On the trade-off between control rate and congestion in single server systems

Krishna Jagannathan, Eytan Modiano, Lizhong Zheng
Massachusetts Institute of Technology
Cambridge, MA 02139
Email: {krishnaj, modiano, lizhong}@mit.edu

***Abstract*—The goal of this paper is to characterize the trade-off between the rate of control and network congestion for flow control policies. We consider a simple model of a single server queue with congestion-based flow control. The input rate at any instant is decided by a flow control policy, based on the queue occupancy. We identify a simple 'two threshold' control policy, which achieves the best possible congestion probability, for any rate of control. We show that in the absence of control channel errors, the control rate needed to ensure the optimal decay exponent for the congestion probability can be made arbitrarily small. However, if control channel errors occur probabilistically, we show the existence of a critical error probability threshold beyond which the congestion probability undergoes a drastic increase due to the frequent loss of control packets. Finally, we determine the optimal amount of error protection to apply to the control signals by using a simple bandwidth sharing model.**

I. Introduction

Congestion control is an essential part of any practical queueing system. There are two basic ways of mitigating congestion in queueing networks: flow control and resource allocation. Flow control involves controlling the rate of the incoming exogenous traffic to a queue, depending on its congestion level, see [1] for example. Resource allocation on the other hand, involves assigning larger service rates to the queues that are congested, and vice-versa [2], [3]. Most systems use a combination of the two methods to control congestion in the network, and to achieve various performance objectives [4], [5].

Since the queue lengths can vary widely over time in a dynamic network, effective congestion control typically requires queue length information to be conveyed to the flow controllers, which adapt their rates to the varying queues. This queue information can be thought of as being a part of the inevitable protocol and control overheads in a network. Historically, Gallager was among the first to address the important question of protocol overhead in communication networks. In his seminal paper [6], he derives information theoretic lower bounds on the amount of protocol information needed for network nodes to keep track of source and destination addresses, as well as message starting and stopping times.

This paper deals with the basic question of how often congestion notifications need to be sent in order to effectively control congestion. The control information in response to congested queues takes the form of resource management (RM) cells in ATM networks. In TCP networks, packet drops occur when buffers are about to overflow, and this in turn leads to a reduction in the window size and packet arrival rate. Active queue management schemes like Random Early Detection (RED) are designed to pro-actively prevent congestion by randomly dropping some packets *before* the buffers reach the overflow limit [7].

Speaking intuitively, if the flow controller has very accurate information about the congestion level in the system, congestion control can be performed very effectively by adapting the input rates appropriately. However, furnishing the controller with accurate queue length information requires significant amount of control. Further, frequent congestion notifications may also have undesirable packet drops in systems that employ TCP-like mechanisms. Therefore, it is of interest to characterize how frequently congestion notifications need to be employed, in order to achieve a certain congestion control objective. We do not explicitly model packet drops in this paper, but instead associate a cost with each congestion notification. This cost is incurred either because of the ensuing packet drops that may occur in practice, or might simply reflect the resources needed to communicate the control signals.

In the first part of this paper, we consider a single server queue with a constant service rate and two possible arrival rates. In spite of being very simple, such a system gives us enough insights into the key issues involved in the flow control problem. The two input rates may correspond to different quality of service offerings of an internet service provider, who allocates better service when the network is lightly loaded but throttles back on the input rate as congestion builds up; or alternatively to two different video streaming qualities where a better quality is offered when the network is lightly loaded.

The arrival rate at a given instant is chosen by a flow control policy, based on the queue length information obtained from a queue observer. We identify a simple 'two-threshold' flow control policy and derive the corresponding tradeoff between the rate of control and congestion probability in closed form. We show that the two threshold policy achieves the best possible decay exponent (in the buffer size) of the congestion probability for arbitrarily low rates of control. An equivalent

This work was supported by NSF ITR grant CCR-0325401, by NSF grant CNS-0626781, and by ARO Muri grant number W911NF-08-1-0238

result for the resource allocation problem was derived earlier in [8]. Although we mostly focus on the two threshold policy owing to its simplicity, we also point out that the two threshold policy can be easily generalized to resemble the RED queue management scheme. Further, our results for the two threshold policy can also be re-derived for the more general RED-like scheme.

In addition to identifying and analyzing the two threshold policy, we make the following contributions in this paper. First, we characterize the impact of control-channel errors on the congestion control performance of the two threshold policy. We assume a probabilistic model for the errors on the control channel, and show the existence of a critical error probability, beyond which the errors in receiving the control packets lead to an exponential worsening of the congestion probability. However, for error probabilities below the critical value, the congestion probability is of the same exponential order as in a system with an error free control channel.

Second, we determine the optimal apportioning of bandwidth between the control signals and the server in order to achieve the best congestion control performance. In particular, if the control channel has a probability of error greater than the critical value, error protection can be applied to the control signals to mitigate the exponential worsening of the congestion probability. However, adding excessive error protection to the control signals can throttle the available service rate by consuming too much bandwidth. We determine the optimal amount of error protection to apply to the control signals by using a simple bandwidth sharing model.

The remainder of this paper is organized as follows. Section II introduces the system model, and the key parameters of interest in the design of a flow control policy. In Section III, we introduce and analyze the two threshold policy. In Section IV, we investigate the effect of control channel errors on the congestion control performance of the two threshold policy. Section V deals with the problem of optimal bandwidth allocation for control signals in a control error prone system.

II. Preliminaries

A. System Description

Let us first describe a simple model of a queue with congestion based flow control. Fig. 1 depicts a single server queue with a constant service rate μ. We assume throughout that the packet sizes are exponentially distributed with mean 1. An observer watches the queue evolution and sends control information to the flow controller, which changes the input rate $A(t)$ based on the control information it receives. The purpose of the observer-flow controller subsystem is to change the input rate so as to control congestion in the queue.

For analytical simplicity, we assume that the input process at any instant is Poisson with one of two distinct possible rates, $A(t) \in \{\lambda_1, \lambda_2\}$, where $\lambda_2 < \lambda_1$ and $\lambda_2 < \mu$. Physically, this model may be motivated by a DSL-like system, wherein a minimum rate λ_2 is guaranteed, but higher transmission rates might be intermittently possible, as long as the system is not congested. The congestion notifications are sent by the observer in the form of information-less control packets. Upon receiving a control packet, the flow controller switches the input rate from one to the other.

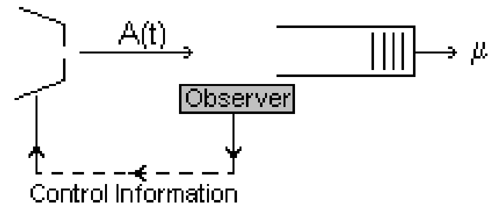

Fig. 1. A single server queue with input rate control.

B. Throughput, congestion and rate of control

There are three key parameters that are of interest while designing a congestion control policy, namely, throughput, congestion probability, and rate of control. There is usually an inevitable tradeoff between throughput and congestion probability in a flow control policy. In fact, a good flow control policy should ensure a high enough throughput, in addition to effectively controlling congestion. In this context, we refer to a related paper [9] wherein the authors derive an optimal flow control policy, which maximizes the long term throughput in a single server queue subject to a congestion probability constraint. In this paper, we assume that a minimum throughput guarantee should be met. Observe that a minimum throughput of λ_2 is guaranteed, whereas any throughput less than $\min(\lambda_1, \mu)$ can be supported, by using the higher input rate λ_1 judiciously. Loosely speaking, a higher throughput is achieved by maintaining the higher input rate λ_1 for a longer fraction of time, with a corresponding tradeoff in the congestion probability.

Let us first consider a single threshold flow control policy, which helps us understand the three key issues mentioned above. In the single threshold policy, the higher input rate λ_1 is used whenever the queue occupancy is less than or equal to some threshold l, and the lower rate is used for queue lengths larger than l. It can be shown that a larger value of l leads to a larger throughput, and vice-versa. Thus, given a minimum throughput requirement, we can determine the corresponding threshold l to meet the requirement. Once the threshold l has been fixed, it can be easily shown that the single threshold policy minimizes the probability of congestion. However, it suffers from the drawback that it requires frequent transmission of control packets, since the system may often toggle between states l and $l+1$. It turns out that a simple extension of the single threshold policy gives rise to a family of control policies, which provide more flexibility with the rate of control, while still achieving the throughput guarantee and ensuring good congestion control performance.

III. The two threshold flow control policy

As suggested by the name, the input rates in the two threshold policy are switched at two distinct thresholds l and

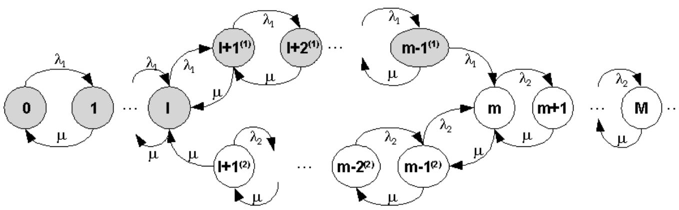

Fig. 2. The Markov process $Y(t)$ corresponding to the two-threshold control policy.

m, where $m \geq l+1$, and l is the threshold determined by the throughput guarantee. The two threshold policy operates as follows.

Suppose we start with an empty queue. The higher input rate λ_1 is used as long as the queue length does not exceed m. When the queue length grows past m, a congestion notification occurs, and the input rate switches to the lower value λ_2. Once the lower input rate is employed, it is maintained until the queue length falls back to l, at which time the input rate switches back to λ_1.

Let $Q(t)$ denote the queue length at time t and $A(t) \in \{\lambda_1, \lambda_2\}$, the input rate. Define $Y(t) = (Q(t), A(t))$ to be the state of the system at time t. It is easy to see that for the two threshold policy, $Y(t)$ is a continuous time Markov process. The state space and transition rates for the process $Y(t)$ are shown in Fig. 2.

Define $k = m-l$ to be the difference between the two queue length thresholds. We use the short hand notation $l+i^{(1)}$ and $l+i^{(2)}$ in the figure, to denote respectively, the states $(Q(t) = l+i, A(t) = \lambda_1)$ and $(Q(t) = l+i, A(t) = \lambda_2)$, $i = 1, \ldots, k-1$. For queue lengths l or smaller and m or larger, we drop the subscripts because the input rate for these queue lengths can only be λ_1 and λ_2 respectively. Note that the case $k=1$ corresponds to the single threshold policy. It can be shown that the throughput of the two threshold policy for $k > 1$ cannot be smaller than that of the single threshold policy. Thus, given a throughput guarantee, we can solve for the threshold l using the single threshold policy, and the throughput guarantee will also be met for $k > 1$.

A. Congestion probability and rate of control

We now define the congestion probability of a control policy.

Definition 1: The congestion probability is defined as the steady state probability that the queue length exceeds M, i.e., $\lim_{t\to\infty} \mathbb{P}\{Q(t) \geq M\}$. We denote it henceforth by $\mathbb{P}\{Q \geq M\}$.

Intuitively, as the gap between the two thresholds $k = m-l$ increases for a fixed l, the probability of congestion should increase, whereas the rate of control packets sent by the observer should decrease. It turns out that we can fully characterize the rate-congestion tradeoff for the two-threshold policy in closed form. We do this by solving for the steady state probabilities in Fig. 2. Define $\rho_2 = \frac{\lambda_2}{\mu}, \rho_1 = \frac{\lambda_1}{\mu}$, and $\eta_1 = 1/\rho_1$. Note that by assumption, we have $\rho_2 < 1$ and $\rho_1 > \rho_2$.

Let us denote the steady state probabilities of the non-superscripted states in Fig. 2 by p_j, where $j \leq l$, or $j \geq m$. Next, denote by $p^{(1)}_{l+i}$ ($p^{(2)}_{l+i}$) the steady state probability of the state $l+i^{(1)}$ ($l+i^{(2)}$), for $i = 1, 2, \ldots, k-1$. By solving for the steady state probabilities of various states in terms of p_l, we obtain:

$$p_i = p_l \eta_1^{l-i}, \quad 0 \leq i \leq l,$$

$$p^{(1)}_{m-j} = \begin{cases} \frac{1-\eta_1^j}{1-\eta_1^k} p_l, & \eta_1 \neq 1 \\ \frac{j}{k} p_l, & \eta_1 = 1 \end{cases}, j = 1, 2, \ldots, k-1,$$

$$p^{(2)}_{l+j} = \rho_1 \frac{1-\rho_2^j}{1-\rho_2} p^{(1)}_{m-1}, j = 1, 2, \ldots, k-1,$$

and

$$p_j = \rho_2^{j-m} \rho_1 \frac{1-\rho_2^k}{1-\rho_2} p^{(1)}_{m-1}, j \geq m.$$

The value of p_l, which is the only remaining unknown in the system can be determined by normalizing the probabilities to 1:

$$p_l = \begin{cases} \left[\frac{k(1-\rho_2\eta_1)}{\eta_1(1-\eta_1^k)(1-\rho_2)} - \frac{\eta_1^{l+1}}{1-\eta_1}\right]^{-1}, & \eta_1 \neq 1 \\ \left[l + \frac{k+1}{2} + \frac{1}{1-\rho_2}\right]^{-1}, & \eta_1 = 1 \end{cases}. \quad (1)$$

Using the steady-state probabilities derived above, we can compute the probability of congestion as

$$\mathbb{P}\{Q \geq M\} = \sum_{j \geq M} p_j = \rho_2^{M-m} \rho_1 \frac{1-\rho_2^k}{(1-\rho_2)^2} p^{(1)}_{m-1}. \quad (2)$$

The average rate of control packets sent can be found by noting that there is one packet transmitted by the observer, every time the state changes from $m-1^{(1)}$ to m or from $l+1^{(2)}$ to l. The rate (in control packets per second) is therefore given by

$$R = \lambda_1 p^{(1)}_{m-1} + \mu p^{(2)}_{l+1}.$$

Next, observe that for a positive recurrent chain, $\lambda_1 p^{(1)}_{m-1} = \mu p^{(2)}_{l+1}$. Thus,

$$R = 2\lambda_1 p^{(1)}_{m-1} = \begin{cases} \frac{2\lambda_1(1-\eta_1)}{1-\eta_1^k} p_l, & \eta_1 \neq 1 \\ \frac{2\lambda_1 p_l}{k}, & \eta_1 = 1 \end{cases}, \quad (3)$$

where p_l was found in terms of the system parameters in (1).

It is clear from (2) and (3) that k determines the trade-off between the congestion probability and rate of control. Specifically, a larger k implies a smaller rate of control, but a larger probability of congestion, and vice versa. Thus, we conclude that for the two threshold policy, the parameter l dictates the minimum throughput guarantee, while k trades off the congestion probability with rate of control packets.

B. Large deviation exponents

In many queueing systems, the congestion probability decays exponentially in the buffer size M. Furthermore, when the buffer size gets large, the exponential term dominates all other sub-exponential terms in determining the decay probability. It is therefore useful to focus only on the exponential rate of decay, while ignoring all other sub-exponential dependencies of the congestion probability on the buffer size M. Such a characterization is obtained by using the so called large deviation exponent (LDE). For a given control policy, we define the LDE corresponding to the decay rate of the congestion probability as

$$E = \lim_{M-l \to \infty} -\frac{1}{M-l} \ln \mathbb{P}\{Q \geq M\}.$$

Here, we define the LDE with respect to $M-l$ getting large, since there is no control applied when $Q \leq l$. Next we compute the LDE for the two threshold policy.

Proposition 1: Assume that k scales with M sub-linearly, so that $\lim_{M-l\to\infty} \frac{k(M)}{M-l} = 0$. The LDE of the two threshold policy is then given by

$$E = \ln \frac{1}{\rho_2}. \quad (4)$$

The above result follows from the congestion probability expression (2) since the only term that is exponential in M is $\rho_2^{M-m} = \rho_2^{M-l-k}$. We pause to make the following observations:

- If k scales linearly with $M-l$ as $k(M) = \beta(M-l)$ for some constant $\beta > 0$, the LDE becomes

$$E = (1-\beta) \ln \frac{1}{\rho_2}.$$

- The control rate (3) can be made arbitrarily small, if $k(M)$ tends to infinity. This implies that as long as $k(M)$ grows to infinity sub-linearly in M, we can achieve an LDE that is constant (equal to $-\ln \rho_2$) for all rates of control.
- As k becomes large, the congestion probability will increase. However, the increase is only sub-exponential in the buffer size, so that the LDE remains constant.

In what follows, we will be interested only in the LDE corresponding to the congestion probability, rather than its actual value. The following theorem establishes the optimality of the LDE for the two threshold policy.

Theorem 1: The two threshold policy has the best possible LDE corresponding to the congestion probability among all flow control policies, for any rate of control.

This result is a simple consequence of the fact that the two threshold policy has the same LDE as an M/M/1 queue with the lower input rate λ_2, and the latter clearly cannot be surpassed by any flow control policy.

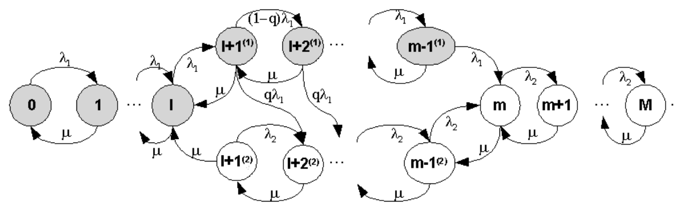

Fig. 3. The Markov process corresponding to the control policy described in subsection III-C that approximates RED.

C. Generalizing the two-threshold policy and relationship to RED

We now point out that the two-threshold policy can be easily extended to resemble the well known RED active queue management scheme [7]. Recall that RED preemptively avoids congestion by starting to drop packets randomly even before the buffer is about to overflow. Specifically, consider two queue thresholds, say l and m, where $m > l$. If the queue occupancy is no more than l, no packets are dropped, no matter what the input rate is. On the other hand, if the queue length reaches or exceeds m, packets are always dropped, which then leads to a reduction in the input rate (assuming that the sender responds to dropped packets). If the queue length is between l and $m-1$, packets are randomly dropped with some probability q.[1]

Consider the following generalization of the two threshold policy, which closely resembles the RED scheme described above:

For queue lengths less than or equal to l, the higher input rate is always used. If the queue length increases to m while the input rate is λ_1, a congestion notification is sent, and the input rate is reduced to λ_2. If the current input rate is λ_1 and the queue length is between l and $m-1$, a congestion notification occurs with probability[2] q upon the arrival of a packet, and the input rate is reduced to λ_2. With probability $1-q$, the input continues at the higher rate. The Markov process corresponding to this policy is depicted in Fig. 3.

It is possible derive the tradeoff between the congestion probability and the rate of congestion notifications for this policy by analyzing the Markov chain in Fig. 3. Once the lower threshold l has been determined from the throughput guarantee, the control rate vs. congestion probability tradeoff is determined by both q and m. Further, since the input rate switches to the lower value λ_2 when the queue length is larger than m, this flow control policy also achieves the optimal LDE for the congestion probability, equal to $\ln \frac{1}{\rho_2}$. We skip the derivations for this policy, since it is more cumbersome to analyze than the two threshold policy, without yielding further qualitative insights. We focus on the two threshold policy in the remainder of the paper, but point out that our methodology can also be used to model more practical queue management

[1]Often, the dropping probability is dependent on the queue length.

[2]We can also let this probability to depend on the current queue length, as often done in RED, but this makes the analysis more difficult

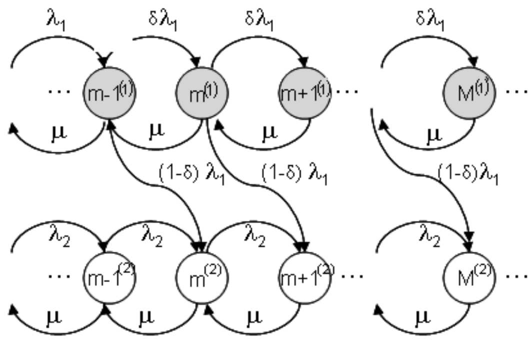

Fig. 4. The Markov process $Y(t)$ corresponding to the error prone two-threshold policy. Only a part of the state space (with $Q(t) \geq M$) is shown.

policies such as RED.

IV. The effect of control errors on congestion

In this section, we investigate the impact of control errors on the congestion probability of the two threshold policy. We use a simple probabilistic model for the errors on the control channel. In particular, we assume that any control packet sent by the observer can be lost with some probability δ, independently of other packets. Using the decay exponent tools described earlier, we show the existence of a critical value of the error probability, say δ^*, beyond which the errors in receiving the control packets lead to an exponential degradation of the congestion probability.

A. The two-threshold policy over an error-prone control channel

As described earlier, in the two threshold policy, the observer sends a control packet when the queue length reaches $m = l + k$. This packet may be received by the flow controller with probability $1-\delta$, in which case the input rate switches to λ_2. The packet may be lost with probability δ, in which case the input continues at the higher rate λ_1. We assume that if a control packet is lost, the observer immediately knows about it[3], and sends another control packet the next time an arrival occurs to a system with at least $m-1$ packets.

The process $Y(t) = (Q(t), A(t))$ is a Markov process even for this error prone two threshold policy. Fig. 4 shows a part of the state space for the process $Y(t)$, for queue lengths larger than $m-1$. Note that due to control errors, the input rate does not necessarily switch to λ_2 for queue lengths greater than $m-1$. Indeed, it is possible to have not switched to the lower input rate even for arbitrarily large queue lengths. This means that the congestion limit can be exceeded under both arrival rates, as shown in Fig. 4.

As in the previous section, we can compute the congestion probability for the error prone two threshold policy by solving for the steady state probabilities from Fig. 4. Since we are interested only in the LDE corresponding to the congestion probability, we only need to consider terms that are exponential in M. We find that the top set of states in Fig. 4 (which correspond to arrival rate λ_1) have steady state probabilities that satisfy

$$p^{(1)}_{m-1+i} = s(\delta)^i p^{(1)}_{m-1}, i = 1, 2, \ldots,$$

where

$$s(\delta) = \frac{1+\rho_1 - \sqrt{(1+\rho_1)^2 - 4\rho_1\delta}}{2}. \tag{5}$$

Similarly, for the bottom set of states in Fig. 4, the steady state probabilities have the form

$$p^{(2)}_{m-1+i} = A\rho_2^i + Bs(\delta)^i, i = 1, 2, \ldots,$$

where A, B are constants that depend on the system parameters ρ_1, ρ_2, and δ. Using the two expressions above, we can deduce that the congestion probability has *two* terms that decay exponentially in the buffer size:

$$\mathbb{P}\{Q \geq M\} = Cs(\delta)^{M-l-k} + D\rho_2^{M-l-k}, \tag{6}$$

where C, D are constants.

In order to compute the LDE, we need to determine which of the two exponential terms in (6) decays slower. It is easy to show that $s(\delta) \leq \rho_2$ for $\delta \leq \delta^*$, where

$$\delta^* = \frac{\rho_2}{\rho_1}(1 + \rho_1 - \rho_2). \tag{7}$$

Thus, for error probabilities less than δ^*, ρ_2 dominates the rate of decay of the congestion probability. Similarly, for $\delta > \delta^*$, the exponential rate of decay is dominated by $s(\delta)$. We thus have the following theorem.

Theorem 2: Consider a two threshold policy in which k grows sub-linearly in M. Assume that the control packets sent by the observer can be lost with probability δ. Then the LDE corresponding to the congestion probability is given by

$$E(\delta) = \begin{cases} \ln \frac{1}{\rho_2}, & \delta \leq \delta^*, \\ \ln \frac{2}{1+\rho_1 - \sqrt{(\rho_1+1)^2 - 4\delta\rho_1}}, & \delta > \delta^*. \end{cases}, \tag{8}$$

where δ^* is the *critical error probability* given by (7).

The theorem shows that the two threshold policy over an error prone channel has two regimes of operation. In particular, for 'small enough' error probability ($\delta < \delta^*$), the exponential rate of decay of the congestion probability is the same as in an error free system. However, for $\delta > \delta^*$, the decay exponent begins to take a hit, and therefore, the congestion probability suffers an exponential increase. For this reason, we refer to δ^* as the critical error probability. Fig. 5 shows a plot of the decay exponent as a function of the error probability δ, for $\rho_1 > 1$. The 'knee point' in the plot corresponds to δ^* for the stated values of ρ_1 and ρ_2.

Remark 1: Large deviation theory has been widely applied to study congestion and overflow behaviors in queueing systems. Tools such as the Kingman bound [10] can be used to characterize the LDE of any G/G/1 queue. Large deviation framework also exists for more complicated queuing systems,

[3]This is an idealized assumption; in practice, delayed feedback can be obtained using ACKS.

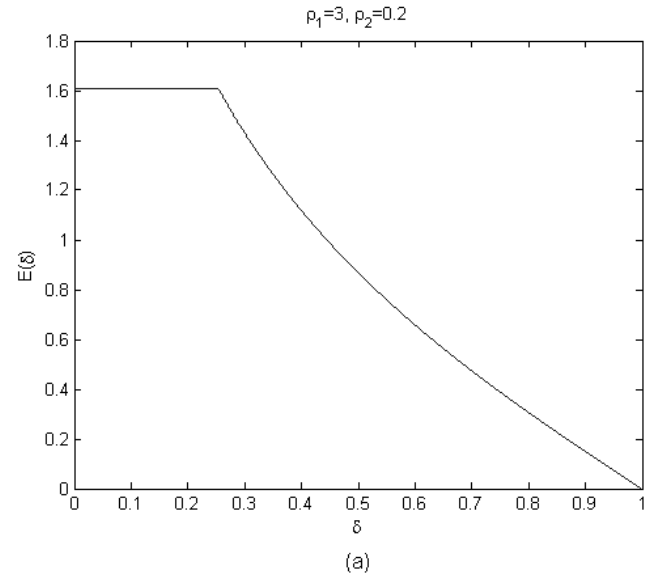

Fig. 5. LDE as a function of δ for $\rho_1 > 1$.

with correlated inputs, several sources, finite buffers etc., see for instance [11]. However, for *controlled* queues, where the input or service rates can vary based on queue length history, simple large deviation formulas do not exist. It is remarkable that for a single server queue with Markovian control, we are able to obtain rather intricate LDE characterizations such as in Fig. 5, just by applying 'brute force' steady state probability computations.

B. Repetition of control packets

Suppose we are given a control channel with probability of error δ that is greater than the critical error probability in (7). This means that a two-threshold policy operating on this control channel has an LDE in the decaying portion of the curve in Fig. 5. In this situation, adding error protection to the control packets will reduce the effective probability of error, thereby improving the LDE. To start with, we consider the simplest form of adding redundancy to control packets, namely repetition.

Suppose that each control packet is transmitted n times by the observer, and that all n packets are communicated without delay. Assume that each of the n packets has a probability δ of being lost, independently of other packets. The flow controller fails to switch to the lower input rate only if all n control packets are lost, making the effective probability of error δ^n. In order to obtain the best possible LDE, the operating point must be in the flat portion of the LDE curve, which implies that the effective probability of error should be no more than δ^*. Thus, $\delta^n \leq \delta^*$, so that the number of transmissions n should satisfy

$$n \geq \frac{\ln \delta^*}{\ln \delta} \tag{9}$$

in order to obtain the best possible LDE of $\ln \frac{1}{\rho_2}$. If the value of δ is close to 1, the number of repeats is large, and vice-versa.

V. Optimal bandwidth allocation for control signals

As discussed in the previous subsection, the LDE operating point of the two-threshold policy for any given $\delta < 1$, can always be 'shifted' to the flat portion of the curve by repeating the control packets sufficiently many times (9). This ignores the bandwidth consumed by the additional control packets.

While the control overheads constitute an insignificant part of the total communication resources in optical networks, they might consume a sizeable fraction of bandwidth in some wireless or satellite applications. In such a case, we cannot add an arbitrarily large amount of control redundancy without sacrificing some service bandwidth. Typically, allocating more resources to the control signals makes them more robust to errors, but it also reduces the bandwidth available to serve data. To better understand this tradeoff, we explicitly model the service rate to be a function of the redundancy used for control signals. We then determine the optimal fraction of bandwidth to allocate to the control packets, so as to achieve the best possible decay exponent for the congestion probability.

A. Bandwidth sharing model

Consider, for the time being, the simple repetition scheme for control packets outlined in the previous section. We assume that the queue service rate is linearly decreasing function of the number of repeats $n-1$:

$$\mu(n) = \mu \left[1 - \frac{n-1}{\Phi}\right]. \tag{10}$$

The above model is a result of the following assumptions about the bandwidth consumed by the control signals:

- μ corresponds to the service rate when no redundancy is used for the control packets ($n = 1$).
- The amount of bandwidth consumed by the redundancy in the control signals is proportional to the number of repeats $n-1$.
- The fraction of total bandwidth consumed by each repetition of a control packet is equal to $1/\Phi$, where $\Phi > 0$ is a constant that represents how 'expensive' it is in terms of bandwidth to repeat control packets.

Thus, with $n-1$ repetitions, the fraction of bandwidth consumed by the control information is $\frac{n-1}{\Phi}$, and the fraction available for serving data is $1 - \frac{n-1}{\Phi}$.

Let us denote by f the fraction of bandwidth consumed by the redundancy in the control information, so that $f = \frac{n-1}{\Phi}$, or $n = \Phi f + 1$. From (10), the service rate corresponding to the fraction f can be written as

$$\mu(f) = \mu[1 - f].$$

In what follows, we do not restrict ourselves to repetition of control packets, so that we are not constrained to integer values of n. Instead, we allow the fraction f to take continuous values, while still maintaining that the error probability corresponding to f is $\delta^{\Phi f+1}$. We refer to f as the 'fraction of bandwidth used for control', although it is really the fraction of bandwidth utilized by the *redundancy* in the control. For example, $f = 0$ does not mean no control is used; instead, it corresponds to each control packet being transmitted just once.

B. Optimal fraction of bandwidth to use for control

The problem of determining the optimal fraction of bandwidth to be used for control can be posed as follows:

Given the system parameters ρ_1, ρ_2 and Φ, and a control channel with some probability of error $\delta \in [0, 1)$, find the optimal fraction of bandwidth $f^*(\delta)$ to be used for control, so as to maximize the LDE of the congestion probability.

Let us define

$$\rho_i(f) = \frac{\lambda_i}{\mu(f)} = \frac{\rho_i}{1-f}, i = 1, 2, \tag{11}$$

as the effective server utilization corresponding to the reduced service rate $\mu(f)$. Accordingly, we can also define the effective knee point as

$$\delta^*(f) = \frac{\rho_2}{\rho_1}\left(1 + \frac{\rho_1 - \rho_2}{1-f}\right), \tag{12}$$

which is analogous to (7), with $\rho_i(f)$ replacing $\rho_i, i = 1, 2$.

First, observe that for the queueing system to be stable, we need the effective service rate to be greater than the lower input rate λ_2. Thus, we see that $\lambda_2 < \mu[1-f]$, or $f < 1-\rho_2$. Next, we compute the LDE corresponding to a given probability of error δ, and fraction f of bandwidth used for control.

Proposition 2: For any $\delta \in [0, 1)$ and $f \in [0, 1 - \rho_2)$, the corresponding LDE is given by

$$E(\delta, f) = \begin{cases} \ln \frac{1}{\rho_2(f)}, & \delta^{\Phi f+1} \le \delta^*(f), \\ \\ \ln \frac{1}{s(f,\delta)}, & \delta^{\Phi f+1} > \delta^*(f), \end{cases} \tag{13}$$

where

$$s(f, \delta) = \frac{1 + \rho_1(f) - \sqrt{(\rho_1(f) + 1)^2 - 4\delta^{\Phi f+1}\rho_1(f)}}{2}.$$

The derivation and expression for $E(\delta, f)$ are analogous to (8), except that ρ_i is replaced with $\rho_i(f), i = 1, 2$, and δ is replaced with the effective probability of error $\delta^{\Phi f+1}$.

Definition 2: For any given $\delta \in [0, 1)$, the optimal fraction $f^*(\delta)$ is the value of f that maximizes $E(\delta, f)$ in (13). Thus,

$$f^*(\delta) = \mathrm{argmax}_{f \in [0,1-\rho_2)} E(\delta, f). \tag{14}$$

Recall that the value of $1/\Phi$ represents how much bandwidth is consumed by each added repetition of a control packet. We will soon see that Φ plays a key role in determining the optimal fraction of bandwidth to use for control. Indeed, we show that there are three different regimes for Φ such that the optimal fraction $f^*(\delta)$ exhibits qualitatively different behavior in each regime as a function of δ. The three ranges of Φ are: (i) $\Phi \le \underline{\Phi}$, (ii) $\Phi \ge \overline{\Phi}$, and (iii) $\underline{\Phi} < \Phi < \overline{\Phi}$, where

$$\begin{aligned} \underline{\Phi} &= \frac{\rho_2 - \delta^*}{\ln(\delta^*)(1 + \rho_1 - \rho_2)}, \\ \overline{\Phi} &= \begin{cases} \frac{1}{\rho_1 - 1} & \rho_1 > 1 \\ \infty & \rho_1 \le 1 \end{cases}. \end{aligned} \tag{15}$$

It can be shown that $\underline{\Phi} < \overline{\Phi}$ for $\rho_2 < 1$.

We shall refer to case (i) as the 'small Φ regime', case (ii) as the 'large Φ regime', and case (iii) as the 'intermediate regime'. We remark that whether a value of Φ is considered 'small' or 'large' is decided entirely by ρ_1 and ρ_2. Note that the large Φ regime is non-existent if $\rho_1 \le 1$, so that even if Φ is arbitrarily large, we would still be in the intermediate regime.

The following theorem, which is our main result for this section, specifies the optimal fraction of bandwidth $f^*(\delta)$, for each of the three regimes for Φ.

Theorem 3: For a given ρ_1 and ρ_2, the optimal fraction of bandwidth $f^*(\delta)$ to be used for control, has one of the following forms, depending on the value of Φ:

(i) Small Φ regime ($\Phi \le \underline{\Phi}$): $f^*(\delta) = 0, \ \forall\ \delta \in (0, 1)$.

(ii) Large Φ regime ($\Phi \ge \overline{\Phi}$):

$$f^*(\delta) = \begin{cases} 0, & \delta \in [0, \delta^*] \\ \hat{f}(\delta), & \delta \in (\delta^*, 1) \end{cases},$$

where $\hat{f}(\delta)$ is the unique solution to the transcendental equation

$$\delta^{\Phi \hat{f}+1} = \frac{\rho_2}{\rho_1}\left(1 + \frac{\rho_1 - \rho_2}{1 - \hat{f}}\right). \tag{16}$$

(iii) Intermediate regime ($\underline{\Phi} < \Phi < \overline{\Phi}$): there exist δ' and δ'' such that $\delta^* < \delta' < \delta'' < 1$, and the optimal fraction is given by

$$f^*(\delta) = \begin{cases} 0, & \delta \in [0, \delta^*] \\ \hat{f}(\delta), & \delta \in (\delta^*, \delta') \\ \tilde{f}(\delta), & \delta \in (\delta', \delta'') \\ 0, & \delta \in (\delta'', 1) \end{cases},$$

where $\hat{f}(\delta)$ is given by (16) and $\tilde{f}(\delta)$ is the unique solution in $(0, 1 - \rho_2)$ to the transcendental equation

$$\delta^{\Phi \tilde{f}+1}[\Phi(1 - \tilde{f}) \ln \delta^* + 1] = s(\tilde{f}, \delta). \tag{17}$$

The proof of the above theorem is not particularly intuitive or interesting, and is omitted for brevity. Instead, we provide some intuition about the optimal solution.

C. Discussion of the optimal solution

1) Error probability less than δ^*: In all three regimes, we find that $f^*(\delta) = 0$ for $\delta \in [0, \delta^*]$. This is because, as shown in Fig. 5, the LDE has the highest possible value of $-\ln \rho_2$ for δ in this range, and there is nothing to be gained from adding any control redundancy.

2) Small Φ regime: In case (i) of the theorem, it is optimal to not apply any control redundancy at all. That is, the best possible LDE for the congestion probability is achieved by using a single control packet every time the observer intends to switch the input rate. In this regime, the amount of service bandwidth lost by adding any control redundancy at all, hurts us more than the gain obtained from the improved error probability. The plot of the optimal LDE as a function of δ for this regime is identical to Fig. 5, since no redundancy is applied.

3) Large Φ regime: Case (ii) of the theorem deals with the large Φ regime. For $\delta > \delta^*$, the optimal $f^*(\delta)$ in this regime is chosen as the fraction f for which the knee point $\delta^*(f)$ equals the effective error probability $\delta^{\Phi f+1}$. This fraction is indeed $\hat{f}$, defined by (16). Fig. 6(a) shows a plot of the optimal fraction (solid line) as a function of δ. In this example, $\rho_1 = 1.2$, $\rho_2 = 0.3$, and $\Phi = 10$. The resulting optimal LDE is equal to $\ln \frac{1-\hat{f}(\delta)}{\rho_2}$ for $\delta > \delta^*$. The optimal LDE is shown in Fig. 6(b) with a solid line.

4) Comparison with naïve repetition: It is interesting to compare the optimal solution in the large Φ regime to the 'naïve' redundancy allocation policy mentioned in equation (9). Recall that the naïve policy simply repeats the control packets to make the effective error probability equal to the critical probability δ^*, without taking into account any service bandwidth penalty that this might entail. Let us see how the naïve strategy compares to the optimal solution if the former is applied to a system with a finite Φ. This corresponds to a network with limited communication resources in which the control mechanisms are employed without taking into account the bandwidth that they consume.

The fraction of bandwidth occupied by the repeated control packets can be found using (9) to be

$$f = \frac{1}{\Phi}\left(\frac{\ln \delta^*}{\ln \delta} - 1\right),$$

where we have ignored integrality constraints on the number of repeats. A plot of this fraction is shown in Fig. 6(a), and the corresponding LDE in 6(b), both using dashed lines. As seen in the figure, the naive strategy is more aggressive in adding redundancy than the optimal strategy, since it does not take into account the loss in service rate ensuing from the finiteness of Φ. The LDE of the naïve strategy is strictly worse for $\delta > \delta^*$. In fact the naïve strategy causes instability effects for some values of δ close to 1 by over-aggressive redundancy addition, which throttles the service rate $\mu(f)$ to values below the lower arrival rate λ_2. This happens at the point where the LDE reaches zero in Fig. 6(b). The naïve strategy has even worse consequences in the other two regimes. However, we point out that the repetition strategy approaches the optimal solution as Φ becomes very large.

5) Intermediate regime: Case (iii) in the theorem deals with the intermediate regime. For $\delta > \delta^*$, the optimal fraction begins to increase along the curve $\hat{f}(\delta)$ exactly like in the large Φ regime (see Fig. 7). That is, the effective error probability is made equal to the knee point. However, at a particular value of error probability, say δ', the optimal fraction begins to decrease sharply from the $\hat{f}(\delta)$ curve, and reaches zero at some value δ''. Equation (17) characterizes the optimal fraction for values of δ in (δ', δ''). No redundancy is applied for $\delta \in (\delta'', 1)$. For this range of error probability, the intermediate regime behaves more like the small Φ regime (case(i)). Thus, the intermediate Φ regime resembles the large Φ regime for small enough error probabilities $\delta < \delta'$, and the small Φ regime for large error probability $\delta > \delta''$. There is also a non empty 'transition interval' in between the two, namely (δ', δ'').

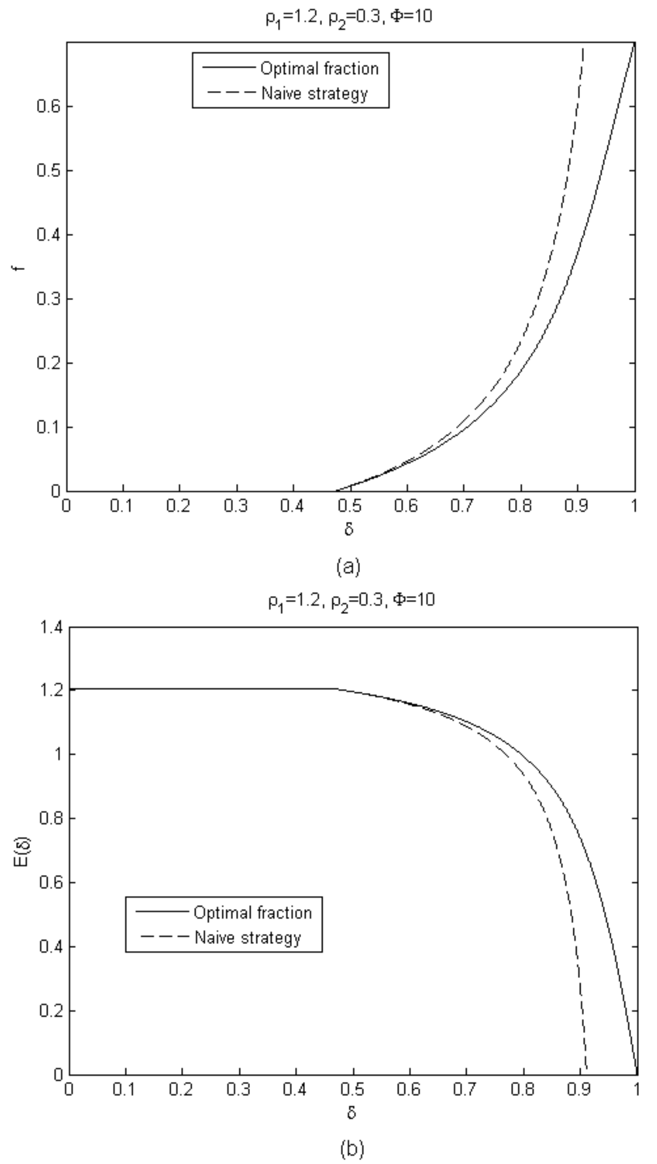

Fig. 6. (a) Optimal fraction $f^*(\delta)$ for the large Φ regime, and the fraction for the naïve strategy (b)The corresponding LDE curves

VI. Conclusions

The goal of this paper was to characterize the tradeoff between the rate of control and network congestion for flow control policies. Specifically, we deal with the question of

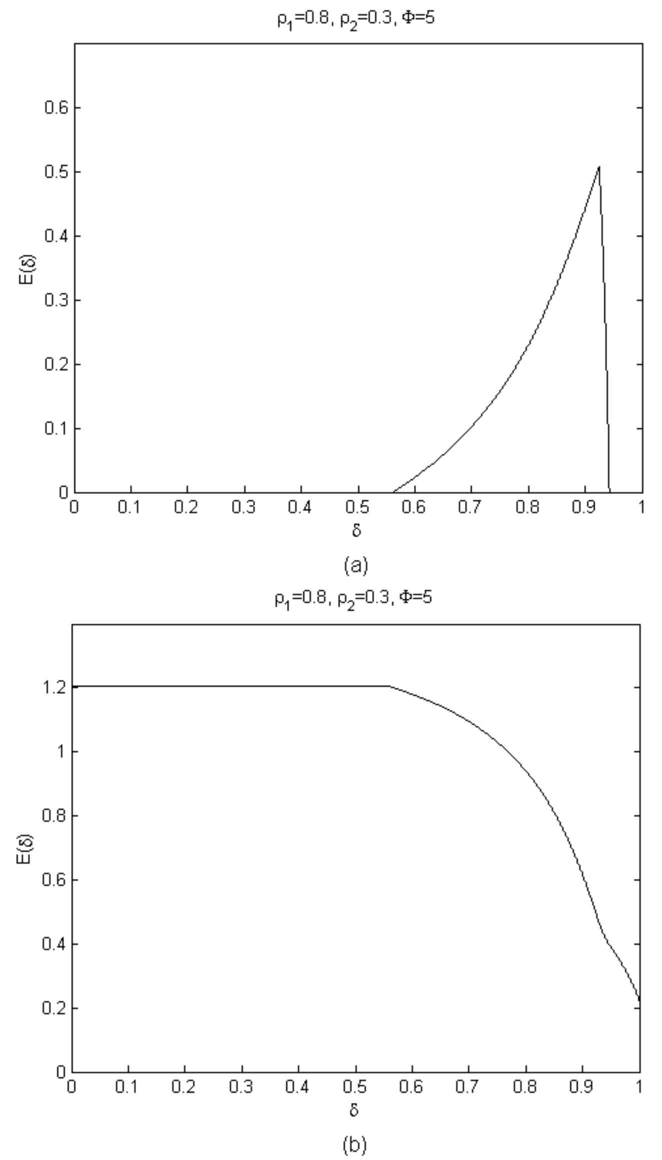

Fig. 7. (a) Optimal control fraction $f^*(\delta)$ for the intermediate regime (b) The corresponding LDE

how often congestion notifications need to be sent in order to effectively control congestion. To our knowledge, this is the first attempt to analytically study this particular tradeoff. Since this tradeoff is difficult to analyze in general networks, we consider a simple model of a single server queue with congestion-based flow control. We learned that in the absence of control channel errors, the control rate needed to ensure the optimal decay exponent for the congestion probability can be made arbitrarily small. However, if control channel errors occur probabilistically, we showed the existence of a critical error probability threshold beyond which the congestion probability undergoes a drastic increase due to the frequent loss of control packets. Finally, we determine the optimal amount of error protection to apply to the control signals by using a simple bandwidth sharing model. For error probabilities larger than the critical value, a significant fraction of the system resources may be consumed by the control signals, unlike in the error free scenario. We also pointed out that allocating control resources without considering the bandwidth they consume, might have adverse effects on congestion.

REFERENCES

[1] T. Bonald, and L. Massoulie, *Impact of Fairness on Internet Performance,* ACM SIGMETRICS Performance Evaluation Review, Volume 29 , Issue 1, June 2001.

[2] L. Tassiulas, A. Ephremides, *Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks.* IEEE Trans. Aut. Contr. **37**, pp. 1936–1948, 1992.

[3] M.J. Neely, E. Modiano, C.E. Rohrs, *Dynamic power allocation and routing for time-varying wireless networks.* INFOCOM Proceedings, 2003.

[4] M. J. Neely, E. Modiano, and C. Li, *Fairness and Optimal Stochastic Control for Heterogeneous Networks,* INFOCOM Proceedings, 2005.

[5] X. Lin, N. Shroff and R. Srikant, *On the Connection-Level Stability of Congestion-Controlled Communication Networks* , IEEE Trans. Inf. Thy. **54**, pp. 2317 - 2338, 2008.

[6] R. G. Gallager, *Basic Limits on Protocol Information in Data Communication Networks*, IEEE Trans. Inf. Thy. Vol. IT-22, No. 4, July 1976, pp. 385-398.

[7] S. Floyd, and V. Jacobson, *Random Early Detection gateways for Congestion Avoidance* V.1 N.4, August 1993, p. 397-413.

[8] K. P. Jagannathan, E. Modiano, L. Zheng, *Effective Resource Allocation in a Single-Server Queue: How Much Control is Necessary?*, Forty-Sixth Annual Allerton Conference on Communication, Control, and Computing, Monticello IL, September 2008.

[9] J R. Perkins, R. Srikant, *The Role of Queue Length Information In Congestion Control and Resource Pricing*, Proceedings of the 38th Conference on Decision & Control, Phoenix, Arizona, USA, December 1999.

[10] R. G. Gallager, *Discrete Stocastic Processes*, Kluwer Academic Publishers, 1996, pp. 234.

[11] A. Ganesh, N. O'Connel, D. Wischik, *Big Queues*, Springer-Verlag, 2004.