elements, running on virtual machines, in order to achieve larger aggregated throughput and improved reliability. The VMSR multistage router architecture [13] was composed by three

Figure 5. ANFORA: investigation of federation and multilayer network scenarios.

internal stages: layer-2 load balancers, Ethernet switches to interconnect components, and standard Linux based software routers. In order to coordinate the internal elements and to allow VMSR to behave externally as a single router, custom-made control and management protocols were designed by PoliTo.

The VMSR experiment in FEDERICA consisted of three L2 balancers (Linux virtual machines running Click on Ethernet Virtual Switch), one Ethernet switch, three L3 routers (i.e., standard PCs running the Linux protocol stack and XORP), and three external nodes that generate and sink traffic. The main advantage of implementing the VMSR experiment in FEDERICA was related to the ability of controlling and monitoring network resources allocation inside the architecture, which would not be possible in a testbed (e.g., on top of the public Internet). Functional tests on the control and management plane protocols were successfully performed. PoliTo measured throughput and latencies on all the involved links while varying the packet size of traffic generators to determine the suitability of the infrastructure supporting high-performance applications.

Last but not least, Fundació i2CAT, the Spanish research foundation, requested a slice in order to perform scalability tests of the Harmony multi-domain and multi-vendor network resource brokering system developed by the European Commission funded project PHOSPHORUS [14]. Harmony defines the architecture for a network service layer between the grid middleware and applications and the network resource provisioning systems (NRPS). Harmony architecture has evolved from a centralized network service plane (NSP) model to a distributed NSP model, passing through a middle stage, the multilevel hierarchical NSP model, as depicted in Fig. 6. Moreover, hybrid architectures of the NSP have also been under consideration for huge high-capacity environments.

The experiment performed over the FEDERICA slice aimed to analyze the performance and scalability of the different Harmony NSP topologies under different workloads. The experiment focused on collecting information about two parameters: the service provisioning time and the average rate of blocking requests. The slice

The PCE database was filled with information of both layers providing a global view of the network to the operator. The interconnection of FEDERICA, OneLab and PASITO provided a unique test environment in order to successfully validate various MTE algorithms.

Figure 6. PHOSPHORUS: Harmony network resource brokering system scalability study.

provisioned by FEDERICA consisted of computational resources (15 virtual machines with 2 Gbytes RAM and 10 Gbytes storage each), where the Harmony NSP entities were deployed.

The performed test cases produced measurements that helped i2CAT to better understand and characterize the performance and scalability of the developed system under higher loads than the ones used in the physical test bed available in the PHOSPHORUS project. Several new results were obtained, such as the comparison of the performance of each of the feasible architectures of the NSP. Moreover, the PHOSPHORUS project participants obtained service provisioning times of very different situations and loads for each test scenario through reconfiguring the FEDERICA slice.

CONCLUSION

In this article we have given a brief overview of the unique features and services of the FEDERICA virtualization capable infrastructure, developed and deployed in the framework of a European Commission funded project. The broad spectrum of Future Internet Research experiments, performed over the virtual slices of FEDERICA, is also illustrated above. Based on the feedback of the user community, it is clearly signaled that the FEDERICA facility is essential for sustainable support of Future Internet Research in Europe and beyond. That is why the European NREN partners of the former project consortium have agreed to continuously support the operation of the infrastructure and are willing to do further development under the recently launched Network Factory task of the GN3 project's Joint Research Activity 2 [15].

ACKNOWLEDGMENT

The FP7 project FEDERICA was partially supported by the European Commission under Grant Agreement no. RI-213107. The authors

acknowledge the contribution of all project partners, user projects, and specifically for their input, Mauro Campanella (GARR, project coordinator of FEDERICA), Peter Kaufmann (DFN), Vasilis Maglaris (NTUA), Frances Cleary, and Eileen Dillon (TSSG).

REFERENCES

- [1] P. Szegedi et al., "With Evolution for Revolution: Managing FEDERICA for Future Internet Research," *IEEE Commun. Mag.*, vol. 47, no. 7, July 2009, pp. 34–39.
- [2] S. Bhardwaj, L. Jain, and S. Jain, "Cloud Computing: A Study of Infrastructure of a Service (IaaS)," *Int'l. J. Eng. and Info. Tech.*, vol. 2, no. 1, Feb. 2010, pp. 60–63.
- [3] J. B. Evans and D. E. Ackers, "Overview of GENI Overview of GENI and Future Internet in the US," 22 May 2007, <http://www.geni.net/>
- [4] H. Harai, "AKARI Architecture Design Project," *2nd Japan EU Symp. New-Generation Network and Future Internet*, Oct. 13, 2009, <http://www.akari-project.jp/>
- [5] FIRE WHITE PAPER, 2009, "FIRE: Future Internet Research and Experimentation," Aug. 2009, <http://www.ict-fire.eu/home/publications/papers.html>
- [6] S. Soltesz et al., "On the Design and Evolution of an Architecture for Federation," *ROADS2007*, July 2007, Warsaw, Poland.
- [7] L. Peterson, "PlanetLab: A Blueprint for Introducing Disruptive Technology into the Internet," Jan. 2004, <http://www.planet-lab.org/>
- [8] FP7-FEDERICA, Deliverable DNA2.3 "FEDERICA Usage Reports," Nov. 2010, <http://www.fp7-federica.eu/>
- [9] P. Hollecsek et al., "Statistical Characteristics of Active IP One Way Delay Measurements," *Int'l. Conf. Net. and Services (ICNS'06)*, Silicon Valley, CA, 16–18 July 2006.
- [10] M. Ruffini, D. O'Mahony, and L. Doyle, "Optical IP Switching: A Flow-Based Approach to Distributed Cross-Layer Provisioning," *IEEE/OSA J. Opt. Commun. and Net.*, vol. 2, no. 8, Aug. 2010, pp. 609–24.
- [11] E. Dillon et al., "PERIMETER: A Quality of Experience Framework," *Int'l. Wksp. Future Internet of Things and Services*, 1 Sept. 2009, Berlin, Germany.
- [12] Platform for Telecommunications Services Analysis (PASITO), <http://www.hpcn.es/projects/pasito/>
- [13] A. Bianco et al., "Multistage Switching Architectures for Software Routers," *IEEE Network*, vol. 21, no. 4, July 2007, pp. 15–21.
- [14] S. Figuerola et al., "PHOSPHORUS: A Single-Step On-Demand Services Across Multi-Domain Networks for e-Science," *Network Architectures, Management, and Applications, Proc. SPIE*, vol. 6784, 2007.

BIOGRAPHIES

PETER SZEGEDI (szegedi@terena.org) received his M.Sc. degree in electrical Engineering from Budapest University of Technology and Economics (Hungary, 2002). He then worked toward a Ph.D. in the Department of Telecommunications. His main research interests include design and analysis of dynamic optical networks, especially optical Ethernet architectures, network virtualization, control, and management processes. He worked for Magyar Telekom (2003–2007) and then joined TERENA in January 2008.

MARCO RUFFINI (marco.ruffini@tcd.ie) is currently a lecturer on optical network architectures at the Department of Computer Science of the University of Dublin, Trinity College, where he obtained his Ph.D. in 2007. He is part of the CTVR telecommunication research centre, and his research interests include low-based optical switching, experimental optical testbeds, and techno-economic studies of next-generation transparent architectures. He worked for Philips Research Laboratories in Aachen (2003–2005), and before that as a technology consultant for Accenture, Milan. He holds a degree in electronic engineering from Università Politecnica delle Marche in Italy (2002).

DONAL O'MAHONY (donal.omahony@cs.tcd.ie) graduated with first class honors in engineering from Trinity College in 1982. After a brief career in industry at Sord Computer Systems in Tokyo and IBM in Dublin, he rejoined Trinity College as a lecturer in computer science in 1984, completing his Ph.D. in the area of software reusability in 1990. At Trinity, he built up a successful research group in networks and telecommunications. He spent the year of 1999 as a Fulbright Fellow at Stanford University, California, before returning to his present position as professor in computer science at Trinity College. In July 2004 he led a team to establish CTVR, a major multi-university research centre established in association with Bell Labs.

JORDI FERRER RIERA (jordi.ferrer@i2cat.net), M.Sc., graduated in computer science from the Technical University of Catalonia (FIB-UPC). He joined the i2CAT foundation as a software engineer in early 2007 for developing the MANTICORE project. In 2008 he started his collaboration in the PHOSPHORUS project. He also collaborates with the GLIF Generic Network Interface Technical Group (GNI-TG), adapting the Harmony Service Interface (HSI) to the GNI specification. In 2010 he started collaborating actively in the GEYSERS project. He is currently a Ph.D. candidate in the Telematics Engineering Department of UPC.

JOAN ANTONI GARCIA-ESPIN (joan.antonigarcia@i2cat.net) holds an M.Sc. from UPC (2007). He is a research project manager at the Network Technologies Cluster of the i2CAT Foundation and Ph.D. candidate at UPC. He wrote his Master's thesis on the design and implementation of TE-enabled, DiffServ-aware MPLS networks for providing end-to-end QoS. He is currently working in European projects such as GEANT3 and GEYSERS and participated in PHOSPHORUS and FEDERICA, as well as several Spanish research projects in the past. He is an active contributor to and editor of the Network Service Interface standard in the Open Grid Forum.

MARKUS HIDELL (mahidell@kth.se) is an assistant professor in communication systems at the Royal Institute of Technology (KTH), Sweden. He received his M.Sc. degree in telecommunication systems and his Ph.D. degree in telecommunication from KTH in 1992 and 2006, respectively. Since January 2008, Markus is an assistant professor at Telecommunication Systems Laboratory (TSLab), KTH. His current research interests include switch and router architectures, protocols, and network architectures. He has co-authored six patents in the area of network resource management and network topologies.

PETER SJÖDIN (psj@kth.se) holds a Ph.D. from Uppsala University in computer science. Since December 2002 he is an associate professor in communications networks at KTH in TSLab of the School of Information and Communication Technology. His current research interests include network systems, protocols and network architectures, virtualization architectures, and residential service management. He holds five patents in the areas of fast lookups and network resource management, and network topologies.

PEHR SÖDERMAN (pehrs@kth.se) received his M.Sc. in engineering from KTH in 2008, specializing in computer science. He joined the TSLab research group at KTH in 2009, where he is currently a Ph.D. student. His current research interests include measurement methodology, experiment repeatability, and security in large-scale research networks.

ANDREA BIANCO (andrea.bianco@polito.it) is an associate professor in the Electronics Department of Politecnico di Torino, Italy. He was Technical Program Co-Chair of HPSR 2003 and 2008, DRCN 2005, and the IEEE ICC 2010 ONS Symposium. He was a TPC member of several conferences, including IEEE INFOCOM, IEEE GLOBECOM, and IEEE ICC. He is editor of the Elsevier *Computer Communications* journal. His current main research interests are in the fields of protocols and architectures for all-optical networks and switch architectures for high-speed networks.

LUCA GIRAUDO (luca.giraudo@polito.it) graduated from Politecnico di Torino with a degree in telematics engineering in April 2007. Between June and December 2007 he was involved in the design and implementation of distributed software router architecture (BORA-BORA project). Since January 2008 he has been with Politecnico di Torino as a Ph.D. student under the supervision of Prof. Andrea Bianco, working on software routers, network virtualization, resource distribution, and centralized control in networks.

CRISTINA CERVELLO-PASTOR (cristina@entel.upc.edu) received her M.Sc. degree in telecom engineering and Ph.D. degree in telecommunication engineering, both from the Escola Tècnica Superior d'Enginyers de Telecomunicació, UPC, Barcelona, Spain. She is currently an associate professor in the Department of Telematics Engineering at UPC, which she joined in 1989, and leader of the optical networks research group within the BAMPLA group.

VÍCTOR LOPEZ (victor.lopez@uam.es) received his M.Sc. degree in telecommunications engineering from Universidad de Alcalá de Henares (2005) and his Ph.D. degree in computer science and telecommunications engineering from Universidad Autónoma de Madrid (UAM) in 2009. In 2004 he joined Telefónica I+D, where he was involved in next-generation networks for metro, core, and access. In 2006 he joined the High-Performance Computing and Networking Research Group at UAM. Currently, he is an assistant professor at UAM, where he is involved in optical metro-core projects (BONE, MAINS). His research interests include Internet services integration over optical networks (OBS solutions and multilayer architectures).

MIGUEL PONCE DE LEON (miguel@tssg.org) is a Waterford Institute of Technology (WIT) graduate with a degree in electronic engineering, and he is currently head of Communication Infrastructure Management, a unit of the TSSG, a research group that focuses on telecommunications software services management and Internet technologies. He has participated in over 35 international research projects focusing on the future Internet, where his research team of 30 staff concentrate on the architecture and design of the future Internet, looking specifically at self-management of virtual resources, threats to communication-based service information, and the Living Labs concept. He was General Chair for TridentCom 2008 and is a member of FITCE.

GEMMA POWER (gpower@tssg.org) holds a first class honors degree from WIT and has extensive industry experience with Apple. Since joining the TSSG in 2007, she has worked on the EU MORE project in the area of wireless sensor networks and validation and test. She is currently working on the EU PERIMETER project where her primary focus is the testbed setup, configuration, and federation, in addition to the code integration and testing, proof of concept, and demonstration aspects of the project.

SUSANNE NAEGELE-JACKSON (susanne.naegele-jackson@rrze.uni-erlangen.de) graduated with an M.Sc. in computer science from Western Kentucky University and the University of Ulm, Germany. She received her Ph.D. (Dr.-Ing.) from the University of Erlangen-Nuremberg in computer science. She has worked at the Regional Computing Center of the University of Erlangen-Nuremberg since 1998 on a variety of national and international research projects such as GTB, Uni-TV, Uni-TV2, VIOLA, MUPBED, EGEE-III, GN3, FEDERICA, and NOVI. She has authored and co-authored over 30 scientific publications, and teaches classes on multimedia networking at the Regional Computing Center.

The European NREN partners of the former project consortium have agreed to continuously support the operation of the infrastructure and are willing to do further development under the recently launched "Network Factory" task of the GN3 project's Joint Research Activity 2.