Directions for a Software-Defined Next Generation Internet

Marco Canini
Université catholique

1. INTRODUCTION

"In short, software is eating the world."
—Marc Andreessen

This is not just a known quote, but also a reality demonstrated by current evidence on the speed at which software is allowing many industries to innovate at a rapid pace and produce solutions that only a few years ago felt far beyond reach. An important example is Cloud computing, a paradigm shift for making compute, storage and network resources available on demand to organizations of all sizes. Other major examples are the many scientific endeavors enabled by the transformative availability of large amounts of data, supported by evolving infrastructure, software and methods to cope with such big data.

In comparison, innovation in networking (not including transmission technologies) has been historically stagnant during the past couple of decades despite much research efforts. Only recently the emergence of technological trends such as Software-Defined Networking, Network Function Virtualization and Programmable Dataplanes, coupled with the ever-demanding requirements of large data centers have spurred a renaissance of interest into bringing novel ideas for designing radically different networks. In additionally, Software-Defined Radios (SDR) have enabled rapid innovation in a domain that used to be dominated by hardware.

It is clear that this process is being enabled by *the ability to innovate through loosely coupled, adaptive software components* (rather than protocols and their standardization) and by *a new generation of programmable hardware* (which is far more open than traditional proprietary networking and amenable to experimentation). Thus, the process towards Next Generation Internet appears having to pass through the emergence of Software-Defined Infrastructure (SDI) and eXchanges (SDXes).

2. SDI AND SDXES

Today's cyber-infrastructure, implemented via hardware with predetermined control systems, is beginning to morph into SDI: fluid, planetary-scale software systems that are highly interconnected, deeply programmable, and virtualized within end-to-end slices across many administrative domains. As such multi-domain SDI systems grow and multiply, SDXes are emerging as the "meet-me" points and marketplaces for resources from multiple SDI domains.

In the coming decade, we envision this transformation will mutate the fundamental structure of the Internet, enabling entirely new classes of research with implications for revolutionary new applications and services. It is conceivable that today's Internet will run in just one slice across this infrastructure, with many other novel services populating other slices.

The above vision, which is becoming increasingly apparent, was the theme for an NSF-sponsored workshop held in February 2016 chaired by Robert Ricci and Nick Feamster. A multitude of research challenges arise with this vision. Workshop participants expressed a high level of agreement on three fundamental research challenges:

1) What are the right abstractions to represent, program, troubleshoot, and reason about SDI systems that incorporate an enormous variety of devices and services?
2) How can we understand, reason about, troubleshoot, and control the dynamics of large-scale SDI systems, and how can we ensure the robustness and resilience of the services they host in the face of unexpected events?
3) How can we understand, reason about, and manage the socio-technical aspects of SDI systems, including the security, privacy, and data-ownership issues that arise in multi-domain systems that weave together many layers of software?

Due to space limit, I refer the reader to the workshop report [1] for the details. In the following, I discuss a possible research agenda on the road to this vision, focusing on the SDXes.

2.1 Research Challenges Ahead

Future networks will need to enable planetary-scale interconnection of multi-domain, geo-distributed, heterogeneous SDI with capabilities far beyond today's interconnection of Clouds so as to incorporate mass-market Internet-of-Things devices and large scale cyber-physical systems including cities, autonomous automobiles and ubiquitous drones providing massive "data torrents" from their mobile, high-bandwidth sensors.

We face several challenges. Rather than trying to be exhaustive, in what follows I highlight the challenges I presented as the keynote address of the NSF workshop [1].

Reliability & Robustness. SDXes will be critical infrastructure that needs to be dependable. Yet this is made challenging by the growing presence of software in networks. How can we reason about these multi-domain systems? How can we drill down to troubleshoot malfunctions, particularly when some measure of opaqueness is required for each domain? According to a recent report [2], many experts from the networking and formal methods communities believe that despite their importance, tools for programming and reasoning about networks are still in a state of infancy.

Security. Future networks will need to ensure far better security than current Internet. There are security issues that must be addressed in the design of an SDI-based network infrastructure including preventing route hijacks, detecting and mitigating DDoS attacks. Beyond DDoS mitigation, SDXes might enable new architectures that can help by design address network attacks. For example, they provide an opportunity to reconsider the line of research on network capabilities as they might be a way to embed costs into traffic, such that costs can act as a deterrent for attackers.

Privacy. Beyond improving inter-domain routing, there is a clear expectation of deploying virtualized network functions at SDXes such as caches, optimizers, packet scrubbers, etc. This raises questions regarding the privacy of processing traffic at exchanges and the neutrality of SDXes. Who controls the network functions? Who specifies what traffic traverses which network functions? Is there any auditing that we can perform?

Socio-Technical Aspects. How do we ensure cross-domain decisions have enough data about each domain, without revealing proprietary information? For instance, generally, policies used by ISPs for their peering agreements and route selection are private. SDXes need to provide rich services that can be consumed by participants through APIs without this causing leakage of any confidential information to other participants at the SDX.

Quality-of-Experience (QoE). SDXes are well suited to fill a role in increasing QoE in the Internet. Existing problems and inefficiencies largely stem from the lack of information exchange between application providers and ISPs, which can be detrimental for the user. SDXes might enable a bi-directional flow of information between these parties and act as a trusted optimizing arbiter in the presence of contrasting objectives.

Marketplace. SDXes prove the unprecedented opportunity to enable an efficient marketplace at the level of transport and routing. There are several forms this can be envisioned: (1) a marketplace for setting up peering on-demand through fast commercial negotiations and connectivity setup, (2) a marketplace for optimizing routing and providing end-to-end guarantees, leveraging the improved network visibility that SDXes have, (3) a marketplace for 3rd party providers of virtualized network functions (e.g., security providers) to tap into the ecosystem.

2.2 End Game

Given the globally disruptive potential for innovation with SDI and SDXes, it is worth asking the question: How can we best connect a rich research agenda on SDI and SDXes into a platform upon which people can build solutions to solve world's problems? Major ongoing trends such as smart grid, smart transportation, smart cities, electronic voting, green economy, online education require better infrastructure as the Internet is ill suited for many new requirements of these technological innovations. Below I suggest that SDXes fit into these themes by giving three examples. I believe that many more exists and that research on next generation networks will be more likely to succeed if the research program is able to "think big" about and connect the research with its potential for enormous impact on our society.

SDXes for resource fluidity. Connecting hundreds of networks each, an interconnection of SDXes can become a key enabler for better sharing of bandwidth and other resources.

Crowd-sourced SDXes. A global network of SDXes can ensure that independent networks can be quickly constructed and operated when times require it (e.g. natural disasters).

SDX vaults. In the IoT era with massive data acquisition capabilities, SDXes stand to act as a neutral common ground providing brokering services for privacy-preserving data aggregation and analysis between producers and consumers of such data.

3. FINAL REMARKS

Meeting these goals will require both fundamental research dealing with formal foundations for reasoning about and programming next generation networks, understanding and controlling their dynamics, as well as applied, interdisciplinary work that blends together system prototyping, deployment, testbed operation with use cases and long term vision about the potential impact of SDI and SDXes. Along side, this research's infrastructure requirements will need to be supported by given researchers access to large-scale SDI and SDXes on which to perform their research.

4. REFERENCES

[1] Report of the NSF Workshop on Software Defined Infrastructures and Software Defined Exchanges, 2016. https://www.flux.utah.edu/beyond-internet-workshops/sdi/.

[2] N. Bjorner, N. Foster, P. B. Godfrey, and P. Zave. Formal Foundations for Networking (Dagstuhl Seminar 15071). *Dagstuhl Reports*, 5(2):44–63, 2015.