

Technical Report
Managing a Queue to a Soft Delay Target

Bob Briscoe*

15 Apr 2019

Abstract

This memo proposes to transplant the core idea of Curvy RED, the softened delay target, into AQMs that are better designed to deal with dynamics, such as PIE or PI2, but that suffer from the weakness of a fixed delay target.

CCS Concepts

•Networks → Cross-layer protocols; Network algorithms; Network dynamics;

Keywords

Data Communication, Networks, Internet, Control, Congestion Control, Quality of Service, Performance, Latency, Responsiveness, Dynamics, Algorithm, Standards, Active Queue Management, AQM, Signalling, Sojourn time, Queue delay, Service time, wait time, Expectation, Estimation, Explicit Congestion Notification, ECN

1 Introduction

All the well-known modern AQMs aim for a constant target queueing delay, e.g. CoDel [NJ12] including fq-CoDel, PIE [PPP⁺13], PI² [DSBTB16], DualPI² [DSBEBT17], and a recent variant of DCTCP's AQM [BCCW16, BCC⁺16]. In the technical report “Insights from Curvy RED” [Bri15] it was proved that the level of loss needed to induce TCP-based load to keep to a fixed delay target has to rise to unacceptably high levels during periods of increased load.

The report makes the point that the time taken to repair losses is itself a source of delay, particularly for short flows. Therefore, it is perverse to hold down queueing delay at the expense of very high loss levels.

*research@bobbriscoe.net,

[Bri15] proposes an adaptation of the RED algorithm [FJ93] called Curvy RED that uses a convex function of queueing delay as a target. This softens (increases) the delay target as load intensifies.

However RED, and by extension Curvy RED, provides no control over queue dynamics, whereas control theoretic AQMs do. This means that during dynamic load excursions, RED and Curvy RED have little control over how much delay overshoots (or undershoots) while trying to bring it back to the target. This allows delay to vary uncontrollably above the target. This is a significant problem because many applications are sensitive to maximum, not average, delay.

This memo proposes to transplant the core idea of Curvy RED, the softened delay target, into AQMs that are better designed to deal with dynamics, such as PIE or PI2, but that suffer from the weakness of a fixed delay target.

It should be emphasized that this combination is only useful when loss is used as the signalling mechanism. By extension that means this combination would also be used with classic ECN [RFB01], which requires any ECN behaviour to be equivalent to loss behaviour. However, this combination would be unnecessary for use with L4S ECN [DSBET17], which is not constrained to be equivalent to loss.

2 Curvy PI²

A Proportional Integral (PI) controller alters the congestion signalling probability dependent on both the distance from the target delay (the error) and the rate of change of the queueing delay. Therefore, it is able to rapidly control load excursions before they cause too much variation in load. Specifically, it uses an control equation of the form:

$$p'(t) = p'(t-T) + \alpha(q(t) - Q_0) + \beta(q(t) - q(t-T)),$$

where $p'(t)$ is the drop probability at time t , $q(t)$ is the queueing delay at time t , α and β are the gain




Figure 1: Soft delay target as a function of p' & p

constants, T is the sampling period and Q_0 is the (constant) target delay.

It would be straightforward to make the target delay a function $Q()$ of the current drop probability $p(t-T)$ rather than a constant, for instance, picking a reasonable formula fairly arbitrarily, one might use:

$$Q(p') = Q_0 + Q_1 p'(t - T), \quad (1)$$

where Q_0 and Q_1 are the min and max values of the soft delay target (at $p' = 0$ and $p' = 1$).

The PI² controller squares the resulting value of $p'(t)$ to determine the drop probability $p = (p')^2$ (see [DSBTB16] for why). Therefore Equation 2 is equivalent to:

$$Q(p) = Q_0 + Q_1 \sqrt{p(t - T)}. \quad (2)$$

Equations 1 and 2 are illustrated in Figure 1

In practice, rather than using the arbitrary formula in Equation 1, it will be possible to determine the optimum compromise between queuing delay and loss from human factors experiments that record the mean opinion score of a 2-D matrix of these two impairments for the popular application that is most sensitive to both, e.g. voice, or perhaps virtual reality (although MOS data is more readily available for voice). Then it should be possible to fit an approximate curve to the contour of optimum pairs that will be amenable to implementation as the soft delay target function.

It might seem of concern that the loss probability depends on a delay target function, which in turn depends on the loss probability, which seems like a circular dependency. However, the target function depends on the loss probability that was output in the previous sampling period (and the second dependency will be much weaker than the first anyway).

It might seem contrary to the goal of a PI controller to allow high load to increase delay. However, there is nothing sacred about the constant delay goal that was first proposed by Hollot *et al* in 2001 [HMTG01a], before designing a solution in the same year [HMTG01b]. A controller can aim for any target that meets human needs, it does not have to be a constant.

3 Variants

By extension, a delay target that itself depends on the level of loss could be used in other, non-control-theoretic AQMs such as CoDel. As before, the intent would be to soften the delay target under high load, so as not to drive loss to extreme levels in pursuit of low queuing delay, given repairing loss itself introduces delay.

In the case of CoDel, the variable called `target` would need to depend on the variable `drop_next`, which determines the interval between drops.

References

- [BCC⁺16] Wei Bai, Kai Chen, Li Chen, Changhoon Kim, and Haitao Wu. Enabling ECN over Generic Packet Scheduling. In *Proc. Int'l Conf Emerging Networking EXperiments and Technologies*, CoNEXT '16, pages 191–204, New York, NY, USA, 2016. ACM.
- [BCCW16] Wei Bai, Li Chen, Kai Chen, and Haitao Wu. Enabling ECN in Multi-Service Multi-Queue Data Centers. In *13th USENIX Symposium on Networked Systems Design and Implementation (NSDI 16)*, pages 537–549, Santa Clara, CA, March 2016. USENIX Association.
- [Bri15] Bob Briscoe. Insights from Curvy RED (Random Early Detection). Technical report TR-TUB8-2015-003, BT, May 2015.
- [DSBEBT17] Koen De Schepper, Bob Briscoe (Ed.), Olga Bondarenko, and Ing-Jyh Tsang. DualQ Coupled AQM for Low Latency, Low Loss and Scalable Throughput. Internet Draft draft-ietf-tsvwg-aqm-dualq-coupled-01, Internet Engineering Task Force, July 2017. (Work in Progress).
- [DSBET17] Koen De Schepper, Bob Briscoe (Ed.), and Ing-Jyh Tsang. Identifying Modified Explicit Congestion Notification (ECN) Semantics for Ultra-Low Queuing Delay. Internet Draft

- draft-ietf-tsvwg-ecn-l4s-id-00, Internet Engineering Task Force, May 2017. (Work in Progress).
- [DSBTB16] Koen De Schepper, Olga Bondarenko, Ing-Jyh Tsang, and Bob Briscoe. PI^2 : A Linearized AQM for both Classic and Scalable TCP. In *Proc. ACM CoNEXT 2016*, pages 105–119, New York, NY, USA, December 2016. ACM.
- [FJ93] Sally Floyd and Van Jacobson. Random Early Detection Gateways for Congestion Avoidance. *IEEE/ACM Transactions on Networking*, 1(4):397–413, August 1993.
- [HMTG01a] C. V. Hollot, Vishal Misra, Donald F. Towsley, and Weibo Gong. A Control Theoretic Analysis of RED. In *Proc. INFOCOM 2001. 20th Annual Joint Conf. of the IEEE Computer and Communications Societies.*, volume 3, pages 1510–19, 2001.
- [HMTG01b] C. V. Hollot, Vishal Misra, Donald F. Towsley, and Weibo Gong. On Designing Improved Controllers for AQM Routers Supporting TCP Flows. In *Proc. INFOCOM 2001. 20th Annual Joint Conf. of the IEEE Computer and Communications Societies.*, pages 1726–1734, 2001.
- [NJ12] Kathleen Nichols and Van Jacobson. Controlling Queue Delay. *ACM Queue*, 10(5), May 2012.
- [PPP+13] Rong Pan, Preethi Natarajan Chiara Piglione, Mythili Prabhu, Vijay Subramanian, Fred Baker, and Bill Ver Steeg. PIE: A Lightweight Control Scheme To Address the Bufferbloat Problem. In *High Performance Switching and Routing (HPSR'13)*. IEEE, 2013.
- [RFB01] K. K. Ramakrishnan, Sally Floyd, and David Black. The Addition of Explicit Congestion Notification (ECN) to IP. Request for Comments 3168, RFC Editor, September 2001.

Document history

Version	Date	Author	Details of change
01	07 Sep 2017	Bob Briscoe	First complete version.
02	15 Apr 2019	Bob Briscoe	Added abstract and keywords.