Report from the EU/US Future Networks Workshop

November 11, 2017

EDITORS: Serge Fdida, Ivan Seskar, Peter Steenkiste, Brecht Vermeulen

Acknowledgments

National Science Foundation (NSF) / European Commission (EC)

Table of Contents

Executive Summary

The EU DG Connect and the US National Science foundation established a joint EU-US committee bring together researchers from both sides of the Atlantic. The committee was asked to discuss the joint development of research-focused networking testbeds and identify joint research opportunities. The committee was organized as two subcommittees, focusing on two broad areas: “Next Generation Internet (NGI)” and “Advanced Wireless Platforms”. The committee met for a workshop at the DG CONNECT premises in Brussels on June 26-27, 2017. The workshop was collocated with the “Net Futures Conference 2017”, which included an “EU-US session on next generation internet” on June 28.

The joint committee built on the well-established US-EU collaboration on network testbeds (GENI/FIRE). The workshop was organized as two parallel tracks, discussing Next Generation Internet and Advanced Wireless Platform testbed infrastructure and joint research opportunities. The NGI track focused on identifying (1) joint research opportunities that can benefit from international collaboration, and (2) enhancements and extension to existing FIRE and GENI testbeds to support that research. The Advanced Wireless Platforms groups focused on joint platform development dedicated to support research in wireless at large.

The two-day workshop resulted in a set of joint recommendations for future collaboration, organized as short, medium and longer term opportunities.

We recommend to develop the collaboration following three time-scales.

- In the short run, the focus of collaboration should be on the exchanges of students, platform design and operation engineers, and faculty members.
- In the medium time scale, the collaboration should include joint projects with lightweight management focusing on specific targets and objectives with the main goal of encouraging development of common advanced wireless platform environments as well as collaborative research projects on NGI.
- In the long run, collaboration should include joint development of control/management frameworks as well as joint development of large experimental wireless and NGI platforms (whether collocated or as a multiple deployments) and architectures.
- The Open Data (ODMP), reproducibility and interoperability features should be considered at a very early stage of any joint design.

1 Introduction and Workshop Rationale

1.1 Workshop Rationale

Digital infrastructures are critical to support the digital transformation of our societies. This infrastructure includes edge networks that support users and devices generating and consuming information, and the core Internet that provides high speed global connectivity. The edge infrastructure is becoming increasingly wireless in order to cope with the fast growth of wireless devices and their mobility requirements. Their design is facing challenging research problems as we seek higher bandwidth, in an environment that is increasingly mobile, diverse and constrained by spectrum. The focus of the core Internet infrastructure is very high bandwidth communication between edge networks, service providers and cloud infrastructures by increasing the capacity of routers and the optical fiber infrastructure. In addition, user quality of experience (QoE) expectations are growing, raising the question of how to best improve QoE in an infrastructure that has traditionally focused on best effort network service.

Platforms and testbeds play an important role in testing and evaluating research contributions both at the edge and in the core of digital infrastructure. Platforms for Future Advanced Wireless Networks provide a scientific instrument in order to assist the design of its components and test the system efficiency and robustness, under realistic conditions, in a controllable and reproducible manner. Such platforms must cover a broad set of research questions, ranging from specific low level wireless concerns to the study of the entire system, including the various verticals (application domains). In the core, wide-area network testbeds are needed both to evaluate new network architecture and protocols designed to address challenges in areas resource management in support of optimizing user QoE, managed network services, traffic engineering, security, etc. In addition, diverse cloud research platforms are needed in support of research in cloud data centers and edge computing (e.g., cloudlets). These infrastructures must be connected, supporting end-to-evaluation under realistic conditions

A joint US-EU partnership in this domain is of utmost importance given the significant effort required to build, instrument and operate these types of platforms. They should be designed in order to provide a remote and open service to the community, so mutualizing the resources as well as sharing best practices and solutions to be integrated is necessary. These shared infrastructures naturally enable joint research. Recent experience with EU-US cooperation in the field of experimentally driven research (like cooperation between GENI and FIRE) are a great example. The goal is to extend US-EU collaboration to more diverse shared experimental platforms and testbeds, including wireless technologies and experiments, and to joint research opportunities that benefit from cross-Atlantic collaboration.

DG Connect and the US National Science foundation established a joint EU-US committee in order to bring together researchers from both sides of the Atlantic. The committee was asked to discuss starting/building/operating/sustaining research-focused networking testbeds and identify joint research opportunities. The committee was organized as two subcommittees, focusing on two broad areas: “Next Generation Internet (NGI)” and “Advanced Wireless Platforms”. Committee members are listed in Appendix A. The committee met for a workshop at the DG CONNECT premises in Brussels on June 26-27, 2017. The workshop consisted of parallel meetings by each subcommittee. In the last session, the Wireless and NGI tracks reported their initial findings. Each subcommittee also had a number of conference calls to prepare for the workshop.

The workshop was collocated with the “Net Futures Conference 2017”, which included an “EU-US session on next generation internet” on June 28. The session opened with the series of invited presentations from academia and industry and concluded with a panel in which the workshop participants presented their recommendations and outcomes. The remainder of this report presents recommendations and summaries of the outcomes of the Wireless and NGI tracks.

1.2 Drivers for Collaborative US-EU Research

The committee identified several compelling drivers for joint EU-US research, i.e., cases for which joint research projects will be more effective than separate projects:

- Many research challenges are inherently Internet-wide: The Internet is a global infrastructure and many challenges cannot be partitioned into per-county or even per-continent challenges. Examples include internet-wide management of CDNs and clouds that must consider latency requirements and geo-diversity, and Internet control functions such as inter-domain routing, traffic engineering, and monitoring. Research in these areas can benefit both from shared research infrastructure and collaborative research.
- Dealing with fundamentally different requirements or constraints: Legal and regulatory requirements with respect to networking and cloud computing differ across countries. Examples include wiretapping laws and rules about user privacy both at the network and application level. Research in how both operators and users can be deal with this diversity naturally benefits from international collaboration.
- Opportunities to learn from different research approaches or contexts: Business models and network deployment models differ across countries, e.g., home networks and IXPs. This has led to different research approaches and opportunity to learn from each other. Along the same lines, the US and EU have focused on different types of future internet architectures (clean slate versus evolutionary), similarly creating opportunities for collaboration.
- Enablers for future research: There is an increasing interest in software-defined infrastructures, making interoperability an important challenge. Defining APIs and open source platforms that are shared between the US and the EU (and more broadly) is an important enabler for future collaborative research. Examples include APIs and platforms for sharing virtual network functions and common ontologies for resource specifications, data integration and big data analysis, and sensors and IoT services.

2 Advanced Wireless Platforms Track

Wireless communications and related technologies have been identified as one of the most important research topics on both sides of the Atlantic, motivating a joint initiative on experiment-driven wireless research with the objective of removing barriers from collaborations. This action is timely as EU and US experimentally driven research communities have matured, structured projects exists in this domain and provide a good source of information. The projects and testbeds listed below demonstrates the vibrant activity this community as well as the diverse set of solutions deployed: In Europe - 5tonic, 5G Labs of Dresden, 5G Labs 5GIC at surrey uni, OneLab, FIT, Fraunhofer 4/5G Testbeds, Bristol is Open, OpenAir, 5G Satellite-Terrestrial Testbed, Fed4FIRE, FIWARE Lab Node, and imec iLab.t testbeds. In the US - Wiser-Lab, WiTEST-Lab, CalIT2, Wireless@VT, WINGSNet, Wireless for Underserved and Under-resourced Communities, PhantomNet and ORBIT.

Following the series of US-EU collaboration planning conference calls and white paper on "Joint partnership between Europe and US on Advanced Wireless Platforms" (http://www.winlab.rutgers.edu/events/euuswws/), a first workshop was held on June 26-28, 2017 in Brussels jointly with a parallel NGI effort. The wireless track had 22 attendees including 16 from academia and 6 representatives from non-profit government agencies (DG Connect and NSF).

The Advanced Wireless Platform track of the workshop was structured to address two goals. The first goal, addressed during the first session in the afternoon of day one, was to share experiences of current large-scale wireless platform developers and operators. In particular, the session focused on new wireless technology challenges for experimental research. The experiences with scope evolution, architectural stability, staffing and management, roles of university/city/government/industry in testbed evolution, user community interactions, sustainability, etc. were also discussed. The second goal, and the focus of the second session of the afternoon, was to introduce new challenges and needs to guide the development of future and innovative wireless experimentation platforms. The discussions continued, during the morning of the second day, primarily focusing on developing recommendations for the collaboration objectives related to wireless testbeds.

The remainder of this section presents the results from the wireless track of the workshop. The material from the workshop including this report, is available at the workshop website at:

http://www.winlab.rutgers.edu/events/euuswws/ .

2.1 Past Experiences and Position Statements

The first day afternoon agenda was split into two sessions:

a. Advanced wireless platforms technical challenges and their suitability for collaborative research
b. Ideas on new models for cross-Atlantic collaboration.

Participants were also asked to briefly point out relevant past experiences and lessons learned from managing/using testbeds as well as from past collaborations involving multiple institutions from both sides of the Atlantic. A number of existing wireless testbeds shown in Table 1 were represented at the workshop and their experiences were used to illustrate various points.

Array of Things (http://arrayoLhings.github.io)	An urban sensing project, a network of interactive, modular sensor boxes installed around Chicago to collect real-time data

	on the city's environment, infrastructure, and activity for research and public use.
"Technology for all" (http://www.techforall.org/)	Community deployed wireless network, created to provide free, secure wireless Internet to 19,000 residents in low-income neighborhood.
WiTEST (https://witestlab.poly.edu/)	The Wireless Implementation Testbed Lab (WITest) in the Department of Electrical and Computer Engineering at NYU Polytechnic School of Engineering conducts research, education, and outreach focused on implementation of, and experimentation with, wireless networking protocols, applications and services
w-iLab.t (http://doc.ilabt.imec.be/ilabt-documentation/wilabfacility.html)	The w-iLab.t is a, generic, heterogeneous wireless testbed deployed in multiple locations offering different wireless technologies: sensors, Wi-Fi, Bluetooth, LTE, SDR, and long-range radios. imec also has a portable testbed that can be deployed anywhere with the same hardware and experiment control features as fixed test facilities.
PHANTOMNET (https://www.phantomnet.org/)	PhantomNet is a mobile networking testbed that provides researchers with a set of hardware and software resources that they can use to develop, debug, and evaluate their mobility ideas
NITOS Testbed (https://nitlab.inf.uth.gr/NITlab/nitos)	Indoor/outdoor testbed focusing on wireless/wired networking and their applications with main technologies including: WiFi, WiMAX, LTE, SDR, mm wave, openflow, cloud, sensors. Federated with most of the EU testbeds through OpenLab/Fed4FIRE/Fed4FIRE+
5TONIC (https://www.5tonic.org/)	5TONIC is an open co-creation laboratory focusing in 5G technologies, founded by Telefónica and IMDEA Networks and based in Madrid
5G BERLIN TESTBED (http://www.5G-Berlin.org)	The 5G Playground enables the 5G ready trial platform, which offers agile MEC/FOG computing capabilities and is connected to multi-access networks within 5G Berlin
WINGS (http://research.cs.wisc.edu/wings/wiki/doku.php)	The WiNGS indoor wireless testbed consists of about 60 nodes, deployed in the Department of Computer Sciences building.
Bristol Is Open (http://www.bristolisopen.com)	City-scale deployment of 144-fiber core network connecting 4 active nodes, full optical switching, flexi optical with Wi-Fi 802.11ac, LTE, mmWave, MM-MIMO, 60GHz backhaul and RF Mesh Network with 8 Fiber-connected lampposts with 1,500 gateways and any-sensor hosting capability.
ORBIT (http://www.orbit-lab.org)	The ORBIT testbed is an indoor 400-node programmable radio grid and an outdoor field trial system of short- and long-range radios. Includes SDR platforms, LTE, MIMO and cloud RAN capabilities
FIT –Future Internet of Things (https://fit-equipex.fr/)	FIT offers large-scale wireless, sensing and mobility infrastructures. FIT platforms are located across France. FIT consists of almost 3000 nodes. You can even plug your own devices in our testbeds and run your tests there as well.

5GIC – University of Surrey (https://www.surrey.ac.uk/5gic)	A large scale and carrier grade open testbed for research and innovation with 44 (indoor and outdoor), 4G and two 5G radio access points with EPC, Soft EPC and Virtualised Flat Distributed Cloud Architecture.
5GTN (http://www.5gtn.fi/)	5GTN is an open nearly carrier grade test network including 5G proof-of-concept (5G-PoC), 4G small cell and macro cell, WiFi, LoRa, NB-IoT connectivity. The environment contains virtualized MEC functions with open APIs for service development as well as IoT data platforms. The core network is realized using both NFV based Nokia EPC solution for operational purposes as well as OpenEPC for more research oriented purposes. The 5GTN resides in Oulu, Finland and is operated by University of Oulu and VTT.

2.2 Advanced Wireless Platforms Technical Challenges

The objectives of the first session were to discuss research issues related to wireless technologies as well as addressing some of the broad fundamental networking challenges: scale, complexity, security i.e., common network challenges also present in mobile and wireless networks.

Mobile mmWave and (distributed) massive MIMO were identified as major physical layer challenges in a number of presentations. It was repeatedly pointed out that real-time SDR implementations brought in by the evolution in signal processing based on standard CPUs are increasingly dominating the research as well as deployment landscape. The limitations of CPU-only implementations in achieving low latencies can be significantly alleviated through the use of hardware accelerators and on-chip high-speed interconnections wrapped by software APIs and use of HW/SW co-design techniques. Closely related to these physical layer issues are support for high-performance fronthaul, mid-haul and backhaul networks and related high-performance fixed networking services (NFV,SDN, etc.) as well as distributed computing and standard off-the-shelf computing platforms (i.e. Intel servers and NVIDIA GPUs) and cloud-computing based services (FOG/MEC/OpenStack, …) and their integration in the cloud radio access networks (cloud RAN).

The softwarization of the wireless technologies ecosystem brought into perspective the issue of difference in design methodologies as well as democratization of the whole development process (i.e. shift from standards driven to open-source approach). This in turn exposes the open source licensing conundrum in which relatively large number of diverse participants (especially major vendors and operators) need to agree on a common licensing model. Experimental validation is required along the full development life cycle including both early limited testing in lab and testing at scale in the field including the need for testbed portability (to better facilitate testing in real-life environments).

Moving further up the protocol stack, a number of participants pointed out the emerging importance of edge microservices/services and overall service integration research especially in the context of services for vertical sectors (i.e, joint optimization of networking and VM placement). The fact that applications quite often drive research and demand, was also repeatedly pointed out and, somewhat related to that, an issue of how to get (a large number of) experimenters and platform users engaged (and how to retain them for the duration of the deployment).

Testbed platforms must be open and hackable supporting rapid innovation and should be shared as much as possible in multiple-deployments. Rapidly evolving these platforms to keep

up with fundamental technology evolution was also recognized as one of the issues that needs to be addressed. Similarly, in order to support sustainability, platforms have to evolve from academic and research use to industrial experimentation that also implies open/unified interfaces and ease-of-use for large-scale experiments in order to reduce support requirements.

In addition to these, a number of other wireless technical issues were discussed including:

- Low-power and battery-less communications with large number of devices (connecting with low/modest bit rates)
- Low latency communications and especially supporting experimentation with cloud RAN based solutions and applications like VR/AR
- High mobility communications including support for high speed V2V and V2I communications
- Visible light communication systems.

A number of operational issues were also pointed out in presentations including:

- Need for common management platform and operations framework
- Addressing the issue of inter-testbed connectivity and federation, or extending the federation concept from testbed-level to (wireless) node level, enabling the creation of heterogeneous testbed infrastructures through mixing platforms owned by different parties.
- Difficulties in staffing with experienced engineers for both design and maintenance/support (in academic environments).
- Need for city/municipality partnership
- Managing the expectation of high level of availability and reliability of the testbed (requiring dedicated support personnel and 24/7 operation)
- The importance of varying levels of user support (i.e. supporting novice, intermediate and advanced users)

And finally, the need for a common data collection standard (e.g. Open Data) and commitment by the whole community was identified as one of the essential requirements.

The collaboration should develop along the four commons:

- Common knowledge: focus should be on the integration of the community, exchange of students and researchers, share of experiences and best practices, organizing common events and other knowledge sharing activities.
- Common tools: in order to broaden the usefulness of individual efforts and to lower the barriers to entry for both experimenters and platform operators, identify common (existing) tools and/or support their development (if they do not exist); connect with other communities (e.g. compute grid, data analytics platforms, etc.)
- Common platforms: jointly develop platforms for experimentation with fully integrated ecosystems covering topics from basic research to start-up launching.
- Common usage/research: develop and carry out cross Atlantic end-to-end technology trials using commonly developed platforms; initiate work on beyond 5G and bridging the gap between SDR and SDN research

3 Next Generation Internet Track

The US and the EU already have a history in sharing research platforms for core Internet research, specifically in the context of the collaboration between FIRE and GENI. As a result, the focus of the NGI track was to develop not only a plan for expanding and improving the sharing of experimental research infrastructure, but also an agenda for collaborative research, especially involving topics that can benefit from cross-Atlantic collaboration. The committee only considered NGI opportunities that are enabled by todays and future communication infrastructures, but we also discussed research enabled by compute and storage resources that are increasingly embedded in the infrastructure.

In preparation of the workshop, the NGI subcommittee had a number of conference calls that were used to identify broad areas of possible NGI US-EU collaboration. Between calls, participants used shared documents to expand on the ideas generated during the conference calls. The workshop sessions on Monday afternoon and Tuesday morning were used to discuss specific research topics in the research areas of interest. The rest of the time was used to generate a draft for the report and recommendations. The rest of this section summarizes the outcomes of the NGI track.

3.1 NGI Opportunities

The Internet has changed dramatically in the last 40 years, and this evolution is continuing at a rapid pace. To end-users, the most visible change is never-ending increases in throughput, but an equally important change is that the Internet has moved away from a “dumb network” to an infrastructure that has significant amount of computational and storage resources embedded in it. Examples include Content Delivery Networks that offer storage and servers for hosted services, and clouds of various sizes (included cloudlets and “fog”) that provide processing and storage resources. These diverse resources can be used to provide new services that can be of benefit for service providers, end-users and networking researchers and developers. The committee discussed three classes of services and capabilities: path aware networking, managed networks, and support for evolutionary and clean-slate architectures.

Path aware networking: Path awareness at higher layers of the stack is an emerging trend in networking. Many endpoints are multiple connected to the Internet (e.g., smartphones via mobile as well as terrestrial networks), and new transport protocols (e.g. multipath TCP) take advantage of this to place different traffic on different paths according to measured or assumed path properties. Another opportunity is for end-points to select paths with specific properties (e.g., bandwidth, latency). While allowing individual clients to request paths with certain quality of service properties may not be practical, this is not out of reach for service providers. More explicit support for awareness of path properties throughout the stack, and the ability to forward traffic along distinct paths based on information other than the destination prefix, will allow for richer optimizations of network treatment per flow and traffic type.

Managed networks: Traditionally, a managed Network consisted of a set of VPNs between different private networks. However, the concept is much broader and extends to creating virtual entities (VMs, Dockers) in the cloud and connecting those with VPN and at the same time instantiate completely different new forward paradigms inside such virtual environment. This greatly benefits applications that need more expressivity towards the networks to provide their service, e.g., content delivery networks. Basically this type of managed network is pioneered as GENI and Fed4FIRE testbeds federated together. Inside the VMs the routing/switching can be totally different than in the normal Internet, one can run his own protocols. Such managed

networks, sometimes called overlays, are not only interesting for research but also for business applications, e.g., providing managed connectivity for distributed applications, or even for services and their users. This would make it necessary that applications can generate and operate overlay networks on the fly as needed. Networks with very specific properties can be generated and surprisingly the QoS can be better for some attributes than the underlying network delivers. Making flexible managed networks a reality requires not only resource allocation, but also mechanisms to customize the management of these resources at runtime.

Evolutionary and clean-slate Internet architectures[1]: The current Internet protocol suite has its roots in the early days of the Internet when the focus was on supporting basic connectivity for a small set of applications. The Internet has changed dramatically, and many research groups are exploring changes to, or even replacement of, IP to better support today's and tomorrow's applications. For example, the US has a number of large "clean slate" internet architecture protects (the eXpressive Internet Architecture, MobilityFirst and Named Data Networking). These architectures often move away from the current host-centric architecture, but since incremental deployment is often a secondary consideration, deployment of these new protocol suites is very difficult. EU research in this space focused on evolutionary research (e.g. 4WARD, POINT, ICN2020, IRATI, PRISTINE, ARCFIRE) and on single core ideas. Examples include work on next-generation Internet-deployable enhancements to the transport layer, e.g. Multipath TCP (Trilogy and Trilogy2 projects), extensions to TCP for low latency (RITE project), and generalized dynamic transport protocol selection (NEAT project). Experimentation with, and incremental deployment of, evolutionary and especially clean slate architectures is very challenging.

3.2 A Research Agenda for EU-US Collaboration

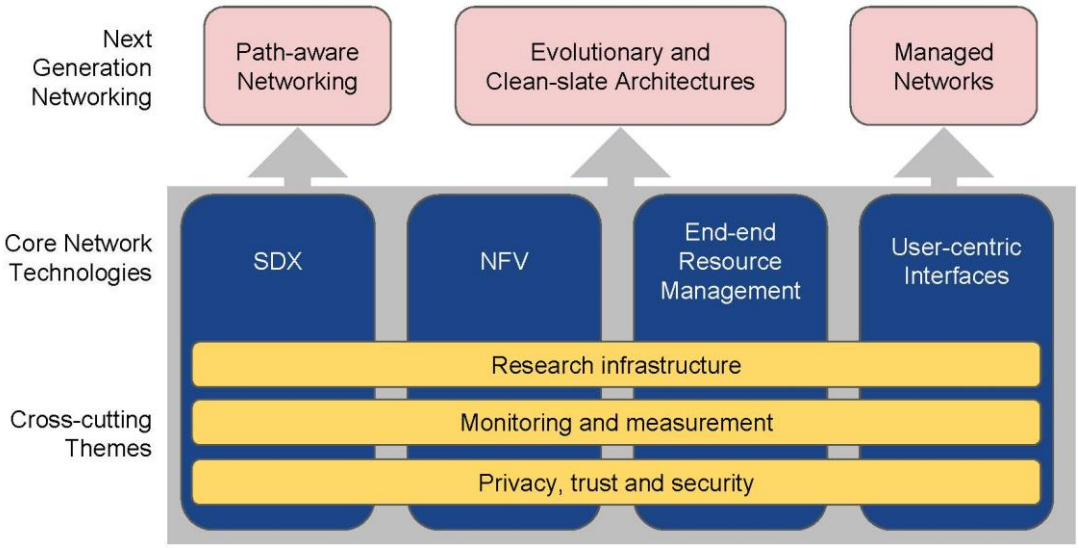

Figure 1: EU-US NGI Research Agenda

The subcommittee identified a research agenda as illustrated in Figure 1. It consists of four core network technology and three cross-cutting areas. The core areas include collaborative research in SDX and NFV, and research on "horizontal' resource management across the Internet and user-centric "vertical" interfaces for propagating user requirements throughout the stack. The

[1] https://cacm.acm.org/magazines/2010/9/98030-future-internet-architecture-clean-slate-versus-evolutionary-research/fulltext

three cross cutting areas are monitoring, shared research infrastructure, and privacy, trust and security. We elaborate on these areas in the next two sections.

3.2.1 Core Network Technology Research

We describe the core research areas identified in Figure 1.

1. Software Defined eXchanges (SDX)

In today's networks, Internet Exchange Points (IXPs) have provided a broad and convenient forum for ISPs to enter into interesting pairwise arrangements. With the advent of software defined controls, IXPs are transforming into Software Defined eXchanges (SDXs) that facilitate programmatic control and enforcement over peering policies and can couple different forwarding technologies together. SDXs are likely to play multiple roles in future networks.[2][3]

- Internet control plane protocols and architectures. The Internet's dated inter-domain routing protocol, BGP, is brittle, inefficient, and insecure. SDXs provide an opportunity to augment or replace BGP with a different and/or more flexible enhancement that addresses these problems from the ground up. By leveraging many Internet-wide SDXs in concert, it could become possible to develop a framework for rolling out such deployments over time. Such changes to BGP are simply not possible today.
- One example opportunity is how the increased flexibility of per-domain SDN can be extended across domains. How are policy needs of operators and users expressed and enforced? What information should flow across the stakeholders and how should it be disseminated? How does this interact with the set of policies that can be expressed and enforced? Can the framework be rich enough to support policies that differ fundamentally across regions (e.g., policies pertaining to legal/regulatory issues in US/EU)? What programming languages can be used to express and reason about these policies.
- Hosting and brokering software defined infrastructure (SDI). Exchange points may host and/or administer resources on behalf of other entities in the Internet. A SDX can act as a inter science DMZ, providing policy-mediated access to resources. Example usage may include (a) remote placement of network functions within an SDX or in SDX-managed resources, (b) incorporation of third-party networking and computing resources, (c) provide storage and data transfer nodes including encryption and transformation functions.
- Enabling new ways for applications and end users to interact with the network. As SDX looks increasing like a public cloud, and it has the potential to scale up in supporting large numbers of customers running customized network functions for specific end-to-end network connections. This fundamentally changed one's relationship with the network from today's "forwarding-only" connectivity to personalized, individually-owned (e.g., via cloud-like leasing agreements) end-to-end networks via SDXs.
- Addressing multi-domain trust issues when enabling end-to-end path provisioning. Analogous to CDNs simplifying the multi-domain QoS engineering problem by decoupling the many-to-many (customers with ISPs) negotiations into a more manageable many-to-one (customers to CDN provider) + one-to-many (CDN provider to ISPs) negotiations, SDXs may have similar contributions when it comes to simplifying multi-domain trust and

[2] "Workshop on Prototyping and Deploying Experimental Software Defined Exchanges (SDXs)," June 2014,http://groups.geni.net/geni/raw-attachment/wiki/SDXandSDIWorkshop/SDX%20Workshop%20Outbrief%20-%20Draft.pdf.

[3] "Software Defined Technologies - What's next?," December 2016, https://ec.europa.eu/futurium/en/system/files/ged/software_defined_technologies_-_whats_next_consultation_report.pdf.

negotiations to achieve end-to-end path provisioning with deterministic properties.

Since these solutions must be applied Internet-wide and business models, laws and policies differ across regions and countries, international collaboration is needed.

2. Network Functions Virtualization (NFV)

Network Functions Virtualization (NFV) decouples Network Functions (e.g., firewalls, load balancers, IDSes, and caches) from proprietary hardware appliances. Virtualized Network Functions (VNFs) can be combined (chained) in a building block-style fashion to deliver full-scale networking communication services. By leveraging standard IT virtualization technologies (VMs, containers, virtual switches), NFV promotes the consolidation of NFs onto industry-grade high-volume servers, switches, and storage reducing OPEX/CAPEX and improving manageability. Besides cost reduction, NFV allows for dynamic, elastic provisioning which can lead to the rapid instantiation of agile services and matching workloads via VNF instance scale out/in or up/down.

However, important research problems remain in order to realize the full potential of NFV and enable new services, such as the next generation network services described above.

- Research topics include (but are not limited to) VNF placement, VNF scaling and load balancing among multiple VNF instances, efficient and seamless VNF state migration, dynamic service function chaining, VNF-specific issues (e.g., security/privacy issues, performance, and portability of specific VNFs), distributed management of VNFs, complete management solutions for NFV/SDN, and research on important use cases such as deploying VNFs to support 5G and IoT. The key objectives of this research direction are to increase scalability, fault tolerance, and performance (latency) of network functions and service function chains (traffic engineered SFC), and enable NFV to support novel use cases (e.g. for immersive AR/VR applications). These challenges are inherently Internet-wide since they must consider geo-diversity and regulatory and legal properties of both the cloud and network infrastructure.
- Interfacing and interoperability are key requirements for chaining functions from different vendors/operators into a single service. The use of open and standardized descriptors for resources, functions and services is crucial, but existing models (e.g., TOSCA, NETCONF/YANG) were not developed to support NFV-specific requirements. Trust and privacy concerns are also important and vary based on the class of VNF (smart forwarding versus deeper packet inspection). Placement of trusted functions (e.g., caching) in untrusted foreign environments introduces additional challenges.
- To promote research on NFV or employing NFV as an enabling technology towards next-generation network services and cloud-based NFV, we need to exploit advances in both the EU and US (e.g., OAI, OpenEPC, vIDS, Clearwater IMS) and make them collectively and readily available to experimenters. Additional synergies may include the promotion (or enrichment) of testbed federations that support research and development in areas related to specific application areas of NFV and SDN (SoftFIRE as an example).
- There are currently different standardization efforts (e.g., IETF SFC WG, ETSI NFV ISG). Collaboration can help US and EU researchers learn about the different use cases for telcos, respectively, in US and EU. It can also help answer questions on the throughput and latency goals for orchestration, state management, scaling to large ISP backbones, and disaggregation of network stacks. We can match US-EU research projects and form academic-industry-network-operators teams. We can also compile a global catalog of NFVs/infrastructure available for research, and share platforms, software, ideas, results, and experience with testbed setup and operation.

3. "Horizontal" Resource Management

The battle between over-the-top (OTT) providers and network operators is an example where tussles between different stakeholders have been played out, with the OTT providers arguably winning. The current situation we are dealing with is OTT content delivered over the Internet by Service Providers without the involvement of a Network Operator in the control or distribution of the content. Each stakeholder independently optimizes resource utilization and service quality through processes to monitor and predict sufficient system state information needed for autonomic and human decision-making. The desired situation we are targeting is for capabilities providing mechanisms for Service Providers and Network Operators to work together in the management of resources required for control and distribution of content and service delivery. This challenge requires research both horizontally (resource allocation across stakeholders – this section) and a vertically (User-centric interfaces – next section).

A first challenge is that the delivery of different types of content and services requires diverse Internet resources that are owned by a very large set of independent providers, e.g., typically multiple ISPs for communication, and compute and storage resources from content providers, CDNs, and cloud providers, each of which has resources distributed across the Internet. The allocation and customization of these resources is complex, in part because the internal resource allocation decisions of individual providers in unpredictable ways. Since there are no interfaces for coordination, a common strategy for OTT providers is to treat other providers as black boxes and to use measurements to predict their performance. This is an ad hoc solution that does not lead to consistent high quality of experience for users.

However, two recent developments in networking, Network Function Virtualization (NFV) and Software Defined Exchanges (SDX) can be the basis for well-architected and widely deployable platforms to support the above services. NFV allows the dynamic configuration of in-network functionality in a cost-effective manner. SDX, and more generally new network control protocols, provides an opportunity for coordinated traffic management across Internet service providers.

Leveraging these two developments require research in two areas. First, we need to define interfaces that allow the different stake holders to coordinate resource allocation decisions explicitly. Challenges include how offered/requested resources are expressed, how to ensure stability, how providers can ensure that internal policy and business requirements are met, and limiting how much internal information providers have to expose. Second, new algorithms are needed for resource allocation, using the above interfaces and also applying data analytics to optimize performance.

4. User-Centric Interfaces

Today's interfaces up and down the Internet protocol stack, and interfaces between the applications at the endpoints and the network elements which forward traffic on their behalf, are very narrow. This narrowness stems from the engineering practice of strict layering and the assumption of a dumb network that offers only best effort service. The narrowness of these interfaces has become a hindrance, by forcing important information about the network (e.g., state) to remain hidden, and thereby limiting control. A great deal of research, e.g., into QoE, has gone into widening these interface, and/or into working around the lack of wider interfaces, e.g. through endpoint protocols that integrate measurement and monitoring. This research has however focused on communication resources, not the diverse resources used by today's services and applications. Another concern is that interfaces need not only express policies from endpoints to control the network, but also allow end points to extract information about network properties (at various granularities, e.g., paths, path segments, etc.).

In addition, there are important current trends that necessitate ground-up research into this topic. Software-defined networks (SDN) and exchanges (SDX), and Network Functions

Virtualization (NFV) provide a basic enabling technology to exercise more flexible control over networks. Also, endpoint software and end protocols themselves are becoming implicitly more aware of network properties (what they can infer) in steering traffic (e.g., across multiple interfaces, or along multiple paths, similar to MPTCP). Finally, the accelerating drive to encrypt user data as well as protocol headers (e.g. QUIC) will render useless existing methods for transferring and maintaining relevant network information which today are implicit in nature (i.e., they use deep packet inspection).

Research agenda: The set of research problems here is rich and includes, but is not limited to:

- representing policies for network treatment of end-to-end flows (e.g. performance and forwarding constraints, etc.) under a set of constraints such as privacy and security;
- finding expressive mechanisms to specify policies, and developing efficient mechanisms for information gathering and policy enforcement; and
- evolving the Internet to widen the interface between applications on endpoints and functions in the network, making the maintenance and transfer of network state explicit.
- improving the quality of state information shared (see monitoring below) between stakeholders to improve the quality of decision making in management processes.

Many of these issues will require standardization (whether through existing SDOs or via an open-source project) because they hinge on incremental changes at endpoints as well as at various points in the Internet core. Furthermore, policies will need to be expressed with globally consistent semantics. In particular, differences in approaches to privacy (i.e., reduction of information radiated about encrypted user traffic by these mechanisms) will be a key point of collaboration and coordination across US and EU. Experimentation with these mechanisms will require global testbeds (i.e., diversity in latency/bandwidth of access networks and in network management policies) to be useful.

3.2.2 Cross-Cutting Research Areas

1. Monitoring

The Internet is highly distributed infrastructure composed of independently managed networks and services, where no single entity has global control. In this context, monitoring within network and service domains as well as end-to-end is essential for network and service management, network and application performance optimization, diagnosis, and security. The management of services and network resources in such systems is challenging. The distribution and complexity makes it impossible to have accurate and timely global state information needed for multi-objective optimization. As such, centralized control will not provide adequate solutions to service availability and performance guarantees, considering the complexity, the uncertainty and the multi-stakeholder nature of the systems.

Monitoring is a cross-cutting topic as it touches multiple layers of the protocol stack as well as multiple networks and networked services and applications. Different stakeholders access different types of information about content and infrastructure resources that they control or have access to in an attempt to maintain sufficient state about resources. Measurement points are established either directly on resources themselves or derived through indirect techniques such as DPI or active probing of network paths where direct access to information is not available. Monitoring is particularly challenging at the network edge, which is extremely heterogeneous, and where it is hard to obtain monitoring vantage points and instrument devices due to portability, privacy and user incentives.

The challenges that emerge include:

- How to capture end-to-end properties (not only about services and infrastructure but also about user demands and experience) given that information is scattered across multiple stakeholders?
- How to ensure the dependability (i.e. correctness, consistency, performance) of state shared between stakeholders (often with competing interests)?
- How to increase completeness and availability of state in communication processes related to resources under the control of other stakeholders whilst considering the need to constrain complexity of optimization processes?
- How to define conditions under which the delegation of direct control of resources between stakeholders is acceptable?
- How to create monitoring systems with the right set of incentives so that they are adopted by a large user population?

Business and political realities of different geographical areas shape how the network is deployed as well as the set of Internet access choices and services available to users. As a result, any study on monitoring the current and future Internet's properties will immensely benefit from cross-Atlantic collaboration to cover different areas of the Internet under different commercial, political realities. Topics that can benefit from cross-Atlantic collaboration:

- The development of measurement methods and knowledge management tools that can work across the large set of scenarios that we'll encounter in both sides of the Atlantic.
- Building, deploying, and maintaining measurement infrastructures across the Atlantic. One example of successful international collaboration is the PerfSonar monitoring system for NRENs.
- Given the role of the Internet as a critical infrastructure, policy makers and regulators in Europe and in the United States are asking similar questions of how to regulate the Internet market and whether and how to enforce the network neutrality, which brings the need and opportunity for joint research on defining standard metrics and shared measurement methods for regulating the Internet market.
- Comparative studies of internet properties through acquisition, integration, reconciliation of distributed data sets across regions

2. Shared research infrastructure

Between US and EU there are already multiple research infrastructure efforts that have been closely interacting and collaborating. These efforts are expected to continue to evolve the research infrastructure in support of the new research agendas. Examples of such infrastructure with current EU-US collaboration include:

- GENI (US) and Fed4FIRE (EU)
- CloudLab (US) with multiple federated EU deployments
- PlanetLab (US) and PlanetLab Europe
- Internet2 (US) and Geant (EU)

The EU SoftFIRE project also takes a federated approach to integrate multiple testbeds to support new research in SDN and NFV. In addition to collaboration in sharing development experiences, having a few stable frameworks/tools/APIs upon which new technologies, features and infrastructure are added is beneficial, so as not to reinvent the wheel when new research testbed needs arise.

There is no need to connect all research infrastructures, but being able to do

experiments/deploy applications on multiple infrastructures (e.g., smart cities) and comparing the experiment data is very beneficial. For other research (e.g., realistic large latency networking, increased scale and interop related), the flexible high bandwidth interconnections are key.

In addition, the need for making available a research facility to support the NGI (and wireless) research is instrumental. This is not only because it should support the discovery process but also as it should provide confidence to the results produced by the community. As such, efforts should be associated towards open data and reproducibility.

3. Privacy, trust, and security

The committee identified a number of opportunities for joint US-EU research in this area.

Privacy across borders: Laws and policies governing Internet security and privacy—ranging from surveillance to data retention—often differ across national borders. This creates many challenges in core networking (e.g., routing), service delivery (e.g., copyright issues), and the use big data. Here are some specific examples:

- New GDPR legislation[4] will enforce the 'Right to be Forgotten' and allow users to request that their data is moved from one organization to another. How can we setup, manage and run Big Data infrastructures that will support applications where users cross national boundaries? What happens when regulation, law and policies change? How can big datasets be interoperable from a privacy point of view, for example, when combining datasets with differing privacy policies. This issue becomes even harder if the inferencing abilities of semantic languages (such as OWL) and machine learning are considered.
- The United States has laws that constrain the surveillance of Internet traffic, including Executive Order 1233 and FISA Section 702. Whether the traffic is collected within the United States or abroad has direct ramifications for the limitations on surveillance[5]. Research by Goldberg *et al.* has demonstrated that today's Internet interdomain routing protocol, BGP, is vulnerable to attacks such as route hijacks, which can cause Internet traffic to detour through different countries. As a result traffic local to the United States might detour through the EU, thus facilitating surveillance that might otherwise have been protected by FISA Section 702[6]. Research is needed in routing protocols that are more robust to such routing hijacks and detours. In addition, mechanisms are needed to provide users with better transparency concerning their traffic's path and privacy.
- The European Union has been drafting copyright laws that threaten to remove safe harbor protections for online service providers (e.g., Google, Facebook) unless they deploy content fingerprinting tools to facilitate automated takedown of copyrighted content[7]. The US Congress is considering similar legislation, for both copyrighted content and for automated detection and takedown of certain forms of speech (e.g., terrorist and hate speech). Since these laws can be misused, the design of networks and systems that protect online speech will become increasingly important in light of these developments. In some countries for example, copyright laws are already being abused to censor political speech[8] or to take down certain content[9]. US and EU researchers have promising initial research in this

[4] https://www.google.be/search?q=gdpr+legislation
[5] http://webpolicy.org/2014/12/03/eo-12333-on-american-soil/
[6] Unrestrained Bulk Surveillance on Americans by Collecting Network Traffic Abroad. Axel Arnbak and Sharon Goldberg. Michigan Telecommunications and Technology Law Review (MTTLR). Vol 21(2), May 2015.
[7] https://www.eff.org/deeplinks/2016/10/upload-filtering-mandate-would-shred-european-copyright-safe-harbor
[8] https://businesstech.co.za/news/internet/162547/south-africas-3-new-proposed-censorship-laws-you-need-to-know-about/
[9] https://cdt.org/blog/pressuring-platforms-to-censor-content-is-wrong-approach-to-combatting-terrorism/

area[10][11] that they can build on to protect online communication and speech in the presence of shifting legal and policy frameworks,

Distributed Ledgers provide a generic technological solution to solving trust issues. We could investigate whether we wish to setup an EU-US blockchain for research, experimentation and innovation. This would provide an experimental space for researchers to setup distributed applications and test along technical (scalability, speed) and non-technical (trust, privacy, usability) dimensions. Distributed Ledgers could also be used to support new ways of sharing research data in a trusted fashion.

Privacy with NFV. Privacy concerns are important with NFV and vary based on the class of the VNF (smart forwarding ('same' data fields processing as a router) versus deeper packet inspection). Again, laws and regulations vary across countries, and so do attitudes towards security and privacy. The VNF instances may be placed in different data centers using particular placement and scaling algorithms. Placement of trusted VNFs (e.g., caches) in untrusted environments, e.g., a data center in another country or a public cloud, introduces additional challenges that require novel solutions. See [http://www.etsi.org/deliver/etsi_gs/NFV-SEC/001_099/004/01.01.01_60/gs_NFV-SEC004v010101p.pdf] for the ETSI security and privacy specification document.

Security and Privacy of IoT Devices. US and EU network infrastructure alike is increasingly under threat from the growing number of insecure Internet-connected special-purpose devices—the so-called "Internet of Things". The past year has seen several high-profile attacks spawned from insecure IoT devices, including a large-scale denial of service attack on the Internet's Domain Name System (DNS) infrastructure from the Mirai botnet. The resource-constrained nature of IoT devices introduces difficult challenges to deploying standard security and privacy solutions, as some functions are heavyweight and require significant power to compute or require significant bandwidth and energy to transmit. A potential solution is using the fog or cloud for certain security functions, e.g., filtering or aggregating traffic, or computing digital signatures. However, user privacy can be violated if the fog or cloud will process sensitive user data. Novel solutions are required to deal with the privacy/security issues. Software Defined Networking (SDN) technologies can potentially help network operators and consumers isolate individual devices on the network to ensure appropriate isolation between devices. Given the extensive body of past work on SDN and security in both the US and EU, applying these technologies to new domains such as the Internet of Things (IoT) in home, enterprise, and industrial networks, presents a rich set of research challenges.

3.3 Topics for Future Consideration

We identified a number of topics that are promising candidates for future EU-US research collaboration that we were not able to explore, either because the committee did not have enough people with the right expertise, or because of lack of time. These topics included:

- *Optical networking* to keep up with ever increasing demand for bandwidth
- *Packet processors* to perform network functions at line rate
- *Securing the network infrastructure*
- *Mobile offloading to clouds and edge computing*, a topic that cuts across the edge and core network infrastructures. Solutions must benefit mobile users as they travel, i.e.,

[10] Liu *et al.*, Tor Instead of IP. https://homes.cs.washington.edu/~tom/pubs/torip.pdf

[11] Hsiao, Hsu-Chun, et al. "LAP: Lightweight anonymity and privacy." *Security and Privacy (SP), 2012 IEEE Symposium on*. IEEE, 2012.

solutions adapt to different business and deployment models used in different countries and regions.

- *Managing edge networks*, specifically with an eye towards wireless edge networks that are traditionally unmanaged, e.g., home networks and hotspots. Similarly, solutions are needed that can adapt to different deployment models.
- *The use of SDN and NFV to optimize and customize wireless edge networks*. The NGI track discussed these topics for core networks, but different solutions will be needed for the wireless, where there is much more diversity.

4 Ideas for Collaboration

We discussed a number of ideas for collaboration models.

Learn from successful examples: It can be useful to look at some examples of successful collaboration between a relatively large number of geographically distributed teams:

The high-performance computing (HPC) community was used as an example of collaboration that we should learn a lot from. Mostly organized in the form of joint proposals, the modalities of collaboration include: joint training, platform sharing, joint platform development and performance prizes.

Within wireless community, existing collaboration on various projects (METIS, mmWave MAGIC, OAI, etc.) was pointed out as one of the success stories despite the fact that funding is quite limited. Most of the joint projects results of individual initiatives and are quite often funded by leveraging existing projects from respective funding agencies. One relatively successful area of collaboration was with student visits and exchanges but even that was quite limited. It was suggested that this modality of collaboration should be expanded to include co-supervision of masters and PhD students and even possibilities of double affiliation.

The success of existing joint platforms like 5TONIC and PlanetLab was used to illustrate the benefits of funded joint projects and platforms. Another example is the cross-Atlantic collaboration between FIRE and GENI, which is starting for possible future collaborative NGI projects outlined in the previous section.

Shared software and hardware platforms: More formalized sharing mechanisms for both software and hardware platforms can benefit the research community in many ways. First, it helps with avoiding duplication of work and support platform harmonization, but it typically requires financial support and considerable engineering. However, the fact that particular project/platform can get access to complementary expertise and experts on the other side of the ocean can be extremely beneficial, and can ultimately reduce cost.

An important concern is related to the availability of Open software tools (control/management/service), Open hardware (function programmable) and Common APIs. The workshop discussed means to organize a community effort that can produce the above. Examples of other environments were given such as ONF. For wireless, the case of OpenAirInterface was presented and discussed. This activity was identified as a core challenge for the partnership. For the NGI area, shared APIs for NFV platform as important example.

A number of platform sharing issues were raised during the presentations. The need for access rules harmonization including addressing fees and quotas was also mentioned. This extends to sharing of user groups and should be pushed beyond typical academic organizations to include

corporations, SMEs and verticals. The discussion also touched base on various other issues related to operation, access and monetization. They are related to hosting equipment, hosting services and hosting people.

Besides hardware, software, and APIs, several other opportunities for sharing were discussed. One low hanging “fruit” for both wireless and NGI is data; network logs, measurement data, as long as IPR issues can be resolved. Another example is shared ontologies, e.g., to represent a wide variety of testbed and platform resources, sensor information, etc.

Other sharing considerations: A number of other topics affecting collaboration were discussed:

- The need for highly focused workshops supporting brainstorming and in-depth discussions on technical subjects was also mentioned as one of the important collaboration opportunities.
- While travel cost was pointed out as main impediment for successful collaboration across the Atlantic, number of other impediments were also mentioned including time zone differences and lack of IPR agreements.
- The point related to access models was illustrated by several practices such as Open Calls (funded), Open access (non-funded), free but limited access to testbed resources (best effort support), Premium access (paid access, guaranteed access to testbed resources, guaranteed support).

5 Conclusions and Recommendations

This section presents the workshop's conclusions and recommendations.

5.1 Software tools, frameworks and platforms

The importance for the community to jointly develop and promote software tools to be part of the platforms as critical was identified. It was mentioned that an **Open source approach** should be adopted for some of the platform components. This could benefit from contributions of existing communities (including both open source and licensed; Spectrum Access Framework). It will certainly require a concerted action towards sharing tool sets. A gap analysis should be carried out for that purpose. In addition, we need to establish the **formal process** to enable this (the example of ONF was mentioned). One contribution could be to mobilize the verticals (domain specific open source) as stakeholders or resource providers.

Overall the **Open Data** (ODMP), reproducibility and interoperability features should be considered at a very early stage of the design.

The roles of platforms need to be clearly identified and they should assist new discoveries in the field of wireless/NGI at large. Their characteristics should be aligned with the challenging research questions in the field as raised by the various relevant communities (from the physical layer in wireless domain to the system design for both wireless and NGI domains). In addition, these platforms are meant to serve a broad set of actors: platform developers/providers, wireless/network researchers, application developers, end-users, etc. Technology-driven testbeds should offer state-of-the-art platforms (generally with a lower maturity level), while application-driven testbeds will include more well-proven technologies (often based on commercially available equipment) with a higher maturity levels. However, each role should be clearly identified and separated as platforms should be designed and operated to serve the needs of the research communities (from academia to industry). Different types of components will be necessary to deploy the platforms: the basic devices and elements of the infrastructure, the software piece that will enable control, programmability and access to the platform, the GUI and the back-end functionalities such as the open data dimension. These platforms should not be isolated as to attract a large set of applications and users. They might materialize through one or several platforms, as well as different level of platforms. However, what is important is to identify their commonalities that justify this global approach.

A joint initiative in this domain might cover various objectives. The following common work items were identified:

- Co-development of platforms
- Co-deployment of platforms

It has been stressed that the initiative should develop in various steps and time-scales according to the engagement of the community, maturity of concepts and availability of funding.

While the main objective is collaborative work on advanced platforms, it will not be possible to have parallel development of new infrastructure and/or large-scale co-deployment of new equipment in the short-term. Thus, the activity should focus on addressing common scientific challenges as well as on aligning effort on main methodologies and tools. Another option is to view this first step as a co-funding opportunity and to add (collaboration) funding to existing projects (platforms) and develop more synergy between the teams working on similar topics by supplementing existing research projects/grants and encourage collaborative experimentation on

the newly created advanced wireless platforms on both sides of the Atlantic. This should be similar to existing collaboration models that leverage existing projects like: Japan-US Network Opportunity (JUNO), Wireless Innovation between Finland and US (WiFiUS), as well as the Korea/EU, Japan/EU, Brazil/EU joint calls to name some.

5.2 Practical collaboration modalities

Common knowledge: It is mostly targeting the integration of the community, exchange of students and researchers, share of experiences and best practices, organizing common events.

- Explore joint research on the common work items identified above, from new physical and massive wireless to agile management.
- Identify target platforms to support the discovery
- Support both emulation and real world usability with a mix of emerging technologies and current state-of-the-art and avoid vertical stack silo only architectures by using open cloud-native and programmable platforms.

Common tools:

The objective is to identify some common development to avoid duplication of efforts and broaden the usefulness.

- Connect with other testbeds and open source communities, compute grid, identity management systems and data analytics platforms.

Common platforms:

This part is more ambitious and requires a more important support but should be targeted in the future. Its ambition is to join efforts to build the platform (and operate).

- Jointly develop platforms for experimentation with fully integrated 5G/NGI ecosystems covering topics from basic research to start-up launching.

Common usage/research:

The ultimate goal is to attract a large set of users from various organizations. Joint and interdisciplinary research/experiments:

- Develop and carry out cross Atlantic end-to-end technology trials that will use commonly developed platforms
- Focus on security and privacy, resilience, low latency of SDN/NFV architectures
- Initiate work on next generation of testbeds (e.g. bridging the gap between SDR and SDN research, QoE optimization for lrtge scale services, etc.)

5.3 Cooperation translated into timescales

As mentioned, this joint action should develop over time with a complementary set of objectives. They are illustrated below.

Short term: Driven mainly by underlying administrative complexities, the short term cooperation recommendations are based on supplemental funding for existing projects and include:

- Expand existing programs for exchanges of students (potentially expanding them to include joint mentoring). The expansion should also include support for exchange of design engineers and testbed operators as well as introduction of support for multi-

timescale faculty engagements in order to better support sharing of research (e.g. support for short and medium term project-related sabbaticals)

- Build on the existing FIRE-GENI collaboration to developed shared experimental NFV platforms, for example by leveraging existing hardware and software infrastructures (GENI, FIRA, OpenCloud). The agenda should include the development of shared APIs that will facilitate remote access and sharing of NFV code, and of shared ontologies to achieve interoperability and information sharing. These efforts will be an enabler for both shorter and longer term research. For example, this platform development can be driven by joint research projects, possible as collaborations between existing US and EU project, on geographic placement of various compute tasks (ranging from network functions to cloud computing services), experimentation with evolutionary clean slate and evolutionary network architectures, harmonization of wireless control frameworks, etc.

Medium term: Joint research proposals but with independent funding on each side

- Collaborate on the develop IXP and SDX research platforms based on shared interfaces and management software. Again this infrastructure could leverage existing nework testbed infrastructure. These efforts could be driven by research in new Internet control protocols and path-aware networking, which involves in horizontal and vertical network interfaces, tools for network monitoring, and privacy and trust issues raised by cross-Atlantic, or more general, cross-country communication. Joint projects (lightweight management), specific targets
- Common platforms environments (mmWave)/tools development
- Research projects on the four NGI core network technology research topics: SDX, NFV, “Horizontal” resource management, and User-centric interfaces, and the three cross-cutting topics: Research infrastructure (see also short term), Monitoring and measurement and privacy, trust and security.
- These research projects can then leverage the common platforms.

Long term: Joint proposals with common funding (like CERN)

- The above NGI testbed infrastructures would provide the basis for joint research for other research topics such as QoE optimization for large scale services.
- Initiatives like PAWR whether a single testbed jointly developed and deployed or two identical testbeds deployed on each side
- Commonly funded research projects on the four NGI core network technology research topics and the three cross-cutting topics.

Appendix A: Workshop Participants

Advanced Wireless Platforms:

Artur Azcorra	IMDEA Networks Institute and University Carlos III of Madrid
Suman Banerjee	University of Wisconsin-Madison
Pete Beckman	Argonne National Laboratory and Northwestern University
Serge Fdida	Université Pierre et Marie Curie (UPMC), Sorbonne University, LIP6 Laboratory & LINCS
Abhimanyu Gosain	Northeastern University
Edward Knightly	Rice University
Raymond Knopp	Eurecom
Tomas Magedanz	Fraunhofer Institute FOKUS and Technische Universität Berlin
Kobus Van Der Merwe	University of Utah
Ingrid Moerman	imec/Ghent University
Ari Pouttu	University of Oulu
Sundeep Rangan	New York University
Ivan Seskar	Rutgers University
Dimitra Simeonidou	University of Bristol
Rahim Tafazolli	University of Surrey
Leandros Tassiulas	Yale University

Next Generation Internet:

Aditya Akella	University of Wisconsin
Michael Boniface	IT Innovation
Cees De Laat	University of Amsterdam
John Domingue	Open University
Chip Elliott	Raytheon BBN Technologies
Sonia Fahmy	Purdue University
Chrysa Papagianni	National Technical University of Athens
George Rouskas	North Carolina State
Peter Steenkiste	Carnegie Mellon University (co-chair)
Renata Cruz Teixeira	Inria
Brian Tramell	ETH Zurich
Brecht Vermeulen	imec/Ghent University (co-chair)
Kuangching ‘KC’ Wang	Clemson University

NSF and DG CONNECT:

Remy Bayou	DG CONNECT
Per Blixt	DG CONNECT
Jack Brassil	NSF
Kenneth L Calvert	NSF
Thyaga Nandagopal	NSF
Georgios Tselentis	DG CONNECT