

On the Elasticity of Marking Functions in an Integrated Network

Yung Yi, *Member, IEEE*, and Sanjay Shakkottai, *Member, IEEE*

Abstract—Much of the research on network modeling and analysis has focused on the design of end controllers and network algorithms with the objective of stability and convergence of the transmission rate. However, a network is typically composed of a mixture of both controlled elastic flows and uncontrolled real-time flows. In this paper, we study the effects of marking elasticity (which characterizes how aggressively the marking function responds to congestion) on queue overflow probability for uncontrolled real-time flows.

First, we derive lower and upper bounds on the queue overflow probability at a router of a single bottleneck system. Using this, we quantify the trade-off between stability for controlled flows and queue overflow probability for uncontrolled real-time flows as a function of marking elasticity. Next, we compare the capacity required at a router with only FIFO scheduling versus a router with priority scheduling (priority given to the real-time flows) for supporting a given queue overflow probability. We quantify the “scheduling-gain” of priority scheduling over FIFO scheduling, as a function of marking elasticity. We show that this scheduling gain decreases with more elastic marking functions.

Index Terms—Congestion control, queueing theory, stochastic processes.

I. INTRODUCTION

THERE has been extensive research on the modeling and analysis of the controlled elastic flows in networks by adopting differential equation based models of source controllers and Active Queue Management (AQM) algorithms. Much of this work has focused on the design of end host controllers and control algorithms (marking functions) at the intermediate routers for (global and local) stable end-to-end operation over the Internet by using control theoretic tools (see [1] and the references therein).

However, a network typically carries a mixture of traffic ranging from controlled non-real-time elastic data traffic to uncontrolled real-time traffic (e.g., voice and multimedia traffic). Uncontrolled real-time flows do not react to network feedback

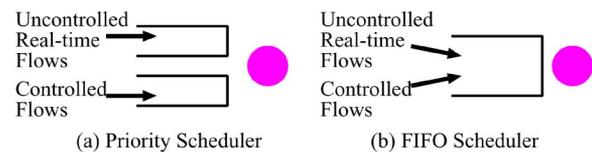


Fig. 1. Priority and FIFO scheduling disciplines.

and require guarantee of queue overflow probability (QOP) as its Quality of Service (QoS), where the QOP requirement stipulates that the probability of packet (from uncontrolled real-time flows) loss should not exceed some threshold. From the perspective of network control and management, real-time flows are admitted into the network only if there are “sufficient” resources to satisfy their QOP requirements. On the other hand, non-real-time sources are always admitted into the network on a best-effort basis, i.e., real-time sources are given higher priority, and the remaining network resources unused by the real-time sources are allocated to the non-real-time sources.

One of the simple architectures for providing differentiated services in the integrated network with controlled and uncontrolled flows is that users can belong to one of two classes, and routers in the network treat (schedule) packets from two classes in a differentiated manner by adopting “priority” based scheduling algorithms (see Fig. 1(a) for an example).

In this paper, we consider a network where resources are shared by uncontrolled real-time and controlled elastic flows, and packets in the router are scheduled in a first-come-first-serve manner (i.e., no differentiation) [see Fig. 1(b)]. Over such a network, the behavior of uncontrolled real-time and controlled flows are coupled together, and the QOP experienced by uncontrolled real-time flows will be affected by the behavior of controlled flows (due to the flows sharing a common link). With this setup, it seems reasonable to believe that by appropriately designing an AQM mechanism (marking function) at intermediate routers, we can potentially provide the required QOP to the uncontrolled real-time flows *without any differentiation between real-time and non-real-time flows* at the routers.

The intuition is the following: an “aggressive” marking function will mark a larger number of controlled flow packets when a burst of packets (which causes congestion) arrive. This will cause the controlled flows to back-off, thus potentially decreasing the delay or packet loss probability experienced by real-time flows. In this paper, we study the trade-off between packet marking [2] for controlled flows and the effect of this marking on the QOP of uncontrolled real-time flows.

We first characterize the “aggressiveness” of a marking function by its *elasticity*.

Manuscript received April 05, 2006; revised June 11, 2007. Current version published February 11, 2009. This work was supported in part by NSF Grants ACI-0305644, CNS-0325788, CNS-0347400, and Korea Research Council of Fundamental Science and Technology. Recommended by Associate Editor I.-J. Wang.

Y. Yi is with the Department of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon 305-701, Korea.

S. Shakkottai is with the Department of Electrical and Computer Engineering, University of Texas at Austin, Austin, TX 78712 USA (e-mail: shakkott@ece.utexas.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TAC.2008.2010697

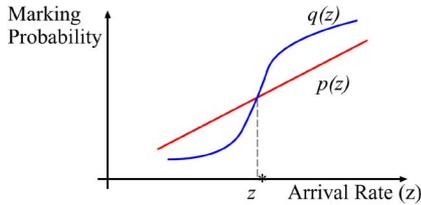


Fig. 2. Elasticity of marking functions.

Definition 1.1: Given any two marking functions $p(z)$ and $q(z)$, we say that $q(z)$ is more elastic than $p(z)$ if for any $z \geq z^*$, we have

$$q(z^*) = p(z^*), \quad q(z) \geq p(z)$$

where z^* is the equilibrium data rate at the router. Thus, the elasticity of a marking function corresponds to how aggressively the marking value changes as the arrival data rate exceeds the equilibrium rate (see Fig. 2). However, note that the marking values at the equilibrium rate are equivalent with both marking functions, i.e., the link utilization at equilibrium is the same.

In this paper, we study and quantify the following two trade-offs:

Stability-elasticity trade-off. This refers to the trade-off between QOP for uncontrolled real-time flows and stability for controlled flows. The key trade-off we quantitatively analyze is that the more elastic the marking function is, the better is the QOP of uncontrolled real-time flows guaranteed. However, this also leads to the negative property of less stability for controlled flows.

Scheduling-elasticity trade-off. This refers to the trade-off between the scheduling algorithm at the router and the elasticity of the marking function, with the performance metric being the QOP achieved for real-time flows. This trade-off is quantified by means of the following: Given the requirement of (a) long-term throughput (equilibrium rate) of controlled flows and (b) QOP for uncontrolled real-time flows, we compute the required link capacity at the router to support (a) and (b).

To satisfy the long-term rate (for controlled flows) constraint (a), it is sufficient that the link capacity exceeds the sum of the mean rates of uncontrolled flows and the equilibrium rates of controlled flows (as long as the scheduling algorithm is work-conserving). However, the magnitude of the excess capacity depends on the given QOP (b) for real-time flows, the marking function elasticity, as well as the scheduling algorithm. In particular, by giving absolute priority to the real-time flows (i.e., priority scheduling), this excess capacity can be minimized, as controlled flows do not affect the QOP for the uncontrolled flows. On the other hand, FIFO scheduling has the advantage of simple implementation (no per-class scheduling required), but at a cost of larger required link capacity. In this study, we quantify the “scheduling-gain” (i.e., the difference in the link capacity required with priority scheduling versus that with FIFO scheduling) of priority scheduling over FIFO scheduling, as a function of the marking elasticity.

The parameters that impact the source dynamics for a controlled flow are the round-trip delay, the elasticity of the marking

function, and the rate of adaptation at the controlled source. In this paper, we first model the dynamics of controlled flows by means of an instant adaptation algorithm, where the sources react to network feedback with no delay and adapt immediately to the equilibrium rate for a given network configuration. The instant adaptation scheme enables us to separate the effect of other parameters and to focus only on the elasticity of marking functions [3], [4]. In Section IV, we extend the discussion to a weighted proportional fair controller [5] with more complex temporal dynamics.

A. Main Contributions and Organization

The main contributions of this paper are the following:

- i) Using the instant adaptation model for source dynamics, we derive lower and upper bounds of the QOP at a router, where a single buffer is shared by controlled and uncontrolled real-time flows. Using these bounds, we quantify the trade-off between stability for controlled flows and QOP-guarantee for uncontrolled real-time flows as a function of the elasticity of the marking function. The results indicate that some marking functions may be “uniformly” better than others. In particular, among the marking functions that we have compared, our bounds indicate that a rate based version of REM [6] seems to provide the largest local-stability region for *any* given QOP requirement.
- ii) We next compare the capacity required at a router with only FIFO scheduling versus a router with priority scheduling for supporting a given QOP requirement. We quantify “scheduling-gain” of priority scheduling over FIFO scheduling, as a function marking elasticity. We show that this scheduling gain decreases with more elastic marking functions. This indicates that by appropriately choosing the marking function and by using only a FIFO queue at the router, we can satisfy the QOP requirement of real-time flows without much over-provisioning.
- iii) We extend the results to the case with a weighted proportional fair controller at the source, and study the trade-off between stability and marking elasticity. Finally, we validate our analytical results using simulations.

In the rest of this paper, we begin with a description of the system model, parameterization of marking elasticity and the problem statement in Section II. Next, in Section III, we derive an upper and lower bound on the QOP, from which we analytically show stability-elasticity trade-off and scheduling-elasticity trade-off with the instant adaptation algorithm. In Section IV, we discuss the stability-elasticity trade-off with weighted proportional fair controllers. In Section V, we provide numerical results and simulation to validate our analysis, and conclude in Section VI.

B. Related Work

The problem of determining the queue overflow probability has been studied extensively for queues [7]–[11] in the context of “open-loop” flows (i.e., there are no controlled flows). From a technical viewpoint, our research differs in that we use a sample path large deviations framework to analyze a system, where the flows react to the link conditions via the congestion controller dynamics.

There has been much recent work on understanding the dynamics of Internet congestion controllers in the presence of uncertainty (e.g., due to short flows, random flow initiations and terminations, and real-time flows). The current literature can be categorized by the *time-scales* we investigate.

First, from the viewpoint of the *router*, as the number of flows and capacity increase, the time-scale of queue becomes faster. In other words, the queueing dynamics become increasingly faster (equivalently, the busy-period duration becomes increasingly smaller) as the system scale size increases. In this regime, by appropriately “slowing-down” time at the router, it has been shown that the uncontrolled arrival processes can be modeled by a Poisson process [12] (formally, using the Poisson limit of renewal processes [13]), and there has been much work [14]–[16] in understanding the queueing dynamics (for instance, in the context of buffer sizing) at this time-scale at the router. We refer to this regime as the “fast time-scale regime,” as this regime captures the time-scale of the router dynamics.

The second important time-scale to consider is that of an end-user. This time-scale is n times “slower” than that of the router dynamics, where n is scaling parameter that determines the number of flows, and the link capacity scaling. The reason for this “slow-down” is that for every packet transmitted by an end-user, there are n packets that arrive at the intermediate router (as there are n flows in the system). From the viewpoint of the end-user, over any fixed time-interval (i.e., time is *not* scaled), one can view the queue as a multiplexer over a large number of random (open-loop) flows, leading to a deterministic limit in the large-number of flows regime. This is essentially the regime that has been studied in [17], where the authors have analyzed the fluid limit of the system from the viewpoint of the controlled flows. In this limit, the open-loop flows (the real-time flows in the context of this paper) appear as a constant rate process (with the rate equal to the expected value of the uncontrolled arrival process), and the closed-loop controlled flows appear as a deterministic delay-differential equation.

The third time-scale is due to that of flow-level arrivals and departures. In this time-scale, the major interest lies in understanding the impact of real-time flows on elastic flows by looking at the system dynamics at the flow-level time-scale and considering dynamic flow arrivals/departures [18]–[21]. A primary research objective at this time-scale is in characterizing flow-level stability (i.e., conditions under which the number of elastic flows remain finite).

Viewed in the above context, the work in this paper can be viewed as trying to model the behavior of a real-time (open-loop) end-user, at the time-scale of the end-user (i.e., the second time-scale discussed above). However, unlike [17] where a deterministic fluid model has been studied (which is of interest when studying the long-term throughput for a controlled flow), we consider the tail probability for open loop flows, as the performance metric here is the queue overflow probability of a real-time flow.

Finally, related work also includes [22], where the authors characterize queue fluctuations using a Central Limit Theorem (CLT) based approximation. This can be interpreted as understanding (i) the variance properties of the router queue as opposed to tail behavior (i.e., the analog of the difference between the large-deviations behavior of an “open-loop” queue and the

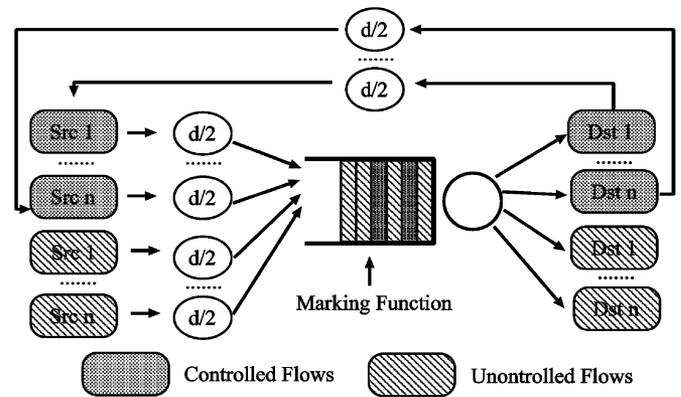


Fig. 3. System model.

diffusion approximation of the queue, however, in the context of a closed-loop system), and (ii) the time-scale of the router queue (see [22, Section 4]).

II. SYSTEM MODEL AND PROBLEM STATEMENT

A. System Model

Consider the system shown in Fig. 3. We consider a *single* discrete time queue with two types of flows: controlled flows and uncontrolled flows. We use the terminology *controlled flows* to refer to flows of data traffic which react and adapt their transmission rates to feedback from the network. *Uncontrolled flows*¹ refer to data flows that do not react to network feedback. Examples of such flows include real-time video/audio. The queue is fed by n independent identically distributed (over flows) stationary, ergodic uncontrolled flows, with each (discrete time) flow being a bounded process, and by n controlled flows (determined by a congestion control algorithm described later). The buffer size is scaled with n , and the link capacity of the corresponding queue is suitably scaled with n so that the queue is stable. Thus, the n -th system has a buffer of size nB , and a capacity of nC . For *queue stability*, we assume that $z^* = x^* + y^* < C$, where x^* is the equilibrium rate of a controlled flow, and y^* is the mean rate of an uncontrolled flow. Further, we assume that each flow experiences a round trip delay d .

From the controlled flows’ point of view, the system we have described above can be thought of as a closed loop system (with delay) and feedback control is applied at the router based on *aggregate arrivals*. A popular modeling and analysis methodology for such closed-loop systems in the Internet context has been through functional differential (or difference) equations based *fluid models* [23], [24].

The router is modeled by a *marking function* (see Section II-B) which signals congestion by marking flows, and receivers detect the marks and inform the respective flow sources to increase or decrease their transmission rate. We model flows by discrete time fluid processes. We denote the fluid rates of individual flows by $\{x_k[i], k = 1, \dots, n\}$, where $x_k[i]$ denotes the number of arrivals² of controlled flow k at

¹Throughout this paper, we use the term “uncontrolled flows” and “real-time flows” interchangeably.

²We use the terms “number of arrivals” and “arrivals” interchangeably. Further, the term “arrival rate” corresponds to the number of arrivals per time-slot.

time i . In this paper, we consider two source rate adaptation algorithms (e.g., [4], [17]), described as follows:

i) *Weighted proportional fair controller.*

$$x_k[i+1] - x_k[i] = \kappa \left(w - x_k[i-d]p \left(\sum_{j=1}^n x_j[i-d] + \sum_{j=1}^n y_j[i-d], nC \right) \right) \quad (1)$$

where $y_k[i]$ denotes the number of arrivals due to an uncontrolled flow k at time i . κ and w are positive constants which determine the rate at which each flow increases or decreases its transmission rate, and the equilibrium point.

ii) *Instant adaptation.*

$$w = x_k[i]p \left(\sum_{j=1}^n x_j[i] + \sum_{j=1}^n y_j[i], nC \right). \quad (2)$$

In the instant adaptation algorithm, congestion controllers adapt to the fixed point of the difference equation in (1) with no delay [3], [4]. In this scheme, as the rate of the uncontrolled flow varies with time, the corresponding equilibrium rate varies appropriately (as determined by the elasticity properties of $p(\cdot)$). The instant adaptation scheme tracks this variation of the equilibrium rate. This allows us to focus purely on the properties of the marking function and to ignore the effects of κ and d .

B. Marking Function

The marking function, $p(z, C)$ represents the fraction of flow to be marked when the total arrivals to the associated router with capacity C is z . In this paper, we study rate-based marking functions, where the router marks incoming packets based on the aggregate instantaneous arrival rate. Another class of widely studied marking functions are those that mark the packets based on the queue length (queue-based marking functions, e.g., RED [2]). Further, in [14], [15], the authors show that under appropriate assumptions, a queue-based marking function can be approximated by an equivalent rate-based marking function. In this paper, we restrict our attention to rate-based marking functions.

We consider the following form of marking functions:

$$p(z, C) = \begin{cases} 0 & \text{if } z \leq \underline{m} \\ \bar{p}(z, C) & \text{if } \underline{m} < z < \bar{m} \\ 1 & \text{if } z \geq \bar{m} \end{cases} \quad (3)$$

where $\underline{m} \in [0, C)$, $\bar{m} \in (0, \infty)$, and $\underline{m} < \bar{m}$. $\bar{p}(z, C)$ is assumed to satisfy the following condition.

Assumption 2.1: We assume that $\bar{p}(z, C)$ is a increasing, Lipschitz continuous, differentiable function in the interval $[\underline{m}, \bar{m}]$ with range $[0, 1]$, that satisfies $\bar{p}(z, C) = \bar{p}(z/C, 1)$. Assumption 2.1 states that the fraction of packets marked depends only on the ratio of the total arrival rate and the link capacity, which

TABLE I
EXAMPLES OF MARKING FUNCTIONS

Type	M	R	L	E	V
$\hat{p}(z, C)$	$(\frac{z}{C})^\alpha$	$\frac{\alpha z}{C - (1-\alpha)z}$	$\alpha(\frac{z}{C} - \eta)$	$1 - e^{-\frac{\alpha}{C}z}$	$(\frac{z-\alpha C}{z})^+$
\underline{m}, \bar{m}	$0, C$	$0, C$	$C\eta, C(1/\alpha + \eta)$	$0, \infty$	$\alpha C, \infty$

is satisfied by typical marking functions such as those in Table I (see [17], [25] for more details)³

In Table I, Type **M** has the interpretation of the queue length exceeding α in an M/M/1 queue with arrival rate z [25]. Type **R** can be used as a rate based model for REM (Random Exponential Marking [6]) for a suitable choice of α [26]. Type **L** is a simple linear marking function. Type **E** is a rate based exponential marking. Finally, type **V** has the interpretation of the fraction of fluid lost when the arrival rate exceeds a certain level, called the “virtual capacity” [27].

Then, from Assumption 2.1, we present the individual flow dynamics at time i by

$$w = x_k[i]p \left(\sum_{j=1}^n x_j[i] + \sum_{j=1}^n y_j[i], nC \right) = x_k[i]p \left(\frac{1}{n} \sum_{j=1}^n x_j[i] + \frac{1}{n} \sum_{j=1}^n y_j[i], C \right). \quad (4)$$

By summing over the flow index k , we then have

$$w = x[i]p(x[i] + y[i]) \quad (5)$$

where $x[i]$ and $y[i]$ are the average arrivals (over flows) at time i , i.e., $x[i] = (1/n) \sum_{j=1}^n x_j[i]$, and $y[i] = (1/n) \sum_{j=1}^n y_j[i]$.

Similarly, with weighted proportional fair controller,

$$x[i+1] - x[i] = \kappa(w - x[i-d]p(x[i-d] + y[i-d])). \quad (6)$$

C. Elasticity of Marking Function: Warping

In this section, we describe how to parameterize the elasticity of marking functions by adopting “warped” marking functions. A warped marking function has a parameter (denoted by β), which determines the elasticity of the marking functions. *The family of warped marking functions enables us to alter the elasticity of the marking function without altering the steady-state utilization.* Prior to describing warping, we first make the following additional assumption on the marking functions considered in this paper.

Assumption 2.2: $1/p(z, C)$ is convex in z over (\underline{m}, ∞) .

The typical marking functions in Table I satisfy Assumption 2.2.

Given any marking function $p(z, C)$ satisfying Assumption 2.1 and 2.2, we construct a family of marking functions parameterized by the warping constant β , $\beta \in (0, \infty)$, defined by:

$$p_\beta(z, C) \triangleq p(\beta z - (\beta - 1)z^*, C)$$

³For notational simplicity, we will omit the second parameter C throughout this paper unless explicitly needed.

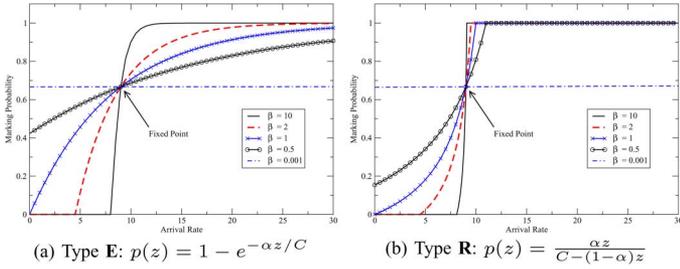


Fig. 4. Examples of warped marking functions: $C = 10$, $w = 3$, and $z^* = 9$.

where z^* is the equilibrium rate at the router with capacity C . Equivalently, $\gamma = z^*/C$ corresponds to the equilibrium utilization at the router, for some fixed value of $\gamma \in (0, 1)$. Note that

$$p_\beta(z^*, C) = p(z^*, C)$$

which ensures that for each fixed nominal marking function $p(z)$ ⁴ and the corresponding family of warped marking functions, $\{p_\beta(z)\}$, $\beta \in (0, \infty)$, the steady-state utilization of the system is independent of β . Then, for $z > z^*$, we have $p_\beta(z) > p(z)$, if $\beta > 1$, and $p_\beta(z) < p(z)$ if $0 < \beta < 1$.

In other words, $\{p_\beta(z)\}$ corresponds to a family of marking functions whose elasticity is varying with respect to the nominal marking function $p(z)$. If $\beta > 1$, $p_\beta(z)$ is *more elastic*, and if $\beta < 1$, $p_\beta(z)$ is *less elastic* from Definition 1.1. We can easily check that for each $\beta > 0$, $p_\beta(z)$ also satisfies Assumption 2.1 and 2.2. Fig. 4 shows the warped marking functions of Type **E** and Type **R** for different values of β .

D. Problem Statement

Our objective is to study the effect of the marking elasticity (using the warped marking functions) on the queue overflow probability (QOP) requirement of uncontrolled flows (i.e., the bound on probability that the queue length exceeds a fixed threshold). It is clear that the QOP for uncontrolled real-time flows will be the “smallest” if such flows are always given strict priority access at the routers (i.e., priority scheduling at the router with priority for uncontrolled real-time flows). We will later use priority scheduling as a reference model to assess the performance of FIFO scheduling (used in Section II-C to study scheduling-elasticity trade-off). With priority scheduling, we assume that two separate queues are used to store data from the controlled and uncontrolled flows, respectively.

We denote the sum of arrivals of n uncontrolled and n controlled flows over the time interval $[i, j]$ by $Y^n[i, j] = \sum_{k=1}^n \sum_{s=i}^{j-1} Y_k[s]$ and $X^n[i, j] = \sum_{k=1}^n \sum_{s=i}^{j-1} X_k[s]$, respectively⁵. We let $Z^n[i, j] = Y^n[i, j] + X^n[i, j]$, to denote the total sum of controlled and uncontrolled arrivals over the same time interval $[i, j]$ in the n -th system.

We consider a discrete time framework, where we suppose that the current time is 0, and the arrival process starts at time

⁴When $\beta = 1$, we use $p_\beta(\cdot, C)$ and $p(\cdot, C)$ interchangeably throughout this paper. Further, for simplicity, we omit the second parameter C in $p(z, C)$ and $p_\beta(z, C)$, unless explicitly needed.

⁵Thus, $X_k[s]$ denotes the random variable corresponding to the number of arrivals from the k th controlled flow at time s , and a similar definition holds for $Y_k[s]$. Finally, we use upper-case letters and lower-case letters to denote random variables and deterministic quantities, respectively.

$-\infty$. Thus, at the current time, the system is in steady state. We denote the queue length at time 0 with FIFO and priority schedulers by Q_0^P (for a queue of uncontrolled real-time flows) and Q_0^F , respectively.

The steady-state behavior of the congestion controllers (i.e., routers accessed by a mixture of controlled and uncontrolled flows) has been studied under fluid models [17], [28], and stability condition has been established [4], [17]. However, our focus here is the transient behavior which leads to queue overflow, and the impacts of marking elasticity on the QOP of real-time flows. Thus, in this paper, our objective is to compute the QOP as a function of the time-scale of the transient phenomenon as well as the marking function elasticity. We assume that for a fixed finite T_I , the system is stable before $-T_I$, i.e., $x[i] = x^*$ and $y[i] = y^*$ for $\forall i < -T_I$; and thus, the queue overflow probability is a function of T_I , the marking function, and scheduling policy. By the queue stability assumption (i.e., $x^* + y^* < C$), the queue length at time $i \in [-\infty, -T_I]$ is 0. Therefore, it suffices to consider the arrival processes only at the time interval over $[-T_I, 0)$.

For a fixed T_I , consider a following non-negative *scaled* (deterministic) arrival vector over the interval $[-T_I, 0)$: $\vec{v}[-T_I, 0) = (v[-T_I], v[-T_I + 1], \dots, v[-1])$. Then, from Loynes’s formula on the queue length process, the queue length at time 0 corresponding to an arrival vector $\vec{v}[-T_I, 0)$ can be defined as:

$$Q(\vec{v}[-T_I, 0)) \triangleq \left[\sup_{0 < T \leq T_I} \left(\sum_{i=-T}^{-1} v[i] - CT \right) \right]^+. \quad (7)$$

Thus, the queue overflows probabilities of priority and FIFO queueing are given by:

$$\begin{aligned} \Pr(Q_0^P \geq nB) &= \Pr \left[Q \left(\frac{1}{n} Y^n[-T_I, 0) \right) \geq B \right] \\ &= \Pr \left[\sup_{0 < T \leq T_I} \left(\frac{1}{n} Y^n[-T, 0) - CT \right) \geq B \right] \end{aligned} \quad (8)$$

$$\begin{aligned} \Pr(Q_0^F \geq nB) &= \Pr \left[Q \left(\frac{1}{n} Z^n[-T_I, 0) \right) \geq B \right] \\ &= \Pr \left[\sup_{0 < T \leq T_I} \left(\frac{1}{n} Z^n[-T, 0) - CT \right) \geq B \right]. \end{aligned} \quad (9)$$

In the large n regime, we derive asymptotic expressions for the queue overflow probabilities using large deviation techniques, which requires the application of the *contraction principle* [29]. Intuitively, the contraction principle facilitates computation of queue overflow probability by considering the queue length at time 0 as a function of the uncontrolled arrival process. For applicability of contraction principle, the queue length at time 0 should be continuous with respect to the uncontrolled arrival process, which we will prove later.

With the instant adaptation, we define the queue length at time 0 for the uncontrolled arrival $\vec{y}[-T_I, 0)$ by:

$$\tilde{Q}(\vec{y}[-T_I, 0)) \triangleq \left[\sup_{0 < T \leq T_I} \left(\sum_{i=-T}^{-1} (x[i] + y[i]) - CT \right) \right]^+ \quad (10)$$

where $\bar{x}[-T_I, 0)$ is determined by (5). Similarly, we use the notation \tilde{Q} to refer to the queue length at time 0 with the weighted proportional fair controller, i.e., $\bar{x}[-T_I, 0)$ is determined by (6).

In Section III, with the instant adaptation, we first prove that $\tilde{Q} : \mathcal{R}^{T_I} \mapsto \mathcal{R}$ is continuous under a suitable topology (which allows the application of a LDP (large deviation principle) and the contraction principle) in Theorem 3.1. Next, by computing Q_0^F as a function of marking elasticity, we will investigate and quantify stability-elasticity and scheduling-elasticity trade-off. In Section IV, we extend the analysis to the weighted proportional fair controller (using the queue length function $\hat{Q} : \mathcal{R}^{T_I} \mapsto \mathcal{R}$).

III. INSTANT ADAPTATION CONTROLLER

A. Continuity of Queue Length and Queue Overflow Probability

We first prove the continuity of \tilde{Q} (with respect to the uncontrolled arrival process) that enables the application of the large deviation results and the contraction principle.

We present a useful lemma, which is used to derive the queue overflow probability and to prove that the queue length at time 0 is continuous with respect to the uncontrolled arrival process. With w and $p(\cdot)$ in (5), let us define a function $h : [\underline{z}, \infty) \mapsto \mathcal{R}^+ \cup \{0\}$, where

$$h(z) \triangleq z - \frac{w}{p(z)}, \quad zp(\underline{z}) = w.$$

Since $p(\underline{z}) > 0$, we have $\underline{z} > \underline{m}$ from (3) on $p(z)$. Thus, $h(z)$ is defined over $[\underline{z}, \infty)$. Also, from the definition of $h(z)$, $h(\underline{z}) = 0$ and $h(z) > 0$ when $z > \underline{z}$. Intuitively, z is the total average arrival rate over flows (i.e., the sum of average uncontrolled and controlled arrivals over flows) at time i , and $h(z)$ is the corresponding uncontrolled arrival rate at time i (see (5)).

Lemma 3.1: Suppose that we have a marking function of the form (3) satisfying Assumption 2.1 and 2.2. Then, h is invertible and concave. Further, by defining $g \triangleq h^{-1}$, g is a convex function satisfying

$$(g(y[i]) - y[i])p(g(y[i])) = w.$$

Proof: Observe that $g(u)$ is the average total sum of uncontrolled and controlled arrivals with respect to the uncontrolled arrival u . Note that since $h(\cdot)$ is continuous, $g(\cdot)$ is also continuous. The formal proof is presented in Appendix. ■

Definition 3.1 (Uniform Norm): We define the uniform norm [30], [36] for a vector $\bar{v}[-T_I, 0)$ as follows:

$$\|\bar{v}[-T_I, 0)\|_u \triangleq \sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=-1} v[i]}{T} \right|.$$

We now prove the following result:

Theorem 3.1: The queue length function (at time 0) $\tilde{Q} : \mathcal{R}^{T_I} \mapsto \mathcal{R}$ is continuous with respect to the uncontrolled arrival process $\bar{y}[-T_I, 0) \in \mathcal{R}^{T_I}$ in the topology endowed with uniform norm. Further, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr \left(Q \left(\frac{1}{n} Z^n[-T_I, 0) \right) \geq B \right) = -I_F(B) \quad (11)$$

where

$$I_F(B) \triangleq \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) : \tilde{Q}(\bar{y}[-T, 0)) \geq B} \mathbf{I}(\bar{y}[-T, 0)) \quad (12)$$

$$\mathbf{I}(\bar{y}[-T, 0)) \triangleq \sup_{\bar{\theta} \in \mathcal{R}^T} \left(\bar{\theta} \cdot \bar{y}[-T, 0) - \log E e^{\bar{\theta} \cdot \bar{Y}[-T, 0)} \right).$$

Note that “ \cdot ” is the inner product of two vectors.

The proof is presented in Appendix.

B. Computation of Bounds on the Rate Function

This section focuses on computation of lower and upper bound on $I_F(B)$, leading to upper and lower bound on asymptotic queue overflow probability, respectively (see (11)). First, we add an additional assumption that the uncontrolled flows are independent and identically distributed *over time* for computational simplicity. The computation of $I_F(B)$ for non-i.i.d arrivals is left as future work. This i.i.d assumption ensures [29] that for any fixed T , we have

$$\mathbf{I}(\bar{y}[-T, 0)) = \sum_{i=-T}^{-1} I(y[i]) \quad (13)$$

where $I(y) \triangleq \sup_{\theta} (y\theta - \log E(e^{\theta Y_1[-1]}))$, and $Y_1[-1]$ is the random variable denoting the number of arrivals from flow “1” at time slot “-1.”

From Theorem 3.1 and (13), the rate function is given by

$$I_F(B) = \inf_{0 < T \leq T_I} I_F^T(B) \quad (14)$$

where

$$I_F^T(B) = \inf_{\bar{y}[-T, 0) \in \mathcal{A}} \sum_{i=-T}^{-1} I(y[i])$$

$$\mathcal{A} = \{\bar{a}[-T, 0) : \tilde{Q}(\bar{a}[-T, 0)) \geq B\}.$$

Then, we have the following result on the upper and lower bound on $I_F(B)$.

Theorem 3.2 (Upper and Lower Bound):

$$\inf_{0 < T \leq T_I} TI \left(C + \frac{B}{T} - \frac{w}{T} \left(\frac{1}{p(B+C)} + \frac{(T-1)}{p(C)} \right) \right) \leq I_F(B) \leq \inf_{0 < T \leq T_I} TI \left(C + \frac{B}{T} - \frac{w}{p\left(\frac{C+B}{T}\right)} \right)$$

We first describe three useful lemmas. Lemmas 3.2, 3.3 and Lemma 3.4 are used for the upper bound and the lower bound in Theorem 3.2, respectively. The proofs of all these lemmas are presented in Appendix.

First, we show that we do not need to optimize over the entire trajectories in the space \mathcal{A} . Instead, it is enough to optimize over the trajectories over $\hat{\mathcal{A}} \subset \mathcal{A}$, defined by

$$\hat{\mathcal{A}} = \left\{ \bar{a}[-T, 0) : \sum_{i=-T}^{-1} g(a[i]) \geq B + CT, a[i] \geq \hat{a} \right. \\ \left. \forall i \in [-T, -1], g(\hat{a}) = C \right\}. \quad (15)$$

Recall that $g(u)$ is the average total sum of uncontrolled and controlled arrivals with respect to the uncontrolled arrival u .

Lemma 3.2: With $\hat{\mathcal{A}}$ defined as (15),

$$\inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \mathcal{A}} \sum_{i=-T}^{-1} I(y[i]) \\ = \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \hat{\mathcal{A}}} \sum_{i=-T}^{-1} I(y[i]).$$

Next, let us define the space \mathcal{B} given by:

$$\mathcal{B} = \left\{ \bar{a}[-T, 0) : \frac{1}{T} \sum_{i=-T}^{-1} a[i] \geq g^{-1} \left(\frac{B}{T} + C \right) \right. \\ \left. a[i] \geq \hat{a}, \forall i \in [-T, -1], g(\hat{a}) = C \right\}. \quad (16)$$

By showing that the space $\mathcal{B} \subset \hat{\mathcal{A}}$, we derive an upper bound of $I_F^T(B)$, described in the following lemma.

Lemma 3.3: $\mathcal{B} \subset \hat{\mathcal{A}}$. Thus,

$$\inf_{\bar{y}[-T, 0) \in \hat{\mathcal{A}}} \sum_{i=-T}^{-1} I(y[i]) \leq \inf_{\bar{y}[-T, 0) \in \mathcal{B}} \sum_{i=-T}^{-1} I(y[i]).$$

Lemma 3.4: Suppose that we have the following optimization problem:

$$\inf_{z[-T, 0) \in \mathcal{C}} \frac{1}{T} \sum_{i=-T}^{-1} f(z[i]) \\ \mathcal{C} = \left\{ \bar{a}[-T, 0) : \sum_{i=-T}^{-1} a[i] \geq B + CT, a[i] \geq C \right\}. \quad (17)$$

Also, suppose that f is increasing and concave. Then, the vector $\bar{z}^*[-T, 0) = (B + C, C, \dots, C)$ is an optimizer.

Proof of Theorem 3.2:

(i) *Upper bound:* From Lemma 3.2 and Lemma 3.3,

$$I_F(B) = \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \hat{\mathcal{A}}} \sum_{i=-T}^{-1} I(y[i]) \\ \leq \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \mathcal{B}} I \left(\frac{1}{T} \sum_{i=-T}^{-1} y[i] \right).$$

Since $I(x)$ is increasing and convex, for $x \geq \hat{a}$, and from the definition of \mathcal{B} (which is a convex set), we have the optimum when $y[i] = g^{-1}(B/T + C)$ for $-T \leq i < 0$. Then, the result immediately follows.

(ii) *Lower bound:* Since $I(x)$ is monotone-increasing over $[\hat{a}, \infty)$ and from Jensen's inequality, we have

$$I_F(B) = \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \mathcal{A}} \sum_{i=-T}^{-1} I(y[i]) \\ \geq \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0) \in \mathcal{A}} TI \left(\frac{1}{T} \sum_{i=-T}^{-1} y[i] \right) \\ \geq \inf_{0 < T \leq T_I} TI \left(\inf_{\bar{y}[-T, 0) \in \mathcal{A}} \frac{1}{T} \sum_{i=-T}^{-1} y[i] \right).$$

Let $z[i] = g(y[i])$ (i.e., $h(z[i]) = y[i]$ from Lemma 3.1). Then, we have the following problem transformation:

$$\inf_{\bar{y}[-T, 0) \in \mathcal{A}} \frac{1}{T} \sum_{i=-T}^{-1} y[i] \Leftrightarrow \inf_{z[-T, 0) \in \mathcal{C}} \frac{1}{T} \sum_{i=-T}^{-1} h(z[i])$$

where \mathcal{C} is defined as (17). Then, from Lemma 3.4 and the definition of h , the result follows. ■

C. Stability-Elasticity Trade-Off

Using the lower and upper bounds on the rate function (i.e., $I_F(B)$) derived in the previous section with the instant adaptation source controller, we study the effect of elasticity of marking functions on the stability (for controlled flows) and QOP (for uncontrolled flows), and their trade-off.

Consider a fixed nominal marking function $p(z)$, and the corresponding family of marking functions $\{p_\beta(z), \beta \in (0, \infty)\}$. Recall that with respect to a nominal marking function, $\beta > 1$ corresponds to a more elastic marking function, and $\beta < 1$ corresponds to a less elastic marking function.

We adopt the following procedure to study the effect of the marking elasticity on the system stability of the closed-loop controlled sources and on the QOP of the open-loop uncontrolled real-time sources:

- (i) For a fixed β , we compute the *smallest* QOP that can be achieved for the uncontrolled real-time sources by assuming that the controlled source adapts and back-off instantly in response to congestion, i.e., the instant adaptation controller is used for QOP analysis. Note that with the instant adaptation scheme, the upper bound on the rate function from Theorem 3.2 provides a lower bound on the QOP. In other words, for a fixed value of β and the corresponding marking function $p_\beta(z)$, we can get *no less QOP* than that given by Theorem 3.2.
- (ii) Using known local stability results for a weighed proportional fair controller from [4], for a fixed β , we compute the *maximum delay* that can be tolerated before the controlled sources go into local (and hence, global) instability.

While (i) and (ii) use different controllers (instant adaptation and proportional fair controller, respectively), our objective here is to illustrate the effect of the marking elasticity (not specific controller mechanisms) on the QOP for the uncontrolled real-time source. Thus, (i) corresponds to the "best-case" scenario for the QOP of the uncontrolled real-time sources (due to the fact that the controlled flows in (i) back-off instantly).

With any other controller, there will be a lag associated with the back-off of controlled flows, thus resulting in larger QOP for uncontrolled real-time flows than that with the instant adaptation controller (see Section IV for the corresponding results when the weighted proportional fair is used for both (i) QOP analysis and (ii) stability analysis, respectively).

To discuss the stability analysis in (ii), we use the local stability condition for a weighted proportional fair controller. For each marking function $p_\beta(z)$, we determine the *maximum round-trip propagation delay* d that the system can tolerate before going into local instability (and thus, global instability). This is given by [4]:

$$\kappa(p_\beta(z^*) + z^*p'_\beta(z^*)) < \sin\left(\frac{\pi}{2(2d+1)}\right). \quad (18)$$

Further, by definition of $p_\beta(x)$, we have

$$p'_\beta(z^*) = f'_\beta(z^*)p'(z^*) = \beta p'(z^*).$$

Thus, for each value of β , the stability condition (18) reduces to

$$\kappa(p(z^*) + \beta z^*p'(z^*)) < \sin\left(\frac{\pi}{2(2d+1)}\right). \quad (19)$$

The trade-off between the QOP for the real-time flows and the stability for the controlled flows is parameterized by β , the elasticity of the marking function. The more elastic the marking function is, the worse is the stability behavior (as β becomes larger in (19)). On the other hand, increasing β *improves* the QOP behavior for the real-time uncontrolled flows. This can be explained by applying the β -elastic marking function $p_\beta(\cdot)$ to the lower and upper bound of the rate function in Theorem 3.2, and by observing that $p_\beta(B+C)$, $p_\beta(C)$, and $p_\beta(C+B/T)$ are increasing with respect to β . For this reason, we refer to this study as *stability-elasticity trade-off*.

From a network provider's viewpoint, the operator would choose a nominal operating load (equilibrium point) and a desired QOP bound. Now, by choosing the appropriate warping, the operator can choose the desired trade-off between QOP for real-time flows and performance of controlled flows (i.e., the amount of variability in the controlled flow rate due to delay). To study this trade-off, we fix an nominal load z^* , and study various marking functions with and without warping. In Section V, we illustrate the results for various values of the nominal load z^* , marking functions, and network environments.

We comment that this loss of local stability phenomenon has been studied in [31], [32]. However, our work differs in that we explicitly consider the QOP of real-time flows and quantifies its impact on the local stability region of elastic flows using a mixture of control theoretic tools and large deviation principle.

D. Scheduling-Elasticity Trade-Off

In this section, we derive bounds on the link capacities needed with priority and FIFO scheduling to support a QOP requirement for the real-time uncontrolled flows, which stipulates that the QOP should not exceed some ϵ .

It is clear that the link capacity required for supporting a fixed QOP with priority scheduling is the smallest (over all scheduling policies), since absolute priority is given to these real-time flows, i.e., the controlled flows do not affect the queue dynamics

for the uncontrolled flows (see Fig. 1). Thus, the required capacity with priority scheduling does not depend on the marking elasticity. On the other hand, with FIFO scheduling, the behavior of uncontrolled flows and controlled flows are coupled together, and thus, the required link capacity for supporting the given QOP is a *function of marking elasticity*, and it will be larger than that with priority scheduling.

With this observation, our objective is to quantitatively study the “scheduling-gain” (see Definition 3.2) of priority scheduling over FIFO scheduling, as a function of marking elasticity, and we show that this gain could be significantly reduced by increasing marking elasticity.

To do so, we adopt the following approach:

- (i) For a fixed marking elasticity, we first determine the per-flow link capacities needed with priority and FIFO scheduling (denoted by C_P and $C_F(\beta)$, respectively) for supporting a fixed QOP ϵ .
- (ii) Using the analysis in (i) we define the following “normalized scheduling-gain” of priority scheduling over FIFO scheduling.

Definition 3.2 (Scheduling Gain):

$$\Delta_C(\beta) \triangleq \frac{C_F(\beta) - C_P}{C_P}. \quad (20)$$

Intuitively, $\Delta_C(\beta)$ quantifies the trade-off between the penalty of choosing “sub-optimal” scheduling algorithm (i.e., FIFO scheduling) in terms of QOP-guarantee and the elasticity of the marking function (i.e., β). We will investigate the behavior of $\Delta_C(\beta)$ as the marking elasticity parameter, β , changes.

To discuss the analysis with priority scheduling in (i) (i.e., computation of the bound on C_P), we use well-known large deviation results [8]. In the large number of flows regime, defining $\delta \triangleq -(1/n) \log \epsilon$, the system would allow the QOP less than ϵ if $I_P(B) \leq \delta$, where

$$I_P(B) = \inf_{0 < T \leq T_r} TI \left(C_P + \frac{B}{T} \right).$$

Observe that this also leads to an *effective bandwidth* characterization for a single server queue accessed by only uncontrolled flows [7], [8], where a sufficient condition for $I_P(B) \geq \delta$ (equivalently, to support a QOP of at most $e^{-n\delta}$) is

$$\frac{\Lambda\left(\frac{\delta}{B}\right)}{\frac{\delta}{B}} \leq C_P \quad (21)$$

where $\Lambda(\theta)$ is the log-moment generating function of a random uncontrolled arrival at a particular time-slot, i.e., $\Lambda(\theta) = \log E[e^{\theta Y_1} - 1]$.

Next, for the analysis with FIFO scheduling (i.e., computation of $C_F(\beta)$) in (i), we have the following proposition:

Proposition 3.1: With FIFO scheduling and with a fixed marking elasticity β , a sufficient condition for $I_F(B) \geq \delta$ (i.e., to support the QOP ϵ) is:

$$\frac{\Lambda\left(\frac{\delta}{B}\right)}{\frac{\delta}{B}} \leq C_F(\beta) - \frac{w}{p_\beta(C_F(\beta))}. \quad (22)$$

The proof is presented in Appendix.

We now apply (21) and Proposition 3.1 to study the effect of marking elasticity on the scheduling gain of priority scheduling (discussed in (ii)). Defining $g(\beta) \triangleq C_F(\beta) - (w/p_\beta(C_F(\beta)))$,

we observe that $g(\beta)$ is increasing in β , since $p_\beta(z)$ is increasing in β for $z \geq z^*$, and $C_F(\beta) \geq z^*$. This implies that we need progressively less $C_F(\beta)$ (to support a given QOP for the uncontrolled real-time flows and queue stability) with increasing values of β .

Further, from the definition of warped marking function discussed in Section II-C, for a fixed $z > z^*$, we have $p_\beta(z) \rightarrow p(z^*)$, as $\beta \rightarrow 0$, and $p_\beta(z) \rightarrow 1$, as $\beta \rightarrow \infty$. Thus, from (22), we have

$$\begin{aligned} C_F(\beta) &\xrightarrow{\beta \rightarrow 0} \frac{\Lambda\left(\frac{\delta}{B}\right)}{\frac{\delta}{B}} + \frac{w}{p(z^*)} \\ C_F(\beta) &\xrightarrow{\beta \rightarrow \infty} \frac{\Lambda\left(\frac{\delta}{B}\right)}{\frac{\delta}{B}} + w. \end{aligned} \quad (23)$$

Note that $(\Lambda(\delta/B)/\delta/B)$ is the minimum required per-flow link capacity with priority scheduling for supporting the given QOP ϵ (which follows from (21)). Then, the results in (23) imply that by increasing the marking elasticity, the scheduling gain can be significantly reduced. We illustrate this by means of an example.

1) *Example 3.1:* Consider a single bottleneck network accessed by 100 uncontrolled and controlled flows (i.e., $n = 100$). Also, let the equilibrium rate for a controlled flow and the mean rate of an uncontrolled flow to be 640 kbps and 120 kbps, respectively. Assuming that each packet is of a fixed size with 1000 bytes, the equilibrium rate of a controlled flow is $x^* = 80$ pkts/sec, and the mean rate of uncontrolled flow is $y^* = 15$ pkts/sec. We model each uncontrolled arrival process by a ON-OFF process with 'ON' probability of 0.1, and 'ON' rate of 150 pkts/sec. The queue buffer size is 800 kbytes, which corresponds to 100 pkts (i.e., $B = 1$). We set the equilibrium marking probability to be 0.03 (i.e., $p(z^*) = p(95) = 0.03$). The QOP bound for uncontrolled flows is set to be $\epsilon = 10^{-6}$ (i.e., $\delta = -(1/100) \log 10^{-6} = 0.1382$). Then, the required per-flow link capacity with priority scheduling is $(\Lambda(0.1382/1)/0.1382/1) = 133.3$ pkts/sec. From the equilibrium analysis of the congestion controller, we have $w = x^* \times p(z^*) = 80 \times 0.03 = 2.4$. Then, with FIFO scheduling, to support the same QOP ϵ , we need the per-flow link capacity $133.3 + (2.4/0.03) = 213.3$ pkts/sec for $\beta = 0$. On the other hand, for $\beta \rightarrow \infty$, we need only $133.3 + 2.4 = 135.7$ pkts/sec.

In terms of queue management and implementation, FIFO scheduling is much simpler than priority scheduling. On the other hand, it is clear that priority scheduling provides less QOP to the uncontrolled real-time flows. However, the results in this section imply that the scheduling gain due to priority scheduling may not become significant in the large scale networks by adjusting the marking elasticity. In other words, the difference in the required capacities with FIFO and priority scheduling for a fixed QOP can be significantly reduced by increasing the marking elasticity. For this reason, we refer to this as *scheduling-elasticity trade-off*.

This trade-off is graphically illustrated in Fig. 6 in Section V, where we numerically present the trade-off.

IV. WEIGHTED PROPORTIONAL FAIR CONTROLLER

In this section, we consider the weighted proportional fair controller described in (1). As in Section III, we again assume that the uncontrolled arrivals are independent over time. However, the analysis with the weighted proportional fair controller is more complicated than that with the instant adaptation due to temporal coupling of the arrival process introduced by the dynamics of each congestion controller, i.e., the total arrivals to the router are not independent over time, *even if* the uncontrolled arrivals are. We begin this section with the proof of continuity of queue length (at time 0) function (i.e., $\hat{Q}(\cdot)$) with respect to the uncontrolled arrival process in the system with controlled flows governed by weighted proportional fair controller, and using this, derive an upper-bound on the rate function to study the stability-elasticity trade-off.

A. Continuity of Queue Length and Queue Overflow Probability

We have the following theorem to show that the queue length at time 0 is continuous with weighted proportional fair controller and round-trip propagation delay d .

Theorem 4.1: The queue length function (at time 0) $\hat{Q} : \mathcal{R}^{T_I} \mapsto \mathcal{R}$ is continuous with respect to the uncontrolled arrival process $\vec{y}[-T_I, 0)$ in the topology endowed with uniform norm. Thus, we have

$$\lim_{n \rightarrow \infty} \frac{1}{n} \log \Pr \left(Q \left(\frac{1}{n} Z^n[-T_I, 0) \right) \geq B \right) = -I_{\hat{F}}(B)$$

where

$$I_{\hat{F}}(B) \triangleq \inf_{0 < T \leq T_I} \inf_{\vec{y}[-T, 0) : \hat{Q}(\vec{y}[-T, 0)) \geq B} \mathbf{I}(\vec{y}[-T, 0)). \quad (24)$$

The proof is presented in Appendix.

B. Stability-Elasticity Trade-Off

In the previous section, we have shown the continuity of queue length with respect to the uncontrolled arrival process. This enables us to apply the contraction principle as in Section III-B. In this section, we briefly study the stability-elasticity trade-off by deriving an upper bound on the rate function in the system with a "one-step" delay. Controllers with delay and the scheduling-elasticity trade-off will be studied in Section V using simulations.

In our analysis, we assume that the controller gain, κ , is small enough to prevent the transmission rate of a controlled flow from becoming negative (note that with a one-step delay, if κ is large, the controller will be unstable irrespective of the marking elasticity).

Theorem 4.2 (Upper Bound): Suppose that $\kappa p(C + B) \leq 1$. Then, we have

$$\begin{aligned} I_{\hat{F}}(B) \leq & \inf_{0 < T \leq T_I} \sum_{i=-T}^{-1} I \left(C + \frac{B}{T} \right. \\ & \left. - \left(x^* - \frac{w}{\hat{p}(T)} \right) (1 - \kappa \hat{p}(T))^{i+T+1} - \frac{w}{\hat{p}(T)} \right), \end{aligned}$$

where $\hat{p}(T) = p(C + B/T)$.

Proof: First, from Theorem 4.1, we have

$$I_{\hat{F}}(B) = \inf_{0 < T \leq T_I} \inf_{\bar{y}[-T, 0]: \hat{Q}(\bar{y}[-T, 0]) \geq B} \sum_{i=-T}^{-1} I(y[i]). \quad (25)$$

Next, since for a fixed T , $\{\bar{y}[-T, 0] \mid x[i] + y[i] = C + B/T, i \in [-T, -0)\} \subset \{\bar{y}[-T, 0] \mid \hat{Q}(\bar{y}[-T, 0]) \geq B\}$, we have

$$I_{\hat{F}}(B) \leq \inf_{0 < T \leq T_I} \inf_{\substack{\bar{y}[-T, 0]: \forall j \in [-T, 0), \\ y[j] + x[j] = C + B/T}} \sum_{i=-T}^{-1} I(y[i]). \quad (26)$$

Then, from (6), the fact that $y[i] + x[i] = C + B/T, \forall i \in [-T, 0)$ implies that $\bar{x}[-T, 0)$ is deterministically fixed by the following form of difference equation:

$$x[i + 1] - x[i] = \kappa(w - x[i]\hat{p}(T)), \quad i \in [-T, 0) \quad (27)$$

where the initial condition $x[-T - 1] = x^*$. Then, by solving the difference (27), and by using $x[i] + y[i] = C + B/T, \forall i \in [-T, 0)$ we have

$$\begin{aligned} x[i] &= \left(x^* - \frac{w}{\hat{p}(T)}\right) (1 - \kappa\hat{p}(T))^{i+T+1} + \frac{w}{\hat{p}(T)} \\ y[i] &= C + \frac{B}{T} - \left(x^* - \frac{w}{\hat{p}(T)}\right) (1 - \kappa\hat{p}(T))^{i+T+1} - \frac{w}{\hat{p}(T)}. \end{aligned}$$

The assumption that $\kappa p(C + B) < 1$ ensures that $0 \leq 1 - \kappa\hat{p}(T) \leq 1$, which eliminates the case when $x[i]$ could be negative and has to be set to 0.

Finally, from (26), $I_{\hat{F}}(B)$ is upper-bounded by:

$$\inf_{0 < T \leq T_I} \sum_{i=-T}^{-1} I\left(C + \frac{B}{T} - \left(x^* - \frac{w}{\hat{p}(T)}\right) \times (1 - \kappa\hat{p}(T))^{i+T+1} - \frac{w}{\hat{p}(T)}\right)$$

To understand the stability-elasticity trade-off, let us apply the marking elasticity parameter β to the upper bound on the rate function. Then, we observe that for a fixed T ,

$$R(\beta) \triangleq C + \frac{B}{T} - \left(x^* - \frac{w}{\hat{p}_\beta(T)}\right) (1 - \kappa\hat{p}_\beta(T))^{i+T+1} - \frac{w}{\hat{p}_\beta(T)}$$

is increasing in β , since $\hat{p}_\beta(T)$ is increasing with respect to β , for a fixed T . Thus, from Theorem 4.2, we observe that as we have an increasingly elastic marking function at the router, we get the larger rate function (and thus, a less QOP). However, from (19), increasing β causes the maximum allowable delay for stability to decrease. The above result points toward a stability-elasticity trade-off (however, note that the stability-elasticity trade-off is based on an upper bound as opposed to the exact values) with the weighted proportional fairness, in a qualitatively similar form to that observed in (19) with the instant adaptation controller.

V. NUMERICAL RESULTS AND SIMULATION

A. Instant Adaptation

1) *Stability-Elasticity Trade-Off:* We first illustrate the stability-elasticity trade-off with the instant adaptation algorithm in

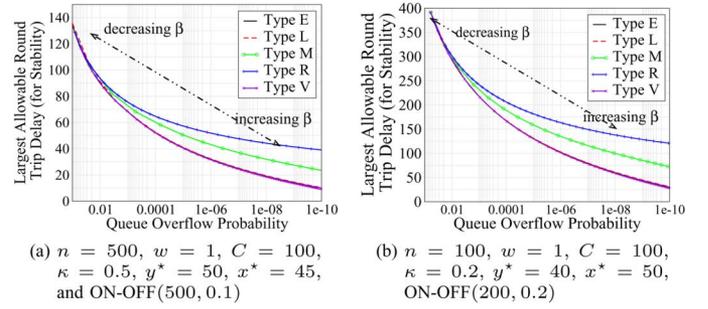


Fig. 5. Stability-elasticity trade-off with instant adaptation. ON-OFF (two state Markov) uncontrolled arrivals, where ON-OFF (a, p) means that the ON rate is a with probability p and the OFF rate is 0 with probability $1 - p$.

Fig. 5. For each marking function in Table I, we plot the trade-off between the largest allowable round-trip delay for stability and (the lower bound on) QOP (computed by the upper bound of the rate function in Theorem 3.2) as a parametric plot of β .

Fig. 5 clearly illustrates the trade-off between the QOP of real-time flows and stability of controlled flows for different values of system parameters such as the total number, the relative volume of each type of flows, and the burstiness of real-time flows.

The results indicate that some marking functions may be “uniformly” better than others. In particular, among the marking functions that we have compared, our bounds indicate that a rate based version of REM (Type **R**) [6], [14] seems to provide the largest local-stability region for *any* given QOP requirement.

An intuitive explanation for this is the following: From Theorem 3.2, it is clear that the QOP for the uncontrolled real-time flows with FIFO scheduling depends on the marking function behavior for arrival rates exceeding the per-flow link capacity C . In particular, the value of the rate function is proportional to the marking function value for arrival rates exceeding C (i.e., $I_F(\cdot) \propto p(z), z > C$). From Fig. 4, we observe that among the example marking functions considered in Table I (which are normalized to have the same fixed point properties), the rate-based REM marking function seems to have the maximum slope for $z > z^*$, which in turn implies a larger marking function value (as all the example marking functions have the same $p(z^*)$). To analytically construct uniformly optimal marking functions is an interesting problem for future research.

2) *Scheduling-Elasticity Trade-Off:* To illustrate the scheduling-elasticity trade-off, Fig. 6 shows the required per-flow link capacity with FIFO and priority scheduling to support a given QOP, for different burstiness of real-time flows with the type **V** marking function. As discussed in Section III-D, the required link capacities for any QOP requirement should satisfy *queue stability condition*, i.e., the link capacities should be large enough such that the queue length is always finite. Thus, Fig. 6 plots the maximum over the capacities governed by the QOP and queue stability condition.

We observe that for a small value β , the difference between the capacities with FIFO and priority queueing is large for all QOP values. This is due to the fact that the controlled flows back-off sluggishly. On the other hand, for more elastic marking

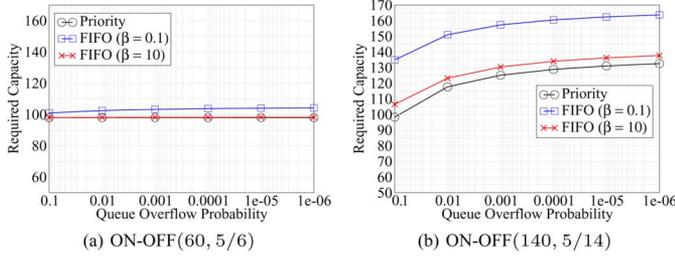


Fig. 6. Scheduling-elasticity trade-off: $w = 5$, $\kappa = 0.2$, $n = 100$, $x^* = 48$, and $y^* = 50$.

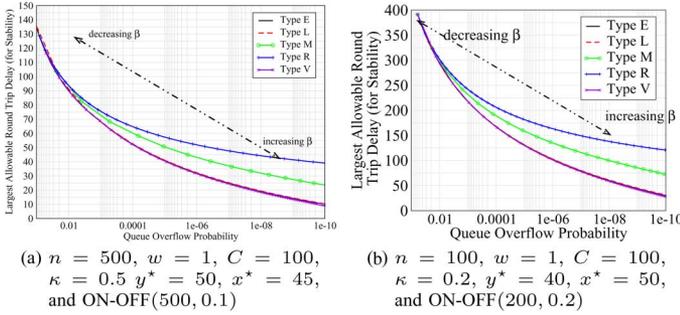


Fig. 7. Stability-elasticity trade-off with weighted proportional fair controller.

functions, the required capacities with both scheduling algorithms are very close.

For a less bursty uncontrolled arrivals [Fig. 6(a)], in priority scheduling, the queue stability condition (i.e., $C_P < z^* = 98$) dominates the QOP condition [i.e., (22)], while for a more bursty arrivals [Fig. 6(b)], the QOP condition is stronger than the queue stability condition. In both cases, we observe that the required capacity with FIFO can be significantly decreased (to a value that is almost the same as that with priority scheduling) by increasing the marking elasticity.

B. Weighted Proportional Fair Controller

1) *Numerical Examples:* Now, we study the trade-offs with the weighted proportional fair controller using numerical examples as well as simulation results. First, as numerical examples, Fig. 7 illustrates the stability-elasticity trade-off with the weighted proportional fair controller. We observe that analogous to the instant adaptation case, as the QOP requirement becomes more strict (i.e., less QOP), the stability region is reduced. Also, importantly, we still observe that the type **R** marking function appears “uniformly better” than other marking functions we have considered.

2) *Simulation Results:* To study the trade-offs with the weighted proportional fair controller in a more practical scenario, we use the *ns-2* [33] packet simulator to validate our analysis. The network topology used in the simulation is same as that discussed in the analysis (see Fig. 3). The number of uncontrolled and controlled flows (i.e., n) are set to be 100 throughout all the simulation results. The per-flow link capacity of the bottle-neck link is 100 pkts/sec (i.e., total capacity is 100×100 pkts/sec). The buffer size of the bottle-neck link is 100 pkts. We use the fixed size of packets (1000 bytes). Uncontrolled flows are modeled by discrete ON-OFF processes, where the burst-time and the idle-time are set to be 100 msec and

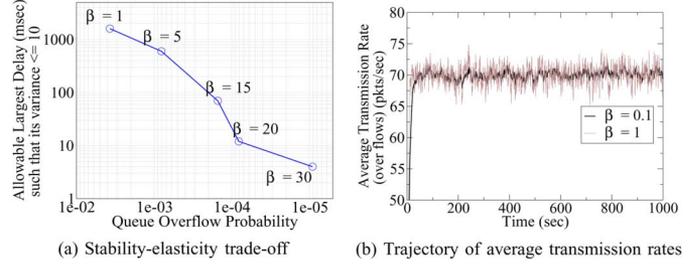


Fig. 8. Stability-elasticity trade-off: The trajectories in (b) show that for two values of β , the average remains the same. However, there is a trade-off between QOP and delay as observed in (a).

900 msec. The transmission rate in ON period is appropriately set such that the specified mean uncontrolled rate is achieved in different simulations. The parameter κ and w with the weighted proportional fair controller are set to be 1 and 5, respectively.

First, Fig. 8 shows the stability-elasticity trade-off for the Type **V** marking function. The equilibrium rate (x^*) of a controlled arrival process, the mean rate (y^*) of an uncontrolled arrival process are set to be 70 pkts/sec and 20 pkts/sec, respectively. Fig. 8(b) shows the trajectories of average transmission rates with two different marking elasticities, from which we clearly observe that as we have larger β (more elastic marking function), the trajectory becomes more fluctuating, thus resulting in less stable behavior. For the plots in Fig. 8(a), we denote a flow to be stable if the transmission rate variance is less than 10 pkts/sec. In Fig. 8(a), we observe that the stability region becomes smaller, as we increase marking elasticity. However, increasing marking elasticity implies lower QOP with larger marking elasticity.

Next, Table II shows the scheduling-elasticity trade-off with the weighted proportional fair controller. In this simulation, (x^* , y^*) are set to be (70,25) pkts/sec, respectively. To see the scheduling-elasticity effect, for a fixed QOP (10^{-5}), we experimentally determine the required per-flow link capacity to satisfy the given QOP. With priority scheduling, it is observed that the queue stability condition is dominant; thus, the required capacities for different β values are all equal to 95 pkts/sec. With FIFO scheduling, we observe that the difference in the required link capacity between $\beta = 0.001$ and $\beta = 20$ is about 10 pkts/sec with all three types of marking functions. Further, the extra capacity that we need with FIFO scheduling to support the given QOP is shown to be about 15 pkts/sec. This validates the analytical result that the required link capacity decreases with increasing values of β .

VI. CONCLUSION

In this paper, we first have quantified the trade-off between stability for controlled flows and guarantee of queue overflow probability (QOP) for uncontrolled real-time flows as a function of marking elasticity. The results indicate that some marking functions may be “uniformly” better than others, where, in particular, among the marking functions that we have compared, our bounds indicate that a rate based version of REM seems to provide the largest local-stability region for *any* given QOP requirement.

Next, we have compared the capacity required at a router with only FIFO scheduling versus a router with priority scheduling

TABLE II
SCHEDULING-ELASTICITY TRADE-OFF: REQUIRED
PER-FLOW LINK CAPACITY FOR 10^{-5} QOP

Marking Functions→	Priority	FIFO		
		Type E	Type L	Type V
$\beta = 0.001$	95	119.2	119.5	118.7
$\beta = 1$	95	118.4	117.2	116.8
$\beta = 20$	95	110	110	108.5

(priority given to the real-time flows) for supporting a given QOP requirement. We have quantified the “scheduling-gain” of priority scheduling over FIFO scheduling, as a function of marking elasticity. We show that this scheduling gain decreases with more elastic marking functions.

In this paper, we have considered only the weighted proportional fair source controller. However, the dominant source controller in the Internet is TCP (Transmission Control Protocol). We conclude this paper by discussing possible extensions of our results to TCP controllers. It is widely known that various types of the source controllers addressed in literature can be modeled in the utility maximization framework [34], where the TCP controller corresponds to the $-1/x$ utility function (with an appropriate setting of κ and w). In order to extend our study to TCP controllers, we need to average n TCP flows to get a one-dimensional averaged equation (analogous to (6) based on which a large deviations problem can be formulated. This seems possible for a collection of TCP flows (modeled using the $-1/x$ utility function) in steady-state and with a common round trip delay. Using the results in [25], in steady state, the controlled flows (in this case, TCP flows) can be represented only by an averaged equation. This will form the initial condition (i.e., at time $-T_I$), and we can then study transient behavior using a large deviations formulation as in this paper. We plan to study this in the future.

APPENDIX

Proof of Lemma 3.1: First, we prove $h(z)$ is strictly increasing (thus, invertible) by showing that $h(z + \delta) - h(z) \geq \delta$, $\forall z \geq \underline{z}, \forall \delta > 0$. Since $p(z)$ is non-zero increasing when $z > \underline{z} > \underline{m}$, we have

$$\begin{aligned} h(z + \delta) - h(z) &= \delta + w \left(\frac{1}{p(z)} - \frac{1}{p(z + \delta)} \right) \\ &\geq \delta. \end{aligned}$$

In addition, $h(z)$ is the sum of two concave functions since $-w/p(z)$ is concave from Assumption 2.2. Thus, $h(z)$ is a concave function.

Let us define $g \triangleq h^{-1} : \mathcal{R}^+ \cup \{0\} \rightarrow [\underline{z}, \infty)$. Then,

$$z[i] = g \left(z[i] - \frac{w}{p(z[i])} \right) = g(y[i]).$$

Applying this to (5), we have $(g(y[i]) - y[i])p(g(y[i])) = w$. The convexity of g follows immediately from the concavity of h . ■

Proof of Theorem 3.1: We note that the authors in [30], [35] have proved that $Q(\bar{z}[-T_I, 0])$ is continuous with respect to $\bar{z}[-T_I, 0] \in \mathcal{R}^{T_I}$ in the topology endowed with uniform norm.

Note that $\bar{z}[-T_I, 0] (= \bar{y}[-T_I, 0] + \bar{x}[-T_I, 0])$ is the function (denoted by H) of $\bar{y}[-T_I, 0]$, since $\bar{x}[-T_I, 0]$ is determined by $\bar{y}[-T_I, 0]$ in the FIFO scheduling. Thus, from the definition of Q , it suffices to show that the function $H : \mathcal{R}^{T_I} \mapsto \mathcal{R}^{T_I}$ is continuous (in the topology endowed with uniform norm) to prove that \hat{Q} is continuous with respect to the uncontrolled arrival $\bar{y}[-T_I, 0]$. We will prove that for any given $\epsilon > 0$, and for two uncontrolled arrival processes, $\bar{y}[-T_I, 0]$ and $\bar{y}_\epsilon[-T_I, 0]$, such that $\|\bar{y}[-T_I, 0] - \bar{y}_\epsilon[-T_I, 0]\|_u < \epsilon$, there exists a function f such that $\|\bar{z}[-T_I, 0] - \bar{z}_\epsilon[-T_I, 0]\|_u < f(\epsilon)$, where $f(\epsilon) \xrightarrow{\epsilon \rightarrow 0} 0$.

First, note that at each time i , $z[i] (= x[i] + y[i])$ depends on only $y[i]$. Then, for a given $\epsilon > 0$, from the assumption that $\|\bar{y}[-T_I, 0] - \bar{y}_\epsilon[-T_I, 0]\|_u < \epsilon$ and the definition of uniform norm, we have

$$\sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=1} (y[i] - y_\epsilon[i])}{T} \right| < \epsilon.$$

Then, the finiteness of T_I , there exists a finite non-negative constant K' (which is a function of T_I), such that $|y[i] - y_\epsilon[i]| < K'\epsilon$, for all $i \in [-T_I, \dots, -1]$.

Now, observe that

$$\begin{aligned} \|\bar{z}[-T_I, 0] - \bar{z}_\epsilon[-T_I, 0]\|_u &= \sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=-1} (z[i] - z_\epsilon[i])}{T} \right| \\ &\leq \sum_{i=-T_I}^{i=-1} |z[i] - z_\epsilon[i]| = \sum_{i=-T_I}^{i=-1} |g(y[i]) - g(y_\epsilon[i])| \\ &< T_I \max_{i=-T_I, \dots, -1} |g(y[i]) - g(y_\epsilon[i])|. \end{aligned} \quad (28)$$

Since $g(y)$ is continuous with respect to y from Lemma 3.1, (28) is arbitrarily small for an arbitrary small ϵ . This completes the proof. Then, the resulting rate function $I_F(B)$ immediately follows from [30, Theorem 9]. ■

Proof of Lemma 3.2: Let us choose $\bar{y}^*[-T^*, 0] \in \mathcal{A}$. Suppose that $g(y^*[k]) < C$, for a $k \in [-T^*, 0)$. It suffices to show that we can find a new trajectory $\hat{y}[-\hat{T}, 0)$ of no larger cost (rate function) than $\bar{y}^*[-T^*, 0)$ such that $g(\hat{y}[i]) \geq C$, $\forall i$, $-\hat{T} \leq i < 0$.

Define a new trajectory $\hat{y}[-\hat{T}, 0)$, $\hat{T} = T^* - 1$ as follows.

$$\hat{y}[i] = \begin{cases} y^*[i] & \text{if } i \in [k + 1, 0) \\ y^*[i - 1] & \text{if } i \in [-T^* + 1, k]. \end{cases}$$

Then, we have constructed a new trajectory with cost as follows: $\sum_{-\hat{T}}^{-1} I(\hat{y}[i]) \leq \sum_{-T^*}^{-1} I(y^*[i])$, and $Q_F(\hat{y}[-\hat{T}, 0)) \geq B$. In addition, if $g(y^*[k]) < C$, for $k_1, k_2, \dots \in [-T^*, 0)$, we can inductively remove k_1, k_2, \dots , and finally construct a new trajectory $\hat{y}[-\hat{T}, 0)$ in $\hat{\mathcal{A}}$, with no larger cost. ■

Proof of Lemma 3.3: It suffices to show that for any $\bar{a}[-T, 0) \in \mathcal{B}$, $\bar{a}[-T, 0) \in \hat{\mathcal{A}}$. We first have

$$\frac{1}{T} \sum_{i=-T}^{-1} a[i] \geq g^{-1} \left(\frac{B}{T} + C \right) \Rightarrow g \left(\frac{1}{T} \sum_{i=-T}^{-1} a[i] \right) \geq \frac{B}{T} + C.$$

Also, since g is convex, we have

$$g\left(\frac{1}{T} \sum_{i=-T}^{-1} a[i]\right) \leq \frac{1}{T} \sum_{i=-T}^{-1} g(a[i]).$$

from Jensen's inequality. This completes the proof. \blacksquare

Proof of Lemma 3.4: First, it is clear that we get the optimum when $\sum_{i=-T}^{-1} z[i] = B + CT$, since f is increasing. Second, we claim that $C \leq z[i] \leq B + C, i \in J$, where $J = \{-T, -T + 1, \dots, -1\}$. Suppose that $z[j] > B + C$ for some $j \in J$. Then, We should have $z[k] < C$ for some $k \in J$, which contradicts the given condition.

Now, let $z^*[-T] = B + C, z^*[-T + 1] = z^*[-T + 2] = \dots = z^*[-1] = C$. Since $C \leq z[i] \leq B + C, i \in J$, we can represent $z[i]$ by the following linear combination of $z^*[i]$.

$$z[-i] = (1 - \alpha_i)C + \alpha_i(B + C), \quad \forall i = 1, \dots, T, \quad (29)$$

where $\sum_{i=-T}^{-1} \alpha_i = 1$ and $0 \leq \alpha_i \leq 1$ for $i \in J$. From (29) and concavity of f ,

$$\begin{aligned} \frac{1}{T} \sum_{i=-T}^{-1} f(z[i]) &= \frac{1}{T} \sum_{i=-T}^{-1} f((1 - \alpha_i)C + \alpha_i(B + C)) \\ &\geq f(C) \sum_{i=-T}^{-1} (1 - \alpha_i) + f(B + C) \sum_{i=-T}^{-1} \alpha_i \\ &= (T - 1)f(C) + f(B + C) = \frac{1}{T} \sum_{i=-T}^{-1} f(z^*[i]). \end{aligned}$$

This completes the proof. \blacksquare

Proof of Proposition 3.1: From Theorem 3.2, and the increasing property of $p_\beta(z)$ with respect to z , we have the following lower bound on $I_F(B)$:

$$\begin{aligned} I_F(B) &\geq \inf_{0 < T \leq T_I} TI \left(\frac{B}{T} + C_F(\beta) \right. \\ &\quad \left. - \frac{w}{T} \left(\frac{1}{p_\beta(B + C_F(\beta))} + \frac{(T - 1)}{p_\beta(C_F(\beta))} \right) \right) \\ &\geq \inf_{0 < T \leq T_I} TI \left(\frac{B}{T} + C_F(\beta) - \frac{w}{p_\beta(C_F(\beta))} \right). \quad (30) \end{aligned}$$

We let

$$\tilde{C}(\beta) = \frac{C_F(\beta) - w}{p_\beta(C_F(\beta))} \quad \tilde{\delta} = \frac{\delta}{B}.$$

Suppose that (22) is true, i.e., $(\Lambda(\tilde{\delta})/\tilde{\delta}) \leq \tilde{C}(\beta)$. Then, we have

$$\begin{aligned} \tilde{C}(\beta)\tilde{\delta} - \Lambda(\tilde{\delta}) &\geq 0 \Rightarrow (\tilde{C}(\beta) + x)\tilde{\delta} - \Lambda(\tilde{\delta}) \geq \tilde{\delta}x, \quad \forall x > 0 \\ &\Rightarrow \inf_{x \in \{B, B/2, \dots\}} \frac{I(\tilde{C}(\beta) + x)}{x} \geq \tilde{\delta}. \quad (31) \end{aligned}$$

Thus, $I_F(B) \geq \delta$ from (30) and (31). \blacksquare

Proof of Theorem 4.1: Similar to Theorem 3.1, $\bar{z}[-T_I, 0)$ is a function of $\bar{y}[-T_I, 0)$ (denoted by \hat{H}), and it suffices to show that the function $\hat{H} : \mathcal{R}^{T_I} \mapsto \mathcal{R}^{T_I}$ is continuous in the topology endowed with uniform norm. To do so, we prove that for any two uncontrolled arrival processes, $\bar{y}[-T_I, 0)$ and

$\bar{y}_\epsilon[-T_I, 0)$, such that $\|\bar{y}[-T_I, 0) - \bar{y}_\epsilon[-T_I, 0)\|_u < \epsilon$, there exists a function f such that $\|\bar{z}[-T_I, 0) - \bar{z}_\epsilon[-T_I, 0)\|_u < f(\epsilon)$, where $f(\epsilon) \xrightarrow{\epsilon \rightarrow 0} 0$.

First, similar to the proof of Theorem 3.1, $\|\bar{y}[-T_I, 0) - \bar{y}_\epsilon[-T_I, 0)\|_u < \epsilon$ implies that $\exists K' > 0$, s.t.,

$$|y[i] - y_\epsilon[i]| < K'\epsilon, \quad \forall i \in [-T_I, 0). \quad (32)$$

We now embed the discrete time trajectory of $\bar{y}[-T_I, 0)$ and $\bar{x}[-T_I, 0)$ in "continuous time," i.e., for $t \in \mathcal{Z}$, we let $x(t) = x[t], y(t) = y[t]$ and use a straight-line approximation to interpolate between the times $t = i, i \in \mathcal{Z}$. Thus, we have the following differential equation to represent the controlled flows dynamics in the continuous time:

$$\begin{aligned} \dot{x}(t) &= \kappa(w - x([t - d])p(x([t - d]) + y([t - d]))) \\ \dot{x}_\epsilon(t) &= \kappa(w - x_\epsilon([t - d])p(x_\epsilon([t - d]) + y_\epsilon([t - d]))) \end{aligned} \quad (33)$$

where $\sup_{-T_I \leq t < 0} \delta(t) < K'\epsilon$.

Then, from [17, Lemma 3.2] and (33) (we can verify that the assumptions needed for [17, Lemma 3.2] to hold are satisfied in a manner identical to that in [17, Theorem 4.1]), we have

$$\begin{aligned} \sup_{-T_I \leq i < 0} |x[i] - x_\epsilon[i]| &\leq \sup_{-T_I \leq t < 0} |x(t) - x_\epsilon(t)| \\ &\leq 2LT_I e^{LT_I} K'\epsilon. \end{aligned} \quad (34)$$

Then, from (32) and (34), we have

$$\begin{aligned} \|\bar{z}[-T_I, 0) - \bar{z}_\epsilon[-T_I, 0)\|_u &= \sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=-1} (z[i] - z_\epsilon[i])}{T} \right| \\ &\leq \sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=-1} (x[i] - x_\epsilon[i])}{T} \right| + \sup_{0 < T \leq T_I} \left| \frac{\sum_{i=-T}^{i=-1} (y[i] - y_\epsilon[i])}{T} \right| \\ &< 2LT_I e^{LT_I} K'\epsilon + K'\epsilon. \end{aligned}$$

Recall that L is the Lipschitz parameter of the marking function. By letting $f(\epsilon) = 2LT_I e^{LT_I} K'\epsilon + K'\epsilon$, the result follows, since ϵ is arbitrary. \blacksquare

REFERENCES

- [1] L. Ying, G. Dullerud, and R. Srikant, "Global stability of internet congestion controllers with heterogeneous delays," in *Proc. Amer. Control Conf.*, Jun. 2004, [CD ROM].
- [2] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Trans. Networking*, vol. 1, no. 4, pp. 397-413, Aug. 1993.
- [3] L. Massouli and J. Roberts, "Bandwidth sharing: Objectives and algorithms," *IEEE/ACM Trans. Networking*, vol. 10, no. 3, pp. 320-328, Jun. 2002.
- [4] R. Johari and D. Tan, "End-to-end congestion control for the Internet: Delays and stability," *IEEE/ACM Trans. Networking*, vol. 9, no. 6, pp. 818-832, Dec. 2001.
- [5] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: Shadow prices, proportional fairness and stability," *J. Oper. Res. Soc.*, vol. 49, pp. 237-252, 1998.

- [6] S. H. Low and D. E. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Trans. Networking*, vol. 7, no. 6, pp. 861–875, Dec. 1999.
- [7] G. de Veciana and J. Walrand, "Effective bandwidths: Call admission, traffic policing and filtering for ATM networks," *Queueing Syst. Theory Appl.*, vol. 20, pp. 37–59, 1995.
- [8] D. D. Botvich and N. G. Duffield, "Large deviations, economies of scale, and the shape of the loss curve in large multiplexers," *Queueing Syst.*, vol. 20, pp. 293–320, 1995.
- [9] C. Courcoubetis and R. Weber, "Buffer overflow asymptotics for a switch handling many traffic sources," *J. Appl. Probab.*, vol. 33, pp. 886–903, 1996.
- [10] G. Kesidis, J. Walrand, and C.-S. Chang, "Effective bandwidths for multiclass Markov fluids and other ATM sources," *IEEE/ACM Trans. Networking*, vol. 1, pp. 424–428, 1993.
- [11] N. Likhanov and R. Mazumdar, "Cell loss asymptotics for buffers fed with a large number of independent stationary sources," *J. Appl. Probab.*, vol. 36, pp. 86–96, 1999.
- [12] J. Cao and K. Ramanan, "A poisson limit for buffer overflow probabilities," in *Proc. IEEE INFOCOM*, New York, NY, Jun. 2002, pp. 994–1003.
- [13] D. Daley and D. Vere-Jones, *An Introduction to the Theory of Point Processes*. New York: Springer-Verlag, 1988.
- [14] Y. Yi, S. Deb, and S. Shakkottai, "Time-scale decomposition and rate-based marking," *IEEE/ACM Trans. Networking*, vol. 14, no. 5, pp. 938–950, Oct. 2006.
- [15] S. Deb and R. Srikant, "Rate-based versus queue-based models of congestion control," in *Proc. ACM SIGMETRICS*, 2004, [CD ROM].
- [16] G. Raina and D. Wischik, "Buffer sizes for large multiplexers: TCP queueing theory and instability analysis," in *Proc. Next Gen. Internet Networks*, 2005, [CD ROM].
- [17] S. Shakkottai and R. Srikant, "Mean FDE models for Internet congestion control under a many-flows regime," *IEEE Trans. Inform. Theory*, vol. 50, no. 6, pp. 1050–1072, Jun. 2004.
- [18] N. Antunes, C. Fricker, F. Guillemin, and P. Robert, "Integration of streaming services and tcp data transmission in the internet," *Performance Evaluation*, vol. 62, no. 1–4, pp. 263–277, 2005.
- [19] F. Delcoigne, A. Proutiere, and G. Regnie, "Modeling integration of streaming and data traffic," *Performance Evaluation*, vol. 55, no. 3–4, pp. 185–209, 2004.
- [20] P. Key, L. Massoulié, A. Bain, and F. Kelly, "Fair internet traffic integration: Network flow models and analysis," *Annales des Telecommun.*, vol. 59, pp. 1338–1352, 2004.
- [21] P. Key and L. Massoulié, "Fluid models of integrated traffic and multipath routing," *Queueing Syst. Theory Appl.*, vol. 53, no. 1–2, pp. 85–98, 2006.
- [22] P. Tinnakornsrisuphap and R. J. La, "Characterization of queue fluctuations in probabilistic AQM mechanisms," in *Proc. ACM Sigmetrics*, 2004, [CD ROM].
- [23] F. Kelly, "Mathematical modelling of the Internet," in *Mathematics Unlimited—2001 and Beyond*, B. Engquist and W. Schmid, Eds. Berlin, Germany: Springer-Verlag, 2001, pp. 685–702.
- [24] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue algorithm for active queue management," in *Proc. ACM SIGCOMM*, San Diego, CA, Aug. 2001, pp. 123–134.
- [25] S. Deb, S. Shakkottai, and R. Srikant, "Asymptotic behavior of internet congestion controllers in a many-flows regime," *Math. Oper. Res.*, vol. 30, no. 2, pp. 420–440, 2005.
- [26] F. P. Kelly, "Models for a self-managed Internet," *Philosophical Trans. Royal Soc.*, vol. A358, pp. 2335–2348, 2000.
- [27] S. Kunniyur and R. Srikant, "End-to-end congestion control: Utility functions, random losses and ECN marks," in *Proc. IEEE INFOCOM*, Tel Aviv, Israel, Mar. 2000, vol. 3, pp. 1323–1332.
- [28] V. Misra, W.-B. Gong, and D. Towsley, "Fluid-based analysis of a network of AQM routers supporting TCP flows with an application to red," in *Proc. ACM SIGCOMM*, 2000, [CD ROM].
- [29] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis*. New York: Chapman and Hall, 1995.
- [30] D. Wischik, "Sample path large deviations for queues with many inputs," in *Proc. Annal. Appl. Probab.*, 2001, vol. 11, no. 2, pp. 379–404.
- [31] L. Massoulié, "Stability of distributed congestion control with heterogeneous feedback delays," *IEEE Trans. Automat. Control*, vol. 47, no. 6, pp. 895–902, Jun. 2002.
- [32] G. Raina, "Local bifurcation analysis of some dual congestion control algorithms," *IEEE Trans. Automat. Control*, vol. 50, no. 8, pp. 1135–1146, Aug. 2005.
- [33] , The Network Simulator—ns-2 2009 [Online]. Available: <http://www.isi.edu/nsnam/ns/>
- [34] R. Srikant, *The Mathematics of Internet Congestion Control*. Basel, Switzerland: Birkhauser, 2004.
- [35] A. J. Ganesh and N. O'connell, "A large deviation principle with queueing applications," *Stochastics: Int. J. Prob. Stoch. Processes*, vol. 73, no. 1, pp. 25–35, 2002.
- [36] N. O'Connell, "A large deviations principle with queueing applications," Hewlett Packard Labs, Bristol, U.K., Tech. Rep. HPL-BRIMS-9705, BRIMS, 1997.



Yung Yi (S'03–M'07) received the B.S. and M.S. degrees from the School of Computer Science and Engineering, Seoul National University, Seoul, Korea, in 1997 and 1999, respectively, and the Ph.D. degree from the Department of Electrical and Computer Engineering, University of Texas at Austin, in 2006.

From 2006 to 2008, he was a Post-Doctoral Research Associate in the Department of Electrical Engineering, Princeton University, Princeton, NJ. Now, he is an Assistant Professor with the Department of Electrical Engineering and Computer Science, Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea. His current research interests include the design and analysis of computer networking and communication systems, especially congestion control, scheduling, and interference management, with applications in wireless ad hoc networks, broadband access networks, and future Internet evolution.



Sanjay Shakkottai (M'01) received the Ph.D. degree from the Electrical and Computer Engineering Department, University of Illinois at Urbana-Champaign in 2002.

He is with The University of Texas at Austin, where he is currently an Associate Professor in the Department of Electrical and Computer Engineering, and the Engineering Foundation endowed Faculty Fellow. His research interests include wireless and sensor networks, stochastic processes and queueing theory.

Dr. Shakkottai received the NSF CAREER Award in 2004.