

Outsourcing the Routing Control Logic: Better Internet Routing Based on SDN Principles

Vasileios Kotronis
ETH Zurich
Zurich, Switzerland
vkotroni@tik.ee.ethz.ch

Xenofontas
Dimitropoulos
ETH Zurich
Zurich, Switzerland
fontas@tik.ee.ethz.ch

Bernhard Ager
ETH Zurich
Zurich, Switzerland
bager@tik.ee.ethz.ch

ABSTRACT

Inter-domain routing is based on a fully decentralized model, where multiple Autonomous Systems (AS) interact via the BGP protocol. Although BGP is the "glue" of the Internet, it faces a lot of problems regarding its fully distributed nature, policy enforcement capabilities, scalability, security and complexity. Due to the widespread adoption of BGP, only incrementally deployable solutions are feasible. Taking this observation into account, we describe a new, backwards-compatible routing model which is based on outsourcing and logically centralizing the routing control plane. We claim that outsourcing enables enhanced inter-domain routing. As multiple ASes outsource their routing control plane to the same outsourcing service contractor, AS clusters are being gradually formed. A logically centralized multi-AS routing control platform based on Software Defined Networking (SDN) principles is the natural point for taking efficient routing decisions, detecting policy conflicts, troubleshooting routing problems, and evolving BGP. We present the technical and financial incentives which support the feasibility of the model and also propose an implementation scheme to facilitate it.

Categories and Subject Descriptors

C.2.1 [Computer-Communication Networks]: Network Architecture and Design—*Distributed networks*; C.2.2 [Network Protocols]: Routing protocols; C.2.3 [Network Operations]: Network management; K.6.0 [Management of Computing and Information Systems]: General—*Economics*; K.6.4 [System Management]: Centralization/decentralization

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Hotnets '12, October 29–30, 2012, Seattle, WA, USA.
Copyright 2012 ACM 978-1-4503-1776-4/10/12 ...\$10.00.

General Terms

Design, Economics, Management

Keywords

Outsourcing, Routing, Management, BGP, SDN

1. INTRODUCTION

BGP is the de-facto inter-domain routing standard today. It keeps the Internet working in a maze of diverse financial interests. ASes use BGP not only to disseminate and acquire reachability information, but also to enforce multifarious policies based on business relationships, traffic engineering schemes, security practices, etc.

BGP-related previous research has thoroughly analyzed the problems which the protocol faces. We summarize the most important of them. First, the fully distributed nature of the protocol is simultaneously a strength and weakness of BGP. Slow convergence times due to path exploration [24] and route flap damping [20] are commonly experienced [16]. Unstable prefixes can cause route oscillation [22] and inter-AS policy conflicts can lead to global instability or even divergence [12].

Furthermore, routing policy enforcement is cumbersome and complex [5]. ASes' policies need to be expressed in a way so that they are totally aligned to the BGP routing mechanism, while the opposite should ideally occur. In addition, operators are forced to overload the embedded BGP mechanisms, e.g., the local preference attribute, in order to implement different policies. Besides, the internal (IGP) and external (BGP) policy mechanisms are not bridged and controlled under a unified scheme leading to further problems [28].

Despite the fact that a large number of proposals aiming towards solving the problems of BGP have been formed, most of these proposals never leave the research stage. The reasons are multiple. First, BGP is a widely adopted protocol implemented by many stakeholders and is therefore very difficult to change. Second, ISPs cannot be easily convinced to take the risk of adopting a proposed improvement unless substantial profit is

imminent. Third, the strict requirement of maintaining backwards compatibility nips many good ideas in the bud.

In this context, SDN offers new opportunities. The key concept of SDN is the separation of the network control plane from the data plane. SDN enables a Network Operating System (NOS) [13,15] which interacts with packet forwarding elements. Control features and applications, including routing algorithms, can be deployed on top of the NOS and run as software modules. The NOS presents a consistent network-wide view to the centralized control logic running on top of it.

One of the most important SDN principles we can exploit to enhance routing is the separation of routing from routers and the logical centralization of the routing control plane. The benefits of these concepts have been shown in previous research, for example in [9], and intra-AS routing platforms based on these principles such as RCP [4] or RouteFlow [25] have been proposed. A centralized routing control plane within an AS is a feasible concept and is capable of simplifying routing management.

In this paper we leverage SDN concepts to improve inter-domain routing. We propose to outsource the routing control plane of an ISP to an external trusted provider, i.e., the service contractor¹. The contractor specializes in routing management and can relieve the ASes of the burden of maintaining expensive, highly-trained staff who manage the cumbersome routing complexity. Since a contractor manages routing for multiple ASes, it can take advantage of this multi-AS level of logical centralization and aggregation in order to improve inter-domain routing, i.e., ASes which are served by the same contractor enjoy more efficient routing. The bird's eye view which each contractor has over its client's domains establishes the logically centralized multi-AS routing plane as the natural place to take efficient routing decisions, detect policy conflicts, troubleshoot inter-domain routing problems and evolve BGP in a legacy-compatible fashion. As a consequence, the model benefits inter-domain routing, while the need to change the core of BGP is obliterated. In addition, each AS preserves its policy-shaping capability, privacy and business identity.

In summary, we make three contributions:

New routing model: We first propose a new routing model that is based on outsourcing the routing control logic of ASes and enables to reap the benefits of centralization beyond AS boundaries.

Incentives: We outline a number of technical and fi-

¹Throughout the paper, we use the term “contractor” to refer to the outsourcing provider, i.e., the one who provides outsourced IT services to multiple clients. The term “outsourcer” is used in order to refer to the outsourcing client, i.e., the one who outsources IT services to an external contractor.

ancial incentives that support the feasibility of the model and the value of our techno-economic framework. **Implementation plan:** We describe a proof of concept implementation plan that maintains legacy compatibility and highlight open research problems.

The rest of the paper is organized as follows. Section 2 presents the basic pillars on which our model is built and discusses technical incentives. In Section 3 we describe details for a proof of concept implementation. In Section 4 we outline economic considerations and in Section 5 we discuss related work and open questions that need to be addressed in future work. Finally, in Section 6 we conclude our paper.

2. A NEW ROUTING MODEL

In a nutshell, our model is based on three pillars, which are described in detail in the following sections:

1. Centralize: Logically centralize routing within ASes to simplify routing management and improve scalability.

2. Outsource: Outsource the routing logic to an external contractor.

3. Expand and evolve: Exploit cumulative outsourcing to centralize *beyond AS boundaries*. Further improve inter-domain routing efficiency and allow to evolve BGP, while preserving legacy compatibility.

2.1 Centralizing routing within an AS

The separation of the network control from the data plane and the consequent logical centralization of the routing control plane can drastically simplify routing management within an AS [4,9] and provide faster routing convergence [10].

The basic concept is that we centralize the control of routing decisions associated with an AS, both for internal and external destinations. By doing so, we gain multiple benefits. We can implement simpler policy-based routing since we are capable of centrally expressing, enforcing and checking routing policies. The policies can be expressed and compiled dynamically by frameworks such as [23]. Logical centralization also helps to improve the scalability of routing within an AS, through lowering the overall management complexity. Additionally, a central SDN control plane simplifies the modification of routing protocols as this can be achieved solely based on software. Therefore, we can have an evolving routing architecture since intra-domain routing protocols can change much more easily. It also becomes easier to clone the control plane to safely deploy configuration changes. Lastly, for load balancing and fail-over purposes we can have multiple routing controllers on the same network.

2.2 Next level: Outsourcing the routing control logic

The art of routing encompasses many more skills than the mere knowledge of how BGP or other routing protocols work. This includes the optimization of traffic flows via traffic engineering, correctly mapping Service Level Agreements (SLAs) to policies, coping with mis-configuration and scalability issues, while at the same time properly securing the network. Each of these oftentimes competing goals requires tuning several knobs in the routing protocol. Optimizing how packets are routed within an ISP to satisfy numerous operational and economic objectives is a difficult research problem. Although a number of advanced traffic engineering techniques have been proposed in the research literature, for example based on integer-programming and multi-commodity-flow optimization [27], network domains are run by practitioners who may not have the required knowledge at hand to optimize their network utilization through advanced traffic engineering or to improve security, e.g., through deploying sBGP [14]. Practitioners are often satisfied with a network that is just running. In addition, the router configuration code an ISP needs to develop, debug, and update is extensive, while the manual configuration of routers requires many administrator work-hours and is an error-prone process. Routing mis-configurations are common and can be very costly, e.g., several well-known incidents took significant parts of the Internet down.

To address these problems we propose that the routing control logic of networks with non-trivial size could be outsourced to a contractor that specializes in routing management, including routing optimization, configuration, troubleshooting, and monitoring. The contractor has extensive knowledge on routing and can therefore provide best routing policies tailored for the requirements of a client. The transition from today's domain-specific situation to an outsourced routing scheme can be handled smoothly by taking several steps. In a first stage, the contractor consults the outsourcer about best practices and together they arrive to a policy plan that satisfies the requirements of the client. Based on the agreed plan, the contractor takes over the handling of traffic and optimizes routing within the network of the client. For networks without intelligent traffic handling, this will have direct performance benefits, e.g., in terms of network load or other performance metrics. Besides, the contractor monitors how network traffic in the customer's network changes with time and enforces corrective traffic engineering actions.

The benefits of outsourcing are multiple, both of technical and economic nature. We first present the technical incentives and then discuss the financial ones in Section 4. Intelligent routing policies and routing optimizations enable to improve the reliability, performance, and

security of a network. Contractors who specialize in routing could therefore provide their knowledge and skills in the context of a new type of routing business relationship. From the perspective of a client ISP, outsourcing enables it to enjoy the benefits of advanced traffic engineering, consulting about best routing practices, policy reconciliation, and network troubleshooting.

The routing logic of the service contractor calculates the proper configuration of the control plane, updates the FIBs, RIBs and state of the network elements of the client AS, and also deals with inter-domain routing with the rest of the world through BGP. In particular, we export the following information to the contractor at the last stage of the outsourcing process.

Routing policies: These are policies of the outsourcer defined by the local administrators or derived based on requirements of SLAs of the outsourcer. They should be enforced and monitored by the contractor. Routing outsourcing does not impede on other services offered by a client AS that may depend on routing. This is because the enforced policies are specified by the client ISP during the consulting phase and take the requirements of offered services into account. In addition, the client ISP can regularly update its routing requirements.

Network's state and monitoring data: The outsourcer exports topology, configuration, and measurement data, e.g. network utilization, bandwidth allocation, loss, delay, jitter, etc. The contractor is a trusted third party that treats these data as well as routing policies confidentially. This information is never shared with other ASes or parties. The model of a trusted third party, although it requires trust, has been very successful in practice for many modern services. Also, SLAs can always specify the level of confidentiality and traffic visibility.

eBGP sessions: The contractor runs the control plane of the client AS and it therefore handles eBGP sessions with other ASes. BGP session messages are redirected from the border gateways of the client AS to a routing control platform operated by the contractor and vice versa.

2.3 Beyond AS boundaries: providing routing outsourcing services to many

Assume that we have formed a centralized and outsourced control plane running all routing-related processes within an AS. Still, the routing decisions are local for each AS and the problems with inter-domain routing at the global level remain. Therefore, we can exploit the benefits of centralization beyond AS boundaries. One of the most interesting aspects of outsourcing is that as more ASes choose the same contractor, AS-clusters are gradually formed. With the term *cluster*, we mean a

group of ASes which have chosen the same contractor, regardless of whether they have specific bilateral agreements with each other². The contractor has the bird’s eye view of the ASes under control. The advantages of a common contractor grow as the size of the clusters increases (horizontal scaling).

Hosting the routing control logic of many ASes technically benefits inter-domain routing in many ways, seeing that the contractor is able to:

1. Optimize inter-domain traffic engineering.

Outsourcing helps to optimize the handling of network traffic beyond AS boundaries. This is because the contractor is aware of the policies, topologies, and monitoring information for all ASes within the cluster. It is therefore the natural point at which inter-domain policy conflicts can be spotted and resolved and routing paths can be optimized. Coordination beyond AS boundaries can yield efficient paths even if ASes have different policies and optimization criteria [19]. This helps to improve routing stability and mitigate path inflation. This benefits *multiple ISPs*, even when they are not part of the cluster, as it will result in shorter and more stable end-to-end paths, thus reducing network load on a larger scale. Even in the case where the ASes within a cluster are not adjacent, the global view of the contractor plays an important role in routing optimization, since conflicting policies may exist beyond the visibility horizon of a single AS and its neighbors. In addition, the contractor can also easily check whether the SLAs between two neighbor ASes are respected, since it is responsible for enforcing the corresponding policies.

2. Evolve inter-domain routing. The contractor is free to adopt new inter-domain routing protocols between the members of its cluster. Innovation inside the cluster can be accelerated, while legacy interfaces with the rest of the Internet can guarantee proper interoperability. Generally, we can have lower convergence times and also decreased churn through centrally controlling the dynamics of intra-cluster routing [10]. Additionally, hierarchical routing, which is good for routing scalability, is implicitly enabled at the inter-AS level thus allowing hierarchical routing schemes to flourish.

3. Implement collaborative security and troubleshooting schemes. In case detailed monitoring data are also exported to the contractor, collaborative security and network troubleshooting are enabled through the mediation of the trusted party. For example, the contractor can easily pinpoint the source of a routing anomaly or failure by analyzing the information it acquires from its clients, and correlating it with other well-known sources [1, 2] for verification and validation purposes. Such benefits can only be leveraged when

aggregating information from multiple ASes, including detection of prefix hijacking [3].

3. IMPLEMENTATION PROPOSAL

In order to remain backwards compatible, the contractor uses legacy control sessions to directly access the FIBs and RIBs of network components. In particular, the routing control platform of the contractor can leverage iBGP sessions or CLIs to control legacy components in the context of SDN as described in [4] and [8], respectively. Legacy APIs [25] provide interoperability within hybrid networks which are gradually adopting OpenFlow [21]. The management of the complexity of this process is the contractor’s task. Compatibility with the rest of the (legacy) world is guaranteed by using the BGP protocol and API, i.e., by accepting and generating BGP messages and maintaining eBGP sessions.

The logically centralized multi-AS routing control platform we propose can be physically distributed or centralized. Clearly the location of the different components affects factors such as resiliency, delay, costs, privacy and transmission overhead. Although the exact functional details of the control platform are outside the scope of this paper a necessary element of the architecture is a local delegate *proxy-controller* as shown in Figure 1. This component is necessary to 1) mitigate the delay overhead between the network elements and the contractor’s platform and to 2) avoid multi-hop control sessions over the Internet and therefore increase resiliency to disruptions. The delegate communicates directly both with the contractor’s platform and the internal elements of the network. For the elements, it acts as the routing controller: it terminates all control sessions, like OpenFlow, Netconf, etc. For the contractor, it acts as a surrogate for the control and management of the network elements. Besides handling the diverse routing configuration channels, it delivers monitoring data, related to the client network’s state and utilization statistics. It is also the end-point of redirected eBGP sessions and relays BGP messages between the contractor and BGP gateways. The delegate can generally host parts of the routing logic of the contractor, acting as its stepping stone inside the client’s local infrastructure.

The contractor’s platform and the delegate communicate using an appropriate API and a secure communication protocol, over which the diverse channels between the client and the contractor are multiplexed. Multiple secure tunnels traversing diverse paths between the parties can be used for resiliency. In summary, the layout of our scheme is composed of the AS infrastructure, a local delegate controller and the remote routing control and management platform of the contractor. The scheme is shown in Figure 1.

Additionally, the contractor can host physically out-

²This term should not be confused with the term ”super-domains” which first appeared in the RFC 1478 and whose context is entirely different.

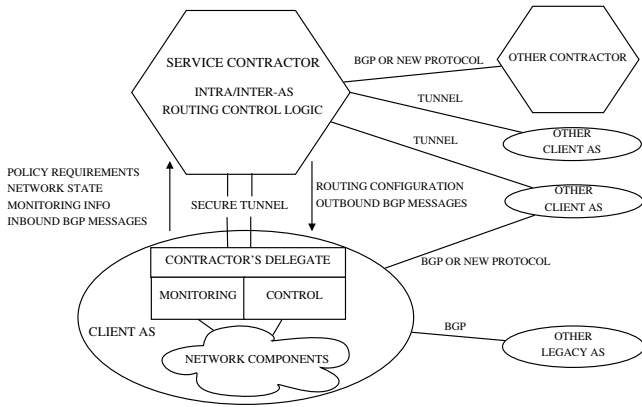


Figure 1: Implementation proposal

sourced control and management components in a private cloud infrastructure. In this way ASes can also enjoy the benefits of cloud computing, i.e., reduced costs, elasticity, scalability, multiplicity, availability and reliability for computing resources.

4. ECONOMIC CONSIDERATIONS

Operators have traditionally viewed the network as their core business. However, declining profit margins have put them under pressure to reduce costs and to launch new, higher margin services. This situation has also pushed operators to streamline their operating expenses (OPEX) and has given rise to an emerging market of managed services, in which the operation and maintenance of the network is outsourced to a third party. In this context, we propose a new model of network outsourcing, i.e., routing outsourcing, which enables the logical centralization of the routing control plane beyond AS boundaries as explained previously. Our model provides the following economic incentives:

- The contractor enjoys an opportunity for an economy of scale as the basic principles of routing optimization are the same across different networks. Economies of scale have been prolific in many computing contexts [18]. We claim that this also holds for routing.
- Outsourcing can reduce network-related OPEX, e.g., by 15% to 30% [6].
- ISPs agonize to diversify by introducing higher margin services in their portfolio. Outsourcing a low-margin service enables more effective use of human resources on higher priority services.

Outsourcing benefits mostly small and medium sized ISPs. Larger, e.g., tier-1, ISPs or third parties that have considerable local routing expertise can become contractors. This is supported by the present status quo.

AT&T, for example, is presently both a tier-1 network operator and among the market leaders in providing outsourced managed network services [6]. Therefore, we envision an ecosystem in which small and medium ISPs outsource their control plane to larger ISPs or third parties.

Smooth transition scheme: To ease the transition, an ISP can take several small steps. First, it can buy solely consulting services from a contractor to have its network optimized. Next, it can decide to hand direct low-level control of FIBs and RIBs to the contractor, including the step to centralize its routing. Last, it can decide to entirely outsource the routing plane, while remaining in control of the high-level policy. In each of these step, it can scrutinize the effects of the changes performed, both in terms of traffic and in terms of expenses, and step back in case it is not satisfied.

Tussle handling: The bird’s eye view enables the contractor to efficiently detect tussles [7] between its clients. The job of the contractor is to allow the tussles to unfold and run as today, based on the choices and policy requirements of each client. The main difference with today’s status quo is that the contractor can easily detect and mediate the resolution of routing problems, like dispute wheels [12], which may stem from these tussles. In addition, the bird’s eye view of the contractor can help find better solutions that meet the policies of each AS than when ASes act alone based on their limited local view. Moreover, new tussles between outsourcing service contractors are introduced. As the contractors start competing for clients, additional tussle dimensions arise, thus enabling a new game between the outsourcing entities.

5. DISCUSSION

Outsourcing network functions is not a new concept. Sherry *et al.* [26] propose the outsourcing of enterprise middlebox processing to the cloud. Gibb *et al.* [11] propose the outsourcing of network functionality to external feature providers. Both studies suggest the export of traffic which is routed in the data plane, while we focus on outsourcing the routing control logic. Lakshminarayanan *et al.* [17] introduce the notion of Routing-as-a-Service, motivated by the resolution of tussles between ISPs and customers over the control of end-to-end paths. The authors propose that ISPs should outsource route computation to customers in the context of virtual link stitching and end-to-end virtual path calculation across multiple domains. Instead, our work focuses on the benefits of outsourcing the entire per-AS routing logic and combining the outsourced inputs from multiple ASes for improving inter-domain routing.

The proposed model opens a number of interesting research questions. In our current work we explore

trade-offs involved in the design and implementation of a multi-AS routing control platform. For example, there is a trade-off between centralization on the one side and geographical proximity between the client and the contractor on the other side. While centralization has been proven useful, the logically centralized architecture should be geographically distributed to increase resiliency and reduce delay. This leads to a requirement of a hierarchical architecture. The main questions in our future research agenda are the following:

Fail-over: How can we keep a network operational even if the client-contractor communication fails?

Interfaces: What is the exact interface between a client network and the service contractor?

Security-Privacy: How can we guarantee security and privacy in the communication between the diverse parties? Which are the new attack vectors that are introduced in the context of the proposed scheme?

Evolved BGP: What other inter-domain communication protocols besides BGP can be used for inter-contractor or inter-client communication?

Market: What is the structure of the new market? Is it evolvable? Can it lead to a new form of multihoming, involving outsourcing the routing logic of different slices of a client network to different contractors?

Legal considerations: What kind of SLAs govern the relationship between the parties? Where do external auditors fit in?

6. CONCLUSION

The SDN paradigm leads us to exploring concepts such as routing logic centralization and new architectures with which we can enhance Internet routing. In this context, we investigate ways in which SDN principles can help perform such enhancements incrementally, while maintaining backwards compatibility. We propose a new routing model and techno-economic framework which introduces additional levels of aggregation in the Internet, based on the concept of a trusted party-mediator. The model holds without sacrificing the autonomy and privacy of ASes or requiring to change the legacy core. It is based on outsourcing and its feasibility is fueled by strong financial incentives. The true value of the proposed model will be empirically proven in follow-up work, since the purpose of this paper is the description of the key ideas, the justification of their potential and the presentation of the associated challenges.

7. ACKNOWLEDGEMENTS

We would like to thank the anonymous reviewers for their helpful feedback. The work presented in this paper is partially funded by the EU FP7 project OFELIA.

8. REFERENCES

- [1] <http://www.routeviews.org/>.
- [2] <http://www.nanog.org/>.
- [3] H. Ballani, P. Francis, and X. Zhang. A study of prefix hijacking and interception in the internet. In *Proc. of ACM SIGCOMM*, 2007.
- [4] M. Caesar, D. Caldwell, N. Feamster, J. Rexford, A. Shaikh, and J. van der Merwe. Design and implementation of a routing control platform. In *Proc. of NSDI*, 2005.
- [5] M. Caesar and J. Rexford. Bgp routing policies in isp networks. *IEEE Network Magazine*, Nov.-Dec. 2005.
- [6] R. Chaudhury and C. Terfloth. The lure of network outsourcing. White paper, Oliver Wyman, 2007.
- [7] D. D. Clark, J. Wroclawski, K. R. Sollins, and R. Braden. Tussle in cyberspace: defining tomorrow's internet. In *Proc. of ACM SIGCOMM*, 2002.
- [8] F. Farias, I. Carvalho, E. Cerqueira, A. Abelem, C. E. Rothenberg, and M. Stanton. Legacyflow: Bringing openflow to legacy network environments. http://changeofelia.info.ucl.ac.be/pmwiki/uploads/SummerSchool/Program/poster_004.pdf, 2011.
- [9] N. Feamster, H. Balakrishnan, J. Rexford, A. Shaikh, and J. van der Merwe. The case for separating routing from routers. In *Proc. of ACM SIGCOMM workshop on FDNA*, 2004.
- [10] J. Fu, P. Sjödín, and G. Karlsson. Intra-domain routing convergence with centralized control. *Computer Networks*, Dec. 2009.
- [11] G. Gibb, H. Zeng, and N. McKeown. Outsourcing network functionality. In *Proc. of ACM SIGCOMM workshop on HotSDN*, 2012.
- [12] T. G. Griffin and G. Wilfong. An analysis of bgp convergence properties. In *Proc. of ACM SIGCOMM*, 1999.
- [13] N. Gude, T. Koponen, J. Pettit, B. Pfaff, M. Casado, N. McKeown, and S. Shenker. Nox: towards an operating system for networks. *ACM SIGCOMM CCR*, Jul. 2008.
- [14] Kent, Lynn, and Seo. Secure border gateway protocol (s-bgp). *IEEE J. Selected Areas in Communications*, 2000.
- [15] T. Koponen, M. Casado, N. Gude, J. Stribling, L. Poutievski, M. Zhu, R. Ramanathan, Y. Iwata, H. Inoue, T. Hama, and S. Shenker. Onix: a distributed control platform for large-scale production networks. In *Proc. of OSDI*, 2010.
- [16] C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian. Delayed internet routing convergence. *IEEE/ACM Trans. Netw.*, Jun. 2001.
- [17] K. Lakshminarayanan, I. Stoica, and S. Shenker. Routing as a service. Tech. rep. ucb-cs-04-1327, UC Berkeley, 2004.
- [18] J.-N. Lee, M. Q. Huynh, R. C.-W. Kwok, and S.-M. Pi. It outsourcing evolution—: past, present, and future. *CACM*, May 2003.
- [19] R. Mahajan, D. Wetherall, and T. Anderson. Negotiation-based routing between neighboring isps. In *Proc. of NSDI*, 2005.
- [20] Z. M. Mao, R. Govindan, G. Varghese, and R. H. Katz. Route flap damping exacerbates internet routing convergence. In *Proc. of ACM SIGCOMM*, 2002.
- [21] N. McKeown, T. Anderson, H. Balakrishnan, G. Parulkar, L. Peterson, J. Rexford, S. Shenker, and J. Turner. Openflow: enabling innovation in campus networks. *ACM SIGCOMM CCR*, Mar. 2008.
- [22] D. McPherson, V. Gill, D. Walton, and A. Retana. Border gateway protocol (bgp) persistent route oscillation condition. IETF RFC 3345 (Informational), 2002.
- [23] C. Monsanto, N. Foster, R. Harrison, and D. Walker. A compiler and run-time system for network programming languages. In *Proc. of POPL*, 2012.
- [24] R. Oliveira, B. Zhang, D. Pei, R. Izhak-Ratzin, and L. Zhang. Quantifying path exploration in the internet. In *Proc. of ACM IMC*, 2006.
- [25] C. E. Rothenberg, M. R. Nascimento, M. R. Salvador, C. N. A. Corrêa, S. Cunha de Lucena, and R. Raszuk. Revisiting routing control platforms with the eyes and muscles of software-defined networking. In *Proc. of ACM SIGCOMM workshop on HotSDN*, 2012.
- [26] J. Sherry, S. Hasan, C. Scott, A. Krishnamurthy, S. Ratnasamy, and V. Sekar. Making middleboxes someone else's problem: Network processing as a cloud service. In *Proc. of ACM SIGCOMM*, 2012.
- [27] R. Teixeira, T. G. Griffin, M. G. C. Resende, and J. Rexford. Tie breaking: tunable interdomain egress selection. In *Proc. of CoNEXT*, 2005.
- [28] R. Teixeira, A. Shaikh, T. Griffin, and J. Rexford. Dynamics of hot-potato routing in ip networks. In *Proc. of ACM SIGMETRICS*, 2004.