

Investigating the Potential of the Inter-IXP Multigraph for the Provisioning of Guaranteed End-to-End Services

Vasileios Kotronis,
Rowan Klöti,
Panagiotis Georgopoulos,
Bernhard Ager
ETH Zurich, Switzerland

Matthias Rost,
Stefan Schmid
TU Berlin, Germany

Xenofontas Dimitropoulos
FORTH, Greece
ETH Zurich, Switzerland

ABSTRACT

In this work, we propose utilizing the rich connectivity between IXPs and ISPs for inter-domain path stitching, supervised by centralized QoS brokers. In this context, we highlight a novel abstraction of the Internet topology, i.e., the inter-IXP multigraph composed of IXPs and paths crossing the domains of their shared member ISPs. This can potentially serve as a dense Internet-wide substrate for provisioning guaranteed end-to-end (e2e) services with high path diversity and global IPv4 address space reach. We thus map the IXP multigraph, evaluate its potential, and introduce a rich algorithmic framework for path stitching on such graph structures.

Categories and Subject Descriptors

C.2.2 [Network Protocols]: Routing Protocols

Keywords

Internet Exchange Point; QoS; Embedding; EuroIX

1. INTRODUCTION

Modern Internet applications, from HD video-conferencing to telesurgery and remote control of power-plants, pose increasing demands on network latency, bandwidth and availability. Presently, ISPs are able to provide certain QoS guarantees only in intra-domain settings based on technologies such as leased circuits and VPN tunnels, e.g., over MPLS. Such services provide an important revenue stream for ISPs. However, despite several research and standardization efforts, providing QoS guarantees at the inter-domain level has seen very limited success so far.

Centralized inter-domain path brokers have been explored in the literature [8, 11, 12] and the industry [1], as an approach to support the requirements of demanding applications across domains. In these setups, ISPs provide QoS-enabled pathlets [6], which are stitched together by an inter-domain routing mediator, e.g., a bandwidth broker mediating the exchange of transit bandwidth between ISPs [12]. This work explores logically centralized inter-domain QoS routing mediators for the provisioning of inter-domain path




Figure 1: The CXP stitches QoS-enabled e2e paths.

guarantees in light of ongoing changes in the Internet ecosystem. In particular, (i) the Internet is becoming denser and more flat [3] due to richer interconnectivity at Internet eXchange Points (IXPs) [4], and (ii) there is a gradual paradigm shift towards network virtualization and Software Defined Networking (SDN), also in the context of IXPs [7]. Our main proposal is that a rich inter-IXP overlay fabric opens interesting traffic engineering flexibilities which can be exploited, e.g., for inter-domain QoS.

2. CXPS: PATH BROKERS OVER IXPS

We propose performing inter-domain QoS routing over a novel abstraction of the Internet topology, in which vertices are IXPs and edges are virtual links connecting two IXPs over an ISP. We call this abstraction the *IXP multigraph* because two IXPs can be connected with multiple edges over different ISPs. The choice of IXPs as switching points exploits their rich connectivity to geographically diverse ISPs, enabling high path diversity and global client reach with a deployment of only a few well-placed switching points as we show in Section 3. We call the routing brokers and controllers that are deployed over IXP multigraphs *Control eXchange Points (CXPs)*. CXPs stitch pathlets [6] across multiple administrative domains to construct global paths. The pathlets are provided by ISPs and are annotated with specific properties, such as bandwidth and latency guarantees. The incentive for ISPs to provide pathlets is the revenue generated by their use for e2e services. As shown in Fig. 1, the ISPs of the source and the destination offer access pathlets to connect to ISP-adjacent data plane anchors, i.e., programmable CXP switching points, while the intermediary ISPs offer transit pathlets between those points.

A CXP performs the following tasks: it (i) handles new requests for QoS-enabled paths (admission control), (ii) computes and sets up suitable paths (embeddings), (iii) monitors pathlet availability and compliance with QoS guarantees, and (iv) performs re-embedding, if required. A client negotiates her request directly with her access ISP, which selects a suitable CXP for establishing the inter-domain route

Permission to make digital or hard copies of part or all 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. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s). Copyright is held by the author/owner(s).

SIGMETRICS'15, June 15-19, 2015, Portland, OR, USA.

Author's copy.




Figure 2: CCDFs of key properties of the IXP multigraph

out of a set of available CXPs. The ISP forwards the client’s request to the chosen CXP, which in turn computes a suitable e2e path. The CXP reserves capacity on the selected pathlets and then configures the respective data plane anchors. Accordingly, the client’s ISP has to configure its network such that the quality sensitive traffic is sent via a pathlet to the correct data plane anchor. A CXP monitors the bandwidth, latency and availability of a path for the duration of the client’s reservation. If the client’s requirements are violated or a pathlet becomes unavailable, the CXP chooses and configures an alternate path for the affected part(s) of the traffic; this can even be a “hot-standby” backup path carrying traffic duplicates. Besides, the CXP may choose to better utilize the available pathlets by re-embedding paths and defragmenting the substrate resources. CXPs and ISPs may use MPLS, SDN or PCE mechanisms for the execution of their associated tasks, such as Traffic Engineering (TE).

3. POTENTIAL OF IXP MULTIGRAPHS

Motivated by the CXP concept, we analyze the properties of the resulting IXP multigraph by extracting ISP membership data for 229 IXPs from the EuroIX dataset [4]. We find 49k inter-connections of type: IXP_A, ISP_X, IXP_B . These form a dense multigraph, where a few pairs of IXPs are interconnected by hundreds of distinct ASes, each of which is in a position to offer one or multiple pathlets between each pair. In Fig. 2, we observe that the resulting pathlet-wise path diversity is one order of magnitude larger than the direct IXP connectivity, assuming that all the member ASes provide one pathlet each per IXP pair. In terms of AS-level path diversity, in our accompanying technical report [9] we show that the application of generalized inter-domain routing policies can lead to an increase of up to 29 times, as compared to classic valley-free practices [5]. More disjoint paths indicate potentially higher availability for demanding applications. In contrast, classic BGP-based routing is known to use a limited number of suboptimal inter-domain paths [10], as they normally cross up to one peering link in an IXP. Furthermore, in terms of IP address coverage, we discovered that even a small deployment (~ 5 IXP anchors) could directly cover a high fraction ($\sim 40\%$) of the Internet IPv4 address space. This increases to $\sim 91\%$ of announced addresses if we also consider the 1-hop customer cone of the IXP members, assuming access pathlets over the member ISPs. This allows an initial deployment of just a few IXPs to serve large sets of IPv4 addresses and enables incremental adoption. Further use of private peering points might selectively extend the graph and augment coverage, where required.

4. PATH STITCHING ALGORITHMS

Given the interesting properties of the IXP-based multigraph, we devised algorithms which leverage its flexibility in the context of CXP-like QoS mediators. The problem is essentially a path embedding problem: CXPs embed e2e paths on dense inter-IXP multigraphs, subject to bandwidth and latency constraints. The objective is to exploit the rich path diversity in order to maximize the number of concurrently embeddable routes. These routes are requested by the client IPv4 endpoints that connect to the multigraph

via their access ISPs. In our accompanying technical report [9], we formally introduce the e2e routing problem considered in this work as the QoS Multigraph Routing Problem (QMRP), together with an optimal offline formulation. This problem is complex for several reasons. (i) Requests from the large client base dynamically arrive over time in a non-predictable manner, necessitating the use of online algorithms. (ii) While a single suitable e2e path can be found in polynomial time, the IXP-based graph offers many choices and requires to carefully select which of the edges between two IXPs is used. In fact, in such multigraphs the number of available edges can be 3 orders of magnitude larger than the number of nodes (IXPs) [9]. (iii) The online selection of e2e paths should reflect multiple conflicting high-level objectives, namely accepting as many requests as possible, avoiding the usage of scarce low-latency, high-bandwidth links, and preventing resource fragmentation.

We thus propose a *sample-select* approach to heuristically tackle the QMRP in an online manner. In particular, given the NP-hardness of the *optimal path* calculations, we employ a sample-select process, where in the first stage, a set of *feasible* paths is sampled (i.e., generated) in polynomial time, and subsequently one of them is selected for the actual embedding. The techniques for the sampling process may range from modified Dijkstra-based approaches to random walks. Moreover, we propose a hybrid online-offline algorithm to support reconfigurations of pre-generated embeddings in order to accommodate further online requests. Using simulations based on the mapped IXP multigraph, we show that our algorithms scale to the sizes of the measured graphs and can serve diverse path request mixes. The full set of online, offline and hybrid algorithmic variants, together with the results and insights stemming from their application on the EuroIX-based multigraph, can be found in our technical report [9]. The software that we used is publicly available [2].

Acknowledgments. This work was partly funded by European Research Council Grant Agreement no. 338402.

5. REFERENCES

- [1] GEANT Bandwidth-on-Demand Provisioning Tool. http://geant3.archive.geant.net/service/autobahn/about_autoBAHN/Pages/About_autoBAHN.aspx.
- [2] Software for CXP-based Experimentation. https://bitbucket.org/vkotronis/cxp_experimentation.
- [3] DHAMDHARE, A., AND DOVROLIS, C. The Internet is Flat: Modeling the Transition from a Transit Hierarchy to a Peering Mesh. In *ACM CONEXT* (2010).
- [4] EUROIX. European Internet Exchange Association. <https://www.euro-ix.net/>. Dataset acquired on 2014-04-09.
- [5] GAO, L., AND REXFORD, J. Stable Internet Routing Without Global Coordination. In *ACM SIGMETRICS* (2000).
- [6] GODFREY, P. B., GANICHEV, I., SHENKER, S., AND STOICA, I. Pathlet Routing. In *ACM SIGCOMM* (2009).
- [7] GUPTA, A., VANBEVER, L., SHAHBAZ, M., DONOVAN, S. P., SCHLINKER, B., FEAMSTER, N., REXFORD, J., SHENKER, S., CLARK, R., AND KATZ-BASSETT, E. SDX: A Software Defined Internet Exchange. In *ACM SIGCOMM* (2014).
- [8] KOTRONIS, V., DIMITROPOULOS, X., AND AGER, B. Outsourcing the Routing Control Logic: Better Internet Routing Based on SDN Principles. In *ACM HotNets* (2012).
- [9] KOTRONIS, V. ET AL. Investigating the Potential of the Inter-IXP Multigraph for the Provisioning of Guaranteed End-to-End Services. Tech. Rep. 360, ETH Zurich, Laboratory TIK, Feb 2015. <ftp://ftp.tik.ee.ethz.ch/publications/TIK-Report-360.pdf>.
- [10] LUMEZANU, C., BADEN, R., SPRING, N., AND BHATTACHARJEE, B. Triangle Inequality and Routing Policy Violations in the Internet. In *PAM* (2009).
- [11] SPRINTSON, A., YANNUZZI, M., ORDA, A., AND MASIP-BRUI, X. Reliable Routing with QoS Guarantees for Multi-Domain IP/MPLS Networks. In *IEEE INFOCOM* (2007).
- [12] VALANCIUS, V., FEAMSTER, N., JOHARI, R., AND VAZIRANI, V. MINT: A Market for INternet Transit. In *ACM CONEXT* (2008).