Latency signals

Piers O'Hanlon

1 Outline

It is clear that to provide for low latency transport we can no longer rely on loss as the primary signal for use in congestion control. However, it is still an important signal to indicate that the network path has been overloaded. We can no longer rely on loss as the main signal since it is only elicited once the buffers on the congested link have filled. The size of such buffers is unknown and if Bufferbloated can lead to excessive delay. By low latency we mean that the end-to-end delay of the flow is maintained close to the minimum RTT of the path, whilst achieving reasonable path utilisation.

To provide for low latency transport between end systems, over the existing infrastructure, we need to develop techniques that allow us to have a good notion of the latency incurred by a flow on the network, and how it changes over time. This latency state information can then allow a congestion control algorithm to set an appropriate send rate.

To allow for some measure of fairness between flows they each need to have a common procurement and interpretation of this state information. Thus each flow should be able to compete in an equal manner with other such flows by using this common interpretation by its congestion control system (as is the case now for loss).

Furthermore to provide for adequate competition with loss-based approaches, such as TCP, it should be possible to relate this state information to TCP's state. Thus the congestion control algorithm needs to assimilate both latency and loss state information so it may be possible to relate it to TCP's behaviour.

In general an adequate response in the face of a loss-based flow means that it needs maintain a reasonable share of the capacity of the path. To achieve this end the congestion control algorithm needs to be able to interpret loss signals in a similar fashion to the loss-based flow and choose a suitably competitive rate. However a key aspect of this functionality is being able to fall back to delay-based operation once a competing loss-based flow has subsided. Another aspect of this dual behaviour is that it is not generally possible for the congestion control system to maintain low delay whilst in direct competition with a loss-based flow due to the fact that loss-based flow will fill the buffers.

Typically such state information may be obtained through measurement of the send and receive times of the actual packets in the flow, though it is also possible that another entity [5] makes such information available. There have been a number of surveys [2, 4] of the various delay based approaches used over the last decade or two. These algorithms start with CARD in 1989, Tri-S and Dual in 1991 and 1992, and proceed through to Vegas in 1994 and FAST in 2003, then on to LEDBAT in 2009.

The majority of the approaches utilise a measurement of the end-to-end delay a compare it to the current or filtered previous version of that delay. Unfortunately it is quite hard to make an reliable assessment of the actual path delay which has lead to a number of issues with some of these protocols. Such problems have been highlighted in the case of LEDBAT [3]. One of the problems is that the algorithms assume that a local minima or maxima is reliable in the face of constant background flows, but unfortunately it can be lead to anomalous behaviour. However some of the original approaches and a number of the newer approaches, such as CAIA Delay Gradient (CDG), Google RRTCC, and most recently Sprout, take a more relative approach and utilise delay gradient information which has less issues with constant background traffic levels providing for more a dynamic approach. Despite these new approaches there still seem to be some outstanding issues as has been highlighted for Google's RRTCC [1].

Thus it seems there are still challenges in finding the right combination of latency and loss information to achieve the suitable behaviour on the Internet today.

References

[1] Luca De Cicco, G. Carlucci, and Saverio Mascolo. An experimental investigation of the google congestion control for real-time flows. In *ACM SIGCOMM 2013 Workshop on Future Human-Centric Multimedia Networking*, 2013.

[2] Sushant Rewaskar, Jasleen Kaur, and Don Smith. Why don't delay-based congestion estimators work in the real-world. Technical report, University of North Carolina at Chapel Hill, 2005.

[3] D. Ros and M. Welzl. Assessing ledbat's delay impact. *Communications Letters, IEEE*, 17(5):1044–1047, 2013.

[4] D. Ros and M. Welzl. Less-than-best-effort service: A survey of end-to-end approaches. *Communications Surveys Tutorials, IEEE*, 15(2):898–908, 2013.

[5] M. Welzl. Coupled congestion control for rtp media. INTERNET-DRAFT draft-welzl-rmcat-coupled-cc, January 2013.