On the smoothing of traffic flows using ON-OFF control

Yong Xiong, Jyh-Charn Liu[+], Wei Zhao
Computer Science Department
Texas A&M University
{yongx,liu,zhao}@cs.tamu.edu

Abstract

This extended abstract reports our on-going work on how to control and ensure smoothness of traffic flows using the sliding mode control theory. The preliminary results suggest that our approach is highly promising, and more results are expected in the near future.

1 General Increase-Decrease Flow Control Model

A key issue in control of packet flows is the *increase-decrease* (I-D) control system model, A network system can be described by

$$\dot{x} = f(x,t) + g(x,t)u, \quad (1)$$

where t is the time variable, $x \in \mathrm{R}^n$, $u \in \mathrm{R}$, f and g is a smooth function. x stands for the state variable of the network dynamics, and u the control input. Eq. (1) defines a very general case. y is a function of x,

$$y = h(x,t). \quad (2)$$

The I-D controllers have three states, *increase*, *decrease* and *maintaining*. Figure 1 shows the general I-D control system model. The overall I-D control rule is defined by:

$$u = \begin{cases} u_-, & if\ y > y_2 \\ u_+, & if\ y < y_1 \\ u_0, & else \end{cases}, \quad (3)$$

where t is the time variable; u is the control input; u_-, u_+ and u_0 are the decrease, increase and maintaining rules. y is the output, y_1 and y_2 are the thresholds of the output. The output y is a function of the network state variables. The network state variables could be the queue length, the changing rate of the queue length and so on. Our objective is to maintain the output y within the interval $[y_1, y_2]$.

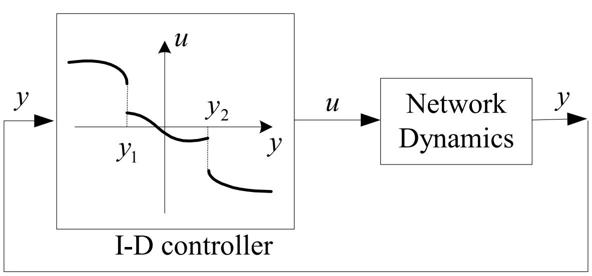

Figure 1: Network system with a general I-D controller.

Designing an I-D controller involves three aspects: 1) choosing the output function $h(x, t)$; 2) setting the thresholds y_1 and y_2; and 3) designing the decrease, increase and maintaining rules, u_-, u_+ and u_0.

I-D control schemes are widely used in communication network congestion control. The most famous example is the *additive-increase-multiplicative-decrease* (AIMD) congestion avoidance scheme of the *Transport Control Protocol* (TCP). Other examples are the "binary feedback" scheme proposed by Ramakrishnan and Jain, binomial congestion control, rate-based AIMD (e.g. the congestion control schemes in RAP [6], LDA+ [8], LTRC [9], etc.) and XON/XOFF flow control introduced in IEEE 802.3x [10] and so on. It is easy to map these algorithms into the general I-D control model.

2 Sliding Mode Control

Asymptotic stability is the basic requirement for any control system. When oscillation occurs, the system has only *boundary* stability. Despite the common belief that I-D control system is inherently stable, the fact is that an I-D control system may be unstable or even divergent. The idea of our method is that we first design the increase and decrease rules (u_+ and u_- in Eq. (3)) to make the system monotonously converge to a subset of the state space—the maintaining state—within a limited time, and then hold in it. The second step is, we design the maintaining rule (u_0 in Eq. (3)) to make the system asymptotically converge to its equilibrium point. By "monotonously converge" we mean the distance to the maintaining state monotonously decreasing. We define

$$S = \begin{cases} y - y_2, & if\ y > y_2 \\ y_1 - y, & if\ y < y_1 \\ 0, & else \end{cases}. \quad (4)$$

and the I-D control law can be rewritten as

$$u = \begin{cases} u_-, & if\ S > 0 \\ u_+, & if\ S < 0 \\ u_0, & else \end{cases}. \quad (5)$$

We call S the switching function and the absolute value of S is the distance to the maintaining state. Monotonously decreasing of $|S|$ means when $S>0$, $\dot{S}<0$ and when $S<0$, $\dot{S}>0$; or equivalently $S\dot{S} < 0$. Furthermore, when inequality (6) is satisfied, the system will reach the maintaining state defined by the condition $S = 0$ within a finite time $\frac{|S_0|}{\eta}$, where S_0 is the initial value of S. η determines the convergence rate to the switching manifold [5].

$$S\dot{S} \le -\eta|S|, \text{ where } \eta>0. \quad (6)$$

[†] Correspondence author

The stability within the maintaining state can be studied by using continuous system method. First, the relative degree of the input-output system defined by (1) and (2) must be one; and second, its *zero dynamics* is asymptotically stable, or equivalently, it is minimum phase. The relative degree is one guarantees that the control input u explicitly appear in the expression of $\dot{y}$. Thus, it is possible to change the sign of $\dot{y}$ through switching the control input.

The zero dynamics of a linear system is defined by its *zeros*. Asymptotical stability of zero dynamics guarantees that the system dynamics in the maintaining state is stable. For the example given in the beginning of this section, it is easy to verify that, when $y = x_1$, it is minimum phase and its relative degree two. Under the limited control (**Error! Reference source not found.**, the attractive zone of the maintaining state is also limited. A divergent maintaining state will finally go beyond the attractive zone. When the system has non-zero delays, measurement errors, or finite control frequency, large control values will aggravate the degree of oscillation caused by these factors, and may even lead to a divergence situation.

3 XON/XOFF flow control

XON/XOFF flow control strategy is introduced in the standard IEEE 802.3x. It is a hop-by-hop flow control, also known as backpressure flow control, implemented in the data-link layer of the OSI reference model. In an IEEE 802.3x network, before buffer gets full, a congested switch temporarily prevents input switches from sending packets to avoid packet dropping. The XOFF flow control message is sent to the upstream node when the buffer exceeds the upper threshold. When a switch receives an XOFF signal, it pauses sending packets until it receives an XON signal from the same switch, or until the XOFF message expires. In our later discussion, we ignore the case the time in XOFF messages expires. The XON signal is triggered when the buffer at the congested switch descends below the lower threshold.

To apply the I-D control model in Eqs.(1)~(3) to the XON/XOFF flow control scheme, y is the queue length, the control input u is the input rate, y_1 and y_2 are the lower and upper thresholds of queue length, u_+ is the capacity of the input link and u_- is zero. u_0 equals to u_- can if the previous flow control message is a XOFF message; otherwise u_0 equals to u_+.

The XON/XOFF scheme is essentially a black-box control system, where the network dynamics is considered to be in a black-box. Such a simple mechanism improves network performance in some situation, but leads to poor performance in others [16]. It responds to congestion much faster than end-to-end schemes. While TCP's congestion control algorithm is highly robust to diverse network conditions, its implicit assumption of end-system cooperation results in a well-known vulnerability to attack by high-rate non-responsive flows.

An important concern about the XON/XOFF flow control scheme is that it makes traffic more bursty. This situation is even more serious when the capacity mismatch between en-route links is large. Figure 3 shows the numerical simulation results of the network topology in Figure 2.

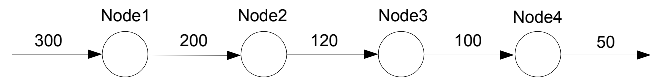

Figure 2: Network topology for simulation

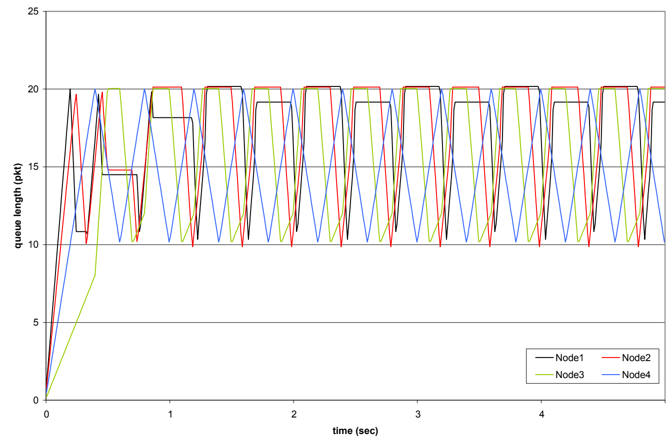

(a) queue length vs. time

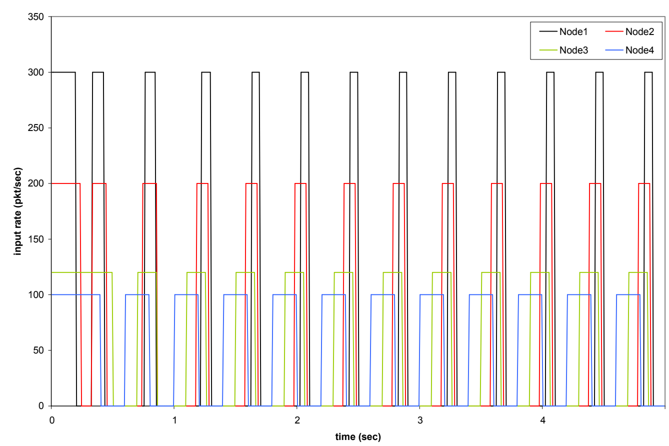

(b) input rate vs. time

Figure 3: Numerical simulation result of XON/XOFF flow control in MAC layer

4 Smooth the Traffic Flow

In this section, we propose a method to improve the performance of XON/XOFF flow control mechanism. In order to smooth the traffic rate, we use the 1st order time differential of the input rate instead of the input rate as the control input. We call our scheme XI/XD/XH scheme. "I" means increase, "D" means decrease, "H" means hold. The XI/XD scheme can be described as follows: XI/XD congestion control is invoked by triggering XI/XD flow control messages. When the queue length plus its 1st order time differential is larger than an upper threshold, the switch sends a XD message to the upstream neighbors; when less than a lower threshold, a XI message;

otherwise, a XH message. When a XD (or XI or XH) message is received, a switch decreases (or increase or holds) its sending rate. Figure 4 shows the numerical simulation results. The simulation results suggest that within a finite (short) time period, one can control the queue length to a stable queue length. This means that one can keep the queue length of a switch running the XON-XOFF scheme within a stable range, so that the queuing time at each switch can maintain a constant. By controlling the queuing delays of connected switches at constant levels, one can scale up constant delay transmissions across large local area networks.

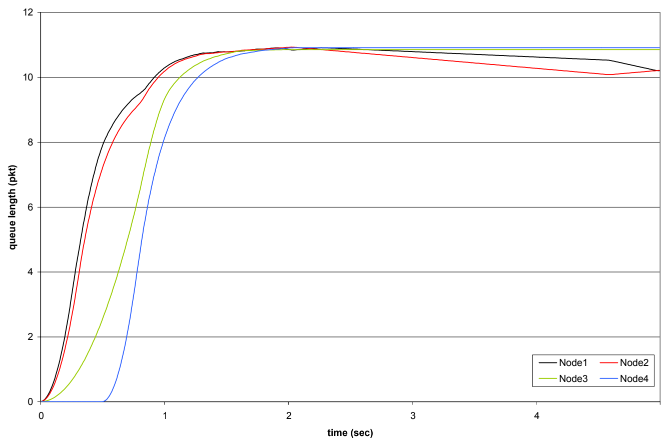

(a) queue length vs time

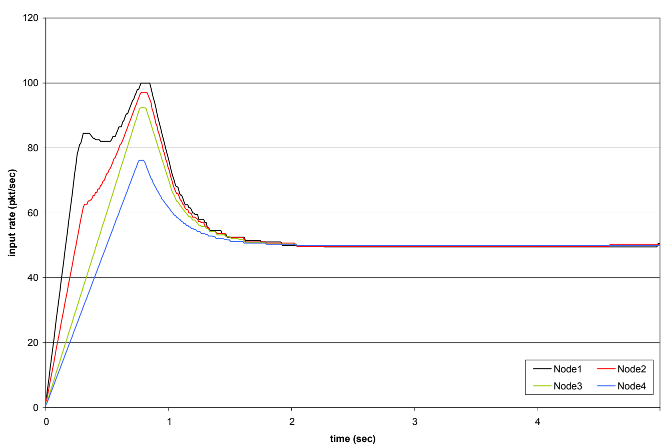

(b) input rate vs. time

Figure 4: Numerical simulation result of XI/XD flow control in MAC layer

5 Future Work

The initial results presented in this work suggest that properly designed feedback control algorithms can be employed on existing congestion control schemes with substantial performance improvements. At the current time, we are working on larger scale simulation and experiments to perfect the control rules. More results are expected in the near futures.

6 REFERENCES

[1] D. Bansal, H. Balakrishnan, “Binomial Congestion Control Algorithms,” *IEEE INFOCOM 2001*.

[2] K.K. Ramakrishnan, Raj Jain, “A Binary Feedback Scheme for Congestion Avoidance in Computer Networks with a Connectionless Network Layer,” *ACM SIGCOMM 1988*.

[3] V. Jacobson, “Congestion avoidance and control,” in *ACM SIGCOMM*, 1988.

[4] V.I. Utkin, “Sliding Modes in Control Optimization,” Springer-Verlag Berlin, Heidelberg 1992.

[5] J. E. Slotine, W. Li, “*Applied Nonlinear Control”*, Prentice Hall, 1991.

[6] Reza Rejaie, Mark Handley, D. Estrin, “RAP: An End-to-end Rate-based Congestion Control Mechanism for Realtime Streams in the Internet,” *IEEE INFOCOM 1999*.

[7] J. Widmer, R. Denda, and M. Mauve, "A survey on TCP-friendly congestion control," *IEEE Network*, vol. 15, no. 3, pp. 28--37, 2001.

[8] D. Sisalem, A. Wolisz, “LDA+ TCP-friendly adaptation: A measurement and comparison study,” in *Proc. International Workshop on Network and Operating Systems Support for Digital Audio and video* (*NOSSDAV*), June 2000.

[9] T. Montgomery, “A loss tolerant rate controller for reliable multicast,” Tech. Rep. NASA-IVV-97-011, West Virginia University, August 1997.

[10] “IEEE 802.3x, Specification for 802.3 Full Duplex Operation”, in *IEEE Standard 802.3*, 1998.

[11] W. Noureddine, F. Tobagi, “Selective back-pressure in switched Ethernet LANs,” *Proceedings of IEEE GLOBECOM*, pp. 1256-1263, 1999.

[12] Duke Lee, Sinem Coleri, Xuanming Dong, Mustafa Ergen, “FLORAX Flow-rate based hop by hop backpressure control for IEEE 802.3x,”

[13] Vincent Guffens, Georges Bastin, “Hop-by-hop congestion control with token buckets in feedback: compartmental analysis and experimental validation with UML,”

[14] C. Ozveren, R. Simcoe, G. Varghese, “Reliable and efficient hop-by-hop flow control,”

[15] K. Yoshigoe, K. Christensen, “Rate control for bandwidth allocated services in IEEE 802.3 Ethernet,” *IEEE 26th Conference on Local Computer Networks*, pages 446-453, 2001.

[16] J. Wechta, A. Eberlein, F. Halsall, “The interaction of the TCP flow control procedure in end nodes on the proposed flow control mechanism for use in IEEE 802.3 switches,” in *HPN*, pages 515-534, 1998.

[17] A. Koubaa, A. Jarraya, Y.Q. Song, “SBM protocol for providing real-time QoS in Ethernet LANs,”