Bringing the Optical Layer of the Internet into Focus for Measurement-based Study

Paul Barford University of Wisconsin – Madison

Measurement-based research enables understanding of how Internet users, systems, protocols and configurations behave in practice, and paves the way for improvements in performance, scalability, manageability, functionality, robustness and security. While measurement-based study of the Internet has focused on different layers of the network protocol stack, the optical layer[1] -- including connectivity, configurations, and dynamics -- has received relatively little attention.

There are several reasons for the dearth of empirical studies of the optical layer of the Internet. First, optical layer issues are typically addressed by an entirely separate community of researchers. This separation of communities is exemplified by the relatively minimal treatment of the optical layer in networking textbooks commonly used in computer science departments. Second, unlike higher layers of the protocol stack, measurement at the optical layer in most cases requires special permission from network operators, who may not be willing or able to collaborate for security or competitive reasons. While a number of enlightening, empirical studies of optical layer characteristics and dynamics have been published in the past, these often come from researchers at companies or through partnerships that are not feasible for the broader community.

I argue that the availability of optical layer measurement data from the wide area Internet would enable a holistic, full-stack understanding of connectivity and behavior and could lead to discoveries and innovations that would otherwise not be possible. With this concept in mind, a myriad of research questions are possible including, how are optical transport devices typically deployed and configured? what are the characteristics of standard optical signals (Q-factor, BER, etc.) and how do these correlate with phenomena such as attacks and outages? are certain types of higher layer traffic (*e.g.,* attacks or services) evident in the optical layer? how do configuration changes in the optical layer impact higher layers of the protocol stack? how can robustness of networks or services be improved through optical layer configurations? how can models developed from empirical studies inform development of new optical protocols and programming paradigms?

To address these (and other similar) research questions, two general types of data are required. The first is data on the deployment and configuration of optical transport devices. Deployment information would have to be provided by operators (these are often illustrated in complex connectivity diagrams), while configuration specifics are available through the Command Line Interfaces (CLIs). The second is data on the dynamic state of optical transport devices, which is typically available through Management Information Bases (MIBs) that can be quired via SNMP. Similar to other types of passive Internet measurements, depending on the level of participation by operators, potentially vast amounts of data could be collected. This would require a large (potentially distributed) storage infrastructure and techniques and tools for analysis that scale similarly.

The primary challenge in bringing the optical layer into focus for empirical research, as noted above, is in enlisting network operators to provide data from their infrastructures, which is typically considered sensitive. I am not aware of "best practices" in dealing with optical layer data from service providers, however the IMC '16 paper by Ghodabi and Mahajan offers a potential framework, which includes obfuscation of certain configuration and measurement details. It has been announced that the FABRIC testbed will "collect measurements and error statistics from the underlying optical/DWDM infrastructure", which is a great step, but a broader set of data from commercial networks would provide a more representative perspective.

The NSF can facilitate empirical research at the optical layer by supporting development of infrastructure for long term collection and sharing of optical layer data in the wide area, and supporting research programs that will benefit directly from that data.

---

[1] For this whitepaper, "optical layer" refers to optical transport equipment connected by fiber links in the wide area Internet. While this is a subset of the physical layer of the protocol stack, optical connectivity carries most Internet traffic on a daily basis.