Optimal Control of Queues in Computer Networks

Mahbub Hassan
School of Computer Science and Engineering
The University of New South Wales, Sydney 2052, Australia.
Email: mahbub@cse.unsw.edu.au
Harsha Sirisena
Department of Electrical and Electronic Engineering
University of Canterbury, Christchurch, New Zealand
Email: sirisehr@elec.canterbury.ac.nz

Abstract— **The design of rate allocation and queue length control in computer networks is treated as a stochastic optimal control problem. The performance index is chosen to achieve the twin objectives of minimising queue length fluctuations and fully utilising the available bandwidth. Simple, practically realisable optimal control schemes are obtained for both LANs and WANs. An adaptive scheme is proposed where the auto-regressive parameters of the traffic, needed for gain calculations, are estimated by an LMS algorithm. Discrete event simulations are carried out to verify the fluid-flow models used in developing the controllers, to compare their performance against PI controllers proposed previously, and to study the effect of self-similar traffic. Two key results are obtained. First, queue-length fluctuations, and hence potentially packet losses, are much smaller for the optimal feedforward controller than for the PI controller. Second, in contrast to uncontrolled queues, the queue length variance decreases with increasing Hurst parameter for self-similar traffic.**

I. Introduction

Network queues can be controlled by dynamically adjusting the transmission rates of the sources feeding traffic into these queues. The source rates should be decreased when the queue grows too much and increased when the queue recedes. Such dynamic control of source rates leads to the typical oscillating queueing behaviour. Large oscillations cause the following undesirable effects : increased buffer overflow (packet loss), buffer underflow (waste of available link bandwidth) and network jitter or delay variance. It is therefore important to minimise queue oscillations (or equivalently queue length variance) in computer networks to minimise the above undesirable effects.

Optimal control of queues aims to minimise queue length variance by optimally computing the source rates. TCP/IP networks do not implement optimal control. The queues in IP routers are loosely controlled by TCP's adaptive congestion control algorithms [7] which mainly focus on preventing congestion collapse and maintaining a stable network. TCP achieves this by reacting to packet loss (buffer overflow) with a severe cut to packet flow and slowly increasing the rate as it starts to receive acknowledgements from the receiver. Large queue oscillations in TCP/IP networks are evidenced by high packet loss rate and large jitter (delay variance)in the Internet.

Before we can achieve optimal control of queues, we must first gain a deep insight into the effect of source rates on the queue-length dynamics. To this end, the control of queues and the source rates can be mathematically modelled as a feedback control system. The basic concept is quite simple: the queue length at the network buffer is monitored periodically and compared with a target threshold. Based on the deviation of the current queue length from the target value, a controller computes an admission rate which is fairly allocated among the competing sources. The feedback control theory allows us to mathematically guarantee the stability of the system.

Although theoretical analysis of source rate allocations and queue length dynamics are rare in TCP/IP networks, there has been great interest in modelling the rate control problem as a feedback control system in the context of ATM networks [1], [5], [9], [11]. In an early deterministic treatment, Benmohamed and Meerkov [5] essentially proposed a PI (*proportional* and *integral*) feedback controller without any disturbance feedforward. In their so called PD scheme, the time derivative of the admission rate is proportional to the queue length and its derivative. Such PI controllers are known to provide queue control with zero offset (the mean queue-length converges to the target value), but are inherently less stable and slower acting than feedforward control. Benmohamed [5] didn't provide an optimal solution to the queue control problem. There have been several other control-theoretic solutions proposed [1], [9], [11] to design stable and robust rate control mechanisms, but none of them address the issue of optimality in terms of minimising queue length variance.

Altman et al. [2] use stochastic optimal control theory to devise a rate control scheme with the twin objectives of minimising queue variance and variance of the rates allocated to the sources. For their optimal solution to work adaptively, it would require to solve a complex non-linear discrete Ricaati equation at every time step. This is too complex for practical implementation (and so they propose a suboptimal solution). Implementation complexity is a crucial factor in the adoption of practical solutions. The widely used ATM ABR algorithms, such as ERICA [8], are simple to implement, but do not provide optimal queue control.

In the present paper, we adopt a different approach than Altman [2]. We drop the requirement to minimise the second term (rate signal to the source) and focus on minimising only the first term (queue variance) because the complexity of the optimal solution is far more important for computer networks than the rate fluctuations. The traditional reason to minimise the second term is to avoid undue mechanical stress in an airplane etc. In computer networks, the "stress" caused by undue oscillations of the rate signal will be, e.g, the effect on TCP that may not necessarily be significant.

Our approach yields a simple, practically realisable optimal control scheme that meets the two objectives. We minimise the queue variance at the earliest time affected by the current control signal, and in so doing achieve the second objective of fully utilising the mean available bandwidth. We find that the optimal control comprises simple proportional control, without the integral term, together with feedforward of the disturbance (background traffic). By using disturbance feedforward our controller is *pro-active* to traffic changes (in addition to being *reactive* to queue length changes).

For our proposed controller, we derive optimal solutions for both LANs (no delay) and WANs (homogeneous delays; all sources are at the same distance from the bottleneck queue). We propose an adaptive algorithm which adaptively estimates the background traffic level and its correlation parameter without requiring any prior knowledge of the dynamics of the background traffic. We perform discrete event simulations with both exponential and self-similar background traffic and compare the performance of our feedforward controller with that of typical PI controllers to assess the efficacy of our controller

We have achieved the following interesting results:

- our analysis for optimal control for both no-delay (LAN scenario) and one round-trip delay (WAN scenario) matches the discrete event simulation results thus confirming their accuracy
- lower queue fluctuations are achieved with self-similar than with exponential background traffic. Such queue behaviour is in contrast to that in uncontrolled queues where self similar traffic is known to have more deleterious consequences than exponential traffic [10]
- the optimal proportional plus disturbance feedforward performs better than a tuned PI controller

The rest of the paper is organised as follows. Section 2 presents the design and analysis of our feedforward controller with optimal solutions for both LANs and WANs. Section 3 describes our simulation model. Simulation results are presented in Section 4 followed by our conclusion in Section 5.

II. Optimal Controller Analysis and Design

A. LAN Scenario

The first scenario considered is of a LAN host, having a connection across a WAN, that has its rate controlled by the WAN gateway. In this case, the latency of the feedback loop is negligible in comparison to packet transmission times.

Let $B(k)$, $R(k)$ and $Q(k)$ denote the k-th samples of, respectively, the bandwidth available to the connection, the controlled source rate and the queue length at the buffer in the gateway. $B(k)$ is the capacity of the outgoing link from the gateway less the (uncontrolled) background traffic rate, and is modelled as the stochastic auto-regressive (AR) process

$$B(k) = \alpha B(k-1) + \beta W(k) + (1-\alpha)\bar{B} \quad (1)$$

where $W(k)$ is unit white noise and $\bar{B}$ is the mean value of $B(k)$. The parameter α captures the time correlation of $B(k)$. The parameter β depends on the variance of $B(k)$, which can be derived from (1) as

$$Var\{B(k)\} = \frac{\beta^2}{(1-\alpha^2)} \quad (2)$$

Note that Equation (1) models short-range dependency, and is exact for on-off sources with exponentially distributed on and off times. Also note that our methodology is readily extended to the case where (1) is replaced by a higher-order AR process. We argue that this is actually a worst-case in the sense that future values of $B(k)$ are less predictable than when $B(k)$ exhibits long-range dependency. This is borne out by the results presented in Section 4 which show decreased queue length variance with Pareto-distributed on and off times.

Denoting the sampling period by T, the queue length $Q(k)$ obeys the following equation

$$Q(k+1) = Q(k) + T[R(k) - B(k)] \quad (3)$$

Our goal now is to find the optimal rate $R(k)$ that minimises the packet loss rate due to buffer overflow while keeping the mean queue length equal to a reference value Qr. To this end, we first determine the variance of $Q(k+1)$ and then choose $R(k)$ to minimise it. This is an optimal minimum-variance control problem [4].

Substituting (1) in (3) gives:

$$Q(k+1) = Q(k) + T[R(k) - \alpha B(k-1) - \beta W(k) - (1-\alpha)\bar{B}] \quad (4)$$

As the mean queue length is required to be Qr,

$$\begin{aligned} Var\{Q(k+1)\} &= E\{Q(k+1) - Q_r\}^2 \\ &= E\{Q(k) - Q_r + T[R(k) - \alpha B(k-1) \\ &\quad -(1-\alpha)\bar{B}] - T\beta W(k)\}^2 \end{aligned} \quad (5)$$

$T\beta W(k)$ is the only stochastic term on the RHS, so $R(k)$ that minimises $Var\{Q(k+1)\}$ must make the deterministic terms vanish. This gives the optimal $R(k)$ as

$$R(k) = \bar{B} + (1/T)[Qr - Q(k)] + \alpha[B(k-1) - \bar{B}] \quad (6)$$

And then the minimum variance of the queue length equals the variance of the stochastic term. Using (2) to substitute for β, this is

$$Var\{Q(k+1)\} = T^2\beta^2 = T^2(1-\alpha^2)Var\{B(k)\} \quad (7)$$

Notice that the optimal control law (6) is the sum of three terms, the first being the *mean* available bandwidth $\bar{B}$. This term ensures that the available bandwidth is fully utilised; it arises from the requirement implicit in Eqn. (5) that $E\{Q(k+1)\} = Qr$. The other two terms in (6) are a *feedback* gain ($= 1/T$) times the difference between the current and reference queue lengths, and a *feedforward* gain ($= \alpha$) times the deviation of the most recent available bandwidth $B(k-1)$ from $\bar{B}$.

0-7803-7097-1/01/$10.00 ©2001 IEEE

In order to implement (6), $B(k-1)$ is estimated from the buffer's drain rate, and $\bar{B}$ is then obtained by low-pass filtering these estimates. The correlation parameter α may be estimated in real time by an LMS (or WLS) algorithm. Details are given in Section 2.3.

Finally, and most importantly, stability of our controller is assured by the standard result [3] that a linear quadratic optimal control law for a controllable system (which Eq. (3) is) is stabilising.

B. WAN scenario

Next we turn to the scenario where source rate control is exerted by a bottleneck node on a WAN. Now there is control loop latency equal to the round-trip time. The buffer equation (3) modified to reflect the latency is:

$$Q(k+1) = Q(k) + T[R(k-1) - B(k)] \tag{8}$$

with the sampling period chosen to be one round-trip time. This choice makes sense because changes in control action should only be made after the effect of previous actions is known to the controller (e.g. as in TCP).

Clearly, the first queue length that can be affected by the controlled rate $R(k)$ is $Q(k+2)$, which by recursive use of (8) and (1) is obtained as

$$\begin{aligned} Q(k+2) &= Q(k) + T[R(k-1) - \alpha B(k-1) - \beta W(k) - \\ &\quad (1-\alpha)\bar{B}] + T[R(k) - \alpha^2 B(k-1) - \alpha\beta W(k) \\ &\quad -\alpha(1-\alpha)\bar{B} - \beta W(k+1) - (1-\alpha)\bar{B}] \end{aligned} \tag{9}$$

Similarly as before, the variance of $Q(k+2)$, i.e. $E\{Q(k+2) - Qr\}^2$, is minimised when the deterministic terms on the RHS of (9) sum to Q_r. This gives the optimal rate control law as

$$R(k) = 2\bar{B} - R(k-1) + (1/T)[Q_r - Q(k)] + (\alpha + \alpha^2)[B(k-1) - \bar{B}] \tag{10}$$

This corresponds to a minimum queue length variance of

$$\begin{aligned} Var\{Q(k+2)\} &= T^2\beta^2[1 + (1+\alpha)^2] \\ &= T^2(1-\alpha^2)[1 + (1+\alpha)^2]Var\{B(k)\} \end{aligned} \tag{11}$$

As in the optimal rate control law (6) for the LAN scenario, the control law (10) requires feedback of the queue length deviation and feedforward of the available bandwidth (but with a gain of $\alpha + \alpha^2$ rather than α). But the term $\bar{B}$ is replaced by $2\bar{B} - R(k-1)$. The latter may be viewed as the addition of the term $\bar{B} - R(k-1)$, i.e. the deviation from the mean available bandwidth of the controlled rate during the previous sampling period. As in the LAN scenario, the term $\bar{B}$ ensures that the available bandwidth is fully utilised in the mean.

C. Parameter Estimation

Implementation of the two control laws (6) and (10) requires real-time estimates of α as well as of $\bar{B}$. The latter is obtained by low-pass filtering (EWMA) the measurements of $B(k)$,

$$\bar{B}(k) = \gamma\bar{B}(k-1) + (1-\gamma)B(k) \tag{12}$$

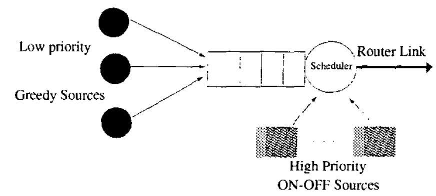

Fig. 1. Network queue configuration for the simulation model

We used $\gamma = 0.95$ in the simulations. Then α can be estimated by applying the normalised Least Mean Square (LMS) algorithm [6] to equation (1):

$$\alpha(k) = \alpha(k-1) + \frac{\delta[B(k) - \alpha(k-1)B(k-1) + (1-\alpha(k-1))\bar{B}(k-1)]}{[\varepsilon + (B(k) - \bar{B}(k))^2]} \tag{13}$$

We set the adaptation gain δ to 0.01. The quantity ε, introduced to avoid division by zero, was set to 0.1 times $Var\{B(k)\}$, estimated by low-pass filtering $[B(k) - \bar{B}(k)]^2$ using an equation similar to (12).

III. Network Simulation

The network queue model simulated is shown in Figure 1. There are three "greedy" low priority sources trying to send traffic through a packet switching router as fast as possible. The router link is also shared by a number of higher priority ON-OFF sources. The link scheduler serves packets from the greedy sources only when there is no traffic from ON-OFF sources.

Traffic from the greedy sources is queued at the router link. This queue is the subject of control and the behaviour of this queue is evaluated in this simulation. The background traffic from ON-OFF sources act as a disturbance to the system, as the available link rate to serve the packets from the queue varies dynamically due to these ON-OFF sources.

The sending rates of the greedy sources are adjusted dynamically according to the rate feedback from the queue controller. Each source gets one third of the optimal rate computed by the controller. The objective of the controller is to keep the queue at a pre-determined target level by adjusting the aggregate admission rate to the queue.

An ON-OFF source transmits traffic at a constant rate during the ON period and does not send any traffic during the OFF period. For exponential background traffic, the ON and OFF periods are exponentially distributed. To simulate self similar background traffic, the lengths of ON and OFF periods are drawn from Pareto distribution as suggested in [13]. In our simulation, ON-OFF sources do not explicitly send any packets and hence no packet queue is explicitly simulated at the scheduler for these high priority sources. Instead, whenever an ON-OFF source is turned on, the scheduler reduces the link capacity by 10 pk/sec and adds the same to the link capacity whenever it is turned off.

In our simulation we set a target queue length Q_r and measure the variance of the current queue length around this target value. The better the control is, the lower the variance should be.

The lengths of all simulation runs were dynamically controlled until the values of the estimated variables (i.e, the queue variance) reached steady state, and 95% confidence interval were achieved with 5% relative precision. With such adaptive length simulations, the number of packets simulated ranged from millions to billions of packets depending on the simulation parameters. The values of simulation parameters are summarised in Table I. Results obtained from the simulations are presented and compared with the analytical results in the following section.

We have simulated both our proportional plus feedforward controller and a tuned PI controller. Benmohamed [5] only showed the existence of a stabilising PI controller, but actually an infinite number of stabilising PI controllers exist. Hence, for simulation purposes, we chose tuning parameters given by Ziegler-Nichols' continuous cycling method [12] which are known to give a good disturbance response.

The discrete PI controller has the transfer function $K_c[1+(\frac{T}{T_i})/(z-1)]$. For RTT = 0 (LAN scenario), the settings are found to be $K_c = 0.9/T$; $T_i = 1.67T$ and for RTT = T (WAN scenario), they are $K_c = 0.45/T$; $T_i = 5T$.

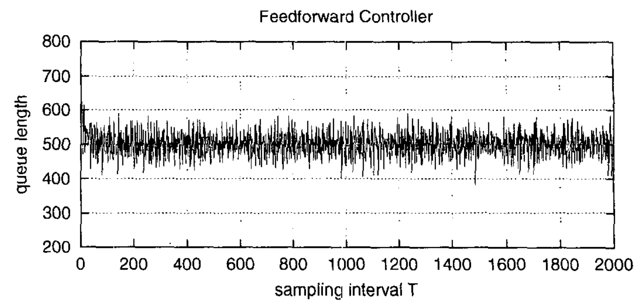

Fig. 2. Queue length dynamics with feedforward controller

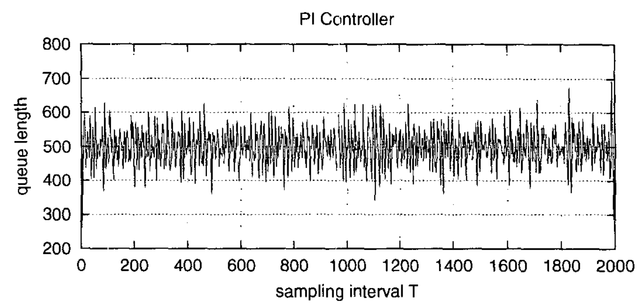

Fig. 3. Queue length dynamics with PI controller

IV. Results

To verify the accuracy of our optimal control analysis of the queue variance, we compare the theoretical values with the ones obtained from the simulation for the ratio $\frac{Var\{Q\}}{Var\{B\}}$. The ratio is compared, instead of the actual queue variance, to avoid the need to compute the variance of the available link capacity. Table II compares this ratio for exponential on and off periods.

For ON-OFF sources with exponential on and off periods with identical mean of 2 seconds, the theoretical value of the correlation parameter α of the background traffic is obtained as 0.7266. This value of α is used to compute values in column 2 (Analysis) of Table II. From Table II we see that the simulation results for $\frac{Var\{Q\}}{Var\{B\}}$ is only slightly different than the ones predicted by our analysis. This may be due to the errors in the estimation of α in the simulation using the LMS algorithm (see Eqn. (13)).

Now we compare the performance of the optimal proportional plus feedforward controller against the PI controller in terms of queue variance. To observe queue length fluctuations, we recorded queue lengths at discrete intervals of T seconds for 2000 samples. Figures 2 and 3 show the queue lengths for the WAN scenario for feedforward and PI controllers respectively. From the figures it is evident that although both controllers successfully maintain the mean queue length around the target value of 500, PI controller performs worse than feedforward controller in terms of the deviation of the queue length from the target value.

Figures 2 and 3 provide visual confirmation that with feedforward control queue lengths are more tightly controlled. For a quantitative evaluation of the queue variance, we run simulations until the queue variance reaches a steady state (many millions of sampling intervals are simulated). For LAN and WAN scenarios, steady state queue variances and the mean queue lengths are shown in Tables III and IV respectively. These tables show results, in separate columns, for background traffic with exponential ON-OFF sources and with Pareto ON-OFF sources. For Pareto ON-OFF sources, the tables show results for two Hurst parameters, H=0.7 and H=0.9.

A glance through these tables reveals several observations:

- in the steady state, both feedforward and PI controllers can successfully maintain the mean queue length at the target value (500 in this case)
- the feedforward controller achieves much lower queue variance than the conventional PI controller
- irrespective of the control law, lower queue variance is observed with self-similar background traffic than with exponential ON-OFF traffic. The higher the Hurst parameter of the self-similar traffic,the lower the queue variance. This result suggests that our analysis, which is based on exponential traffic, actually provides results for worst case scenarios.

V. Conclusion

We have designed a simple and practically realisable adaptive rate allocation algorithms for optimal control of queues in both LANs and WANs. Using our algorithms, queue oscillations can be minimised to reduce packet loss and network jitter, and increase bandwidth utilisation. No such simple, explicit optimal solutions havebeen proposed before. Discrete event simulations were carried out to ver-

TABLE I

VALUES OF SIMULATION PARAMETERS.

Parameter	*Value*
Number of ON-OFF priority sources	90
Transmission rate during ON period for VBR source	10 pk/sec
Mean length of ON and OFF period for VBR source	2 sec
Capacity of router link	900 pk/sec
Sampling interval T	0.5 sec
Target queue-length Q_r	500 packets

Scenario	*Analysis*	Simulation
LAN	0.118	0.116
WAN	0.470	0.541

TABLE II

ANALYSIS AND SIMULATION RESULTS FOR $\frac{Var\{Q\}}{Var\{B\}}$, FOR FF CONTROLLER WITH EXPONENTIAL BACKGROUND TRAFFIC IN LAN SCENARIO.

Type of Control	*qlength variance*			*mean qlength*		
	Exponential	*H=0.7*	*H=0.9*	*Exponential*	*H=0.7*	*H=0.9*
Feedforward	155.81	111.45	86.89	500.54	500.50	500.45
PI	223.45	155.07	112.03	499.97	499.98	499.97

TABLE III

SIMULATION RESULTS FOR LAN SCENARIO

Type of Control	*qlength variance*			*mean qlength*		
	Exponential	*H=0.7*	*H=0.9*	*Exponential*	*H=0.7*	*H=0.9*
Feedforward	744.98	515.22	365.64	501.09	501.15	501.08
PI	2173.33	1306.80	880.10	499.99	500.16	500.13

TABLE IV

SIMULATION RESULTS FOR WAN SCENARIO

ify the fluid-flow models used in developing our algorithms and to compare their performance against non-optimal PI controllers proposed previously. The results showed that queue oscillations are much smaller for our optimal algorithms than for the PI controllers.

We also studied the effect of self-similar traffic on both types of controller by performing simulations with aggregated ON-OFF sources having Pareto-distributed on and off times. It was found that the queue length variance *decreases* with increasing Hurst parameter for self-similar traffic. This result implies that if a queue control system was designed assuming exponential background traffic, the system performance will not deteriorate if the actual traffic is of a self-similar nature. This is in stark contrast with the well known problem of *uncontrolled* queues, where self-similar traffic has a more deleterious effect on the queue length.

REFERENCES

[1] E. Altman, T. Basar, and R. Srikant. Robust Rate Control for ABR Sources. In *IEEE INFOCOM*, 1998.

[2] E. Altman, T. Basar, and R. Srikant. Congestion Control as a Stochastic Control Problem with Action Delays. In *14th IFAC World Congress*, Beijing, China, July 1999.

[3] Anderson and Moore. *Optimal Control.* Prentice Hall, 2nd edition, 1989.

[4] K.J. Astrom and B. Wittenmark. *Computer Controlled Systems.* Prentice Hall, second edition, 1990.

[5] L. Benmohamed and S.M. Meerkov. Feedback Control of Congestion in Packet Switching Networks: The Case of a Single Congested Node. *IEEE/ACM Transaction on Networking*, 1(6):693–708, 1993.

[6] S. Haykin. *Adaptive Filter Theory.* Prentice Hall, 3rd edition, 1998.

[7] V. Jacobson. Congestion Avoidance and Control. *ACM Computer Communication Review*, 18:314–329, August 1988.

[8] S. Kalyanaraman, R.Jain, S. Fahmy, R Goyal, and B. Vandalore. The ERICA Switch Algorithm for ABR Traffic Management in ATM Networks. *IEEE/ACM Transactions on Networking*, February 2000.

[9] A. Kolarov and G. Ramamurthy. A Control-Theoretic Approach to the Design of an Explicit Rate Controller for ABR Service. *IEEE/ACM Transactions on Networking*, 7(5):741–753, 1999.

[10] W.E. Leland, M.S. Taqqu, W. Willinger, and D.V. Wilson. On the Self-Similar Nature of Ethernet Traffic (Extended Version). *IEEE/ACM Transactions on Networking*, 2:1–15, 1994.

[11] C.E. Rohrs. A Control Engineers Look at ATM Congestion Avoidance. In *IEEE GLOBECOM*, 1995.

[12] F.G. Shinskey. *Process Control.* McGraw-Hill, 2nd edition, 1979.

[13] W. Willinger, M.S. Taqqu, R. Sherman, and D.V. Wilson. Self-Similarity Through High-Variability: Statistical Analysis of Ethernet LAN Traffic at the Source Level. *IEEE/ACM Transactions on Networking*, 5(1):71–86, 1997.