Dynam-IX: a Dynamic Interconnection eXchange

Pedro Marcos
UFRGS and FURG

Marco Chiesa
KTH Royal Institute of Technology

Lucas Müller
UFRGS and CAIDA/UCSD

Pradeeban Kathiravelu
INESC-ID and
Université catholique de Louvain

Christoph Dietzel
DE-CIX / TU Berlin

Marco Canini
KAUST

Marinho Barcellos
UFRGS

CCS CONCEPTS

• **Networks** → **Network management**;

KEYWORDS

Peering, Internet eXchange Point

ACM Reference Format:

Pedro Marcos, Marco Chiesa, Lucas Müller, Pradeeban Kathiravelu, Christoph Dietzel, Marco Canini, and Marinho Barcellos. 2018. Dynam-IX: a Dynamic Interconnection eXchange. In *SIGCOMM Posters and Demos '18: ACM SIGCOMM 2018 Conference Posters and Demos, August 20–25, 2018, Budapest, Hungary.* ACM, New York, NY, USA, 3 pages. https://doi.org/10.1145/3234200.3234218

1 INTRODUCTION

Internet connectivity is changing [8, 11]: Autonomous Systems (ASes) can now reach hundreds of networks directly through Internet eXchange Points (IXPs) [1], while reducing latency, and improving traffic delivery performance and competitiveness [2]. Despite these benefits, any pair of ASes needs first to agree on exchanging traffic. Establishing interconnections is mostly a manual and lengthy process that is heavily influenced by personal relationships and brand image. As a result, ASes miss interconnection opportunities and prefer long-term agreements, at the expense of a potential mismatch between actual delivery performance and current Internet traffic dynamics. We posit that IXPs have a large unexplored potential to improve wide-area traffic delivery performance as they offer a rich path diversity that could be leveraged to enable responsiveness to traffic dynamics (e.g., traffic surges [3, 4, 16, 19], link failures [12]). Facilitating interconnection via IXPs poses two major challenges: (*i*) How to quickly negotiate an agreement? (*ii*) How to decide which networks can be trusted to route traffic? To address these challenges, we propose Dynam-IX. Differently from previous academic work [5, 20, 21] and industry initiatives [9, 13, 17, 18], our approach addresses both aspects while keeping the privacy of the interconnection policies, which is a concern for most network operators [7].

2 DYNAM-IX

Our goal is to improve wide-area traffic delivery performance by empowering operators to exploit the rich interconnection opportunities at IXPs quickly. To facilitate adoption, we design our approach to complement the existing practices, leading us to the following high-level requirements: *expressive interface*: an operator should be able to specify its interconnection policies, including the traditional interconnection models as well as future ones; *confidentiality*: no information considered private about an interconnection agreement (e.g., interconnection policies) should be leaked to unauthorized parties as operators are reluctant to sharing interconnection policy-related information with third parties [6]; *mechanism to build trust*: network operators should be able to identify partners deemed reliable systematically.

To achieve such requirements, we design Dynam-IX in a decentralized manner with four main components: (*i*) a protocol to automate the interconnection process; (*ii*) a legal framework to digitally handle contracts; (*iii*) an intent abstraction to specify interconnection policies; and (*iv*) a tamper-proof distributed ledger to enable ASes to build trust cooperatively. A potential design would be to rely on the IXP to offer a service to automate the interconnection process. To preserve confidential information, the service could be engineered to guarantee strong security properties (e.g., using secure multi-party computations [15] or trusted execution environments [10]). However, this raises the complexity of the

SIGCOMM Posters and Demos '18, August 20–25, 2018, Budapest, Hungary

ACM ISBN 978-1-4503-5915-3/18/08…$15.00
https://doi.org/10.1145/3234200.3234218

solution and incurs processing overheads. Instead, Dynam-IX achieves confidentiality by keeping all private information locally stored on the ASes. The ledger is distributed to prevent IXPs from influencing the interconnection decisions of their members. Figure 1(a) depicts an overview of Dynam-IX.

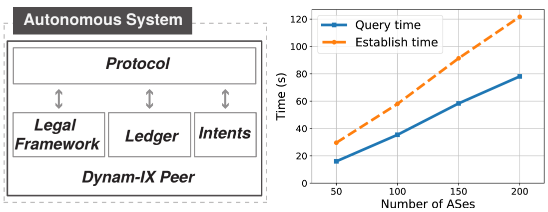

(a) Dynam-IX overview (b) Average response time

Figure 1: Dynam-IX overview and performance results.

Protocol. We define a protocol that resembles the current method for establishing interconnection agreements and works as follows. First, an AS willing to form an interconnection agreement queries the ledger to identify providers that may offer connectivity to the intended destination (i.e., the IP prefix). Then, the AS submits a request to each provider for interconnection proposals with specific desired interconnection properties (e.g., bandwidth, latency). When a provider receives a query for an interconnection offer, it decides whether to answer or not. The decision could be made automatically by an algorithm based on the provider's business policy or be delegated to a human. Then the AS selects one of the offers and sends an agreement proposal to the provider of the selected offer. The provider verifies that the proposal corresponds to a valid offer (each offer has an expiry date). Given a legitimate offer, the provider creates a contract using the legal framework, digitally signs it, and sends it to the customer. The customer verifies the provider's signature and the contract terms. If the signature is valid and the contract terms are as expected, the customer digitally signs the contract and sends it to the provider. In turn, the provider verifies the customer's signature and, assuming it is valid, proceeds to register the interconnection agreement on the ledger. The registration only includes public information about the agreement and it will be used to associate scores with valid agreements. Once the data is recorded in the distributed ledger, both ASes update their BGP configurations and start exchanging traffic. When an interconnection agreement expires, besides tearing down the BGP configuration, both ASes store on the ledger a score reflecting their experience. This score information is used as the basis to compute an overall reputation rank for each participating AS.

Interconnection intent abstraction. We define an *intent abstraction* as the relevant technical and business information associated with an interconnection offer. Each intent consists of a *target*, i.e., the traffic destination considered within the intent, and a set of attributes that carry information about the interconnection offer. These are divided into four categories: *routing*, Service Level Agreement (*SLA*), *pricing*, and *time*.

Legal framework. We mitigate the challenges regarding legal procedures by adopting a Legal Framework. It involves defining a *general contract template(s)* that is stored on the ledger, and digitally signed by every AS that joins Dynam-IX. The model contains standard clauses related to the interconnection agreement and empty fields to be completed with the specific properties and to be digitally signed by both ASes when an interconnection agreement is established.

Ledger. We use a tamper-proof distributed ledger to store information related to interconnection agreements to allow ASes to decide whether or not to interconnect instead of merely relying on personal relationships and brand recognition. When an agreement ends, each AS invokes a procedure to provide a score about the interconnection agreement.

3 PRELIMINARY EVALUATION

We built a prototype of Dynam-IX using Hyperledger Fabric 1.0.5 [14], a permissioned blockchain, as a distributed tamper-proof ledger. Using the prototype, we answer the question *how long does it take to establish an interconnection agreement?*

We measure the time to perform a query and the time to establish an agreement. The query time is the elapsed time between an AS sending a query to a potential provider and the response with an interconnection offer. The establishment time is measured from the moment an AS sends an interconnection proposal to the moment the agreement is established. We determine the limits of Dynam-IX with a throughput test: N ASes flood a single AS with queries and establishing interconnection agreements proposals. We evaluate this scenario using up to 200 AWS EC2 instances, each hosting a single AS. During the experiment, each instance executes the protocol 30 times at maximum rate. Figure 1(b) presents the average response times in the number of ASes.

Even with response times in the order of a few dozens of seconds, the average number of established agreements per second is 2.4 (50 ASes) and 1.4 (200 ASes), meaning that an AS can establish more than 80 interconnection agreements within a minute. While Dynam-IX performs well even under high loads, we observe that under more relaxed conditions it can establish a single agreement in less than 10 seconds.

4 SUMMARY AND FUTURE WORK

Dynam-IX is a framework to improve wide-area traffic delivery performance by allowing operators to exploit the rich connectivity opportunities at IXPs quickly, while achieving privacy. As future work, we plan to investigate the impact of Dynam-IX on storage requirements and network traffic, and to compare its performance using different ledgers.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their valuable feedback. This research is (in part) supported by European Union's Horizon 2020 research and innovation program under the ENDEAVOUR project (grant agreement 644960) and by the project Mapping Interconnection in the Internet: Colocation, Connectivity, and Congestion (NSF CNS-1414177 grant).

REFERENCES

[1] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger. Anatomy of a Large European IXP. In *SIGCOMM '12*, 2012.

[2] A. Ahmed, Z. Shafiq, H. Bedi, and A. Khakpour. Peering vs. Transit: Performance Comparison of Peering and Transit Interconnections. In *IEEE ICNP*, 2017.

[3] C. Arhtur. iOS 5 update causes massive internet traffic spike - to users' frustration, 2011. Available at https://www.theguardian.com/technology/2011/oct/13/ios-5-update-internet-traffic-spike.

[4] J. Brodkin. iOS 7 downloads consumed 20 percent of an ISP's traffic on release day, 2013. Available at https://arstechnica.com/information-technology/2013/11/ios-7-downloads-consumed-20-percent-of-an-isps-traffic-on-release-day/.

[5] I. Castro, A. Panda, B. Raghavan, S. Shenker, and S. Gorinsky. Route Bazaar: Automatic Interdomain Contract Negotiation. In *USENIX HotOS 2015*, 2015.

[6] M. Chiesa, D. Demmler, M. Canini, M. Schapira, and T. Schneider. Internet Routing Privacy Survey, 2017. Available at https://six-pack.bitbucket.io/media/privacy-survey-2017.pdf.

[7] M. Chiesa, D. Demmler, M. Canini, M. Schapira, and T. Schneider. SIXPACK: Securing Internet eXchange Points Against Curious onlooKers. In *CoNEXT '17*, 2017.

[8] Y.-C. Chiu, B. Schlinker, A. B. Radhakrishnan, E. Katz-Bassett, and R. Govindan. Are We One Hop Away from a Better Internet? In *IMC 2015*, 2015.

[9] Console. Console Connect - Interconnection made easy, 2018. Available at https://www.consoleconnect.com/.

[10] V. Costan and S. Devadas. Intel SGX Explained. Cryptology ePrint Archive, Report 2016/086, 2016. http://ia.cr/2016/086.

[11] A. Dhamdhere and C. Dovrolis. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *CoNEXT 2010*, 2010.

[12] S. Duncan. Australian internet slows to a crawl after undersea cable cut, 2017. Available at http://www.dailymail.co.uk/news/article-5146795/Aussie-internet-slows-crawl-undersea-cable-cut.html.

[13] Epsilon. Epsilon Telecommunications – Connectivity Made Simple, 2018. Available at www.epsilontel.com/.

[14] T. L. Foundation. Hyperledger Fabric, 2017. Available at https://www.hyperledger.org/projects/fabric.

[15] O. Goldreich, S. Micali, and A. Wigderson. How to Play any Mental Game or A Completeness Theorem for Protocols with Honest Majority. In *STOC'87*, pages 218–229. ACM, 1987.

[16] J. Mcgee-Abe. Apple devices behind DE-CIX Frankfurt 5.88Tbps data traffic rate, 2017. Available at http://www.capacitymedia.com/Article/3751343/Apple-devices-behind-DE-CIX-Frankfurt-588Tbps-data-traffic-rate.

[17] Megaport. Megaport - We make connectivity easy, 2018. Available at http://megaport.com/.

[18] PacketFabric. PacketFabric, 2018. Available at https://www.packetfabric.com/.

[19] Sandvine. FIFA 16 - The Beautiful Game?, 2015. Available at http://www.internetphenomena.com/2015/09/fifa-16-the-beautiful-game/.

[20] V. Valancius, N. Feamster, R. Johari, and V. Vazirani. MINT: A Market for INternet Transit. In *ReArch 2008*, 2008.

[21] T. Wolf, J. Griffioen, K. L. Calvert, R. Dutta, G. N. Rouskas, I. Baldin, and A. Nagurney. ChoiceNet: Toward an Economy Plane for the Internet. *SIGCOMM Comput. Commun. Rev.*, 2014.