The Role of Information Update in Flow Control

E. Altman, T. Başar, and N. Malouch

***Abstract*—A common feature of congestion control protocols is the presence of information packets used to signal congestion. We address here the question of how frequently such protocols need to generate information packets in order to *optimize* their performance. Through a number of congestion control models, we identify and quantify different types of effects of the frequency of generating information packets. We consider both TCP-type protocols, in which controlling the frequency of information packets is done through static or dynamic delayed ACK options, as well as ATM type flow control, where the optimal time spacing between the generation of network management packets is computed. We show how the spacing between information packets influences the throughput and the stability of the system.**

***Index Terms*—Congestion control, flow control, information rates.**

I. Introduction

CONGESTION and flow control protocols make use of special control packets to convey to the traffic sources actions to be taken in order to adapt the transmission rates to the available bandwidth. In TCP/IP, these are the acknowledgement packets; in the ABR (Available Bit Rate) class of ATM, these are the resource management packets. Even when the size of these packets is significantly smaller than that of the data packets, they still compete over network resources with the data packets and thus decrease the amount of resources available to them. In many cases, they may require an amount of resource much larger than what their size would suggest. For example, when using the IEEE 802.11 MAC protocol, each data as well as ACK packet of the TCP flow requires the same (large) overhead as of the three link layer packets (RTS, CTS and a link layer ACK). Even in the absence of such overheads, the processing time of an ACK at the destination might require an additional overhead that could be pretty large with respect to the transmission time when very high speed networks are considered. We therefore raise the question, in this paper, of at what frequency flow control protocols should send control packets. Our aim is to answer this question using analytical tools, supported by numerical and simulation studies.

Paper approved by T.-S. P. Yum, the Editor for Packet Access and Switching of the IEEE Communications Society. Manuscript received October 6, 2005; revised January 25, 2007 and October 13, 2007. A preliminary version of this paper has appeared in the proceedings of Networking'04, Athens, May 9–14, 2004.

E. Altman is with INRIA, 2004 Route des Lucioles B.P. 93, 06902 Sophia Antipolis, France (e-mail: Eitan.Altman@sophia.inria.fr). E. Altman was supported by ECOS Nord contract V04M01.

T. Başar is with the University of Illinois, 1308 West Main Street, Urbana, IL 61801-2307, USA. Research of this author was supported in part by NSF Grant CCR 00-85917 ITR.

N. Malouch is with LIP6/CNRS, Université Pierre et Marie Curie - Paris 6. 104 Avenue du Président Kennedy, 75016 Paris, France.

Digital Object Identifier 10.1109/TCOMM.2008.050389.

Note that TCP/IP already has the "delayed ACK" option that allows it to reduce the ACK frequency from one ACK for every received packet to one ACK every $d = 2$ received packets. For Ad-Hoc networks using the IEEE 802.11 MAC protocol, it has been established through simulations that further improvement can be obtained when using TCP with $d > 2$ [6], [7]. Note that ACKs could also be filtered within the network (see e.g. [9] and references therein).[1]

We introduce three frameworks in which to study the optimization of the frequency of control packets. The first is the case of TCP/IP traffic sources with routers using drop tail queues. A simple mathematical model is derived for optimizing the amount of ACK thinning at the destination so as to maximize the system's throughput. A simulation study validates the conclusions we arrive at using the mathematical model. In the second framework, we study the dynamics of an Additive-Increase Multiplicative-Decrease (AIMD) flow control interacting with a RED type buffer. We model the system's dynamics through a system of delay-differential equations, and study the stability of the system as a function of the frequency of ACKs. The third framework we adopt is that of optimal rate control with sampled state information, similar to models used for rate control in the ABR class in ATM [2]. We formulate a linear-quadratic control model with two different views of uncertainty (stochastic, and deterministic but worst case), and we optimize the time between two successive packets in information feedback by trading off length of time against performance.

The three frameworks introduced above are covered respectively in the three sections to follow, Sections II, III, and IV, which also include simulation and numerical results. The paper ends with the concluding remarks of Section V, and an appendix.

II. A Fixed Point Method to Model ACK Thinning

In this section, we use the expression for TCP throughput developed in the literature and standard queueing models to quantify the impact of the delay factor d on the throughput. The throughput T of a TCP connection can be approximated by [21]

$$\left[RTT\sqrt{\frac{2\,d\,p}{3}} + RTO\,\min\left(1,\,3\sqrt{\frac{3\,d\,p}{8}}\right)p\left(1+32p^2\right)\right]^{-1} \tag{1}$$

[1]The advantage of thinning ACKs at the destination using the delayed ACK option is that if d packets have not yet arrived at the destination but some timer has expired, the destination will generate an ACK thus avoiding situations in which the source will interpret the lack of ACKs as a loss of a packet detected through the source's timeout expiration.

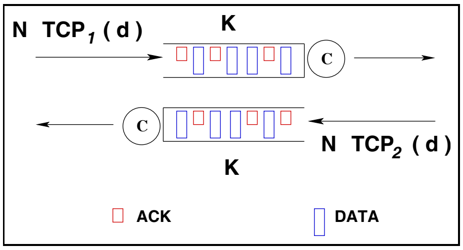

Fig. 1. A network model for ACK thinning.

where p is the loss probability of TCP packets, RTT is the round-trip delay experienced by the TCP connection, and RTO is the retransmission timeout.

The above formula was established under the assumption that the TCP sender increases its congestion window for each arriving ACK. It is worthwhile to notice that the TCP sender can rather increase its congestion window using the number of previously unacknowledged bytes each ACK covers [24], [25]. In this case, the factor d disappears from the formula, and the throughput increases. Since in this paper our goal is to find the optimal value of d that *maximizes* the total throughput of the TCP connections in the network, we keep the factor d in (1). However, our study could also serve as a worst-case analysis for the other scenario mentioned.

Remark 2.1: We use the term "throughput" in the same sense as in [17], [18], i.e. the average number of packets sent in a time unit. Note however that "throughput" is used in some papers (see, e.g. [21], [17], [18]) in the meaning of "goodput", which is the average number of packets received error-free in a time unit. The term "send-rate" is then used in the sense of our meaning of "throughput" [21]. In any case, for small loss probabilities, throughput is a good approximation of goodput.

We start by modeling the network as a bidirectional link. Each direction of the link is modeled by a queue system. Two sets of N symmetric TCP sources send data from both end-points of the link (Figure 1). Each source of the first set connects to a receiver that does not belong to the second set and vice versa. Thus, in each queue TCP packets and ACK packets from different connections are multiplexed in the same queue and served by the same server. Using this model, we assume that losses that occur in the system are only due to congestion, i.e. buffer overflow.

We denote by α the "effective" size of ACK packets and by Z the size of TCP packets. The parameter α would model not only the actual size of ACK packets but also the eventual overheads introduced in the processing time. Denote by $b(t)$ the service time distribution, which can be expressed as

$$b(t) = \begin{cases} \dfrac{\alpha}{Z\,C} & \text{if an ACK is in service at time } t \\[2ex] \dfrac{1}{C} & \text{if a TCP packet is in service at time } t \end{cases} \tag{2}$$

where C is the capacity of the link in TCP packets per unit of time. In order to develop a tractable model, we need a simple formula that relates the throughput at the buffers to the losses that will be experienced there. To that end we shall assume that the packet arrival process at each queue can be approximated by a Poisson process. (We note that the validity of this approximation in a similar context has been discussed and examined in [1], [8].) Then, the packet loss probability is the loss probability of an $M/G/1/K$ system. Note that the loss probabilities seen by ACK packets and TCP packets are theoretically the same because of the PASTA property.

Below we propose two modeling approaches for the service time: the exponential service time which provides a simple expression for the losses but gives a rough approximation, and the deterministic service time (whose duration varies according to whether it is an ACK or a data packet) which gives a better approximation but with a more complex expression for the losses. In the deterministic model there are thus two possible values of service times: α/ZC or $1/C$.

A. The Exponential Service Time Case

In this case, the service time of packets is exponentially distributed. The loss probability is then given by the loss probability of an $M/M/1/K$ system:

$$p = \rho^K \frac{1-\rho}{1-\rho^{K+1}} \tag{3}$$

where ρ is the load of the system, which is computed as follows:

$$\begin{aligned} \rho &= \left(NT + \frac{NT}{d}\right)\left(\frac{1}{d+1}\frac{\alpha}{ZC} + \frac{d}{d+1}\frac{Z}{ZC}\right) \\ &= \frac{1}{C}(NT + NT\frac{\alpha}{dZ}) \end{aligned} \tag{4}$$

We use the fixed point method to solve numerically the system of equations (1), (3) and (4). The advantage of this simple model is that we can compute the loss probability and thus the throughput for large values of K and ρ.

B. The Deterministic Service Time Case

Here, we modify only the assumption that the service times are exponentially distributed and thus only equation (3) is replaced by the expression of the loss probability of an $M/G/1/K$ queue [22], [23]:

$$p = \frac{1 + (\rho - 1)\,f}{1 + \rho\,f}\,, \text{ where } f = \frac{1}{2\pi\,i}\oint_{D_r} \frac{1}{G(s)\,s^{K-1}}\,\mathrm{d}s \tag{5}$$

Here, D_r is any circle in the complex plane with center 0 and with radius r chosen small enough so that all zeros of the function $G(s)$ are outside the circle, i.e. $r < |G(z)|\ \forall\ z$ such that $G(z) = 0$. The complex-valued function $G(s)$ is defined as

$$G(s) = LST(b(\lambda(1-s))) - s\,,$$

where $LST(b(\cdot)) = \int_0^\infty b(t)\,e^{-st}\,dt$ is the Laplace Stieltjes transform of the service time distribution:

$$LST(b(s)) \quad = \quad \underbrace{\frac{1}{d+1}e^{-(\alpha s/ZC)}}_{ACK} + \underbrace{\frac{d}{d+1}e^{-s/C}}_{TCP}$$

The parameter λ is the total arriving rate at the entrance of each queue, which is equal to $(NT + NT/d)$. Computation of f is detailed in the Appendix.

Again, we solve the three equations (1), (5) and (4) numerically, and compare the results to those obtained using the exponential time distribution, as presented next.

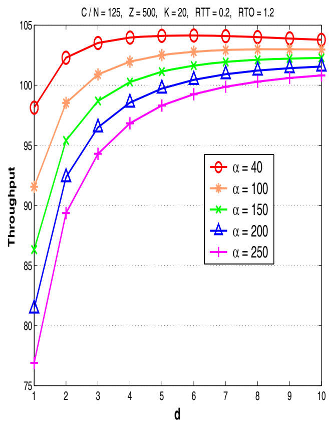

Fig. 2. Effect of d on the throughput for various ACK sizes: the deterministic service time case.

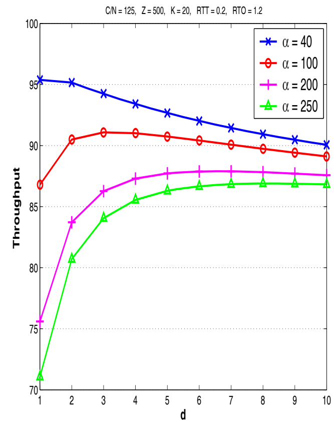

Fig. 3. Effect of d on the throughput for various ACK sizes: the exponential service time case.

C. Numerical Results and Simulations

In this section we use our model to study numerically the trade-off controlled by the delay factor d. We will show that a relative gain in the throughput, ranging from 5% to 50%, can be achieved by setting $d > 1$. We also performed simulations using NS-2 to evaluate the accuracy of the two models at capturing the impact of d on the throughput.

Figure 2 depicts plots of the throughput of TCP as a function of the delay factor d. We consider an "effective" size of TCP data packet of 500 bytes and various "effective" sizes of ACKs, ranging from $\alpha = 40$ to 250. (As mentioned in the introduction, the difference between actual and effective sizes of a packet is that additional overhead may be added to its real size due to other protocols of other layers; in addition, processing an ACK at some nodes may take longer than its relative size with respect to a TCP data packet.) The queue size is 20 packets, the round-trip time RTT is 200 ms and the retransmission timeout RTO is 1.2 s. We set C to 125 TCP packets/s, and we set N to 1 since the throughput depends only on the fraction C/N. The plots were generated numerically using the more precise model of deterministic service times.

The figure shows that for the small ACK size of 40, spacing the ACK (using $d > 1$) results in a small improvement in the performance (6%). However, the throughput is maximized when $d = 4$. For higher values of "effective" ACK size ($\alpha \geq 100$), we see an improvement of 12% to 31%, with the optimum obtained for example for $\alpha = 100$ at $d = 5$. Using a value of $d > 2$, which is the default value in TCP, results in a relative gain of around 10%. We have observed similar trends with larger link capacities C (larger congestion window) and with much smaller buffer sizes K (higher loss probabilities). When $K \geq 60$, the throughput approaches the value obtained simply by solving for T from (4) with ρ set equal to 1, which yields

$$T = \frac{C}{N(1 + \alpha/(dZ))},$$

and the relative gain can reach 50%.

Figure 3 depicts plots of the throughput versus d with the same parameters as in the previous scenario, but now using the exponential service time approximation. In this case, the packet size can represent the average of the "effective" packet size in the network. This model is useful when the packet size in the backbone is variable and the exact distribution is unknown. The figure shows that the trend in the throughput is almost the same as before, except for the smallest ACK size of 40 where spacing the ACK results in deterioration of the performance for all values of $d > 1$. For larger ACK sizes ($\alpha \geq 200$), once again, the gain of delaying ACKs goes beyond 20%.

We can explain this behavior as follows. When increasing d, it is seen from (1) that the throughput decreases; this is the effect of the explicit dependence of the throughput on d. But d has yet an indirect opposite impact on the throughput through its impact on p: the probability of loosing a TCP data packet is expected to decrease with d since the load decreases with d, see eq. (4). This indirect impact of d on the throughput clearly becomes smaller as the ACK size becomes smaller. This impact is even smaller for the deterministic packet size. This is due to the fact that the loss probability p in the exponential model (M/M/1/K) overestimates the loss probability computed by the deterministic model (M/D+D/1/K). We conclude that for all sufficiently small ACK size we can expect the throughput to deteriorate when d becomes larger than 1. Moreover, with exponentially distributed service time, this will already occur with larger ACK size than those for which it would occur in the case of deterministic service time.

Next, we performed two sets of simulations to study the robustness of the models described above. First, we study the effect of the number of connections N. Then, we look at the effect of the receiver timeout[2] which is not captured by the two models. We use a similar network configuration as the one described in the model of Figure 1, and we add N access links of capacity 1 Mbs each to the bottleneck link. We set the propagation delays such that $RTT \approx 200$ ms and $RTO \approx 1.2$ s.

In the first set of simulations, we vary the number of connections N from 30 to 200, and we vary the bottleneck capacity C in order to keep the ratio C/N constant at 125. We fix the ACK size to $\alpha = 100$. In this set of simulations, we use the default value of the receiver timeout which is set to 100 ms. Figure 4 plots the average TCP throughput versus the delayed factor d, using the two models and simulation traces for $N = 30$, 50, 100 and 200.

The key observation we make here is that when N increases, the relative error induced by the deterministic model decreases notably. This is because the resulting process of multiplexing the TCP connections approaches a Poisson process [8]. Surprisingly, the model with the exponential service time approximation predicts very well the throughput when $N = 30$. The combination of two phenomena is responsible for that.

(i) First, we note that whereas the queueing models yield the same throughput when multiplying both the number of connections and the capacity by the same constant, the simulation results do not: when the number of connections in the simulation is small, the capacity cannot be fully utilized and the throughput of single connections decreases. Another

[2]The receiver timeout is used to send an ACK for arriving packets even before d packets arrived at the receiver if the time since the unacknowledged packet arrived exceeds the timeout.

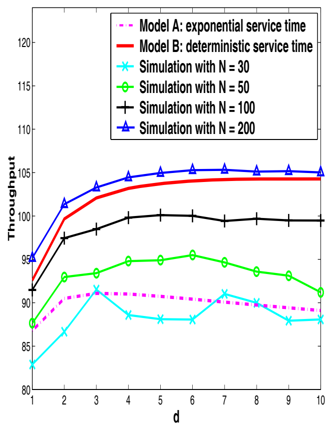

Fig. 4. Simulation results versus numerical results.

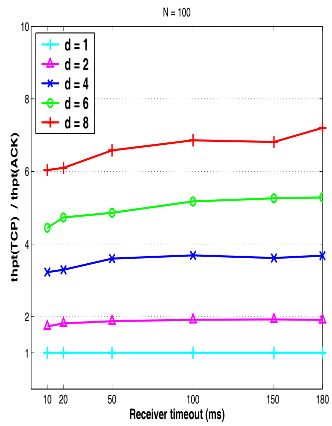

Fig. 5. Effect of the receiver timeout on the number of generated ACKs when $N = 100$.

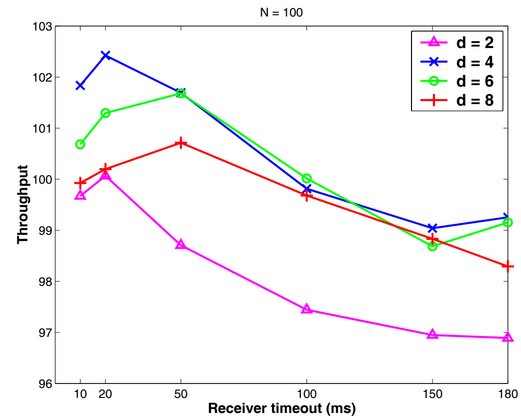

Fig. 6. Effect of the receiver timeout on the TCP throughput.

reason for the fact that the throughput of the simulation is lower for smaller N is that there is less multiplexing so the packet arrival process is more bursty than the Poisson model.

(ii) Second, the exponential service model yields lower throughput than the deterministic service one (this is due to the fact that the exponential model (M/M/1/K) over-estimates the loss probability with respect to the one computed by the deterministic model (M/D+D/1/K)). Combining this with the fact that the simulated throughput is lower than the deterministic model can explain the better fit of the simulation with the exponential model for smaller values of N. However, it is important to notice that this observation do not lead to the conclusion that exponential distribution provides a bound for small N.

If we examine closely the plots in Figure 4, we can observe that when d is large, the throughput in the simulations decreases (as opposed to the one estimated by the deterministic model). The explanation is that when d is large, the receiver timeout expires more often especially when the congestion window size is small. Thus, the number of ACKs increases and the bandwidth consumed by ACKs becomes larger than NT/d. Figure 5 illustrates this behavior when $N = 100$, which depicts the ratio between the TCP and the ACK throughputs. This ratio is seen to increase when the timeout value increases.

After examining the influence of the receiver timeout on the ratio of throughputs, we now turn to examine its impact directly on the TCP throughput. We use the same parameter values as in the previous simulations, and set $N = 100$. We vary the receiver timeout from $t_1 = 10$ ms to $t_2 = 180$ ms. Note that t_1 is just larger than the minimum inter-arrival time of two consecutive TCP packets sent in the same window, and t_2 is just less than RTT.

Figure 6 depicts plots of the average TCP throughput as a function of the receiver timeout for various values of d. We see clearly that when the timeout is large, then the throughput is reduced. In the cases $d = 2$ and $d = 4$, the throughput is maximized when the timer is equal to 20 ms. For $d = 6$ and $d = 8$, the throughput is maximized at a timer of 50 ms.

Here, the gain in the throughput is relatively small ($\approx 3\%$), but this is because the average window size is also small ($\approx$ 20). For larger window sizes, for example when the delay-bandwidth product is large, the gain is more significant. However, setting a small value for the timeout is risky since the packet inter-arrival time depends on the cross traffic along the path and hence could vary over time.

More generally, it is difficult to find a constant timeout that is adequate for all network scenarios. Another alternative is to compute dynamically the timeout using a similar method as the one used by the sender to compute the retransmission timeout.

In concluding this section, we note that it is worth to use a delay factor $d > 2$, particularly for long TCP connections and when the window size is large. In fact, the improvement obtained from reducing the number of ACKs in the network is more significant than the decrease in the throughput due to the lack of information update. Besides, by choosing an adequate receiver timeout, we can improve further the gain in the throughput.

III. A Dynamic Control Model of TCP Traversing a RED Buffer

After focusing on the quantitative impact of the rate of information on the throughput, we study in this section a more qualitative property, the *stability*, and show how stability conditions are influenced by the rate of information packets.

We assume that N symmetric persistent TCP connections share a bottleneck link of capacity C, located close to the source, assumed to operate in the congestion avoidance regime. Let $W(t)$ be the window size of a connection at time t, and R be the round trip delay (including queueing delay), which is assumed to be a constant (this assumption holds when queueing delays are much smaller than propagation delays). Let $p(t)$ be the loss (or marking) probability of the RED buffer at time t. Let $q(t)$ be the amount of buffered traffic at the queue at the bottleneck link.

A standard way to study the interaction between TCP/IP and the RED buffer is to consider a fluid approximation of the window size of TCP. The window size evolution is then given by

$$\frac{\mathrm{d}W}{\mathrm{d}t} = \frac{1}{R} - \frac{W(t)W(t-R)}{2R}p(t-R),$$

see [15]. This equation has been obtained under the assumptions that (i) the delayed ACK mechanism is not used; and (ii) a new ACK is generated with each TCP packet that arrives at the destination, resulting in an increase of the window size by one unit every round trip time.

We shall now consider the possibility of using a (dynamic) delayed ACK approach in which an ACK is generated for every d TCP packets that arrive at the destination. d will be considered as a control variable and will thus be allowed to be time dependent. Both the increase rate as well as the decrease rate are divided by a factor d since the rate of ACKs that arrive at the sources is d time smaller; In particular indications for decreasing the rate (we assume that ACKs have marks indicating congestion) return less frequently. The window size evolution then becomes

$$\frac{\mathrm{d}W}{\mathrm{d}t} = \frac{1}{Rd(t-R)} - \frac{W(t)W(t-R)}{2Rd(t-R)}p(t-R).$$

The queue dynamics is given in [15] by

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \frac{N}{R}W(t) - C.$$

We shall assume that not only the TCP packets have to queue but also the ACKs. To model the difference between the size of an ACK and that of a TCP packet, we assume that an ACK requires a fraction γ of buffer space required by a TCP packet. The queue dynamics above is then modified to

$$\frac{\mathrm{d}q}{\mathrm{d}t} = \frac{N}{R}W(t) - C + \frac{\gamma W(t-R)N}{d(t-R)R}.$$

The tradeoff that influences the choice of the control d will be the following: on one hand, when the window size is small, we may wish d to be small so that the window size can grow quickly so as to achieve higher throughput. On the other hand, when the window size is large then we may wish to increase d so as to limit the congestion due to ACKs in the bottleneck queue.

We shall analyze in this section a linear control mechanism in which d has the form $d(t) = \zeta(1+\beta W(t))$.

Finally, we shall consider the RED buffer marking probability (ignoring the averaging of the queue size) in its linear operation regime: $p(t) = \eta_1 q(t) - \eta_2$.

We summarize the system's overall dynamics below:

$$\begin{aligned}
\frac{\mathrm{d}W}{\mathrm{d}t} &= \frac{1}{Rd(t-R)} - \frac{W(t)W(t-R)}{2Rd(t-R)}p(t-R) && (6)\\
\frac{\mathrm{d}q}{\mathrm{d}t} &= \frac{N}{R}W(t) - C + \frac{\gamma W(t-R)N}{d(t-R)R} && (7)\\
d(t) &= \zeta(1+\beta W(t)),\ \zeta>0, \beta \geq 0 && (8)\\
p(t) &= \eta_1 q(t) - \eta_2 && (9)
\end{aligned}$$

Equilibrium: The equilibrium point (identified by a subscript 'o') is obtained by equating the right-hand sides of (6) and (7) to 0, which leads to

$$\begin{aligned}
&\frac{N}{R}W_o + \frac{N\gamma W_o}{d_o R} = C \Rightarrow W_o = \frac{CR}{N(1+\gamma/d_o)}\\
&d_o = \zeta(1+\beta W_o)\\
&W_o^2 p_o = 2 \Rightarrow p_o = 2/W_o^2\\
&p_o = \eta_1 q_o - \eta_2 \Rightarrow q_o = (p_o+\eta_2)/\eta_1
\end{aligned}$$

Hence the throughput is given by:

$$T = \frac{W_o}{R} = \frac{C}{N(1+\gamma/d_o)},$$

which is seen to increase with d_o.

A linearization of the dynamical system in a neighborhood of the equilibrium point yields

$$\begin{aligned}
\frac{\mathrm{d}\delta W}{\mathrm{d}t} &= -\frac{1}{RW_o d_o}\delta W(t) - \frac{1}{RW_o d_o}\delta W(t-R)\\
&\quad -\frac{\eta_1 W_o^2}{2Rd_o}\delta q(t-R)\\
\frac{\mathrm{d}\delta q}{\mathrm{d}t} &= \frac{N}{R}\delta W(t) + \frac{\zeta\gamma N}{d_o^2 R}\delta W(t-R),
\end{aligned}$$

where δ stands for the shifted version of variables in which the equilibrium value is subtracted (e.g. $\delta W := W - W_o$). Taking the Laplace transform of these equations we obtain

$$\begin{aligned}
s\delta W(s) &= -\frac{1}{RW_o d_o}\delta W(s) - \frac{e^{-sR}}{RW_o d_o}\delta W(s)\\
&\quad -\frac{\eta_1 W_o^2}{2Rd_o}e^{-sR}\delta q(s)\\
s\delta q(s) &= N\left(\frac{1}{R} + \frac{\zeta\gamma}{d_o^2 R}e^{-sR}\right)\delta W(s)
\end{aligned}$$

With $z = sR$, the stability condition is then given by requiring that all zeros of $g(z) = 0$ have strictly negative real parts [27], where $g(z)$ is given by

$$z^2 + \frac{z}{W_o d_o}(1+e^{-z}) + \frac{\eta_1 W_o^2 N}{2d_o}(1 + \frac{\zeta\gamma}{d_o^2}e^{-z})e^{-z} \qquad (10)$$

A. Solution of $g(z) = 0$

Let $x := \Re(z), y := \Im(z)$. Then $g(z) = 0$ is equivalent to

$$\begin{aligned}
0 &= x^2 - y^2 + 2\imath xy + \frac{x+\imath y}{W_o d_o}\\
&\quad + \left(\frac{x+iy}{W_o d_o} + \frac{\eta_1 W_o^2 N}{2d_o}\right)e^{-x}(\cos y - \imath \sin y)\\
&\quad + \frac{\eta_1 W_o^2 \zeta\gamma N}{2d_o^3}e^{-2x}(\cos 2y - \imath\sin 2y)
\end{aligned}$$

or equivalently,

$$\begin{cases}
x^2 - y^2 + \frac{x}{w_o d_o} + \left(\frac{x}{W_o d_o} + \frac{\eta_1 W_o^2 N}{2d_o}\right)e^{-x}\cos y\\
\quad + \frac{y}{W_o d_o}e^{-x}\sin y + \frac{\eta_1 W_o^2\zeta\gamma N}{2d_o^3}e^{-2x}\cos 2y = 0\\
\\
2xy + \frac{y}{W_o d_o} - (\frac{x}{W_o d_o} + \frac{\eta_1 W_o^2 N}{2d_o})e^{-x}\sin y\\
\quad + \frac{y}{W_o d_o}e^{-x}\cos y\\
\quad - \frac{\eta_1 W_o^2 \zeta\gamma N}{2d_o^3}e^{-2x}\sin 2y = 0
\end{cases}$$

The goal. One may now identify two possible goals: (i) maximize the system throughput while maintaining stability, and (ii) for a given desired throughput, make the system "as stable as possible" by which we mean to choose the parameters so as to have the real part of the largest zero of $g(z)$ as negative as possible.

Note that whereas the throughput only depends on d_o, γ, N and C, and not directly on the values of $\zeta, \beta, \eta_1, \eta_2$, the stability regime does depend directly on ζ and η_1, and hence these parameters do enter into the optimization under goal (i). For the second goal, we may first compute d_o and then optimize stability with respect to the other parameters.

We first need to verify (numerically) that there are no zeros of g with $x \geq 0$. We show below that we can in fact restrict the numerical search to a bounded domain.

Assume $x \geq 0$. Then,

$$|g(z)| \geq |z|^2 - \frac{2|z|}{w_o d_o} - \frac{\eta_1 w_o^2 N}{2d_o}\left(1 + \frac{\zeta\gamma}{d_o^2}\right) =: f(|z|)$$

A necessary condition for $g(z) = 0$ is then $f(|z|) \leq 0$. Note that $f(|z|) = (|z| - z_1)(|z| - z_2)$, where

$$z_{1,2} = \frac{1}{w_o d_o} \pm \sqrt{\frac{1}{(w_o d_o)^2} + \frac{\eta_1 w_o^2 N}{2d_o}\left(1 + \frac{\zeta\gamma}{d_o^2}\right)}$$

Thus, for $x \geq 0$, a necessary condition for $g(z) = 0$ is that $|z| \leq z_1$ (where z_1 is the positive zero).

Note: in a similar way, we can show that to find the zeros of $g(z) = 0$ in the region $x \geq v$ (where v may be negative or positive), it suffices to consider $|z| \leq z(v)$ where

$$z(v) := \frac{1}{2}\left(\frac{1 + e^{-v}}{w_o d_o} + \sqrt{\left(\frac{1 + e^{-v}}{w_o d_o}\right)^2 + 4\frac{\eta_1 w_o^2 N e^{-v}}{2d_o}(1 + \frac{\zeta\gamma e^{-v}}{d_o^2})}\right)$$

B. A numerical example

This numerical example is picked to show that the system could be unstable for meaningful choices of the values of the parameters.

$\eta_1 = 0.001$, $\eta_2 = 0.02$, $q_{min} = 20$, $q_{max} = 60$. Further let $\zeta = 1$, $d_0 = 3$, $N = 5$, $\gamma = 250/500 = 0.5$, $RC >> q_o$ so we take $RC = 100$. Then,

$$W_o = \frac{CR}{N(1 + \gamma/d_o)} = 100/(5 \times 1.166) = 17.142857$$

Hence $p_o = 2/W_o^2 = 0.006806$, $\beta = (d_0/\zeta - 1)/W_o = 0.11666$, and $q_0 = \frac{p_0 + \eta_2}{\eta_1} = 26.80555$, which is indeed small with respect to RC. We can verify that $z = 0.09964430039 + 0.4704656808 * I$ is a zero of (10), and hence the system is unstable.

C. Impact on Stability

In this section we present conclusions drawn from the results presented in previous subsections.

Examining the form of the equation $g(z) = 0$ (whose solutions provide the stability condition), we observe the following.

(i) The stability condition is not a function of the round trip delay. This is seen directly from (6)-(9): if we scale time so that a time unit corresponds to a round trip time, we get a new system of equations that does not involve R.

(ii) From the form of the expression for $g(z)$, we see that for the same values of $C, R, N, \gamma, \eta_1, \eta_2$ and for the same value d_o at equilibrium, the stability region can change according to the choice of the parameters ζ and β. The choice $\beta = 0$ corresponds to a non-dynamic value of d (i.e. a value that does not change with W) in which case $\zeta = d_o$. (Other advantages of dynamic d have already been illustrated in [6] in the context of mobile communications.)

TABLE I
IMPACT OF DYNAMIC CONTROL ON STABILITY

d_o	ζ	$\Rightarrow \beta$	The system is
3	1	0.33 ($W \nearrow$ $d \nearrow$)	stable
3	2	0.08 ($W \nearrow$ $d \nearrow$)	stable
3	3	0 (no control)	unstable
3	4	-0.04 ($W \nearrow$ $d \searrow$)	unstable

(iii) We also see that for fixed $C, R, N, \gamma, \zeta, \beta$ and a fixed queue size q_o, at equilibrium, $g(z)$ will be influenced by η_1 (and hence the stability region). In fact $g(z)$ does not depend on η_2, but note that since we assume that R includes the mean queueing delay (which is proportional to q_o), this means that fixing q_o and η_1 already determines η_2.

In the following, we choose a scenario that shows clearly the impact of using a dynamic delayed ACK factor d on the stability of RED. We fix the parameters used in the numerical example. We set $N = 14$, and vary ζ from 1 to 4. Since we fix d_0, β is deduced. Table I presents the results concerning the stability of the system.

The first observation we make is that when $\zeta = 3 \ \ (= d_0)$, the system is unstable, which means that when there is no dynamic variation of the delay factor d, the system is unstable. However, for $\zeta = 1$ or 2, the system is stable. Besides, β is positive, which means that the linear control is the correct one ($d(t)$ is a non-decreasing function of $W(t)$). When $\zeta = 4$, β is negative and the system is unstable. These results corroborate the fact that with an increase of the TCP window, we should increase the delay factor d. Moreover, using an adequate adaptive control of d, we can improve the stability of the system.

IV. Linear Quadratic Approaches for Flow Control

In this section we study another facet of the interaction between the network and control data, namely, the problem of *optimum* choice (in a sense to be defined later in this section) of the time interval between successive transmissions of information in ATM-like networks. However, unlike the previous study of TCP we do not consider here a precise model for an ATM protocol such as ABR but we employ a specific control theoretic approach that is representative of the ABR rate control mechanism [5]. Our approach uses continuous variables and lies within a linear quadratic framework. Such frameworks have frequently been used in rate-based flow control where the nonlinearity at queue boundaries (corresponding to empty or full buffers) are avoided through tight control. More precisely, we make three simplifying (but realistic) modeling assumptions, and a fourth one on the nature of the information flow:

1. *Fluid approximation.* We replace a discrete number of packets by a continuous fluid. The fluid approximation is justified by the fact that in today's technology, buffering

capabilities are very large (in the order of several thousands) in terms of the number of packets they can store, so that the error introduced by replacing integers by real numbers is small (in fact, negligible) relative to the size of the buffers. This type of approximation is common both in the design of controllers in high-speed telecommunication networks (see e.g. [13]) as well as in the performance evaluation of existing controllers [19], [20].

2. *Linearized dynamics.* We assume that the network has linearized dynamics for the control of queue length; see (11) below. This amounts to neglecting losses due to buffer overflow and also neglecting the boundary effect of an empty queue. The latter is a reasonable simplification given the fact that the controllers that we derive operate in a region close to full throughput utilization, so that the queue length will almost never hit zero. Full utilization is indeed common in the control of ABR switches, see e.g., [16], and can be assured by *regulating* the (controlled) input rate and adapting it to the available capacity. As described below, we set some desirable threshold on the queue length which we attempt to track, precisely so as to avoid large queues (which might lead to losses) or empty queues (which might result in loss of potential throughput). When a control mechanism has a full utilization, then the nonlinearity around zero disappears. For similar models with a single controller, simulations have confirmed [4], [3] that *controlled* linearized models lead to trajectories that are very close to the original one. The other assumption– that buffer overflow will not happen– can be justified through similar arguments, since the optimal control to be derived will be shown to be symmetric with respect to positive or negative deviations around the target queue value.

3. *Bottleneck assumption.* We assume that all performance measures (such as throughput, delays, loss probabilities, etc.) are determined essentially by a bottleneck node. This assumption admits theoretical as well as experimental justifications; see [12].

4. *Information flow.* We assume that information is sent to the controller on the queue length periodically, every τ seconds.

We next introduce the model, where we assume that there is a single bottleneck link, which is used by both the information packets and the data packets, with the former having priority over the latter. The link capacity, C, available to data packets thus depends on the rate, $1/\tau$, of the information packets:

$$C(\tau) = C_{\text{total}} - a/\tau \, ,$$

where C_{total} is the total capacity and a is a proportionality constant. Part of C is also used by uncontrolled inputs, whose total input rate is assumed to be $C_1 + w(t)$ where C_1 is a constant and w is some uncertainty, which could be modeled as a stochastic process with independent increments and zero mean (capturing the stochastic nature of the rate of uncontrolled sources), or some unknown deterministic process with average value zero (we will consider both possibilities in the sequel). The average rate of the controlled source is then $C_2 = C(\tau) - C_1$, and we denote the rate of change around this nominal value to be $u(t)$, which is the control variable. Letting $q(t)$ denote the queue length at the bottleneck link, we then have the queue length dynamics given by

$$dq = u \, \mathrm{d}t + \mathrm{d}w \tag{11}$$

which is *idealized* because the end-point effects have been ignored. The objectives of the flow controller are to ensure that (i) the bottleneck queue size stays around some desired level $\overline{Q}$, and (ii) variations in the rate remain small. The choice of $\overline{Q}$ and the variability around it have a direct impact on loss probabilities and throughput.

We next introduce a shifted version of q: $x(t) := q(t) - \overline{Q}$, and note that (11) can equivalently be written as

$$dx = u\mathrm{d}t + \mathrm{d}w \tag{12}$$

An appropriate local cost function that is compatible with the objectives stated above would be the one that penalizes variations in $x(t)$ and $u(t)$ around *zero* — a candidate for which is the weighted quadratic cost function: $x^2 + ku^2$, where k is some positive weighting parameter.

We now seek a control policy that is *optimal* (in a sense to be clarified shortly) among those which choose $u(t)$ as a function of the queue length at the time instants when new information becomes available; that is, with μ denoting such a policy,

$$u(t) = \mu(x(0), x(\tau), ..., x(n\tau), t),$$

for $t \in [n\tau, \, (n+1)\tau), \;\; n = 0, 1, 2, ...$

We consider two different approaches toward the characterization of the optimum μ, depending on whether w is stochastic or deterministic (but unknown). The former case leads to the LQG (linear quadratic Gaussian) formulation, and the latter to the H$^\infty$ formulation. Both cases are discussed below.

A. LQG model

Here we assume that w is a zero mean Brownian motion with incremental variance intensity r. The expected average cost corresponding to a given policy μ and initial state x is then defined as

$$J(x, \mu, \tau) = \limsup_{T \to \infty} \frac{1}{T} E_x^\mu \left[\int_0^T \left(x^2(t) + k \, u^2(t) \right) dt \right]$$

where E_x^μ is the expectation with respect to the probability measure induced by a policy μ and an initial state x. We first seek to obtain the optimal policy, μ^*, and the corresponding value of the cost, J^*, as a function of x and τ:

$$J^*(x, \tau) := \min_\mu J(x, \mu, \tau) = J(x, \mu^*, \tau). \tag{13}$$

The following theorem whose proof is in the Appendiix states that result.

Theorem 4.1: The optimal cost value of the LQG problem is independent of the initial state x and is given by

$$J^*(\tau) = \frac{\tau}{2} r + r\sqrt{k} \, , \tag{14}$$

and the unique policy that attains the minimum is given by

$$u^*(t) = \mu^*(x(n\tau), t) = -\frac{1}{\sqrt{k}} \, e^{-\frac{t - n\tau}{\sqrt{k}}} \, x(n\tau) \tag{15}$$

for $t \in [n\tau, \, (n+1)\tau)$.

We note that as long as τ^2 is much smaller than $\sqrt{k}$, the value is quite insensitive to changes in the spacing τ. On the other hand, when it is much larger than $\sqrt{k}$, we see that the spacing of information packets has a big impact on the performance: the cost grows quadratically in the spacing.

Optimal spacing of information for the LQG model

Having determined the optimal control policy for a given parameter τ of spacing of information packets, we next turn to finding the *optimum* value of τ. In order to bring in some tradeoff in the choice of τ, in addition to the optimum cost $J^*(\tau)$ given by (14), we assume that the flow control has some utility $U(\tau)$ for the average throughput $C_{\text{total}} - a/\tau$ available for data packets. We will consider two types of functions:

1) A linear utility: $U(\tau) = C_{\text{total}} - a/\tau$.
2) A logarithmic utility: $U(\tau) = \log(C_{\text{total}} - a/\tau)$.

The overall cost to be minimized in both cases will be taken as $Z(\tau) = J^*(\tau) - \beta\, U(\tau)$, where β is some positive weighting parameter that characterizes the level of tradeoff between sending data packets and information packets. Under both types of utility functions, the function $Z(\tau)$ is a strictly convex function of τ on $(0, \infty)$, and admits a unique minimum. This solution is given in the theorem below.

Theorem 4.2: The optimal information spacing, τ^*, is given by the following under each of the two utility functions:
(i) under linear utility: $\tau^* = \sqrt{2\beta a/r}$
(ii) under logarithmic utility:

$$\tau^* = \frac{a}{2C_{\text{total}}}\left[1 + \sqrt{1 + 8(\beta/ar)C_{\text{total}}}\right]$$

Proof. The expressions for τ^* follow by setting the derivative of Z to zero, and solving for the positive roots of the resulting quadratic equations. ◇

B. A worst-case (H^∞) approach

The second modeling paradigm does not ascribe any statistics to the uncertainty process, w, but instead associates an adversary (player) with the uncertainty, who chooses the process in such a way so as to maximize its effect on the deviations from the nominal queue length and the nominal rate, which is the framework of H^∞-optimal control [11]. More precisely, taking the initial deviation on the queue length as $x(0) = 0$, and letting $v(t) := dw(t)/dt$, we define

$$V(\mu, \tau) = \sup_v \frac{L(\mu, v, \tau)}{\|v\|^2} \tag{16}$$

where $L(\mu, v, \tau) = \int_0^\infty \left(x^2(t) + ku^2(t)\right) dt$, and where μ is a controller of the type defined earlier, and $\|v\|$ is the L^2 norm of v over the interval $[0, \infty)$, that is, $\|v\|^2 = \int_0^\infty v^2(t)\, dt$. One then wishes to find the μ that minimizes $V(\mu, \tau)$ over the class of policies that depend on sampled (in time) values of x, as identified in the previous subsection; denote the infimum by $(\gamma^*)^2$, that is $\inf_\mu V(\mu, \tau) =: (\gamma^*)^2$.

For $\gamma > \gamma^*$, let us introduce the γ-parameterized soft-constrained cost function $L_\gamma(\mu, v, \tau) := L(\mu, v, \tau) - \gamma^2 \|v\|^2$, and consider a two-player dynamic game where L_γ is to be minimized by Player 1 (controlling μ) and maximized by Player 2 (controlling v). For this game, we can take $x(0) = x_0$, where x_0 is known but not necessarily zero. If there exists some policy μ^* that minimizes $V(\mu, \tau)$, then it has the property [11]:

$$\sup_v L_{\gamma^*}(\mu^*, v, \tau) = \inf_\mu \sup_v L_{\gamma^*}(\mu, v, \tau),$$

The quantity above is the upper value of the zero-sum dynamic game with kernel L_{γ^*}, which is in fact a quadratic function of x_0, and hence equal to zero if $x_0 = 0$. It can actually be shown that for any $\gamma > \gamma^*$, the upper value of the game with parameterized kernel L_γ is also quadratic in x_0 (and hence is equal to zero if $x_0 = 0$), and for $\gamma < \gamma^*$, its upper value is infinite. Hence, γ^* is the "smallest" positive scalar γ for which the zero-sum game with kernel L_γ has a finite upper value. Furthermore, in this result, and the computation of the upper value, what information v is endowed with does not play any role; it could be closed loop, or sampled data, or open loop.

In view of the above result, instead of obtaining μ^*, we will in fact solve for a parameterized class of controllers, $\{\mu^\gamma,\ \gamma > \gamma^*\}$, where μ^γ is obtained from $\sup_v L_\gamma(\mu^\gamma, v, \tau) = \inf_\mu \sup_v L_\gamma(\mu, v, \tau)$. The controller μ^γ will clearly have the property that it ensures a performance level no worse than γ^2 for the index adopted in (16), i.e. the attenuation is bounded by

$$\frac{\{L(\mu^\gamma, v, \tau)\}^{1/2}}{\|v\|} \le \gamma \quad \text{for all } v. \tag{17}$$

This solution is now given in the theorem below.

Theorem 4.3: For the problem formulated above, given a fixed $\tau > 0$ and a fixed attenuation level $\gamma > \sqrt{k}$, an admissible control policy μ^γ that leads to satisfaction of the disturbance attenuation bound (17) is

$$\begin{aligned} u^\gamma(t) &= \mu^\gamma(x(n\tau), t) \\ &= -\frac{\gamma}{k\sqrt{(\gamma^2/k) - 1}}\, e^{-\frac{1}{\gamma}\sqrt{\frac{\gamma^2}{k} - 1}\,(t - n\tau)} x(n\tau) \end{aligned}$$

for $t \in [n\tau,\ (n+1)\tau)$, $n = 0, 1, 2, \ldots$, provided that the following inequality (tradeoff) between τ and γ holds:

$$\frac{\tau}{\gamma} + \arctan \frac{1}{\sqrt{(\gamma^2/k) - 1}} \le \frac{\pi}{2} \tag{18}$$

The corresponding worst-case uncertainty process, v^γ, is given in closed-loop form by

$$v^\gamma(t) = \nu^\gamma(x(t)) = \frac{1}{\gamma\sqrt{(\gamma^2/k) - 1}}\, x(t),\ t \ge 0$$

and under (μ^γ, ν^γ) the queue dynamics

$$\frac{dq}{dt} = u + v,\ q(0) = q_0$$

is asymptotically stable, with $q(t)$ exponentially converging to $\bar{Q}$ for all q_0.

The largest value of τ (that is the coarsest spacing of information) that leads to attainment of a given level, $\gamma > \sqrt{k}$, of disturbance attenuation is

$$\tau = \frac{\pi}{2}\gamma - \gamma \arctan \frac{1}{\sqrt{(\gamma^2/k) - 1}}, \tag{19}$$

and conversely, for a given $\tau > 0$, the optimum level of disturbance attenuation, $\gamma^*(\tau)$, is obtained as the unique

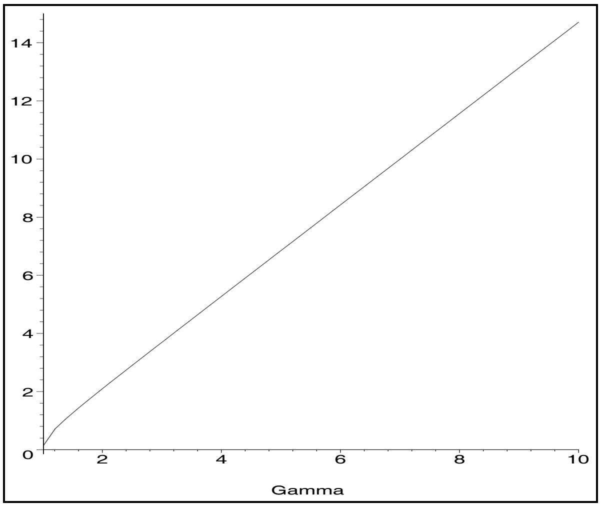

Fig. 7. Maximum value of τ from (19) as a function of γ.

solution of (19), and the corresponding controller is $\mu^{\gamma*}$, which is well defined as the limit of μ^{γ} as $\gamma \downarrow \gamma^*$.

Proof. See the Appendix. $\diamond$

Remark. It is worth noting that for each $\tau > 0$ the limit $\lim_{\gamma\to\infty} \mu^{\gamma} =: \mu^{\infty}$ is a well-defined controller, and is precisely the controller (15) obtained for the Gaussian model. $\diamond$

Remark. It directly follows from H^{∞}-optimal control theory ([11], [10]) that τ in (19) is a monotonically increasing function of γ–a property that can also be proven directly by differentiating the right-hand-side of (19) with respect to γ, and seeing that it is indeed positive:

$$\frac{d\tau}{d\gamma} = \frac{\pi}{2} - \arctan\frac{1}{\sqrt{(\gamma^2/k)-1}} + \frac{1}{\sqrt{(\gamma^2/k)-1}} > 0$$

Figure 7 depicts this property for $k = 1$. The spacing (y-axis) is seen to be concave increasing in the required attenuation, and asymptotically it grows linearly with it.

Optimum spacing of information packets under H-infinity Given the constraint (19) linking τ and γ, we want to find an *optimal* value of τ under a criterion that shows tradeoff between minimization of the level of disturbance attenuation, γ, and maximization of utility for capacity available for data traffic, that is $C_{\text{total}} - a/\tau$. Hence, the objective function to be minimized, subject to (19), is $Z(\tau;\gamma) = F(\gamma) - \beta U(\tau)$ where $F(\cdot)$ is an increasing differentiable function on $(0,\infty)$, and $U(\cdot)$ the utility function, which is strictly concave increasing, and differentiable everywhere on $(0,\infty)$. Furthermore, $\beta > 0$ is a tradeoff parameter.

Differentiating $Z(\tau;\gamma)$ with respect to τ, by also (implicitly) taking into account the dependence of γ on τ through the constraint, we have $F'(\gamma)\,\gamma'(\tau) = \beta U'(\tau)$ where 'prime' ($'$) denotes derivative with respect to the argument of the corresponding function. Note that because of the conditions imposed on F and U, and the one-to-one relationship between $\gamma > \sqrt{k}$ and $\tau > 0$ through the constraint (19), the stationarity condition above is not only necessary but also sufficient. Further, total-differentiating the constraint equation (19) with respect to τ, we obtain: $\left(\frac{\tau}{\gamma} + \frac{1}{\sqrt{(\gamma^2/k)-1}}\right)\gamma'(\tau) = 1$. Using the expression for $\gamma'(\tau)$ from the preceding equation in the stationarity condition leads to $F'(\gamma) = \beta\left(\frac{\tau}{\gamma} + \frac{1}{\sqrt{(\gamma^2/k)-1}}\right)U'(\tau)$. This will have to be solved, along with (19), for the optimum value of τ. Note that τ from (19) can be substituted into the equation above, to lead to a single equation for γ, which can then be solved (numerically) for the optimum γ (and in turn for the optimum τ) corresponding to different choices of U (as in the LQG case) and F.

V. Conclusion

We have studied in this paper the problem of *optimum* information transfer in communication networks, using three specific models. This study provides us with insight into the interaction (and tradeoffs) between network performances and control information that should be sent back to the sources in congestion control mechanisms. We have deliberately chosen simple models in order to obtain tractable results. Indication on the validity and accuracy of our first model has been obtained using simulations as well as numerical computations involving more precise models. We consider both IP-based and ATM-based networks.

In IP-based networks, we have shown that the rate of information can have quantitative impacts on the system, as conveying information consumes resources which would otherwise be available for the data packets. In particular, delaying control information (ACK packets) results in increasing TCP throughput by 5% to 50% in our numerical computations. However, above some spacing threshold, the gain is worthless. The rate of information can also have qualitative effects, in the sense that it may impact the stability of a congestion control algorithm. Especially, we have found that adapting dynamically the frequency (spacing) of control information might be necessary to avoid network instability.

In ATM-based networks, control information are importantly deployed in the rate control mechanism of ABR class of service. We have used a specific control theoretic approach that is representative of an ABR rate control in order to study the optimum frequency. We have shown that the optimal spacing between control information can be found by solving numerically a trade-off equation that relates the cost of control information and the capacity available to data. We have shown how to derive this optimal spacing τ^* for LQG and H-infinity design problems. In the former case, we have found an explicit expression for τ^*, and in the latter case as a root of the trade-off equation.

We are currently applying similar methodologies in order to investigate the gain we can possibly obtain from separating control and data information. So far, we have considered that control and data share the same queue. Preliminary results show that it is worthy to separate between them only under high load.

VI. Appendix

A. Computation of f for Section II

Let $s = re^{\imath\theta} = r\cos\theta + \imath r\sin\theta$, and $\beta = \zeta/Z$ and $L = \lambda/C$, then we can write

$$f \;=\; \frac{1}{2\pi\imath}\oint_{D_r}\frac{1}{G(s)\,s^{K-1}}\,\mathrm{d}s$$

$$= \frac{1}{2\pi\, r^{K-2}} \int_0^{2\pi} \frac{e^{-\imath\theta(K-2)}\, \overline{G(\theta)}}{|G(\theta)|^2}\, \mathrm{d}\theta \qquad (20)$$

where

$$G(s) = \frac{1}{d+1} e^{-\frac{\zeta\lambda(1-s)}{ZC}} + \frac{d}{d+1} e^{-\frac{\lambda(1-s)}{C}} - s,$$

$$\begin{aligned}
&e^{-\imath\theta(K-2)}\, \overline{G(\theta)} = \\
&\cos(\theta K - 2\theta) \left(\frac{e^{-\beta L(1-r\cos\theta)} \cos(\beta Lr\sin\theta)}{d+1} \right. \\
&\left. + \frac{de^{-L(1-r\cos\theta)} \cos(Lr\sin\theta)}{d+1} - r\cos\theta \right) \\
&+ \sin(\theta K - 2\theta) \left(- \frac{e^{-\beta L(1-r\cos\theta)} \sin(\beta Lr\sin\theta)}{d+1} \right. \\
&\left. - \frac{de^{-L(1-r\cos\theta)} \sin(Lr\sin\theta)}{d+1} + r\sin\theta \right) \\
&+ \imath \left[-\sin(\theta K - 2\theta) \left(\frac{e^{-\beta L(1-r\cos\theta)} \cos(\beta Lr\sin\theta)}{d+1} \right. \right. \\
&\left. + \frac{de^{-L(1-r\cos\theta)} \cos(Lr\sin\theta)}{d+1} - r\cos\theta \right) \\
&+ \cos(\theta K - 2\theta) \left(- \frac{e^{-\beta L(1-r\cos\theta)} \sin(\beta Lr\sin\theta)}{d+1} \right. \\
&\left. \left. - \frac{de^{-L(1-r\cos\theta)} \sin(Lr\sin(\theta))}{d+1} + r\sin\theta \right) \right]
\end{aligned}$$

In view of this,

$$\begin{aligned}
|G(\theta)|^2 = -\Bigg(& -d^2 e^{2L(-1+r\cos(\theta))} - r^2 \\
&+ 2\, e^{\beta L(-1+r\cos(\theta))} \cos(\beta Lr\sin(\theta))\, r\cos(\theta) \\
&+ 2\, e^{\beta L(-1+r\cos(\theta))} \sin(\beta Lr\sin(\theta))\, r\sin(\theta) \\
&- 2\, e^{L(-1+r\cos(\theta))(\beta+1)} \sin(\beta Lr\sin(\theta))\, d\sin(Lr\sin(\theta)) \\
&- 2\, e^{L(-1+r\cos(\theta))(\beta+1)} \cos(\beta Lr\sin(\theta))\, d\cos(Lr\sin(\theta)) \\
&+ 2\, e^{\beta L(-1+r\cos(\theta))} \cos(\beta Lr\sin(\theta))\, r\cos(\theta)\, d \\
&+ 2\, d^2 e^{L(-1+r\cos(\theta))} \cos(Lr\sin(\theta))\, r\cos(\theta) \\
&+ 2\, de^{L(-1+r\cos(\theta))} \cos(Lr\sin(\theta))\, r\cos(\theta) \\
&+ 2\, e^{\beta L(-1+r\cos(\theta))} \sin(\beta Lr\sin(\theta))\, r\sin(\theta)\, d \\
&+ 2\, d^2 e^{L(-1+r\cos(\theta))} \sin(Lr\sin(\theta))\, r\sin(\theta) \\
&+ 2\, de^{L(-1+r\cos(\theta))} \sin(Lr\sin(\theta))\, r\sin(\theta) \\
&- e^{2\beta L(-1+r\cos(\theta))} - r^2 d^2 - 2\, r^2 d \Bigg) \Bigg/ (d+1)^2
\end{aligned}$$

We can check that the real parts of $e^{-i\theta(K-2)}\,\overline{G(\theta)}$ and $|G(\theta)|^2$ are even functions in θ, and that the imaginary part of $e^{-i\theta(K-2)}\,\overline{G(\theta)}$ is odd. Then, (20) is reduced to:

$$f = \frac{1}{\pi\, r^{K-2}} \int_0^{\pi} \frac{Real(e^{-i\theta(K-2)}\,\overline{G(\theta)})}{|G(\theta)|^2}\, \mathrm{d}\theta \qquad (21)$$

Choosing a radius r for the integration. Note that when $\beta = 1$, i.e. $\zeta = Z$, then $L = \rho$ and the queueing system is reduced to $M/D/1/K$. It is possible to compute a radius in this case [26]. Let $r(\rho)$ denote such a radius. We can show that $r(\rho)$ can also be used to compute the integral in the case when $\beta < 1$. $\diamond$

B. Proof of Theorem 4.1

Let us first consider the version of the problem with finite-horizon cost

$$J(x,\mu;T) = \frac{1}{T}\int_0^T E\big(x^2(t) + ku^2(t)\big)\mathrm{d}t$$

where T is a multiple of τ. Under the given scalar dynamics (12), $J(x,\mu;t_f)$ can equivalently be written as [14]

$$\frac{1}{T}\left[\int_0^T E|u(t) + \frac{1}{k}p(t;T)x(t)|^2 k\mathrm{d}t + m(0;T)\right],$$

where p and m satisfy (and are unique solutions of) the ODEs: $\dot p + 1 - \frac{1}{k}p^2$, $\;p(T,T) = 0$ and $\dot m + pr = 0$, $\;m(T,T) = 0$. Taking $T = N\tau$ for some integer N, a further equivalent expression for $J(x,\mu;T)$ is

$$\sum_{i=0}^{N-1} \frac{1}{N\tau}\int_{n\tau}^{(n+1)\tau} E|u + \frac{1}{k}px|^2 k\mathrm{d}t + \frac{1}{T}m(0;T).$$

Given that on the time interval $[n\tau,(n+1)\tau)$, the control u depends on $x(n\tau)$ (as well as the past sampled values of x), $x(t)$ can be written for $t \in [n\tau,(n+1)\tau)$ as $x(t) = \hat x(t) + e(t)$, where $\frac{\mathrm{d}\hat x(t)}{\mathrm{d}t} = u(t)$, $\;\hat x(n\tau) = x(n\tau)$, $\mathrm{d}e(t) = \mathrm{d}w(t)$, $\;e(n\tau) = 0$ and $\hat x$ and e are statistically independent. In view of this, $J(x,\mu;T)$ can equivalently be written as

$$\sum_{i=0}^{N-1} \frac{1}{N\tau}\int_{n\tau}^{(n+1)\tau} E\Big[|u + \frac{1}{k}p\hat x|^2 k + p^2(t,T)e^2(t)\Big]\mathrm{d}t + \frac{1}{T}m(0;T).$$

It readily follows from the above form of $J(x,\mu;T)$ that the unique minimizing control for the finite-horizon sampled-data stochastic control problem is

$$u(t) = \mu(x(n\tau),t) = -(1/k)p(t,N\tau)\,\hat x(n\tau),$$

and the corresponding minimum value is

$$\frac{1}{N}\sum_{i=0}^{N-1}\frac{1}{\tau}\int_{n\tau}^{(n+1)\tau} p^2(t,N\tau)(t-n\tau)\mathrm{d}t + \frac{1}{N\tau}m(0;N\tau).$$

As $N \to \infty$, $p(t,N\tau)$ converges uniformly in t to the constant $\bar p = \sqrt{k}$, and $(1/N\tau)m(0,N\tau) \to \bar p r$. Furthermore,

$$\frac{1}{N}\sum_{i=0}^{N-1}\frac{1}{\tau}\int_{n\tau}^{(n+1)\tau} p^2(t,N\tau)(t-n\tau)\mathrm{d}t \to \frac{\tau}{2}r$$

Given that with $u = -(1/\sqrt{k})\hat x(t)$, $\hat x(t) = e^{-\frac{t-n\tau}{\sqrt{k}}}\,x(n\tau)$, $t \in [n\tau,(n+1)\tau)$, Theorem 4.1 readily follows. $\diamond$

C. Proof of Theorem 4.3

This result is a direct application of the general theory in [10] on H^∞-optimal control of linear systems under deterministic sampled-data information. With instantaneous state feedback, a control that guarantees a level of disturbance attenuation (DA) γ is given (for the scalar system at hand) by: $u_\gamma(t) = -(1/k)Z_\gamma x(t)$, where Z_γ is the positive solution of the generalized algebraic Riccati equation: $\big(\frac{1}{k} - \gamma^{-2}\big)Z^2 - 1 = 0$, provided that $\gamma^2 > k$. If $\gamma < \sqrt{k}$, then no control can guarantee such a level of DA, and the upper value of the associated soft-constrained game is infinite. When the upper value is bounded, then the worst-case disturbance is

$v_\gamma(t) = \gamma^{-2} Z_\gamma x(t)$. If the information available is sampled data (say, with frequency $1/\tau$), then the achievable level of DA is higher, and is determined by the condition of solvability of a series of finite-horizon Riccati differential equations (RDEs). On the time interval $[n\tau, (n+1)\tau)$, the relevant RDE is

$$\dot{S} + \gamma^{-2} S + 1 = 0\,, \quad S((n+1)\tau) = Z_\gamma$$

We require that this RDE has no finite escape on the given interval, for which the condition is precisely (18), that is

$$\frac{\tau}{\gamma} + \arctan \frac{1}{\sqrt{(\gamma^2/k) - 1}} \leq \frac{\pi}{2}\,.$$

If this condition holds, then the RDE admits a unique positive solution which is finite over $t \in (n\tau, (n+1)\tau)$, and if the left-hand side is larger than $\pi/2$, then the RDE will have a finite escape before the solution reaches $n\tau$. Note that the condition is independent of n, and thus of the specific interval on which the RDE is defined, and hence it applies to all such RDEs. The final observation is that, again following [10], a sampled-data control that guarantees the achievable level of DA is obtained from $u_\gamma(t) = -(1/k) Z_\gamma x(t)$ by solving for $x(t)$ on the interval $(n\tau, (n+1)\tau]$ in terms of $x(n\tau)$, from

$$\dot{x} = u_\gamma(t) + v_\gamma(t) = \left(\gamma^{-2} - \frac{1}{k}\right) Z_\gamma x(t)\,,$$

and substituting into $u_\gamma(t)$, which leads to the $\mu^\gamma(x(n\tau), t)$ given in the theorem. This completes the proof of the theorem, including asymptotic stability of the resulting queue dynamics (which directly follows from [10]). $\diamond$

REFERENCES

[1] K. Abratchenkov, U. Ayesta, E. Altman, P. Nain, and C. Barakat, "The effect of router buffer size on the TCP performance," in *Proc. LONIIS workshop on Telecommunication Networks and Teletraffic Theory*, St. Petersburg, Russia, pp. 116–121, Jan. 2002.

[2] E. Altman and T. Başar, "Multi-user rate-based flow control," *IEEE Trans. Commun.*, vol. 46, no. 7, pp. 940–949, 1998.

[3] E. Altman, F. Baccelli, and J. C. Bolot, "Discrete-time analysis of adaptive rate control mechanisms," *High Speed Networks and Their Performance*, H. G. Perros and Y. Viniotis Eds., North Holland, pp. 121–140, 1994.

[4] E. Altman and T. Başar, "Optimal rate control for high speed telecommunication networks," in *Proc. 34th IEEE Conf. on Decision and Control*, New Orleans, LA, Dec. 1995.

[5] E. Altman, T. Başar, and R. Srikant, "Robust rate control for ABR sources," in *Proc. IEEE INFOCOM*, San Francisco, CA, USA, March 1998.

[6] E. Altman and T. Jiménez, "Novel delayed ACK techniques for improving TCP performance in multihop wireless networks," in *Proc. PWC 2003 Personal Wireless Communications*, Sept. 2003, Venice, Italy.

[7] R. de Oliveira and T. Braun, "A dynamic adaptive acknowledgment strategy for TCP over multihop wireless networks," in *Proc. IEEE Infocom*, Miami, Mar. 2005.

[8] U. Ayesta, K. Avrachenkov, E. Altman, C. Barakat, and P. Dube, "Simulation analysis and fixed point approach for multiplexed TCP flows," in *Proc. PFLDnet'03: Workshop on Protocols for Fast Long-Distance Networks*, CERN-Geneva, Switzerland, Feb. 2003.

[9] C. Barakat and E. Altman, "On ACK filtering on a slow reverse channel," in *Proc. First International Workshop on Quality of Future Internet Services (QOFIS)*, Berlin, Germany, Sept. 2000.

[10] T. Başar, "Optimum H$^\infty$ designs under sampled state measurements," *Systems & Control Lett.*, vol. 16, pp. 399–409, 1991.

[11] T. Başar and P. Bernhard, *H^∞-Optimal Control and Relaxed Minimax Design Problems: A Dynamic Game Approach*, 2nd edition. Boston, MA: Birkhäuser, 1991.

[12] J.-C. Bolot, "End-to-end packet delay and loss behavior in the Internet," in *Proc. ACM Sigcomm '93*, pp. 289–298, San Francisco, CA, Sept. 1993.

[13] J.-C. Bolot, A. U. Shankar, "Analysis of a fluid approximation to flow control dynamics," in *Proc. IEEE INFOCOM '92*, pp. 2398–2407, Florence, Italy, May 1992.

[14] W. H. Fleming and H. M. Soner, *Controlled Markov Processes and Viscosity Solutions*. New York: Springer Verlag, 1993.

[15] C. Hollot, V. Misra, D. Towsley, and W.-B. Gong, "A control theoretic analysis of RED," CMPSCI Technical Report TR 00-41, July 2000. Also in *Proc. IEEE INFOCOM'01*.

[16] R. Jain, "ERICA+," contribution to ATM Forum/95-1346.

[17] Y. L. Raphael Rom and M. Sidi, "Analysis of discarding policies in high-speed networks," *IEEE J. Select. Areas Commun.*, vol. 16, no. 5, pp. 764–777, June 1998.

[18] M. A. Labrador and S. Banerjee, "Performance analysis of generalized selective packet discarding schemes," *Telecommun. Systems*, vol. 21, no. 1, pp. 87–101 Sept. 2002.

[19] H. Ohsaki, M. Murata, H. Suzuki, C. Ikeda, and H. Miyahara, "Rate-based control for ATM Networks," *Computer Commun. Review, ACM-Sigcomm,* special issue in ATM, ed. R. Jain and K. Y. Siu, vol. 25, no. 2, pp. 60–71, 1995.

[20] M. Ritter, "Network buffer requirements of the rate-based control mechanism for ABR services," in *Proc. IEEE INFOCOM '96*, San Francisco, USA, Mar. 1996.

[21] J. Padhye, V. Firoiu, D. Towsley, and J. Kurose, "Modeling TCP throughput: a simple model and its empirical validation," *IEEE/ACM Trans. Networking*, vol. 8, no. 2, pp. 133–145, Apr. 2000.

[22] J. W. Cohen *The Single Server Queue*. North Holland Publishing Company, 1982.

[23] E. Altman, C. Barakat, and V. M. Ramos, "Queuing analysis of simple FEC schemes for Voice over IP," *Commun. Networks*, vol. 39, no. 2, pp. 185–206, June 2002.

[24] M. Allman, "TCP congestion control with appropriate byte counting (ABC)," request for comments: 3465, Feb. 2003.

[25] M. Allman, V. Paxson, and W. Stevens, "TCP congestion control," request for comments: 2581, Apr. 1999.

[26] S. Alouf, "Parameter estimation and performance analysis of several network applications, Ph.D. thesis, Department of Sciences, University of Nice Sophia-Antipolis, Nov. 2002.

[27] G. F. Franklin, J. D. Powell, and A. Emami-Naeini, *Feedback Control of Dynamic Systems*. Addison-Wesley, 1995.

Eitan Altman received the B.Sc. degree in electrical engineering (1984), the B.A. degree in physics (1984) and the Ph.D. degree in electrical engineering (1990), all from the Technion-Israel Institute, Haifa. In (1990) he further received his B.Mus. degree in music composition in Tel-Aviv university. Since 1990, he has been with INRIA (National research institute in informatics and control) in Sophia-Antipolis, France. His current research interests include performance evaluation and control of telecommunication networks and in particular congestion control, wireless communications and networking games. He is in the editorial board of the scientific journals: WINET, JDEDs and JEDC, and served in the journals Stochastic Models, COMNET, SIAM SICON. He has been the general cairman and the (co)chairman of the program committee of several international conferences and workshops (on game theory, networking games and mobile networks). He has published more than 140 papers in international refereed scientific journals. More informaion can be found at http://www.inria.fr/mistral/personnel/Eitan.Altman/me.html

Tamer Başar (S'71-M'73-SM'79-F'83) is with the University of Illinois at Urbana-Champaign (UIUC), where he holds the positions of Swanlund Endowed Chair, Center for Advanced Study Professor of Electrical and Computer Engineering, and Research Professor at the Coordinated Science Laboratory and the Information Trust Institute. He received the B.S.E.E. degree from Robert College, Istanbul, in 1969, and the M.S., M.Phil, and Ph.D. degrees from Yale University during the period 1970-1972. He joined UIUC in 1981 after holding positions at Harvard University and Marmara Research Institute (Turkey). He has published extensively in systems, control, communications, and dynamic games, and has current research interests in modeling and control of communication networks; control over heterogeneous networks; resource allocation, management and pricing in networks; mobile computing; security issues in computer networks; and robust identification, estimation and control.

Dr. Başar is the Editor-in-Chief of *Automatica*, Editor of the *Birkhauser Series on Systems & Control*, Managing Editor of the *Annals of the International Society of Dynamic Games* (ISDG), and member of editorial and advisory boards of several international journals in control, wireless networks, and applied mathematics. He has received several awards and recognitions over the years, among which are the Medal of Science of Turkey (1993); Distinguished Member Award (1993), Axelby Outstanding Paper Award (1995), and Bode Lecture Prize (2004) of the IEEE Control Systems Society (CSS); Millennium Medal of IEEE (2000); Tau Beta Pi Drucker Eminent Faculty Award of UIUC (2004); the Outstanding Service Award (2005) and the Giorgio Quazza Medal (2005) of the International Federation of Automatic Control (IFAC); and the Richard E. Bellman Control Heritage Award of the American Automatic Control Council (2006). He is a member of the National Academy of Engineering (of USA), a member of the European Academy of Sciences, a Fellow of IEEE, a Fellow of IFAC, a past president of CSS, and the founding president of ISDG.

Naceur Malouch is an assistant professor at the University Pierre et Marie Curie Paris 6, France since 2003, and a researcher in the Network and Performance Analysis team at LIP6. He received the Ph.D. in computer science from the University of Nice Sophia-Antipolis, France in January 2003. During his ph.d studies, he was with the MISTRAL group at INRIA Sophia-Antipolis. From 2002 to 2003, he was assistant professor at the "Ecole Superieure en Sciences Informatiques" (ESSI), University of Nice Sophia-Antipolis. He is the responsible of the European CELTIC BOSS project and the national MosoMuso project for LIP6. He has been the responsible of the Research program of the Master degree from 2005 to 2007. His publications are in the area of transport protocols, overlay multicast, low complexity high scalability Quality of Service (QoS) and performance over wireless Mesh networks. Doctor Malouch was a cooperative visitor at IBM T.J. Watson Research Center from June to August 2001, Naceur Malouch has served as a TPC member in the ad-hoc and wireless networks symposium of GLOBECOM2007. He is a TPC memeber for ICCCN'2008, PIMRC'2008 and GlobeCom'2008. He is a publicity co-chair of sigmetrics'2008.