Measurement-based Protocol Design

Gorry Fairhurst
University of Aberdeen, Scotland
Mirja Kühlewind
Networked Systems Group, ETH Zurich, Switzerland
Diego Lopez
Telefonica I+D, Spain

***Abstract*—The increasing public concerns about the interference of Internet traffic have led to a rapidly expanding deployment of encryption to protect end-user privacy, in protocols like QUIC. At the same time, network operators and access providers, especially in mobile networks, have come to rely on the in-network functionality provided by middleboxes both to enhance performance and support network operations. This presents a need for architectural changes and new approaches to the way network transport protocols are designed. This poster will explore the opportunities to use an experimentally-driven measurement-based approach to facilitate this network architecture evolution.**

I. Introduction

Middleboxes are prevalent in current generation mobile networks [1]. This can involve multiple layers of NAT, complex firewall policies, a range of performance enhancing proxies, and an assortment of methods to support mobile network operations. On the one hand, this raises questions about whether any new protocol header would be passed though a network path, and what the implications are of deploying new protocol headers. On the other hand, future Internet protocols (such as QUIC [2]) enabling large-scale encryption assists the restoration of the end-to-end nature of the Internet by returning complex processing to the endpoints. Middleboxes cannot modify what they cannot see.

At the same time, there has been renewed interest in methods that have the potential to significantly improve network performance, notably reducing latency, by introducing more explicit feedback between network equipment and endpoint devices. Active Queue Management (AQM) and Explicit Congestion Notification (ECN) therefore have the potential to replace some middlebox functions with mechanisms that can effectively work with encrypted traffic. Network Function Virtualization (NFV) also can increase the flexibility in the design and placement of functions that previously were located in dedicated middleboxes.

In current networks, operators have access to a body of information by examining the headers of network packets [**?**]. This enables operations staff to diagnose performance-related issues experienced by users, and to verify whether procedures to engineer solutions were successful in addressing any identified issues (e.g. to use TCP sequence numbers to measure the RTT across a path, or to examine if flows make progress in their transfers). The expanding deployment of encryption to protect end-user privacy will necessitate new approaches to operational support for users, since protocol headers would no longer be necessarily visible (among the options for design of QUIC, there is the possibility to encrypt all transport information, including sequence numbers, etc). This raises questions about how this impacts existing deployed infrastructure, what options may exist to design new protocols, and what form of operational support would need to be offered when these new protocols are deployed.

The need for in-network monitoring of performance is particularly important as new mobile architectures are being developed and new link technologies start to emerge (utilising higher frequency bands, opportunistic frequency use, etc) - resulting in a greater range of supported network capacity, but also greater variability in performance, and heterogeneity. Networks continue to be built on a mixture of deployed and new equipment designs, e.g., even as new network technologies are deployed, operators need to continue to be able to utilise existing deployed eNodeBs and other legacy equipment.

II. Measurement as part of the design Process

A common approach to protocol design has been to synthesise solutions from design requirements, with measurements often featuring towards the end of the design process, as a tool for evaluating the new design. Initial design can therefore often be based on assumptions about what mechanisms and protocol headers seem to be deployable, but given the wide range of complexity in the current Internet, these assumptions are often not fully understood.

In contrast QUIC has adopted a measurement-based approach, where instrumentation in prototype code has been used to inform decisions as the protocol evolves. Measurement lies at the heart of modern protocol development, and following this trend the authors previously developed measurement tools focussed on the core of the Internet (e.g., using the *PATHspider* [3]), to measure Internet path transparency. *PATHspider* is open source and publicly available on GitHub[1]. The current release, is also available as A Debian package.

In our current work we move the focus to use of a European-scale platform with multi-homing capabilities for measurements of a variety of operational 3G/4G Mobile Broadband (MBB) networks [4]. This enables experimentation with Internet protocol innovations on MBB networks. Endpoints (MONROE Nodes) are connected to up to three MBB providers, and often also to WiFi. Information about network, time and location for experiments, MONROE Nodes collect metadata from externally connected modems such as cell ID, signal strength and connection mode. Data from such experiments can help understand the key characteristics of the

[1] See https://pathspider.net/ for the *PATHspider* codebase and documentation

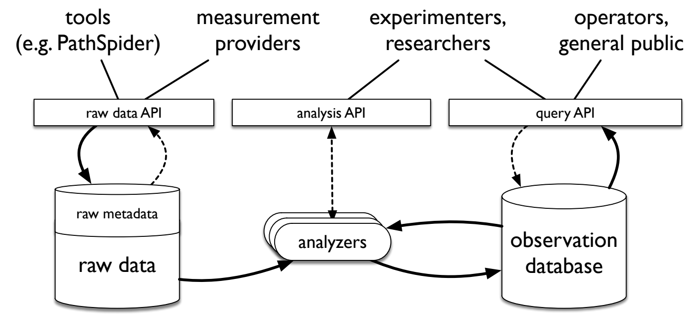

Fig. 1. Functions and dependencies of PTO components.

path, including any modification to packet by middleboxes. This complements more direct measurement of the path with tools such as Tracebox [5].

The Path Transparency Observator Path Transparency Observatory (PTO) [6] will be used as a repository for measurement data. PTO was developed by authors of this poster, for collection of large raw measurement data sets from diverse sources, tools, and measurement campaigns; and for the analysis of this data to derive *observations* on path transparency. An observation is an assertion that a given *condition* was observed on a given path at a given time; e.g. that ECN was successfully negotiated, or that an experimental TCP Fast Open (TFO) cookie was seen.

As illustrated in in Figure 1, the PTO maintains two data bases, one for unstructured raw measurement data, and one for structured observations. Analyzer modules derive observations from raw data in a given format. Beside an interface to upload raw data and run analysis for advanced users, the PTO also provides a graphical web front-end[2] for querying and aggregating observation data to answer questions such as "What proportion of observed web servers negotiate ECN?"

III. Protocol Design for Measurement

Availability of large scale measurement data enables a new approach to protocol design that can enable new techniques to be incrementally deployable by different actors (users, application developers, platform developers, equipment vendors, and network operators). Maps of middlebox manipulation within the Internet, such as those that could be derived from the PTO, will provide background for design decisions about protocol engineering and evolution.

Availability of measurement data throughout the design process is need to increase the likelihood that new protocols will be deployable across the entire Internet. Including the range of effects from middlebox manipulation on various packet headers is important because middleboxes in the Internet can have long deployment times. As next generation networks emerge, we need to continue to measure them. Indeed, the need to support in-network performance measurement needs to be explicitly designed into next generation networks that by default may be designed to encrypt all end-to-end protocol information.

In such a world, there is a question about what if any information should be made available to the network. This information exposure has to happen under explicit endpoint control. If immutable data is exposed to network devices (e.g. to measure the RTT), this needs to be protected from change within the network by end-to-end authentication. A design should follow the principle of least exposure, while balance with the advantages of making information available to, e.g., measurement tools: in each use case, this should define the minimum amount of information exposed by endpoints and middleboxes required by the proposed mechanism to solve the identified problem. And "trust-by-verify" should be applicable to all information exposed under the assumption that two endpoints have a trust relation for integrity protection and encryption but there is generally no requirement for an explicit trust relationship with network devices. The Path Layer UDP Substrate (PLUS) [7] proposes a framework for information exposure to the network. A range of new transport encrypted protocols can be designed over PLUS. At the same time, tools can emerge that utilise the exposed PLUS information to support measurements and diagnostics in a transport-protocol-independent way.

IV. Acknowledgments

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421, and was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0268. The opinions expressed and arguments employed reflect only the authors' views. The European Commission is not responsible for any use that may be made of that information. Further, the opinions expressed and arguments employed herein do not necessarily reflect the official views of the Swiss Government.

References

[1] Z. Wang, Z. Qian, Q. Xu, Z. M. Mao, and M. Zhang, "An untold story of middleboxes in cellular networks," in *ACM SIGCOMM*, 2011.

[2] R. Hamilton, J. Iyengar, I. Swett, and A.Wilk, "QUIC: A UDP-Based Multiplexed and Secure Transport," IETF, Internet-Draft draft-hamilton-quic-transport-protocol-00, Jul. 2016.

[3] G. Fairhurst, "The impact of transport header encryption on operation and evolution of the internet," Working Draft, IETF Secretariat, Internet-Draft draft-fairhurst-tsvwg-transport-encrypt-01, June 2017. [Online]. Available: https://www.ietf.org/id/draft-fairhurst-tsvwg-transport-encrypt-01.txt

[4] I. Learmonth, B. Trammell, M. Kühlewind, and G. Fairhurst, "PATHspider: A tool for active measurement of path transparency," in *First ACM/IRTF Applied Networking Research Workshop*, Berlin, Germany, Jul 2016.

[5] O. Alay, A. Lutu, D. Ros, R. Garcia, V. Mancuso, A. F. Hansen, A. Brunstrom, and M. A. M. H. Lonsethagen, "MONROE: Measuring mobile broadband networks in europe," in *IRTF & ISOC Workshop on Research and Applications of Internet Measurements (RAIM)*, 2015.

[6] G. Detal, B. Hesmans, O. Bonaventure, Y. Vanaubel, and B. Donnet, "Revealing Middlebox Interference with Tracebox," in *Proceedings of the 2013 Conference on Internet Measurement Conference*, ser. IMC '13. Barcelona, Spain: ACM, 2013, pp. 1–8. [Online]. Available: http://doi.acm.org/10.1145/2504730.2504757

[7] S. Neuhaus, K. Edeline, B. Donnet, and E. Gubser, "Towards an observatory for network transparency research," in *Proceedings of the Applied Networking Research Workshop (ANRW'16)*. ACM, 2016. [Online]. Available: https://mami-project.eu/wp-content/uploads/2015/10/anrw16-final2.pdfhttps://irtf.org/anrw/2016/anrw16-final2.pdf

[2] currently in alpha testing; a public access URL will appear in the camera-ready, and is available from the authors upon request.