Efficient Admission Control for EDF Schedulers *

Victor Firoiu vfiroiu@cs.umass.edu
Jim Kurose kurose@cs.umass.edu
Don Towsley towsley@cs.umass.edu

Department of Computer Science
University of Massachusetts
Amherst, MA 01003-4610 USA

Abstract

In this paper we present algorithms for flow admission control at an EDF link scheduler when the flows are characterized by peak rate, average rate and burst size. We show that the algorithms have very low computational complexity and are easily applicable in practice. The complexity can be further decreased by introducing the notion of flex *classes. We evaluate the penalty in efficiency that the classes incur to the EDF scheduler. We find that this efficiency degradation can be made arbitrarily small and is acceptable even for a small number of classes.*

1 Introduction

The demand for real-time communication in data networks such as Internet has grown rapidly in recent years. Two important examples are voice and video communication over the Internet – applications that require timely delivery of data packets. To be able to guarantee such delay requirements, the network must (at least implicitly, if not explicitly) reserve resources at the links on the path of the given real-time flow. Several flow setup protocols that convey end-to-end user delay requirements to the links have been proposed and are in the process of standardization; these include RSVP [2] for the Internet, and ATM signaling [1] for ATM networks.

The problem of providing delay guarantees at a network link is the focus of much current research. A part of this work focuses on the issue of packet scheduling – determining the order in which queued packets are forwarded over outgoing links at switches and routers. This order determines the packets' waiting time in the link's queue, and ultimately the delay that the link scheduler can guarantee. Several analytical models for link scheduling have been proposed in the literature. A variant of Weighted Fair Queuing (WFQ) [4] (also known as Generalized Processor Sharing (GPS) [14]) was proposed in [15] to guarantee a maximum queuing delay by reserving a certain amount of link bandwidth for the given flow. Although simple, this policy is known to be sub-optimal. Another discipline, Earliest Deadline First (EDF) [13] associates a per-hop deadline with each packet and schedules packets in the order of their assigned deadlines. EDF has been proven to be an optimal scheduling discipline in the sense that if a set of tasks is schedulable under any scheduling discipline (i.e., if the packets can be scheduled in such a way that all of their deadlines are met), then the set is also schedulable under EDF. Also, Rate-Controlled EDF [16] was proven to outperform GPS in providing end-to-end delay guarantees in a network [8]. In the present work we adopt the Rate-Controlled model where the EDF scheduler of each link has an independent contribution to the end-to-end delay guarantee of a flow. The end-to-end admission control is thus reduced to EDF schedulability verifications at each link.

Sufficient conditions for the EDF schedulability of flows have been proposed for some particular cases of flow characterizations [9, 17]. Recently, a set of necessary and sufficient conditions for flow schedulability has been put forward by [7, 12], using a general characterization of flows. The fact that EDF is an optimal scheduling policy and that there exist necessary and sufficient conditions for schedulability makes EDF an attractive choice for providing delay guarantees for real-time flows. There are, however, two important concerns about the practicality of EDF scheduling. First, the implementation of EDF scheduling requires a search of $O(logQ)$ time in the list of packets (ordered by their deadlines) waiting in the queue of length Q for transmission. This issue has been successfully addressed in [11], where the search time is brought to constant ($O(1)$) time by discretizing the range of packet dead-

* This material is based upon work supported in part by the National Science Foundation under Grant NCR-95-08274 and CDA-95-02639. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

line values. The second issue is that, although the EDF schedulability conditions in [12] can be expressed simply, the algorithms to perform these schedulability tests can be computationally complex, or, in the general case, require an unbounded number of values that must be checked.

In this paper we consider the problem of simplifying the computation of EDF flow schedulability conditions and present simple and computationally efficient algorithms for performing flow admission at links using EDF scheduling. We take advantage of the particular flow characterization (peak rate, mean rate and burst size) proposed in the emerging standards of Internet Integrated Services [15] and ATM signaling [1]. We find that our algorithms have low complexity ($O(N)$), where N is the number of admitted flows at the EDF scheduler at the moment of the algorithms' invocation. We further simplify these algorithms significantly by discretizing the range of values for certain flow parameters (the flex points) to a predefined set of values (classes). We obtain a very significant improvement in the execution time (two orders of magnitude speedup), with the additional benefit of the execution time being no longer dependent on the number of flows N. We examine the relative performance degradation (in terms of the number of flows admitted) incurred by the introduction of classes and find the tradeoff to be small.

The remainder of the paper is organized as follows. In Section 2 we describe the requirements imposed by IP and ATM flow setup protocols on the local (link) admission control. In Section 3 we derive simple admission control algorithms for flows characterized by peak rate, average rate, and burst size. In Section 4 we evaluate by simulation the EDF admission control with classes. Section 5 concludes the paper.

2 Flow Admission Control in Networks: EDF Schedulers

Flow setup protocols for real-time flows such as ATM signaling and RSVP with Guaranteed Services have certain requirements for flow admission control algorithms at a link. In this section, we examine these requirements; in Section 3 we present specific admission control algorithms meeting these requirements.

Consider a source that wishes to establish a flow f to a destination using ATM signaling. It sends a SETUP message to the destination, including information such as the flow's traffic characteristics (maximum cell rate, sustained cell rate, maximum burst size [1]), and the maximum allowable end-to-end delay, d_f. At each link l along the path from source to destination, the minimum delay that link l can guarantee to f, $\overline{d}_l$ is computed, and added to d_c, the cumulative delay included in the SETUP message. If at some node the cumulative delay exceeds the maximum allowable delay, the flow cannot be accepted, and a RELEASE message is returned. Otherwise, at the end of the first pass (at the destination node), $d_f \geq d_c$ and the flow is accepted. A CONNECT message is returned on the same path to the source, assigning a delay $d_{fl} \geq \overline{d}_l$ to flow f at link l on path P, such that $\sum_{l \in P} d_{fl} \leq d_f$ according to some delay division policy (see for example [6]).

A similar mechanism is used by RSVP with Guaranteed Services [15] where the minimum end-to-end delay can be accumulated in the D_{tot} term of the TSpec, and the receiver's delay requirement can be specified in the delay slack term S of the RSpec [5].

We see that each of the above flow setup protocols requires that a local admission control procedure be invoked at each link l with the following capabilities:

- given a flow f and its characteristics, provide the delay bound that link l can guarantee to f, $\overline{d}_l$, based on the current state (set of reserved flows) at the local scheduler;
- given a flow f, its characteristics, and a requirement $d_{fl} \geq \overline{d}_l$, update the current state of the local scheduler following the reservation of f.

In the following we examine how to provide these capabilities in the case of EDF scheduling.

[12, 7] have given flow schedulability conditions at EDF schedulers for flows characterized by minimum envelopes, or rate-controlling functions. Consider a data flow f with the amount of arrivals (measured in bits/second) in the time interval $[t_1, t_2]$ denoted by $A_f[t_1, t_2]$. The flow is characterized by the minimum envelope A^*_f, an upper bound on the flow's arrival pattern:

$$A_f[t, t+\tau] \leq A^*_f(\tau) \quad \forall t \geq 0, \forall \tau \geq 0$$

We take $A^*_f(t) = 0, \forall t < 0$. Note that, in order to provide a better intuition, in this paper we measure the traffic in number of data units (bits) rather than transmission time (seconds), as in [12].

Let $\mathcal{N} = \{1, 2, ..N\}$ be a set of flows, where flow $i \in \mathcal{N}$ is characterized by the envelope A^*_i. The stability condition for a work-conserving scheduler (thus including the EDF scheduler) is ([12], eq. (5)):

$$\lim_{t \to \infty} \frac{\sum_{i \in \mathcal{N}} A^*_i(t)}{ct} < 1 \qquad (1)$$

where c is the constant rate of the server (bits/second). Assuming a preemptive EDF scheduler or negligible packet sizes (as in the case of ATM cells), we give the following variant of the schedulability condition proposed in [12] for the set $\mathcal{N}$ of flows.

Theorem 1 (Liebeherr,Wrege,Ferrari 1994) *Let $\mathcal{N}$ be a set of flows, stable by (1), where flow $i \in \mathcal{N}$ is characterized by the envelope A^*_i and has a maximum packet delay* *of d_i. The set $\mathcal{N}$ is EDF-schedulable if and only if:*

$$ct \geq \sum_{i \in \mathcal{N}} A_i^*(t - d_i), \quad \forall t \geq 0 \tag{2}$$

We say that the set $(A_i^*, d_i)_{i \in \mathcal{N}}$ is schedulable if (1) and (2) are satisfied. [12] provides schedulability conditions, but does not provide algorithms for schedulability testing.

3 EDF Admission Control for (C, σ, ρ) Token Bucket Flows

3.1 Analysis of EDF Schedulability Conditions

Let us consider flows that are characterized by the following type of envelope, referred to as (C, σ, ρ) envelope, used in both IP [15] and ATM [1] networks:

$$A_i^*(t) = \begin{cases} 0 & t < 0 \\ C_i t & 0 \leq t < a_i \\ \sigma_i + \rho_i t & a_i \leq t \end{cases} \tag{3}$$

where

- C_i is the peak rate of the flow (bits/second);
- $\sigma_i \geq 0$ is the maximum burst size (bits);
- $\rho_i > 0$ is the average rate of the flow (bits/second), and $\rho_i < C_i$;
- $a_i = \sigma_i/(C_i - \rho_i)$ (seconds) is the maximum duration of the flow's burst at peak rate.

Figure 1 shows an example of a (C, σ, ρ) envelope. We shall refer to the point (a, Ca) of A^* as the *flex* point of A^*.

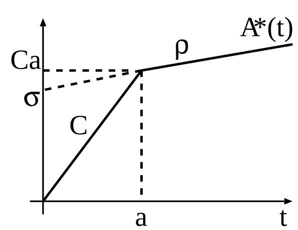

Figure 1: An illustration of (C, σ, ρ) envelope

Let $\mathcal{N}$ be a set of flows, flow i being characterized by the envelope A_i^* of the form given in (3) and having a maximum packet delay requirement of d_i. The stability condition (1) becomes:

$$\sum_{i \in \mathcal{N}} \rho_i < c \tag{4}$$

Let us consider next the schedulability conditions (2). Defining

$$F(t) = ct - \sum_{i \in \mathcal{N}} A_i^*(t - d_i)$$

the schedulability condition (2) becomes $F(t) \geq 0 \quad \forall t \geq 0$, which in turn is equivalent to $F(u) \geq 0$ for all $u \geq 0$ that are proper local minima for F. (u is a proper local minimum for F if $F(t) \geq F(u)$ in a vicinity of u and F is not constant in any vicinity of u.) Given that A_i^* has the form in (3) for all $i \in \mathcal{N}$, it is easy to see that all proper local minima of F are included in the set $\{d_i + a_i | i \in \mathcal{N}\} \cup \{0\}$. Hence, schedulability condition (2) is equivalent to $F(d_i + a_i) \geq 0$ $i \in \mathcal{N}$ and $F(0) \geq 0$. $F(0) \geq 0$ is equivalent to $d_i \geq 0$, $i \in \mathcal{N}$. Let us assume, without loss of generality, that the flows in $\mathcal{N} = \{1, 2, ..N\}$ are ordered such that:

$$i < j \Rightarrow d_i + a_i \leq d_j + a_j \quad \forall i, j \in \mathcal{N} \tag{5}$$

Since the form of A_i^* is given in (3), $F(t)$ becomes:

$$\begin{cases} 0, & t < 0 \\ ct - \sum_{\substack{i \in \mathcal{N} \\ d_i < t}} C_i(t - d_i), & 0 \leq t < d_1 + a_1 \\ ct - \sum_{\substack{i \in \mathcal{N} \\ i \leq j}} (\sigma_i + \rho_i(t - d_i)) & \\ \quad - \sum_{\substack{i \in \mathcal{N} \\ i > j, d_i < t}} C_i(t - d_i), & d_j + a_j \leq t < d_{j+1} + a_{j+1} \\ ct - \sum_{i \in \mathcal{N}} (\sigma_i + \rho_i(t - d_i)), & d_N + a_N \leq t \end{cases} \tag{6}$$

Thus the schedulability conditions (in addition to $d_i \geq 0$, $i \in \mathcal{N}$) are:

$$F(d_j + a_j) = c(d_j + a_j) - \sum_{\substack{i \in \mathcal{N} \\ i \leq j}} (\sigma_i + \rho_i(d_j + a_j - d_i)) \\ - \sum_{\substack{i \in \mathcal{N} \\ i > j, d_i < d_j + a_j}} C_i(d_j + a_j - d_i) \geq 0, \quad j \in \mathcal{N} \tag{7}$$

Suppose now that a new flow, f, arrives at the EDF scheduler. Let f be characterized by (C_f, σ_f, ρ_f) and have a delay guarantee d_f, and assume (without loss of generality) that there is $b \in \mathcal{N}$ such that $d_{b-1} + a_{b-1} < d_f + a_f \leq d_b + a_b$. By inserting the flow f with $d_f + a_f$ in the set $\mathcal{N}$ that is ordered by $(d_i + a_i)_{i \in \mathcal{N}}$, we obtain $\mathcal{N}' = \mathcal{N} \cup \{f\} = \{1, 2, ..N + 1\}$, as in Figure 2. If $d_f + a_f \leq d_1 + a_1$ or $d_f + a_f > d_N + a_N$, f is inserted as the first or last element of $\mathcal{N}'$ respectively.

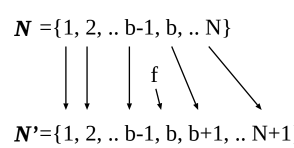

Figure 2: The mapping of $\mathcal{N}$ to $\mathcal{N}'$

The set $\mathcal{N}'$ is schedulable iff (for $i \in \mathcal{N}$):

$$d_i \geq 0 \quad (8)$$
$$d_f \geq 0 \quad (9)$$
$$F(d_i + a_i) \geq 0, \quad d_i + a_i \leq d_f \quad (10)$$
$$F(d_i + a_i) - C_f(d_i + a_i - d_f) \geq 0, \quad d_f < d_i + a_i \leq d_f + a_f \quad (11)$$
$$F(d_f + a_f) - C_f a_f \geq 0 \quad (12)$$
$$F(d_i + a_i) - (\sigma_f + \rho_f(d_i + a_i - d_f)) \geq 0, \quad d_f + a_f < d_i + a_i \quad (13)$$

3.2 Admission control algorithms for EDF schedulers

Let us consider the problem of computing the minimum delay $\overline{d}$ that can be guaranteed to a flow f characterized by (C_f, σ_f, ρ_f) at an EDF scheduler that has allocated a schedulable set $\mathcal{N}$ of flows. This reduces to the problem of computing the minimum value for d_f that satisfies the constraints (8)-(13). In the following we explain intuitively the solution to this problem. The formal solution is given in Lemma 1 and Theorem 2.

From (6) it follows that F has the general form shown in Figure 3: it is continuous, linear on intervals, concave on the intervals $(0, d_1 + a_1)$, $(d_1 + a_1, d_2 + a_2)$,.. $(d_{N-1} + a_{N-1}, d_N + a_N)$ and convex at $d_1 + a_1$, $d_2 + a_2$,.. $d_N + a_N$. Given the flow f with envelope A^*_f as in Figure 1, the

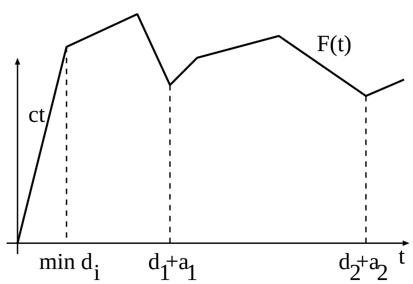

Figure 3:

problem of finding the minimum value for d_f (let us call this minimum value $\overline{d}$) that can be guaranteed to f reduces to determining the leftmost position for A^*_f such that it is below the graph of F for all $t \geq 0$, as in Figure 4, 5 and 6. Three sets of constraints are imposed by F on A^*_f:

- The first segment of $A^*_f(t - \overline{d})$, for $t \in (\overline{d}, \overline{d} + a_f)$ must lie below any local minimum of F that is less than $C_f a_f$. This is expressed by (11), is depicted in Figure 4, and is considered in Lemma 1.1. By defining y_i to be a lower bound on $\overline{d}$ imposed by the local minimum in $d_i + a_i$ of F on the first part of A^*_f, we have that $\overline{d} \geq \max_i y_i$.
- The second segment of $A^*_f(t - \overline{d})$, for $t \in (\overline{d} + a_f, \infty)$ must lie below any local minimum of F that is greater than $C_f a_f$. This is expressed by (13), is depicted in Figure 5, and is considered in Lemma 1.1. By defining x_i to be a lower bound on $\overline{d}$ imposed by the local minimum in $d_i + a_i$ of F on the second part of A^*_f, we have that $\overline{d} \geq \max_i x_i$.
- Finally, the flex point $(\overline{d} + a_f, C_f a_f)$ of $A^*_f(t - \overline{d})$ must lie below F within any concavity interval of F. This is expressed by (12), is depicted in Figure 6, and is considered in Lemma 1.2. By defining $\overline{d}_b$ to be a lower bound on $\overline{d}$ imposed by F on the flex point of A^*_f, we have that $\overline{d} \geq \overline{d}_b$.

In the following lemma and theorem we give the formal solution for computing the minimum delay that can be guaranteed to a flow. The proofs can be found in [5].

Lemma 1 *Let $\mathcal{N}$ be a schedulable set of flows $(C_i, \sigma_i, \rho_i, d_i)_{i \in \mathcal{N}}$ sorted in increasing order of $(d_i + a_i)$, and let (C_f, σ_f, ρ_f) characterize a new flow f such that the stability condition $\sum_{i \in \mathcal{N}} \rho_i + \rho_f < c$ is satisfied.*

1. *Let $(y_i)_{i \in \mathcal{N}}$ such that*
$$F(d_i + a_i) = C_f(d_i + a_i - y_i) \quad i \in \mathcal{N} \quad (14)$$
and let $m_y = \max(0, \max_{\substack{i \in \mathcal{N} \\ F(d_i + a_i) < C_f a_f}} y_i)$.
Let $(x_i)_{i \in \mathcal{N}}$ such that
$$F(d_i + a_i) = \sigma_f + \rho_f(d_i + a_i - x_i) \quad i \in \mathcal{N} \quad (15)$$
and let $m_x = \max(0, \max_{\substack{i \in \mathcal{N} \\ F(d_i + a_i) \geq C_f a_f}} x_i)$.
Let $m = \max(m_x, m_y)$ and b, $1 \leq b \leq N + 1$, such that
$$d_{b-1} + a_{b-1} < m + a_f \leq d_b + a_b \quad (16)$$
where $d_0 + a_0 = -\infty$ and $d_{N+1} + a_{N+1} = \infty$. Then b exists and is unique.

2. *If*
$$F(d_b + a_b) > C_f a_f \quad (17)$$
and
$$F(d_{b-1} + a_{b-1}) < C_f a_f \quad (18)$$
Let $\overline{d}_b$ such that
$$d_{b-1} + a_{b-1} < \overline{d}_b + a_f < d_b + a_b \quad (19)$$
and
$$F(\overline{d}_b + a_f) = C_f a_f \quad (20)$$
Then $\overline{d}_b$ exists, is unique and
$$F(t) < C_f a_f \quad t \in (d_{b-1} + a_{b-1}, \overline{d}_b + a_f) \quad (21)$$
$$F(t) > C_f a_f \quad t \in (\overline{d}_b + a_f, d_b + a_b) \quad (22)$$

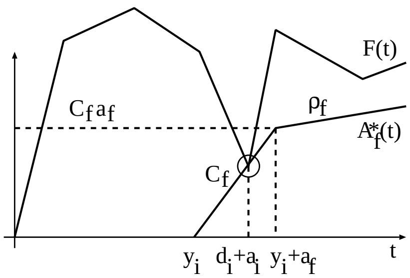

Figure 4:

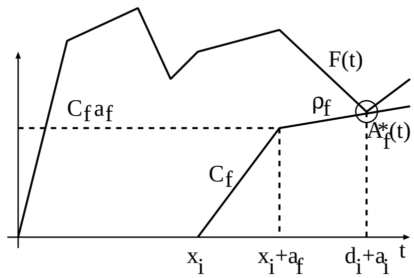

Figure 5:

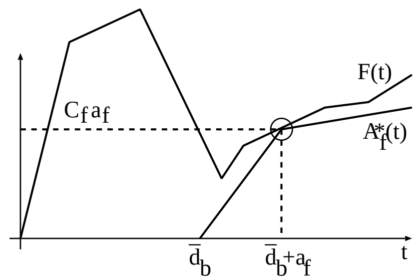

Figure 6:

3. *Let $\overline{d}$ such that:*

$$\overline{d} = \begin{cases} \max(m, \overline{d}_b) & \textit{if (17) and (18)} \\ m & \textit{otherwise} \end{cases} \tag{23}$$

Then $\overline{d} \geq 0$ and

$$d_{b-1} + a_{b-1} < \overline{d} + a_f \leq d_b + a_b \tag{24}$$

Theorem 2 *$\overline{d}$ defined in Lemma 1 is the minimum delay that can be guaranteed to flow f.*

We can see that the computation of $\overline{d}$ has $O(N)$ complexity. Although this may be acceptable for small values of N, even this level of computation can be problematic when there is a large number of flows present (e.g., thousands of flows on an OC12 link). Thus, in the next section we explore a technique for further reducing the flow admission computation time.

3.3 Admission control algorithms for class-based EDF schedulers

We have seen that the admission control of flow f needs N computations $F \geq A_f^*$, one for each flex point $d_i + a_i$. We can simplify this computation by limiting the range of values that flex points can take to a predefined set $\{e_i | i \in \mathcal{L}\}$. If we want to reserve sufficient resources to guarantee the delay d_f to flow A_f^* and to have the flex point $B \in \{e_i | i \in \mathcal{L}\}$, we can reserve the envelope $\overline{A}_f^*$, as in Figure 7. After the reservation, $F'(t) = F(t) - \overline{A}_f^*(t)$, and if we assume that F had all of its relative minima in $\{e_i | i \in \mathcal{L}\}$, then the same is true for F'. The resulting schedulability conditions are $F(e_i) \geq 0$ for $i \in \mathcal{L}$, which results in an $O(L)$ complexity.

Observe that there is a tradeoff associated with the introduction of flex classes: there may be situations where $F(t) \geq A_f^*(t) \quad \forall t$, but $F(t) \geq \overline{A}_f^*(t)$ is not true for some t, i.e., the envelope A_f^* can be admitted, but not $\overline{A}_f^*$. In Section 4 we investigate this tradeoff.

Since all local minima of F are included in the set $(e_i)_{i \in \mathcal{L}}$, the schedulability conditions in (10)-(13) become

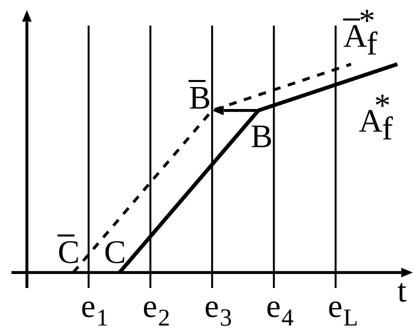

Figure 7:

$(i \in \mathcal{L})$:

$$F(e_i) \geq 0 \quad e_i \leq \overline{d} \tag{25}$$

$$F(e_i) - C_f(e_i - \overline{d}) \geq 0 \quad \overline{d} < e_i \leq \overline{d} + a_f \tag{26}$$

$$F(e_i) - (\sigma_f + \rho_f(e_i - \overline{d})) \geq 0 \quad \overline{d} + a_f < e_i \tag{27}$$

Observe that there is no inequality corresponding to (12), as the constraint is included in (26) since the class-based EDF scheduler mandates $\overline{d} + a_f \in (e_i)_{i \in \mathcal{L}}$. Thus the computation of $\overline{d}_b$ in Lemma 1.2 is no longer needed.

In the next theorem we give the formal solution for computing the minimum delay that can be guaranteed to a new flow.

Theorem 3 *Let a class-based EDF scheduler have a set $\mathcal{L} = \{1, 2, ..L\}$ of classes of flex points, $e_1 < e_2 < ..e_L$. Let $\mathcal{N}$ be a schedulable set of flows $(C_i, \sigma_i, \rho_i, d_i)_{i \in \mathcal{N}}$, sorted in increasing order of $d_i + a_i$ and $d_j + a_j \in (e_i)_{i \in \mathcal{L}}$, and let (C_f, σ_f, ρ_f) characterize a new flow f such that the stability condition $\sum_{i \in \mathcal{N}} \rho_i + \rho_f < c$ is satisfied.*

1. *Let $(y_i)_{i \in \mathcal{L}}$ such that*

$$F(e_i) = C_f(e_i - y_i) \quad i \in \mathcal{L} \tag{28}$$

and let $m_y = \max(0, \max_{\substack{i \in \mathcal{L} \\ F(e_i) < C_f a_f}} y_i)$.

Let $(x_i)_{i \in \mathcal{L}}$ such that

$$F(e_i) = \sigma_f + \rho_f(e_i - x_i) \quad i \in \mathcal{L} \tag{29}$$

and let $m_x = \max(0, \max_{\substack{i \in \mathcal{L} \\ F(e_i) \geq C_f a_f}} x_i)$.

Let $m = \max(m_x, m_y)$ *and* b, $1 \leq b \leq L+1$, *such that*

$$e_{b-1} < m + a_f \leq e_b \quad (30)$$

where $e_0 = -\infty$ *and* $e_{L+1} = \infty$. *Then* b *exists and is unique.*

2. *If* $b = L + 1$, *the flow* f *cannot be scheduled by the class-based EDF scheduler. Otherwise* f *can be scheduled and* $\overline{d} = e_b - a_f$ *is the minimum delay that can be guaranteed to flow* f *by the class-based EDF scheduler.*

Using Theorem 3 we can give the algorithms for admission control for (C, σ, ρ) flows at class-based EDF schedulers. In Figure 8, we give an algorithm to compute the minimum delay that can be guaranteed to a new flow (C_f, σ_f, ρ_f), using a set of pre-computed parameters $W_i = F(e_i)$ and $B = \sum_{i \in \mathcal{N}} \rho_i$.

MINIMUM_DELAY_CLASS(input: $(W_i)_{i \in \mathcal{L}}$, B, $(e_i)_{i \in \mathcal{L}}$, (C_f, σ_f, ρ_f);
output: $\overline{d}$)

if $B \leq \rho_f$
  then exit "cannot accept flow f"
$m_x \leftarrow 0$; $m_y \leftarrow 0$
for $i = 1$ **to** L **do**
  if $W_i \geq C_f a_f$
    then $x_i \leftarrow e_i - \frac{W_i - \sigma_f}{\rho_f}$
      $m_x \leftarrow \max(m_x, x_i)$
    else $y_i \leftarrow e_i - \frac{W_i}{C_f}$
      $m_y \leftarrow \max(m_y, y_i)$
$m \leftarrow \max(m_x, m_y)$
find b **s.t.** $e_{b-1} < m + a_f \leq e_b$
if $b = L + 1$
  then exit "cannot accept flow f"
$\overline{d} \leftarrow e_b - a_f$
return $\overline{d}$

Figure 8: An $O(L)$ algorithm for computing the minimum delay for (C, σ, ρ) flow at class-based EDF scheduler

Figure 9 shows an algorithm to update the parameters $(W_i)_{i \in \mathcal{L}}$ and B whenever a flow $(C_g, \sigma_g, \rho_g, e_p)$ joins the class-based EDF scheduler.

We can easily see that we have an overall $O(L)$ complexity algorithm for admission control of (C, σ, ρ) flows at class-based EDF schedulers.

4 Evaluation of admission control algorithms through simulations

We will evaluate the benefit of class-based EDF admission control over the non-class-based algorithm by comparing their respective runnning times in a simulation environment. We have seen that the class-based admission

JOIN_UPDATE(input: $(W_i)_{i \in \mathcal{L}}$, B, $(e_i)_{i \in \mathcal{L}}$, $(C_g, \sigma_g, \rho_g, e_p)$;
output: $(W_i)_{i \in \mathcal{L}}$, B)

$B \leftarrow B - \rho_g$
/* update of $(W_i)_i$ */
for $i \leftarrow 1$ **to** L **do**
  if $(e_p \leq e_i)$
    then $W_i \leftarrow W_i - \rho_g(e_i - e_p) - C_g a_g$
    else if $(e_p - a_g < e_i)$
      then $W_i \leftarrow W_i - C_g(e_i - e_p + a_g)$

Figure 9: Updating parameters after a flow join

control has the drawback of admitting less flows. We will evaluate this tradeoff by measuring the link blocking probability yielded by the two algorithms through simulation.

We consider a link that forwards ATM traffic according to the EDF scheduling policy. The characteristics of the flows to be serviced at this link are generated randomly and are intended to cover a wide range of traffic patterns. In our simulations we take $\rho = 10^p Kb/s$ where p is uniformly distributed in $[1, 3]$, that makes ρ cover the range $[10Kb/s, 1Mb/s]$. From multiple video and audio traces we have observed that both C and σ are correlated with ρ. In our simulations we take $C = q * \rho Kb/s$, where q is uniformly distributed in $[2, 5]$. Similarly, $\sigma = r * \rho Kb$ where r is uniformly distributed in $[0.8, 1.6]$. Observe that the range of generated traffic patterns include a typical MPEG video source (sequence of advertisements presented in [10]) with peak rate $C = 1Mb/s$, mean rate $\rho = 500Kb/s$, burst size $\sigma = 500Kb$, and a typical packetized voice source (see e.g., [3]) with peak rate $C = 32Kb/s$, mean rate $\rho = 10Kb/s$, burst size $\sigma = 8Kb$. Flows are created according to a Poisson process with parameter α and their duration is exponentially distributed with mean $1/\beta$. The ratio α/β characterizes the load offered to the link, i.e., the average number of flows that would exist at any time at a link with no capacity limitation. Each flow has a delay requirement $d = 10^s * 30ms$, where s is uniformly distributed in $[0, 1.52]$, thus d ranging in $[30ms, 1s]$. After a flow is generated with the above parameters, its EDF schedulability is verified by our admission control algorithms. We generate 100000 flows in one simulation run, and we are interested in the link blocking probability, i.e., the ratio between the number of rejected flows and the total number of generated flows. We take the link blocking probability for an admission control algorithm as an indication of its performance. In our simulations, we use the method of independent replications to generate 90% confidence intervals for the link blocking probability.

In the following we compare the computational performance of class-based admission control algorithms (having 13 classes) with non-class-based algorithm when both

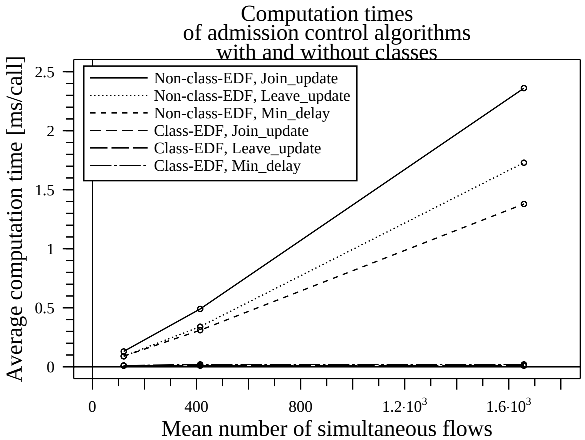

Figure 10:

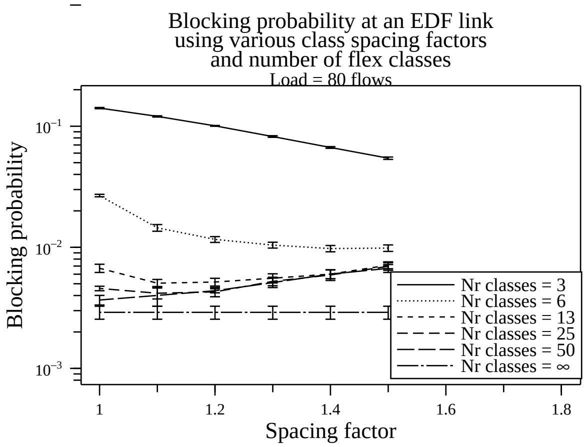

Figure 11:

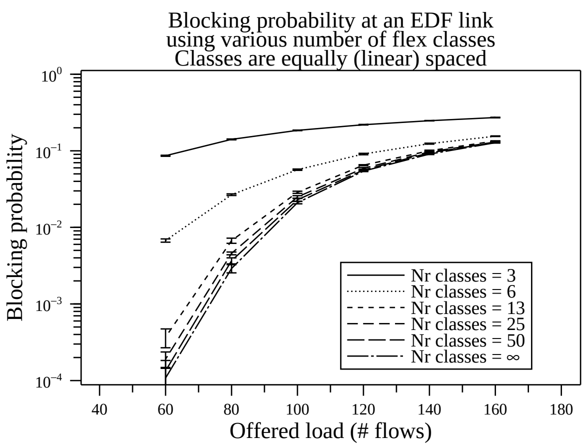

Figure 12:

operate in the same environment. Both algorithms input the same series of flows under three scenarios: link capacity $45Mb/s$ (T3) and offered load 120 flows; link capacity $155.52Mb/s$ (OC3) and offered load 414 flows; link capacity $622.08Mb/s$ (OC12) and offered load 1658 flows.

The offered loads have been chosen to incur the same blocking probability (0.05) in all three scenarios. Given this low rejection probability, the average number of flows N reserved at the link at any time is approximately equal to the offered load. The average computation time has been measured with the GNU code profiler $gprof$ on a DECalpha 347. We know that the admission control algorithm without classes has an asymptotic computation complexity of $O(N)$, which is confirmed by the linear shape of the plots of non-class-based EDF MIN_DELAY, JOIN_UPDATE, LEAVE_UPDATE in Figure 10. The figure confirms also that the execution time $(0.01$ms/function call) of admission control with classes is independent of the number of flows N. Most importantly, the figure shows the very large gain in computation time when using classes: 240 times faster for an OC12 link having an average load of 1658 flows.

For the rest of our simulations we consider a T3 link $(45Mb/s)$. In the following we evaluate the increased blocking penalty when using admission control with classes. Recall that the class-based algorithms in Section 3.3 take their flex point values from a finite set $\{e_i | i \in \mathcal{L}\}$. A large spacing between flex classes (coarse granularity of classes) implies a significant over-reservation for a flow, that would translate in fewer flows being admitted (higher blocking probability). A small spacing between classes, on the other hand, results in a large number of classes and consequently a higher overhead for the admission control algorithms. In the following we address two issues. First, for a fixed number of classes, what is a good policy for choosing the spacing between classes? Second, given that we have found a good spacing policy, what is a number of classes that is sufficient for good link performance and small enough for low computational overhead.

One possibility for class spacing is equal (linear) spacing:

$$e_2 - e_1 = e_3 - e_2 = .. = e_L - e_{L-1}$$

Another possibility is to have the classes geometrically spaced:

$$\frac{e_3 - e_2}{e_2 - e_1} = \frac{e_4 - e_3}{e_3 - e_2} = .. = \frac{e_L - e_{L-1}}{e_{L-1} - e_{L-2}} = \text{spacing factor}$$

This latter spacing policy results in a smaller over-reservation for a small (flex value) request compared to the linear policy, due to a smaller space the request falls in.

In Figure 11 we plot the results of our simulations for values of spacing factor between 1 and 1.5, value 1 corresponding to linear spacing. The graph "Nr classes $= \infty$" corresponds to the exact admission control algorithm (that does not use classes), which forms the base case for our comparison. First, we note that with less that 13 classes, the blocking probability is unacceptably high, compared to the base case. For the rest (more than 13 classes), we see that the linear class spacing policy can provide link performance close to or better than that given by the geometric spacing policy, with any spacing factor. For this scenario, the linear spacing is the solution of choice due to its simplicity and near optimal performance.

In Figure 12 we plot the results of simulation experiments with algorithms using linear spacing and various number of classes. We can see that 13 classes are sufficient to provide link utilization within 10% off the optimal (compare the offered load for the same link blocking probability).

5 Conclusion and future work

In this paper we have proposed practical solutions to the problem of admission control for real-time flows with delay guarantees at an EDF scheduler, as a part of end-to-end flow admission control in IP and ATM networks. We applied the admission control conditions put forward by [12] to flows characterized by peak rate, mean rate and burst size. We developed a first set of algorithms with a computation complexity of $O(N)$, where N is the number of flows admitted in the EDF scheduler at the time of algorithm invocation. A second set of algorithms places the delay requirements of flows into a predefined set of values (classes), thus reducing the computational complexity of admission control to $O(L)$, where L is the number of predefined delay classes. A set of simulation experiments showed that the performance improvement achieved by introducing classes was indeed very important (240 times faster for an OC12 link) and that the algorithms' execution time was independent on the number of flows admitted. We have seen, using simulation, that the relative link performance degradation that the delay classes incur is less that 10%, while using a small number of classes (13). Taken together, these results suggest that the algorithms we have studied in this paper form the basis of a practical and highly efficient solution for the problem of admission control of real-time flows with EDF schedulers.

Our present work can be extended in several ways. First, we can generalize our results to take into consideration packet sizes at non-preemptive EDF schedulers. Second, it is relatively straightforward to extend this work to flows characterized by multiple (σ, ρ) pairs (i.e., envelopes consisting of multiple linear segments).

References

[1] ATM Forum. ATM User-Network Interface Specification. Version 3.1, September 1994.

[2] R. Braden, L. Zhang, D. Estrin, S. Herzog, and S. Jamin. Resource ReSerVation Protocol (RSVP) - Version 1 Functional Specification. Internet draft-ietf-rsvp-spec-12.ps, May 1996.

[3] P. T. Brady. A statistical analysis of on-off patterns in 16 conversations. *Bell System Technical Journal*, 47:73–91, January 1968.

[4] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queuing Algorithm. *Internetworking: Research and Experience*, 1, 1989.

[5] V. Firoiu, J. Kurose, and D. Towsley. Efficient Admission Control for EDF Schedulers. Technical Report TR 96-46, Univ. of Massachusetts, Amherst, 1996. ftp://ftp.cs.umass.edu /pub /techrept /techreport /1996 /UM-CS-1996-046.ps.

[6] V. Firoiu and D. Towsley. Call Admission and Resource Reservation for Multicast Sessions. In *IEEE INFOCOM*, 1996.

[7] L. Georgiadis, R. Guérin, and A. Parekh. Optimal multiplexing on a single link: Delay and buffer requirements. In *IEEE INFOCOM*, 1994.

[8] L. Georgiadis, R. Guérin, V. Peris, and K. Sivarajan. Efficient Network QoS Provisioning Based on per Node Traffic Shaping. In *IEEE INFOCOM*, 1996.

[9] D. Kandlur, K. Shin, and D. Ferrari. Real-time Communications in Multi-hop Networks. In *Proceedings of the 11-th Int'l Conference on Distributed Computing Systems*, 1991.

[10] E. Knightly, D. Wrege, J. Liebeherr, and H. Zhang. Fundamental Limits and Tradeoffs of Providing Deterministic Guarantees to VBR Video Traffic. In *ACM SIGMETRICS'95*, 1995.

[11] J. Liebeherr and D. Wrege. Design and Analysis of a High-Performance Packet Multiplexer for Multiservice Networks with Delay Guarantees. Technical Report CS-94-30, University of Virginia, 1994.

[12] J. Liebeherr, D. Wrege, and D. Ferrari. Exact Admission Control for Networks with Bounded Delay Services. Technical Report CS-94-29, University of Virginia, 1994.

[13] C. Liu and J. Layland. Scheduling Algorithms for Multiprogramming in a Hard-Real-Time Environment. *Journal of ACM*, 20(1), January 1973.

[14] A. K. Parekh. *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks*. PhD thesis, Department of Electrical Engineering and Computer Science, MIT, February 1992.

[15] S. Shenker, C. Partridge, and R. Guerin. Specification of Guaranteed Quality of Service. Internet draft-ietf-intserv-guaranteed-svc-06.txt, August 1996.

[16] H. Zhang and D. Ferrari. Rate-Controlled Service Disciplines. *Journal of High Speed Networks*, 3(4), 1994.

[17] Q. Zheng and K. Shin. On the Ability of Establishing Real-Time Channels in Point-to-Point Packet-Switched Networks. *IEEE Transactions on Communications*, 42(2/3/4), 1994.