

Making latency matter

23rd June 2013
Toke Høiland-Jørgensen
Roskilde University

Fixing latency issues is important if networked applications are to become even more prevalent than they are today. However, current internet connections are not always providing a quality of experience sufficient to make internet communication and online applications really seamless: interruptions and delays are prevalent, to the point that people expect this to be the norm. At the same time, latency is often not even mentioned when discussing internet service; latency doesn't *matter*.

To a very large extent, the technology exists to fix bufferbloat and the accompanying latency under load issues. While there's still no perfect solution that can be deployed across all connection technologies, what we have available from a technical standpoint (in the form of mainly better queue management algorithms), is orders of magnitude better than the current status quo. Hence deployment is now the mayor hurdle to overcome to dramatically improve users' general internet experience.

However, the perception that poor latency characteristics are the norm, is a barrier to deployment of better technology. And this is exacerbated by service and equipment providers using bandwidth as the main (or even the only) figure in marketing material and even in technical specifications. Hence, creating demand for better technology requires users to become aware both that a better experience is possible, and that it is within reach. Getting to that point requires a multi-pronged approach, and I believe the following elements can help work towards that goal:

1. Establishing a test for latency behaviour (including the presence of bufferbloat) that is easily accessible to users. Currently there exists testing tools to evaluate latency under load, but they are research tools more suitable for experts or, at best, technically inclined users. Coming up with an accessible test (and the infrastructure to support it), as well as a way to effectively convey the results (e.g. a "bufferbloat score") would go a long way towards making this issue visible to users.
2. Making latency behaviour a visible metric in marketing material and benchmarking of both equipment and internet connection services. Internet customers currently have no way of knowing the latency characteristics of the products they purchase, and thus no way of picking the better-performing one. Latency having a prominent place will help visibility and adoption, and can be pursued by influencing both vendors and regulatory bodies.
3. Having application providers focus on latency in user-facing support and marketing messages can help visibility of the issue and is a good way to make users aware of the impact of poor latency characteristics. This will be even more effective if the application vendors can point to a well-established benchmark to test the user's connection (see point 1 above).

While the above three points are by no means exhaustive, I believe they constitute a starting point towards making latency matter in the minds of users, and hence create the demand for deployment of technology with better latency characteristics.