

On the Elasticity of Marking Functions: Scheduling, Stability, and Quality-of-Service in the Internet

Yung Yi and Sanjay Shakkottai¹
 Department of Electrical and Computer Engineering
 University of Texas at Austin
 e-mail: {yi, shakkott}@ece.utexas.edu

Abstract — Much of the research on Internet 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, the Internet is composed of a mixture of both (controlled) elastic flows and (uncontrolled) real-time flows. Uncontrolled real-time flows do not react to network congestion as well as they require a certain level of QoS guarantees. In this paper, we study the effects of marking elasticity (which characterizes how quickly the marking level changes during transients) on the QoS for uncontrolled real-time flows, at a router accessed by both uncontrolled real-time and controlled 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 QoS guarantee for uncontrolled real-time flows as a function of marking elasticity. 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 seems to provide the largest local-stability region for *any* given QoS requirement.

Next, as a function of the marking function elasticity, we quantify the excess capacity required at the router with FIFO scheduling that results in the same queue overflow probability if priority scheduling was used instead. We show that the difference in the required capacities with FIFO and priority queueing decreases with more elastic marking functions. In other words, the gains due to scheduling decreases with increasing marking elasticity.

I. INTRODUCTION

There has been extensive research on the modeling and analysis of the controlled elastic flows in the Internet by adopting differential equation based models of source controllers and AQM (Active Queue Management) 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 adopting control theoretic tools [1, 2, 3, 4, 5, 6, 7].

However, the Internet 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 and requires tight QoS (Quality of Service) guarantees. From a network control and management point of view, real-time sources are admitted into the network only if there are sufficient resources to satisfy their QoS requirements. On the other hand, non-real-time sources are always admitted into the network with the understanding that the resources in the network would be allocated to them on a best-effort basis (i.e., real-time sources are given higher priority and whatever bandwidth is



Figure 1: Priority and FIFO Queueing Disciplines

left unused by the real-time sources is allocated to the non-real-time sources). One of the proposed architectures for providing differentiated QoS in the Internet is “Differentiated Service” model (DiffServ) [8], where users can belong to one of a small number of classes, and QoS (such as delay, loss ratio, and throughput) for a user’s data flow will be class-dependent. To implement such a service, routers in the Internet treat (schedule) packets from various classes in a differentiated manner depending on the class QoS specifications by adopting “priority” based scheduling algorithms (see Figure 1).

On the other hand, it seems reasonable to believe that by appropriately designing an AQM mechanism (marking function) at intermediate routers, we can potentially provide the required QoS to the uncontrolled real-time flows *without any differentiation* at the routers. The intuition is the following: an “aggressive” marking function will mark a larger number of controlled flow packets (for instance, those controlled by TCP) when a burst of packets arrive. This will cause the controlled flows to back-off, thus potentially decreasing the delay or packet loss probability experienced by real-time flows with their link utilization being sustained equivalent. In this paper, we study the trade-off between packet marking [9] for controlled flows and the effect of this marking on the QoS of uncontrolled real-time flows.

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). Over such a network, the behavior of real-time and controlled flows are coupled together, and the QoS experienced by real-time flows will be affected by the behavior of controlled flows due to their flows sharing a common link. For example, a large burst of QoS-sensitive packets from a real-time flow could potentially encounter a significant delay or loss at the router due to the controlled flows sharing the link.

With this setup, we first characterize the “aggressiveness” of a marking function by its *elasticity*.

Definition 1.1. Given any two marking functions $p_1(z)$ and $p_2(z)$, we say that $p_2(z)$ is more elastic than $p_1(z)$ if for any $z \geq z^*$, we have

$$\begin{aligned} p_1(z^*) &= p_2(z^*) \\ p_2(z) &\geq p_1(z), \end{aligned}$$

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 deviates from the equilibrium rate (see Figure 2).

¹Yi and Shakkottai are with the Wireless Networking and Communications Group. This research was partially supported by NSF Grants ACI-0305644, CNS-0325788, and CNS-0347400.

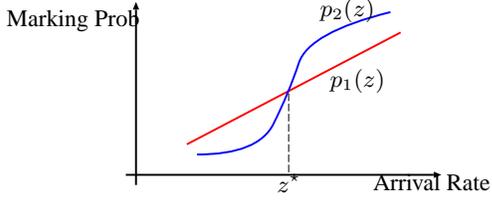


Figure 2: Elasticity of Marking Functions

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 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 [10, 2].

We consider queue overflow probability as the QoS parameter experienced by real-time flows. The problem of determining the queue overflow probability has been studied extensively for queues [11, 12, 13, 14, 15] in the context of “open-loop” flows (i.e., there are no controlled flows). This research has been done, primarily using a large deviation framework, which is known to be appropriate to analyze performance of the system with a large number of flows. In this paper, we focus on a router where the resource is shared by real-time flows as well as controlled flows (with FIFO queueing). Thus, the controlled sources react and adapt to the dynamics at the router, and affects the QoS of real-time flows. We derive the queue overflow probability by adapting large-deviation technique to such a shared system by controlled and real-time flows.

In this paper, we study and quantify the following two trade-offs related with marking elasticity: (i) stability-elasticity trade-off and (ii) scheduling-elasticity trade-off. First, stability-elasticity trade-off refers to the trade-off between QoS-provisioning for real-time flows and stability for controlled flows. The key trade-off we explore is the following: *the more elastic the marking function is, the better is the QoS experienced by real-time flows. However, this also leads to the negative property of less stability for controlled flows (stability-elasticity trade-off).*

Next, we consider the case where the real-time flows are given absolute priority over the controlled flows (priority scheduling). In this case, the delay or loss perceived by the real-time flows (QoS) will not be affected by the behavior of the controlled flows. However, by simply *increasing the capacity of the link and suitably adjusting marking elasticity (parameters of marking functions), it seems possible to give the same perceived QoS to the real-time flows in the network with only FIFO scheduling.* We call this *scheduling-elasticity trade-off* in this paper.

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 queue overflow probability at a router, where a single buffer is shared by controlled and real-time flows. Using these bounds, we quantify the trade-off between stability for controlled flows and QoS-guarantee for uncontrolled real-time flows as a function of the elasticity of the marking function (see Figure 4). 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 [7] seems to provide the largest local-stability region

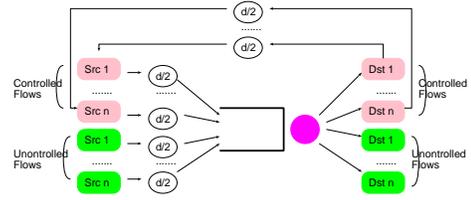


Figure 3: System Model

for *any* given QoS requirement.

- (ii) We next quantify the required excess capacity (the difference in the required link capacities with FIFO and priority queueing), as well as the marking function elasticity that results in the same queue overflow probability with FIFO scheduling and priority scheduling (see Figure 5). We show that the difference in the required capacities with FIFO and priority queueing *decreases* with stricter QoS requirements (queue overflow probability). In other words, the gains achieved by scheduling¹ at the router decrease with increasing QoS. This indicates that by appropriately choosing the marking function and by using only a FIFO queue at the router, we can satisfy the QoS requirements of real-time flows without much over-provisioning.

II. SYSTEM MODEL AND PROBLEM STATEMENT

II.A SYSTEM MODEL

Consider the system shown in Figure 3. We consider a *single* discrete time queue with two types of flows: (i) controlled flows and (ii) uncontrolled flows. We use the terminology “controlled flows” (also called elastic flows) to refer to flows of data traffic which react and adapt their transmission rates to feedback from the network. An example of such a flow is a TCP (Transmission Control Protocol) flow. “Uncontrolled flows” refer to data flows that do not react to network feedback. Examples of such flows include real-time video/audio, which usually require guarantee of real-time data transfer. The queue is fed by n independent identically distributed (over flows) stationary, ergodic uncontrolled flows and by n controlled flows determined by a congestion control algorithm described later. The buffer size is scaled with n , and the output capacity of the corresponding queue is suitably scaled with n so that the queue is stable. Thus, n^{th} system has a buffer of size nB , and a capacity of nC . For system stability, we assume that $x^* + y^* < C$, where x^* is the fixed point of a controlled flow and y^* is the expected value of an uncontrolled flow.

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* [16, 3].

The router is modeled by a *marking function* (see Section 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 fluid processes. We denote the fluid rates of individual flows by $\{x_k[i], k = 1, \dots, n\}$,

¹Any class based scheduling policy will result in poorer QoS to real-time flows than strict priority queueing. Thus, a priority queueing system provides an upper bound on the QoS given real-time flows with any class-based scheduling policy employed at the router.

Table 1: Examples of Marking Functions

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

where $x_k[i]$ denotes the number of arrivals² of controlled flow k at time i . In this paper, we consider *instant adaptation* algorithm, which is represented by:

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

In the instant adaptation algorithm, congestion controllers adapt to the fixed point of the difference equation of weighted proportional fair controller [17, 18, 10, 2] with no delay [10, 2]. In this scheme, as the rate of the uncontrolled flow varies with time, the corresponding equilibrium rate varies appropriately (as determined by the elastic properties of $p(\cdot)$), and the instant adaptation scheme tracks this variation of the equilibrium rate. This allows us to focus purely on the properties of the marking function.

II.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 . We consider the following form of marking functions

$$p(z, C) = \begin{cases} 0 & \text{if } 0 \leq z \leq \underline{m}, \\ \hat{p}(z, C) & \text{if } \underline{m} < z < \overline{m}, \\ 1 & \text{if } z \geq \overline{m}, \end{cases} \quad (2)$$

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

Assumption II.1. We assume that $\hat{p}(z, C)$ is a increasing, Lipschitz continuous, differentiable function with range $[0, 1]$, that satisfies $\hat{p}(z, C) = \hat{p}(z/C, 1)$.

Assumption II.1 says that the fraction of packets marked simply depends on the ratio of the total arrival rate and the link capacity, which is satisfied by typical marking functions such as those in Table 1 (see [17, 19] for more details)³.

Examples of marking functions $\hat{p}(z, C)$ we consider in this paper (satisfying Assumption II.1) are shown in Table 1: Type **M** has the interpretation of the queue length exceeding B in an M/M/1 queue with arrival rate z . Type **R** can be used as a rate based model for REM (Random Exponential Marking [7]) for a suitable choice of α [20]. Type **L** is a linear marking function, and models the simplified RED (Random Early Detection [9]). 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” [21].

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

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

Then, the individual flow dynamics at time i of instant adaptation can be represented as follows by using Assumption II.1 and summing over the flow index k ,

$$\begin{aligned} w &= x_k[i]p\left(\frac{1}{n}\sum_{j=1}^n x_j[i] + \frac{1}{n}\sum_{j=1}^n y_j[i], C\right) \\ &= x[i]p(x[i] + y[i], C), \end{aligned} \quad (3)$$

where $x[i]$ and $y[i]$ are the average arrivals (over flows) at time i .

II.C PROBLEM STATEMENT

A widely used QoS parameter (for the uncontrolled real-time flows) is the probability that the queue length exceeds some threshold. The queue overflow probability can be extended to analyze the delay of a typical packet [22]. It is clear that the QoS performance for uncontrolled flows will be the “best” if such flows are always given strict priority access at the routers (i.e., priority scheduling at the router). We will later use priority scheduling as a reference model to assess the performance of FIFO scheduling (used in Section D 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 consider a discrete time framework, and denote the queue length at time 0 with FIFO and priority schedulers by Q_0^F (for a queue of uncontrolled real-time flows) and Q_0^P , respectively. We consider the queueing process over the time interval $[-T_I, 0)$, where $T_I < \infty$. We assume that the system is stable (thus, queue length is 0) over the time interval $(-\infty, -T_I)$, meaning that the transmission rates of all controlled flows and all uncontrolled flows are x^* and y^* , respectively, where x^* is the fixed point of controlled flows and y^* is the expected value (average rate) of a uncontrolled flow.

In other words, in this paper we are interested in the transient behavior of the system. We assume that the system is in “steady-state” until time $-T_I$, and our objective is to compute the buffer overflow probability as a function of the time-scale of the transient phenomenon (i.e., T_I) as well as the marking function. However, we observe that the analysis in this paper holds even if T_I is not finite as we can show that with instant adaptation, the queueing process is continuous (see Theorem III.1) with respect to the arrivals even over arbitrarily larger interval of time.

We denote the sum of arrivals of n uncontrolled and n controlled flows over the time interval $[i, j)$ by $Y^n[i, j) = \sum_{k=1}^n Y_k[i, j)$ and $X^n[i, j) = \sum_{k=1}^n X_k[i, j)$, respectively⁴. 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.

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 Loyne’s formula on the queue length process, the queue size function corresponding to an arrival vector $\vec{v}[-T_I, 0)$ can be defined as:

$$Q(\vec{v}[-T_I, 0)) \triangleq \sup_{0 < T \leq T_I} \left(\sum_{i=-T}^{-1} v[i] - CT \right) \quad (4)$$

⁴Thus, $X_k[i, j)$ denotes the random variable corresponding to the number of arrivals from the k^{th} controlled flow over the time interval $[i, j)$, and a similar definition holds for $Y_k[i, j)$. In addition, we use $X_k[i]$ to denote $X_k[i, i + 1)$, and this notation is applied for other random arrivals. Finally, we use upper-case letters and lower-case letters to denote random variables and deterministic quantities, respectively.

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

$$\Pr(Q_0^P \geq nB) = \Pr\left(\sup_{0 < T \leq T_I} \left(\frac{1}{n} Y^n[-T, 0] - CT\right) \geq B\right) \quad (5)$$

$$\Pr(Q_0^F \geq nB) = \Pr\left(\sup_{0 < T \leq T_I} \left(\frac{1}{n} Z^n[-T, 0] - CT\right) \geq B\right) \quad (6)$$

II.D ELASTICITY OF MARKING FUNCTIONS: WARPING

In this section, we parameterize the elasticity of marking functions by adopting ‘‘warped’’ marking functions. *An warped marking function has a parameter (denoted by β), which determines the elasticity of the marking functions by shifting ($0 < \beta < 1$) and twisting ($\beta \geq 1$) the original marking functions.*

Prior to describing warping, first we make the following additional assumption on the marking function.

Assumption II.2. $1/p(z, C)$ is convex over (z_0, ∞) , where $z_0 = \sup\{x : p(x, C) = 0, x \geq 0\}$

The typical marking functions in Table 1 satisfy Assumption II.2.

Given any marking function $p(z)$ satisfying Assumption II.1 and II.2, we construct a family of marking functions $\{p_\beta(z)\}$, which are parameterized by β and are defined by $p_\beta(z) \triangleq p(f_\beta(z))$, where

$$f_\beta(z) = \begin{cases} \beta(z - \gamma) & \text{if } 0 < \beta < 1, \\ \gamma z^\beta & \text{if } \beta \geq 1. \end{cases}$$

For a given system (with a mixture of controlled and uncontrolled arrivals), let the equilibrium rate at the router be denoted by z^* . For each value of β , the parameter γ (in the definition of $f_\beta(z)$) is chosen such that at this equilibrium rate z^* , $f_\beta(z^*) = z^*$. This definition ensures that we have the equivalent steady-state marked volume of data at a router over the set of marking functions $\{p_\beta(\cdot)\}$ ($p_\beta(z^*) = p(z^*)$, $\forall \beta > 0$.) leading to *invariability of steady-state utilization of the system*. Further, for $z > z^*$, we have

$$\begin{aligned} p_\beta(z) &> p(z) && \text{if } \beta > 1, \\ p_\beta(z) &< p(z) && \text{if } 0 < \beta < 1. \end{aligned}$$

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 I.1. Note that the function $f_\beta(z)$ is constructed such that for each β , $p_\beta(z)$ satisfies Assumption II.1 and II.2.

III. STABILITY-ELASTICITY AND SCHEDULING-ELASTICITY TRADE-OFF

III.A QUEUE OVERFLOW PROBABILITY

From (5), (6), and (4), the queue overflow probability with priority and FIFO scheduling can be expressed as $\Pr(Q(\frac{1}{n} Y^n[-T_I, 0]) \geq B)$ and $\Pr(Q(\frac{1}{n} Z^n[-T_I, 0]) \geq B)$, respectively. In the large n regime, we can derive asymptotic expressions for the queue overflow probabilities using large deviation techniques. This requires the application of the Gartner-Ellis Theorem, as well as the contraction principle [23]. Applicability of contraction principle depends on the continuity of the queue size function with respect to the arrival process *from the uncontrolled flows* (denoted by \tilde{Q} in Theorem III.1). However, with FIFO scheduling, the arrival process to

the queue consists of the sum of arrivals from the controlled and uncontrolled flows. Thus, we need to prove: (i) the queueing process is continuous with respect to the total arrival process, and (ii) the controlled arrival process (determined by the dynamics of the congestion controller, the marking function, and the uncontrolled flows) is a continuous function of the uncontrolled arrival process. In [24], the author proved that $Q(\cdot)$ (the queue size function with only stochastic uncontrolled arrivals) is continuous (i.e., no controlled flows are present).

We now prove the continuity of the queue length at time 0 with respect to the uncontrolled flows $\tilde{\mathbf{y}}[-T_I, 0]$, with FIFO scheduling.

Theorem III.1. *With the instant adaptation algorithm, the queue size function $\tilde{Q} : \mathcal{R}_+^{T_I} \mapsto \mathcal{R}$ is continuous with respect to the uncontrolled arrival process $\tilde{\mathbf{y}}[-T_I, 0]$ in the topology endowed with supremum norm, where \tilde{Q} is defined as*

$$\begin{aligned} \tilde{Q}(\tilde{\mathbf{y}}[-T_I, 0]) &\triangleq Q(\tilde{\mathbf{y}}[-T_I, 0] + \tilde{\mathbf{x}}[-T_I, 0]) \\ &= \sup_{0 < T \leq T_I} \left(\sum_{-T}^{-1} (x[i] + y[i]) - CT \right), \end{aligned}$$

and $\tilde{\mathbf{x}}[-T_I, 0]$ is determined by (3) (i.e., is function of $\tilde{\mathbf{y}}[-T_I, 0]$). Thus, we have

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

where $I_F(B)$ is defined as

$$I_F(B) \triangleq \inf_{\tilde{\mathbf{y}}[-T_I, 0]} \left\{ \mathbf{I}(\tilde{\mathbf{y}}[-T_I, 0]) : \tilde{Q}(\tilde{\mathbf{y}}[-T_I, 0]) \geq B \right\}, \quad (7)$$

where $\mathbf{I}(\tilde{\mathbf{y}}[-T_I, 0])$ is the rate function of the vector $\tilde{\mathbf{y}}[-T_I, 0] \in \mathcal{R}_+^{T_I}$ [23].

Proof. The proof is presented in [25]. \square

III.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. First, we add an additional assumption that an uncontrolled flows are independent and identically distributed *over time* for simplicity. The computation of $I_F(B)$ for non-i.i.d arrivals is left as future work. This i.i.d assumption ensures [23] that for any fixed T , we have

$$\mathbf{I}(\tilde{\mathbf{y}}[-T, 0]) = \sum_{i=-T}^{-1} I(y[i]), \quad (8)$$

where $I(\cdot)$ is defined as

$$I(y) = \sup_{\theta} \left(y\theta - \log \mathbb{E}(e^{\theta Y_1[-1]}) \right),$$

and $Y_1[-1]$ is the random variable denoting the number of arrivals from flow ‘1’ at time slot ‘-1’.

From Theorem III.1 and (8), the rate function is given by:

$$I_F(B) = \inf_{0 < T \leq T_I} I_F^T(B), \quad (9)$$

where

$$I_F^T(B) = \inf_{\tilde{\mathbf{y}}[-T, 0] \in \mathcal{A}} \sum_{i=-T}^{-1} I(y[i]), \quad \mathcal{A} = \left\{ \tilde{\mathbf{y}}[-T, 0] : \tilde{Q}(\tilde{\mathbf{y}}[-T, 0]) \geq B \right\}$$

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

Theorem III.2 (Upper and lower bound).

$$\begin{aligned} \inf_{0 < T \leq T_I} TI \left(C + \frac{B}{T} - \frac{w}{T} \left(\frac{1}{p(B+C)} + \frac{(T-1)}{p(C)} \right) \right) &\leq I_F(B) \\ &\leq \inf_{0 < T \leq T_I} TI \left(C + \frac{B}{T} - \frac{w}{p(C+B/T)} \right) \end{aligned} \quad (10)$$

Proof. The proof is presented in [25] \square

III.C STABILITY-ELASTICITY TRADE-OFF

Using the lower and upper bounds on the rate function derived in the previous section, we study the effect of elasticity of marking functions on the stability (for controlled flows) and QoS (for uncontrolled flows), and their trade-off.

We fix a nominal marking function $p(z)$, and consider the family of marking functions $\{p_\beta(z)\}$ that are correspondingly generated for various values of β , the elasticity parameter. Recall that $\beta > 1$ corresponds to a more elastic marking function, and $\beta < 1$ corresponds to a less elastic marking function.

For the stability analysis, we use the local stability condition for the weighted proportional fair controller from [2, 4], and determine for each marking function $p_\beta(z)$, 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 [2]:

$$\kappa(p_\beta(z^*) + z^* p'_\beta(z^*)) < \sin \left(\frac{\pi}{2(2d+1)} \right), \quad (11)$$

where κ is the gain constant that determines the rate of source adaptation to network feedback (see [4] for details). Further, by definition of $p_\beta(x)$, we can show that

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

Thus, for each value of β , the stability condition (11) reduces to

$$\kappa(p(z^*) + \beta z^* p'(z^*)) < \sin \left(\frac{\pi}{2(2d+1)} \right) \quad (12)$$

On the other hand, with the instant adaptation scheme, the upper bound on the rate function from Theorem III.2 provides a lower bound on the queue overflow probability. In other words, for a fixed value of β and the corresponding marking function $p_\beta(z)$, we can get *no better QoS* than that given by Theorem III.2. With a finite value of κ , (and thus, non-instant source adaptation), the QoS will be worse (as the controlled source will take longer to adapt to a burst).

To summarize, for each value of β , we compare the best QoS that can be provided by an instantly adapting source (an ideal scheme), and the corresponding largest round-trip delay that can be tolerated and still lead to system stability, if the same marking function were used with a proportional fair controller. Such a trade-off 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 (12)). On the other hand, increasing β *improves* the QoS behavior for the real-time uncontrolled flows. For this reason, we refer to this study as *stability-elasticity trade-off*.

We illustrate this trade-off in Figure 4. For each marking function in Table 1, we plot the trade-off between largest allowable round-trip delay for stability and queue overflow probability as a parametric plot of β . The parameter settings for Figure 4 are that n , C , w , and the link utilization are set to be 500, 100, 5, and 95%. Uncontrolled arrivals are modeled by bursty two state-Markov ON-OFF process, where packets arrive at the rate of 500 pkts/unit-time in the ON state, and with ON probability being 0.1 (this is denoted by ON-OFF(500,0.1)). Thus, the scaled expected uncontrolled arrival rate is 50, and the fixed

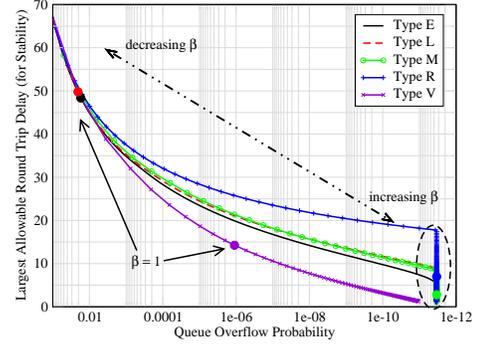


Figure 4: Stability-Elasticity Trade-off: $n = 500$, $w = 5$, $c = 100$, $\kappa = 0.2$ and ON-OFF(500, 0.1)

point of scaled controlled arrival rate is 45 due to the setting of 95% link utilization, leading to the marking probability at the fixed point of $5/45$. The plot clearly illustrates the trade-off between QoS for real-time flows and stability for controlled flows. *The results also 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 for the fixed point considered in Figure 4 (i.e., z^*), a rate based version of REM [7] seems to provide the largest local-stability region for any given QoS requirement.* To analytically construct uniformly optimal marking functions is an interesting problem for future research.

In addition, we see different sensitivities to marking elasticity for different marking functions. The reason why we have vertical lines in the rate based version of REM (Type **R**) and M/M/1 (Type **M**) marking function (in the dotted elliptical region) is that their original (non-warped) marking value $p(z)$ is 1, when $x > C$ (see Table 1). Thus, the queue overflow probability in this case decreases only until some threshold $\bar{\beta}$ and stays constant after this threshold.

III.D SCHEDULING-ELASTICITY TRADE-OFF

In this section, we derive the required capacity (which results in the same queue overflow probability with FIFO scheduling and priority scheduling) as a function of marking elasticity (i.e., β). It is clear that the capacity required for supporting some fixed queue overflow probability L (for real-time flows) with priority scheduling is the smallest (over scheduling policies) since absolute priority is given to these real-time flows (see Figure 1). With FIFO scheduling, the important question to address is: how much *extra capacity* is needed to support a given queue overflow probability L . In this section, we *quantitatively show that this extra capacity can be significantly decreased by appropriately changing the marking elasticity without changing the equilibrium traffic rates.*

The lower bound on the rate function in Theorem III.2 provides an upper bound on the queue overflow probability (which is a function of the capacity at the router). Thus, Theorem III.2 can be used to derive the (upper bound) capacity required to support a given QoS.

With priority queueing (where only stochastic uncontrolled flows are considered, since controlled flows do not affect the queue dynamics for the uncontrolled flows), a sufficient condition of required capacity (in the large number of flows regime) for a given queue overflow probability [12] and for the queue stability is given by:

$$\frac{\Lambda(\delta/B)}{\delta/B} < C_P, \quad (13)$$

$$x^* + y^* < C_P \quad (14)$$

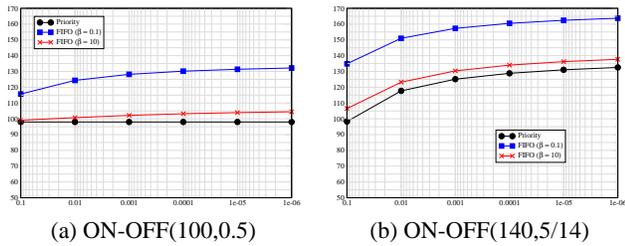


Figure 5: Scheduling-Elasticity Trade-off: $n = 100$, $w = 5$, $z^* = 98$ and $B = 1$

where $n \times C_P$ is the capacity required to support a queue overflow probability of $\exp(-nI_P(B))$, when n flows are present. Note that either (13) or (14) can be the dominant condition of C_P , depending on the given queue overflow probability.

Proposition III.1. A sufficient condition for $I_F(B) > \delta$ and the queue stability is:

$$\frac{\Lambda(\delta/B)}{\delta/B} < C_F - \frac{w}{p(C_F)},$$

where

$$\Lambda(\theta) = \log E[e^{\theta Y_1[-1]}].$$

Proof. The proof is presented in [25]. \square

Based on Proposition III.1 and (13), Figure 5 shows the required scaled capacity (i.e., C_F and C_P) with FIFO and priority scheduling to support a given QoS, for two values of the elasticity parameter β and for different queue overflow probabilities. Note that the parameters of marking functions are automatically determined by choosing link utilization, the fixed point, and the expected value of uncontrolled and controlled arrivals.

First, we observe that for a small value β , the difference between the capacities with FIFO and priority queueing is large for all values of the queue overflow probability. This is due to the fact that the controlled flows back-off sluggishly. On the other hand, for more elastic marking functions, the required capacities with both scheduling algorithms are very close.

For a less bursty uncontrolled arrivals (Figure 5-(a)), in priority scheduling, the queue stability condition (i.e., $C_P < z^* = 98$) dominates the QoS condition (14), while for a more bursty arrivals (Figure 5-(b)), the QoS condition is stronger than the queue stability condition. In both cases, we observe that the required capacity with FIFO can be significantly decreased (almost same as that with priority scheduling) by increasing the marking elasticity.

REFERENCES

- [1] L. Ying, G. Dullerud, and R. Srikant, "Global stability of internet congestion controllers with heterogeneous delays," in *Proceedings of American Control Conference*, June 2004.
- [2] R. Johari and D. Tan, "End-to-end congestion control for the Internet: Delays and stability," *IEEE/ACM Transactions on Networking*, vol. 9, no. 6, pp. 818–832, December 2001.
- [3] S. Kunniyur and R. Srikant, "Analysis and design of an adaptive virtual queue algorithm for active queue management," in *Proceedings of ACM SIGCOMM*, San Diego, CA, August 2001, pp. 123–134.
- [4] F. P. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, pp. 237–252, 1998.
- [5] C. V. Hollot and Y. Chait, "Nonlinear stability analysis for a class of TCP/AQM schemes," in *Proceedings of the IEEE Conference on Decision and Control*, December 2001.
- [6] G. Vinnicombe, "On the stability of end-to-end congestion control for the Internet," 2001, University of Cambridge Technical Report.
- [7] S. H. Low and D. E. Lapsley, "Optimization flow control, I: Basic algorithm and convergence," *IEEE/ACM Transactions on Networking*, pp. 861–875, December 1999.
- [8] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services," December 1998, RFC 2475.
- [9] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance," *IEEE/ACM Transactions on Networking*, vol. 1, no. 4, pp. 397–413, August 1993.
- [10] L. Massouli and J. Roberts, "Bandwidth sharing: objectives and algorithms," *IEEE/ACM Transactions on Networking*, vol. 10, no. 3, pp. 320–328, 2002.
- [11] G. de Veciana and J. Walrand, "Effective bandwidths: Call admission, traffic policing and filtering for ATM networks," *Queueing Systems Theory and Applications*, vol. 20, pp. 37–59, 1995.
- [12] D. D. Botvich and N. G. Duffield, "Large deviations, economies of scale, and the shape of the loss curve in large multiplexers," *Queueing Systems*, vol. 20, pp. 293–320, 1995.
- [13] C. Courcoubetis and R. Weber, "Buffer overflow asymptotics for a switch handling many traffic sources," *Journal of Applied Probability*, vol. 33, pp. 886–903, 1996.
- [14] G. Kesidis, J. Walrand, and C-S. Chang, "Effective bandwidths for multiclass Markov fluids and other ATM sources," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 424–428, 1993.
- [15] N. Likhanov and R. Mazumdar, "Cell loss asymptotics for buffers fed with a large number of independent stationary sources," *Journal of Applied Probability*, vol. 36, pp. 86–96, 1999.
- [16] F.P. Kelly, "Mathematical modelling of the Internet," in *Mathematics Unlimited - 2001 and Beyond (Editors B. Engquist and W. Schmid)*, Berlin, 2001, pp. 685–702, Springer-Verlag.
- [17] S. Deb, S. Shakkottai, and R. Srikant, "Stability and convergence of TCP-like congestion controllers in a many-flows," in *Proceedings of IEEE INFOCOM*, San Francisco, CA, 2003.
- [18] S. Shakkottai and R. Srikant, "How good are deterministic fluid models of Internet congestion control?," in *Proceedings of IEEE INFOCOM*, New York, NY, June 2002, vol. 2, pp. 497–505.
- [19] S. Shakkottai and R. Srikant, "Mean FDE models for Internet congestion control," To appear in *IEEE Transactions on Information Theory*, June, 2004.
- [20] F. P. Kelly, "Models for a self-managed Internet," *Philosophical Transactions of the Royal Society*, vol. A358, pp. 2335–2348, 2000.
- [21] S. Kunniyur and R. Srikant, "End-to-end congestion control: utility functions, random losses and ECN marks," in *Proceedings of IEEE INFOCOM*, Tel Aviv, Israel, March 2000, vol. 3, pp. 1323–1332.
- [22] S. Shakkottai and R. Srikant, "Many-sources delay asymptotics with applications to priority queues," *Queueing Systems: Theory and Applications*, vol. 39, pp. 183–200, October 2001.
- [23] A. Shwartz and A. Weiss, *Large Deviations for Performance Analysis*, Chapman and Hall, New York, NY, 1995.
- [24] D. Wischik, "Sample path large deviations for queues with many inputs," *Annals of Applied Probability*, 2000.
- [25] Y. Yi and S. Shakkottai, "On the elasticity of marking functions: Scheduling, stability, quality-of-service in the internet," Technical Report, 2004.