Inter-domain Networking Innovation on Steroids: Empowering IXPs with SDN Capabilities

Pre-print Version

Marco Chiesa[1] Christoph Dietzel[2,7] Gianni Antichi[3] Marc Bruyère[4]
Ignacio Castro[5] Mitch Gusat[6] Thomas King[2] Andrew W. Moore[3]
Thanh Dang Nguyen[1] Philippe Owezarski[4] Steve Uhlig[5] Marco Canini[1] *

Abstract

While innovation in inter-domain routing has remained stagnant for over a decade, Internet Exchange Points (IXPs) are consolidating their role as economically advantageous interconnection points for reducing path latencies and exchanging ever increasing amounts of traffic. As such, IXPs appear as a natural place to foster network innovation and assess the benefits of Software-Defined Networking (SDN), a recent technological trend that has already boosted innovation within data-center networks.

In this paper, we give a comprehensive overview of use cases for SDN at IXPs, which leverage the superior vantage point of an IXP to introduce advanced features like load-balancing and DDoS mitigation. We discuss the benefits of SDN solutions by analyzing real-world data from one of the largest IXPs. We also leverage insights into IXP operations to not only shape benefits for members but also for operators.

1 Introduction

The growth of demands for high performance online services is posing tremendous challenges on the independent networks, i.e., Autonomous Systems (ASes), that carry the traffic of these services and form the Internet. Pressing requirements, such as lower latencies and higher bandwidth, have pushed ASes to move from the strict hierarchical, transit-based interconnection model of the early commercial Internet towards a denser and flatter structure.

Internet eXchange Points (IXPs) are playing a leading role during this transition phase by providing a simple layer 2 broadcast domain to which members connect and exchange IP traffic, possibly with any other member (see Figure 1 left). After establishing peering relationships, IXP members exchange routing information by means of the Border Gateway Protocol (BGP), the de-facto standard inter-domain routing protocol. IXPs have become central to the Internet peering ecosystem by attracting increasing numbers of members and, consequently, traffic. Nowadays, about 80% of the address space is reachable through more than 350 existing IXPs [1]. The largest IXPs interconnect hundreds of ASes and carry traffic volumes comparable to those of Tier 1 transit providers [2].

Over a decade of work has gone into proposing modifications to the BGP routing control plane to improve its security and make it easier to manage and troubleshoot [3]. As these proposals require substantial global changes in BGP, unfortunately there has been no significant adoption. As such, inter-domain routing still suffers from well-known shortcomings of BGP such as its coarse-grained control of traffic based on just destination IP prefixes, and indirect control of how remote networks forward traffic.

As others have argued [4], we also deem IXPs are

* [1] Université catholique de Louvain, [2] DE-CIX, [3] University of Cambridge, [4] LAAS-CNRS, [4] Queen Mary University of London, [6] IBM Research, [7] TU Berlin

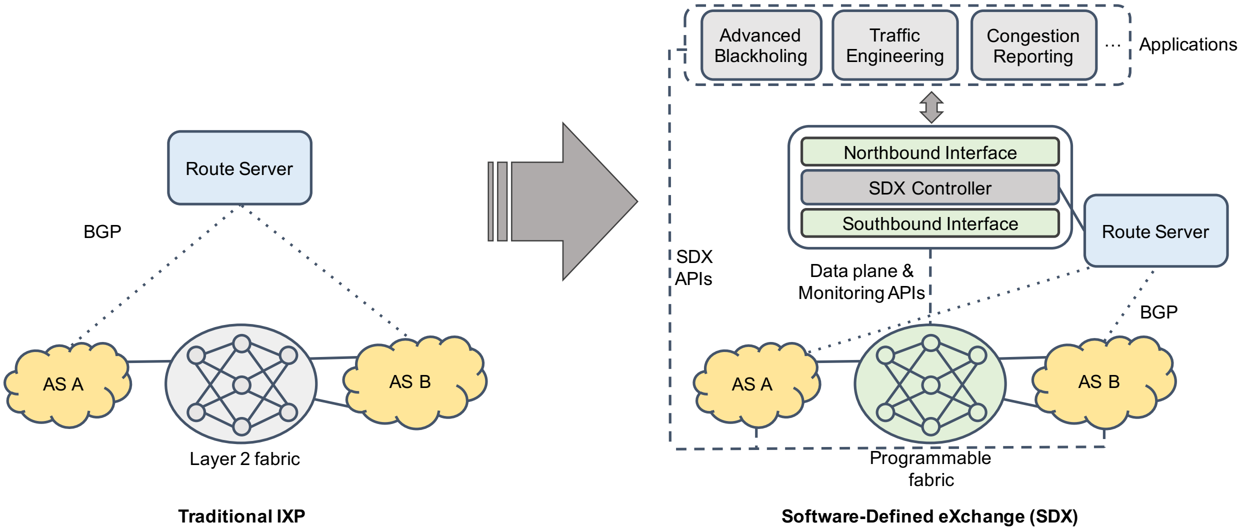

Figure 1: The evolution from a traditional IXP fabric to a richer SDN-enabled IXP: i.e., a Software-Defined eXchange (SDX). The SDX exposes APIs to IXP members through which they can consume advanced services. For compatibility, IXP members exchange routes using BGP via a route server. Applications establish high-level policies and behaviors for the underlying IXP fabric. The SDX Controller mediates access to the programmable IXP fabric and exposes monitoring information.

an ideal place to spur innovation in the Internet ecosystem. First, IXPs are convergence points for a large number of ASes. Any improvement that can be deployed at a large IXP has direct impact on hundreds of members. Second, IXPs' physical networks are fairly limited in size. This eases any physical migration towards novel network architectures and infrastructure. Third, IXPs are interested in offering services that go beyond simple layer 2 connectivity. To simplify peering, many IXPs already operate route servers, which allow IXP members to peer with many other ASes via a single BGP session to a route server. Finally, IXPs have strong economic incentives to embrace innovation for reducing the operating costs in face of continuous traffic growth in their networks.

Software-Defined Networking (SDN) recently emerged as a new paradigm that enables network programmability to facilitate management and enhance flexibility [5]. SDN supports logically centralized control of a collection of programmable switches that can act on traffic at a fine level of granularity. In contrast, BGP, which has been the standard for exchanging inter-domain reachability information for decades, only provides limited and indirect control knobs to network operators. BGP operates at a coarse-grained level based on IP destination prefixes. Despite of the SDN potential, its success has been so far restricted to environments such as intra-domain routing and data-centers [6].

In this paper, we advocate that IXPs offer an exciting opportunity for inter-domain networking to benefit from the advantages of SDN deployment in neutral, dense environments; thus giving rise to the concept of Software-Defined eXchanges (SDXes). By looking at real data from one of the largest IXPs, we demonstrate the need of enhanced network capabilities. Further, we show the potential benefits of SDXes by presenting use cases. We advance existing work [4, 7–9] by proposing use cases that benefit not only the IXP members but also the IXP management itself. Also, we show evidence of the need of a breakthrough in network management at the IXP. We believe that SDN programmability and its finer-grained control capabilities at large IXPs will lead to novel peering arrangements, greater responsiveness, and easier network management.

The work presented in this paper is part of an ongoing effort within the ENDEAVOUR project[1], a research and innovation action funded by the European Union's Horizon 2020 program, which began in January 2015. By focusing on SDN deployment at IXPs, the project aims to create a flexible SDN ecosystem, which can support a service marketplace composed of Cloud/data-centers, networked applications, and the underlying interconnection fabric.

Figure 1 depicts the architectural evolution that we envision as IXPs transition to SDN. At the data plane level, programmable SDN switches forward traffic according to fine-grained forwarding and filtering rules generated by the control plane (i.e., the SDX Controller) in accordance to declarative high-level goals established by applications. In this scenario, a number of different applications can be instantiated to accomplish various use cases (e.g., Traffic Engineering, Advanced Blackholing). The SDX controller then translates these goals in well-defined rules and programs the underlying data plane through the southbound interface, e.g., OpenFlow.

The rest of the paper is organized as follows. In Section 2, we provide an overview of those customized applications, including fine-grained traffic engineering (Sec. 2.1.1), mitigation of Denial-of-Service (DoS) attacks (Sec. 2.1.2), broadcast message reduction (Sec. 2.2.1), layer 2 label switching (Sec. 2.2.2), and other emerging use cases (Sec. 2.3). Section 3 discusses some of the challenges and the road ahead for SDXes, and finally, Section 4 concludes.

2 Use-Cases

This section presents some of the benefits we believe an SDX approach brings to the current generation of IXPs from both operators' and members' perspective.

2.1 IXP Member's Benefits

2.1.1 Traffic Engineering

Traffic-Engineering (TE) refers to the steering of traffic flows by network operators to improve performance. Achieving these goals depends on the accuracy of the network state estimation, the algorithms used to compute the optimal routing paths, and the specific features of the routing scheme adopted (e.g., per-destination routing, hash-based load balancing).

At an IXP, ASes exchanging large traffic volumes connect with multiple ports (e.g., multiple 100GE). These ASes employ TE to load-balance traffic through these ports and avoid congestion while attaining high port bandwidth utilization. The outcome of the TE operations is the result of the interplay between the inbound TE policies of the traffic receiver and the outbound TE policies of the originator. While the inbound policies specify what type of traffic can be received through each port alongside information about the route used to forward that traffic, outbound policies involve a mechanism to compute where to send the originated traffic.

Well-known limitations of BGP hinder what TE goals network operators can achieve. Operators must resort to indirect BGP configuration mechanisms, e.g., AS path prepending, communities, and selective announcements. It remains to hope that configuration changes have the desired effect, e.g., incoming traffic is evenly split among the IXP member ports. Certain types of control are simply not possible. This problem is particularly acute at IXPs where the wide range of independent and inconsistent peering policies might clash. The lack of network programmability poses an additional problem that can lead to human mistakes and reduces AS responsiveness to network events.

Despite the aforementioned pitfalls of current solutions, TE is widely performed by IXP members. Figure 2 shows how members with multiple ports (roughly 15% of the IXP members) load-balance their inbound traffic across them. Each IXP member is depicted by a vertical bar and the partitioning within the same bar representing the relative traffic volume per port. Observe that 27 out of 103 IXP members

[1] https://www.h2020-endeavour.eu/

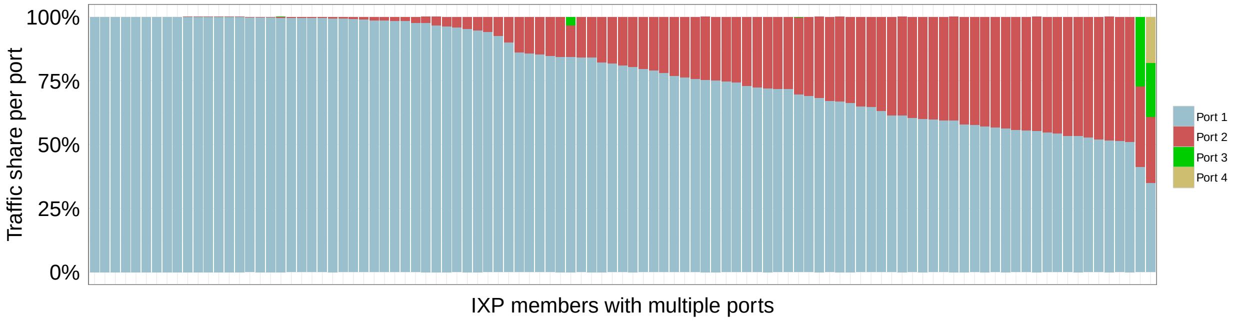

Figure 2: Inbound traffic load-balancing in a large IXP.

receive more than 99% of the traffic through a single port, probably using the second one as a backup port. The remaining members perform inbound TE more intensively. We advocate that TE operations should be simplified by means of more intuitive primitives and easier network programmability.

We believe that SDN substantially enhances TE operations. First, it allows network operators to express their TE routing policies with a specification language that is more fine-grained than the legacy BGP-based one. Operators can effectively use this greater expressiveness to effectively steer the flows of traffic according to layer 2-4 attributes. Second, to perform TE, IXP members have to estimate their incoming traffic volumes. IXPs can leverage SDN monitoring capabilities to provide an interface for accessing a global coherent view of the state of the network, which also include incoming traffic volumes statistics. In addition, TE can directly be outsourced to the IXP, which applies the policies specified by the IXP members. Automating TE engineering leads to less human errors whereas outsourcing the process frees IXP member resources. Third, SDN programmable networks allow operators to choose TE algorithms based on their performance goals and scalability limits. For instance, SDN-based load-balancing can be performed both in static and dynamic manner. The former leverages weighted hash-based per-flow load-balancing mechanisms, which is available in the OpenFlow 1.3 standard. It allows an IXP network to spread the incoming traffic on behalf of the receiver member without any knowledge of the traffic volumes. The latter enables each IXP member to further optimize TE whenever the former approach does not achieve the desired optimization goals, as it is the case whenever a few large flows of traffic are hashed to the same port. In that case, fine-grained monitoring and routing capabilities can be used to detect and reroute those flows of traffic that cause traffic unbalances.

2.1.2 Advanced Blackholing

Distributed Denial of Service (DDoS) attacks are a serious threat to the Internet ecosystem.[2] To cope with this threat some IXPs offer blackholing services to their members.[3] Blackholing in general is a technique that allows an AS to ask its neighbors to drop packets destined towards a certain IP prefix. DDoS attacks, which typically affect a limited number of IP destination addresses, cannot be mitigated by withdrawing the BGP announcement of the entire IP prefix that contains the target of the attack, since this would affect also large portions of legitimate traffic. With blackholing, an AS first has to detect that it is receiving malicious traffic targeting a specific small contiguous set of IP prefixes. To stop the incoming flow, it sends a special BGP announcement for the affected IP prefix to its upstream neighbor that orig-

[2] `http://blog.cloudflare.com/the-ddos-that-almost-broke-the-internet`

[3] `www.de-cix.net/products-services/de-cix-frankfurt/blackholing/`

inates this traffic. Upon receiving such message, each packet destined to that IP prefix is deflected towards a dummy MAC address by means of Address Resolution Protocol (ARP) mechanisms. All packets with this destination MAC are discarded via layer 2 filtering rules implemented via access control lists at the IXP switching fabric.

Current blackholing implementations at IXPs are too coarse-grained [10]. Although AS operators are capable of filtering on layer 4 ports within their own ASes, fine-grained blackholing for multiple AS border routers is severely limited by ACL hardware constraints, e.g., due to static nature and limited number of rules. Also, it does not protect the AS peering link at the IXP, i.e., it can be overwhelmed and congested. Furthermore, blackholing cannot distinguish between legitimate and malicious traffic: all packets destined for the blackholed IP prefix are dropped, thus disconnecting its upstream networks. To make things more complex, blackholing lacks programmability and must be manually triggered after the DDoS attack is detected. Since the AS operators cannot observe the traffic volumes (e.g., attack termination) and patterns (e.g., destination port mix), identifying the beginning and termination of an attack is cumbersome.

We envision SDN-enabled blackholing to overcome those limitations. SDN, and OpenFlow in particular, allows operators to specify fine-grained drop policies and eases the blackholing process, hence minimizing the risk of misconfiguration. Using OpenFlow, ASes can detect a DDoS attack by monitoring the traffic properties and define fine-grained drop rules to discard the unwanted packets. The IXP can then provide an interface, e.g., an API, so that members can express their very own precise drop rules and have them automatically implemented when an attack is detected. The IXP can also provide insights in the blackholed traffic by monitoring the corresponding flows through the OpenFlow flow counters.

While some traffic might still be unintentionally blackholed, OpenFlow rate-limiting capabilities can alleviate the problem. By limiting the traffic (according to specific header fields) towards the attacked members to a non-critical volume, the legitimate traffic can still stand a chance to reach its destination.

2.2 IXP Operator's Benefits

2.2.1 Controlling Broadcast

IXPs interconnect member's routers through a shared layer 2 broadcast domain. As the scale of IXPs keep increasing (the biggest IXPs nowdays count more than 600 members), traditional address resolution mechanisms that rely on broadcast solutions, e.g., ARP or Neighbor Discovery (ND), pose a challenge for faultless operation and stability. Broadcast packets are needlessly processed by all the routers connected at the IXP, consuming an excessive amount of each router's CPU capacity. Thus, broadcast mitigation and filtering at the IXP's edge becomes crucial. Figure 3 shows an increase of a 10% in the volume of broadcast traffic (i.e., location discovery) at a large IXP in only 15 months. This growth continues to happen in spite of IXPs' strict rules for the configuration of the members' routers[4] and techniques such as ARP sponge.[5]

In contrast to a traditional switch, an SDN switch by default can drop all packets not matching any of the installed forwarding rules. IXP operators can thus program the IXP fabric to solve the problems exposed above. We observe that the topology of an IXP fabric is fairly stable over time, with the addition or removal of members happening on a sufficiently low frequency. Because the location of all members' routers is stable and well-known to the IXP operators, broadcast traffic can be eliminated by just programming the switching SDN-enabled fabric to exclusively transport packets to the requested destination by ARP and ND packets.

2.2.2 Layer 2 Label Switching

Large IXPs, especially those with multiple points of presence need complex and robust infrastructures to satisfy members' requirements, such as robustness and scalability. These IXPs use a transparent layer 3 infrastructure for internally addressing these require-

[4] `https://www.euro-ix.net/networks/configuration-samples/`

[5] `https://ams-ix.net/technical/specifications-descriptions/controlling-arp-traffic-on-ams-ix-platform`

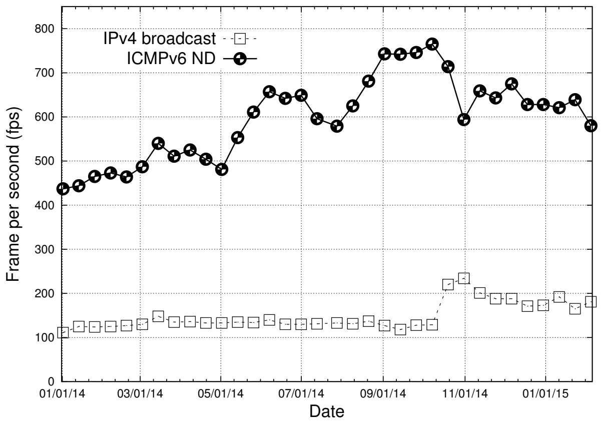

Figure 3: ARP requests and ICMPv6 ND frames per second at a large IXP.

ments, while providing an emulated layer 2 fabric to their members.[6] Unfortunately, exploiting path diversity within an IXP fabric remains a problem. Traditional solutions such as MPLS and VPLS help but come at the cost of increased complexity. Despite multiple recent attempts, the two main standards bodies, i.e, IEEE 802.1aq[7] and TRILL,[8] have not yet converged to a uniquely and universally accepted solution. As this divergence effectively precludes the interoperability across vendors, a solution is required.

We believe that SDN can help to simplify this complexity while preserving the benefits of layer 3 protocols, i.e., leveraging the IXP architecture and reducing the protocol and management overhead. Fine-grained label switching capabilities such as those provided by OpenFlow can be used to route the traffic efficiently over multiple paths across the switching fabric. Ingress packets can be matched with installed flow rules and annotated with a label, either with an MPLS label or encoded in any arbitrary header field [7], which shall be removed at the IXP's egress interface.

2.3 Emerging Novel Services

2.3.1 Port-Based Congestion Reporting

Congestion and its management is an essential factor for an IXP and its members. However, in spite of experiencing significant and persistent congestion at multiple peering links, ASes and IXPs have no primary means of controlling congestion. That is, as the traffic sources and destinations are beyond its domain, an IXP cannot rely on the traditional congestion notification mechanisms such as Explicit Congestion Notification (RFC3168).

Nowadays, congestion at the level of a member's port is a significant problem. To illustrate, Figure 4 shows our measurements from one of the largest IXPs during a seven day period, where we monitor using a 5-minute interval the traffic volume flowing from the IXP into each IXP member's port. We found that 43 ($\sim$ 6%) out of the 760 IXP member's ports suffered at least once from congestion, i.e., more than 100% utilization. Once a member's egress port reaches 100% utilization, excess traffic starts filling buffers and is dropped once buffers are full. By accounting for all traffic that traverses the IXP, our measurement capture the severity of this congestion. For instance, a 200 Mbits and a 10 Gbits port experienced traffic rates at 473% and 214% of their capacity, respectively. Moreover, we found that utilization at 114 ports exceeds 85%.

Empowering IXPs with SDN, and in particular with OpenFlow capabilities, would ease both the per port utilization and packet loss monitoring process for each member. Furthermore, SDN enables highly dynamic solutions wherein the ASes can be informed whenever a specific congestion level is reached. For instance, since version 1.5.1, OpenFlow includes push-based counters monitoring triggered by predefined thresholds. IXP members could act upon congestion reports and apply different routing policies based on network conditions.

2.3.2 IXP as Transport Marketplace

IXPs have dynamized peering interconnections by enabling ASes to automatically peer through the route server with those ASes willing to peer with any other

[6] `https://ams-ix.net/technical/ams-ix-infrastructure/the-ams-ix-mplsvpls-infrastructure`

[7] `http://www.ieee802.org/1/pages/802.1aq.html`

[8] `http://tools.ietf.org/html/rfc6325`

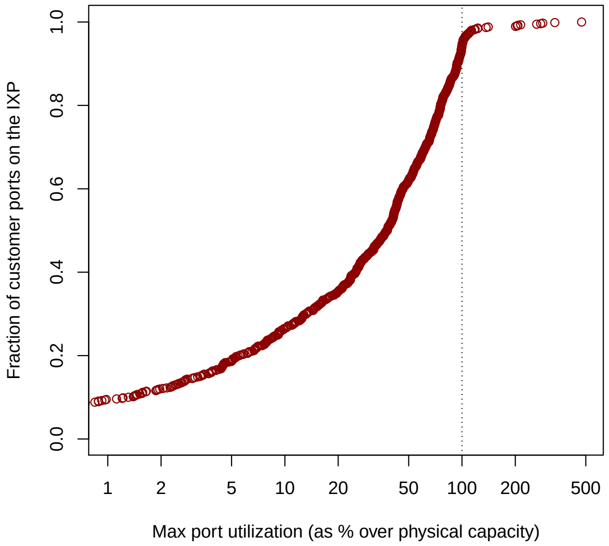

Figure 4: CDF of port utilization in percent showing that 43 ports exceed 100% utilization.

network. Furthermore, going beyond their original goal of being the central place for peering relationships, IXPs have expanded the transport opportunities. IXPs are frequently interconnected by remote peering providers, which offer layer 2 transport between IXPs [1]. This interconnection between IXPs is frequently directly promoted by the IXPs themselves. There is even anecdotal evidence of transit providers offering their transport services at IXPs.

Building upon this increasing diversity, IXPs could become a marketplace for transport. Benefiting from the co-location of ASes, the existing infrastructure and transparency standards, the IXP could help supply to meet the demand of transport over the Internet in all its different flavors.

Despite these developments, current interconnection practices are still characterized by a cumbersome and slow contract negotiation process. We believe that SDN can play a key-role in supporting IXPs to meet those challenges. As with other SDN applications, APIs at the IXP can be leveraged to facilitate the choice of how and whether to interconnect. For instance, an AS might declare its peering conditions through the API and all networks meeting them can automatically peer.

2.3.3 Service Chaining

Increasingly, middleboxes (e.g., firewalls, VPN gateways, proxies, intrusion detection systems, WAN optimizers or scrubbers) play a critical role for the performance and security of many networks. Operators commonly want to chain together multiple middleboxes to create a processing path of desired network functions, also called a *network service* [11]. Service chaining is emerging as an appealing solution to enable operators to dynamically deploy network services without having to make hardwire physical middleboxes. However, using traditional protocols like BGP or MPLS, service chains are difficult to deploy and change. Operators must carefully configure protocols to “hijack” wanted traffic and steer it through the sequence of middleboxes. Moreover, such mechanisms steer traffic based only on IP destination addresses, limiting the overall service granularity. In competitive markets, with rapid innovation at the application layer, this limits operators’ ability to address emerging use cases and business models. Instead, an SDN approach is a well suited alternative to enable service chaining given its logically centralized management and configurable forwarding rules at fine granularity. Thus, inter-domain routing convergence points like IXPs appear as promising locations to deploy service chains.

3 Discussion and Outlook

In the previous section, we presented a selection of appealing use cases that highlight the potential for innovation with SDXes. We now turn our discussion on broad research challenges towards making SDN at IXPs within grasp.

Scalability. SDXes will interconnect hundreds of networks and expose rich service APIs. This creates scalability challenges that will need to be addressed. For instance, given the size of the global Internet routing, which counts over 500K IP prefixes, naively supporting arbitrary SDN applications may result in solutions that cannot even cope with the scale of medium-sized IXPs. Moreover, the growing interest towards remote peering [1] will pose even more challenges in designing scalable solutions for large-sized IXPs. In fact, despite the SDN fine-grained control capabilities over the traffic, the hardware forwarding table sizes are a constrained resource. These considerations demand techniques such as those explored in the iSDX design [7] to efficiently compute and encode flexible routing policies and to minimize the number of data plane updates due to changes in inter-domain reachability information.

Reliability. A recent incident at AMS-IX in May 2015,[9] wherein a forwarding loop affected the IXP for a brief time, demonstrated how configuration mistakes or failures can dramatically impact the network. In this regards, the rise of new applications enabled by SDN will pose new challenges in ensuring a stable and reliable network operation. According to a recent report [12], 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. We believe this also applies to SDXes specifically.

Security. Increased Internet security is a goal highly sought after and a desired use case by operators, in particular, for detecting and preventing DDoS. Beyond DDoS mitigation, SDXes might enable new architectures that can help address broad classes of network attacks by design. For example, these architectures can provide an opportunity to reconsider the line of research on network capabilities, where embedding costs into traffic could act as a deterrent for attackers [13].

Privacy. SDXes will provide a multitude of services beyond a traditional layer 2 interconnection including network functions (such as caches, optimizers, packet scrubbers) that will be deployed at SDXes. This raises questions regarding the privacy of processing traffic at exchanges and the neutrality of SDXes [14]. Who controls the network functions? Who specifies what traffic traverses which network functions? How will it be regulated and what auditing will be performed?

Business Confidentiality. Members will interact with SDXes' control software through APIs to consume services. However, the correct or efficient consumption of these services might require exchange of information that is traditionally considered proprietary due to its business-critical nature, such as peering and routing policies. Therefore, new solutions (e.g., see [15]) are needed to align conflicting objectives such as service consumption while avoiding leakage of confidential information. This is very challenging as possible approaches such as secure multiparty computations might be prohibitively compute intensive in practical settings. Beyond the technical solutions, we believe that to a large extent the problem lies in the vague legal framework for IXP-data disclosure rather than on whether the IXP is a reliable partner for neutral information sharing. A clearer legal framework will definitely help foster innovation at IXPs.

Transitioning to SDN. Despite the advantages that SDN offers for innovating wide-area traffic delivery, SDXes will require a clear migration path from existing systems to deploy SDN software and hardware. We observe that a first step has been made to show the viability of deploying SDN hardware side-by-side with existing production equipment at a public IXPs [16, 17], and also best-practices exist for migration in enterprise data-center, campus [18], and carrier networks.[10] However, actual experience with production deployments remain limited and research that further spurs early adoption in operational networks will be crucial towards reaching a wide-spread deployment.

[9] https://ams-ix.net/newsitems/195

[10] https://www.opennetworking.org/images/stories/downloads/sdn-resources/use-cases/Migration-WG-Use-Cases.pdf

4 Summary

We explored the benefits of casting SDN concepts into Internet eXchange Points for fostering innovation in the Internet peering ecosystem. Based on multiple discussions with operators from the IXP community and an analysis of data traces from one of the largest IXP in the world, we illustrated several use cases that stand to benefit from advanced SDN capabilities at IXPs. Finally, we discussed open problems towards reaching the goal of wide-spread SDN deployment at IXPs.

Acknowledgments: We thank Josh Bailey, Yatish Kumar, David Meyer, and Jennifer Rexford for providing advisory guidance to the project. This research is (in part) supported by European Union's Horizon 2020 research and innovation programme under the ENDEAVOUR project (grant agreement 644960).

References

[1] I. Castro, J. C. Cardona, S. Gorinsky, and P. Francois, "Remote Peering: More Peering without Internet Flattening," in *CoNEXT*, 2014.

[2] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," in *SIGCOMM*, 2012.

[3] K. R. B. Butler, T. R. Farley, P. McDaniel, and J. Rexford, "A survey of BGP security issues and solutions," *Proceedings of the IEEE*, vol. 98, no. 1, pp. 100–122, 2010. [Online]. Available: http://dx.doi.org/10.1109/JPROC.2009.2034031

[4] N. Feamster, J. Rexford, S. Shenker, D. Levin, R. Clark, and J. Bailey, "SDX: A Software Defined Internet Exchange," in *Open Networking Summit*, 2013.

[5] H. Kim and N. Feamster, "Improving Network Management with Software Defined Networking," *Communications Magazine, IEEE*, vol. 51, no. 2, pp. 114–119, February 2013.

[6] D. Kreutz, F. M. 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, 2015.

[7] A. Gupta, R. MacDavid, R. Birkner, M. Canini, N. Feamster, J. Rexford, and L. Vanbever, "An Industrial-Scale Software Defined Internet Exchange Point," in *NSDI*, 2016.

[8] 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 *SIGCOMM*, 2014.

[9] S. Whyte, "Project CARDIGAN An SDN Controlled Exchange Fabric," 2012, https://goo.gl/eTTgxG.

[10] C. Dietzel, A. Feldmann, and T. King, "Blackholing at IXPs: On the Effectiveness of DDoS Mitigation in the Wild," in *PAM*, 2016.

[11] S. Palkar, C. Lan, S. Han, K. Jang, A. Panda, S. Ratnasamy, L. Rizzo, and S. Shenker, "E2: A Framework for NFV Applications," in *SOSP*, 2015, pp. 121–136.

[12] N. Bjorner, N. Foster, P. B. Godfrey, and P. Zave, "Formal Foundations for Networking (Dagstuhl Seminar 15071)," *Dagstuhl Reports*, vol. 5, no. 2, pp. 44–63, 2015.

[13] C. A. Shue, A. J. Kalafut, M. Allman, and C. R. Taylor, "On Building Inexpensive Network Capabilities," *SIGCOMM Comput. Commun. Rev.*, vol. 42, no. 2, pp. 72–79, Apr. 2012. [Online]. Available: http://doi.acm.org/10.1145/2185376.2185386

[14] J. Sherry, C. Lan, R. A. Popa, and S. Ratnasamy, "BlindBox: Deep Packet Inspection over Encrypted Traffic," in *SIGCOMM*, 2015, pp. 213–226.

[15] M. Chiesa, D. Demmler, M. Canini, M. Schapira, and T. Schneider, "Towards Securing Internet eXchange Points
Against Curious onlooKers," in *Applied Networking Research Workshop (ANRW '16)*, Jul. 2016.

[16] J. Stringer, Q. Fu, C. Lorier, and C. E. Rothenberg, "Cardigan: Deploying a Distributed Routing Fabric," in *HotSDN*, 2013.

[17] LightReading, "Pica8 Powers French TOUIX SDN-Driven Internet Exchange," Jun. 2015, http://ubm.io/1Vc0SLE.

[18] M. Canini, A. Feldmann, D. Levin, F. Schaffert, and S. Schmid, "Software-Defined Networks: Incremental Deployment with Panopticon," *Computer*, vol. 47, no. 11, pp. 56–60, Nov 2014.

Biographies

***Marco Chiesa** is a Postdoctoral Researcher at Université catholique de Louvain, Belgium. His research interests include routing optimization and security. Chiesa received a PhD in computer science and engineering from Roma Tre University. He is a member of IEEE. Contact him at marco.chiesa@uclouvain.be.*

***Christoph Dietzel** is member of the DE-CIX R&D team since June 2014 where he is responsible for several research efforts and also involved in numerous projects funded by the public sector. Moreover, he is a PhD student in the INET group, advised by Anja Feldmann at Technische Universität*

Berlin, since the end of 2014. His ongoing research interests focus on Internet measurements and security, routing, and traffic classification. Dietzel is also highly interested in IXP-related aspects of the Internet ecosystem. Contact him at christoph.dietzel@de-cix.net.

Gianni Antichi *is a Senior Researcher at the Network and Operative System group in the Computer Laboratory at University of Cambridge, United Kingdom. His research interests include software-defined networking, network measurements and hardware accelerated networking systems. Antichi received a PhD in computer science and engineering from the University of Pisa in 2011. Contact him at gianni.antichi@cl.cam.ac.uk.*

Marc Bruyère *is a Ph.D. student at the LAAS CNRS in Toulouse, France. He designed and deployed the first European OpenFlow IXP fabric for the TouIX. He started his career in 1996 working for Club-Internet.fr, and for Cisco, Vivendi Universal, Credit Suisse First Boston, Airbus/Dimension Data, Force10 Networks, and Dell. He is a Cisco Certified Internetwork Expert. He has been involved in the NetFPGA project for a few years, and his Ph.D. thesis is about measurements in an IXP OpenFlow/SDN environment. Contact him at marc.bruyere@laas.fr.*

Ignacio Castro *received a Ph.D. degree from the Internet Interdisciplinary Institute (IN3) while researching at the Institute IMDEA Networks in 2015. He is currently a postdoctoral researcher at Queen Mary University of London. His research interests focus on the economics of Internet interconnections. Contact him at i.castro@qmul.ac.uk.*

Mitch Gusat *is a Researcher at the IBM Zurich Research Laboratory. His current focus is on datacenter performance, distributed deep learning and feedback. He has contributed to the design and standardization of CEE, IBA and RapidIO - while also advising Master and PhD students from several European universities. He is member of ACM, IEEE, and holds a few dozen patents related to SDN, transports, HPC architectures, switching and scheduling. Contact him at mig@zurich.ibm.com.*

Thomas King *received an M.Sc. degree (Diplom) in Computer Science and Business Administration from the University of Mannheim, Germany in 2004, and a Ph.D. degree at the chair of Computer Networks from the University of Mannheim in 2008. Thomas was Head of the Research & Development department at DE-CIX until the end of 2015. In 2016, Thomas King has been promoted to the newly-created position of CIO of DE-CIX. Contact him at thomas.king@de-cix.net.*

Andrew W. Moore *is a Reader in systems at the University of Cambridge Computer Laboratory in England, where he is part of the Systems Research Group working on issues of network and computer architecture. His research interests include enabling open-network research and education using the NetFPGA platform; other research pursuits include low power energy-aware networking, and novel network and systems data-center architectures. He is a member of IEEE. Contact him at andrew.moore@cl.cam.ac.uk.*

Thanh Dang Nguyen *is Postdoctoral Researcher at Université catholique de Louvain, Belgium. His research interests include software-defined networking and large-scale distributed computing. He obtained his Ph.D. in Information Science from Japan Advanced Institute of Science and Technology (JAIST) in 2015. He is a member of ACM. Contact him at thanh.nguyen@uclouvain.be.*

Philippe Owezarski *is Director of Research at CNRS (the French center for scientific research), working at LAAS (Laboratory for Analysis and Architecture of Systems), in Toulouse, France. He got a PhD in computer science in 1996 from Paul Sabatier University, Toulouse III, and habilitation for advising research in 2006. His main interests deal with next generation Internet, more specifically taking advantage of IP networks monitoring and machine learning for enforcing Quality of Service and security. Contact him at owe@laas.fr.*

Steve Uhlig *received a Ph.D. degree in applied sciences from the University of Louvain (2004). He is currently a Professor of Networks at Queen Mary University of London. His research interests are focused on the large-scale behaviour of the Internet, Internet measurements, software-defined networking, and content delivery. He is a member of ACM. Contact him at steve.uhlig@qmul.ac.uk.*

Marco Canini *is an Assistant Professor in the Institute of Information and Communication Technologies, Electronics, and Applied Mathematics (ICTEAM) at Université catholique de Louvain, Belgium. His research interests include software-defined networking and large-scale and distributed cloud computing. Canini received a PhD in computer science and engineering from the University of Genoa. He is a member of IEEE and ACM. Contact him at marco.canini@uclouvain.be.*