

## A Case for Long-Term Statistics

Michael Welzl

The common model of network monitoring, where a user or administrator interactively runs a test and views the result, is probably not the most suitable approach to derive and present useful metrics to end users. In the face of network problems, most end users only have long term solutions at their disposal, such as: changing the Internet Service Provider (ISP); buying new equipment (e.g., a WiFi Access Point, or end user device), or updating its firmware; changing or configuring the Operating System of end user devices. It makes sense to support such long-term measures with long-term statistics, in the order of hours or days.

Consider the following (real) example: in a residential building, Internet access is generally relatively fast and stable. The apartment owners association starts a debate about switching to an ISP which offers fiber-to-the-home. Arguments brought forward in this debate include “fiber is faster”, as well as “a friend has this ISP, and his network appears to be less stable than ours”. So, will the switch really be worthwhile? Understanding this would require to know:

- When the rates of transfers exceed the path’s capacity, how often, on average, are they bottlenecked at the access link of users in the building? (separate for uplink/downlink)
- How much overall gain would it be to remove this bottleneck: how often do transfers experience a bottleneck at all? E.g., web transfers often finish early, and long-term transfers such as video streams are typically bursty and do not behave like a “greedy” flow.

As another example, with interactive real-time media, it would often be useful to understand where quality fluctuations come from. Sudden lag or interruptions are common with video conferencing applications or in online games, but what causes them? Is it a bottleneck on the user’s side, at one of the peers, or at the Application Service Provider’s server?

Again, a long-term measurement (in this case, over the course of the application usage: minutes, not hours or days) could help – it would be useful to be able to consult a tool after a conversation and understand what happened. With consent of the user, this tool could, for example, passively listen on a Wi-Fi network during a conversation and determine metrics such as physical rate changes, signal quality problems and transport-layer retransmissions. From these data, it could derive a general statement about whether the local network (Wi-Fi and the modem up- and downlink) is likely the problem or not. To some degree, such longer-term statistics are already available – e.g., in the “Wireless Diagnostics” application on OS X – but they are not typically associated with applications.