

On the Role of Queue Length Information in Network Control

Krishna Jagannathan, Eytan Modiano, *Senior Member, IEEE*, and Lihong Zheng, *Member, IEEE*

Abstract—We study the role played by queue length information in the operation of flow control and server allocation policies. We first 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 exponential scaling for the queue congestion probability, for any rate of control. We show that when the control channel is reliable, the control rate needed to ensure the optimal decay exponent for the congestion probability can be made arbitrarily small. However, if control channel erasures occur probabilistically, we show the existence of a critical erasure probability threshold beyond which the congestion probability undergoes a drastic increase due to the frequent loss of control packets. We also determine the optimal amount of error protection to apply to the control signals by using a simple bandwidth sharing model. Finally, we show that the queue length based server allocation problem can also be treated using this framework and that the results obtained for the flow control setting can also be applied to the server allocation case.

Index Terms—Buffer overflow, congestion control, large deviations, queue length information, resource allocation.

I. INTRODUCTION

NETWORK control plays a crucial role in the operation of modern communication networks, and consists of several functions, such as scheduling, routing, flow control, and resource allocation. A primary goal of network control is to ensure stability [10], [12]. Furthermore, network control is used to provide QoS guarantees on metrics such as throughput, delay and fairness. These objectives are usually accomplished through a combination of resource allocation and flow control. The control decisions often take into account the instantaneous channel quality of the various links, the queue backlogs, and the quality of Service (QoS) requirements of the different flows. In this paper, we study the role played by queue length information in flow control and resource allocation policies.

Flow control and resource allocation play an important role in keeping congestion levels in the network within acceptable

limits. Flow control involves regulating the rate of the incoming exogenous traffic to a queue, depending on its congestion level, see [9] for example. Resource allocation, on the other hand, involves assigning larger service rates to the queues that are congested, and vice-versa [10], [12]. Most systems use a combination of the two methods to avoid congestion in the network, and to achieve various performance objectives [7], [9].

The knowledge of queue length information is often useful, sometimes even necessary, in order to perform these control tasks effectively. Almost all practical flow control mechanisms base their control actions on the level of congestion present at a given time. For example, in networks employing TCP, 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 such as Random Early Detection (RED) are designed to pro-actively prevent congestion by randomly dropping some packets *before* the buffers reach the overflow limit [2]. On the other hand, resource allocation policies can either be queue-blind, such as round-robin, first come first served (FCFS), and generalized processor sharing (GPS), or queue-aware, such as maximum weight scheduling. Queue length based scheduling techniques are known to have superior delay and queue overflow performance than queue-blind algorithms such as round-robin and processor sharing [1], [8], [13].

Since the queue lengths can vary widely over time in a dynamic network, queue occupancy based flow control and resource allocation algorithms typically require the exchange of control information between agents that can observe the various queue lengths in the system, and the controllers which adapt their actions to the varying queues. This control information can be thought of as being a part of the inevitable protocol and control overheads in a network. Gallager’s seminal paper [3] on basic limits on protocol information was the first to address this topic. 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 control messages need to be sent in order to effectively control congestion in a single server queue. We separately consider the flow control and resource allocation problems, and characterize the rate of control necessary to achieve a certain congestion control performance in the queue. In particular, we argue that there is an inherent tradeoff between the rate of control information, and the corresponding congestion level in the queue. That is, if the controller has very accurate information about the congestion level in the system, congestion control can be performed very effectively by adapting the input/service rates appropriately. However, furnishing the controller with accurate queue

Manuscript received September 02, 2009; revised October 25, 2010; accepted March 23, 2011. Date of current version August 31, 2011. This work was supported in part by NSF ITR Grant CCR-0325401, in part by NSF Grant CNS-0626781, and in part by ARO MURI Grant number W911NF-08-1-0238. The material in this paper was presented in part at INFOCOM, Rio de Janeiro, Brazil, April 2009.

The authors are with the Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: krishnaj@mit.edu; modiano@mit.edu; lizhong@mit.edu).

Communicated by R. Yates, Associate Editor for Communication Networks. Digital Object Identifier 10.1109/TIT.2011.2162155

length information requires significant amount of control. Further, frequent congestion notifications may also lead to undesirable retransmissions in packet drop based systems such as TCP. 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 the 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.

We consider a single server queue with congestion based flow control. Specifically, the queue is served at a constant rate, and is fed by packets arriving at one of 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. 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.

Next, we consider a model where erasures may occur on the control channel, possibly due to wireless transmission. We characterize the impact of control channel erasures on the congestion control performance of the two-threshold policy. We assume a probabilistic model for the erasures on the control channel, and show the existence of a critical erasure probability, beyond which the losses in receiving the control packets lead to an exponential worsening of the congestion probability. However, for erasure probabilities below the critical value, the congestion probability is of the same exponential order as in a system with an erasure-free control channel. Moreover, we determine the optimal apportioning of bandwidth between the control signals and the server in order to achieve the best congestion control performance.

Finally, we study the server allocation problem in a single server queue. In particular, we consider a queue with a constant input rate. The service rate at any instant is chosen from two possible values depending on the congestion level in the queue. This framework turns out to be mathematically similar to the flow control problem, so that most of our results for the flow control case also carry over to the server allocation problem. Earlier versions of this work appeared in [5], [6].

The rest of the 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 intro-




Fig. 1. Single server queue with input rate control.

duce and analyze the two-threshold policy. In Section IV, we investigate the effect of control channel erasures on the congestion control performance of the two-threshold policy. Section V deals with the problem of optimal bandwidth allocation for control signals in an erasure-prone system. The server allocation problem is presented in Section VI; and Section VII concludes the paper.

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 that the packet sizes are exponentially distributed with mean 1. Exogenous arrivals are fed to the queue in a regulated fashion by a flow controller. 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.

The controller chooses the input process at any instant t to be a Poisson processes of rate $A(t)$. We assume that the input rate at any instant is chosen to be one of two distinct possible values, $A(t) \in \{\lambda_1, \lambda_2\}$, 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. Moreover, [11] showed that in a single server system with flow control, where the input rate is allowed to vary in the interval $[\lambda_2, \lambda_1]$, a “bang-bang” solution is throughput optimal. That is, a queue length based threshold policy that only uses the two extreme values of the possible input rates, is optimal in the sense of maximizing throughput for a given congestion probability constraint.

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. We focus on Markovian control policies, in which the input rate chosen after an arrival or departure event is only a function of the previous input rate and queue length. We will show that Markovian policies are sufficient to achieve the optimal asymptotic decay rate for the congestion probability.

B. Markovian Control Policies

We begin by defining notions of Markovian control policies, and its associated congestion probability.

Let $t > 0$ denote continuous time. Let $Q(t)$ and $A(t)$ respectively denote the queue length and input rate (λ_1 or λ_2) at time t . Define $Y(t) = (Q(t), A(t))$ to be the state of the system at




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

time t . We assign discrete time indices $n \in \{0, 1, 2, \dots\}$ to each arrival and departure event in the queue (“queue event”). Let Q_n and A_n respectively denote the queue length and input rate just after the n^{th} queue event. Define $Y_n = (Q_n, A_n)$. A flow control policy assigns service rates A_n after every queue event.

Definition 1: A control policy is said to be Markovian if it assigns input rates A_n such that

$$\mathbb{P}\{A_{n+1}|Q_{n+1}, Y_n, \dots, Y_0\} = \mathbb{P}\{A_{n+1}|Q_{n+1}, Y_n\} \quad (1)$$

$\forall n = 0, 1, 2, \dots$

For a Markovian control policy operating on a queue with memoryless arrival and packet size distributions, it is easy to see that $Y(t)$ is a continuous time Markov process with a countable state space, and that Y_n is the imbedded Markov chain for the process $Y(t)$. A Markovian policy is said to be stabilizing if the process $Y(t)$ is positive recurrent under the policy. Thus, if a policy is stabilizing, the steady-state distribution of the queue length, denoted by Q , exists. In other words, Q is the distributional limit of $Q(t)$ as $t \rightarrow \infty$. The congestion probability under a stabilizing Markovian policy is defined using the limiting distribution Q as follows.

Definition 2: Let $M > 0$ be a congestion limit. The congestion probability is defined as $\mathbb{P}\{Q \geq M\}$.

C. Throughput, Congestion and Rate of Control

We will focus on three important parameters of a flow 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 paper, we assume that a minimum throughput guarantee γ 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.

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. In particular, the throughput obtained under the single threshold policy is given by

$$\begin{aligned} & \frac{(\rho_1^{l+1}-1)(1-\rho_2)}{\rho_1^{l+1}(\rho_1-\rho_2)-(1-\rho_2)}(\lambda_1-\lambda_2)+\lambda_2 \quad \text{for } \rho_1 \neq 1 \\ & \frac{l+1}{l+1+\frac{1}{1-\rho_2}}(\lambda_1-\lambda_2)+\lambda_2 \quad \text{for } \rho_1 = 1, \end{aligned} \quad (2)$$

where $\rho_2 = \frac{\lambda_2}{\mu}$, and $\rho_1 = \frac{\lambda_1}{\mu}$. Thus, given the throughput requirement γ , we can determine the corresponding threshold l to meet the requirement. Once the threshold l has been fixed, it can be 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. 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 m , where $m \geq l+1$, and l is the threshold determined by the throughput guarantee. As we shall see, the position of the second threshold gives us another degree of freedom, using which the rate of control can be fixed at a desired value. 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 , 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 .

We will soon see that this “hysteresis” in the thresholds helps us tradeoff the control rate with the congestion probability. The two-threshold policy is easily seen to be Markovian, and the state space and transition rates for the process are shown in Fig. 2. In what follows, we assume that the congestion limit M is much larger than the thresholds l and m . Since the two-threshold policy switches to the lower rate for queue lengths larger than m , it is clear that the system is stable, and the steady-state queue lengths exists.

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 $(Q(t) = l + i, A(t) = \lambda_2)$, $i = 1, \dots, k - 1$. 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 a single threshold policy. As we shall see later, 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 throughput expression (2) for the single threshold policy, and the throughput guarantee will also be met for $k > 1$. We now explain how the parameter k can be used to tradeoff the rate of control and the congestion probability.

A. Congestion Probability Versus Rate of Control Tradeoff

Intuitively, as the gap between the two-thresholds $k = m - l$ increases for a fixed l , the rate of control packets sent by the observer should decrease, while the probability of congestion should increase. It turns out that we can 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 nonsuper-scripted states in Fig. 2 by p_j , where $j \leq l$, or $j \geq m$. Next, denote by $p_{l+i}^{(1)}$ ($p_{l+i}^{(2)}$) the steady-state probability of the state $l + i^{(1)}$ ($l + i^{(2)}$), for $i = 1, 2, \dots, k - 1$. Let us now solve for the steady-state probabilities of various states in terms of p_l . For the states $0 \leq i \leq l$, the local balance equations imply that

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

Next, for the set of states $l + i^{(1)}$, $i = 1, 2, \dots, k - 2$, the balance equations read¹

$$p_{l+1+i}^{(1)} - (1 + \rho_1)p_{l+i}^{(1)} + p_{l+i-1}^{(1)}\rho_1 = 0, \quad i = 1, 2, \dots, k - 2.$$

For the state $m - 1^{(1)}$, we have

$$p_{m-1}^{(1)}(1 + \rho_1) = p_{m-2}^{(1)}.$$

Using the last two relations, we can recursively obtain the steady-state probabilities $p_{l+i}^{(1)}$, $i = 1, 2, \dots, k - 1$ in terms of p_l

$$p_{m-j}^{(1)} = \begin{cases} \frac{1-\eta_1^j}{1-\eta_1^k} p_l, & \eta_1 \neq 1 \\ \frac{j}{k} p_l, & \eta_1 = 1 \end{cases} \quad j = 1, 2, \dots, k - 1.$$

Similarly, the states $l + i^{(2)}$, $i = 1, 2, \dots, k - 1$, satisfy

$$p_{l+1}^{(2)}(1 + \rho_2) = p_{l+2}^{(2)}$$

and

$$p_{l+1+i}^{(2)} - (1 + \rho_2)p_{l+i}^{(2)} + p_{l+i-1}^{(2)}\rho_2 = 0, \quad i = 1, 2, \dots, k - 1.$$

From the last two relations, we obtain

$$p_{l+j}^{(2)} = \frac{1 - \rho_2^j}{1 - \rho_2} p_{l+1}^{(2)}, \quad j = 1, 2, \dots, k.$$

¹We use the convention $p_i^{(j)} = p_i$ for $j = 1, 2$.

To write the above in terms of p_l , note that $p_{l+1}^{(2)}\mu = p_{m-1}^{(1)}\lambda_1$, and that $p_{m-1}^{(1)}$ is already known in terms of p_l . Therefore

$$p_{l+j}^{(2)} = \rho_1 \frac{1 - \rho_2^j}{1 - \rho_2} p_{m-1}^{(1)}, \quad j = 1, 2, \dots, k.$$

Finally, for the states beyond m , straightforward local balance yields

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

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

$$p_l = \begin{cases} \left[\frac{k(1-\rho_2\eta_1)}{\eta_1(1-\eta_1^k)(1-\rho_2)} - \frac{\eta_1^{l+1}}{1-\eta_1} \right]^{-1}, & \eta_1 \neq 1 \\ \left[l + \frac{k+1}{2} + \frac{1}{1-\rho_2} \right]^{-1}, & \eta_1 = 1. \end{cases} \quad (3)$$

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

$$\mathbb{P}\{Q \geq M\} = \sum_{j \geq M} p_j = \rho_2^{M-m} \rho_1 \frac{1 - \rho_2^k}{(1 - \rho_2)^2} p_{m-1}^{(1)}. \quad (4)$$

We define the control rate simply as the average number of control packets transmitted by the queue observer per unit time. Since 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 given by

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

Since $\lambda_1 p_{m-1}^{(1)} = \mu p_{l+1}^{(2)}$, we have

$$R = 2\lambda_1 p_{m-1}^{(1)} = \begin{cases} \frac{2\lambda_1(1-\eta_1)}{1-\eta_1^k} p_l, & \eta_1 \neq 1 \\ \frac{2\lambda_1 p_l}{k}, & \eta_1 = 1 \end{cases} \quad (5)$$

where p_l was found in terms of the system parameters in (3).

It is clear from (4) and (5) that k determines the tradeoff 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.

As we mentioned earlier, the threshold l is computed based on the throughput guarantee γ , using the throughput expression (2) for the single threshold policy. The throughput obtained under the two-threshold policy can be computed using the steady-state probabilities derived above. The throughput expression (for $\eta_1 \neq 1$) is given by

$$\frac{(k(1 - \eta_1) - \eta_1^{l+1}(1 - \eta_1^k)) \eta_1(1 - \rho_2)}{k(1 - \eta_1)(1 - \rho_2\eta_1) - \eta_1^{l+2}(1 - \eta_1^k)(1 - \rho_2)} (\lambda_1 - \lambda_2) + \lambda_2. \quad (6)$$

By direct computation, it can be seen from the above expression that the throughput for $k > 1$ is no smaller than the throughput obtained under the single threshold policy (i.e., $k = 1$). Thus,




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

the two-threshold policy satisfies the throughput guarantee for all $k > 1$.

B. Large Deviation Exponents

In many simple 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

$$E = \lim_{M \rightarrow \infty} -\frac{1}{M} \log \mathbb{P}\{Q \geq M\}$$

when the limit exists. We now compute the LDE for the two-threshold policy.

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

$$E = \log \frac{1}{\rho_2}. \quad (7)$$

The above result follows from the congestion probability expression (4) since the only term that is exponential in M is $\rho_2^{M-m} = \rho_2^{M-l-k}$. Note that once l has been determined from the throughput requirement, it does not scale with M . When k grows sub-linearly in M , it should not be surprising that the two-thresholds are “invisible” in the large deviation limit. We pause to make the following observations:

- If k scales linearly with M as $k(M) = \beta M$ for some constant $\beta > 0$, the LDE becomes $E = (1 - \beta) \log \frac{1}{\rho_2}$, for $\rho_1 \geq 1$ and $E = (1 - \beta) \log \frac{1}{\rho_2} + \beta \log \frac{1}{\rho_1}$ for $\rho_1 < 1$.
- The control rate (5) 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 $-\log \rho_2$) 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 cannot be surpassed by any flow control policy.

C. More General Markovian Policies and Relationship to RED

In the previous section, we analyzed the two-threshold policy and concluded that it has the optimal congestion probability exponent for any rate of control. This is essentially because the input rate switches to the lower value deterministically, well before the congestion limit M is reached. In this subsection, we show that the two-threshold policy can be easily modified to a more general Markovian policy, which closely resembles the well known RED active queue management scheme [2]. Furthermore, this modification can be done while maintaining the optimal exponent behavior for the congestion probability.

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 host responds to dropped packets). If the queue length is between l and m , packets are randomly dropped with some probability q .²

Consider the following flow control 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 , a congestion notification occurs with probability³ 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.

We can 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

²Often, the dropping probability is dependent on the queue length.

³We can also let this probability to depend on the current queue length, as often done in RED, but this makes the analysis more difficult.

been determined from the throughput guarantee, the control rate versus 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 $\log \frac{1}{\rho_2}$. 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 model more practical queue management policies like RED.

IV. EFFECT OF CONTROL ERASURES ON CONGESTION

In this section, we investigate the impact of control erasures on the congestion probability of the two-threshold policy. We use a simple probabilistic model for the erasures 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 erasure probability, say δ^* , beyond which the erasures in receiving the control packets lead to an exponential degradation of the congestion probability.

A. Two-Threshold Policy Over an Erasure-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,⁴ and sends another control packet the next time an arrival occurs to a system with at least $m - 1$ packets.

In practice, the erasures on the control channel would affect the rate increasing control packets as well as the rate decreasing ones. However, in the interest of simplicity, we ignore the erasures in the rate increasing control packets, as these erasures do not affect the asymptotic behavior of the congestion probability. Thus, we assume in what follows that the rate decreasing control packets are erased with probability δ , while the rate increasing control packets are always received correctly.

The process $Y(t) = (Q(t), A(t))$ is a Markov process even for this erasure-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 erasures, 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. The following theorem establishes the LDE of the erasure-prone two-threshold policy, as a function of the erasure probability δ .

⁴This is an idealized assumption; in practice, delayed feedback can be obtained using ACKS.




Fig. 4. Markov process $Y(t)$ corresponding to the erasure-prone two-threshold policy. Only a part of the state space (with $m - 1 \leq Q(t) \leq M$) is shown.

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 queue is stable for all $\delta < 1$, and the LDE corresponding to the congestion probability is given by

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

where δ^* is the critical erasure probability given by

$$\delta^* = \frac{\rho_2}{\rho_1} (1 + \rho_1 - \rho_2). \quad (9)$$

Before we give a proof of this result, we pause to discuss its implications. The theorem shows that the two-threshold policy over an erasure-prone channel has two regimes of operation. In particular, for “small enough” erasure probability ($\delta < \delta^*$), the exponential rate of decay of the congestion probability is the same as in an erasure-free system. However, for $\delta > \delta^*$, the decay exponent starts decreasing, and, therefore, the congestion probability suffers an exponential increase. For this reason, we refer to δ^* as the critical erasure probability. Fig. 5 shows a plot of the decay exponent as a function of the erasure 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 .

Proof: We first need to prove that the Markov chain is positive recurrent for all $\delta < 1$, so that the congestion probability is well defined. To do this, we will write out the balance equations, and obtain one possible set of “candidate” steady-state probabilities. According to Theorem 3 in [4, Chapter 5], if a set of positive, normalized numbers that satisfy the balance equations are found, then the solution is unique, and the irreducible Markov chain (with countable state-space) is positive recurrent. We will also be able to obtain the desired LDE expression, once we derive the steady-state probabilities.

As before, we will first express all the steady-state probabilities in terms of p_i . The balance equations for the states with queue length less than m are as in the erasure-free case. Thus

$$p_i = p_l \rho_1^{l-i}, \quad 0 \leq i \leq l. \quad (10)$$




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

Next, for the set of states $l + i^{(1)}$, $i = 1, 2, \dots, k - 1$

$$p_{l+i+1}^{(1)} - (1 + \rho_1)p_{l+i}^{(1)} + p_{l+i-1}^{(1)}\rho_1 = 0, \quad i = 1, 2, \dots, k - 1.$$

Using the above, we can express $p_{l+i}^{(1)}$, $i = 2, \dots, k$ in terms of $p_{l+1}^{(1)}$ and p_l as

$$p_{l+i}^{(1)} = p_{l+1}^{(1)} \frac{1 - \rho_1^i}{1 - \rho_1} - p_l \frac{\rho_1 - \rho_1^i}{1 - \rho_1}, \quad i = 2, \dots, k. \quad (11)$$

Similarly, for the states $l + i^{(2)}$, $i = 2, \dots, k$, we have as before

$$p_{l+i}^{(2)} = \frac{1 - \rho_2^i}{1 - \rho_2} p_{l+1}^{(2)}, \quad i = 2, \dots, k. \quad (12)$$

The balance equations for the top set of states in Fig. 4 can be written as

$$p_{i+1}^{(1)}\mu - (\lambda_1 + \mu)p_i^{(1)} + \delta\lambda_1 p_{i-1}^{(1)} = 0, \quad i = m, m+1, \dots \quad (13)$$

Solving the second order recurrence relation above, we find that the top set of states in Fig. 4 (which correspond to arrival rate λ_1) have steady-state probabilities of the form

$$p_{m-1+i}^{(1)} = As(\delta)^i + Br(\delta)^i, \quad i = 1, 2, \dots \quad (14)$$

where

$$s(\delta) = \frac{1 + \rho_1 - \sqrt{(1 + \rho_1)^2 - 4\rho_1\delta}}{2} \quad (15)$$

and

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

It is easy to show that $s(\delta) < 1 < r(\delta)$ for all $\delta < 1$. Since we are looking for probabilities that decay to zero, the growing

term $r(\delta)^i$ is inadmissible. Therefore, we set $B = 0$ in (14), and obtain the form

$$p_{m-1+i}^{(1)} = s(\delta)^i p_{m-1}^{(1)}, \quad i = 1, 2, \dots \quad (16)$$

which satisfies (13). We now use the relation $p_m^{(1)} = s(\delta)p_{m-1}^{(1)}$ in (11) to write

$$p_{l+1}^{(1)} \frac{1 - \rho_1^k}{1 - \rho_1} - p_l \frac{\rho_1 - \rho_1^k}{1 - \rho_1} = s(\delta) \left(p_{l+1}^{(1)} \frac{1 - \rho_1^{k-1}}{1 - \rho_1} - p_l \frac{\rho_1 - \rho_1^{k-1}}{1 - \rho_1} \right)$$

from which we can obtain $p_{l+1}^{(1)}$ in terms of p_l :

$$p_{l+1}^{(1)} = p_l \left[\frac{1 - \eta_1^{k-1} - s(\delta)(\eta_1 - \eta_1^{k-1})}{1 - \eta_1^k - s(\delta)(\eta_1 - \eta_1^k)} \right]. \quad (17)$$

We next use the relation $p_l = \eta_1(p_{l+1}^{(1)} + p_{l+1}^{(2)})$ to obtain $p_{l+1}^{(2)}$ in terms of p_l

$$\eta_1 p_{l+1}^{(2)} = p_l \left[\frac{(1 - \eta_1)(1 - \eta_1 s(\delta))}{1 - \eta_1^k - s(\delta)(\eta_1 - \eta_1^k)} \right]. \quad (18)$$

Finally, let us deal with the states $p_i^{(2)}$, $i = m+1, m+2, \dots$. The balance equations for these states read

$$p_{i+1}^{(2)}\mu - (\lambda_2 + \mu)p_i^{(2)} + \lambda_2 p_{i-1}^{(2)} + (1 - \delta)\lambda_1 p_{i-1}^{(1)} = 0$$

for $i \geq m$. Since $p_m^{(2)}$, $p_{m-1}^{(2)}$, and $p_{l+1}^{(2)}$ have already been determined in terms of p_l , we can use the above recursion to obtain

$$p_{m+i}^{(2)} = \begin{cases} p_m^{(2)} \rho_2^i - p_{l+1}^{(2)} s(\delta) \frac{s(\delta)^i - \rho_2^i}{s(\delta) - \rho_2}, & s(\delta) \neq \rho_2 \\ p_m^{(2)} \rho_2^i - p_{l+1}^{(2)} i \rho_2^i, & s(\delta) = \rho_2. \end{cases} \quad (19)$$

Thus, among (10), (17), (18), (11), (12), (16), and (19), we have expressed the probabilities of all states in terms of p_l . The value of p_l can finally be obtained by normalizing the probabilities. Since $\rho_2 < 1$, and $s(\delta) < 1$ for $\delta < 1$, the infinite summations converge, and p_l is obtained as a positive number. We have thus obtained a set of normalized numbers that, by construction, satisfy the balance equations. Using Theorem 3 in [4, Chapter 5], we conclude that the chain is positive recurrent, and that the steady-state probabilities that we obtained above are unique.

Let us now proceed to obtain the LDE of the congestion probability. From (16) and (19), we can deduce that the congestion probability has *two* terms that decay exponentially in the buffer size

$$\mathbb{P}\{Q \geq M\} = Cs(\delta)^{M-l-k} + D\rho_2^{M-l-k} \quad (20)$$

where C, D are constants. In order to compute the LDE, we need to determine which of the two exponential terms in (20) decays slower. It is seen by direct computation that $s(\delta) \leq \rho_2$ for $\delta \leq \delta^*$, where δ^* is given by (9). Thus, for erasure probabilities less than δ^* , ρ_2 dominates the rate of decay of the congestion probability. Similarly, for $\delta > \delta^*$, we have $s(\delta) > \rho_2$, and the LDE is determined by $s(\delta)$. This proves the theorem. \square

B. Repetition of Control Packets

Suppose we are given a control channel with probability of erasure δ that is greater than the critical erasure probability in (9). 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 erasure, 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 erasure 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 erasure should be no more than δ^* . Thus, $\delta^n \leq \delta^*$, so that the number of transmissions n should satisfy

$$n \geq \frac{\log \delta^*}{\log \delta} \quad (21)$$

in order to obtain the best possible LDE of $\log \frac{1}{\rho_2}$. 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 section, 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 (21). 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 erasures, 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 $n - 1$

$$\mu(n) = \mu \left[1 - \frac{n-1}{\Phi} \right]. \quad (22)$$

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 (22), 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 erasure 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 erasure $\delta \in [0, 1)$, 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

$$\rho_i(f) = \frac{\lambda_i}{\mu(f)} = \frac{\rho_i}{1-f}, i = 1, 2 \quad (23)$$

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

$$\delta^*(f) = \frac{\rho_2}{\rho_1} \left(1 + \frac{\rho_1 - \rho_2}{1-f} \right) \quad (24)$$

which is analogous to (9), with $\rho_i(f)$ replacing ρ_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 erasure 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

$$E(\delta, f) = \begin{cases} \log \frac{1}{\rho_2(f)}, & \delta^{\Phi f + 1} \leq \delta^*(f) \\ \log \frac{1}{s(f, \delta)}, & \delta^{\Phi f + 1} > \delta^*(f) \end{cases} \quad (25)$$

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 erasure $\delta^{\Phi} f^{+1}$.

Definition 3: For any given $\delta \in [0, 1)$, the optimal fraction $f^*(\delta)$ is the value of f that maximizes $E(\delta, f)$ in (25). Thus

$$f^*(\delta) = \operatorname{argmax}_{f \in [0, 1 - \rho_2]} E(\delta, f). \quad (26)$$

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) $\Phi \leq \underline{\Phi}$, (ii) $\Phi \geq \bar{\Phi}$, and (iii) $\underline{\Phi} < \Phi < \bar{\Phi}$, where

$$\begin{aligned} \underline{\Phi} &= \frac{\rho_2 - \delta^*}{\log(\delta^*)(1 + \rho_1 - \rho_2)} \\ \bar{\Phi} &= \begin{cases} \frac{1}{\rho_1 - 1} & \rho_1 > 1 \\ \infty & \rho_1 \leq 1 \end{cases}. \end{aligned} \quad (27)$$

It can be shown that $\underline{\Phi} < \bar{\Phi}$ 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 nonexistent if $\rho_1 \leq 1$, 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 ($\Phi \leq \underline{\Phi}$): $f^*(\delta) = 0, \forall \delta \in (0, 1)$.
- (ii) Large Φ regime ($\Phi \geq \bar{\Phi}$):

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

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

$$\delta^{\Phi} \hat{f}^{+1} = \frac{\rho_2}{\rho_1} \left(1 + \frac{\rho_1 - \rho_2}{1 - \hat{f}} \right). \quad (28)$$

- (iii) Intermediate regime ($\underline{\Phi} < \Phi < \bar{\Phi}$): 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 (28) and $\tilde{f}(\delta)$ is the unique solution in $(0, 1 - \rho_2)$ to the transcendental equation

$$\delta^{\Phi} \tilde{f}^{+1} [\Phi(1 - \tilde{f}) \log \delta^* + 1] = s(\tilde{f}, \delta). \quad (29)$$

The proof of the above theorem is not particularly interesting, and is postponed to the appendix. Instead, we provide some intuition about the optimal solution.

C. Discussion of the Optimal Solution

1) *Erasure 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 $-\log \rho_2$ 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 erasure 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 erasure probability $\delta^{\Phi} f^{+1}$. This fraction is indeed \hat{f} , defined by (28). 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 $\log \frac{1 - \hat{f}(\delta)}{\rho_2}$ 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 (21). Recall that the naïve policy simply repeats the control packets to make the effective erasure 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 (21) to be

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

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 Fig. 6(b), both using dashed lines. As seen in the figure, the naïve 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







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

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

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 erasure probability is made equal to the knee point. However, at a particular value of erasure probability, say δ' , the optimal fraction begins to decrease sharply from the $\hat{f}(\delta)$ curve, and reaches zero at some value δ'' . Equation (29) characterizes the optimal fraction for values of δ in (δ', δ'') . No redundancy is applied for $\delta \in (\delta'', 1)$.

For this range of erasure 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 erasure probabilities $\delta < \delta'$, and the small Φ regime for large erasure probability $\delta > \delta''$. There is also a non empty “transition interval” in between the two, namely (δ', δ'') .

VI. QUEUE LENGTH INFORMATION AND SERVER ALLOCATION POLICIES

In this section, we discuss the role of queue length information on server allocation policies in a single server queue. We mentioned earlier that queue aware resource allocation policies tend to allocate a higher service rate to longer queues, and vice-versa. Intuitively, if the controller is frequently updated with accurate queue length information, the service rate can




Fig. 8. Single server queue with service rate control.

be adapted to closely reflect the changing queue length. However, if the queue length information is infrequently conveyed to the controller, we can expect a larger queue length variance, and hence, a higher probability of congestion. We study this tradeoff between the probability of congestion and the rate of queue length information in a single server queue.

Fig. 8 depicts a single server queue with Poisson inputs of rate λ . An observer watches the queue evolution and sends control information to the service rate controller, which changes the service rate $S(t)$ based on the control information it receives. The purpose of the observer-controller subsystem is to assign service rates at each instant so as to control congestion in the queue.

For analytical simplicity, we assume that the service rate at any instant is chosen to be one of two distinct values: $S(t) \in \{\mu_1, \mu_2\}$, where $\mu_2 > \mu_1$ and $\mu_2 > \lambda$. The control decisions are sent by the observer in the form of information-less packets. Upon receiving a control packet, the rate controller switches the service rate from one to the other. As before, we only focus on Markovian control policies, which are defined analogously to (1).

Note that if there is no restriction imposed on using the higher service rate μ_2 , it is optimal to use it all the time, since the congestion probability can be minimized without using any control information. However, in a typical queueing system with limited resources, it may not be possible to use higher service rate at all times. There could be a cost per unit time associated with using the faster server, which restricts its use when the queue occupancy is high. Alternately, one could explicitly restrict the use of the faster server by allowing its use only when the queue occupancy is over a certain threshold value. In this paper, we impose the latter constraint, i.e., when the queue length is no more than some threshold l , we are forced to use the lower service rate μ_1 . If the queue length exceeds l , we are allowed to use the higher rate μ_2 without any additional cost until the queue length falls back to l .⁵

It turns out that this model is, in a certain sense, dual to the flow control problem considered earlier in this paper. In fact, for every Markovian flow control policy operating on the queue in Fig. 1, it is possible to identify a corresponding server allocation policy which has identical properties. For example, we can define a two-threshold server allocation policy analogously to the flow control policy as follows:

The service rates are switched at two distinct queue length thresholds l and m . Specifically, when the queue length grows past m , the service rate switches to μ_2 . Once the higher service

rate is employed, it is maintained until the queue length falls back to l , at which time the service rate switches back to μ_1 . The Markov process corresponding to the two-threshold server allocation policy is depicted in Fig. 9.

Evidently, the Markov chain in Fig. 9 has the same structure as the chain in Fig. 2, and therefore, can be analyzed in the same fashion. In particular, we can derive the control rate versus congestion probability tradeoff, along the same lines as (4) and (5). The following result regarding the two-threshold server allocation policy, can be derived along the lines of Proposition 1 and Theorem 1.

Theorem 4: Suppose that k goes to infinity sub-linearly in the buffer size M , in a two-threshold server allocation policy. Then, the LDE can be maintained constant at

$$E = \log \frac{\mu_2}{\lambda}$$

while the control rate can be made arbitrarily small. Further, the two-threshold policy has the largest possible congestion probability LDE among all server allocation policies, for any rate of control.

The above result shown the optimality of the two-threshold policy with respect to the congestion probability exponent. Next, if the control signals that lead to switching the service rate are subject to erasures (as detailed in Section IV), we can show that the LDE behaves exactly as in Theorem 2.

In essence, we conclude that both the flow control and resource allocation problems in a single server queue lead to the same mathematical framework, and can thus be treated in a unified fashion.

VII. CONCLUSION

The goal of this paper was to study the role played by queue length information in flow control and resource allocation policies. Specifically, we deal with the question of how often queue length information needs to be conveyed 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 in which the control decisions are based on the queue occupancy. We learned that in the absence of control channel erasures, the control rate needed to ensure the optimal decay exponent for the congestion probability can be made arbitrarily small. However, if control channel erasures occur probabilistically, we showed the existence of a critical erasure 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 erasure probabilities larger than the critical value, a significant fraction of the system resources may be consumed by the control signals, unlike in the erasure-free scenario. We also pointed out that allocating control resources without considering the bandwidth they consume, might have adverse effects on congestion. We also observed that the server allocation problem and the flow control problem can be treated in a mathematically unified manner.

⁵Recall that in the flow control problem, the threshold l was derived from a throughput constraint.




Fig. 9. Markov process corresponding to the two-threshold server allocation policy.

APPENDIX
PROOF OF THEOREM 3

Given Φ and δ we want to find the fraction f that satisfies (26). As shown in Fig. 5, the LDE curve for $f = 0$ is flat and has the highest possible value of $-\log \rho_2$ for $\delta \in [0, \delta^*]$. Indeed, for δ in the above range, using any strictly positive fraction f would reduce the LDE to $\log \frac{1-f}{\rho_2}$. This implies that the optimal fraction

$$f^*(\delta) = 0, \delta \in [0, \delta^*].$$

Thus, the problem of finding the optimal $f^*(\delta)$ is nontrivial only for $\delta \in (\delta^*, 1)$. We begin our exposition regarding the optimal fraction with two simple propositions.

Proposition 3: For any given $\delta \in (\delta^*, 1)$, the optimal fraction $f^*(\delta)$ is such that the effective erasure probability $\delta^{\Phi f^*+1}$ cannot be strictly lesser than the knee point of the curve $E(\delta, f^*)$. That is, $\delta^*(f^*) \leq \delta^{\Phi f^*+1}$, $\delta \in (\delta^*, 1)$.

Proof: Suppose the contrary, i.e. for some $\delta \in (\delta^*, 1)$, the optimal f^* is such that $\delta^*(f^*) > \delta^{\Phi f^*+1}$. The optimal LDE would then be $E(\delta, f^*) = -\log \rho_2(f^*) = \log \frac{1-f^*}{\rho_2}$. Continuity properties imply that $\exists \xi > 0$ for which $\delta^*(\hat{f}^* - \xi) > \delta^{\Phi(f^*-\xi)+1}$. Thus, if we use the smaller fraction $f^* - \xi$ for control, the LDE would be $E(\delta, f^* - \xi) = \log \frac{1-f^*+\xi}{\rho_2}$. Since this value is greater than the ‘‘optimal value’’ $E(\delta, \hat{f}^*)$, we arrive at a contradiction. \square

Proposition 4: For a given $\delta \in (\delta^*, 1)$, there exists a unique fraction $\hat{f}(\delta) \in (0, 1 - \rho_2)$ such that the knee point $\delta^*(\hat{f})$ equals the effective erasure probability $\delta^{\Phi \hat{f}+1}$. Furthermore, the optimal fraction $f^*(\delta)$ lies in the interval $[0, \hat{f}(\delta)]$.

Proof: The knee point corresponding to any fraction f is given by (24). Therefore, if there exists a fraction \hat{f} for which $\delta^*(\hat{f}) = \delta^{\Phi \hat{f}+1}$, then \hat{f} satisfies

$$\delta^{\Phi \hat{f}+1} = \frac{\rho_2}{\rho_1} \left(1 + \frac{\rho_1 - \rho_2}{1 - \hat{f}} \right). \quad (30)$$

The transcendental equation in (30) has a solution in $(0, 1 - \rho_2)$ for any given $\delta \in (\delta^*, 1)$. This can be shown by applying the intermediate value theorem to the difference of the right and left hand side functions in the equation. The uniqueness of \hat{f} follows from the monotonicity properties of the right and left hand side functions.

To prove the second statement, suppose that $f^* > \hat{f}$. Since the knee point (24) is monotonically strictly increasing in f , we

have $\delta^*(f^*) > \delta^*(\hat{f}) = \delta^{\Phi \hat{f}+1} > \delta^{\Phi f^*+1}$. This contradicts Proposition 3. \square

The above proposition shows that the optimal fraction lies in the interval $[0, \hat{f}(\delta)]$. Thus, for a given $\delta > \delta^*$, we seek $f^*(\delta) \in [0, \hat{f}(\delta)]$ for which $s(f, \delta)$ (defined in Proposition 2) is minimized. In particular, if $s(f, \delta)$ is monotonically decreasing in $[0, \hat{f}(\delta)]$ for some δ , then clearly, $f^*(\delta) = \hat{f}(\delta)$. The following proposition asserts the condition under which $s(f, \delta)$ is monotonically decreasing.

Proposition 5: For some $\delta > \delta^*$, suppose the following inequality holds:

$$\delta^{\Phi \hat{f}+1} [\Phi(1 - \hat{f}) \log \delta + 1] \leq \frac{\rho_2}{1 - \hat{f}}. \quad (31)$$

Then, $f^*(\delta) = \hat{f}(\delta)$.

Proof: Fix $\delta > \delta^*$. By direct computation, we find that

$$s'(f, \delta) \leq 0 \iff \delta^{\Phi f+1} [\Phi(1 - f) \log \delta + 1] \leq s(f, \delta).$$

However, since the left side of the inequality above is strictly increasing in f , we find that $s(f, \delta)$ is monotonically decreasing whenever

$$\delta^{\Phi \hat{f}+1} [\Phi(1 - \hat{f}) \log \delta + 1] \leq s(\hat{f}(\delta), \delta) = \frac{\rho_2}{1 - \hat{f}}$$

where the last equality follows from the definition of \hat{f} . Thus, if (31) is satisfied for a particular δ , $s(f, \delta)$ is decreasing in f , and hence, the optimal fraction is given by $f^*(\delta) = \hat{f}(\delta)$. \square

Now, suppose that $\Phi > \bar{\Phi}$. Upon rearrangement, this implies that $\delta^*[\Phi \log \delta^* + 1] < \rho_2$. In other words, (31) is satisfied with strict inequality, at $\delta = \delta^*$. By continuity, we can argue that there exist a range $\delta \in (\delta^*, \delta')$ for which (31) is satisfied. By Proposition 5, we have $f^*(\delta) = \hat{f}(\delta)$, $\delta \in (\delta^*, \delta')$, which partially proves part (iii) of the theorem. Note that δ' is the smallest value of the erasure probability, if any, for which the strict monotonicity of $s(f, \delta)$ (as a function of f) is compromised. As argued in Proposition 5, this implies (31) holds with equality. A simple rearrangement yields a transcendental equation for δ'

$$\frac{\rho_1}{\rho_1 - \rho_2 + 1 - \hat{f}(\delta')} = 1 + \Phi(1 - \hat{f}(\delta')) \log \delta'. \quad (32)$$

Using basic calculus, it is possible to show that there always exists a solution $\delta' < 1$ to (32) if $\rho_1 < 1$. However, if $\rho_1 > 1$, there exists a solution iff $\Phi < \frac{1}{\rho_1 - 1} = \bar{\Phi}$. In particular, if $\Phi >$

$\bar{\Phi}$, we have that $s(f, \delta)$ is monotonically decreasing in f for all δ , so that $f^*(\delta) = \hat{f}(\delta)$ for all $\delta \in (\delta^*, 1]$. This proves part (ii) of the theorem.

Finally, let us consider the case $\underline{\Phi} < \Phi < \bar{\Phi}$. (It is straightforward to show that $\underline{\Phi} < \bar{\Phi}$). Then a solution $\delta' < 1$ to (32) exists, and for $\delta > \delta'$, the optimal fraction is no longer equal to $\hat{f}(\delta)$. It can also be shown that the existence of $\delta' < 1$ guarantees the existence of another $\delta'' > \delta'$ such that $s(f, \delta)$ is increasing in f for each $\delta \in (\delta'', 1)$. In such a case, $f^*(\delta)$ would be equal to zero.

Proposition 6: For some $\delta > \delta^*$, suppose the following inequality holds:

$$\delta[\Phi \log \delta + 1] \geq \frac{1 + \rho_1 - \sqrt{(1 + \rho_1)^2 - 4\delta\rho_1}}{2}. \quad (33)$$

Then, $f^*(\delta) = 0$.

The proof is similar to Proposition 5. The value of δ'' is obtained as a solution to the transcendental equation

$$\delta''[\Phi \log \delta'' + 1] = \frac{1 + \rho_1 - \sqrt{(1 + \rho_1)^2 - 4\delta''\rho_1}}{2}. \quad (34)$$

Again using basic calculus, we can show there exists a solution $\delta'' < 1$, to (34) if $\Phi < \bar{\Phi}$, which is the same condition as for the existence of δ' . Thus, in the intermediate regime, there exists (i) $\delta' \in (\delta^*, 1)$ such that $f^*(\delta) = \hat{f}(\delta)$ for $\delta^* < \delta < \delta'$, and (ii) $\delta'' > \delta'$ such that $f^*(\delta) = 0$ for $\delta \geq \delta''$. For $\delta \in (\delta', \delta'')$, the function $s(f, \delta)$ has a minimum in $(0, \hat{f})$, so that the optimum fraction is obtained by setting $s'(f, \delta)$ to zero. This condition is the same as (29), and part (iii) of the theorem is thus proved.

We finally show that it is optimal to not use any redundancy in the small Φ regime.

Proposition 7: For $\Phi \leq \underline{\Phi}$, $s(f, \delta)$ is monotone increasing in $f \in [0, \hat{f}(\delta)]$ for each $\delta > \delta^*$.

Proof: Simple rearrangement shows that the condition $\Phi \leq \underline{\Phi}$ is equivalent to saying that $\delta'' \leq \delta^*$. Since no redundancy is used for values of erasure probability greater than δ'' , it follows that no redundancy is used in the small Φ regime. \square

REFERENCES

- [1] D. Bertsimas, I. Paschalidis, and J. Tsitsiklis, "Asymptotic buffer overflow probabilities in multiclass multiplexers: An optimal control approach," *IEEE Trans. Autom. Control*, vol. 43, no. 3, pp. 315–335, Mar. 1998.
- [2] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Netw.*, vol. 1, no. 4, pp. 397–413, Apr. 1993.
- [3] R. Gallager, "Basic limits on protocol information in data communication networks," *IEEE Trans. Inf. Theory*, vol. IT-22, no. 4, pp. 385–398, Jul. 1976.
- [4] R. Gallager, *Discrete Stochastic Processes*. Boston, MA: Kluwer, 1996.
- [5] K. Jagannathan, E. Modiano, and L. Zheng, "Effective resource allocation in a single-server queue: How much control is necessary," presented at the 46th Annu. Allerton Conf. Communication, Control, and Computing, Monticello, IL, 2008.
- [6] K. Jagannathan, E. Modiano, and L. Zheng, "On the trade-off between control rate and congestion in single server systems," in *Proc. IEEE INFOCOM*, 2009, pp. 271–279.
- [7] X. Lin, N. Shroff, and R. Srikant, "On the connection-level stability of congestion-controlled communication networks," *IEEE Trans. Inf. Theory*, vol. 54, no. 5, pp. 2317–2338, May 2008.
- [8] M. Neely, "Order optimal delay for opportunistic scheduling in multi-user wireless uplinks and downlinks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 5, pp. 1188–1199, May 2008.
- [9] M. Neely, E. Modiano, and C. Li, "Fairness and optimal stochastic control for heterogeneous networks," *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 396–409, Feb. 2008.
- [10] M. Neely, E. Modiano, and C. Rohrs, "Dynamic power allocation and routing for time varying wireless networks," presented at the IEEE INFOCOM 22nd Annu. Joint Conf. IEEE Computer and Communications Societies, 2003.
- [11] J. Perkins and R. Srikant, "The role of queue length information in congestion control and resource pricing," in *Proc. IEEE Conf. Decision and Control*, 1999, vol. 3, pp. 2703–2708.
- [12] L. Tassiulas and A. Ephremides, "Stability properties of constrained queueing systems and scheduling policies for maximum throughput in multihop radio networks," *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, Dec. 1992.
- [13] L. Tassiulas and A. Ephremides, "Dynamic server allocation to parallel queues with randomly varying connectivity," *IEEE Trans. Inf.*, vol. 39, no. 2, pp. 466–478, Feb. 1993.

Krishna Jagannathan obtained his B.Tech. in Electrical Engineering from the Indian Institute of Technology Madras, India in 2004, and the S.M. and Ph.D. degrees from the Massachusetts Institute of Technology in 2006 and 2010, respectively. He worked as a consultant at the Mathematical Sciences Research Center, Bell Laboratories, Murray Hill, NJ in 2005, an engineering intern at Qualcomm, Campbell, CA, in 2007, and a summer research fellow at the Indian Institute of Science (IISc), Bangalore, India, in 2003. His research interests lie in communication networks, queueing theory, and network control.

Eytan Modiano (S'90–M'93–SM'00) received his B.S. degree in Electrical Engineering and Computer Science from the University of Connecticut at Storrs in 1986 and his M.S. and Ph.D. degrees, both in Electrical Engineering, from the University of Maryland, College Park, MD, in 1989 and 1992, respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Post Doctoral Fellow during 1992–1993. Between 1993 and 1999, he was with MIT Lincoln Laboratory where he was a project leader for MIT Lincoln Laboratory's Next Generation Internet (NGI) project. Since 1999, he has been on the faculty at MIT; where he is a Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research is on communication networks and protocols with emphasis on satellite, wireless, and optical networks. He is the co-recipient of the Sigmetrics 2006 Best paper award for the paper "Maximizing Throughput in Wireless Networks via Gossiping," and the Wiopt 2005 best student paper award for the paper "Minimum Energy Transmission Scheduling Subject to Deadline Constraints."

He is currently an Associate Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. He served as Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY and as guest editor for IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS special issue on WDM network architectures; the *Computer Networks Journal* special issue on Broadband Internet Access; the *Journal of Communications and Networks* special issue on Wireless Ad-Hoc Networks; and for the IEEE JOURNAL OF LIGHTWAVE TECHNOLOGY special issue on Optical Networks. He was the Technical Program co-chair for IEEE Wiopt 2006, IEEE Infocom 2007, and ACM MobiHoc 2007.

Lizhong Zheng (S'00–M'02) received the B.S. and M.S. degrees from the Department of Electronic Engineering, Tsinghua University, Beijing, China, in 1994 and 1997, respectively, and the Ph.D. degree from the Department of Electrical Engineering and Computer Sciences, University of California (UC) Berkeley, in 2002.

Since 2002, he has been with the Department of Electrical Engineering and Computer Sciences, Massachusetts Institute of Technology (MIT), Cambridge, where he is currently an Associate Professor. His research interests include information theory, wireless communications, and wireless networks.

Dr. Zheng received the Eli Jury Award from UC Berkeley in 2002, the IEEE Information Theory Society Paper Award in 2003, the National Science Foundation (NSF) CAREER award in 2004, and the AFOSR Young Investigator Award in 2007.