Proceedings of the SwitchOn Workshop

University of Sao Paulo Medical School, Sao Paulo, Brazil
October 15-16, 2015

TABLE OF CONTENTS

PRELIMINARY PROGRAM

Table of Contents

DAY 0 (Wednesday, 14 October 2015)

20:00-22:00 Reception (Renaissance Marriott Sao Paulo)

DAY 1 (Thursday, 15 October 2015)

8:30-8:50 Welcome: Luis Fernandez Lopez (USP) and John Brassil (NSF)

8:50-9:20 Report on First SwitchOn Workshop in Miami, Julio Ibarra and Jason Liu (FIU)

9:20-10:05 Keynote: Lisandro Zambenedetti Granville (UFRGS), Session Chair: Julio Ibarra (FIU)

TITLE: Reviewing Brazilian International Collaboration Programs on IT: Challenges and Lessons Learned

ABSTRACT: In the last years, the Brazilian government has been financially supporting a number of research projects selected from joint international open calls, especially with Europe. Important areas have been covered in such projects, such as Cloud Computing, HPC, and Future

Internet. Joint projects with international partners lead to unique opportunities to Brazilian researchers, although they also pose challenges that are not often found in traditional, national projects. Opportunities and challenges are also faced by funding agencies, that need to assist researchers to achieve the full potential of international collaborations. In this talk, we will review the last international joint calls with Brazil, discussing their benefits, experiences, and lessons learned. We will also observe the faced difficulties and how such difficulties have been addressed, and finally we will provide considerations about future opportunities of international collaborations.

SPEAKER'S BIO: Lisandro Zambenedetti Granville is director of the Center for Research and Development in Digital Technologies for Information and Communications (CTIC/RNP), president of the Brazilian Computer Society (SBC), associate professor on Computer Networks at the Institute of Informatics of the Federal University of Rio Grande do Sul (UFRGS), chair of the Committee on Network Operations and Management (CNOM) of the IEEE Communications Society, co-chair of the Network Management Research Group (NMRG) of the Internet Research Task Force (IRTF), and member of the Brazilian Internet Steering Committee (CGI.br). As Director of CTIC/RNP, Lisandro is responsible for running the processes of advertising, selecting, assisting, and managing research projects developed by Brazilian universities and research centers. As a researcher, Lisandro's topics of interest include network management, dynamic circuit networks, software-defined networking, and network functions virtualization.

10:05-10:30 Coffee Break

10:30-12:30 PICO Session, Moderator: Heidi Morgan (FIU)

12:30-14:00 Lunch

14:00-16:00 Technical Session (I), Session Chair: TBA

1. Expanding US/Brazil FIDC Testbed Federation, by Mark Berman (BBN)

ABSTRACT: Future Internet and distributed cloud (FIDC) testbed researchers in Brazil and the US are among the first participants in a developing worldwide trend toward greater federation of research cyberinfrastructure. Federation of FIDC testbeds offers a promising value proposition for both operators and end users, offering scale, geographic diversity, and resource heterogeneity, while maintaining local control over resources, policies, and cost.
Federation of FIDC testbeds occurs at several levels, but the first (and often most challenging) step is the human-to-human and organization-to-organization process of establishing trust and resource sharing agreements and the associated implementation policies. Once this step is addressed, the technical implementation may include some or all of AAA (authentication, authorization, and accounting) federation, control plane federation, and data plane federation. Established or emerging standards exist for each of these aspects of technical federation, at various levels of maturity.
Thanks to strong researcher-to-researcher relationships, there are a number of very compelling examples of federation between Brazilian and US testbeds, supporting a variety of research collaborations. There is great potential in pursuing two promising paths for expanding the power of these federations. One path is to agree upon and implement policies that simplify human-level processes, moving toward routine, automated instantiation of federated infrstructure configurations. A second path is to continue to expand research and infrastructure collaborations to a broader set of participating countries and regions.

SPEAKER'S BIO: Mark Berman is Vice President for Technology Development at BBN and Project Director for NSF's GENI project, with overall responsibility for GENI's technical direction and successful implementation. At BBN, Mark works to bring technical innovations into practical use. Mark's research interests are in the area of complex distributed systems and their usability. As VP of BBN's Intelligent Computing business, he exercised management oversight for numerous research and technology transition efforts, across a diverse set of disciplines and advocated for an open source approach to technology transition. Mark holds A.B. and S.M. degrees in computer science from Harvard.

2. Chameleon: A Large-scale, Reconfigurable Experimental Environment for Next Generation Cloud Research, by Kate Keahey (U. Chicago)

ABSTRACT: While cloud services have become ubiquitous to all major 21st century economic many open questions remain on how to best assimilate and leverage this new technology. In particular, many open research questions concern the relationship between cloud computing and high performance computing, the suitability of cloud computing for data-intensive applications, and its position with respect to emergent trends such as Software Defined Networking. A persistent barrier to further understanding of those issues has been the lack of a large-scale testbed where they can be explored. This is partly because a suitable experimental platform would require large investment in hardware and support for deep reconfigurability to provide the required support for reliable and reproducible research.
With funding from the National Science Foundation (NSF), the Chameleon project provides such platform to the open research community allowing them to explore transformative concepts in deeply programmable cloud services, design, and core technologies. The testbed, deployed at the University of Chicago and the Texas Advanced Computing Center, will ultimately consist of almost 15,000 cores, 5PB of total disk space, and leverage 100 Gbps connection between the sites. While the primary investment in hardware deployed so far consists of primarily of homogenous hardware to support large-scale experiments, future capabilities will support heterogeneous units allowing experimentation with high-memory, large-disk, low-power, GPU, and co-processor units. To support a broad range of experiments, the project provides a configuration system allowing full control of the software stack, from provisioning of bare metal and network interconnects to delivery of fully functioning cloud environments. This talk will describe the goals, the building, and the capabilities of the testbed.

SPEAKER'S BIO: Kate Keahey is one of the pioneers of infrastructure cloud computing. She created and leads the development of the Nimbus project, recognized as the first open source Infrastructure-as-a-Service implementation, and engages in many application leveraging cloud computing capabilities for science. She also leads the Chameleon project, a distributed experimental platform for cloud computing research. Kate is a Scientist at Argonne National Laboratory and a Senior Fellow at the Computation Institute at the University of Chicago.

3. Chameleon: A Self-Adaptative Multi-Cloud Testbed, by Carlos R. Senna and Edmundo R. M. Madeira (UNICAMP)

ABSTRACT: Nowadays the cloud computing technology allows us to integrate local resources (networks, grids and private clouds) to intensive computing centers (data centers), forming a hybrid and heterogeneous environment accessible through the Internet. However, new emerging technologies, notably Internet of Things (IoT), bring with them new requirements to support mobility, geographical distribution, location awareness and low latency, requiring a new platform. To address these new challenges the cloud is migrating to the network edge. This evolution of the cloud demands a new paradigm that extends the cloud computing services to the

edge of the network, to meet a new generation of services and applications. In this talk, I will present Chameleon, the LRC (Computer Network Laboratory)/UNICAMP multicloud testbed, designed to meet the IoT requirements. Chamaleon consists of a hardware infrastructure with the capacity to hostmultiple cloud technologies and a software infrastructure to orchestrate the execution of distributed applications, integrate technologies and facilitate cooperation between the various levels of multicloud environments. Several applications that use Chamaleon will be commented, such as image processing applications, scientific workflows, sensor network applications, vehicular applications and Hadoop applications.

SPEAKER'S BIO: Edmundo R. M. Madeira is a Full Professor at the University of Campinas (UNICAMP), Brazil. He received his Ph.D. in Electrical Engineering from UNICAMP in 1991. He has published over 150 papers in national and international conferences and journals. He was the General Chair of the 7th Latin American Network Operation and Management Symposium (LANOMS'11), and he was a Technical Program Co-chair of the IEEE LatinCloud'12. He is a member of the editorial board of Journal of Network and Systems Management (JNSM), Springer. His research interests include network management, future Internet and cloud computing.

4. BAMBU: A Metropolitan Innovation Testbed for Promoting Future Internet Research, by Leobino N. Sampaio, Luciano Rebouças, and Marcos Barreto (UFBA), Romildo Martins and Allan E. S. Freitas (IFBA)

ABSTRACT: A broad number of large-scale network testbeds have been proposed in the last couple of years as a mean of better supporting Future Internet research. This paper present Bambu, a metropolitan innovation testbed for promoting experimental researches in the city of Salvador, Bahia.

SPEAKER'S BIO: Leobino N. Sampaio is Assistant Professor at Federal University of Bahia (UFBA) since 2010. He holds Bachelors in Computer Science from Salvador University (UNIFACS) in 1996, and Master Science from Salvador University (UNIFACS) in 2002 and PhD in Computer Science by Federal University of Pernambuco (UFPE) in 2011. He has a large experience in computer networks, in particular, network measurements and performance evaluation. Currently, his research interests are software-defined networking, mobile communication, content-centric networks, and energy efficiency in data center networks.

5. From Federated Software Defined Infrastructure to Future Internet, by Kuang-Ching Wang (Clemson)

ABSTRACT: Since 2009, Clemson University embarked on a series of NSF projects that led its researchers and IT engineers through an exciting learning curve for software defined networking (SDN) and cloud computing technologies. From academic explorations of NSF GENI, CC-NIE, FutureCloud to production adoption of SDN-based data center fabric, we have worked on SDN and cloud programming and integration at data center, campus, regional, national, and international scope. It is our belief that software defined infrastructures (SDIs) around the world will be increasingly federated and interoperable, allowing applications to be launched from anywhere, with highly customizable communication with other services near and far. As a result, their logical requirements, both locally and globally, will be driving the federated SDI architecture, which will shape the future Internet.

SPEAKER'S BIO: Dr. KC Wang is Networking CTO and Associate Professor of Electrical and Computer Engineering at Clemson University. From 2012 to 2015, KC was also a Member of the Technical Staff at Big Switch Networks developing SDN and OpenStack. As a researcher, KC has been PI for multiple NSF projects on SDN and future cloud computing systems. As the CTO, KC drives the roadmap and implementation of Clemson's production network and data center infrastructure. In 2012, KC was on sabbatical with Stanford's Open Networking Lab and jointly developed ON.Lab's SDN-BGP peering solution.

6. Network Virtualization and Testbeds at AmLight SDN, by Jeronimo Bezerra (FIU)

ABSTRACT: In 2014, OpenFlow was enabled at AmLight network to create a Software-defined Networking infrastructure, focusing on provisioning optimization and network virtualization support. The virtualization feature has been deployed to support experimental testbeds, providing a platform for innovation for researchers in the U.S. and Latin America. Even though AmLight is a production infrastructure, experimentation has been widely supported, and researchers encouraged to use it for prototyping their applications. Since 2014, more than eight testbeds were created in parallel with production applications. This presentation plans to introduce AmLight SDN resources available to support new network testbeds between the U.S. and South America, as well describe future actions to enhance academic collaborations.

SPEAKER'S BIO: MSc in Mechatronics and BS in Computer Science by the Federal University of Bahia/Brazil, Jeronimo Bezerra has been involved with academic networks for the last 12 year. Currently, he is the AmLight Senior Network Engineer (www.amlight.net), responsible for the AmLight Software Defined Network and all its applications, protocols and performance requirements.

16:00-16:30 Coffee Break

16:30-18:30 Technical Session (II), Session Chair: TBA

1. Looking Beyond the Internet - Next Steps, by Chip Elliott (BBN)

ABSTRACT: The technical structure of the Internet is now starting to undergo a major transformation into its next generation, which many are calling Software Defined Infrastructure (SDI). This transformation is opening up major new opportunities for research. We will briefly discuss this transformation as well as the next steps being planned from a United States perspective.

SPEAKER'S BIO: Chip Elliott is Principal Investigator for GENI, a nationwide suite of experimental infrastructure being created by the National Science Foundation for at-scale research in future internet architectures, services, and security. He is Chief Scientist at Raytheon BBN Technologies, Adjunct Professor at Dartmouth College, a Fellow of the AAAS, ACM, and IEEE, and an active inventor with 90 patents. Mr. Elliott has served on many national panels and has held visiting faculty positions at Tunghai University in Taiwan and the Indian Institute of Technology, Kanpur.

2. Advanced Infrastructure and Applications for Interoperable and Interconnected Smart Cities, by Glenn Ricart (US Ignite)

ABSTRACT: The NSF has recently initiated a grant to US Ignite to add a locavore infrastructure in 15 gigabit US cities and to create a self-sustaining set of innovative, interoperable, and interconnected smart city applications made feasible by that infrastructure. The outline of that infrastructure will be presented and an invitation issued to Brasilian cities, universities, and funding agencies, to research, develop, deploy, and interconnect a Brasilian-customized version. Especially valuable will be projects which involve both US and Brasilian researchers working on bleeding-edge applications and/or the infrastructure to support them which would be researched, developed, deployed and tested in both countries. Other potential international partners including Europe and Japan are also considering such joint projects.

SPEAKER'S BIO: Glenn Ricart is founder and CTO of US Ignite, a nonprofit creating an ecosystem of innovative applications and services in smart gigabit communities enabled by software-driven infrastructure. US Ignite brings together corporations, cities and regions, government, nonprofits, and academic to make and deploy these innovative applications and services. Dr. Ricart is also Adjunct Professor in the School of Computing at the University of Utah. Previously, Glenn was CEO of National LambdaRail, Managing Director of PricewaterhouseCoopers, Executive Vice President and CTO of Novell during its heyday, Program Manager at DARPA, and Assistant Vice Chancellor of Academic Affairs and academic CIO at the University of Maryland. Dr. Ricart was inducted into the Internet Hall of Fame as an Internet Pioneer in 2013. He is also the founder or co-founder of five startups raising a combined total of $190M.

3. Flexible Infrastructure Management for Smart Societies, by Carlos Alberto Kamienski (UFABC)

ABSTRACT: Smart cities require a variety of new services to be available to citizens of large urban spaces, based on data coming from different sources. This will pose high demands for dynamic allocation of existing computing and communication infrastructure that currently are based on datacenters. We advocate that the management of dynamic configuration mechanisms for cloud elasticity, SDN and NFV is a key feature for enabling such new scenarios. Encouraging results with dynamic datacenter reconfiguration will be used for providing services for smart cities, such as dealing with high and varying density of people in particular locations.

SPEAKER'S BIO: Carlos Alberto Kamienski received his Ph.D. in computer science from the Federal University of Pernambuco (Recife PE, Brazil) in 2003. He is an associate professor of computer science at the Federal University of ABC (UFABC) in Santo André SP, Brazil, where currently he holds the position of director of international relations. His current research interests include cloud computing, software-defined networks, Future Internet, network function virtualization, analysis of online social networks and smart cities.

4. Cybersecurity and Privacy and the Human Factors Behind It, by Daniela Oliveira (UFL)

ABSTRACT: Computer systems are everywhere. Society is dependent on networked computer systems and software for everything from running the smart grid that powers your house to social networking with your family and friends. The societal dependency has raised new responsibilities for scientists to develop sound methods and practices to protect computer systems from attack, as well as protect the user against privacy threats created by ubiquitous computer systems. The methods and practices must consider both threats from attackers looking to compromise systems, as well as the privacy and human factors considerations new systems bring to users. Cyber security scientists and engineers are also faced with the daunting task of not just securing systems against lone individuals, but also nation states using computer systems to carry out their political and military agendas. In this talk, I will discuss these challenges from the viewpoint of implementing security across layers of abstraction, and designing systems that protect user privacy and that take into account the social, behavioral and economic implications of security.

SPEAKER'S BIO: Daniela Oliveira is an Associate Professor in the Department of Electrical and Computer Engineering at the University of Florida. She received her B.S. and M.S. degree in Computer Science from the Federal University of Minas Gerais in Brazil. She then earned her Ph.D. in Computer Science from the University of California at Davis. Her main research interest is interdisciplinary computer security, where she employs successful ideas from other fields to make computer systems more secure. Her current research interests include employing warfare deception strategies to protect OS kernels. She is also interested in understanding the nature of software vulnerabilities and social engineering attacks, especially how these attacks can be prevented from a Psychological perspective. She received an NSF CAREER Award in 2012 for her innovative research into operating systems' defense against attacks using virtual machines and the 2014 NSF Presidential Early Career Award for Scientists and Engineers (PECASE).

5. NovaGenesis Project: Overview, Current Activities and Research Opportunities, by Antonio Marcos Alberti (INATEL)

ABSTRACT: This talk provides an overview of the NovaGenesis convergent information project, summarizing its current research and development activities. It also presents a list of research opportunities – open problems – that could be addressed by our team together with other institutions in joint research initiatives.

SPEAKER'S BIO: Antônio Marcos Alberti received the degree in Electrical Engineering from Santa Maria Federal University (UFSM), Santa Maria, RS, Brazil, in 1986, and the M.Sc. and Ph.D. degrees in Electrical Engineering from Campinas State University (Unicamp), Campinas, SP, Brazil, in 1998 and 2003, respectively. In February 2004, he joined National Institute of Telecommunications (INATEL), Brazil, as an Adjunct Professor. He was a visiting researcher at Future Internet Department at Electronics and Telecommunications Research Institute (ETRI), South Korea, from March 2012 up to February 2013. He is the head of the INATEL's ICT Lab and chief architect of NovaGenesis convergent information architecture project. His current interests include Future Internet Design, Internet of Things, Cloud Networking, Service-Defined Architectures, 5G, Cognitive Radio and Cognitive Radio Networks.

6. ANSP advanced infrastructure, by Luis Fernandez Lopez (USP)

ABSTRACT: The ANSP project (an Academic Network at São Paulo), as determined by a decision of FAPESP's Board of Trustees, "provides the State of São Paulo research community with state of the art computer networking connectivity.ANSP develops and maintains the infrastructure, Internet services and data communication in general which, in conjunction with ANSP's Technical Reserve Program, offers the education and research community in the State of São Paulo the technological means needed to access information throughout the world, to share knowledge, to develop collaborative projects and for innovation on a large scale.

SPEAKER'S BIO: Prof. Lopez holds PhD in Mathematical Physics. Currently he is a professor at USP (Medicine School, University of São Paulo) and at FIU (Florida International University – Miami, USA). He is also the NARA Coordinator (Center for Advanced Networking Applications) of USP and Principal Investigator of the Project ANSP (Academic Network at São Paulo), funded by FAPESP (Foundation for Research Support of the State of São Paulo) and NSF (National Science Foundation).

20:00 Group Dinner

DAY 2 (Friday, 16 October 2015)

8:30-9:15 Keynote: Christos Papadopolous (Colorado State U.), Session Chair: Jason Liu (FIU)

TITLE: Named Data Networking: An Internet Architecture for the Future

ABSTRACT: The current Internet names the hosts, leaving it to the application to locate the host with the desired data. However, with the emergence of technologies such as CDNs and the cloud, and trends such as mobility and IoT, the need to associate data with an IP address has become a hindrance. This misalignment requires enormous corrective effort at the expense of application complexity and robust security.
Named Data Networking (NDN) transforms the current network of hosts into a network of data objects. In the process, application complexity is substantially reduced and data is easily secured. A named data network is an efficient content distributor that can natively support communication models such as multicast and anycast, routing models such as multipath and data operations such as publishing, discover and caching. In general, NDN offers a network service model that aligns better with user needs, building a stronger foundation for current and future applications.
This talk will present NDN and some illustrative applications.

SPEAKER'S BIO: Christos Papadopoulos is currently a professor at Colorado State University. He received his Ph.D. in Computer Science in 1999 from Washington University in St. Louis. In 2002 he received a NSF CAREER award to explore router services as a component of the Internet architecture. He currently works on future Internet architectures, network security and measurements. Active projects include Named Data Networking (NDN), Supporting Climate Applications over NDN, Netbrane, PREDICT, Retro-Future, WIT-II and Making Internet Routing Data Accessible to All. He is a senior IEEE member and has served on numerous ACM and IEEE conference program committees.

9:15-10:15 Panel on US and Brazil Funding Opportunities Moderator: Julio Ibarra (FIU)

10:15- 10:45 Coffee Break

10:45-12:15 Breakout Discussion, Moderator: Julio Ibarra (FIU)

12:15-13:45 Lunch

13:45-14:15 Breakout Group Reports, Moderator: Jason Liu (FIU)

14:15-16:15 Technical Session (III), Session Chair: TBA

1. Some Research Problems in Future Networking, by Deep Medhi (UMKC)

ABSTRACT: In this talk, I'll touch briefly on a number of issues. Here, I highlight two:

Federated networking in the future presents a number of challenges. In a new collaborative project between Brazil, US, and Belgium, we are investigating a wide-area federated network environment where applications require deterministic network services (such as high quality video conferencing services). We are exploring research as well as technical challenges in establishing federation and coordination between three different sites located in three different continents. In particular, we are exploring how software-defined networking can be used in this federated environment to give service guarantees.

Virtual machine migration is an important requirement in a data center networking to allow resource flexibility. We want to understand how migration works with the increase the speed of underlying networks in a data center. Our initiation assessment is that scheduling algorithms needs to be network-speed aware since the network speed reduces the time taken to migration VMs. In this collaborative work, we are considering new scheduling algorithms that are network-aware to show how performance can be improved. We are further continuing this work to consider the topological structure of the data center networks in addition to network speed.

SPEAKER'S BIO: Deep Medhi is Curators' Professor in the Department of Computer Science & Electrical Engineering at the University of Missouri-Kansas City, USA, and honorary professor in the Department of Computer Science & Engineering at the Indian Institute of Technology-Guwahati, India. He obtained B.Sc. in Mathematics from Cotton College, Gauhati University, India, M.Sc. in Mathematics from the University of Delhi, India, and his M.S. and Ph.D. in Computer Sciences from the University of Wisconsin-Madison, USA. Prior to joining UMKC in 1989, he was a member of the technical staff at AT&T Bell Laboratories. He is the Editor-in-Chief of Springer's Journal of Network and Systems Management, and is on the editorial board of IEEE/ACM Transactions on Networking, IEEE Transactions on Network and Service Management, and IEEE Communications Surveys & Tutorials. His research has been funded by NSF, DARPA, and industries. He is co-author of the books, Routing, Flow, and Capacity Design in Communication and Computer Networks (2004) and Network Routing: Algorithms, Protocols, and Architectures (2007), both published by Morgan Kaufmann Publishers/Elsevier Science.

2. Towards an Emulator for Software Defined Wireless Networks, by Christian Esteve Rothenberg (UNICAMP)

ABSTRACT: In this talk, we will present Mininet-WiFi as a tool to emulate wireless OpenFlow/SDN scenarios allowing high-fidelity experiments that replicate real networking environments. Mininet-WiFi augments the well-known Mininet emulator with virtual wireless stations and access points while keeping the original SDN capabilities and the lightweight virtualization software architecture. We will present a demo use case in a mobile video streaming scenario to showcase the ability of Mininet-WiFi to emulate the wireless channel in terms of bandwidth, packet loss, and delay variations as a function of the distance between the communicating parties.

SPEAKER'S BIO: Christian Esteve Rothenberg is an Assistant Professor in the Faculty of Electrical and Computer Engineering at University of Campinas (UNICAMP), where he received his Ph.D. in 2010 and currently leads the INTRIG (Information & Networking Technologies Research & Innovation Group). From 2010 to 2013, he worked as Senior Research Scientist in the areas of IP systems and networking at CPqD Research and Development Center in Telecommunications (Campinas, Brazil), where he was technical lead of R&D acitivities in the field of OpenFlow/SDN such as the RouteFlow project, the OpenFlow 1.3 Ericsson/CPqD softswitch, or the ONF Driver competition. Christian holds the Telecommunication Engineering degree from Universidad Politecnica de Madrid (ETSIT – UPM), Spain, and the M.Sc. (Dipl. Ing.) degree in Electrical Engineering and Information Technology from the Darmstadt

University of Technology (TUD), Germany, 2006. Christian holds two international patents and has over 50 publications including scientific journals and top-tier networking conferences such as SIGCOMM and INFOCOM. Since April 2013, Christian is an ONF Research Associate.

3. Scalable Emulation of SDN Applications with Simulation Symbiosis, by Cesar Marcondes (UFSCAR) and Jason Liu (FIU)

ABSTRACT: Mininet is a popular container-based emulation environment built on Linux for testing OpenFlow applications. Using Mininet, one can compose an experimental network using a set of virtual hosts and virtual switches with flexibility. However, it is well understood that Mininet can only provide a limited capacity, both for CPU and network I/O, due to its underlying physical constraints. We propose a method for combining simulation and emulation to improve the scalability of network experiments. This is achieved first by applying the symbiotic approach to effectively integrate emulation and simulation for hybrid experimentation. For example, one can use Mininet to directly run OpenFlow applications on the virtual machines and software switches, with network connectivity represented by detailed simulation at scale. We also propose a method for using the symbiotic approach to coordinate separate Mininet instances, each representing a different set of the overlapping network flows. By effectively distributing network emulation among separate machines, one can significantly improve the scalability of the network experiments.

SPEAKER'S BIO: Jason Liu is an Associate Professor at Florida International University in Miami, Florida, USA. He received a B.A. degree in Computer Science from Beijing University of Technology in China in 1993, an M.S. degree in Computer Science from College of William and Mary in US in 2000, and a PhD degree in Computer Science from Dartmouth College in US in 2003. His research interests include parallel discrete-event simulation, high-performance modeling and simulation of communication networks and computer systems. Dr. Liu is currently a steering committee member for the ACM SIGSIM-PADS conference and an associate editor for ACM TOMACS and SIMULATIONS journals. He was a general chair for PADS'12, MASCOTS'10, and SIMUTools'11, and a program chair for SIMUTools'10 and PADS'08. He also served on the technical committees for many conferences. Dr. Liu is an NSF CAREER Awardee in 2006 and an ACM Distinguished Scientist in 2014.

4. Abstract Topology and Cost Maps for Software-Defined Inter-Domain Circuits, by Malathi Veeraraghavan (UVA) and Christian Esteve Rothenberg (UNICAMP)

ABSTRACT: BGP software and policy configuration are widely acknowledged to be highly complex. A research question of interest in whether simplifications are possible to inter-domain routing protocols for best-effort IP service in the new SDN paradigm.
There is an effort called Application Layer Traffic Optimization (ALTO) underway in the IETF to define a protocol that enables applications to request and receive information such as network topology, path costs, link availability, routing policies, and end-host properties. Providers can offer network information in abstracted form since the ALTO protocol offers topology hiding. The ALTO protocol was proposed for IP-routed best-effort services. We propose to explore whether the ALTO protocol can be extended to support advance-reservation dynamic L2-path and L1-circuit services.

SPEAKER'S BIO: Malathi Veeraraghavan is a Professor in the Charles L. Brown Department of Electrical & Computer Engineering at the University of Virginia (UVa). Dr. Veeraraghavan received her BTech degree from Indian Institute of Technology (Madras) in 1984, and MS and PhD degrees from Duke University in 1985 and 1988, respectively. After a ten-year career at Bell

Laboratories, she served on the faculty at Polytechnic University, Brooklyn, New York from 1999-2002, where she won the Jacobs award for excellence in education in 2002. She served as Director of the Computer Engineering Program at UVa from 2003-2006. Her current research work on optical and IP networks is supported by NSF and DOE. She holds twenty-nine patents, has over 100 publications, and has received five Best-paper awards. She is currently serving as TPC Co-Chair of IEEE BlackSeaCom 2016. She served as the Technical Program Committee Chair for IEEE ICC 2002, and Associate Editor for the IEEE/ACM Transactions on Networking.

5. The Capacity Crunch: Avoidance, Adaptation and Reform, by Helio Waldman (UNICAMP)

ABSTRACT: Our research adresses the need to provide more capacity to accommodate the ever increasing demand for bandwidth generated by Internet traffic. At the current stage, we investigate the ability of elastic networking to provide a response to this need. The next stage will be the deployment of a new infrastructure based on a new generation of optical fibers that support space division multiplexing (SDM). Networking solutions for SDM are under investigation, as well as business and regulatory issues related with such emerging infrastructures.

SPEAKER'S BIO: Helio Waldman is a Collaborating Professor at Universidade Estadual de Campinas (UNICAMP), where he retired in 2006 as a Full Professor (Professor Titular). Between 1986 and 1990, he served as the first Research Provost of UNICAMP, after a 4-year term as the Director of its old Campinas Engineering School (FEC), which originated its current Schools (Faculdades) of Computer and Electrical Engineering (FEEC), Mechanical Engineering and Chemical Engineering. In 2006, he joined the staff of the newborn Universidade Federal do ABC (UFABC), where he reached compulsory retirement in 2014, at the age of 70, after serving a 4-year term as President (Reitor). He received an Engineer degree from Instituto Tecnológico de Aeronáutica (ITA) in 1966, and M.Sc. and Ph.D. degrees from Stanford University in Stanford, California, USA, in 1968 and 1971 respectively. His research interests include modelling of the routing and spectrum assignment (RSA) in emerging elastic optical networks (EON's), survivability of optical networks, modelling of optical buffers, and application of Game Theory to the modelling of competitive and regulatory strategies in the Internet and other public utility environments.

6. RNP national network infrastructure and support for Future Internet research, by Michael Stanton (RNP)

ABSTRACT: We will describe the network infrastructure used to provide services for RNP's growing user base throughout Brazil. These services are based on a national backbone network linking points of presence in every state, with international connectivity to the rest of Latin America, to the United States and to Europe. This structure provides network access to over 1200 research and higher education sites, providing them with quality connectivity to both academic and commercial sites both in Brasil and throughout the world. RNP also provides specific support for using the network as a laboratory for both present and future Internet experimentation, and this will also be presented.

SPEAKER'S BIO: Michael Stanton is Director of Research and Development at RNP, the Brazilian national research and education network. After a PhD in mathematics at Cambridge University in 1971, he has taught at several universities in Brazil, since 1994 as professor of computer networking at the Universidade Federal Fluminense (UFF) in Niterói, Rio de Janeiro state, until his retirement in 2014. Between 1986 and 1993, he helped to kick-start research and education networking in Brazil, including the setting-up and running of both a regional network in Rio de Janeiro state (Rede-Rio) and RNP. He returned to RNP in 2001, with responsibility for R&D and RNP involvement in new networking and large-scale collaboration projects. Since 2012, Michael's new international activities have included representing RNP on the Global CEO Forum's Global Network Architecture (GNA) group, engaged in the design of the architecture of the future academic Internet, and joining the Council of the Research Data Alliance in 2014.

16:15-16:30 Coffee Break

16:30-17:50 Technical Session (IV), Session Chair: TBA

1. Experiences supporting data intensive science flows on campus using a Science DMZ, by Julio Ibarra (FIU)

ABSTRACT: The network of a campus or a research enterprise supports multiple business operations: Network traffic associated with administration,   The network must be operated with security that protects the financial and personnel data of the organization. At the same time, these networks also support scientific research processes. Researchers depend on the network infrastructure to share, store, and analyze research data from many different external sources. Campus networks are normally not engineered to support the data movement requirements of data intensive science. Data intensive science applications experience poor performance on general purpose networks. Campuses serve multiple communities of interest. As a result, they are faced with the following challenge: How can campus networks be adapted to optimize flows from science applications without impacting the operation of their production networks?
This talk explores the concept of the Science DMZ as a methodology to reengineer the campus network to support the needs of multiple communities of interest, in particular, the science research community conducting data intensive science, alongside production operation. The implementation of a Science DMZ at FIU's campus is described. It is a component of a project called FlowSurge: a coordinated and coherent approach to bridging among different cyberinfrastructures by building a Science DMZ at FIU that can be tailored to support the needs of high-performance science applications and data-intensive flows. Researchers, from their laboratory, will be able to (1) generate science flows to shared CI or CI- connected resources on the campus; at other campuses; or at regional, national or international sites. (2) Orchestrate and participate in "at-scale" experiments by using FlowSurge to tailor the end-to-end path for optimized high-performance science applications and data-intensive flows. FlowSurge will enable IT and network staff at FIU to partition the campus network infrastructure to support Science flows without causing impact to general-purpose production flows on the campus.
SPEAKER'S BIO: As the Assistant Vice President for Technology Augmented Research at FIU, Dr. Julio Ibarra is responsible for furthering the mission of the Center for Internet Augmented Research and Assessment (CIARA) – to contribute to the pace and the quality of research at FIU through the application of advanced Cyberinfrastructure. He is responsible for strategic planning and development of advanced research networking services, including the development and management of the AMPATH International Exchange Point for Research and Education networks, in Miami, Florida. He holds B.S. and M.S. in Computer Science from FIU, and Ph.D. in Telematics and Information Technology from Twente University.

2. US West-coast Future Internet Infrastructure; Pacific Research Platform, Wave SDX & PEERING, by John Silvester (USC)

ABSTRACT: This year marks important developments in International Research and Education Network Infrastructure for Future Internet and High-Performance Data-Driven Science on the

United States West Coast. Pacific Wave (which operates the West Coast International R&E Exchange points) has recently been awarded funding from the US National Science Foundation to upgrade its capabilities and deploy SDX functionality across its facilities in Seattle, Los Angeles, and Sunnyvale (San Francisco Bay Area) and to add access points in several other locations. This is an integral component of the international effort to interconnect research and education networks using Software Defined Networking (SDN). In addition, NSF has funded a $5 million, five-year award to UC San Diego and UC Berkeley and their collaborators to establish a Pacific Research Platform (PRP), a science-driven high-capacity data-centric “freeway system” extending the notion of a DMZ to a large regional scale. This project plans to utilize the Pacific Wave infrastructure to expand the boundaries of campus DMZ’s across a wide geographic area and eventually to international partners. The PRP will give participating universities and other research institutions the ability to move data 1,000 times faster compared to speeds on today’s inter-campus shared Internet and to establish (software defined and managed) virtual networks to support multiple science domains communities across the network.
In this talk, we will discuss these developments and the associated measurement infrastructure that is being deployed to monitor network performance and troubleshoot bottlenecks. We will also describe some of the international projects that utilize the infrastructure.

SPEAKER’S BIO: John Silvester has been with the University of Southern California since 1979 and is currently a Professor of Electrical Engineering. He was Director of the Computer Engineering Division from 1984-1990, Vice-Provost of Academic Computing from 1994-1997, and Vice-Provost of Scholarly Technology from 1997-2006. He received an MA in Mathematics and Operations Research from Cambridge University, MS in Statistics and Computer Science from West Virginia University and Ph.D. in Computer Science from UCLA. His academic interests are in high speed, optical, and wireless networking. He has supervised more than 25 Ph.D. students and published over 150 technical papers. For the last 10 years, he has been actively involved in planning, management and implementation of advanced networking for the research and education community at the State, National, and International level. He was a member of the Board of Directors of CENIC (Consortium for Educational Networks in California), was Vice-Chair from its founding in 1997 to 1999 and Chair from 1999 to 2006. From 2005-2010 he was Principal Investigator of the TransLight/PacificWave project (part of the International Research Networks Connections Program of the US National Science Foundation), that operates the research and education exchange points on the US West Coast and, in partnership with the Australian Academic and Research Network, supports advanced connectivity through Hawaii to Australia. He continues as a consultant for CENIC on international projects with a particular focus on the Pacific Rim and Latin American partners that utilize Pacific Wave (the main academic exchange on the US West Coast), which recently received funding from the NSF to upgrade its facilities to 100G and to deploy a parallel SDX capability. He is co-PI on the NSF Network Infrastructure project TEN-II (Trojan Network II) which is experimenting with SDN capabilities and implementing a “DMZ” to support high performance networking at USC. Participants in TEN-II are part of the Pacific Research Platform based at UCSD (recently funded by the NSF to expand DMZ functionality across multiple campuses to support multiple disciplinary science domains).

3. The Policy Space for Software Defined eXchanges, by Joaquin Chung Miranda, Joaquin Chung, Russell Clark, Henry Owen (GaTech)

ABSTRACT: An Exchange Point is a common place where different organizations collaborate and share computing, storage and networking resources. Recently, Software Defined Networking technologies are being included into Exchange Points to optimize resource sharing and allocation. In this talk, we explore the policy space for a Software Defined eXchange, focusing on the applications that these policies support and the options that we have to express them.

SPEAKER'S BIO: Joaquin Chung received both his B.S. in Electrics and Communications Engineering (2007) and his M.Sc. in Communication Systems Engineering with Emphasis in Data Networks (2013) from University of Panama, Panama. He is a Fulbright scholar and currently pursuing his Ph.D. in Electrical and Computer Engineering under the supervision of Dr. Henry Owen and Dr. Raheem Beyah at Georgia Institute of Technology, GA. His research interests include Software Defined Networking and Network Security.

4. Exploring research collaborations between University of Puerto Rico and Brazil, by José Ortiz-Ubarri (University of Puerto Rico)

ABSTRACT: For the last few years my research group has dedicated most of the time in educational projects to improve CS curriculum and the way computer science is taught at the University of Puerto Rico, and cyber infrastructure to support big data transfers. However our main research interests are computer and network security, and the application of high performance technologies to solve computationally intensive scientific problems.

SPEAKER'S BIO: José Ortiz-Ubarri is an Associate Professor at the University of Puerto Rico, Rio Piedras, USA. He received a B.S. degree in Computer Science from the University of Puerto Rico (UPR), Rio Piedras in 2003, and a PhD degree in Computer Science and Engineering from the University of Puerto Rico, Mayagüez in 2010. His research interests include computer and networks security, high performance computing, and computer science education. He is the director of the NSF supported project Perimeter Network to Expedite the Transmission of Science (PR-NETS) at the University of Puerto Rico. José was a systems and network administrator for seven years at the High Performance Computing facility of the UPR. He is a member of the technical committee of the International Conference on Malicious and Unwanted Software (2013-2015).

5. Brocade SDN/OpenFlow Update, by Norival Figueira (Brocade)

ABSTRACT: This presentation covers Brocade's latest SDN/OpenFlow feature developments, including the Brocade SDN Controller, the Brocade Flow Optimizer application, and new router OpenFlow features. The presentation also briefly discusses related research work by the author on SDN/NFV Policy Framework and Architecture.

SPEAKER'S BIO: Norival Figueira received Ph.D. and M.S. degrees in Computer Science from the University of California, San Diego (UCSD). Dr. Figueira is currently a distinguished member of the Office of the CTO at Brocade, where he leads projects in the areas of SDN and network policies and security. Previously, Dr. Figueira was on Brocade's product management team, where he led the development of SDN/OpenFlow and other advanced features on Brocade router products. Prior to Brocade, Dr. Figueira served in senior leadership positions with Foundry Networks, Hammerhead Systems, Nortel Networks, and Bay Networks. Dr. Figueira is accredited with the issuance of 18 US patents in networking technology.

17:50 Closing Remarks

Expanding US/Brazil FIDC Testbed Federation

Mark Berman
GENI Project Office
Cambridge, MA, USA

***Abstract*—Future Internet and distributed cloud (FIDC) testbed researchers in Brazil and the US are among the first participants in a developing worldwide trend toward greater federation of research cyberinfrastructure. Federation of FIDC testbeds offers a promising value proposition for both operators and end users, offering scale, geographic diversity, and resource heterogeneity, while maintaining local control over resources, policies, and cost.**

Federation of FIDC testbeds occurs at several levels, but the first (and often most challenging) step is the human-to-human and organization-to-organization process of establishing trust and resource sharing agreements and the associated implementation policies. Once this step is addressed, the technical implementation may include some or all of AAA (authentication, authorization, and accounting) federation, control plane federation, and data plane federation. Established or emerging standards exist for each of these aspects of technical federation, at various levels of maturity.

Thanks to strong researcher-to-researcher relationships, there are a number of very compelling examples of federation between Brazilian and US testbeds, supporting a variety of research collaborations. There is great potential in pursuing two promising paths for expanding the power of these federations. One path is to agree upon and implement policies that simplify human-level processes, moving toward routine, automated instantiation of federated infrstructure configurations. A second path is to continue to expand research and infrastructure collaborations to a broader set of participating countries and regions.

Chameleon: A Large-Scale, Reconfigurable Experimental Environment for Cloud Research

K. Keahey, J. Mambretti, D.K. Panda, P. Rad, W. Smith, D. Stanzione

Cloud computing services have become critical to all major 21st century economic activities – yet, we are only beginning to understand this new important paradigm. Questions persist regarding applicability of the cloud platform to the emergent data-intensive and sensor-based applications, its suitability for high performance computing (HPC) applications, and its potential to leverage major emergent technologies such as Software Defined Networking (SDN). Answering those questions requires the ability to perform experiments at scale – in other words, an experimental testbed that would support experimentation with Big Data and Big Compute problems.

The Chameleon project has been built to provide such a large-scale platform to the open research community, allowing them to explore transformative concepts in deeply programmable cloud services, design, and core technologies. It's reconfigurability allows users to explore a broad range of problems ranging from the creation of Software as a Service to kernel support for virtualization. The broad range of supported research includes areas such as developing Platform-as-a-Service solutions, creating new and optimizing existing Infrastructure-as-a-Service components, investigating software-defined networking, and optimizing virtualization technologies.

The Chameleon testbed, deployed at the University of Chicago (UC) and the Texas Advanced Computing Center (TACC), will ultimately consist of 650 multi-core cloud nodes (~14,500 cores total), over 5PB of total disk space, and leverage 100 Gbps connection between the sites. The part of the testbed deployed to date consists of homogenous hardware to support large-scale experiments. The deployed hardware comprises 12 racks in total, two at University of Chicago and ten at TACC. Each deployed rack contains 46 Xeon Haswell processors (42 compute and 4 storage servers) with OpenFlow-enabled switches and 128 TB of storage. In addition, one of the racks is equipped with Infiniband network. In addition to distributed storage nodes, Chameleon has 3.6 PB of central storage currently being built for a persistent object store and shared filesystems. The testbed will also support heterogeneous units consisting of Atom microservers, ARM microservers, as well as a mix of servers with high RAM, FPGAs, NVidia K40 GPUs, and Intel Xeon Phis to allow experimentation with high-memory, large-disk, low-power, GPU, and co-processor units. In its initial phase, the project leveraged existing FutureGrid hardware at the University of Chicago and the Texas Advanced Computing Center configured to provide a transition period for the existing FutureGrid community of experimental users. This part of the testbed, called FutureGrid@Chameleon, is now being phased out.

To support a broad range of experiments, Chameleon supports a graduated configuration system allowing full user configurability of the software stack, from provisioning of bare metal and network interconnects to delivery of fully functioning cloud environments. The software that supports it, CHI, has been built by leveraging existing widely supported tools. Specifically, Chameleon leverages mainstream OpenStack tools such as Nova, Ironic, and Ceilometer to provide resource management, bare metal reconfiguration and monitoring. The OpenStack team has also extended OpenStack to provide advance reservations via Blazar. CHI also leverages Grid'5000 registry services to provide fine grained, verifiable information about compute resources.

In addition, to facilitate experiments and provide a "one stop shopping" for experimental artifacts, Chameleon will ultimately support a set of services designed to meet researchers needs, including repositories of trace and workload data of production cloud workloads based on both commercial and scientific clouds. The project also works on providing innovative ways of integrating testbeds into the educational pipeline by designing and publishing new educational artifacts such as deployment ready Chameleon appliances.

The project is led by the Computation Institute at the University of Chicago and partners from the Texas Advanced Computing Center at the University of Texas at Austin, the International Center for Advanced Internet Research at Northwestern University, the Ohio State University, and University of Texas at San Antonio, comprising a highly qualified and experienced team. The team includes members from the NSF supported FutureGrid project and from the GENI community, both forerunners of the NSFCloud solicitation under which this project is funded. Chameleon will also form a set of partnerships with commercial and academic clouds, such as Rackspace, CERN and Open Science Data Cloud (OSDC), and will partner with other testbeds, notably GENI and INRIA's Grid'5000 testbed.

Chameleon: a self-adaptative multi-cloud testbed

Carlos R. Senna and Edmundo R. M. Madeira
Institute of Computing - IC, University of Campinas – UNICAMP
{crsenna, edmundo}@ic.unicamp.br

Over the last years good results have been achieved in parallel and distributed computing initiatives, culminating in computing clouds, turning the Internet into a virtual plant of processing. At this time the technology allows us to integrate local resources (networks, grids and private clouds) to intensive computing centers (data centers), forming a hybrid and heterogeneous environment accessible through the Internet. However, new emerging technologies, notably Internet of Things (IoT) [1], bring with them new requirements to support mobility, geographical distribution, location awareness and low latency, requiring a new platform. To address these new challenges the cloud is migrating to the network edge, where routers can become a virtualized infrastructure by themselves. This evolution of the cloud, called Fog Computing [2], is a paradigm that extends the cloud computing services to the edge of the network, to meet a new generation of services and applications. The characteristics of Fog Computing are: low latency and location independence, wide geographic distribution, mobility, a very large number of nodes (billion), predominantly wireless access, strong presence of streaming media (streaming), real-time applications and heterogeneity. The Fog can be seen as the nearest floor of the cloud to users and devices. There is a good integration between the two technologies especially in the big data analytics and management. However, the current cloud technology has serious limitations to support the Fog. Most solutions for computing cloud are proprietary and do not support interoperability, federation and heterogeneity. In addition, the Cloud aims to serve millions of users/services while the Fog must meet billions (IoT).

In order to support the development of solutions that meet the IoT requirements, we have built a multi-cloud testbed based on the concepts of Fog Computing. Chameleon, the LRC (Computer Network Laboratory) multi-cloud testbed, consists of a hardware infrastructure with the capacity to host multiple cloud technologies and a software infrastructure to orchestrate the execution of distributed applications, integrate technologies and facilitate cooperation between the various levels of multi-cloud environments.

The Chameleon (Figure 1) is formed by the Orion, ViC and Cepheus clouds and provides access to resources of CloudLab, Amazon and Azure. It allows experiments in private clouds (Orion, ViC or Cepheus), in regional community clouds (Orion, ViC and Cepheus), in geographically distributed community clouds (Orion, ViC, Cepheus and CloudLab), hybrid clouds using combinations of Orion, ViC, Cepheus and Azure, and multi-clouds using all clouds simultaneously. Access to the testbed is done through the desktops of the local network of the LRC and through the Internet. The testbed also has an optical link with FAPESP Kyatera wide area network. The Orion is a cloud to execute general applications, Cepheus is prepared for Big Data analytics projects, while ViC is configured for sensor and vehicular applications.

The scope and youth of the IoT universe have not allowed the creation of models able to address all the potential requirements of this paradigm. What we have today are solutions that comply with niches of specific applications, such as sensor networks, vehicular computing and applications that process large amounts of data, among others. However, the IoT universe encompasses endless possibilities.

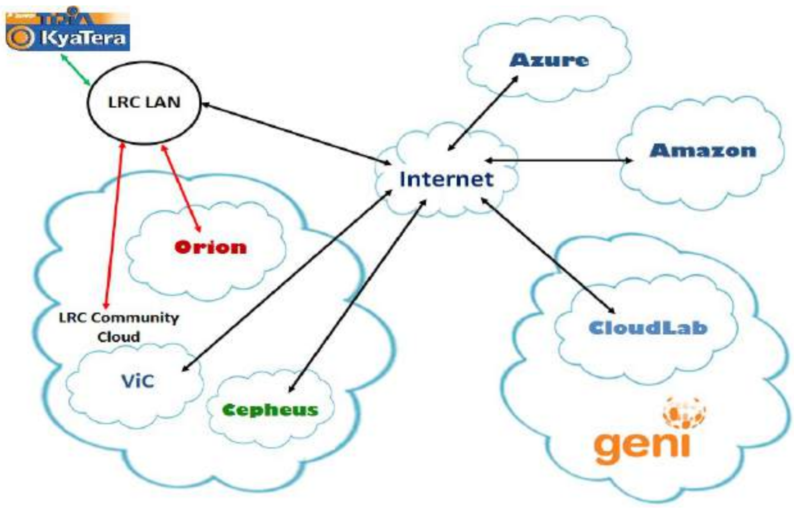

Figure 1. Chameleon – The LRC multi-cloud testbed.

The designing of the Chameleon addresses some aspects which we believe are fundamental to support research on IoT: bring computing to the generating points of information, providing interoperability between applications and adding scale and dynamism to the computing environment. The Chameleon facilities have been used in imaging applications, scientific workflows, sensor networks, vehicular applications and Hadoop applications [3], among others.

References

[1] L. Atzori, A. Iera, and G. Morabito,” The Internet of Things: A survey”. Comput. Netw. 54, 15 (October 2010), 2787-2805.
[2] F. Bonomi, R. Milito, J. Zhu, and S. Addepalli. “Fog computing and its role in the internet of things”. In Proceedings of the First Edition of the MCC Workshop on Mobile Cloud Computing, MCC ’12, pages 13–16, New York, NY, USA, 2012. ACM.
[3] C. R. Senna, L. G. C. Russi, and E. R. M. Madeira, “An Architecture for Orchestrating Hadoop Applications in Hybrid Cloud”, 14th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, May 2014.

BAMBU: A Metropolitan Innovation Testbed for Promoting Future Internet Research

Leobino N. Sampaio, Luciano Rebouças, and Marcos Barreto
Universidade Federal da Bahia (UFBA)
{leobino, lrebouca, marcosb}@ufba.br

Romildo Martins and Allan E. S. Freitas
Instituto Federal de Educação, Ciência e Tecnologia da Bahia (IFBA)
{romildo,allan}@ifba.edu.br

ABSTRACT
A broad number of large-scale network testbeds have been proposed in the last couple of years as a mean of better supporting Future Internet research. This paper present Bambu. A metropolitan innovation testbed for promoting experimental researches in the city of Salvador, Bahia.

Keywords
Network Testbeds, Future Internet, Software-Defined Networks.

1. INTRODUCTION

Network testbeds have supported many of the current initiatives of Future Internet research over the recent years, attracting attention from researchers of academia and industry [1-3]. The Future Internet testbeds/experimentation between BRazil and Europe (FIBRE) project[1] is an example. It consists of joint effort made by Brazilian and European research institutions to build a federated large-scale experimental Future Internet facility. The project has ended its activities in October 2014. However, the computing resources reserved for experimentation from the Brazilian side of FIBRE (FIBRE-BR) have been used by Brazilian institutions to support the communities interested in promoting network innovations.

In FIBRE-BR, each participant institution hosts and manages an experimentation node (also named “island”) that provides a set of programmable network elements to support experiments using wired and wireless technologies. Users access these experimentation resources through the FIBRE network (i.e., FIBRENET), that enables them to allocate a slice of computing resources (i.e., network links, router ports, CPU, storage) across the whole network involved in a given experiment [4]. This initiative certainly presents great benefits for the Brazilian research community, however, there is room for improvement due to some reasons. Firstly, despite it consists of a far-reaching initiative, students and researchers from non-participant institutions still lack resources for network experimentation. Secondly, since such environments provides resources for conducting experiments related to national and large scale networks, use cases concerning small contexts cannot be well studied and analyzed through these facilities. Thirdly, in regional contexts, students and researchers still need incentives and assistances for knowledge acquisition about technologies used in testbeds. For instance, those technologies adopted to implement the Software-Defined Networks [5] concept, such as the Openflow protocol.

Motivated by the aforementioned discussion, the Bambu project has been submitted in response to FAPESB's 013/2015 call for research project proposals and approved for funding. Its main goal is to design and implement a metropolitan innovation testbed for promoting practical Future Internet research in the city of Salvador, Bahia. Salvador is the third-biggest metropolis of Brazil, hosting a great number of academic institutions. Its metropolitan network, named Remessa[2], comprises a multitude of organizations, from municipality, State Government and Federal Government, to Healthcare Institutions, to Technology Park of Bahia, to Military Units, and to Academic Institutions. Remessa is connected to Brazilian Academic Network through RNP[3], connecting around of other 32 institutions and almost 100 sites and relies on Layer 2 transport and on redundancy through metro Ethernet rings. Through the Bambu testbed, a bigger number of students and researchers will benefit from experimentation facilities. The next sections describe the Bambu project, its goals and future plans.

2. THE BAMBU PROJECT

The Bambu project resulted from a partnership between Universidade Federal da Bahia (UFBA), Instituto Federal da Bahia (IFBA), and Centro de Pesquisa Gonçalo Moniz – Bahia (Fiocruz-Ba). In addition to that, it has been supported by Rede Nacional de Ensino e Pesquisa (RNP), Universidade Federal do Espírito Santo (UFES), Laboratório Nacional de Computação Científica (LNCC), Florida International University (FIU), and Phillips Research North America. Representative members from these institutions started this initiative in order to build a flexible, robust and scalable[4] network to support innovation and training, and to meet high-volume traffic demands from advanced applications used by the project partners. Bambu will run as an extension of the FIBRE-BR project in order to enable researchers from other institutions to use its facilities. Therefore, this testbed will also benefit researchers with large-scale real network conditions in their activities. Such potential for scope expansion opens opportunities for students and outside researchers to collaborate with academic institutions from the city of Salvador.

Besides the testbed itself, this project also aims to: i) deploy an experimental “island” at UFBA in order to make available computing resources to other FIBRE members; ii) adapt FIBRE’s control planes to be adopted by the testbed to be built in the project; iii) build a local virtual laboratory for students and researchers from academic institutions of Salvador, which help them in designing and testing new network applications; iv)

[1] http://www.fibre-ict.eu

[2] http:// http://www.pop-ba.rnp.br/Remessa

[3] The Brazilian National Research and Education Network (NREN) http://www.rnp.br

[4] These requirements motivated the name chosen for the testbed, since Bambu (in Portuguese) is a plant that enables the innovation of a diverse set of creative new products with such properties.

develop and build new Openflow-based network optimization techniques; v) create a more realistic and reliable training environment for knowledge acquisition in programmable networks; vi) demonstrate the testbed benefits through pilot applications or showcases to be deployed on top of the available network infrastructure. The last goal concerns the demonstration of Bambu's potentials. We will carry it out through three pilot applications related to ongoing projects, as following explained.

1) "Probabilistic linkage of public healthcare data: the 100 million Brazilian cohort". This project aims to integrate disparate databases from the Brazilian government and perform epidemiological and statistical studies over the integrated data. It created a cohort comprised by a 100 million records from *Cadastro Único*, a database kept by the Ministry of Social Development that stores socioeconomic data of people enrolled in social programmes offered by the Brazilian government. The project includes the set up of a secure data center at SENAI CIMATEC in order to use the Yemoja supercomputer to perform parallel, probabilistic correlations of all these databases. The data marts generated by linkage routines will be accessed from several sites within the Federal University of Bahia (UFBA) and Fiocruz-ba. We expect to use the Bambu's facilities to allow for secure, fast, and reliable access and file transfer among all sites.

2) "High-volume video data transferring". This project results from a partnership between UFBA and the Companhia de Governança Eletrônica (COGEL), in Salvador, to analyze thousands of videos coming from camera networks spread in the city, without going out of the lab. Traffic detection is almost mandatory in big cities, nowadays. Not only because of the great number of cars in these cities, which demands controlling urban traffic jam, but also to make automatic and adaptive the way of managing traffic lights. Using Computer Vision and Image Pattern Recognition, this project aims to detect traffic jam by the texture produced by the vehicles in a urban scenarios. With the estimated density and the optical flow of the vehicle clouds, it is possible to predict the degree of traffic jam in a possibilistic way (Fuzzy). Since these solutions generate high-volume video data, we plan to exploit the Bambu facilities for dynamic network circuit provisioning.

3) "An Internet of Things framework for supporting Smart cities". This project aims to build a software framework that enables the human activity optimization, thereby minimizing the high complexity of smart cities management. Through this framework, it is expected that the existence of too much network traffic among people and objects (Internet of Things). Therefore, the testbed will be used to support studies and investigations about the dynamic creation of network circuits to assist such traffic in an user interactive manner.

3. BAMBU TESTBED

We will build the Bambu testbed through an overlay network on the Remessa infrastructure. Remessa is already equipped with dark fibres and some obsolete switches. Hence, our plan is to connect the project partners through an optical ring by combining the use of VLANs on existing switches with available dark fibres. To this aim, we will acquire Openflow switches equipped with 10Gbps network interfaces, thereby updating the Remessa's communication infrastructure and taking advantage of unused optical fibres.

Figure 1 illustrates the initial Bambu topology and its relation with Remessa network. As it shows, Bambu will connect its partners coupled with the Remessa's SENAI CIMATEC unit. The latter hosts the Yemoja supercomputer. For this reason, its integration with Bambu is very important for the project, in particular, to carry out one of our pilot application and to allow for future related projects using the super computer's power of computation.

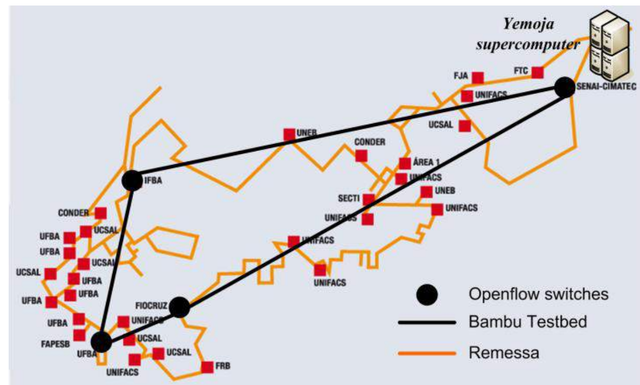

Figure 1. The Bambu network and its relation with Remessa.

4. CONCLUDING REMARKS

The Bambu project has just been approved for funding by Fapesb. The project's kick off meeting will take place by the end of 2015, but until then we will be discussing ideas and possibilities to deploy the most appropriate infrastructure for promoting Future Internet research in the city of Salvador. Our team comprises researchers from multidisciplinary fields, students, network engineers, and application developers.

5. REFERENCES

[1] Ibarra, J.; Bezerra, J.; Morgan, H.; Fernandez Lopez, L.; Stanton, M.; Machado, I.; Grizendi, E.; Cox, D.A., "Benefits brought by the use of OpenFlow/SDN on the AmLight intercontinental research and education network," in IM 2015 IFIP/IEEE International Symposium May 2015.

[2] Berman, M.; Chase, J. S.; Landweber, L.; Nakao, A.; Ott, M.; Raychaudhuri, D.; Ricci, R. & Seskar, I. GENI: A Federated Testbed For Innovative Network Experiments Computer Networks, 2014, 61, 5-23.

[3] Anastasius Gavras, Arto Karila, Serge Fdida, Martin May, and Martin Potts. 2007. Future internet research and experimentation: the FIRE initiative. SIGCOMM Comput. Commun. Rev. 37, 3 (July 2007), 89-92.

[4] Julio Ibarra, Jeronimo Bezerra, Heidi Alvarez, Donald A. Cox, III, Michael Stanton, Iara Machado, and Eduardo Grizendi: Benefits brought by the use of OpenFlow/SDN in the AmLight intercontinental research and education network, IM 2015 Ottawa, Canada, 11-15 May, 2015.

[5] D. Kreutz, F. Ramos, P. Esteves Verissimo, C. Esteve Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," Proceedings of the IEEE, vol. 103, no. 1, pp. 14–76, Jan 2015.

From Federated Software Defined Infrastructure to Future Internet

Kuang-Ching Wang
Department of Electrical and Computer Engineering
Center of Excellence on Next Generation Computing and Creativity
Clemson University
kwang@clemson.edu

Since 2009, Clemson University embarked on a series of NSF projects that led its researchers and IT engineers through an exciting learning curve for software defined networking (SDN) and cloud computing technologies. From having researchers set up and experiment with Clemson's very first OpenFlow network over a few cross-campus VLANs as part of the GENI meso-scale testbed (NSF GENI, 2009), to having IT engineers deploying and operating a twenty-building production OpenFlow research network parallel to the legacy campus network (NSF CC-NIE, 2012), to building "CloudLab" (NSF FutureCloud, 2014), to hosting academic collaborators' SDN solutions (BGPMux, 2011, ONOS, 2015), to acquiring commercial SDN fabric based cloud data center pods for Clemson's production computing (2015), the tightly weaved academic-IT team has had an unprecedented opportunity to have closely navigated the technology evolution of SDN. During this time, significant efforts have been devoted to learning the practical challenges of deploying and programming SDN in both green field and brown field environments, from campus to regional, national, and international networks. The experiences altogether have proven fruitful from both academic and operation perspectives, giving us bearings for SDN's truths and myths, pros and cons of its variants, open challenges and opportunities.

The Clemson team has been exploring OpenFlow at multiple granularities. We explored wide area parallel data transfer (SOS [1]), campus networking, mobile device seamless handover (GENI WiMAX [2]), data center networking (VNGuard [3]), distributed video gateway and content delivery (GENI Cinema [4]). We are applying SDN in a range of research projects in domains such as genomics [5], connected vehicle [6], and smart grid [7]. For federation, we work closely with international partners such as University of Kaiserslautern and the EU GEANT Testbed Service [8] to address short term and long term challenges. With CloudLab, our interest reaches into Network Function Virtualization, where computing end hosts (virtual or physical servers in a data center) can be flexibly deployed with network services running also in virtual machines within the data center. Our journey revealed to us a software defined infrastructure (SDI) that is taking shape across the globe. It is our belief that SDIs around the world will be increasingly federated and interoperable, allowing applications to be launched from anywhere, with highly customizable communication with other services near and far. While there remains development cycles to catch up for the SDI infrastructure to be mature, it is our belief that the applications and their logical requirements, both locally and globally, will be driving the federated SDI architecture, which will shape the future Internet.

References

[1] A. Rosen and K.-C. Wang, "Steroid OpenFlow Service: Seamless Network Service Delivery in Software Defined Networks," in Proceedings of the First GENI Research and Educational Experiment Workshop, pp.1~5, March 2012.

[2] Ke Xu, Ryan Izard, Kuang-Ching Wang, and Jim Martin , "Cloud-based Handoff as a Service for Heterogeneous Vehicular Networks with OpenFlow", in Proceedings of GENI Research and Education Experimentation Workshop 2013, pp. 1-5, March 2013.

[3] Juan Deng, et al., "VNGuard: An NFV/SDN Combination Framework for Provisioning and Managing Virtual Firewalls", in Proceedings of IEEE NFV-SDN'15.

[4] Qing Wang, et al., "GENI Cinema: A SDN-Assisted Scalable Live Video Streaming Service," in Proceedings of IEEE ICNP 2014.

[5] Stephen Ficklin, et al., "CIF21 DIBBS: Tripal Gateway, a Platform for Next-Generation Data Analysis and Sharing", NSF ACI-1443040.

[6] James Martin, Hongxin Hu, Kuang-Ching Wang, and Mashrur Chowdhury, "US Ignite: Track 1: Enabling Connected Vehicle Applications through Advanced Network Technology", NSF CNS-1531127.

[7] G. Venayagamoorthy, et al., "I/UCRC for Real-Time Intelligence for Smart Electric Grid Operations (RISE)," NSF IUCRC-1464637.

[8] GEANT, GEANT Testbeds Service, http://gts.geant.net.

US Ignite Position Paper

Glenn Ricart
15 September 2015

1 How is today's Internet inadequate?

There are entire classes of end-user applications which are not feasible on today's widely-deployed Internet. These currently not-feasible applications require one or more of the following:

- Highly-responsive performance. Today's Internet is tuned for cloud-based web responses which can take hundreds of milliseconds adding up to 1-3 seconds per web page. Applications which need more responsiveness may not be possible.

- Security by isolation. Given the track record of hackers being able to bypass encryption, highly sensitive information such as healthcare, financial, or public safety information needs to be protected by being isolated from the general Internet. Although enterprises can set up VLANs and MPLS, these tools are not available for home and small business users.

- Highly-responsive and reliable packet delivery to support cyberphysical systems. Cyberphysical systems operate in both the real world and in cyberinfrastructure simultaneously. Today's Internet is often too slow or too unreliable to support real-time cyberphysical systems such as cloud-based brains for robotics or 3D printers or home medical monitors.

- Integrating information from multiple Internet of Things sensors simultaneously. To make good decisions, timely information often needs to be integrated from multiple sensors. Can we depend on deterministic networks that will deliver all information in a timely way?

- Connecting an Industrial Internet together in ways that maintain the synchronism of assembly line inspection cameras, pick-and-place robotics, and automated test procedures. All of these are time critical in ways that today's ordinary Internet may not honor, especially when there is other competing traffic.

- The capability to continue to operate when disaster strikes and remote clouds are no longer reachable either due to severed fiber routes or overloaded circuits.

- The ability to move big data when required, both upstream and downstream.

2 THE US IGNITE RESPONSE

US Ignite proposes to address the above issues by giving cities their own digital town square with local cloud servers and at least one exchange point where purely local traffic can be exchanged. This architecture, called *locavore* architecture, is a kind of city-edge architecture with the following components:

1. Abundant bandwidth provided by local fiber backhauls. Several vendors are now selling 10Gbps GPON ONTs. For longer distances, multimode fiber can easily handle tens of terabytes. Big bandwidth is needed for photo-realistic virtual reality controlled in real-time by the end-user.

2. Local cloud computing provided by local servers. The local servers can handle and generate big bandwidth and/or very high responsiveness because there are right there in the city. In addition, they are still operational even if the city is cut off from other national infrastructure such as Twitter or Facebook or those services are overloaded.

3. One or more exchange points that serve purely local traffic. Local carriers can interchange traffic and keep it local, reliable, and low latency. This allows multiple high-speed (possibly gigabit) providers to operate competitively (while cooperatively) in the same city. It also allows multiple providers of local clouds to operate competitively (while cooperatively) in the same city.

4. Software Defined Networking to create within-city isolated flows for sensitive information. Eventually multi-city SDN flows (mediated by SDXes) will also be needed.

5. Lower-latency layer 2 inter-city circuits for minimum latency multi-city applications such as medical video conferencing without the delays.

6. Virtualization so that the above capabilities can be deployed dynamically into a slice as needed to serve a specific application or user. The slice of virtualized resources can also serve as a unit of infrastructure for billing purposes. It is suggested that the bill for the infrastructure be assigned to the applications / services it supports and the end-user simply see one charge for the use of the application or service.

3 THE US TESTBED

Fifteen US Ignite cities will develop applications previously infeasible. Each of those cities will have an interoperable and interconnected locavore infrastructure as described above. Initially, GENI racks will play the role of local cloud servers. An in-city exchange point will improve response times for all users for all local applications. For example, a major employer or university will be able to give their employees the same experience working from home that they get on campus or at the office. This is a three year project beginning September 1, 2015. The first cities will install locavore components in early 2016.

4 IGNITE MUNDI

This methodology can be executed in other countries as well. Early discussions between the European Commission, the U.S., Canada, and Japan have suggested the overall project title of "Ignite Mundi" to talk about the parallel international implementations. Sister international cities could collaborate on their implementations and perhaps run some international services with security and latency not previously possible.

Is there interest in Brasil in deploying one or more cities or communities with a locavore architecture and participating in Ignite Mundi?

Flexible Infrastructure Management for Smart Cities

Carlos Alberto Kamienski
Universidade Federal do ABC (UFABC), Brazil, cak@ufabc.edu.br

Abstract

Smart cities require a variety of new services to be available to citizens of large urban spaces, based on data coming from different sources in the Internet of Things (IoT). This will pose high demands for dynamic allocation of existing computing and communication infrastructure that currently are based on datacenters. We advocate that the dynamic management of centralized and distributed cloud-based and IoT infrastructures with elasticity, SDN and NFV is a key feature for enabling such new scenarios. Results with dynamic datacenter reconfiguration will be used for providing services for smart cities, such as dealing with high and varying density of people in particular locations.

1 Introduction

Cloud computing, software-defined networks (SDN) and network function virtualization (NFV) make it possible to dynamically manage a cloud datacenter for providing online computing resources to a variety of applications and users [1][2]. Flexibility is the key attribute associated to this concept, which needs adequate management in order to be useful.

On the other hand, the building of smart societies where people fully exert their rights and duties is intensively based on the use of information and communication technologies [3]. The ideal of smart cities, together with the new developments in the Internet of Things area can be leveraged by the new concepts introduced by urban computing [4] to make it possible the building of smarter societies. Urban computing may be defined as a process of acquisition, integration, and analysis of big and heterogeneous data generated by different sources, such as sensors, devices, vehicles, buildings, and human, to deal with important challenges in our major cities such as traffic congestion, mobility and high and varying people density in some places [4].

This position paper presents a research area that aims at developing and analyzing the performance of mechanisms and technologies for the flexible management of distributed cloud-based and IoT infrastructures in scenarios for improving the life quality of citizens in smarter societies. For example, such technology may be used for agile configuration of networked resources for supporting flash crowd events in our cities. Previous and current results in managing elasticity for cloud computing, dynamic network reorganization for SDN and on demand creation and reconfiguration of network functions can be used together to support society needs providing adequate quality of experience.

2 Flexible Datacenter Management

Some previous and current research aimed at analyzing the behavior of dynamic reconfiguration of resources in datacenters. Here we present two examples of using elasticity for managing virtual machines and one example of SDN reconfiguration.

Figure 1 depicts a time series of the response time in a smart grid scenario where the request arrival rate increases and decreases over time [5]. We increased workload from zero to 3500 simultaneous requests and then back to zero and during the experiments up to 10 virtual machines have been allocated. We used response time as a metric for managing elasticity, with a policy defining higher and lower thresholds for triggering elasticity actions, i.e., VM creation and teardown. Our SLA for the response time is defined as 10 seconds and one can observe that the mean and median are mostly kept under 3 seconds and the 90th percentile is under 9 seconds.

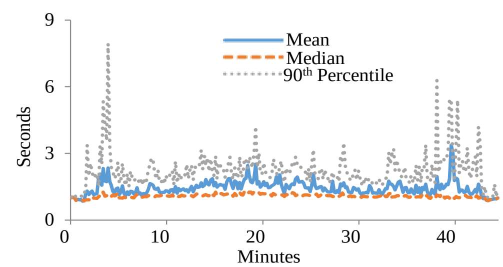

Figure 1 – Elasticity management in a private cloud

In a different scenario, we analyzed the effectiveness of using elasticity for VM failure management, instead of traditional recovery mechanisms. Figure 2 shows results comparing the use of elasticity only (no recovery) and recovery only (no elasticity) for two metrics, namely, transfer time per MB and request error rate. What can be seen is that the two mechanisms are equivalent because the confidence intervals overlap for both metrics.

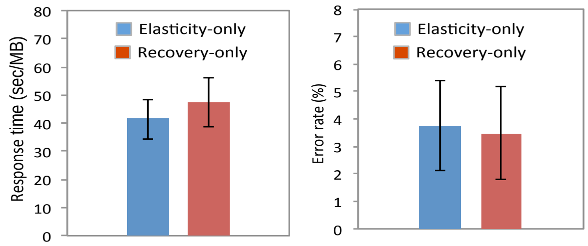

a) Response Time b) Error rate

Figure 2 – Elasticity for failure recovery

Figure 3 presents results of a SDN reconfiguration scenario, aimed at moving automatically all VMs of cloud tenants to a place in datacenter that is as close as possible to each other. Similar to a disk defrag, the idea is that the elasticity mechanism may place VMs in different racks in a datacenter, so that from time to time a reconfiguration might improve performance. The picture shows that the throughput between VM #1 and the other five VMs improves considerably after a reconfiguration event.

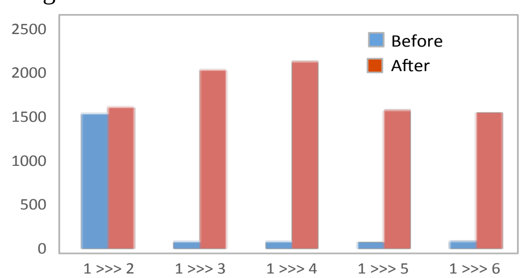

Figure 3 – SDN reconfiguration (network defrag)

3 Scenarios for Smarter Societies

We believe that the management of dynamic configuration mechanisms in datacenters and distributed IoT and cloud infrastructures is a key feature for making it possible new scenarios for improving life quality in large urban spaces. Different situations that occur in the everyday life of citizens and in special occasions require different digital services, including network services, to be available with a given configuration. Usually they are statically created and configured, which on the one hand wastes resources when the utilization level is low and on the other hand restricts service offerings to some users when utilization is high. In order to figure out and better understand the need for dynamic and flexible configuration of functions frequently used in service offering for citizens, new scenarios, environments and possibilities that make use of computing resources that are available in datacenters need to be introduced, modeled and analyzed. They involve the creation and teardown of services composed of different network functions and customized applications, as well as its on demand expansion and contraction (i.e. elasticity).

Of particular interests are events that generate flash crowds, or simply the unexpected concentration of people in a large city, no matter whether the location is prepared for that. Examples are music concerts in public places, games in sports stadiums, traffic jams and demonstrations. In all cases, different services must be available with adequate security, robustness and scalability levels, including Internet connectivity allowing access to information that may help avoiding turmoil, chaos or disturbing the peace.

Technologies that will be considered in the datacenter are elasticity, SDN and NFV. Also, careful event identification is needed, which can be done by a variety of different sensors, but also by searching the Web for programmed events and social networks for instant ones. Crowdsourcing and crowdsensing are everywhere in different applications for smartphones in play a important role in providing smarter services for Smart Cities. Analytics will be needed using typical big data technics. Figure 4 shows the big picture of the technological components to be considered.

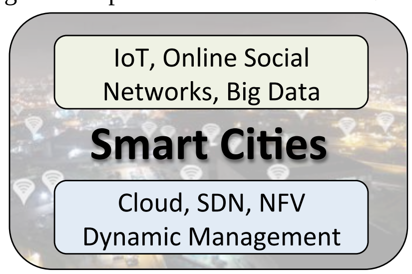

Figure 4 – Technological components

4 Conclusions

Smart cities may become a reality with the new trend in urban computing, which requires data to the acquired from different sources, stored and processed, and finally services to be deployed for the citizens. We believe that the flexible datacenter management will play a key role in providing instant resources for a variety of services in urban spaces. Previous and past results confirm this view and support future work towards making technology work for smart cities.

References

[1] Armbrust, M. et al., “A View of Cloud Computing”, Communications of the ACM, 53(4), pp. 50-58, 2010.
[2] ONF, "OpenFlow-enabled SDN and Network Functions Virtualization", ONF Solution Brief, February 2014.
[3] Caragliu, A. et. al, “Smart Cities in Europe”, Journal of Urban Technology, 18:2, pp. 65-82, August 2011.
[4] Zheng, Y., et. al, "Urban Computing: Concepts, Methodologies, and Applications", ACM Transactions on Intelligent Systems and Technology, 6(2), July 2014.
[5] Simões, R., Kamienski, C., “Elasticity Management in Private and Hybrid Clouds”, IEEE Cloud 2014.

NovaGenesis Project: Overview, Current Activities and Research Opportunities

Antonio Marcos Alberti

Abstract—**This paper presents an overview of the NovaGenesis convergent information project, summarizing its current research and development activities. It also presents a list of research opportunities – open problems – that could be addressed by our team together with other institutions in joint research initiatives.**

Index Terms—**Future Internet, NovaGenesis, architecture design, experimentation, performance evaluation and comparison.**

I. Introduction

NOVAGENESIS (NG) project emerged in 2008 as the result of a feeling of discomfort: Is it possible to do something new in communications networks research? Is there any novel ideas for the Internet? What if we redesign it from scratch? How it could be? This is the aim of NovaGenesis project: to rethink current architectures for information exchanging, processing and storage, including the Internet, towards proposing a convergent information approach, which synergistically integrates many emerging paradigms for ICT.

As the Internet was opened to the general public and began to be used for an increasing diversity of applications, a complex agglomerate of patchwork solutions were applied to extend its scope – creating some inconsistencies that now began to be questioned. The Internet scales and roles changed considerably from its original purposes [1]. Many researchers started wondering whether the current stack can support or not the multifaceted exponential growths we are experiencing on number of devices, mobility, interactivity, content and traffic. Concerned with this situation, several initiatives emerged worldwide to reshape the Internet under the banner of the so called "**future Internet architecture**" (FIA) design [2], [3].

NovaGenesis is a **clean slate convergent information architecture** (CSCIA), which covers information processing, storage, exchanging, and visualization. Its scope is not limited to the original Internet scope: a robust network with flexible applications to exchange information among computers. Instead, a CSCIA embraces not only networking aspects, but also computing, storage, and visualization. Thus, it looks for redesigning both communication mechanisms and protocols implementations for enabling a program running anywhere to address messages to programs anywhere else with acceptable performance and portability levels. Also, it looks for the convergence of computing and communications, merging technologies like cloud computing, mobile computing, Internet of things (IoT), service frameworks, software-defined networking (SDN), network function virtualization (NFV) and distributed systems. Therefore, to design a convergent information architecture is more broad than designing protocols and their stack. It includes the computing, storage, and services/applications aspects.

ICT Lab, Instituto Nacional de Telecomunicações - INATEL. João de Camargo 510, Centro, Santa Rita do Sapucaí, CEP 37540-000, Minas Gerais, Brazil. Email: alberti@inatel.br

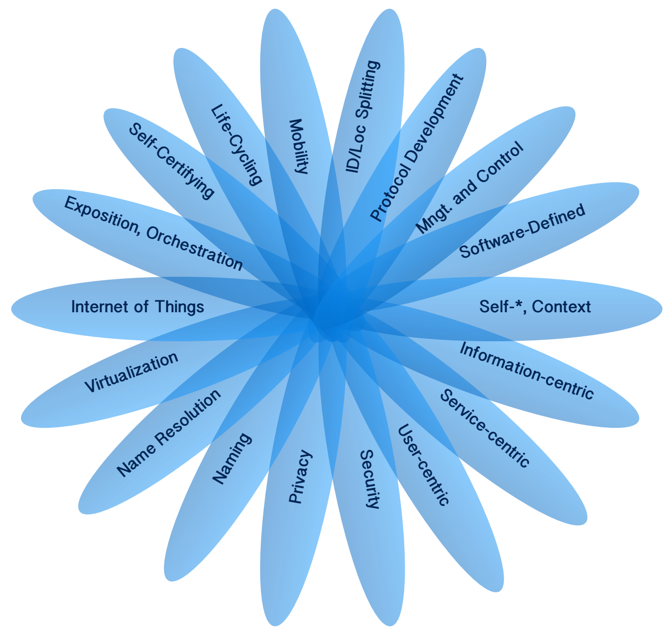

Fig. 1. Key ingredients for Future Internet.

Figure 1 illustrates key CSCIA ingredients we tried to converge in NovaGenesis. We started from naming, name resolution, and life-cycling of contents and services. The idea was to create a generic framework that could deal with naming, life-cycling, persistent identification, location-independency, joint content, services, and "things" coordination (orchestration). In other words, we tried to merge information-centric networking (ICN) with service-oriented architecture (SOA) and service-centric networking (SCN), providing ID/Loc splitting for mobility.

We started by supporting not only natural language names (NLNs), but also self-verifying naming (SVN). We created a distributed facility to store bindings among names and/or contents, creating a huge web of names. NG supports unlimited namespaces with generic name resolution. Next, we advanced towards software-defining the architecture, which we defined as service-defined networking (SDN)[4], aiming to accommodate not only SDN, but also software-defined cloud resources and "things". We also advanced towards protocols-implemented-as-a-service (PIaaS) paradigm [4].

The service-oriented nature of NG is aligned with virtualization, e.g. NFV. Service chaining is supported by the contract-based, self-verifying nature of NG. Finally, the design also embraced IoT and its self-organization/management/operation. We proposed a proxy/gateway/controller approach to represent “things” on NG [5]. Autonomic services are the final piece to be added in next version.

II. Current Activities

In 2012, we implemented a C++ prototype of NovaGenesis to run at user space of Linux Kernel (see Figure 2). This 40k line codes prototype already runs without TCP/IP, using a raw socket directly over Ethernet or Wi-Fi.

Fig. 2. NovaGenesis prototype running at ICT Lab, Inatel.

Current activities on NovaGenesis project include:

- NG architecture for adaptive Wi-Fi with real-time, multi-sector embedded spectrum sensing. NG services have been developed to coordinate contract-based spectrum sensing, access point configuration, active antenna control, and Wi-Fi frame forwarding [6].
- NG integration to OpenFlow [4]. We developed a service that represents python OpenFlow controller inside NG. Next step is to extend this service to help OpenFlow switches to forward NG messages.
- Representation of ”things” on NG cloud. We are developing an embedded proxy/gateway/controller to integrate heterogeneous “things” with NG services [5].
- We are employing NG as an opportunistic, socially-driven, self-organizing, cloud networking architecture. This work is being done together with Lusofona University, Unisinos, Technical Univ. of Bucharest, and Hankuk University of Foreign Studies, in South Korea [7].
- Scalability testing of NG services at ICT Lab, GENI, and other facilities.
- NovaGenesis comparison to current Internet and other FIA proposals. Right now, we are comparing NG to domain name system (DNS), expressive Internet architecture (XIA) and recursive internetwork architecture (RINA).
- Applying NG concepts to 5G. This work is being done under the scope of the Radiocommunication Reference Center, funded by Finep/Funttel in Brazil.

III. Research Opportunities

Each FIA, 5G, IoT, NFV, SDN, cognitive radio initiative has a particular focus and its own set of design requirements, principles, and choices. They progress in several aspects when compared to the current Internet. However, the scientific method requires two things for all of them: (i) to prove they are better than current approaches; and (ii), to outperform their competitors.

In this scenario, what features need to be evaluated and compared? What are the point of views to be considered? What can be learned from these comparisons? Despite the importance of these questions, they remain largely unanswered. We, as a community, need to develop multidisciplinary, multi-dimensional methodologies for emerging architectures evaluation/comparison. Up to now, this effort is limited to ICN.

Experimental facilities could have a big role on this matter. Right now, we are starting our scalability tests on GENI. Let’s start working on more deep scenarios and methodologies for emerging architectures comparison. Mathematical modeling and simulation are also welcome.

Another topic would be the synergistic integration of the key emerging architectures’ ingredients, like IoT, SDN, NFV, SOA, SCN, ICN, etc. Architectures should allow for the application of different paradigms in a neat way. Each of these key ingredients strongly contributes to advance a specific designing dimension. When two or more of these ingredients are synergistically integrated, there is a “cross fertilization”, a catalyzing effect, which favor global architectural advances instead of local ones. How to evaluate this point is an open issue.

Acknowledgment

This work was partially supported by Finep/Funttel Grant No. 01.14.0231.00, under the Radiocommunication Reference Center (Centro de Referência em Radicomunicações- CRR) project of the National Institute of Telecommunications (Instituto Nacional de Telecomunicações - Inatel), Brazil.

References

[1] T. Leighton, “Improving performance on the internet,” *Commun. ACM*, vol. 52, no. 2, pp. 44–51, Feb. 2009. [Online]. Available: http://doi.acm.org/10.1145/1461928.1461944

[2] J. Pan, S. Paul, and R. Jain, “A survey of the research on future internet architectures,” *Communications Magazine, IEEE*, vol. 49, no. 7, pp. 26 –36, july 2011.

[3] A. M. Alberti, “A conceptual-driven survey on future internet requirements, technologies, and challenges,” *Journal of the Brazilian Computer Society*, vol. 19, no. 3, pp. 291–311, 2013.

[4] A. Alberti, V. de O Fernandes, M. Casaroli, L. de Oliveira, F. Pedroso, and D. Singh, “A novagenesis proxy/gateway/controller for openflow software defined networks,” in *Network and Service Management (CNSM), 2014 10th International Conference on*, Nov 2014, pp. 394–399.

[5] A. Alberti and D. Singh, “Developing a novagenesis architecture model for service oriented future internet and iot: An advanced transportation system scenario,” in *Internet of Things (WF-IoT), 2014 IEEE World Forum on*, March 2014, pp. 359–364.

[6] L. de Oliveira, A. Alberti, M. Favoreto Casaroli, E. Raimundo-Neto, I. Feliciano da Costa, and A. Cerqueira Sodre, “Service-oriented, name-based, and software-defined spectrum sensing and dynamic resource allocation for wi-fi networks using novagenesis,” in *Telecommunications (IWT), 2015 International Workshop on*, June 2015, pp. 1–8.

[7] A. M. Alberti, W. Moreira, R. da Rosa Righi, F. J. P. Neto, C. Dobre, and D. Singh, “Towards an opportunistic, socially-driven, self-organizing, cloud networking architecture with novagenesis,” in *ESaaSA 2015 - Proceedings of the 2nd International Workshop on Emerging Software as a Service and Analytics, Lisbon, Portugal, 20-22 May, 2015.*, 2015, pp. 27–36. [Online]. Available: http://dx.doi.org/10.5220/0005526300270036

Research Problems in Future Networking

Deep Medhi
University of Missouri-Kansas City (UMKC), USA
dmedhi@umkc.edu

Collaborators:
Edmundo Madeira, University of Campinas (UNICAMP), Brazil; edmundo@ic.unicamp.br
Lisandro Zambenedetti Granville, and Luciano Paschoal Gaspary, Federal University of Rio Grande do Sul (UFRGS), Brazil; {granville, paschoal}@inf.ufrgs.br
Sejun Song Baek-Young Choi, University of Missouri-Kansas City (UMKC), USA, {sjsong, choiby}@umkc.edu

October 2015

In this position paper, we present several on-going activities in our research team that facilitate essential components for building Future Internet.

Existing Collaborations

Deep Medhi has served as a leading international research visitor under the Brazilian Science Mobility Program during 2013-2015 with the University of Campinas (UNICAMP), Brazil as his host institution. Under this program, he has made multiple visits (each of about four to five weeks) to UNICAMP, where he has collaborated with Edmundo Madeira and Nelson Fonseca. Furthermore, with the Science Mobility Program visits to Brazil, he has also visited the Federal University of Rio Grande de Sul (UFRGS), Brazil in July 2013, July 2014, and July 2015, where he has initiated research collaboration with Lisandro Zambenedetti Granville and Luciano Paschoal Gaspary. The research areas of these collaborations span data center networking, multilayer networking, and resilient embedding in network virtualization and SDN, which are very much in alignment with the scope of the US-Brazil workshop.

Potential Collaboration Projects

Hadoop-in-a-Hybrid Cloud

Hadoop is typically used in a closed environment at a single physical site. In this research project between UMKC and UNICAMP, we consider the situation where a site runs short on Hadoop resources and requires resources from one or more sites in different geographic regions. This situation arises, for example, if one has access to a private cloud, and wants to use a public cloud on demand in a different geographic area. In this project, we are currently exploring setting up Hadoop in a hybrid cloud environment and how the network latency and bandwidth availability may be taken into account to allocate Hadoop processes in a remote cloud, depending on the applications using the Hadoop.

Impact of Network Speed in Data Center Networking

Virtual machine migration is an important requirement in a data center networking to allow resource flexibility. We want to understand how migration works with the increase the speed of underlying networks in a data center. Our initiation assessment is that scheduling algorithms needs to be network-speed aware since the network speed reduces the time taken to migration VMs. In this collaborative work between UMKC and UNICAMP, we are considering new scheduling algorithms that are network-aware to show how performance can be improved. We are continuing this work to consider the topological structure of the data center networks in addition to network speed.

Federated Network and Deterministic Networking Services

This is a new collaborative project between UMKC, UFRGS, and the University of Antwerp, Belgium, where we are investigating a wide-area federated network environment where applications require deterministic network services (such as high quality video conferencing services). We are exploring research as well as technical challenges in establishing federation and coordination between three different sites located in three different continents. In particular, we are exploring how software-defined networking can be used in this federated environment to give service guarantees.

Cloud Auditability: An efficient facility for storing and mining cloud event logs

Cloud computing has been appealing its highly scalable, convenient, efficient, and cost effective services to become an alternative to the traditional infrastructure-oriented services. Despite its promising benefits, the wide adoption of cloud services, however, has been hesitated by users mainly because of the lack of ***security***, ***accountability,*** and ***auditability***. For deploying the practical cloud, in addition to building firewalls to close up ports and enhancing the authentication to regulate the access, ***an effective event tracking and recording facility*** is one of the most fundamental and inevitable components. Cloud event logs could be designed as free-form texts with a minimal structure. Because of sheer volume, diversity, and complexity, the typical recording process may require significant storage and network resources. In addition, examination required by an auditor or on-going troubleshooting investigation could cause heavy system performance overheads.

We propose a new paradigm of cloud event record and storage system architecture that automatically transforms low-level minimally-structured event logs into meaningful and prioritized high-level cloud service events, using powerful data mining techniques tailored to the problem domain. Through this research, we will identify potential performance bottlenecks and scalability issues that may arise in large-scale deployments. Our goal is to provide service providers with actionable insights from the event logs—in near real-time—to ensure secure and accountable usage of cloud resources. The long term vision of the research is not limited to the data description technology development and the enhancement, but 1) to replace the traditional system-wide protection concepts with the service-oriented concepts, 2) to make seamless event processing with true cloud virtualization, and eventually, 3) to equip the secure cloud computing infrastructure.

Cyberinfrastructure Federation: It can be a new recipe for a pandemic disaster

Despite many foreseeable benefits of Cyberinfrastructure (CI) federations, such as performance synergy and resource balancing, it should be clearly noticed that CI federation also means a federation of problems. It may also synergize disaster spreading. Although the interoperability policies of CI federation may include the problem resolution among the entities, they often overlook global problem resolution policies and associated facilities required to ensure such policies. A new problem can be created when multiple CIs are federated. There is very little facility to perform the root cause analysis over the multiple federated environment. The root cause can be also contributed partially or accumulatively by each entity. Hence, root cause analysis and testing facilities for the federated CI should be approached as a totally different problem domain, instead of one-to-one functionality mapping.

We list several important research aspects for next-generation, federated CIs:

- Physical bypassing and logical tunneling over other federated CIs should be treated differently;
- When a new function is added within a CI, its impact should be tested in the federated environment. The testing facility should be available;
- Each CI should have a sandbox for check-pointing and an abstraction layer for limiting internal problems within a domain;

- A facility to identify and contain problems in an intermediate area should be available;
- A root cause investigation framework should be transparent to each CI. For that, each CI should provide sufficient information on a designated area.

Software-Defined Networking (SDN) Scalability

Despite SDN's promises of flexibility and simplicity, the abstractions toward the remote and centralized control tend to extend the legacy network management's inaccurate and unreliable problems into the control plane. SDN also imposes excessive control traffic overheads in order for the controller to acquire global network visibility. More significantly, overheads will be further increased by traditional network management events as well as the application specific control traffic, as they may use the same physical network paths, buffers, and I/O channels at the same time. Overheads will be even worsened if the control uses an in-band network sharing with the data plane. If overheads are not controlled properly, they can cause various scalability problems on networking devices, controllers, and network itself, including slow message processing, potential message drops, delayed root cause analysis, and late responses against urgent problems.

We propose to design a comprehensive classification and prioritization facility that serves as a fundamental approach to improve the scalability of SDN. The proposed facility intends to reduce the detrimental impact of an overloaded system by incorporating the vitality of individual control messages. ***Selective process facility*** reduces control message processing overheads on the controller by identifying the essential messages to be processed among the related messages. ***Delegation and registration facility*** expedites the response time for the urgent issues by delegating actions to the immediate controllers, servers, or switches. ***Correlation facility*** enhances the root cause analysis capability by providing intelligences for the classified and prioritized control messages.

Software-Defined Networking (SDN) High Availability (HA)

High Availability (HA) is one of the most critical requirements in a real network operation. Thus, in traditional network systems, an HA architecture consists of many distributed protocols (for redundancy control, synchronization, and failure detection, notification, and mitigation) built-in dedicated network devices according to the physical network topologies. However, with the recent trend of Software-Defined Networking (SDN) and Network Function Virtualization (NFV), the control plane is logically centralized and decoupled from the data plane, and network functions are virtualized from dedicated hardware to general software. It poses many critical challenges on the existing HA mechanisms to achieve the same HA Service Level Agreement (SLA) for services in SDN and NFV environments.

We identify that the new architecture poses more complex HA issues by creating a new connection network between control and data planes, as well as HA issue becomes more critical at the 'controllers' than at the 'switches', as controllers are responsible for the intelligent decision of virtualization and switch policies. Consequently, we propose several control path HA algorithms and protocols that align logical redundancy along with physical redundant topologies between control and data planes, virtualize the controller cluster structure, and expedite failure detection and mitigation.

Towards an Emulator for Software-Defined Wireless Networks

Ramon dos Reis Fontes*, Christian Esteve Rothenberg*
*University of Campinas, São Paulo, Brazil
ramonrf,chesteve@dca.fee.unicamp.br

I. Context

Network emulation [1] has been widely used in performance evaluation, testing and debugging protocols and also in several related topics on research in computer networks. In this talk, we will present Mininet-WiFi as a tool to emulate wireless OpenFlow/SDN [2, 3] scenarios, providing experiments in Software-Defined Wireless Networks [4, 5], also aiming high-fidelity experiments that replicate real networking environments. Mininet-WiFi augments the well-known Mininet [6] emulator with virtual wireless stations and access points while keeping the original SDN capabilities and the lightweight virtualization software architecture. We will present a demo use case of a mobile video streaming scenario showcasing the ability of Mininet-WiFi to emulate the wireless channel in terms of bandwidth, packet loss, and delay variations as a function of the distance between the communicating parties.

II. Mobility Demonstration

This section illustrates the use of mobility inside of Mininet-WiFi and its ability to considering signal propagation aspects taking into account the distance between a station and the associated access point(s) in addition to the potential benefits of more than one simultaneous association (Figure 1). In this experiment, we use a mobility model called *Random Direction*, where stations move towards to the limit of a defined area also varying its velocity and an OpenFlow controller operating in reactive mode. The mobility models are important in the context of wireless networks, since they describe the behavior of the nodes moving on the space and its properties, like direction and speed.

The video published in [7] shows that the quality of video varies according the station movement and when it is connected simultaneously on two access points. This case is inspired on a related research work [8] carried in 2011 by the OpenFlow research group at Stanford University. In their use case, the station (or laptop) had three wireless interfaces (two WiFi and one WiMAX). In our replicated experiments, we carried similar experiments but using only the two WiFi intefaces. In all scenarios, the station is able to associate automatically to any access point that it can reach, so there are moments that the station should be associated either to one or two access points.

In future work, we intend to support existing and consolidated propagation models and the impact of distance in bandwidth, packet loss, and delay. The latest implementation is available in the Mininet-WiFi code repository.*https://github.com/intrig-unicamp/mininet-wifi*

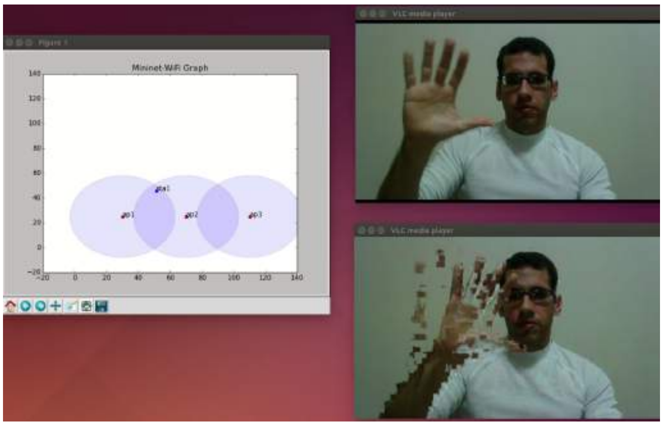

Fig. 1. Practical demonstration of using all wireless networks around us within Mininet-WiFi.

References

[1] K. R. Fall, "Network emulation in the vint/ns simulator," in *Proceedings of International Symposium on Computers and Communications (ISCC)*, pp. 244–250, 1999.

[2] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, pp. 69–74, Mar. 2008.

[3] D. Kreutz, F. M. V. Ramos, P. Veríssimo, C. E. Rothenberg, S. Azodolmolky, and S. Uhlig, "Software-defined networking: A comprehensive survey," *CoRR*, vol. abs/1406.0440, 2014.

[4] N. A. Jagadeesan and B. Krishnamachari, "Software-defined networking paradigms in wireless networks: A survey," *ACM Comput. Surv.*, vol. 47, pp. 27:1–27:11, Nov. 2014.

[5] M. Yang, Y. Li, D. Jin, L. Zeng, X. Wu, and A. V. Vasilakos, "Software-defined and virtualized future mobile and wireless networks: A survey," *CoRR*, vol. abs/1409.0079, 2014.

[6] B. Lantz, B. Heller, and N. McKeown, "A network in a laptop: Rapid prototyping for software-defined networks," in *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, Hotnets-IX, (New York, NY, USA), pp. 19:1–19:6, ACM, 2010.

[7] R. Fontes, "Using all the wireless networks around us within Mininet-WiFi." https://www.youtube.com/watch?v=_C4H2gBdyQY, 2015. [Online; accessed 07-October-2015].

[8] K. M. K.K. Yap, Te-Yuan Huang and Y. Yiannis, "Using all the wireless networks around us." http://archive.openflow.org/wp/uninterrupted-streaming-from-moving-golf-cart-with-openflow-wireless, 2011. [Online; accessed 07-October-2015].

Scalable Emulation of SDN Applications with Simulation Symbiosis

Jason Liu[1] and Cesar Marcondes[2]
[1] Florida International University, Miami, Florida, USA
[2] Federal University of São Carlos, São Carlos, SP, Brazil

I. Overview

Software Defined Networks (SDN) is an inspiring technology that promises to make networks both easier and cheaper to operate. Therefore, SDN makes networks flexible and programmable, by using a language so simple that average software engineers can comprehend and use to quickly reconfigure centrally the network, without dealing with distributed protocols manually configured that are prone to errors and time-consuming.

The capability of studying and experimenting with large-scale Future Internet SDN applications is of significant importance. For example, the Global Environment for Network Innovations (GENI) has been a community-based effort for building a collaborative and exploratory network experimentation platform for studying future network applications [1]. Follow-up efforts include various cyber-infrastructure design, development, and build-out projects, such as NSFCloud [2], [3], for building mid-scale cloud-computing testbeds in the U.S. There are similar attempts made in European Union, Japan, Brazil, and other nations.

While all these efforts would pave the way for the network researchers (as well as the network engineers) to validate design and implementation issues directly on the cyber-infrastructure testbeds, one needs to understand the deficiencies of solely relying on real-world implementation and physical deployment in network studies. We illustrate this important issue through a few hypothetical examples:

- A new robust map-reduce algorithm [4] needs to be evaluated for multi-tenant cloud computing environments. The performance of the algorithm depends on the job characteristics (such as the distribution on the number of jobs and the individual job sizes), as well as the configuration and stability of the available resources of the cloud platform. One would find it extremely time-consuming to explore the entire algorithmic parameter space on physical testbeds; let alone the highly diverging cloud configurations.
- A novel enterprise network traffic engineering solution based on OpenFlow [5], which uses opportunistic traffic load balancing and multi-path schemes to increase the throughput of heavy-hitter flows, has been proposed. Important questions remain unanswered—for example, whether this algorithm is robust under various traffic conditions, whether the algorithm would perform well due to partial deployment with varying proportions of non-cooperative entities, and whether the algorithm could scale out to a larger number of ISPs.
- A data center transport-layer protocol has been proposed (similar to [6]), which is expected to both reduce flow completion time and increase data throughput. The algorithm has been implemented and tested in a small-scale homespun DCN testbed; one needs to know whether it is ready for deployment in a production data center. Before that, however, one would like to investigate the algorithm's optimal performance conditions for the large data center with high bisection network capacity and also with various traffic loads with known stochastic properties.

These examples highlight some of the intrinsic limitations of cyber-infrastructure testbeds. Another popular method is to use emulation. For example, the current widespread use of container-based emulation combined with a software switch SDN capable have change the way research is conducted in SDN. Mininet is a popular container-based emulation environment built on Linux for testing OpenFlow applications. Using Mininet, one can create network experiments using a set of virtual hosts and virtual switches connected as an arbitrary network. With Mininet, it is relatively easy to build and execute experiments. However, it is well-known that Mininet only provides a limited capacity for both CPU and network I/O. Consequently, it does not work well on large scenarios and topologies with large volume of traffic, even if used in a cluster environment.

We propose a method for applying a symbiotic approach that combines simulation and emulation for improving the scalability of network experiments. The essential aspect of our approach is to remove a large portion of traffic from emulation. Rather than re-creating each network packet generated from applications, we can capture in real time the aggregate traffic demand of these applications and simulate the corresponding effect on the network queues (effective bandwidths, packet loss, and packet delays), that can affect other applications.

With the symbiotic approach, we achieve two objectives. First, we can effectively integrate emulation and simulation so that one can conduct hybrid network experiments with flexibility and scale. For example, we can use Mininet to directly run SDN applications using the virtual machines and software switches, whereas the connections between the virtual machines can be integrated with large-scale network models performed in simulation. In this case, the SDN applications can

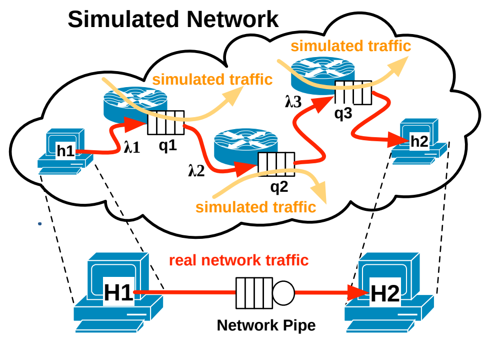

Fig. 1. A symbiotic network experiment.

be tested under diverse virtual network scenarios. Second, the symbiotic approach can be used to coordinate separate Mininet instances, each representing a set of different yet possibly overlapping network flows. In this case, we can significantly increase the scalability of the network emulation experiments in a cluster environment.

II. Symbiotic Simulation

We propose a symbiotic approach that can effectively combine simulation and emulation [7]. A network experiment consists of a virtual network with an arbitrary topology, potentially with a large number of hosts and routers. For a specific experiment, we can examine a subset of network protocols and applications by directly running them on an emulation testbed. Fig. 1 shows an example where the real applications are running on two physical hosts, H1 and H2. To test them, we specify a simulated network, which contains virtual hosts, h1 and h2, that correspond to the two physical hosts. We "modulate" the real network traffic between the physical hosts using statistics collected from the simulated network. More specifically, we use a facility, called the "network pipe", to represent the sequence of network queues *supposed* to be traversed by the real network traffic if it were placed on the simulated network. The example shows one network pipe consisted of three simulated network queues: q_1, q_2, and q_3. (For brevity, we focus only on the forward traffic from H1 to H2 and ignore the traffic in the reverse direction.) It is important to know that with symbiotic simulation, the network packets generated between the physical hosts, from H1 to H2, do not need to be captured and simulated individually as in real-time simulation. Instead, the symbiotic system captures only the traffic demand at the physical hosts and then sends this information to the simulator so that the simulator can regenerate the same traffic and model its effect over the simulated network (e.g., packet delays and losses).

The network pipe is a mechanism used by the emulator to reflect the traffic conditions in the simulated network so that the packet delays and losses can be applied to the real traffic. In [7], we derived a closed-form solution, for which we only capture the main results below.

In general, let $q_1, q_2, \cdots, q_n$ be the list of network queues in simulation that are supposed to be traversed by the real network traffic. In simulation, we collect three measurements for each queue q_i and periodically report them to the emulator: 1) We measure p_i, which is the average drop probability due to buffer overflow; 2) We measure λ_i, which is the arrival rate of the regenerated emulated network flow; and 3) We measure w_i, the average packet queuing delay. Once these measurements are propagated to the emulator, we can calculate the packet drop probability for the network pipe from individual loss drop probabilities of the constituent queues. And we can calculate the service rate (i.e., the bandwidth) of the network pipe:

$$\mu = \frac{\lambda_p(\Delta T + W_2 - W_1)}{\Delta T\left(1 + W_1\lambda_p - \sqrt{1 + W_1^2\lambda_p^2}\,\right)}$$

where $\lambda_p = \min_{1\leq i\leq n}\{(1-p_i)\lambda_i\}$, which is the minimum effective arrival rate at all queues; ΔT is the sample interval (say, 100ms), which is also the interval at which the simulator updates the emulator with the measurements; $W_1 = \sum_{1\leq i\leq n} w_i$ is the total queuing delay through the n queues measured in simulation; and W_2 is the average packet queuing delay through the corresponding network pipe measured in emulation.

After calculating p and μ, we can apply them at the network pipe in the emulator, which is essentially a first-in-first-out queue installed between the physical hosts. The network pipe will randomly drop packets according to the set probability p and process packets according to the given bandwidth μ, which will effectively add queuing delays to the packets as they go through the network pipe. In this paper, we will show how to implement the network pipe using the Linux traffic control (tc) facility.

In [7], we conducted extensive experiments to show that this symbiotic approach is able to produce accurate results. Using this approach, we can test the real applications running on the physical environment with different network scenarios—such as running on different network topologies, testing with diverse traffic intensity, and using different workload and user demand. In this way, we can enable high-fidelity high-performance large-scale network experiments by combining both simulation and emulation testbed, using simulation for the full-scale detailed network representation and using emulation testbed for directly executing network applications for real. On the one hand, emulation testbeds can execute real applications, operate with real systems, accept real input, produce real output, and respond to real network conditions. They provide the operational realism and fidelity usually unattainable by modeling and simulation. On the other hand, simulation is expedient for constructing and testing models to obtain "the big picture", which would be highly valuable especially when a good understanding of the system's complex behavior is absent. Simulation makes it easy for prototyping, for exploring the design space, for assessing the performance in diverse network settings, and for investigating what-if scenarios.

III. Mininet Symbiosis

We describe our design for integrating the symbiotic approach with Mininet. A detailed discussion can be found in [8]. Our goal is to execute the target network applications in Mininet containers while creating an illusion that these applications are running on an arbitrary network. Our approach starts by first having the user to specify a network model, which includes a simulated network topology (on which the target real applications are expected to run), as well as network protocols and applications, and how they are engaged during the experiment.

Next, the user can identify a subset of hosts to be emulated in Mininet (we call them *emulated hosts*). They will be instantiated as containers and therefore capable of directly running the target network applications. To reduce overhead, we also ask the user to identify flows that will be generated between the emulated hosts during the experiment (we call *emulated flows*). This can significantly reduce the facilities that need to be maintained for symbiosis.

Afterwards, we invoke a process, called *downscaling*, in which the original full-scale network simulation model together with the identified emulation traffic is processed to produce an reduced emulation model for Mininet. The downscaling process first prunes the original network model and remove all hosts, routers, and links not traversed by emulated traffic, since they are not needed in emulation involving real traffic. It then compresses the pruned topology by combining the intermediate nodes and links visited by the same set of emulated flows into a single network pipe.

Our symbiotic system consists of a simulation system and an emulation system running side by side. The simulation system is a real-time network simulator (we use PrimoGENI [9] for our prototype implementation), and the the emulation system consists of one or more Mininet instances, potentially running on separate machines (see Fig. 2). Communication between the real-time network simulator and the Mininet instances is achieved via TCP connections, whereas the simulator functions as the server and each Mininet instance as a client. The real-time network simulator runs the original full-scale network; as such, it needs to implement necessary network elements (such as routers, hosts, network interfaces and links) and common network protocols (such as IP, TCP, UDP, and others). In addition, two components are added to the simulator to facilitate synchronization with the Mininet instances: a traffic monitor and a traffic generator. The traffic monitor is used to collect measurements at each queue q_i traversed by the emulated flows, which include the packet drop probability p_i, the arrival rate of emulated flows λ_i, and the queuing delay w_i. These measurements are collected periodically every ΔT units of time and then sent to the corresponding Mininet instances. The traffic generator receives information from Mininet about the traffic demand d_k from applications for each emulated flow k in terms of the number of bytes requested to be sent during the last interval. Upon receiving this information, the simulator generates the emulated flows by initiating the corresponding TCP or UDP sessions in simulation with the same demand size accordingly.

In Mininet, the emulated hosts are instantiated as Linux containers with separate network namespaces, and the switches

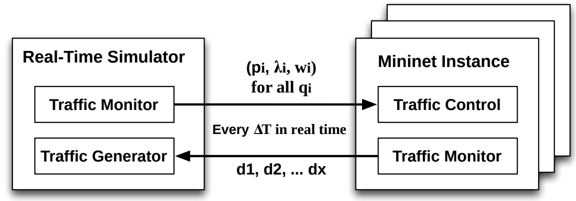

Fig. 2. Mininet symbiosis setup.

are represented by OVS instances. The virtual Ethernet (`veth`) pairs are used to represent the links augmented with the Linux traffic control (`tc`) for managing the link properties. Linux `tc` is a set of tools (included since kernel 2.2) to allow users to have fine-grained control over the packet transmission. Linux `tc` consists of different queuing mechanisms, easily composable for handling more complex situations (including packet mangling, IP firewalling, and bandwidth metering). We use `tc` for setting the link bandwidth, the packet delay, and the random packet loss probability. More specifically, we statically set the link delay as the cumulative propagation delay of the links between the consecutive queues that constitute the network pipe. We modify the packet loss probability and the link bandwidth dynamically during the experiment using the measurements from simulation.

Note that such symbiotic approach can support distributed emulation, where multiple Mininet instances can operate in parallel, each handling a different set of emulated flows. In this case, the state of common network pipes is mirrored on different instances, which will be controlled by the simulator with the identical link properties.

References

[1] NSF Global Environment for Network Innovations (GENI), http://www.geni.net/.
[2] CloudLab, https://www.cloudlab.us/.
[3] Chameleon - A configurable experimental environment for large-scale cloud research, https://www.chameleoncloud.org/.
[4] J. Dean and S. Ghemawat, "Mapreduce: simplified data processing on large clusters," in *Proceedings of the 6th Symposium on Operating System Design and Implementation (OSDI'04)*, 2004.
[5] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: Enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, 2008.
[6] A. Munir, G. Baig, S. M. Irteza, I. A. Qazi, A. Liu, and F. Dogar, "Friends, not foes - synthesizing existing transport strategies for data center networks," in *Proceedings of the 2014 ACM SIGCOMM Conference (SIGCOMM)*, 2014.
[7] M. A. Erazo and J. Liu, "Leveraging symbiotic relationship between simulation and emulation for scalable network experimentation," in *Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS)*, 2013, pp. 79–90.
[8] J. Liu, C. Marcondes, M. Ahmed, and R. Rong, "Toward scalable emulation of future internet applications with simulation symbiosis," in *Proceedings of the 19th IEEE/ACM International Symposium on Distributed Simulation and Real Time Applications (DS-RT)*, 2015, to appear.
[9] N. V. Vorst, M. Erazo, and J. Liu, "PrimoGENI for hybrid network simulation and emulation experiments in GENI," *Journal of Simulation*, vol. 6, no. 3, pp. 179–192, 2012.

Abstract Topology and Cost Maps for Software-Defined Inter-Domain Circuits

Malathi Veeraraghavan*, Christian Esteve Rothenberg†
* University of Virginia, Charlottesville, VA 22904–4743,mv5g@virginia.edu
† University of Campinas, São Paulo, Brazil, chesteve@dca.fee.unicamp.br

I. Problem Statement

Inter-domain routing protocols, such as Border Gateway Protocol (BGP), are critical to the scalability of the Internet, enabling address reachability between routers belonging to different Autonomous Systems (ASes). BGP software and policy configuration are very complex. Service providers pay equipment vendors substantial annual fees for the licensing and maintenance of route-processor software – including BGP and associated management/analytic tools. Under the Software-Defined Network (SDN) paradigm, such per-router annual software costs are expected to be replaced with presumably lower-cost software for centralized SDN controllers.

A research question of interest in whether simplifications are possible to inter-domain routing protocols for best-effort IP service in this new SDN paradigm. Further, for new services, such as dynamic rate-guaranteed Layer-2 (L2) path services and dynamic optical Layer-1 (L1) circuit services, new inter-domain, inter-SDN-controller protocols are called for. Methods for inter-domain routing to support Path Computation Engine (PCE) consider tradeoffs between hierarchical link-state protocols and path-vector protocols [1]. Proposals to share available bandwidth to reduce call blocking probability for inter-domain paths have been proposed [2]. Commercial providers would prefer to share abstracted topologies and available resources without disclosing technology/topology specifics. An open question is whether abstracted available bandwidth information is still useful/feasible in path selection.

II. Solution Approach

There is an effort called Application Layer Traffic Optimization (ALTO) underway in the IETF [3] to define a protocol that enables applications to request and receive information about network topology, path costs, link availability, routing policies, and end-host properties. Providers can offer network information in abstracted form since the ALTO protocol offers topology hiding. ALTO cost maps can be used to signal changes in the network state. Dynamic cost maps can reflect information gathered through real-time monitoring.

The ALTO cost abstraction allows for costs to be any metric. The cost variable does not need to be available bandwidth (or its inverse function), or real monetary costs. The costs may be relative, which makes it technology independent. Parameters such as AS-Path, Multi-Exit Discriminator (MED) used in BGP can be encoded as ALTO costs. A method for carrying link-state and traffic engineering information in BGP messages to support ALTO has been proposed in the IETF [4].

The ALTO protocol was proposed for IP-routed best-effort services. We propose to explore whether the ALTO protocol can be extended to support advance-reservation dynamic L2-path and L1-circuit services.

In a recent project, we deployed Open Exchange Software Suite(OESS) [5] and On-Demand Secure Circuits and Advance Reservation System (OSCARS) [6] controllers in 8 university campuses. These campus SDNs were connected via static L2 paths to Internet2's Advanced Layer 2 Service (AL2S) OpenFlow-switch based network. AL2S runs OSCARS and OESS to support dynamic L2 service. This multi-domain setup allowed us to experiment with dynamic advance-reservation L2-path service. The OSCARS server within a domain uses its topology service to push the topology of its domain to the perfSONAR (pS) Topology Service, which then allows OSCARS in other domains to request topology information when needed. While this open topology sharing approach works in the R&E community, it is not suitable for commercial providers. Therefore, as part of this work, we propose to study the feasibility of using the ALTO protocol for sharing inter-domain information required for advance-reservation dynamic L2-path and L1-circuit services.

Our starting point will be an implementation of ALTO in OpenDaylight. Our evaluation approach will combine simulations and experiments on the GENI testbed.

References

[1] M. Chamania and A. Jukan, "A survey of inter-domain peering and provisioning solutions for the next generation optical networks," *Communications Surveys Tutorials, IEEE*, vol. 11, pp. 33–51, First 2009.

[2] D. Siracusa, S. Grita, G. Maier, A. Pattavina, F. Paolucci, F. Cugini, and P. Castoldi, "Domain sequence protocol (DSP) for PCE-based multi-domain traffic engineering," *Optical Communications and Networking, IEEE/OSA Journal of*, vol. 4, pp. 876–884, Nov 2012.

[3] R. Alimi, R. Penno, Y. Yang, S. Kiesel, S. Previdi, W. Roome, S. Shalunov, and R. Woundy, "Application-Layer Traffic Optimization (ALTO) Protocol." RFC 7285 (Proposed Standard), Sept. 2014.

[4] A. F. H. GredlerJ. Medved, S. Previdi, "North-bound distribution of link-state and te information using BGP." draft-ietf-idr-ls-distribution-11.

[5] "Open exchange software suite (OESS)." http://globalnoc.iu.edu/sdn/oess.html.

[6] "On-demand secure circuits and advance reservation system (OSCARS)." http://www.es.net/OSCARS/docs/index.html.

POSITION PAPER

THE CAPACITY CRUNCH: AVOIDANCE, ADAPTATION AND REFORM

Helio Waldman

The exponential growth of Internet traffic, in the range of 50-100% a year, makes it a system in permanent disequilibrium, always approaching collapse in less than a decade, but often averting it. The capacity crunch has been avoided now for many decades, leading many people to believe that it will be avoided forever, just like the end of Moore's Law. Unfortunately, the physics, the engineering and the economics behind the Internet and data processing tell us that these beliefs are overly optimistic.

In the 1970's, communication networks were already facing a "transmission bottleneck" as the capacity of metallic cables reached saturation in the face of ever increasing (mostly voice) traffic. This early "crunch" was overcome in the 1980’s by the deployment of the first single-mode optical fibers and the development of advanced lasers. However, the multi-THz capacity of the optical fibers remained almost unexplored, as each fiber would carry only a single wavelength modulated by a few Gbps. An effort to fill this huge gap was then started under the name of "bandwidth mining". The foremost bandwidth mining technology of the 1990's was wavelength division multiplexing (WDM), which allowed the propagation of several wavelengths in the fiber, each one modulated with ten or a few tens of Gbps. These developments generated a capacity growth of a thousand times per decade, much more than the 100 times per decade afforded by Moore's Law to the data processing capacity of computers. By the end of the decade and century, then, the transmission bottleneck had been replaced by an "electronic bottleneck", as the electronic nodes were overwhelmed by a flood of data traffic, which had become dominant in communication networks.

In the past decade, (partial) solution of the electronic bottleneck demanded a thorough reshaping of the network architecture, reducing the number of instructions per packet so that critical functionalities, such as routing, may now be implemented in the optical domain with the help of emerging photonic capabilities. Today, bandwidth mining has advanced to the point where the WDM frequency grid effectively occupies the whole 4 THz gain bandwidth of the erbium-doped fiber amplifiers (EDFA’s) used to keep the signal power level along the fiber transmission lines. So now the next bandwidth mining step, - and probably the last that can be supported by the current physical plant, - is to improve the spectral efficiency, i.e. the number of bps per Hz of the spectral assignment.

Our research investigates the maximization of spectral efficiency in emerging Elastic Optical Networks (EON’s). In EON’s, users (most often service providers) do not request wavelengths from a fixed grid, but rather an amount of spectrum (specified by a number of contiguous frequency slots) that is strictly sufficient to meet their requirements, thus avoiding waste of spectral resources. The focus of our research is the routing and spectrum assignment (RSA) problem, which looks for the “best” possible assignment of route and spectral placement. By “best” one normally means most efficient in the overall algorithmic occupation of spectral resources. However, most efficient solutions are not fair among traffic classes, so we are looking for a good compromise between efficiency and fairness. The multiclass nature of EON-bound traffic is the single feature that makes elastic networking much more complex than plain wavelength routing of old WDM networks, thus demanding more intelligence in management and control.

Elastic networking should extract more capacity from currently installed fibers, but it is expected to saturate them before 2025. So the next step might be the building of a new infrastructure based on a new generation of multimode fibers where each wavelength carries a number of independent spatial modes, each one carrying a high bitrate. Mode coupling may occur both in signal modulation and propagation, but signals may be restored with sophisticated signal processing, albeit at the cost of losing important add-and-drop capabilities. So we envision the investigation of emerging networking problems associated with these prospective new infrastructures. Since their deployment is sure to demand huge investments, though, it is likely to subject the current Internet business model to some stress.

Finally, it is important to realize that spatial mode multiplexing may be a solution for sending more information per fiber, but not with the same power, as each spatial mode will require the same amount of power as current wavelengths in WDM, or spectral chunks in EON’s, for the same bitrates and distances. Therefore, space-division multiplexing (SDM) is not a solution for the energy bottleneck, which is likely to bring the ultimate capacity crunch in the 2030’s. When this happens, disaster avoidance will require architectural and business constraints that inhibit wasteful transmission of redundant data and/or fancy virtual environments, as well as some regulatory compromise between service universalization and differentiation.

In the long run, dealing with impending capacity crunches in the coming decades is bound to demand innovative solutions in technology, business and regulation.

The Policy Space for Software Defined eXchanges

Joaquin Chung*, Russell Clark^, Henry Owen*
*School of Electrical and Computer Engineering, ^School of Computer Science
Georgia Institute of Technology
Atlanta, GA. USA
joaquin.chung@gatech.edu, russ.clark@gatech.edu, henry.owen@ece.gatech.edu

***Abstract*—An Exchange Point is a common place where different organizations collaborate and share computing, storage and networking resources. Recently, Software Defined Networking technologies are being included into Exchange Points to optimize resource sharing and allocation. In this position paper we explore the policy space for a Software Defined eXchange, focusing on the applications that these policies support and the options that we have to express them.**

Keywords—Software Defined Exchange; inter-domain SDN; policy space

I. Introduction

In traditional Research and Education (R&E) networks, exchange points are used to share computing, storage and networking resources. For instance, data generated in the Large Hadron Collider in Europe could be consumed by scientists in Japan. Likewise, pictures of the night sky taken by a telescope in South America could be processed by a supercomputer in North America. The interconnection of different National Research & Education Networks (NRENs) allows this kind of collaboration[1], [2].

The commercial space has similarities to the R&E case. Internet eXchange Points (IXPs)[3] allow Internet Service providers (ISPs) and Content Providers (CPs) to exchange Internet traffic by peering in a common facility. Similarly, Hybrid Clouds allow users to move workloads from private to public clouds when the capacity of the private cloud is exhausted.

A Software Defined eXchange (SDX) seeks to introduce Software Defined Networking (SDN) and Cloud Orchestration technologies into the Exchange Point to optimize resource sharing and allocation. Depending on the resources shared and the technology used; the definition of and SDX can go from a Layer-3 Network SDX all the way to a Federated Cloud Service. In this position paper we will focus on Network Exchange Points and more specifically in the policy space that allows us to define more granular behavior in the network.

II. SDX Applications

The SDX policy space is closely related to the applications deployed over the network fabric. Gupta et al. [4] proposed four applications for a Software Defined Internet eXchange Point: application-specific peering, inbound traffic engineering, wide-area load balancing, and redirection through middle boxes. The authors used SDX policies to augment the capabilities of BGP policies in an IXP.

In the context of federated testbeds, the FELIX project[1] defines six applications classified into two main groups: Data Domain and Infrastructure Domain. Under the Data Domain the authors define: data-on-demand, data preprocessing, and high-quality media transmission over long-distance networks. The application for the Infrastructure Domain are data mobility for Inter-cloud use, follow the sun (or moon) principles for Datacenter, and disaster recovery by IaaS migration.

Finally, we ask the participants of the Southern Crossroads (SoX), the Research and Education Network for the U.S. Southeastern region (www.sox.net), to identify the types of challenges they face with peering in order to prioritize our efforts. Among all the suggestions, we can mention: bandwidth calendaring, time- and load-based resource reservation, and a visual interface for resource reservation.

III. SDX Policies

Considering the SDX application examples presented in the previous section, we can classify SDX policies into two groups: the ones based on packet header fields and those based on external data.

A. Policies Based on Packet Header Fields

In general, the policies based on packet header fields will match TCP or UDP source and destination ports, source and destination IP address or source and destination MAC addresses, and apply actions accordingly. If we use OpenFlow 1.3, the policy will be able to match other fields such as DSCP and MPLS labels.

For instance, the application specific peering example defined in [4] matches the TCP destination port of packets entering the IXP. In the scenario proposed by the authors (see Fig. 1), AS A wants to forward all the traffic destined to TCP port 80 to AS B, and all traffic destined to TCP port 4321 and 4322 to AS C regardless of the BGP decisions. To achieve this behavior, the operator can express a Pyretic policy in the SDX controller as follows:

```
match(dstport = 80) >>  fwd(B)
+ match(dstport=4321/4322) >> fwd(C)
```

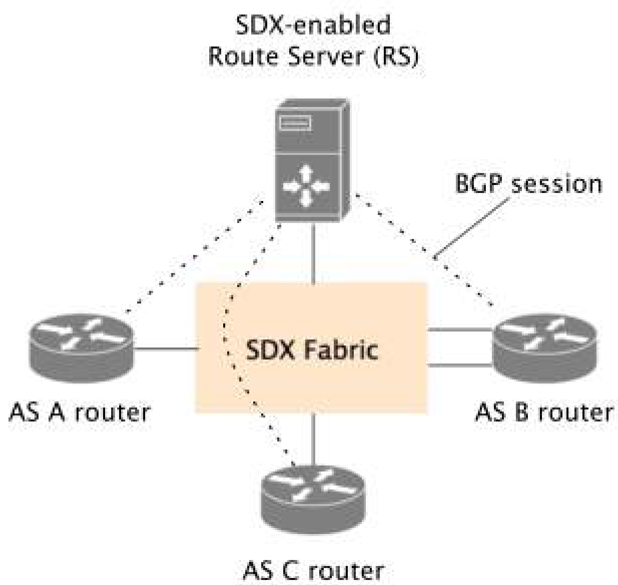

Fig. 1. SDX application specific peering

B. Policies Based on External Data

Policies based on external data collect information from other systems such as: network monitoring systems, user databases, DNS or NTP server and take actions accordingly. In this case, the policies will match parameters such as network latency, bandwidth, user name, domain name, date and time. In the follow the sun/moon application, the load on a datacenter could be moved to another datacenter depending on the time. Data-on-demand and high-quality video applications will require latency and bandwidth information to select the less congested path in the network.

IV. ONGOING WORK

High level network policy languages like Pyretic could be used to define SDX policies. However, Pyretic is limited to match packet header fields. NetAssay[5] extended Pyretic to match domain names from DNS servers and user names from LDAP or Active Directory services. On the other side, well known interfaces such as the RESTful API or JSON could be used to express SDX policies. Furthermore, the SDN field is currently moving towards the intent-based network paradigm, in which the operators define what they want the network to do, instead of defining how they want the network to do it. In our future work, we will evaluate all these approaches to express SDX policies seeking the most optimal solution.

ACKNOWLEDGMENT

The authors will like to thank NSF, GENI, GT-RNOC and SoX. Additionally, special thanks to the network operators from the SoX member community for the suggestions in new applications and policies for the Software Defined eXchange.

REFERENCES

[1] G. Carrozzo, R. Monno, B. Belter, R. Krzywania, K. Pentikousis, M. Broadbent, T. Kudoh, A. Takefusa, A. Vieo-Oton, C. Fernandez, B. Puvpe, and J. Tanaka, “Large-scale SDN experiments in federated environments,” in *2014 International Conference on Smart Communications in Network Technologies (SaCoNeT)*, 2014, pp. 1–6.

[2] J. Ibarra, J. Bezerra, H. Alvarez, M. Stanton, I. Machado, E. Grizendi, and L. F. Lopez, “Benefits brought by the use of OpenFlow/SDN in the AmLight intercontinental research and education network,” in *INTERNATIONAL SYMPOSIUM OF INTEGRATED MANAGEMENT OF NETWORKS OF THE IEEE*, 2015, pp. 1–6.

[3] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, “There is more to IXPs than meets the eye,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 5, pp. 19–28, 2013.

[4] A. Gupta, E. Katz-Bassett, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, and R. Clark, “SDX,” *ACM SIGCOMM Comput. Commun. Rev.*, vol. 44, no. 4, pp. 551–562, Aug. 2014.

[5] S. Donovan and N. Feamster, “Intentional Network Monitoring: Finding the Needle without Capturing the Haystack,” in *Proceedings of the 13th ACM Workshop on Hot Topics in Networks (HotNets)*, 2014, pp. 1–7.

Position Paper: Exploring research collaborations between University of Puerto Rico and Brazil

José Ortiz-Ubarri
Computer Science Department
University of Puerto Rico - Rio Piedras
jose.ortiz23@upr.edu

In the past few years there has been an emergence in big data science research supported by many federal agency in the United States. Big data research is a big international topic in natural, computer, and social sciences, and because of the current international collaborations it is not a surprise that there is an upsurge of projects that aim to create or strengthen cyberinfrastructure to support big data research.

For the last few years my research group has dedicated a lot of time in educational projects to improve our Computer Science (CS) curriculum and the way CS is taught at the University of Puerto Rico [1,2,3], and cyber infrastructure to support big data transfers [4], but our main research interests are computer and network security and the application of high performance technologies to solve computationally intensive scientific problems. In the past we have applied different high performance computing technologies to find secure codes that are suitable for applications to multiple target recognition, optical orthogonal codes, and digital watermarking [5,6,7,8]. These kind of codes (information theory codes) can be constructed using algebraic constructions, for which exists codes for infinite sizes, or by exhaustive search that reveal codes but with finite sizes. In our research we have applied exhaustive search algorithms to enumerate codes upto computationally feasible sizes, to answer open questions about the existence of such codes, or research whether there is an algebraic construction that generates codes for infinite sizes.

At the same time our research group has been working on Toa, a web based application for network situational awareness at scale [9, 10], and also to perform network forensics using network flows. We work with network flows because it allows us to monitor network traffic at the network, host and service levels, and it also permits to monitor network traffic from one network to another network efficiently. Network flows data allows us to monitor at those levels without storing all the traffic that pass over a large network, but only a summary of the connections. Currently we are working on adding IPv6 support to Toa and integrating with the CERT silk toolkit. As research institutions are increasing their bandwidth to support big data research, our project aims to provide monitoring for such networks efficiently.

Brazil has a big Information Theory group with which our group in Puerto Rico could potentially perform collaborations. Our group could also apply its expertise in high performance computing to collaborate in current big data scientific projects. Collaborations to research the security of the new network infrastructure could also arise from this workshop. Specifically aligned to the topics areas of this workshop, we could create or join a collaboration to research the use of network flow data analysis to guide the configuration or auto configuration of Software-defined networks (SDN)[11].

References:

[1] E. Orozco, R. Arce-Nazario, J. Ortiz-Ubarri and H. Ortiz-Zuazaga. A Curricular Experience With Parallel Computational Thinking: A Four Years Journey. In Proceedings of EduPDHPC, Denver, Colorado, USA, 2013.

[2] J. Ortiz-Ubarri, R. Arce-Nazario, I. Rubio. Development of engaging and readily transferable laboratory experiences for the introductory programming course. National Science Foundation under Grant No. DUE-1245744.

[3] J. Ortiz-Ubarri, H. Ortiz-Zuazaga, R. Arce-Nazario, P. Ordoñez. Academics and Training for the Advancement of Cybersecurity Knowledge in Puerto Rico (ATACK-PR). National Science Foundation under Grant No. DUE-1438838.

[4] J. Ortiz-Ubarri, H. Ortiz-Zuazaga, R. Arce-Nazario. Perimeter Network to Expedite the Transmission of Science (PR-NETS). National Science Foundation under Grant No. ACI-1340959.

[5] J. Ortiz-Ubarri. New families of asymptotically optimal doubly periodic arrays with ideal correlation constraints. Cryptography and Communications(2015): 1-12.

[6] J. Ortiz-Ubarri, O. Moreno, A . Tirkel, R. Arce-Nazario, S. Golomb. Algebraic symmetries of generic (m+1) dimensional periodic Costas arrays. IEEE Transactions on Information Theory 59(2):1076-1081 (2013).

[7] R. Arce-Nazario, J. Ortiz-Ubarri. Multidimensional Costas arrays and their enumeration using GPUs and FPGAs. International Journal of Reconfigurable Computing (2012).

[8] R. Arce-Nazario, J. Ortiz-Ubarri. Enumeration of Costas Arrays in FPGAs and GPUs. International Conference on ReConFigurable Computing and FPGAs, Cancun, Mexico, 2011.

[9] J. Ortiz-Ubarri, H. Ortiz-Zuazaga, A. Maldonado, E. Santos, J. Grullón. Toa: A Web-Based NetFlow Data Network Monitoring System at Scale. Proceedings of the IEEE Big Data Congress, New York, USA, 2015

[10] J. Ortiz-Ubarri, H. Ortiz-Zuazaga, A. Maldonado, E. Santos, J. Grullón. Toa: A Web-Based NetFlow Data Network Monitoring System. In Proceedings FloCon 2015, Portland Oregon. January 2015.

[11] J.Ortiz-Ubarri, W. Cabán. Private communication.

Software Defined Measurement for Scalable Network Monitoring

Yan Luo
Department of Electrical and Computer Engineering
University of Massachusetts Lowell
yan_luo@uml.edu , http://acanets.uml.edu/~yluo/

Network measurement is instrumental in quantifying the performance and status of network infrastructure, understanding network usage activities, and helping network operators to maintain networks and respond to network issues. Network measurement, however, faces several challenges: (1) the ever increasing line rate. It is non-trivial to provide detailed per-flow measurement data at a line rate of 40+Gbps; (2) the diversity and dynamics of measurement tasks. The measurement target could change rapidly, so do the interested flow metrics; and (3) deriving knowledge from measurement data. The sheer amount of measurement data has introduced a big data challenge on how we store and process them to learn interesting patterns and predict trends. We are motivated to address these challenges with new hardware architecture and software framework.

We believe hardware accelerated flow measurement is the effective approach to scale with the link speed while achieving fine grained measurement results. In particular, a network processor (NP), which is highly optimized for both performance and programmability, suits best in flow classification and metrics collection on high speed network flows. Our NSF IRNC project is leveraging Netronome NFP6xxx processors to accelerate the flow measurement tasks at the network interface card level, where the majority of low level flow statistics can be obtained at close to the packet fast path without the overhead of moving packets across the system interconnects such as PCIe. As a result, the many-core architecture of the NP can be leveraged to process the network flows in parallel at a speed up to 100Gbps. Only the interested flow statistics, as opposed to the packet headers and payloads, are pushed to the flow feature processors (often x86 based) for further analysis. Such decoupling of measurement plane (using NPs) and control plane (e.g. x86 architecture) makes it easier to focus on different parts of the measurement tasks and enables scalable measurement on high-speed network flows.

Programming a network measurement task has become increasing important and interesting, because the measurement objectives typically vary from time to time in scope, granularity, conditions, and actions. Using fixed measurement metric for every single flow is not necessary or feasible as one may need to zoom in or zoom out on a particular flow. The measurement may have to start when a given condition is met. There is a strong need to support the flexibility, and a software-defined approach is very promising. Our research in Software Defined Measurement (SDM) proceeds in two areas: (a) SDM data plane and (b) a control plane for SDM.

(a) SDM Data Plane: We are working to define a set of APIs for instantiating a measurement task with its measurement target, metrics, algorithms, and conditions of start/stop it. The APIs will be implemented on programmable measurement instrument (based on NP or x86), meeting the diverse and dynamic needs of network operators. These APIs will work closely with network domain specific languages such as P4 and programmable hardware architectures such as NP. Our goal is to provide a suite of libraries and applications that carry out representative network measurements, with the extension to future new measurement tasks;
(b) SDM Control Plane: The SDM control plane is to orchestrate large scale collaborative measurement given a growing set of measurement devices such as perfSONAR hosts. As the first step, we extend the widely deployed perfSONAR tools with a control plane supporting software defined measurement. We also leverage hardware accelerators to achieve ultra-high speed flow-granularity measurement for in-depth analytics of network flows.

The analysis of network flow data obtained from measurement instruments is challenged with the data volume and a variety of flow features. This calls for new systems to handle the streaming data, new algorithms to derive knowledge from data, and new control schemes to take advantages of the learned network usage patterns. In this direction, we have designed a highly scalable heterogeneous architecture with CPU/GPU called HeteroSpark to speed up the machine learning algorithms applied to the large amount streaming measurement data. We plan to build network traffic models based on the measurement data and provide tools for network monitoring and trend prediction.

Reference

[1] Peilong Li, Xiaobing Huang, Tian Zhao, Yan Luo, and Yu Cao. Sparkling: Identification of task skew and speculative partition of data for spark applications. In Spark Summit, San Francisco, CA, USA, 2014.

[2] Peilong Li, Yan Luo, and Jun Yang. Transformer: Run-time reprogrammable hetero- geneous architecture for transparent acceleration of dynamic workloads. Journal of Parallel and Distributed Computing, 86:45 – 61, 2015.

[3] Peilong Li, Yan Luo, Ning Zhang, and Yu Cao. HeteroSpark: A heterogeneous cpu/gpu spark platform for machine learning algorithms. In Networking, Architecture and Storage (NAS), 2015 IEEE International Conference on, pages 347–348, Aug 2015.

[4] Thomas Calloway, Peilong Li, Yan Luo and Imran Vakil. Exploration of memory hierarchy for heterogeneous computing architecture. In The 3rd Workshop on SoCs, Heterogeneous Architectures and Workloads, New Orleans, LA, USA, 2012.

[5] Chen Xu, Xiaoban Wu, Yan Luo, Brian Tierney and Jeronimo Bezerra, PSCP: PerfSONAR Control Plane for Software Defined Measurement and Troubleshooting. In submission.

Communications-aware job scheduling in hybrid SDN datacenters

Malathi Veeraraghavan*, Nelson Fonseca†, Andrea Fumagalli‡
*University of Virginia (UVA), Charlottesville, VA 22904, USA
†Institute of Computing of The University of Campinas, Campinas, Brazil
‡The University of Texas at Dallas, Dallas, TX 75080, USA

Topic areas: Software-defined networking (SDN), and Cloud computing and data center networking

Background: Optical-switched networks consume lower levels of energy when compared to electronic-switching networks. Hybrid electronic-optical networks have been proposed for datacenters, such as Helios [1], c-Through [2], and OSA [3]. However, owing to the cost of DWDM networks, each lightpath (wavelength circuit) is typically operated at 40 or 100 Gbps. This constraint allows only a few top-of-rack (TOR) switch (typically Ethernet or InfiniBand) pairs to be interconnected via optical circuits at any given time. A new technology that supports lower circuit rates while simultaneously consuming lower energy, called Digital Subcarrier Multiplexing (DSCM), was proposed [4]. With DSCM, 100 subcarriers can be multiplexed on a single wavelength. Assuming that 160 wavelengths can be carried on a single fiber, DSCM can support 16,000 channels per fiber. With OFDM, 512 subcarrier channels are carried on a 40Gb/s system, which implies that circuit rates on the order of tens and hundreds of Mbps can be supported. Digital Subcarrier Cross-Connect (DSXC) switching times are in the sub-milliseconds range. In summary, the digital subcarrier technology has the potential to support higher-granularity circuits, while simultaneously consuming low energy, and is hence considered in this study for use in hybrid datacenter software-defined networks (SDNs).

Problem statement: In commercial data centers, applications use Hadoop MapReduce for analyzing large datasets. In scientific data centers, computing jobs that require thousands of cores are submitted to batch schedulers such as PBS, LSF, or SLURM. These schedulers assign the jobs corresponding to a single application to different hosts, which could be distributed across several racks. Highly parallelized applications are often limited by network throughput when using 1000s of cores. The problem statement of this study is to improve communication performance within datacenter networks, while simultaneously reducing energy consumption.

Solution approach: We plan to use an approach proposed by Wang et al. [5] in which application-awareness is leveraged to determine circuit allocations between ToR electronic (packet) switches. Specifically, network-aware job scheduling is proposed for Hadoop jobs with "rack-based bin-packing placement to aggregate tasks on to a minimum number of racks." Such a job placement method then allows for ToR switches to be interconnected via an optical circuit for faster inter-task communications. An SDN controller allows for an easier integration of the Hadoop job scheduler with the scheduling and provisioning of optical circuits. Our DSCM technology allows for a greater number of racks to be interconnected simultaneously. Our project will study Hadoop job communication patterns and design a solution that integrates a job scheduler, SDN controller, and DSCM/WDM optical switches. A recent paper by Ren et al. [6] offers a public dataset characterizing Hadoop workloads, which will serve as a starting point for our communication pattern analysis. We also plan to run our own Hadoop Mapreduce software for network traffic analysis as was done by Lee et al. [7], and collect network usage statisics.

Other research topics: We are also working on the other topic areas of the SwitchOn workshop, Cyber-infrastructure and GENI and other federated testbeds. Specifically, an NSF MRI grant was used to deploy a distributed Dynamic Network System (DYNES). Leveraging the DYNES equipment, and equipment deployed through the NSF CC-NIE program grants, we are creating inter-domain dynamic Layer-2 (L2) virtual circuits between campuses. High-speed file transfer applications are being tested. We are also developing a solution for inter-domain Layer-1 (L1) end-to-end optical circuits and will be testing 10 Gb/s transfers on an Internet2 supported testbed. We would like to explore creating inter-domain L2, and if possible L1, circuits from UVA campus to Campinas campus.

References

[1] N. Farrington, G. Porter, S. Radhakrishnan, H. Bazzaz, V. Subramanya, Y. Fainman, G. Papen, and A. Vahdat, "Helios: a hybrid electrical/optical switch architecture for modular data centers," in *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4. ACM, 2010, pp. 339–350.

[2] G. Wang, D. Andersen, M. Kaminsky, K. Papagiannaki, T. Ng, M. Kozuch, and M. Ryan, "c-Through: Part-time optics in data centers," in *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4. ACM, 2010, pp. 327–338.

[3] K. Chen, A. Singla, A. Singh, K. Ramachandran, L. Xu, Y. Zhang, X. Wen, and Y. Chen, "OSA: An optical switching architecture for data center networks with unprecedented flexibility," in *Proceedings of the 9th USENIX conference on Networked Systems Design and Implementation*. USENIX Association, 2012, pp. 18–18.

[4] R. Hui, W. Huang, Y. Zhang, M. Hameed, M. Razo, M. Tacca, and A. Fumagalli, "Digital subcarrier cross-connects (DSXCs)," in *Transparent Optical Networks (ICTON), 2012 14th International Conference on*, July 2012, pp. 1–6.

[5] G. Wang, T. E. Ng, and A. Shaikh, "Programming your network at run-time for big data applications," in *Proceedings of the First Workshop on Hot Topics in Software Defined Networks*, ser. HotSDN '12. New York, NY, USA: ACM, 2012, pp. 103–108. [Online]. Available: http://doi.acm.org/10.1145/2342441.2342462

[6] K. Ren, Y. Kwon, M. Balazinska, and B. Howe, "Hadoop's adolescence: An analysis of Hadoop usage in scientific workloads," *Proc. VLDB Endow.*, vol. 6, no. 10, pp. 853–864, Aug. 2013. [Online]. Available: http://dx.doi.org/10.14778/2536206.2536213

[7] Y. Lee and Y. Lee, "Toward scalable internet traffic measurement and analysis with Hadoop," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 1, pp. 5–13, Jan. 2012. [Online]. Available: http://doi.acm.org/10.1145/2427036.2427038

Simulation and Emulation the Future Internet and Applications
A Research Collaboration between FIU and UFSCar

Jason Liu* and Cesar Marcondes+
* Florida International University (FIU)
+ Universidade Federal de São Carlos (UFSCar)

The paper briefly describes the ongoing collaboration between Prof. Jason Liu from Florida International University, Miami, Florida, US and Prof. Cesar Marcondes from Universidade Federal de São Carlos, São Carlos-SP, Brazil. Started in August 2014, the collaboration has been primarily focused on the Future Internet test bed development and on the transport-layer support for software-defined networks (SDNs).

Jason Liu's research includes applying parallel and distributed discrete-event simulation and high-performance modeling techniques for simulation of large-scale complex systems, including computer systems and communication networks. His overall research has three focuses. The first research focus is to design efficient and scalable parallel algorithms and develop efficient tools to enable either time-critical or large- scale simulations on high-performance computing (HPC) platforms [1][2]. The techniques can be applied to modeling large-scale complex systems, such as the Internet, cellular and mobile wireless networks, computer systems, transportation networks, mission-critical systems, and disaster-response systems. The second focus is to apply parallel simulation for real-time interactive simulation and emulation of large-scale computer networks [3]. A hybrid approach combining simulation, emulation, and physical systems [4][5] has been developed and has applied as part of the NSF's Global Environment for Network Innovations (GENI), a community-wide network experimentation platform for Future Internet research. The third focus is to develop high-performance multi-resolution models that can drastically improve the computational efficiency and scalability of simulation and emulation [6][7].

Cesar Marcondes research includes lines of research from congestion control [8][9], AI-enabled networks [10] and DTN [11], and particularly nowadays, focusing on SDN and NFV space. In the area of SDN, the research objective is to re-design important Internet problems like congestion control, multicast, and routing using a fresh new clean slate approach. For example, incorporating prefetching peer-to-peer dissemination algorithms directly in OpenFlow networks [12], or calculating in-advance Steiner-trees for all sources in multicast that allows fast switching of video sources [13], and in routing, incorporate a mix of content distribution and normal routing to optimize deliver, and also incorporating congestion signals directly in header fields driven by OpenFlow controllers. Another area of interest is on Experimental test beds like FIBRE or GENI, the research objective is to enable a scalable and reconfigurable backbone overlaid on top of NREN networks [14]. Finally, in the area of NFV, the research includes network control theoretical models to allocate on demand virtual network functions [14] and also proposals of a new system calls driven SDN and NFV architecture.

The collaborative efforts between Marcondes and Liu can be summarized in three main areas. The first area is a project, called Conflux, which aims to create a scalable and realistic symbiotic emulation and simulation test bed where computational resources are allocated to: a) simulate complex cross traffic rapidly in simulation, b) emulate advanced real applications based on Mininet, and c) provide an optimal combination of abstraction and realism by combining distributed simulation and emulation with a single API. The second area is on optimizing emulation between virtual machines and obtain greater throughput, by considering that emulation traffic is not content sensitive and therefore can be abstracted in the physical transmission. The third line of research is on an evolvable congestion control scheme that can be directly driven by the network using SDN, which provides both a distributed mechanism for measurement and control, and a centralized mechanism for sophisticated optimization algorithms. In addition, The congestion control scheme needs to be designed carefully as to benefit from the SDN mechanisms to achieve scalability and efficiency in order for controlling individual flow rates.

REFERENCES:

[1] LIU, J. ; RONG R. Hierarchical composite synchronization. In *Proceedings of the 26th Workshop on Principles of Advanced and Distributed Simulation (PADS 2012),* Zhangjiajie, China, July 2012, pp. 3-12.

[2] NICOL, D.M. ; LIU, J. ; LILJENSTAM, M. ; YAN, G. Simulation of large-scale networks using SSF. In *Proceedings of the 2003 Winter Simulation Conference (WSC'03),* New Orleans, LA, December 2003, pp. 650-657.

[3] LIU, Y. ; LIU, J. ; RANGASWAMI, R. Real-time network simulation support for scalable routing experiments. *International Journal of Simulation and Process Modelling, Special Issue on Parallel and Distributed Simulation*, 5(2):146-156, 2009.

[4] VAN VORST, N. ; ERAZO, M. ; LIU, J. PrimoGENI for hybrid network simulation and emulation experiments in GENI. *Journal of Simulation,* 6:179-192, 2012.

[5] ERAZO, M.A. ; LIU, J. Leveraging symbiotic relationship between simulation and emulation for scalable network experimentation. In *Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS)*, Montreal, Canada, May 2013, pp. 79-90.

[6] LIU, J. ; LIU, Y. Parallel hybrid network traffic models. *Simulation: Transactions of the Society for Modeling and Simulation International,* 85(4):271-286, 2009.

[7] LIU, J. ; LIU, Y. ; DU, Z. ; LIU, T. GPU-assisted hybrid network traffic model. In *Proceedings of the 2014 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (SIGSIM-PADS)*, Denver, CO, May 2014, pp. 63-74.

[8] CAINI, C ; FIRRINCIELI, R ; LACAMERA, D ; DECOLA, T ; MARCHESE, M ; MARCONDES, C. A. C. ; SANADIDI, M ; GERLA, M. Analysis of TCP live experiments on a real GEO satellite testbed? *Performance Evaluation Journal*, 10, 2008.

[9] ANDREW, L. ; FLOYD, S. ; MARCONDES, C. A. C. ; Lawrence Dunn ; GUILLIER, R. ; Gang, W. ; EGGERT, L. ; HA, S. ; RHEE, I. Towards a Common TCP Evauation Suite. In *Proceedings of the 6th International Workshop on Protocols for FAST Long-Distance Networks*, 2008, Manchester. PFLDNET 2008.

[10] MARCONDES, C. A. C. ; SHIMONISHI H. ; SANADIDI M. ; GERLA M. Combining TCP Westwood and Adaptive Reno: A Scalable and Safe Congestion Control Proposal. In *Proceedings of the IEEE International Conference on Communications* (ICC 2008).

[11] POMA, L. ; MARCONDES, C. A. C. ; HERMES S.; ARANTES L. Applying Machine Learning to Reduce Overhead in DTN Vehicular Networks. *XXXII SBRC* 2014.

[12] ASSUMPÇÃO J. ; SALVADOR, M.R.; ROTHENBERG, C.E. ; MARCONDES, C. A. C. Load Balance by Splitting a Flow and P2P Prefetching Algorithms. In *Proceedings of the P2P Networks Workshop, SBRC,* 2013.

[13] MARCONDES, Cesar; SANTOS, T., GODOY A.; VIEL, C. and TEIXEIRA C. CastFlow: Clean-slate multicast approach using in-advance path processing in programmable networks. ISCC 2012.

[14] ABELEM, A. ; STANTON, M. ; MACHADO, I. ; SALVADOR, M. ; MAGALHÃES, L. C. S. ; FERNANDES, N. ; CORREA, S. ; CARDOSO, K. V. ; MARCONDES, C. ; MARTINS, J. ; MONTEIRO, J. A. S. ; REZENDE, J.F. FIT@BR - a Future Internet Testbed in Brazil. In *Proceedings of the APAN (Asia Pacific Advanced Network) Network Research Workshop,* 2013.

[15] BAREA, E. ; MARCONDES, C. NFV Resource Scheduling using a PID Control Model. A MSc Proposal approved by UFSCar Computer Science Dept. 06/2014.

UCSC's Internetwork Research Group (i-NRG)'s Research Collaborations With Brazil

Katia Obraczka
Department of Computer Engineering
UC Santa Cruz, CA, USA
Email: katia@soe.ucsc.edu

I. Introduction

This paper outlines current and prospective collaborations between UCSC's Internetwork Research Group (i-NRG) and research groups in Brazil. Existing and potential collaborative efforts span different fields within computer neyworking including: software-defined networking (SDN), wireless sensor networks, disruption-tolerant networking (DTN), to name a few.

II. Software-Defined Networking

The idea of "programmable networks" has been proposed as a way to facilitate network evolution. In particular, Software Defined Networking (SDN) is a new networking paradigm in which the forwarding hardware is decoupled from control decisions. It promises to dramatically simplify network management and enable innovation and evolution. The main idea is to allow software developers to rely on network resources in the same easy manner as they do on storage and computing resources. In SDN, the network intelligence is logically centralized in software-based controllers (the control plane), and network devices become simple packet forwarding devices (the data plane) that can be programmed via an open interface. The field of software-defined networking is quite recent, yet growing at a very fast pace and attracting significant attention from both academia and industry.

I have been collaborating with Prof. Christian Esteve Rothemberg and his group at University of Campinas (UniCamp) on developing a distributed, logically decentralized SDN framework, called D-SDN, for hybrid network environments (e.g., [1], [2], [3], [4], [5], [6]). This work was part of Dr. Mateus Santos' PhD thesis during his internship in my lab in 2013-14. Dr. Santos is now a pos-doc researcher working with Dr. Rothemberg at UniCamp and participating in our collaborative effort.

III. Wireless Sensor Networks

Another notable research collaboration is a joint effort with Prof. Valmir Carneiro Barbosa (http://www.cos.ufrj.br/ valmir/), at the Federal University of Rio de Janeiro, who specializes in distributed and parallel algorithms. The project is related to Wireless Sensor Networks, an area of research on which I have done considerable work (see http://inrg.cse.ucsc.edu for pointers to our publications in the areas of sensor network protocols and energy models). More specifically, we will investigate distributed algorithms for optimal placement of sensor nodes on 2.5D terrain. We are targeting applications such as surveillance of outdoor areas using sensors that require line-of-sight (e.g., cameras). Deployment algorithms for wireless sensor networks aim to find node locations that optimize network characteristics such as energy usage, network lifetime, and sensor coverage. These algorithms often constrain node locations to a region of a plane (2D) or to a region of space (3D) with or without obstructions. Such 2D and 3D scenarios can model many real-world network deployments, but they may only approximate the characteristics of networks that are deployed over non-planar terrain (2.5D). Most sensors that do not physically touch a target must have a clear line of sight (CLOS) to sense their target's haracteristics. In 2.5D deployments, the CLOS requirement is a constraint. For example, sensors that detect visible and near-visible light, such as cameras and motion detectors, cannot see through or around terrain. Other sensors that require CLOS include those that detect ultrasonic acoustic waves or microwave electromagnetic waves. We plan to design, develop, and evaluate distributed algorithms for sensor networks whose nodes reside on 2.5D terrain. We will use centralized deployment algorithms as performance baseline.

IV. Efficient Routing and Congestion Control for Disruption Tolerant Networks

DTNs are networks where end-to-end connectivity cannot be guaranteed at all times. Due to the potential for intermittent connectivity, the resulting end-to-end delays may be arbitrarily large. DTNs find a variety of applications ranging from habitat and environmental monitoring, disaster recovery and special operations, surveillance, vehicular communications, to name a few.

I have also been collaborating with Prof. Celso Hirata (ITA), Prof. Marcos Nogueira (UFMG), and Prof. Marcelo Hilario (UFMG) on the topic of efficient congestion control for Disruption Tolerant Networks (DTNs). This is the topic of Aloizio Pereira da Silva's PhD work who has been doing an internship in my lab since June 2013 under the Brazilian Science without Borders Program. Our other collaborator is Dr. Scott Burleigh from NASA's Jet Propulsion Lab in Pasadena, California. To-date, we have published two papers on our work ([7], [8]); more recent results have been reported in two other papers that are under submission.

V. References

[1] B. Nunes, M. A. S. Santos, B. T. de Oliveira, K. Obraczka, T. Turletti, C. B. Margi, “Software-defined-networking-enabled capacity sharing in user-centric networks,” Communications Magazine, IEEE , vol.52, no.9, pp.28,36, September 2014.

[2] M. A. S. Santos, B. Nunes, B. T. de Oliveira, K. Obraczka, T. Turletti, C. B. Margi, “Decentralizing s Control Plane”, 39th Annual IEEE Conference on Local Computer Networks (LCN), 2014.

[3] B. Nunes, Marc Mendonca, Xuan-Nam Nguyen, K. Obraczka, Thierry Turletti,”A Survey of Software-Defined Networking: Past, Present, and Future of Programmable Networks”, IEEE Communications Surveys and Tutorials, vol.16, no.3, 3rd. quarter 2014.

[4] B. Nunes, Marc Mendonca, K. Obraczka, Thierry Turletti,”Software Defined Networking for Heterogeneous Networks”, IEEE COMMUNICATIONS SOCIETY MULTIMEDIA COMMUNICATIONS, IEEE COMSOC MMTC E-Letter, Vol. 8, No. 3, May 2013, pgs 36-39.

[5] Mateus A. S. Santos, B. Nunes, Cintia B. Margi, Bruno T. de Oliveira, K. Obraczka, Thierry Turletti,”Software-Defined Networking Based Capacity Sharing in Hybrid Networks”, Workshop on Capacity Sharing (CSWS’13), with the IEEE International Conference on Network Protocols (ICNP13), 2013.

[6] Marc Mendonca, Katia Obraczka, and Thierry Turletti, “The Case for Software-Defined Networking in Heterogeneous Networked Environments”, ACM CoNEXT Student Workshop, Nice, France, December 2012.

[7] Aloizio P. Silva, Scott Burleigh, Celso M. Hirata, and K. Obraczka,”A Survey on Congestion Control for Delay and disruption Tolerant Networks”, Ad Hoc Networks, Elsevier, August 2014.

[8] Aloizio P. Silva, Scott Burleigh, Celso M. Hirata, and K. Obraczka,”Congestion Control in Disruption-Tolerant Network: A Comparative Study for Interplanetary Networking Applications”, ACM CHANTS 2014, September 2014.

Position Paper Elaborated By Prof. Dr. Alessandro Anzaloni.

Instituto Tecnológico de Aeronaútica-São José dos Campos-SP/Brazil.

Title:Software Defined Networks applied to SMART GRIDS.

This position paper addresses some topics and issues that we believe that are important to the introduction and development of SDN into the Smart Grid.

Considering the incremental importance of electrical energy in ours lives, the electrical grids are evolving to the so called Smart Grids. In order to be Smart the electric grid needs a communication system that can guarantee reliability, efficiency, security and the possibility to introduce innovation. A deep view in Smart Grids reveals the necessity of two main control systems:

a) A Control System applied to the communication system that support the management/control of electrical substations, to give the required reliability to the Smart Grid.
b) A control system that provide a \`smart\`distribuition of the available energy between users in a limited area, like ´Condominios\`in Brazil.

Considering (a) our proposal is to :

i) Study the advantages and disadvantages of SDN in relation to the previous solutions.
ii) Particularly investigate the problem of failover of links and switches in SDN solutions for Smart Grid (for instance, evaluate the time necessary to recover from a link fail). Early work was down in this subject in [1] and [2].

The result of this studies will have strong impact in SDN equipment providers interested in Smart Grid Market.

Considering (b) we can raise two important issues:

i) Develop a Control System that allows the users (proconsumers) to sell/buy energy from their neighbor considering their own available energy (including solar energy, wind energy etc..) [3].

ii) Considering that interoperability between different equipment providers can be a good property the developed control system will use the functionalities offered by the standard IEC 61499 [4] considering the holonic architecture proposed in [3].

iii) An important activity will be the matching of the above functionalities to a SDN solution (the controller of the SDN network).

References.

[1] Thomas Pfeiffenberer and Jia Lei Du, `Evaluation of Software Dedined Networking for Power Sytems`. IEEE International Conference on Intelligent Energy and Power Systems(IEPS). IEEE 2014.

[2] Pedro Bittencourt Arruda, ´An Application using OpenFlow for Fault Tolerant Multicast in Energy Substations`. Graduation Project at ITA. Advisers: Alessandro Anzaloni and Thomas Pfeiffenberer.2014.

[3] Ebisa Negri et al. `Holonic Architecture of Smart Grid `.Smart Grid and and Renewable Energy- Scientific Research 2013.

[4] James H. Christensen et al. `The IEC 61499 Function Block Standard : Overview of the Second Edition`.ISA Automation Week 2012.

POSITION PAPER:
Advanced Peering with a Software-Defined Knowledge Plane

Christian Esteve Rothenberg

Information & Networking Technologies Research & Innovation Group (INTRIG)
School of Electrical and Computer Engineering (FEEC)
University of Campinas (UNICAMP), Campinas, Sao Paulo, Brazil

chesteve@dca.fee.unicamp.br

Abstract

This position paper presents our research vision on advanced peering scenarios between Software Defined Networking (SDN) domains by means of a so-called Knowledge Plane. Ongoing research efforts are devoted to the investigation of SDN-SDN communication options that allow SDN networks in different administrative domains to achieve advanced peering agreements beyond pure packet routing, for instance by integrating best placement of content and applications. Among the identified use cases, we are working on enabling distributed NFV (Network Functions Virtualization) offerings where SDN peering mechanisms allow resource trading (computation, storage, network) between different domains to support optimized VNF Infrastructures-as-a-Service models.

1 VISION

Our research project seeks to research, design and experimentally validate a Knowledge Plane approach, inspired by the seminal work by Clark *et al* [1], to Software Defined Networking (SDN) [2] to provide advanced peering capabilities between SDN controlled domains beyond pure packet routing (see Fig.1). The resulting architecture will be exercised in the context of Internet Exchange Points (IXP) scenarios where participants implement SDN technologies inside their domains and wish to collaborate beyond the peering capabilities of the traditional BGP routing protocol [3]. One target scenario we envision is to extend peering by integrating best placement of content, applications, and virtualized network processing functions (NFV - Network Functions Virtualization).

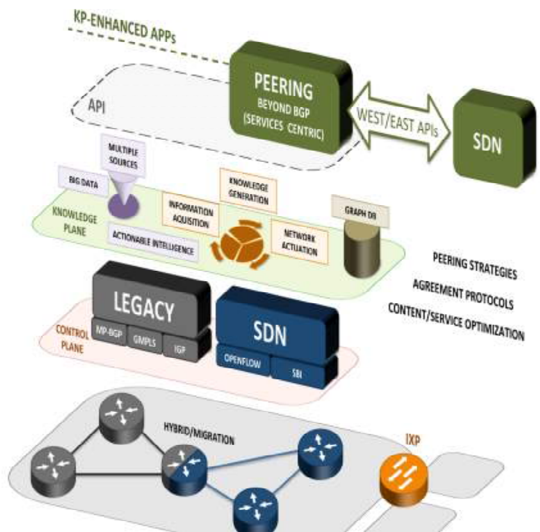

Figure 1: Big picture of SDN architecture with knowledge plain for advanced peering

2 RESEARCH CHALLENGES

There are a number of challenging research issues when aiming at interconnecting SDN islands and trying to conceive and introduce the notion of a (shared) Knowledge Plane. Security, economics and performance are candidate objectives at IXPs worth to explore under the possibilities of SDN and Big Data technologies.

In this context, among the main research questions is how to allow rich peering among SDN domains to differentiate by optimizing content/application placement between and by enabling new peering services to exchange resources beyond best-effort packet routing in a trading-like as a service model. One candidate set of resources to be exchange among SDN domains are the infrastructure capabilities to host virtual network functions (VNF) to effectively deliver VNF-as-a-service across peering domains.

This research scenario unfolds in a number of research questions such as how should SDN networks in different administrative domains communicate with the objective of enabling richer peering beyond plain packet routing? There are a number of open questions towards operating an Internet of SDN networks, including the rethinking of peering policies, the incentives to share "knowledge", and the underlying mechanisms that allow sharing per-domain abstractions in a usable and trustful way while allowing to keep and do not expose.

A core research challenge itself is the notion and representation of "knowledge" of a SDN domain, including possible embodiments in suitable data structures (e.g. graph-oriented databases) and adequate inter-domain proto cols towards the optimization of content and application placement.

ACKNOWLEDGEMENTS

This work is being supported by the Innovation Center, Ericsson Telecomunicações S.A., Brazil.

REFERENCES

[1] D. D. Clark, C. Partridge, J. C. Ramming, and J. T. Wroclawski. A knowledge plane for the internet. In Proceedings of SIGCOMM '03, pages 3–10, New York, NY, USA, 2003. ACM

[2] Diego Kreutz, Fernando M. V. Ramos, Paulo Verissimo, Christian Esteve Rothenberg, Siamak Azodolmolky, Steve Uhlig. "Software-Defined Networking: A Comprehensive Survey." In Proceedings of the IEEE, Vol. 103, Issue 1, 2015.

[3] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett. Sdx: A software defined internet exchange. In Proceedings of SIGCOMM '14, pages 579–580, New York, NY, USA, 2014. ACM.

LERIS Research Activities

Fábio Luciano Verdi
Department of Computing
Federal University of São Carlos
Sorocaba-SP
verdi@ufscar.br

***Abstract*—This paper presents the research activities currently done in the Networking, Innovation and Software Laboratory (LERIS). The main research fields include Software Defined Networking, Datacenters, Multipathing and Cloud.**

Keywords*—*Multipathing, SDN, datacenters.

I. Description of the research activities

This position paper is divided in four small sections each one representing the activities currently done in our research group. The first section presents the usage of Multipathing protocols together with OpenFlow for forwarding flows in the network. The second section shows a discussion on how to deal with elephant flows by splitting them into mouse flows in the end-hosts. The third section presents the research on cloud monitoring. Finally, in the last section we present our research on the integration of mathematical modelling with the deployment of SDN rules in a network.

A. Multipathing and OpenFlow Networks

Today's network research are focused on new routing techniques for improving end-to-end transfer [1], [2]. Also, transfer protocols have been redesigned for enabling multipath forwarding [3]–[5]. These techniques are used to improve the resilience and the network throughput. As an example, large WANs such as the Internet use best-effort routing to forward packets and traffic to the same destination tend to use the same route although there are multiple different paths to reach their endpoints. Another example is related to data centers that may use multiple pathways to forward data traffic inside its network. However, few protocols provide better alternatives to use different and disjointed paths to increase the throughput between servers [6]. Those issues motivate the use of network protocols to allow benefits of routing through multiple paths.

Recently, MPTCP was proposed as a TCP extention for multipath forwarding [7]. The main idea is to use a single TCP connection and separate it into many derivative flows, called subflows. A subflow is a conventional TCP connection derived from a main MPTCP connection. The network operator sets automated network rules to forward distinct subflows among each avaliable network interfaces card (NIC). When the host has multiple NICs and the subflows are forwarded through them, the end-to-end throughput is improved due to the sum of throughput. The idea is that each NIC be connected with different switch-ports. However, sometimes the host has only one NIC and the subflows are forwarded to the same path. We call such cases as shared bottlenecks. A shared bottleneck occurs when many subflows use a single network link, what implies in over-utilization of the bandwidth in these links. The problematic issue about MPTCP is when there is over-utilization of links due to the increase of subflows genereted by this protocol.

In this research topic our main proposal is to improve the throughput in shared bottlenecks by forwarding subflows from a same MPTCP connection through multiple paths. Our intention is to guarantee that the subflows of the same MPTCP connection are sent using different routes in the network. By doing this, the first advantage is to bring resilience to the original MPTCP connection since the subflows are spread in the network using disjointed paths. A second advantage is to increase the final throughput by distributing the subflows among several paths. This fine-grained controll of every MPTCP subflow is only possible because of OpenFlow that allows to have a total programability of the network elements. We develope Multiflow as a component for the POX controller. This component maps multiple subflows from the same MPTCP connection to multiple pathways.

B. Management of Elephant Flows using Multipathing Protocols

Data Centers need links with high bandwidth in order to meet the requirements necessary to establish their interconnections [8]. Several highly redundant network topologies are proposed by many researchers to meet the demand of existing services. But a traffic management is required to use effectively the division of bandwidth that is available for these topologies. Through these topologies pass many streams that carry a large amount of data some times causing network delays. These are called Elephant Flows and need to be detected and treated.

This field of research aims to develop an architecture capable of managing Elephant Flows in the end-host by monitoring the TCP flows, enabling better and more efficient visibility of the features of these flows. This management consists in monitoring the Elephant streams and split them into small streams in the network.

The idea is to use a module to perform end-host traffic management by detecting and treating Elephant Flows. We will use the MPTCP protocol to force the splitting of these flows. Then, when an Elephant Flow is detected, the module triggers the MPTCP protocol and the original TCP connection is divided into several smaller flows (mice) that are scattered over the network through multiple paths of the topology. We expect to obtain better results with this solution so that we minimize the impact of Elephant Flows in the network.

C. Cloud Monitoring for Improving Virtual Machine Migration

The monitoring of cloud computing environments mainly when the subject is related to the monitoring of the resources for Virtual Machines (VMs) migration is a key point to assure the availability of the services provided by any datacenter [9], [10]. By having a detailed information about how the environment as a whole is being used in terms of CPU, RAM, HD and network interfaces allows to better coordinate the VM migrations among physical servers. So, the idea of this research field is to develop a monitoring system by which the information about the current usage of the resources of the datacenter is obtained on-line and then helps to find the most appropriate physical server to receive any given VM. Also, the system must take into account the variation of the collected information so that ideas from the control theory may be used to add value before taking the decision to where a VM should move. In this project, we assume that the datacenter network is totally OpenFlow-enabled and then, the monitoring system must be integrated with the OpenFlow controller in order to prepare the network for sending the VM image from a machine to another considering that VM images are typically of huge sizes.

The whole architecture must have an strong relation with the resource monitoring in the cloud datacenter as well as with the SDN controller. The resource monitoring will depend on the IaaS being used in the datacenter and our solution must be integrated with such IaaS for collecting resources usage information.

D. Mathematical Modelling for SDNs

In this research we refresh that computational modelling is the key to describe, evaluate and analyse the most diverse computational problems before its prototyping. However, computational modelling for SDN is still superficially researched in the literature. Based on this, we are developing MILPFlow (**M**ixed **I**nteger **L**inear **P**rogramming with Open**Flow**)[1], a toolset for the integration of computational modelling and deployment of SDN networks using the OpenFlow protocol. Our approach is two-fold: MILP modelling and the deployment of data paths in the SDN network.

In the first step, the MILP modelling describes in higher level the SDN topology and its traffic matrix for the specified data center. In this matrix, we define the amount of traffic between every pair of source Top of Rack (ToR) and destination ToR.

We use these inputs to evaluate the feasibility for submitting traffic through the established routes without losses of packets in the network. So, we optimize the usage of physical servers and, at the same time, accommodate the traffic inside the data center topology without packet losses. Our goal is to obtain well distributed data traffic with high throughput in the whole topology without over-utilization of links.

In the second step, the deployment of data paths uses a proactive methodology where the data paths are set before the data is sent through the network. Due to the large amount of rules that are commonly necessary to establish data paths between servers of SDN topologies, we define a methodology for creation and management of these rules according to the mathematical modelling. These rules are used to set the data paths for the flows.

We are validating our approach by using Mininet emulator. MILPFlow is used to generate MILP models about the data center resources, map the solution in OpenFlow rules and deploy these rules in data paths of SDN networks. This whole mapping is important to simplify the management of OpenFlow rules by data center administrators for large SDN topologies.

[1] MILPFlow is published as open-source software at: https://github.com/milpflow/milpflow.

References

[1] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "OpenFlow: Enabling Innovation in Campus Networks."

[2] C. Raiciu, S. Barre, C. Pluntke, A. Greenhalgh, D. Wischik, and M. Handley., "Improving Datacenter Performance and Robustness with Multipath TCP," *In Proceedings of the ACM SIGCOMM 2011 conference (SIGCOMM '11)*, 2011.

[3] M. Arye, E. Nordstrom, R. Kiefer, J. Rexford, and M. Freedman, "A Formally-verified Migration Protocol for Mobile, Multi-homed Hosts," in *Network Protocols (ICNP), 2012 20th IEEE International Conference on*, Oct 2012, pp. 1–12.

[4] K.-J. Grinnemo, A. Brunstrom, and J. Cheng, "Using Concurrent Multipath Transfer to Improve the SCTP Startup Behavior for PSTN Signaling Traffic," in *Advanced Information Networking and Applications Workshops (WAINA), 2014 28th International Conference on*, May 2014, pp. 772–778.

[5] B. Arzani, A. Gurney, S. Cheng, R. Guerin, and B. T. Loo, "Impact of Path Characteristics and Scheduling Policies on MPTCP Performance," in *Advanced Information Networking and Applications Workshops (WAINA), 2014 28th International Conference on*, May 2014, pp. 743–748.

[6] L. Ming, A. Lukyanenko, S. Tarkoma, and A. Yla-Jaaski, "MPTCP Incast in Data Center Networks," *Communications, China*, vol. 11, no. 4, pp. 25–37, April 2014.

[7] A. Ford, C. Raiciu, M. Handley, and O. Bonaventure, "TCP Extensions for Multipath Operation with Multiple Addresses," *IETF RFC*, 2013.

[8] B. R. N. H. M. Al-Fares, S. Radhakrishnan and A. Vahdat, "Hedera: Dynamic flow scheduling for data center networks," *NSDI*, pp. 19 – 33, 2010.

[9] C. W. S. De Chaves, R. Uriarte, "Toward an architeture for monitoring private clouds," *IEEE Communications Magazine*, vol. 49, no. 12, pp. 130 – 137, December 2011.

[10] M. Z. L. R. Z. Zhang, L. Xiao, "Mvmotion: a metadata based virtual machine migration in cloud," *Journal Cluster Computing*, vol. 17, no. 2, pp. 441 – 452, June 2014.

HPC-NFV: Empowering Network Functions Virtualization

Hermes Senger

December 2nd, 2014

1 Introduction

Current networks are made of dedicated devices that execute a myriad of operations including packet forwarding, routing, NAT, firewall, policy enforcement (e.g., for QoS, security, etc), deep packet inspection, and many others. Traditionally, such networks implement a control plane that decides how to handle network traffic, and a data plane that forwards traffic according to the decisions made by the control plane. These two planes are bundled inside the networking devices, reducing flexibility and hindering innovation and evolution of the networking infrastructure.

Recently, the emerging concept of Software-Defined Networking (SDN) [McKeown et al. 2008] has broken the vertical integration of traditional networking, separating the role of network control logic (the control plane) and the role of traffic forwarding (the data plane). Such separation of the control and data planes turns network switches into simple forwarding devices while the control logic is implemented in a logically centralized controller. This separation simplifies policy enforcement and network (re)configuration and evolution, leveraging the new opportunities for innovation and the evolution of networking technologies. Under the SDN approach, the controller exercises direct control over the data plane that is implemented by the switches through well defined APIs such as OpenFlow [McKeown et al. 2008, ONF 2014].

More recently, another important concept emerged that is the Network Function Virtualization (NFV). Complementary to the promise of Software Defined Networks (SDN) to support automatic orchestration of networking resources, NFV aims at automating the deployment and control of network functions, that are moved from dedicated (and often expensive) appliances to platforms based on virtualized commodity servers. This concept has been shown prominent mainly for telecom operators, because of its potential to reduce operating costs, efforts, and time to deliver new technologies to the market. NFV paves the way to deploy the scalable and dynamic management of network services with high availability on the top of shared infrastructure of networking and computational resources with high performance [Martins et al. 2014].

2 High Performance NFV

Originally, the purpose of NFV was focused on the virtualization of a wide range of middleboxes including firewalls, NAT, load balancers, and other simple and less computationally intensive functions. The main rationale behind this choice comes from the observation that it is possible to chain elementary and virtualized functions to achieve high throughputs in the order of ten gigabits per second [Martins et al. 2014]. However, more complex network functions with higher computational cost can also be implemented in efficient manner on the top of both specialized hardware such as GPUs and MICs [Zhou et al. 2014a], or even on the top of commodity multicore servers [Zhou et al. 2014b]. Examples of more complex functions include traffic classification, intrusion detection and preventions. Most of these functions are internally implemented by algorithms such as neural networks, statistical, and machine learning classification, clustering, feature selection, support vector machines, and naive Bayesian Classifiers [MIT 2014]. In this work we will investigate how such functions can be implemented and orchestrated in an efficient, scalable, and flexible manner on computing infrastructures such as GPU, MIC, FPGA, and commodity virtualized servers such as on the cloud.

References

[Martins et al. 2014] Martins, J., Ahmed, M., Raiciu, C., Olteanu, V., Honda, M., Bifulco, R., and Huici, F. (2014). Clickos and the art of network function virtualization. In *11th USENIX Symposium on Networked Systems Design and Implementation (NSDI 14). Seattle, WA: USENIX Association*, pages 459–473.

[McKeown et al. 2008] McKeown, N., Anderson, T., Balakrishnan, H., Parulkar, G., Peterson, L., Rexford, J., Shenker, S., and Turner, J. (2008). Openflow: enabling innovation in campus networks. *ACM SIGCOMM Computer Communication Review*, 38(2):69–74.

[MIT 2014] MIT (2014). In *Website of the Lincln Lab. at MIT*. [Online] Available at: http://www.ll.mit.edu/mission/cybersec/softwaretools/tools.html.

[ONF 2014] ONF (2014). In *Open networking foundation*. [Online] Available at: https://www.opennetworking.org/.

[Qu and Prasanna 2014] Qu, Y. R. and Prasanna, V. K. (2014). Compact hash tables for high-performance traffic classification on multi-core processors. In *International Symposium on High Performance Computing (SBAC-PAD 2014)*, pages 1–8.

[Zhou et al. 2014] Zhou, S., Nittoor, P. R., and Prasanna, V. K. (2014). High-performance traffic classification on gpu. In *International Symposium on High Performance Computing (SBAC-PAD 2014)*, pages 1–8.

Exploring sensor and Intelligent Vehicles data storage and control in Cloud Computing for Cyber-Physical Systems

Alex Sandro Roschildt Pinto† and Kalinka R L J Castelo Branco*

*Institute of Mathematics and Computer Science, University of São Paulo São Carlos-SP, Brazil
Email: kalinka@icmc.usp.br

†Department of Control Engineering and Automation, Federal University of Santa Catarina
Blumenau-SC, Brazil a.r.pinto@ufsc.br

The Internet of Things is a new paradigm where smart embedded devices and systems are connected to the Internet. In this context, Wireless Sensor Networks (WSNs) are becoming an important alternative for sensing and actuating critical applications like industrial automation, remote patient monitoring and domotics. Moreover, intelligent vehicles (i.e. Unmanned Aerial Vehicles, Unmanned Ground Vehicles) can be used to monitor and help users to detect events.

Over the last years, aerial and ground vehicles as well as mobile robot systems have been receiving an increased number of electronic components, connected through wireless networks and running embedded software. This strong integration between dedicated computing devices, the physical environment and networking, composes a Cyber-Physical System (CPS).

Wireless Sensor Networks (WSNs) are composed by tiny nodes that can sense the environment and communicate through wireless media [1-3]. Several applications can use these technologies like: smart homes, health care, environmental monitoring, military and security [4].

Furthermore, as processing power increases and software becomes more sophisticated, these vehicles gain the ability to perform complex operations, becoming more autonomous, efficient, adaptable, comfortable, safe and usable. These are known as Intelligent Vehicles (IVs) [5-6].

CPS is a novel trend in WSN and IVs technologies that combine sensing, actuating and mobility in order to assist people to monitor and actuate their facilities. Thus, CPS can significantly improve quality of life of millions of people that can remote monitor their situation to achieve a better health condition.

However, there are several challenges that CPS imposes which should be faced in order to implement these kind of applications like:

- Large amount of sense data: CPS devices produce lots of sensor data, which must be collected and stored. This way, Cloud Computing can help CPS applications to store data from different sources and locations.
- Communication Optimization: the communication between sensors-gateway, IVs-Cloud, Cloud-IVs, sensors-cloud and cloud-sensors must be optimized in order to reduce the amount of data.
- Real-time requirements: sensed data must be delivered to users before its deadline. On the other hand, old data can be eliminated and;
- Information Diffusion: after the data had been collected and stored, important information must be detected and distributed to users.

In this context, the main challenge of this project is to explore Cloud facilities to ease data diffusion, communication and data storage in CPS applications. The second challenge is to optimize data communication of sensors and IVs based on data fusion techniques. Finally, Cloud Computing will help to integrate several kinds of sensors and IVs based on different sensing and communication technologies.

Figure 1 shown the main idea of this project challenge.

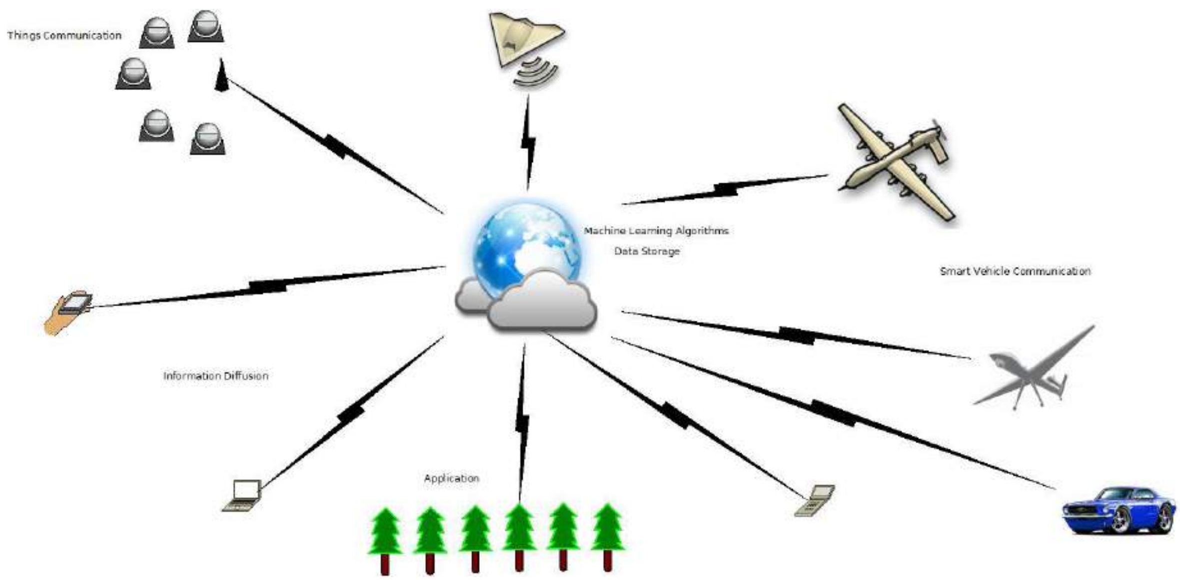

Figure 1: Main idea project challenge.

References

[1] A.R. Pinto, C. Montez, G. Araújo, F. Vasques, P. Portugal, An approach to implement data fusion techniques in wireless sensor networks using genetic machine learning algorithms, Information Fusion, Volume 15, January 2014, Pages 90-101.

[2] Schrickte, L. F., et al. "Integration of Wireless Sensor Networks to the Internet of Things Using a 6LoWPAN Gateway." Computing Systems Engineering (SBESC), 2013 III Brazilian Symposium on. IEEE, 2013.

[3] Lavratti, Felipe, et al. "A transmission power self-optimization technique for wireless sensor networks." ISRN Communications and Networking 2012 (2012): 1.

[4] A. Chaves, R. La Scalea, A. Colturato, C. Kawabata, E. Furtado, and K. Branco, “Using uavs and digital image processing to quantify areas of soil and vegetation,” in Journal of Physics: Conference Series, vol. 633, no. 1. IOP Publishing, 2015, pp. 12 112–12 115.

[5] Branco, Kalinka RLJC, Jorge Marques Pelizzoni, Luciano Oliveira Neris, Onofre Trindade Jr, Fernando Santos Osório, and Denis Fernando Wolf. "Tiriba-a new approach of uav based on model driven development and multiprocessors." In Robotics and Automation (ICRA), 2011 IEEE International Conference on, pp. 1-4. IEEE, 2011.

[6] Fernandes, Leandro C., Jefferson R. Souza, Gustavo Pessin, Patrick Y. Shinzato, Daniel Sales, Caio Mendes, Marcos Prado et al. "CaRINA intelligent robotic car: Architectural design and applications." Journal of Systems Architecture 60, no. 4 (2014): 372-392.

A FRAMEWORK FOR METEOROLOGICAL DATA MANAGEMENT AND ATMOSPHERIC MODEL EXECUTION

Eugênio Sper de Almeida
Ivo Koga
National Institute for Space Research (INPE)
Center for Weather Forecasting and Climate Research (CPTEC)
Rodovia Presidente Dutra, Km 40 - SP - 12630-000 - Cachoeira Paulista - SP - Brazil
{eugenio.almeida, ivo.koga}@cptec.inpe.br

Abstract

The research in meteorology requires High Performance Computing (HPC) resources, meteorological data and the configuration of a wide variety of parameters to execute meteorological models. This work has the objective of investigating meteorological data management methods to support research in meteorology, allowing the visualization and the manipulation of this data on the Web. It will address the specification and the development of a computational framework to combine data and atmospheric model execution. The results also include the creation of mechanisms to access historical meteorological data and the improvement of the BRAMS execution portal.

Keywords: High Performance Computing (HPC), eScience, data management, HPC Cluster, numerical atmospheric models, BRAMS.

1. Introduction

Modern atmospheric numerical weather prediction (weather, climate and environmental) uses atmospheric numerical models, complex programs that represent the movement and physical processes of the atmosphere through mathematical equations (Pielke, 2013). They require different types of data including observational data, vegetation maps, data derived from satellite images and outputs from low-resolution atmospheric models.

The atmospheric numerical simulation has a high computational cost due to the large amount of data processed by equations embedded within the models. Each simulation generates files (matrices) representing the weather forecast time, physical variables and atmospheric levels.

The regional models operate similarly, using the global models outputs, and produce outputs in a higher spatial resolution and lower geographical coverage. The weather forecast reliability depends on processes parameterization (physical and chemical) and spatial-temporal resolutions. Once computation is complex, they demand increasing HPC computing power.

Conducting research in meteorology is not a trivial task. It requires customization of physical parameters, compilation and execution of models in HPC clusters through command line, handling a large amount of data, and visualization of results through scientific visualization systems. To perform these tasks it is required a specialized knowledge, causing difficulties to conduct the experiments.

In order to facilitate the execution of atmospheric models a common approach encapsulates the experiments execution in meteorological workflows (Asvija et al., 2010), with access through scientific gateways. So, allowing researchers to keep the focus on their research and reduce the time spent on experiments.

Usually meteorological data are difficult to locate and most of the time inaccessible to external users, since they are dispersed within the storage systems of the institutions. Additionally, researchers still faces technical difficulties related to the heterogeneity of formats and access protocols.

Creating facilities for data access and execution of experiments is a way to overcome this challenge and enable the creation of a collaborative community of scientists from different fields to generate a high quality forecast.

A collaborative infrastructure for data access also enables the integration with other infrastructures providing a more accurate forecast of global events. Thus, the combination of models and tools, available data and scientists with appropriate knowledge will more accurately predict disasters (Shukla et al., 2010).

The Numerical Weather Prediction Center and Climate Studies (CPTEC) of the National Institute of Spatial Research (INPE), has the technological and scientific knowledge to develop and operate atmospheric numerical models. One is the Brazilian Regional Atmospheric Modeling System (BRAMS), the result of a research project to produce a tuned version of the Regional Atmospheric Modeling System model (RAMS) (Tremback and Walko, 1997) for the tropics.

This work will provide an infrastructure necessary for the provision and handling of this meteorological data on the Web. Once it is developed, an unique database of INPE/CPTEC data for research will be accessible contributing to the data democratization. Thus, an increasing number of users of atmospheric numerical models can use the infrastructure and data to their experiments.

2. Related Work

Lawrence et al. (2014) verified that scientific gateways perform an important role in research and education process, since scientists generally rely on Web applications to access data and use computer resources. Wilkins-Diehr (2011) points out the importance of gateways for applications execution in HPC clusters and grids.

Scientific gateways have expanded to meet the demands of high-performance distributed applications. In the US, they support users of atmospheric sciences, earthquakes mitigation, bioinformatics, drug discovery, biophysics, computational chemistry, social computing, earth sciences and biology applications (https://portal.xsede.org/science-gateways). Similar solutions exist in Europe and Asia.

The atmospheric model Weather Research and Forecasting (WRF) is widely used for regional forecast (Done et al., 2004). Scientific gateways were developed to simplify its use: LEAD (Christie and Marru, 2007), WRF and WRF4SG (Blanco et al, 2013; Smith et al., 2007). The Distributed Research Infrastructure for Hydro-Meteorology project (DRIHM) has developed an infrastructure of e-science to provide hydro-meteorological services aimed at the scientific community of atmospheric and earth sciences (Dagostino et al., 2014).

We observed that the data are generally of low spatial resolution or high resolution restricted to a particular geographic area. Thus, to the best of our knowledge, they handle less data volume than we propose to deal in our infrastructure.

Our experience in this area originated with the G-BRAMS project (Grid for the BRAMS), with the objective of generating a 10 years climatology of BRAMS, using computational grids (Souto et al., 2007). Later, we deployed a computational grid and developed a scientific gateway to predict the air quality (Almeida et al., 2008) within the project South American Emissions, Mega-cities and Climate (SAEMC) project (SAEMC, 2015).

In these previous projects we observed the impact of the data transfer time on a computational grid, related to the meteorological input/output data size, and the communication networks bandwidth. As the data volume increases, currently coming to petabytes, it is more efficient to send the user's query to the data than send data to the user (Hey et al., 2009). Thus, one of our objectives is to reduce data transfer and process data locally.

Few works deal with computing environment to execute atmospheric numerical models. Studies are limited to a particular set of meteorological data for specific use cases.

3. Environment for Atmospheric Numerical Simulation (ENAS) framework

Our previous efforts focused on the development of a scientific gateway to execute the BRAMS simulations. This scientific gateway facilitates and simplifies the execution of simulations with the BRAMS model and minimizes setup difficulties, area definition, access to computing resources, and use of scientific data visualization systems.

The scientific gateway has functionalities related to the application and infrastructure. The infrastructure-related functionalities deal with security, submission mechanisms, and tasks control/monitoring. The application-related functionalities include the presentation and handling of the information contained in RAMSIN file (parameters for the BRAMS execution), execution of the meteorological workflow and visualization of meteorological maps produced at the end.

One of our current efforts includes the specification and implementation of a generic framework called Environment for Numerical Atmospheric Simulation (ENAS) to allow the execution of weather forecast models; the development of algorithms for combined data query, enabling integration of different sources with different spatial and temporal scales; the redesign and implementation of scientific gateway to enable implementation of meteorological models and access to different data types on the web; and to make experiments data available through efficient data transfer mechanisms.

The development of ENAS will enable users to execute an atmospheric model or determine if there are data available. Once the period is defined, the user will be informed of the availability of such data. The ENAS initially will include the implementation and consultation of the BRAMS data, allowing the model parameterization and execution, and providing the results to users of the scientific gateway.

This project will started with the BRAMS model that requires: NDVI generated from MODIS satellite images, soil moisture, surface temperature of the sea (climatological or weekly), soil texture classes, land use and topography, data from Atmospheric General Circulation Model (AGCM) from INPE/CPTEC (Bonatti, 1996). The data organization will be the starting point, including mechanisms to keep updated the meteorological data for the experiments (historical and current data).

Another initiative that integrates the meteorological collaboration context is the Science DMZ [10], a network concept developed by ESnet engineers tailored to high-performance scientific applications. It is a portion of network focused in optimizing scientific data transfers including high-volume data transfer, remote experiment control, and data visualization.

4. Final considerations

Research in meteorology using atmospheric numerical models is not a trivial task, since it deals with High Performance Computing (HPC) environments, data management, data visualization and interaction with computing resources. An important contribution of this work is to bridge the gap between users and researchers, enabling users to easily access high-performance simulation environments.

The proposed framework will decrease manual data manipulation and include mechanisms to facilitate the use of atmospheric models: parameterization definition, model execution monitoring, output visualization and experiment data download. We will start with the BRAMS, but the framework will be generic and we have plans to include atmospheric numerical models.

References

[1] Almeida, E. S., Delgado, R., Rodrigues, L. F., Baeza, C., Coura, L. H. R., Gallardo, L., Longo, K., and Freitas, S. (2008). SAEMC_GRID: South america megacities emissions and climate grid. In Proc. of Latin American Grid (LAGrid) workshop, Campo Grande, MS, Brazil.

[2] Asvija, B., Shamjith, K., Sridharan, R., and Chattopadhyay, S. (2010). Provisioning the mm5 meteorological model as grid scientific workflow. In Intelligent Networking and Collaborative Systems (INCOS), 2010 2nd International Conference on, pages 310–314.

[3] Blanco, C., Cofino, A., and Fernandez-Quiruelas, V. (2013). Wrf4sg: A scientific gateway for the weather research and forecasting model. In Information Communication Technology Electronics Microelectronics (MIPRO), 2013 36th International Convention on, pages 172–176.

[4] Christie, M. and Marru, S. (2007). The lead portal: a teragrid gateway and application service architecture. Concurrency and Computation: Practice and Experience, 19(6):767–781.

[5] Dagostino, D., Clematis, A., Galizia, A., Quarati, A., Danovaro, E., Roverelli, L., Zereik, G., Kranzlmuller, D., Schiffers, M., Gentschen Felde, N., Straube, C., Caumontz, O., Richard, E., Garrote, L., Harpham, Q., Jagers, H., Dimitrijevic, V., Dekic, L., Fiorizz, E., Delogu, F., and Parodi, A. (2014). The drihm project: A flexible approach to integrate hpc, grid and cloud resources for hydro-meteorological research. In High Performance Computing, Networking, Storage and Analysis, SC14: International Conference for, pages 536–546.

[6] Done, J., Davis, C. A., and Weisman, M. (2004). The next generation of nwp: explicit forecasts of convection using the Weather Research and Forecasting (WRF) model. Atmospheric Science Letters, 5(6):110–117.

[7] Hey, A. J. G., Tansley, S., and Tolle, K. M. (2009). The Fourth Paradigm: Data-Intensive Scientific Discovery. Microsoft Research.

[8] Lawrence, K., Wilkins-Diehr, N., Wernert, J., Pierce, M., Zentner, M., and Marru, S. (2014). Who cares about science gateways? a large-scale survey of community use and needs. In Gateway Computing Environments Workshop (GCE), 2014 9th, pages 1–4.

[9] Pielke, Sr, R. A. (2013). Mesoscale meteorological modeling. Academic press.

[10] SAEMC (2015). South American Emissions Mega–cities and Climate project – SAEMC & ADAPTE: Higlights. http://saemc.cmm.uchile.cl. Accessado em 2015-05-21.

[11] Smith, J. S., Govett, M. W., & McCaslin, P. T. (2007). WRF Portal and Domain Wizard: Graphical tools for WRF. 23rd Conf. on IIPS, San Antonio, TX. Amer. Met. Soc.

[12] Shukla, J., Palmer, T. N., Hagedorn, R., Hoskins, B., Kinter, J., Marotzke, J., Miller, M., and Slingo, J. (2010). Toward a new generation of world climate research and computingf acilities. Bulletin of the American Meteorological Society, 91:1407–1412.

[13] Souto, R. P., Ávila, R. B., Navaux, P. O. A., Py, M. X., Maillard, N., Diverio, T. A., Velho, H. F. C., Stephany, S., Preto, A. J., Panetta, J., Rodrigues, E. R., and Almeida, E. S. (2007). Processing mesoscale climatology in a grid environment. In Cluster Computing and the Grid, 2007. CCGRID 2007. Seventh IEEE International Symposium on, pages 363–370. IEEE.

[14] Tremback, C. J. and Walko, R. L. (1997). The regional atmospheric modeling system (rams): development for parallel processing computer architectures. In 3rd RAMS Users Workshop.

[15] Wilkins-Diehr, N. (2011). A history of the teragrid science gateway program: A personal view. In Proceedings of the 2011 ACM Workshop on Gateway Computing Environments, GCE ’11, pages 1–12, New York, NY, USA. ACM.

Cloud computing with distributed data centers

Luiz F. Bittencourt

Institute of Computing, University of Campinas, Brazil
bit@ic.unicamp.br

Abstract. Cloud computing has been consolidated as a platform to process and store data in large scale data centers. This configuration can fulfill many applications requirements, resulting in scalable computation with lower cost than in-house processing/storage. Thus, cloud data centers are able to provide the necessary quality requirements for applications that do not need low latency. On the other hand, applications with low latency requirements can suffer of degradation when using cloud computing to process or store data. To contour this problem, scattered data centers can be deployed to reduce access latency, similarly to content delivery networks (CDNs). The architecture of such data center distribution has influence on resource management and networking requirements. This position paper discusses some options and aspects involved when building such an infrastructure to bring data centers to the edges of the network, closer to the end users.

1 Cloud computing at the edges

The consolidation of the cloud computing paradigm brought a diversity of online applications to different kinds of devices. Data and processing offloading is nowadays very common in connected devices, providing higher processing and storage capacity as well as ubiquitous access to data. As a consequence, users became more and more dependent of connectivity to properly run their (often mobile) applications. Moreover, the number of devices used to access cloud services is in expansion as a consequence of the aforementioned applications but also the increasing interest in Internet of Things (IoT). The amount of devices accessing cloud services is expected to reach 24 billion by 2020 [1].

The predicted scenario is capable of generating unprecedented amount of data in the network edges, urging solutions for data transport and management. More than potentially higher latencies, centralized cloud data centers can also have impact in the internet service providers networks when traffic demand increases. This is the expected scenario as a consequence of the increasing amount of devices implementing data/processing offloading, potentially leading to unpractical or inefficient data transport and processing.

Devices in the near future scenario can generate a lot of data that do not need to be delivered to cloud data centers to be stored in its raw form. Instead, these data can be processed at the edges and aggregated to reduce the amount of data transferred to the cloud. Considering that many devices – ranging from small fixed sensors to mobile, smart devices – have low processing capacity and/or battery constraints, they need to offload data and processing. A distributed cloud data center can help in reducing turnaround time for offloaded processing, and also contribute to reduce traffic among network edge and clouds data centers.

Architectural and technical aspects of a distributed cloud deployment presents interesting challenges. The main objective is to provide seamless access for mobile or fixed devices maintaining the level of quality according to the applications requirements. Some challenges are currently being addressed by a number of emerging technologies and approaches – such as Fog Computing [2], C-RAN [3], and mobile clouds [4], that attempt to move data and processing closer to the user, thereby moving cloud provisioning from centralized data centers to edge servers [5].

With distributed data centers, devices can extend their computing capacity using servers that are a single hop away, instead of accessing potentially distant centralized cloud data centers. These servers can process/store offloaded application data, and then forward the relevant data/processing to be done in the centralized data centers. This decision-making depends on several factors, among which application delay requirements play a central role. If we consider mobile users with interactive applications, one important challenge is to make offloaded user's data and processing to follow the user, i.e., migrate user's relevant data/processing across servers in the distributed data center to keep them close to where they will be needed. This can involve mobility prediction and techniques to identify relevant content to be migrated according to the (predicted) user mobility. This data movement can be supported by networking technologies, such as SDN and/or network virtualization, to provide time guarantees when a migration is necessary. Virtualization, which enables sharing of infrastructure amongst users with software and, potentially, hardware isolation, can have a central role in both data migration and network guarantees in this context. Another challenging network aspect is to define how distributed data center servers communicate: dedicated links or through the core network.

References

1. J. Gubbi, R. Buyya, S. Marusic, and M. Palaniswami, "Internet of things (iot): A vision, architectural elements, and future directions," *Future Generation Computer Systems*, vol. 29, no. 7, pp. 1645 – 1660, 2013.
2. F. Bonomi, R. Milito, J. Zhu, and S. Addepalli, "Fog computing and its role in the internet of things," in *MCC workshop on mobile cloud computing.* ACM, 2012, pp. 13–16.
3. A. Checko, H. Christiansen, Y. Yan, L. Scolari, G. Kardaras, M. Berger, and L. Dittmann, "Cloud ran for mobile networks: A technology overview," *Communications Surveys Tutorials, IEEE*, vol. 17, no. 1, pp. 405–426, Firstquarter 2015.
4. N. Fernando, S. W. Loke, and W. Rahayu, "Mobile cloud computing: A survey," *Future Generation Computer Systems*, vol. 29, no. 1, pp. 84 – 106, 2013.
5. I. Petri, J. Diaz-Montes, O. Rana, Y. Rezgui, M. Parashar, and L. Bittencourt, "Coordinating data analysis & management in multi-layered clouds," in *EAI International Conference on Cloud, Networking for IoT systems*, 2015.

Brown University Systems and Networking Lab

Planck

Software-defined networking replaces the distributed, per-switch control planes of traditional networks with a (logically) centralized control plane that programs the forwarding behavior of all the switches in a given network. This single centralized control plane, executed on a controller, can act as a control loop that (i) gathers traffic information and other measurements from the network and (ii) uses the gathered information to compute and install modified forwarding behavior in switches. Although there are two logical components to this control loop---measurement and control---the focus of the vast majority of SDN research and development has focused on control.

OpenFlow, the dominant protocol used to implement SDN, relies on two measurement techniques to provide a global view of the network: packet_in messages and per-port/per-rule counters. In practice, neither of these measurement mechanisms enable a scalable and low-latency measurement system because they update on the order of seconds. This limits the control loop of a controller to operating at a speed that is too slow to detect anything but the absolute largest flows.

In collaboration with IBM Research we have been working on a system called Planck [1] which leverages a novel mechanism for packet sampling that can be used in switches today. By dedicating one port per switch to monitoring we are able to obtain dynamic sampling rates, depending on traffic, between 1:1 and 1:N packets (where N is the number of ports on a switch). These rates are drastically higher than what is presently available by methods such as sFlow. By collecting these samples across the entire network we are able to get a global view of the network's congestion points in almost real-time. We are then able to use this information for traffic engineering or network policy conformance detection purposes to strategically place flows on the network to increase overall throughput or fairness in the network.

Participatory Networking (PANE)

Network configuration and management today is a complex and error-prone task, done by expert administrators, who implement policies as a brittle composition of independent control rules on heterogeneous devices. Networks become stifled, as changes to configuration have hard-to-predict consequences in the resulting policies. At the same time, there is pressure in the opposite direction for more dynamic networks, as increasingly distributed applications can benefit from network flexibility, and actually have information about the semantics of their traffic that can inform network configuration.

The goal of this work, PANE, is to enable end-user applications to take active part in the configuration, provisioning, and management of the network, and, in turn, have more visibility into the network state to better reconfigure themselves [2]. Our approach aims to be a unifying and extensible framework to expose network control mechanisms and state to end user applications. It leverages Software Defined Networking, with its logically centralized and programmable control plane, and, in an analogy to Operating Systems, provides the "user level system calls to the network”. Applications request current and future service qualities, gain visibility into network properties, and provide hints about future demands to the network.

FlowLog

High-level SDN programming languages can ease controller implementation. Our Flowlog language [3] takes a new approach to SDN programming: we call Flowlog "tierless" because it abstracts out both data storage and OpenFlow rule generation, allowing programmers to write their entire (stateful) program as if every packet passes through the controller. Flowlog's underlying runtime proactively installs switch rules to match the current controller state and automatically ensures consistency. This eliminates some common bugs in controller/switch interaction, simplifies the programming experience, and also enables automated program verification.

Flowlog allows programmers to verify program-behavior properties, but it can also perform differential program analysis: given two versions of a controller program, Flowlog can generate a set of concrete scenarios that fully illustrate how the programs' behavior differs. This approach allows network developers to exploit the power of formal methods tools without having to be trained in formal logic or property elicitation.

Simon

It can be frustrating to understand, test, and debug SDN controller behavior, even in a relatively simple environment such as Mininet. For instance, a test that focuses only on connectivity, such as testing via ping, can disguise efficiency problems such as unnecessary flooding. Static checking tools, while powerful, are limited in scope. Operators need tools that can inspect more than just high-level program behavior or low-level forwarding rules, and which can help determine whether the network respects their goals. Our SIMON (Scriptable Interactive Monitoring) tool [4], an interactive debugger for SDN controller programs, is a next step in that direction. SIMON is independent of the controller platform used, and does not require code annotations or knowledge of the controller software being run. Operators may compose debugging scripts both offline and interactively at SIMON’s debugging prompt, taking advantage of the rich set of reactive functions SIMON provides as well as the full power of Scala.

[1] “Planck: Millisecond-scale Monitoring and Control for Commodity Networks” Jeff Rasley, Brent Stephens, Colin Dixon, Eric Rozner, Wes Felter, Kanak Agarwal, John Carter, Rodrigo Fonseca. SIGCOMM 2014

[2] “Participatory Networking: An API for Application Control of SDNs” Andrew D. Ferguson, Arjun Guha, Chen Liang, Rodrigo Fonseca, and Shriram Krishnamurthi. SIGCOMM 2013

[3] “Tierless Programming and Reasoning for Software-Defined Networks” Tim Nelson, Andrew D. Ferguson, Michael J. G. Scheer and Shriram Krishnamurthi. Networked Systems Design and Implementation (NSDI) 2014

[4] “Simon: Scriptable Interactive Monitoring for SDNs” Tim Nelson, Da Yu, Yiming Li, Rodrigo Fonseca, and Shriram Krishnamurthi. Symposium on SDN Research (SOSR) 2015

Communication, Security and Management using Software-Defined Wireless Sensor Networks

Cíntia Borges Margi
Escola Politécnica da Universidade de São Paulo
Email: cintia@usp.br

***Abstract*—Wireless sensor networks are a valuable technology for the development of applications in several areas, such as environmental, industrial and urban monitoring, health and agriculture. Given the challenges of different protocols and technologies on the link layer, resource constrained devices and security requirements, current research problems address secure communication between wireless sensor networks devices and network resource management. The main goal in my current work is the design and validation of a framework to provide basic communication services, security and resource management for heterogeneous wireless networks, as well as wireless sensor networks, and considering the software-defined network paradigm in order to provide a flexible solution, which facilitates the adaptation of wireless heterogeneous networks to potential topology changes and reconfiguration to perform new tasks, ensuring the associated security services.**

I. Introduction

Wireless Sensor Networks (WSN) have been used to support several different applications, mainly related to monitoring, detection and tracking. Nodes in a WSN are typically battery-powered and resource constrained (i.e. limited amount of memory, processing and communication), and communicate through a multihop ad hoc network [1]. There is not a unique definition for the Internet of Things (IoT), but most of them agree that it is composed of devices capable of sensing/actuation, communication and processing [2]. Thus, WSN became a key technology for IoT and research on WSN should provide good basis to develop IoT.

Software-Defined Networking (SDN) has been envisioned as a way to reduce the complexity of network configuration and management, initially focused on wired networks. The main approach to SDN is OpenFlow [3], which focus on wired networks.

SDN for WSN and IoT imposes different challenges and requirements, but provides several opportunities. Among the challenges we highlight the limited resources: energy, processing, memory and communication. Requirements are related to the applications characteristics (e.g. data frequency and size), as well to the nodes behavior due to duty-cycles, operating systems and programming approach. On the other hand, opportunities provided by SDN include: to improve resource reuse, to implement node retasking, node and network management, as well as to enable experiments with new protocols and to ease transition to standard protocols for deployed networks.

II. Related Work

The literature presents two main approaches to SDN applied to WSN and IoT.

First approach is to consider the IoT/sensor nodes as non-SDN components, which must rely on other SDN-capable devices [4], [5]. Flauzac et al. propose an SDN based architecture for IoT to improve security [4]. Authors assume that IoT or WSN devices cannot have embedded SDN capacities, but that they will be in the neighborhood of an SDN device. Tadinada [5] describes their SDN implementation and use cases. One of the use cases focuses on IoT, and proposes the use of an IoT gateway based on their SDN switch.

The second approach to SDN applied to WSN and IoT considers that nodes are SDN capable. Flow-Sensor [6] proposes sensor nodes with the main features of OpenFlow [3], aiming to obtain more reliable nodes in comparison with typical sensor nodes since packets and sensor nodes can be easily monitored and reconfigured, enabling a dynamical management from a perspective of cost, network performance and energy-effectiveness. Sensor OpenFlow [7] and SDWN [8] propose a clear separation between data plane and control plane, a centrally controlled communication protocol between these two planes and some data plane features; enabling dynamic sensor retasking, promoting sharing of implemented functions and hardware resources between different deployed applications. Gante et al. [9] propose a framework to apply SDN to WSN management. Besides the benefits highlighted by previous works, the authors include accurate localization and topology discovery as advantages of SDN usage in WSN.

TinySDN [10] addresses common WSN characteristics, such as possible disruption and delay in communication, low energy supply, reduced data frame length and in-band control by enabling multiple controllers on the software-defined wireless sensor networks, improving on previous work. Furthermore, the authors present the protocol for communication between controllers and WSN nodes, and discuss implementation issues and approaches for TinyOS [11].

III. Challenges and Trends

WSN and the IoT are composed of devices capable of sensing/actuation, communication and processing, which could improve our daily life, but standardization has been considered one of the main challenges to their deployment. Given the SDN centralized control characteristics and match bitmasking, it is flexible enough to adapt compatible wireless sensor networks (e.g., creating gateways nodes) to communicate with older protocols. Thus, TinySDN could integrate new standardized WSNs and non-standardized/legacy WSNs, assisting in the transition to 6LoWPan and RPL standards.

An interesting feature that could be achieved through SDN enabled devices is node and resource management. For instance, when the controller determines a route to be used, it could consider the energy available in a given node (or set of nodes) to determine which route will provide the best network lifetime. Also, the controller could determine when a given sensor node fails, since no further messages are received. Furthermore, WSN nodes are usually considered disposable and cheap devices, which could be deployed for a specific task. But consider smart cities, where sensor nodes should collect, process and transmit different types of data for different applications. If these sensor nodes and other devices collecting data could be managed by the SDN paradigm, one could achieve a much better usage of the underlying infrastructure through dynamic node retasking and routing. In addition, one WSN node could collect data and forward it to different sink nodes, according to the requirements of the different applications running over the WSN shared infrastructure.

Another important issue that could be addressed through SDN is security [4]. Since the SDN controller knows the network topology, this information could be used to support node admission decisions, using and approach similar to the capacity sharing proposal [12]. Once nodes are admitted to an SDN domain, a key distribution mechanism can be used to support security services, such as confidentiality and authenticity.

IV. CURRENT RESEARCH

In our current research, we leverage from the SDN paradigm to reduce complexity in WSN configuration, management and re-tasking. The SDN controller, through its Management plane, has a local representation of the network state including nodes, links and existing associated information. This information enables resource management, improving routing decisions, as well as node programming and re-tasking.

Another issue that remains open is how one could use the WSN local knowledge to improve the decisions taken by the SDN controller. We explore the use of an architecture with distributed and hierarchical controllers, where local controllers leverage from local information to reply to nodes in its cluster area, and a global controller oversees the whole network, exchanging information with local controllers.

Another goal is to include secure node admission control, confidentiality, data authenticity and integrity both for end users as well as SDN devices. Therefore, we could achieve a flexible solution, which eases the heterogeneous wireless networks adaptation to topology changes and to task reconfiguration, while assuring the associated security services.

V. PRIOR RESEARCH

My research focus has been on Wireless Sensor Networks. I concluded my PhD in June 2006 with research focused on *Energy Consumption Trade-offs in Power Constrained Networks*, and my advisor was Professor Katia Obraczka at University of California Santa Cruz.

After my PhD, the research focused on security and its implementation in WSN [13]. Then our research focused on software-defined wireless networks in general, and we have published works on *Software-Defined Networking Based Capacity Sharing in Hybrid Networks* [12], *Decentralizing SDN's control plane* and *Software-defined-networking-enabled capacity sharing in user-centric networks* in collaboration with Professor Katia Obraczka (University California Santa Cruz).

We also looked at software-defined wireless sensor networks. We have specified, implemented and evaluated TinySDN [10]. We discussed *Software-Defined Wireless Sensor Networks and Internet of Things Standardization Synergism* [14] and executed experiments on *Energy Consumption and Memory Footprint Evaluation of RPL and CTP in TinyOS*.

ACKNOWLEDGMENT

This work is funded by São Paulo Research Foundation (FAPESP) under grant #2013/15417-4 and #2014/06479-9.

REFERENCES

[1] D. Culler, D. Estrin, and M. Srivastava, "Overview of sensor networks," *Computer Magazine*, vol. 37, no. 8, pp. 41–49, 2004.

[2] R. Want, B. N. Schilit, and S. Jenson, "Enabling the Internet of Things," *Computer*, vol. 48, no. 1, pp. 28–35, Jan 2015.

[3] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner, "Openflow: enabling innovation in campus networks," *SIGCOMM Comput. Commun. Rev.*, vol. 38, no. 2, pp. 69–74, Mar. 2008. [Online]. Available: http://doi.acm.org/10.1145/1355734.1355746

[4] O. Flauzac, C. Gonzalez, A. Hachani, and F. Nolot, "SDN Based Architecture for IoT and Improvement of the Security," in *Advanced Information Networking and Applications Workshops (WAINA), 2015 IEEE 29th International Conference on*, March 2015, pp. 688–693.

[5] V. Tadinada, "Software Defined Networking: Redefining the Future of Internet in IoT and Cloud Era," in *Future Internet of Things and Cloud (FiCloud), 2014 International Conference on*, Aug 2014, pp. 296–301.

[6] A. Mahmud and R. Rahmani, "Exploitation of openflow in wireless sensor networks," in *Computer Science and Network Technology (ICCSNT), 2011 International Conference on*, vol. 1, 2011, pp. 594–600.

[7] T. Luo, H.-P. Tan, and T. Q. S. Quek, "Sensor openflow: Enabling software-defined wireless sensor networks." *IEEE Communications Letters*, vol. 16, no. 11, pp. 1896–1899, 2012.

[8] S. Costanzo, L. Galluccio, G. Morabito, and S. Palazzo, "Software defined wireless networks: Unbridling sdns," in *Proceedings of the 2012 European Workshop on Software Defined Networking*, ser. EWSDN '12. Washington, DC, USA: IEEE Computer Society, 2012, pp. 1–6. [Online]. Available: http://dx.doi.org/10.1109/EWSDN.2012.12

[9] A. D. Gante, M. Aslan, and A. Matrawy, "Smart wireless sensor network management based on software-defined networking," in *Communications (QBSC), 2014 27th Biennial Symposium on*, June 2014, pp. 71–75.

[10] B. T. de Oliveira, C. B. Margi, and L. B. Gabriel, "TinySDN: Enabling multiple controllers for software-defined wireless sensor networks," in *Communications (LATINCOM), 2014 IEEE Latin-America Conference on*, Nov 2014, pp. 1–6.

[11] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister, "System architecture directions for networked sensors," *SIGPLAN Notices*, vol. 35, no. 11, pp. 93–104, 2000.

[12] M. A. S. Santos, B. T. de Oliveira, C. B. Margi, B. Nunes, T. Turletti, and K. Obraczka, "Software-defined networking based capacity sharing in hybrid networks," in *Proceedings of Capacity Sharing Workshop (CSWS'13), In conjunction with ICNP'13*, May 2013.

[13] C. B. Margi, B. T. de Oliveira, G. T. de Sousa, M. A. Simplício, P. S. L. M. Barreto, T. C. M. B. Carvalho, M. Näslund, and R. Gold, "Impact of operating systems on wireless sensor networks (security) applications and testbeds," in *Computer Communications and Networks (ICCCN), 2010 Proceedings of 19th International Conference on*, Aug. 2010, pp. 1–6.

[14] B. T. de Oliveira, R. C. A. Alves, and C. B. Margi, "Software-Defined Wireless Sensor Networks and Internet of Things Standardization Synergism," in *IEEE CSCN 2015 - Software Defined Sensors Networks and IoT: perspective and proposals for new standardization activities*, Tokyo, JP, Oct 2015.

Network Multi-layer Optimization Strategies

Andrea Fumagalli (andreaf@utdallas.edu), Marco Tacca, and Miguel Razo

White Paper – October 15, 2015

An increasingly large portfolio of network equipment is becoming Software Defined Networking (SDN) enabled. At the same time, a number of SDN controllers are entering the market, each one able to control a specific network layer, ranging from the DWDM layer (L0), OTN layer (L1), Ethernet layer (L2) all the way to the IP/MPLS network layer (L3). Unless carefully orchestrated, these SDN controllers can only achieve local optimization of their respective layer and cannot provide more desirable global optimization outcomes across all layers. Most of the current research efforts are focused on defining orchestrator-layer interactions, while little attention is being paid to designing performing optimization strategies.

We are still facing interesting challenges when it comes to designing multi-layer optimization strategies that can be successfully deployed to ensure well-orchestrated actions across all of the network layers and their respective SDN controllers. These strategies must be carefully designed to account for a number of factors, including the coexistence of a number of independent layer-specific SDN controllers and their optimization engines. Another challenging task is to leverage at its best the latest progress in coherent optical technology while controlling and optimizing network resources at the lowest (DWDM) layer. If properly applied, the performance advantages achievable through coherent optical networking (e.g., more flexible, robust, and predictable optical signal quality) can be propagated across the upper layers all the way to benefit network applications and end-users.

More specifically, the DWDM layer offers unique opportunities for cross-sublayer optimization strategies, which must account for two sides: the communication side (quality of the optical signal in the form or Optical Signal to Noise Ratio or OSNR and Bit Error Rate or BER) and the networking side (performance of the network as a whole). Interesting tradeoffs can be found and should be investigated in this arena realizing that multiple optimization algorithms must be combined to achieve common objectives, i.e., lower latency, increased network capacity, matching application's requirements. Unlike other conventional wired layers (e.g., Ethernet, IP/MPLS), optical networks must account for signal quality, which depends on a number of factors including signal modulation scheme and power level, noise power level, interfering signal power level, signal wavelength or frequency, and power loss due to various Network Elements (NEs). There is an interesting analogy here with wireless networks, in that transmission rates can be automatically adjusted based on the received SNR at the receiver. However, optical networks present an additional challenge when compared to wireless networks, which is represented by their multi-hop, multi-link signal propagation. While in wireless networks the signal is transmitted to a neighboring node (single-hop) where it is processed and regenerated electronically, optical networks allow circuits to be routed over multiple fiber links without requiring signal regeneration at the intermediate nodes. In the latter case, the OSNR at the receiving end-node is thus affected by each link along the signal propagation path, making signal power optimization a much harder and network-wide problem to tackle. An additional issue that requires further investigation is the efficient handling of the optical spectrum when pursuing flexible grid optics. In flex grid networks, optical circuits can be assigned variable portions of the spectrum in order to efficiently match the application bandwidth requirement. A side effect that must be carefully taken into account is the resulting fragmentation of the spectrum availability in every fiber link as well as across network paths.

In summary, while vying for SDN solutions, researchers must look for innovative approaches that account for the concurrent optimization strategies operating at each layer and sub-layer of the network stack. While in today's products, most of these optimization strategies are still operating independently at their layer, it will be beneficial to identify solutions that can orchestrate these strategies for improved global optimization.

Management of The Future Internet With SDN

Antônio Abelém - abelem@ufpa.br

1. Introduction

The Internet was originally designed to provide network services to a closed community, but has become critical to everyday life in domains as diverse as education, health, defense, commerce, travel and entertainment.

Specific adaptations have historically been proposed and implemented at the emergence of new demands. This approach, despite having attended to momentary needs, has generated increasing complexity and cost of maintaining the Internet. Moreover, the higher the number of these adjustments, the greater is the complexity of the resulting architecture, making it more difficult to overcome future challenges, in a situation commonly referenced as Internet ”ossification”, increasingly resistant to structural changes [1].

In this context, Software Defined Networking (SDN) is today one of the most relevant solutions for Future Internet, and the OpenFlow protocol implementation is the best-known method [2].

The Research Group on Computer Networks and Multimedia Communications (GERCOM), coordinates by me, which among other things, was responsible for the FIBRE Project (Brazilian coordinator), which the main goal was the design, implementation and validation of a shared Future Internet research facility, supporting the joint experimentation of European and Brazilian researchers, has been investigating several aspects of the Future Internet, especially within the context of Software Defined Networks and Global access to the Internet for all using wireless technologies.

2. Ongoing Research

- Management SDN using Graph Databases

Current SDN controllers still stores their data in a traditional manner, using relational databases, however this organization is not gainful to the controller or to the networking itself. We are working a proposal for modeling and administrating SDN using a graph-oriented database. The network's abstraction data in the graph database represents its own global view, using graph concepts such as nodes and relationships. Additionally, the data stored by the controller in this approach preserves the data plane's view within these relationships, enabling economy of scale in the control plane usage, facilitating the search for the SDN information.

- Network Virtualization Proxy

The combination of Network Virtualization Function (NFV) and Software Defined Net- working (SDN) can improve the control and utilization of network resources. However, this issue still requires proper solutions to virtualize large-scale networks, which would allow the use of SDN and Virtualization in real environments. we proposed a virtualization architecture for SDN that relies on a proxy-based approach. The NVP (Network Virtualization Proxy) is a virtualization proxy that intercepts messages exchanged between controllers and switches SDN enabling network virtualization. An implementation of the proposal was developed as a proof of concept and load testing was performed showing that the solution can provide network virtualization in a scalable manner, using less than 2.5 MB of memory to manage 100 switches performing simultaneous requests, whereas Flowvisor requires more then 200 MB.

- Software Defined Wireless Networks

SDN also presents as one of the main alternatives to solve problems of traditional technologies of mobile and wireless networks. However, many challenges to the use SDN for wireless networks are still open, including the end to end architecture, design and implementation of wireless technologies programmable and virtualized. Wireless technology is a promising solution to the problem of digital exclusion and can be instrumental in democratizing access to the Internet by unfettering developing communities from the encumbering constraints of infrastructure. It is an alternative and a useful solution to the North region in Brazil (45.27% of Brazilian territory), including the Amazon rainforest, where traditional wired communications and power technologies are frequently difficult or impossible to use. We proposed an extension to the OpenFlow Protocol to adapt it, both the environments of wireless network as to facilitate the integration of these with the wired networks. We defined minimum architecture to the requirements of SDMWN (Software-Defined Multihop Wireless Network) and create new messages in the OpenFlow protocol, extensions for header and a new set of actions for flows tables to contemplate the peculiarities of the wireless network environment.

3. Cyber Infrastructure Experience

A number of global and local researchers (in the US, EU, BR and Asia) are looking at future network architectures and building testbeds to evaluate new protocols and systems based on these new ideas.

- FIBRE Project

FIBRE (http://www.fibre.org.br) was developed as an FP7 project part of cooperation between Brazil and Europe. FIBRE is a heterogeneous facility, multi-domain and multi-technology of intercontinental scope based on the principles of virtualization and federation [3]. FIBRE is currently operated by RNP and it is aimed to be expanded and maintained as a national laboratory service. FIBRE is an OpenFlow-based network that interconnects different universities in Brazil.
Currently, FIBRE operates using two distinct control frameworks, one called OCF to control OpenFlow networks and another called OMF to control wireless devices. We are working in project to propose extensions to the FIBRE control framework in order to support and integrate new technologies with the current infrastructure using XOS, a service orchestration layer built on top of OpenStack and ONOS.

References:

[1] Alcober, J., Hesselbach, X., de la Oliva, A., Garcia-Saavedra, A., Roldan, D., Bock, C.;Internet future architectures for network and media independent services and protocols, 15th International Conference Transparent Optical Networks(ICTON), 2013 On, pp. 14 (2013)

[2] Open Networking Foundation. Software-Defined Networking: The New Norm for Networks. White paper, Open Networking Foundation, Palo Alto, CA, USA (April 2012).

[3] Sallent, S., Abelem, A., Machado, I., Bergesio, L., Fdida, S., de Rezende, J.F., Azodolmolky, S., Salvador, M., Ciuffo, L.N., Tassiulas, L.(2012); Fibre project: Brazil and europe unite forces and testbeds for the internet of the future. In: Korakis, T., Zink, M., Ott, M. (eds.)