

Quality of experience must be measured from end-to-end to establish ground truth

Peter Boothe

September 1, 2015

Quality of experience is an end-to-end property. If any piece of the relevant network parts is not working, then the quality of experience quickly drops to zero. End-to-end properties must be measured from end-to-end, because network modeling and the development of techniques for combining piecewise network measurements has always lagged network reality.

Therefore, in order to measure quality of experience on an evolving Internet, we must measure in an end-to-end manner. It is only by measuring from end-to-end that we can establish ground truth.

Quality of Experience is End to End

When we speak of quality of experience, we must ask “Quality of experience for whom?” We can’t simply mean that the local link works well for the most technologically proficient, but that the applications are successfully delivered in a way that pleases the users of the network. Users of network-delivered services have a high-quality experience if every aspect of the network path works well, along with the client and server¹ at either end. It is possible for any piece in the middle to degrade the experience and cause a low-quality experience, which means any assessment of “quality of experience” needs to take into account as much of the end-to-end path as possible.

¹ Both client and server here may be multiple entities as in the case of CDNs or distributed protocols, but even in that case the service delivered is end-to-end for a nuanced definition of “end”.

Abstraction Works, so Measurement is Hard

Because abstraction was an Internet design goal ahead of measurement², the abstraction boundaries on the Internet are measurement-hostile. As traffic transits networks, each network’s behavior must be regarded as a black box, because there is no way of knowing what underlying technology within each network is being used to forward along the IP packets³. This is a good thing, and it is by design, but it does make it a theoretical impossibility to model other networks’ properties in the abstract.

² Abstraction and measurement fall under goals 2 and 7, respectively, in *The Design Philosophy of the DARPA Internet Protocols* by Clark in SIGCOMM CCR in 1988.

³ There are frequent attempts at network telemetry to break down these abstraction boundaries. The fact that this remains an active area of research indicates that these abstraction boundaries are quite robust.

When faced with a theoretical impossibility, we in computer science know what to do: route around the problem. In this case we can route around the impossibility of prediction by measuring the network and seeing what the resulting performance is instead of trying to predict it.

Piecewise Measurements are Fundamentally Brittle

Piecewise network measurements can only be combined in a coherent way into an end-to-end assessment to the degree that the model used for combining those pieces is reflective of reality. Models break when new network technologies get deployed, and new models take a long time to develop. Even in the case of TCP, it was deployed and working for many years before a model was developed of its performance in terms of network characteristics⁴. Without a model that accurately predicts the effects of combining network properties of the systems that lie along a communication path, measured piecewise properties can not be combined into an overall prediction.

As long as the network is continuing to evolve, all network models will remain brittle and will continue to lag behind network reality. When network models are brittle, combining piecewise measurements can, without warning, yield an overall performance prediction that bears no relation to reality. In the presence of such uncertainties, the only reliable way of assessing how well a service is being delivered is to either instrument the service, or to have a measurement system which, as much as possible, does what the service does.

End to end measurements of network performance provide a ground truth from which models can be developed for piecewise measurement. Piecewise measurement can never provide an authoritative ground truth until the network stops evolving. In the presence of an evolving Internet, measuring end-to-end properties requires end-to-end measurements.

Summary

Quality of experience can be degraded by any network element in the service path, and as such is an end-to-end property. On an evolving Internet, end-to-end properties can only be measured in an end-to-end manner. Therefore quality of experience must be measured end-to-end. This end-to-end measurement may be augmented by piecewise measurements, but it is the end-to-end measurement that establishes ground truth.

⁴ Most influentially in *The Macroscopic Behavior of the TCP Congestion Avoidance Algorithm* by Mathis et al. in SIGCOMM CCR in 1997, although that paper does cite some earlier efforts.