

## Monitoring and Managing QoE in the Wide Area: Programmable Networks to the Rescue?

Nick Feamster, *Princeton University*

Interdomain routing plays an important role in the quality of experience that users ultimately experience. The paths that traffic takes “from content to eyeballs” often involves multiple independently operated networks. Given the role that wide-area networks play in the ultimate performance of QoE, we have surprisingly poor tools for both understanding how wide-area networks influence QoE and mechanisms that we might use at these interconnection points that would ultimately help us better manage traffic at these locations. The state-of-the-art mechanisms that users have to locate these problems involve executing end-to-end performance measurements and drawing (often indirect) inferences about locations along each path that might be responsible for introducing problems.

Part of the difficulty stems from the outdated, baroque technology that is deployed in these networks, particularly at interconnection points. In many cases, the best information that these networks can offer comes from sampled IPFIX/NetFlow/sFlow records and SNMP interface counters, neither of which provide useful information about the QoE that users might be experiencing or what the root causes of poor QoE might be. Worse yet, the mechanisms that current networking technology offers to help operators manage this traffic—essentially BGP-based traffic engineering—is coarse-grained (i.e., only on the granularity of IP prefix) and indirect (i.e., only using BGP’s clumsy knobs), making it exceedingly difficult to fix any problems with QoE, even if they could be detected in the first place.

The advent of programmable networking may offer some hope to both monitoring and management of QoE in the wide area. The remainder of this white paper describes two such possibilities: (1) How the deployment of SDN at IXPs might facilitate traffic management practices that improve QoE; (2) How the emergence of in-band network telemetry may ultimately make it easier to detect and locate portions of the network that ultimately degrade QoE. Although significant questions remain, I believe that our ongoing research on SDN-based Internet Exchange Points (SDX) and on programmable hardware may improve the current state of affairs for monitoring and managing QoE in the wide area.

### Flexible, Application-Specific Traffic Management

The future looks brighter: My work with colleagues (and lead Ph.D. student Arpit Gupta) on a [Software-Defined Internet Exchange Point \(SDX\)](#) makes it possible for network operators to control traffic on much finer granularity. Specifically, the SDX enables the following functions:

- *More direct inbound traffic control.* Whereas BGP makes it difficult for a network to control how traffic reaches its network (particularly from remote locations), the SDX makes it possible for a network to exercise direct control over how traffic enters its network.
- *Ability to route traffic for certain applications along different paths.* The SDX makes it possible to exercise fine-grained traffic control for traffic in different parts of flowspace, making it possible to route traffic destined for one port along a completely different path than the default BGP path. Such capabilities make it possible to implement more fine-grained mechanisms and policies, including peering for specific applications.

We have recently released a prototype of SDX that is based on the Ryu SDN controller that incorporates some of this basic functionality. Yet, challenges remain: Specifically, performing traffic management for specific applications in the middle of the network requires a mechanism to identify the application that the traffic belongs to. Such application classification may be increasingly difficult as applications begin to deploy end-to-end encryption more pervasively. To take full advantage of the mechanisms offered by the SDX in the wide-area, we need to develop mechanisms that can map application traffic to portions of flowspace that these switches can act upon.

Another opportunity for traffic management involves the deployment of fair queuing and prioritization mechanisms at Internet Exchange Points and interconnection points. These types of mechanisms have been deployed at other locations in the network to improve QoE for certain application traffic (e.g., [in the home router](#), at the [CMTS](#)), but in many cases these deployment locations are too far downstream to have

any appreciable effect on application performance. I believe that it is worth exploring whether these types of mechanisms could be deployed at interconnection points such as the SDX, and how these types of fair queuing mechanisms might be deployed.

In summary, our work thus far on SDX has primarily involved “plumbing”—the implementation of the SDX functionality, and the mechanisms to make it scale to a large number of participants, routes, and traffic flows. I believe that a potentially compelling use case for this type of technology concerns opportunities for traffic management that will ultimately improve QoE for applications that depend on more advanced traffic management mechanisms in the wide-area than BGP can offer.

### **Localizing QoE Impediments in the Wide Area**

Another challenge involves pinpointing the causes of poor performance. Often the home wireless network is the cause for poor performance, and [our ongoing research](#) (in collaboration with the FCC’s Measuring Broadband America program) has developed a mechanism that can help shed light on the extent of bottlenecks in home wireless networks. Yet, pinpointing performance problems in the wide area remains a significant challenge. Part of the difficulty is that, because each portion of the network is independently operated, it can be difficult to obtain measurements that shed light on performance problems that may exist in other parts of the network.

The state of the art for localizing problems in wide-area networks often involves end-to-end measurements that traverse disjoint paths, a process often referred to as “network tomography”. Unfortunately, these techniques are generally inaccurate, and they often require large numbers of vantage points to draw any meaningful conclusions. One of the problems, of course, is that the probe traffic that tomography generates must be synchronized in certain ways, and the probe traffic that is used to draw conclusions about performance often bears little resemblance to the performance that users ultimately experience with real traffic.

Ongoing work on [in-band network telemetry](#) (INT), however, may enable remote networks to draw more meaningful conclusions about the locations of performance problems that ultimately affect QoE. Briefly, in-band network telemetry involves network elements augmenting data packets with certain state (e.g., timestamps, buffer occupancy levels), for the flow in question, at which point this data is ultimately transported to the endpoint. This type of information, if it becomes available in modern chipsets, might make it possible to obtain extremely fine-grained timing and congestion information about path segments in the network. Of course, INT alone may not shed light on exact locations of problems that affect QoE, particularly because a congested network may not be willing to reveal such fine-grained information about the performance of individual links and routers inside its own network. Yet, if INT is deployed at Internet Exchange Points (e.g., at SDXes), it may be possible for a network to obtain timing information at multiple points along an end-to-end path—possibly even on either side of a congested transit network. Important research remains to better understand how these emerging technologies might shed light on QoE—either directly, or in conjunction with other existing methods.

As a side note, it is possible that mechanisms such as INT may be useful not only the the wide area but also in customer premises equipment (CPE). For example, precise timing information at this point in the path, in conjunction with similar information at interconnection points, may help users and ISPs better determine whether impediments to QoE are in the home network, access link, or further afield.

### **Coupling Wide-Area Metrics with QoE Metrics**

All of the mechanisms mentioned above, of course, involve network monitoring and management using metrics that do not correspond directly to application QoE. A critical piece to this problem is thus determining how network metrics that can be measured in the wide area (or in the home network) ultimately map to QoE metrics that correspond to how users ultimately experience application performance. Various studies have been performed to this end, much work remains. For example, we do not have a good understanding of how network-level metrics along a path might correlate with video rebuffering or Web page load times. As our ability to instrument the network to obtain fine-grained information at multiple points along an end-to-end

path become more likely, new possibilities for mapping these features to end-to-end QoE may also emerge. Our ongoing research is looking at how such network-level features in home networks can map to QoE-relevant metrics (e.g., video buffering), but the bigger picture might eventually incorporate the wide area, once wide-area instrumentation becomes possible.