Deterministic End-to-End Delay Guarantees with Rate Controlled EDF Scheduling

Vijay Sivaraman

CSIRO ICT Centre, Australia. Email: Vijay.Sivaraman@csiro.au

Fabio M. Chiussi

Invento Networks, USA. Email: fabio@inventonetworks.com

Mario Gerla

University of California at Los Angeles, USA. Email: gerla@cs.ucla.edu

Abstract

Quality of Service (QoS) in terms of end-to-end delay guarantees to real-time applications is an important issue in emerging broadband packet networks. Earliest Deadline First (EDF) scheduling, in conjunction with per-hop traffic shaping (jointly called Rate Controlled EDF or RC-EDF) has been recognised as an effective means of end-to-end deterministic delay provisioning. This paper addresses the issue of identifying RC-EDF shaping parameters that realize maximal network utilizations. We first prove that finding "optimal" shapers is in general infeasible, and then propose a heuristic choice derived from the flow's hop-length. Our choice varies gracefully between known optimal settings for the limiting values of the hop-length, and outperforms shaper selections proposed previously in the literature.

Key words: EDF scheduling, per-hop traffic shaping, end-to-end delay guarantee.

1 Introduction

Quality of Service (QoS) in terms of end-to-end transfer delays for real-time communication services such as voice and video is a key issue in emerging broadband packet-switched networks. Providing this QoS requires packet scheduling schemes more sophisticated than First-In-First-Out (FIFO) at switches in the network. The Earliest Deadline First (EDF) [3,12] scheduling discipline, that associates a per-hop deadline with each packet and schedules packets in order of deadlines, is an attractive choice due to its proven optimality [5,8] characteristics at a single node.

In a multi-hop network, a per-hop traffic shaper [14] in conjunction with the EDF scheduler (together called *Rate Controlled EDF* or *RC-EDF*), can support end-to-end delay guarantees [6]. The schedulable region associated with RC-EDF is sensitive to the selection of shaping parameters. Earlier work such as [6] presented ad-hoc choices and [10,13] presented results for restricted settings, but to the best of our knowledge there has been no systematic study of shaping parameters with a view to realizing maximal schedulable regions.

Our work studies the problem of identifying "good" RC-EDF shaping parameters. Our contributions are two-fold. First, we establish analytically that except in trivial cases, it is infeasible to identify "optimal" flow shapers, independent of the other traffic in the network. Second, we propose a heuristic choice of shaper parameters derived from the flow's hop-length. Our heuristic is simple to compute, and varies gracefully between the limiting cases on the hop-length (for which results are already established). Simulation of a realistic traffic scenario show that our heuristic allows RC-EDF to realize significantly larger admissible regions than before.

The rest of this paper is organised as follows: section 2 provides background on EDF scheduling and traffic shaping. The infeasibility of optimal RC-EDF shaping is established in section 3, and our heuristic choice is proposed and its performance demonstrated in section 4. Concluding remarks are presented in section 5.

2 Background

Under the EDF scheduling discipline [3,12], flow i at switch m is associated with a *local* delay bound d_i^m; then, a packet of flow i arriving to the scheduler at time t is stamped with a deadline $t + d_i^m$, and packets in the scheduler are served by increasing order of deadline. EDF is known to be the optimal scheduling policy at a single switch [5], namely has the largest *schedulable region*. For flow i, let $A_i(t)$ denote the traffic envelope, i.e. an upper bound on the flow traffic for every interval of size t. In this work we use the well-known *multiple-leaky-bucket* [2] descriptor $(\sigma_k, \rho_k)_{k=1,\dots,K_i}$ corresponding to the envelope $A_i(t) = \min_{1 \le k \le K_i}\{\sigma_k + \rho_k t\}$. Given N flows, where flow i has delay guarantee d_i, [5,8] show that EDF can support the largest set of delay guarantees, in particular those satisfying the constraint:

$$\sum_{i=1}^{N} A_i(t - d_i) \le Ct, \qquad \forall t > 0 \tag{1}$$

where C denotes link rate and $A_i(t) = 0$ for $t < 0$. To account for non-preemptive scheduling with non-negligible packet sizes, the delay guarantee for each flow is augmented by L/C, where L denotes maximum packet size. Equation (1) directly yields a single-node CAC mechanism.

For a multi-node network, RC-EDF reshapes the traffic at each node. References [6,10] derive fundamental results on shaping and end-to-end delays under RC-EDF. Let $E_i^m(t)$ denote the shaper envelope for flow i at the m-th node on its path, and the symbol $\wedge$ represent concatenation (i.e., series placement) of shapers. This work assumes all traffic and shaper envelopes to be concave, increasing, piecewise-linear functions with a finite number of slopes; thus their multiple-leaky-bucket descriptor $(\sigma_k, \rho_k)_{k=1,\ldots,K}$ satisfies $\sigma_1 < \sigma_2 < \cdots < \sigma_K$ and $\rho_1 > \rho_2 > \cdots > \rho_K$.

Lemma 1 *[6, Proposition 2] Consider flow i traversing nodes $1, \ldots, M$. The RC-EDF discipline that uses shaper E_i^m at node m offers no better end-to-end delay guarantees than the RC-EDF discipline using shaper $E_i' = \bigwedge_{m=1}^{M} E_i^m$ at each node.*

Lemma 2 *[6, Proposition 3] Given flow i with envelope $A_i(t) = \min_{1\le k\le K_i}\{\sigma_{i,k} + \rho_{i,k}t\}$, the* smallest *shaper $E_i(d)(t)$ which guarantees shaping delay $D(A_i \| E_i(d)) \le d$, where $0 \le d \le \sigma_{i,K_i}/\rho_{i,K_i}$, is unique and given by*

$$E_i(d)(t) = \begin{cases} A_i(\tau_{i,k^*})t/(\tau_{i,k^*} + d), & \textit{if } 0 \le t < \tau_{i,k^*} + d \\ A_i(t-d), & \textit{if } t \ge \tau_{i,k^*} + d \end{cases}$$

where $\tau_{i,1} = 0$, $\tau_{i,k} = (\sigma_{i,k} - \sigma_{i,k-1})/(\rho_{i,k-1} - \rho_{i,k})$ for $2 \le k \le K_i$, and $k^* = \min_{1\le k\le K_i}\{k \,:\, A_i(\tau_{i,k}) - \rho_{i,k}(\tau_{i,k} + d) \ge 0\}$.

Lemma 3 *[6, Proposition 4] Consider flow i traffic with envelope $A_i(t)$ traversing nodes $1, \ldots, M$ with corresponding output link speeds C^m. Then given an RC-EDF discipline that uses shaper envelope $E_i(d)(t)$, there is an RC-EDF discipline using shaper envelope $E_i(d')(t)$, $d' \ge d$, which guarantees the same end-to-end delays to all flows and whose peak rate $p' \le \min_{1\le m\le M}\{C^m\}$.*

Lemma 1 states that it suffices to restrict attention to RC-EDF disciplines that for a flow use identical shapers at each node. Lemma 2 shows that a "good" shaper can be characterised by a single parameter: the shaping delay. Lemma 3 further restricts the family of "good" shapers to be those with peak rate no larger than the link rate at any of the switches on the path of the flow. Given such a "good" shaper $E_i(t)$, the end-to-end delay bound d_i for flow i is:

$$d_i = d_i^{sh} + \sum_{m=1}^{M} d_i^m \tag{2}$$

where $d_i^{sh} = D(A_i \| E_i)$ denotes the maximum shaping delay and d_i^m is the local scheduler delay bound at the m-th switch for the flow. The maximum shaper delay is incurred only *once*, and is independent of the number of nodes on the path. Together, (1) and (2) lead to an end-to-end CAC framework [6,1]: for incoming flow

i, a shaping delay and associated shaper are picked; the remaining delay is split among the schedulers on the path of the flow, and the flow is admitted only if the single-node CAC at every switch along the path admits the flow. The schedulable region under RC-EDF depends critically on the choice of shaper E_i, and identifying shapers that yield maximal schedulable regions is the subject of this paper. Before we embark on our study, we note the following result [10, Theorem 4.3] showing that in the restricted setting of single-leaky-bucket envelopes, complete smoothing of traffic from flows with "sufficiently" large hop-length is optimal.

Theorem 1 *[10, Theorem 4.3] Consider RC-EDF wherein traffic shapers are restricted to the single-leaky-bucket form, carrying flow i with envelope $A_i(t) = \sigma_i + \rho_i t$ traversing M hops with link rates $C^1, \ldots, C^M$. Then, if $\sum_{m=1}^{M} \rho_i / C^m \geq 1$, the schedulable region of the RC-EDF discipline is not reduced if the shaper with envelope $E_i(t) = \sigma_i' + \rho_i t$, where $0 \leq \sigma_i' \leq \sigma$, is used for flow i at every switch.*

However, two important questions remain unanswered: 1) In the general setting of multiple-leaky-bucket shapers, can "optimal" shapers be identified?, and 2) Can a general shaping strategy be derived which applies to all flows (not just ones with large hop-lengths) and realizes large RC-EDF schedulable regions? To the best of our knowledge these questions have not been addressed in the literature.

3 Infeasibility of Optimal RC-EDF Shaping

In this section, we address the issue of identifying "optimal" shapers for use with RC-EDF disciplines that guarantee end-to-end deterministic delay bounds to flows.

Definition 1 *An* optimal *shaper E_i for flow i is such that the RC-EDF discipline that uses shaper envelope $E_i(t)$ for flow i guarantees end-to-end delays to all flows no lower than the RC-EDF discipline that uses any envelope $E_i'(t)$ for flow i.*

The shaper envelope for the flow is typically chosen at call-setup, and not modified during the lifetime of the flow. Moreover, the choice is made independent of the other flows in the network, since 1) the number of flows in the network is typically too large, and 2) the set of flows in the network varies dynamically as flows enter and leave the network. Therefore, it is reasonable to restrict our focus to *network-state-independent shapers*, i.e., shapers that are constructed independent of the exogenous traffic in the network. We then ask if there exist shapers that are *both* optimal and network-state-independent, and show that in general they do not, except in the trivial cases of constant bit rate or single-hop flows.

Consider a link of rate C employing EDF scheduling, and let the workload $\mathcal{W}$ consist of N flows, where flow i has concave piecewise linear envelope $A_i(t)$ and maximum delay requirement d_i. Then we define the following:

Definition 2 *The* service demand $D_{\mathcal{W}}(t)$*, the* residual capacity $F_{\mathcal{W}}(t)$ *and the* effective residual capacity $R_{\mathcal{W}}(t)$ *corresponding to the workload* $\mathcal{W}$ *are defined by*

$$D_{\mathcal{W}}(t) = \sum_{i=1}^{N} A_i(t - d_i), \quad t \geq 0 \tag{3}$$

$$F_{\mathcal{W}}(t) = Ct - D_{\mathcal{W}}(t), \quad t \geq 0 \tag{4}$$

$$R_{\mathcal{W}}(t) = \min_{t' \geq t} F_{\mathcal{W}}(t'), \quad t \geq 0 \tag{5}$$

Further define $\{R_{\mathcal{A}} \succeq R_{\mathcal{B}}\} \equiv \{\forall t \geq 0 : \ R_{\mathcal{A}}(t) \geq R_{\mathcal{B}}(t)\}$, and $\{R_{\mathcal{A}} \succ R_{\mathcal{B}}\} \equiv \{R_{\mathcal{A}} \succeq R_{\mathcal{B}} \text{ and } \exists t \geq 0 : \ R_{\mathcal{A}}(t) > R_{\mathcal{B}}(t)\}$. Then we establish the following lemmas (proofs in appendix A):

Lemma 4 *Let* $\mathcal{W}$ *denote the workload at the EDF scheduler. A disjoint workload* $\mathcal{U}$ *is admissible at the EDF scheduler if and only if* $\forall t \geq 0 : \ D_{\mathcal{U}}(t) \leq R_{\mathcal{W}}(t)$*.*

Lemma 5 *Let workloads* $\mathcal{A}$ *and* $\mathcal{B}$ *yield effective residual capacities* $R_{\mathcal{A}}(t)$ *and* $R_{\mathcal{B}}(t)$ *respectively, and let* $\lim_{t\to\infty} R_{\mathcal{A}}(t)/t > 0$ *and* $\lim_{t\to\infty} R_{\mathcal{B}}(t)/t > 0$*. Then* $R_{\mathcal{A}} \succeq R_{\mathcal{B}}$ *if and only if, for every workload* $\mathcal{U}$*,* $\mathcal{B} \cup \mathcal{U}$ *is feasible implies that* $\mathcal{A} \cup \mathcal{U}$ *is also feasible.*

We are now ready to prove the main claim of this section, which states that a shaper cannot be both network-state-independent and optimal. Recall that it suffices to focus on shapers which, for any flow, are identical at each node (lemma 1), have the smallest envelope for a given shaping delay (as per the construction in lemma 2), and have peak rate no larger than the link rate (lemma 3).

Theorem 2 *Consider flow* i *with multiple-leaky-bucket input traffic arrival envelope* $A_i(t) = \min_{1\leq k\leq K}\{\sigma_{i,k} + \rho_{i,k}t\}$ *(*$K \geq 2$*) that traverses nodes* $1, \ldots, M$ *(*$2 \leq M < \infty$*) with corresponding output link speeds* C^m*. Further, let the peak rate of flow* i *be no more than the minimum link speed along the path of the flow (i.e.,* $\sigma_{i,1} = 0$ *and* $\rho_{i,1} \leq \min_{1\leq m\leq M}\{C^m\}$*). Then, there does not exist a network-state-independent shaper* $E_i(d_i^{sh})$ *that is optimal, in the sense of guaranteeing that the RC-EDF discipline employing shaper envelope* $E_i(d_i^{sh})(t)$ *provides end-to-end delays to all flows no worse than the RC-EDF discipline that uses shaper envelope* $E_i(d_i^{'sh})(t)$ *for arbitrary* $0 \leq d_i^{'sh} \leq \delta_{i,K}/\rho_{i,K}$*.*

Proof: By contradiction. Let $d_i > 0$ denote the end-to-end delay requirement of flow i, and assume that there exists a value of d_i^{sh} such that the shaper envelope $E_i(d_i^{sh})(t)$ is network-state-independent and optimal. Thus, irrespective of the cross-traffic at the various switches, the RC-EDF discipline that uses shaper envelope $E_i(d_i^{sh})(t)$ guarantees end-to-end delays to all flows no worse than the RC-EDF discipline that uses any shaper $E_i(d_i^{'sh})(t)$ where $0 \leq d_i^{'sh} \leq \delta_{i,K}/\rho_{i,K}$. Consider the two cases:

Case I – $d_i^{sh} > 0$: Choose the cross-traffic at each switch m to be a single flow with dual-leaky-bucket envelope $A^m(t) = \min\{C^m t, \frac{d_i \rho_{i,1}}{M} + (C^m - \rho_{i,1})t\}$ and hop-length 1. From 1), the envelope $E_i(d'^{sh})(t)$ where $d'^{sh} = 0$ can guarantee an end-to-end delay bound of d_i to flow i (by guaranteeing delay bound d_i/M at each node) while simultaneously providing a delay bound of 0 to each of the other flows. The shaper envelope $E_i(d^{sh})(t)$, however, cannot guarantee these delay bounds, since $d^{sh} > 0$ implies that at least one node m on the flow's path has to guarantee a delay bound lower than d_i/M to the flow i envelope $E_i(d^{sh})(t)$, but simultaneously providing a delay bound of 0 to the cross-traffic at node m is not feasible.

Case II – $d_i^{sh} = 0$: Select the cross-traffic at each switch m to be a single flow with dual-leaky-bucket envelope $A^m(t) = \min\{C^m t, \frac{d_i - \epsilon}{M} \frac{\rho_{i,1}}{1+\epsilon(\rho_{i,1}-\rho_{i,2})/\sigma_{i,2}} + (C^m - \frac{\rho_{i,1}}{1+\epsilon(\rho_{i,1}-\rho_{i,2})/\sigma_{i,2}})t\}$ (where $0 < \epsilon < \min\{d_i, \frac{\sigma_{i,2}}{\rho_{i,2}}\}$) and hop-length of 1. The shaper envelope $E_i(d'^{sh})$ where $d'^{sh} = \epsilon$ guarantees end-to-end delay bound d_i to flow i (by guaranteeing a delay bound of $\frac{d_i - \epsilon}{M}$ at each hop) as well as delay 0 to the cross-traffic. However, $E_i(d^{sh})(t)$ where $d^{sh} = 0$ cannot simultaneously provide a delay bound of 0 to the cross-traffic while providing an end-to-end delay bound of d_i to flow i, as that would require at least one of the nodes m on the path to provide a delay bound no larger than d_i/M, and this is not feasible for $M > 1$. □

That shapers cannot be network-state-independent and yet optimal, may seem to be in apparent contradiction with existing results. Reference [13] claims that smoothing is beneficial for homogeneous flows if and only if the hop-lengths are larger than a critical value; however, this does not hold when the flow homogeneity restriction is removed. In the single-leaky-bucket setting, [10, Theorem 4.3] (reproduced as theorem 1 above) claims optimality; in what follows we show how it reconciles with our theorem 2. Note first that the single-leaky-bucket envelopes in theorem 1 are outside the family of "good" shapers, since (a) for non-zero bucket size, the single-leaky-bucket has infinite peak rate, whereas theorem 2 requires peak shaping rates to be no greater than the link rate, and (b) the single-leaky-bucket form does not allow the smallest shaper for a given shaping delay budget (as in lemma 2) to be realized, and is thus sub-optimal. In what follows we show that the apparent advantage of shaping in the single-leaky-bucket case is due to poor shaper description (by virtue of the single-leaky-bucket restriction) rather than inherent advantage of smoothing. We first establish the following lemmas (proofs in appendix A):

Lemma 6 *Consider an arbitrary workload at an EDF scheduler operating at rate C. A flow f with envelope $A(t) = \sigma + \rho t$ and delay bound d is admissible if and only if flow f' with envelope $A'(t) = \min\{Ct, \sigma(1 - \frac{\rho}{C}) + \rho t\}$ and delay bound $d' = d - \sigma/C$ is admissible.*

Lemma 7 *For an arbitrary workload at an EDF scheduler operating at rate C, if the flow f with envelope $A(t) = \sigma + \rho t$ can be guaranteed delay bound d, then the flow f' with envelope $A'(t) = \sigma' + \rho t$, where $\sigma' \leq \sigma$, can be guaranteed delay bound $d' = d - \frac{\sigma - \sigma'}{C}$.*

To see why it may not be possible to provide a delay bound tighter than d' to flow f' in the above lemma, consider cross traffic with envelope $E(t) = Cd - \sigma - \frac{\epsilon\sigma}{C} + \epsilon t$. Flow f' is not admissible for any delay bound lower than $d - \frac{\sigma - \sigma'}{C - \epsilon}$, which, for small ϵ, can be arbitrarily close to d'. This prepares us for the following theorem:

Theorem 3 *Consider a flow i with single-leaky-bucket input traffic envelope $A(t) = \sigma_i + \rho_i t$ that traverses nodes $1, \ldots, M$ with corresponding output link speeds C^m. Then the RC-EDF discipline that employs for flow i the network-state-independent shaper $E_i^m(t) = \min\{C^m t, \sigma_i(1 - \frac{\rho_i}{C^m}) + \rho_i t\}$ at each node $m = 1, \ldots, M$ guarantees end-to-end delay bounds no worse than* any *RC-EDF discipline that employs for flow i* only *single-leaky-bucket network-state-independent shapers.*

Proof: If flow i uses shaper $E(t) = A(t)$, let the corresponding delay at node m be denoted by d_i^m, and the total end-to-end delay $d_i = d_i^1 + d_i^2 + \ldots + d_i^M$.

Consider first the RC-EDF discipline that is restricted to single-leaky-bucket shaping. Denote by $E_i(t) = \sigma_i' + \rho_i t$ the shaper envelope (where $\sigma_i' \geq 0$ is picked independent of the cross-traffic in the network), and by δ_i^m the delay bound at node m. Then the shaping delay $\delta_i^{sh} = \frac{\sigma_i - \sigma_i'}{\rho_i}$, and the end-to-end delay bound for flow i is $\delta_i = \delta_i^{sh} + \delta_i^1 + \ldots + \delta_i^M$. From lemma 7, $\delta_i^m = d_i^m - \frac{\sigma_i - \sigma_i'}{C}$ (note that by virtue of the network-state-independent property tighter delay bounds cannot be guaranteed). Thus $\delta_i = \frac{\sigma_i - \sigma_i'}{\rho_i} + \sum_{m=1}^{M}(d_i^m - \frac{\sigma_i - \sigma_i'}{C^m})$, i.e.,

$$\delta_i = d_i + \frac{\sigma_i - \sigma_i'}{\rho_i}\left[1 - \sum_{m=1}^{M}\frac{\rho_i}{C^m}\right] \tag{6}$$

The above equation incidentally proves theorem 1 by showing that $\sum_{m=1}^{M}\frac{\rho_i}{C^m} > 1$ implies $\delta_i < d_i$, i.e. smoothing can be beneficial irrespective of exogenous traffic.

Now consider the RC-EDF discipline that enforces the peak rate at each node, i.e. uses shaper envelope $E_i^m(t) = \min\{C^m t, \sigma_i(1 - \frac{\rho_i}{C^m}) + \rho_i t\}$ at node m. The total shaping delay is $\theta_i^{sh} = \sigma / C^{\min}$ where $C^{\min} = \min_{1 \leq m \leq M}\{C^m\}$. From lemma 6, the delay bound at node m is $\theta_i^m = d_i^m - \sigma_i / C^m$. The end-to-end delay bound is thus $\theta_i = \sigma_i / C^{\min} + \sum_{m=1}^{M}(d_i^m - \sigma_i / C^m)$, i.e.,

$$\theta_i = d_i + \frac{\sigma_i}{C^{\min}} - \sum_{m=1}^{M}\frac{\sigma_i}{C^m} \tag{7}$$

Since $0 \leq \sigma_i' \leq \sigma$ and $\rho_i \leq C^{\min}$, it can be seen that $\theta_i \leq \delta_i$ for any choice of σ_i'. $\square$

This shows that RC-EDF disciplines that use a naive dual-leaky-bucket shaper can realize better end-to-end delays than RC-EDF disciplines restricted to single-leaky-bucket shapers. The apparent advantage of smoothing in theorem 1 thus arises from poor shaper description rather than an inherent benefit of smoothing.

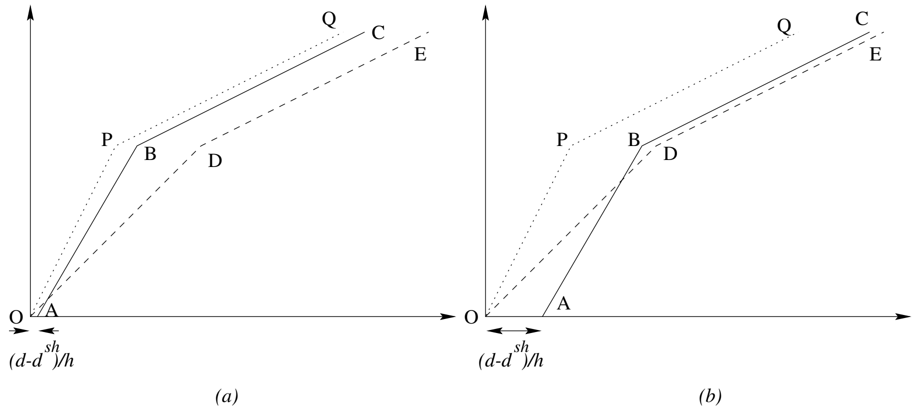

Fig. 1. Service demands for high and low choices of d^{sh} for (a) large h, (b) small h

4 Heuristic Shaper Choice

Though we cannot hope to realize network-state-independent shapers that are *optimal* (as established in theorem 2), we can still identify *heuristic* choices that can be expected to perform well for reasonably realistic traffic mixes. In general, smoothing can be expected to be useful for flows with large hop-lengths and detrimental when the hop-lengths are small. To illustrate this with an example, consider a flow f with (p, σ, ρ) dual-leaky-bucket ingress traffic. The envelope $A(t) = \min\{pt, \sigma + \rho t\}$ is depicted by OPQ in figure 1. Let d (where $d \leq \sigma/\rho$) denote the end-to-end delay requirement and h the hop-length of the flow. Further, assume that once a shaping delay d^{sh} has been selected for the flow, the remaining scheduling delay $d - d^{sh}$ is split equally among the hops.

Consider first the case when h is reasonably large. Figure 1(a) shows, at a switch, the service demand ABC when d^{sh} is very small (i.e., very little smoothing) as also the service demand ODE when $d^{sh} = d$ (complete smoothing). Though ABC $\not\succeq$ ODE, the service demand ODE lies below the service demand ABC for the most part. Therefore it seems reasonable to expect that smoothing will yield considerable benefits when h is high. On the other hand, when h is low, the service demand ABC corresponding to small d^{sh} as shown in figure 1(b) is preferable in general over the smoothed case ODE. Thus a small value for d^{sh} can be expected to yield better performance when h is low. These observations lead us to propose the following heuristic choice of the shaping delay:

$$d^{sh} = \min\left\{ d\left(1 - \frac{1}{h}\right), \frac{\sigma_K}{\rho_K} \right\} \tag{8}$$

The shaper envelope corresponding to this choice of shaping delay can be computed using lemma 2. Such a choice of shaping delay conforms to the intuition presented above, and varies gracefully between the limiting cases. When $h = 1$, d^{sh} computes to zero, consistent with earlier results showing that smoothing of single-hop flows

Table 1
Four-Segment Characterisation for Six MPEG-Coded Movie Traces

Movie	σ_1	ρ_1	σ_2	ρ_2	σ_3	ρ_3	σ_4	ρ_4
Advertisements	0	1.6	800.0	0.8000	1333.0	0.6000	1600.0	0.5330
Jurassic	0	4.0	133.3	1.0540	400.0	0.8533	1066.0	0.7619
Mtv	0	6.0	266.6	2.3565	933.3	1.9730	1866.6	1.8666
Silence	0	4.0	266.6	0.6665	533.3	0.6000	1133.0	0.5000
Soccer	0	5.0	266.6	2.5000	1000.0	1.2380	2133.3	1.0666
Terminator	0	3.4	133.3	0.7878	266.6	0.5866	800.0	0.3666

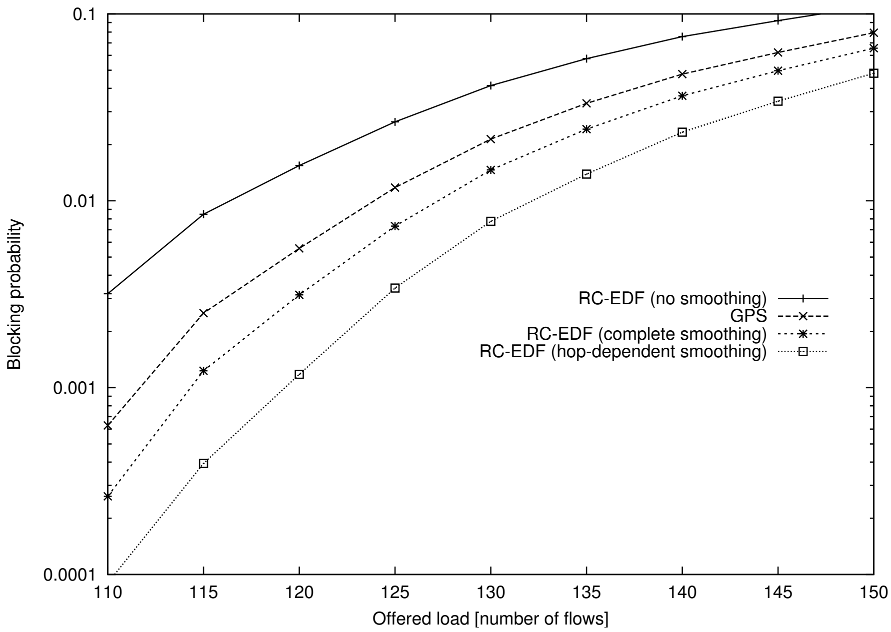

Fig. 2. Call blocking probabilities under RC-EDF (no smoothing), GPS, RC-EDF (complete smoothing) and RC-EDF (hop-length dependent smoothing)

is detrimental to network performance [7]. On the other extreme, $h \to \infty$ yields $d^{sh} = d$, signifying that the end-to-end delay budget is better used towards traffic smoothing rather than being fragmented as scheduler delay among the hops. This is again consistent with [13] showing that in the restricted case where all traffic is homogeneous, smoothing is beneficial is and only if the hop-lengths are *above* a critical value.

To quantify the advantages of our proposed "hop-length dependent" shaper choice, we perform simulations of an OC-3 ATM switch operating at 155 Mbps and multiplexing a traffic mix consisting of six types of video flows, with traffic characteristics shown in table 1. Each row represents a four-segment leaky-bucket characterisation $(\sigma_k, \rho_k)_{k=1,\cdots,4}$ of a movie trace, where the σ's are in Kbits and the ρ's in

Mbits/s. These characterisations are borrowed from [4], and have been derived as four-segment covers of the empirical envelopes of traces of MPEG-1 coded movies in [11]. Flow arrivals are generated according to a Poisson process with parameter α and their durations are exponentially distributed with mean $1/\beta$. The ratio α/β characterises the load offered to the link. Each flow has traffic characteristics chosen randomly from the characteristics of the six types shown in table 1. The end-to-end delay requirement d (excluding propagation delays) of the flow is uniformly distributed in [100ms, 1.5s], and its hop-length uniformly chosen in [1,5]. After a flow is generated with the above parameters, shaper envelopes for the flow are selected as per the three shaping strategies: 1) no smoothing ($d^{sh} = 0$), 2) complete smoothing ($d^{sh} = \min\{d, \sigma_4/\rho_4\}$), and 3) hop-length dependent smoothing per (8). The remaining delay $d - d^{sh}$ is then split equally among the hops, and the EDF call acceptance test is performed at the switch to determine if the flow can be accepted into the network (it is assumed that the switch of interest is the bottleneck, and hence determines if the flow can be accepted into the network or not). We generate 100,000 flows in each simulation run, and plot the call blocking probability in figure 2 under 1) Generalized Processor Sharing (GPS) scheduling [9], 2) RC-EDF with no smoothing, 3) RC-EDF with complete smoothing, and 4) RC-EDF with the hop-length dependent smoothing method of (8), as the offered load is varied from 110 to 150 calls. It is seen that RC-EDF employing our hop-length dependent shaper significantly outperforms the other methods, and hence has the potential to offer larger schedulable regions in real networks.

5 Conclusions

EDF scheduling with per-hop traffic shaping (called RC-EDF) has been proposed as an attractive mechanism for supporting traffic with deterministic end-to-end delay requirements [6,1]. In this paper, we have addressed the question of identifying traffic shapers that realize large schedulable regions under RC-EDF. We have shown that identifying the "optimal" shaper is in general infeasible, and have proposed a heuristic choice which is simple to compute and yields increased schedulable regions for realistic traffic scenarios. Future work will look into characterising the performance of RC-EDF in the statistical setting.

References

[1] F. M. Chiussi and V. Sivaraman. Achieving High Utilization in Guaranteed Services Networks Using Early-Deadline-First Scheduling. In *IWQoS '98*, pages 209–217, Napa, CA, May 1998.

[2] R. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Trans. Inf. Theory*, 37:114–121, Jan 1991.

[3] D. Ferrari and D. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE Jrnl. Selected Areas in Comm.*, 8(3):368–379, Apr 1990.

[4] V. Firoiu, J. Kurose, and D. Towsley. Efficient Admission Control of Piecewise Linear Traffic Envelopes at EDF Schedulers. *IEEE/ACM Trans. Netw.*, 6(5):558–570, Oct 1998.

[5] L. Georgiadis, R. Guérin, and A. Parekh. Optimal Multiplexing on a Single Link: Delay and Buffer Requirements. *IEEE Trans. Inf. Theory*, 43(5):1518–1535, Sep 1997.

[6] L. Georgiadis, R. Guérin, V. Peris, and K. Sivarajan. Efficient Network QoS Provisioning Based on per Node Traffic Shaping. *IEEE/ACM Trans. Netw.*, 4(4):482–501, Aug 1996.

[7] E. Knightly and P. Rossaro. On the Effects of Smoothing for Deterministic QoS. *Dist. Sys. Eng. Jrnl: Special Issue on Quality of Service*, 4(1):3–15, Mar 1997.

[8] J. Liebeherr, D. Wrege, and D. Ferrari. Exact Admission Control for Networks with a Bounded Delay Service. *IEEE/ACM Trans. Netw.*, 4(6):885–901, Dec 1996.

[9] A. Parekh and R. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case. *IEEE/ACM Trans. Netw.*, 1(3):344–357, Jun 1993.

[10] V. Peris. *Architecture for Guaranteed Delay Service in High Speed Networks*. PhD thesis, University of Maryland, College Park, 1997.

[11] O. Rose. Statistical Properties of MPEG Video Traffic and their Impact on Traffic Modeling in ATM Systems. Technical Report TR 101, University of Würzburg, Würzburg, Germany, Feb 1995. The MPEG coded video traces are available at ftp://ftp-info3.informatik.uni-wuerzburg.de/pub/MPEG/.

[12] D. Verma, H. Zhang, and D. Ferrari. Guaranteeing Delay Jitter Bounds in Packet Switching Networks. In *Proc. TRICOMM*, pages 35–46, Chapel Hill, NC, Apr 1991.

[13] T. Wu and E. W. Knightly. Buffering vs. Smoothing for end-to-end QoS: Fundamental Issues and Comparison. In *Perf. '99*, Istanbul, Turkey, Oct 1999.

[14] H. Zhang and D. Ferrari. Rate-Controlled Service Disciplines. *Jrnl. High Speed Netw.*, 3(4):389–412, 1994.

A Proofs of Lemmas

Lemma 4: From (1), the admissibility criteria for workload $\mathcal{U}$ is seen to be $\forall t \geq 0 : D_{\mathcal{U}}(t) \leq F_{\mathcal{W}}(t)$. Since all traffic envelopes $A_i(t)$ are non-decreasing functions of t, so is the service demand $D_{\mathcal{U}}$ of workload $\mathcal{U}$. Thus the above condition is equivalent to $\forall t \geq 0 : D_{\mathcal{U}}(t) \leq \min_{t' \geq t} F_{\mathcal{W}}(t')$. The quantity on the right is nothing but $R_{\mathcal{W}}(t)$; this proves the result. $\square$

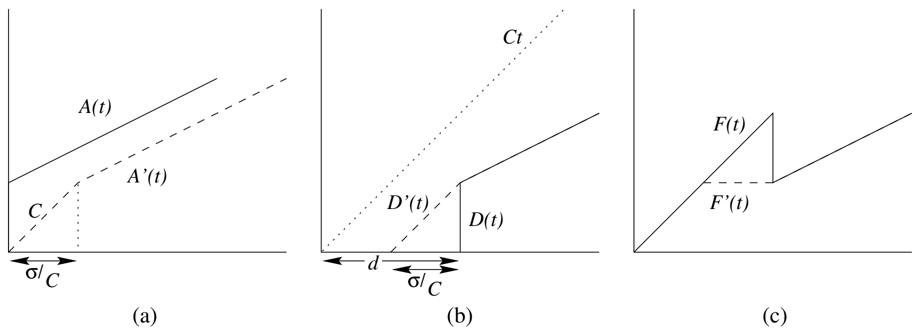

Fig. A.1. The (a) envelopes, (b) service demands and (c) residual capacities for flows f (solid lines) and f' (dashed lines) in the proof of lemma 6.

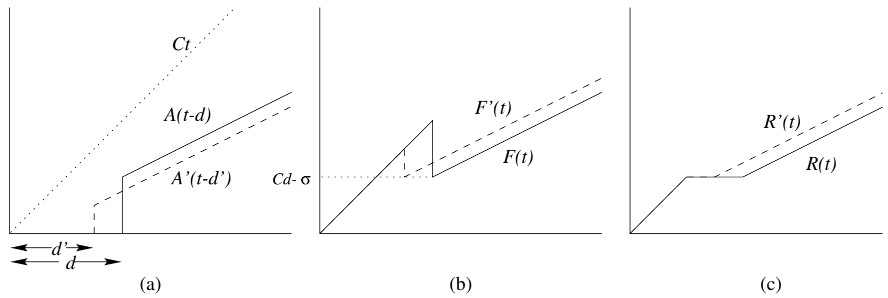

Fig. A.2. The (a) service demands, (b) residual capacities and (c) effective residual capacities for flows f (solid lines) and f' (dashed lines) in the proof of lemma 7.

Lemma 5: *If part*: Say $R_{\mathcal{A}} \succeq R_{\mathcal{B}}$. Then any workload $\mathcal{U}$ satisfying $\forall t \geq 0: \; D_{\mathcal{U}}(t) \leq R_{\mathcal{B}}(t)$ (i.e., the workload $\mathcal{B} \cup \mathcal{U}$ is feasible) also satisfies $\forall t \geq 0: \; D_{\mathcal{U}}(t) \leq R_{\mathcal{A}}(t)$, and so the workload $\mathcal{A} \cup \mathcal{U}$ is also feasible. *Only if part*: Say $R_{\mathcal{A}} \not\succeq R_{\mathcal{B}}$. Then $\exists \tau : \; R_{\mathcal{B}}(\tau) - R_{\mathcal{A}}(\tau) = 2\delta > 0$. Consider workload $\mathcal{U}$ consisting of a single flow with envelope $A(t) = R_{\mathcal{B}}(\tau) - \delta + \epsilon t$ and delay requirement τ. Since $R_{\mathcal{B}}(t)$ is monotonically non-decreasing and $\lim_{t \to \infty} R_{\mathcal{B}}(t)/t > 0$, ϵ can be chosen small enough such that $\forall t \geq 0: \; D_{\mathcal{U}}(t) \leq R_{\mathcal{B}}(t)$, thus making the workload $\mathcal{B} \cup \mathcal{U}$ feasible. But the workload $\mathcal{A} \cup \mathcal{U}$ is not feasible, since $D_{\mathcal{U}}(\tau) = R_{\mathcal{B}}(\tau) - \delta \not\leq R_{\mathcal{A}}(\tau)$. The existence of a workload $\mathcal{U}$, such that $\mathcal{B} \cup \mathcal{U}$ is admissible while $\mathcal{A} \cup \mathcal{U}$ is not, completes the proof. □

Lemma 6: The envelope $A(t)$, service demand $D(t)$ and residual capacity $F(t)$ for the workload consisting of flow f with delay bound d are shown in figure A.1, as are the corresponding quantities $A'(t)$, $D'(t)$ and $F'(t)$ for the workload consisting of flow f' with delay bound d'. It is easy to see that the residual effective capacities are equivalent in both cases, and therefore by lemma 5 the result follows.

Lemma 7: Consider the effective residual capacities $R(t)$ when the workload consists of the flow f with delay bound d, and $R'(t)$ when the workload consists of the flow f' with delay bound $d' = d - \frac{\sigma - \sigma'}{C}$. The envelopes, residual capacities and effective residual capacities are depicted graphically in figure A.2, and it is easily verified that $R' \succeq R$. From lemma 5, it follows that for any workload, if flow f is feasible, so if flow f'. Thus f' is admissible if f is. □