Aggregate Rate Control for TCP Traffic Management

Jae Chung and Mark Claypool
Computer Science Department
Worcester Polytechnic Institute, USA
{goos|claypool}@cs.wpi.edu

TCP, the de-facto Internet transport protocol, has a successful end-host congestion control mechanism that has largely been effective in managing Internet congestion. Yet, with TCP alone it is hard to achieve efficient congestion control mainly due to the inefficiency of congestion estimation at the end-host. Active queue management (AQM) promises to overcome the limitations of end-host only congestion control by providing congestion feedback information before router buffers overflow.

The most promising approaches view AQM as a feedback controller on a time-delayed response system and apply control engineering principles to design an efficient controller for TCP traffic [2, 3, 4]. In modern control systems, proportional integral derivative (PID) designs dominate due to their simplicity and effectiveness. Without exception, this applies to recent active queue management developments and has altered AQM research from basic framework design into detailed controller design and practical implementation issues.

Among PID principles, only the proportional integral (PI) feedback control approach is primarily considered for AQM since the effect of the derivative control is often insignificant under practical Internet environments. While the PI control approach seems promising, a critical deployment challenge is the configuration of PI control parameters in a time-delayed feedback system, i.e., the Internet. There are no simple and effective PI control parameter configuration available for time-delayed system [5]. The existing PI control-based AQM mechanisms such as the PI controller [3] or Adaptive Virtual Queue (AVQ) [4] lack complete configuration guidelines, making their practical deployment difficult.

We propose a practical, rate-based AQM mechanism offering aggregated rate control (ARC) for TCP traffic. ARC is a reduced parameter PI controller, founded on classical control theory and a sound understanding of PI behavior for the Internet traffic control domain, offering easy configuration while keeping to proven system stability characteristics. We model a TCP-ARC feedback control system using a linear TCP model [3] and develop practical yet effective ARC configuration guidelines. The guidelines cover issues in choosing a target stable boundary system for ARC configuration and provide a method for selecting control parameters that help avoid system instability even when the system is out of the stability boundary. The guidelines also address the effects of the rate sampling interval on system stability, a consideration often neglected in other AQM studies.

Controllers with the same underlying principle design can result in noticeably different implementations, both in terms of complexity and performance depending on the way in which control information is obtained and feedback is processed. AQM requires information on the incoming traffic load to make accurate congestion control decisions, information which can be obtained in two different ways: derivation of queue samples or incoming traffic rate over the service rate. ARC takes a rate-based control information acquisition approach. For a small amount of data collection overhead, rate-based data acquisition reduces sampling noise that can significantly degrade the accuracy of congestion measurement. In addition, rate-based mechanisms can more effectively react to impending congestion, making control decisions before outbound queue buildup and can also be tuned to minimize queuing delay for a small loss in link utilization, allowing enhanced support for the quality of service (QoS) needs of various Internet applications.

Through an extensive simulation study, we have evaluated ARC and similar AQM mechanisms including the PI controller [3], AVQ [4] and SFC [2], and drop-tail queue management over a wide range of network and traffic conditions including Web flash crowd and multiple bottleneck cases. Our simulations show that ARC effectively handled network congestion in all the tested traffic conditions, outperforming other mechanisms in terms of queuing delay, link utilization, packet loss, and object response time for Web traffic. For details, see [1].

REFERENCES

[1] J. Chung and M. Claypool. Aggregate Rate Control for Efficient and Practical Congestion Managment. Technical Report WPI-CS-TR-04-03, CS Department, Worcester Polytechnic Institute, Feb. 2004. ftp://ftp.cs.wpi.edu/pub/techreports/pdf/04-03.pdf.

[2] Y. Gao and J. Hou. A State Feedback Control Approach to Stabilizing Queues for ECN-Enabled TCP Connections. In *Proceedings of IEEE INFOCOM*, San Francisco, CA, USA, Apr. 2003.

[3] C. V. Hollot, V. Misra, D. Towsley, and W.-B. Gong. On Designing Improved Controllers for AQM Routers Supporting TCP Flows. In *Proceedings of IEEE INFOCOM*, pages 1726–1734, 2001.

[4] S. Kunniyur and R. Srikant. Analysis and Design of an Adaptive Virtual Queue (AVQ) Algorithm for Active Queue Management. In *Proceedings of ACM SIGCOMM*, San Diego, CA, USA, August 2001.

[5] G. Silva, A. Datta, and S. P. Bhattacharyya. PI Stabilization of first-order Systems with time delay. *Automatica*, December 2001.