Exact Admission Control for Networks with Bounded Delay Services *

Jörg Liebeherr † *Dallas E. Wrege* † *Domenico Ferrari* ††

†Department of Computer Science
University of Virginia
Charlottesville, VA 22903

††University of California at Berkeley
and
International Computer Science Institute
1947 Center St., Suite 600
Berkeley, CA 94704-1105

Abstract

To support the requirements for the transmission of continuous media, such as audio and video, multiservice packet switching networks must provide service guarantees to connections, including guarantees on throughput, network delays, and network delay variations. For the most demanding applications, the network must offer a service which can provide deterministic guarantees for the maximum delay of packets from all connections, referred to as *bounded delay service*. The admission control functions in a network with a bounded delay service must have available *schedulability conditions* that detect violations of delay guarantees in a network switch. In this study, exact schedulability conditions are presented for packet switches which transmit packets based on an Earliest-Deadline-First (EDF) or a Static-Priority (SP) algorithm. The schedulability conditions are given in terms of a general traffic model, making the conditions applicable to a large class of traffic specifications. A comparison of the new schedulability conditions with existing, less accurate, conditions show the efficiency gain obtained by using exact conditions. Examples are presented that show how the selection of a particular traffic specification and a schedulability condition impact the efficiency of a bounded delay service.

Key Words: Multiservice Networks, Real-time Networks, Bounded Delay Service, Multiplexing, Quality of Service, Packet Scheduling, Admission Control, Static-Priority, Earliest-Deadline-First.

*This work is supported in part by the National Science Foundation under Grant No. NCR-9309224.

1 Introduction

Recent technology trends have dramatically advanced the state-of-the art of computer and communication hardware, and have enabled the design of multiservice packet-switching networks that support the transmission of continuous media, such as audio and video. Traditionally, transmission of continuous media was based on circuit-switching technology. By statistically multiplexing network traffic, packet-switching networks are superior to a circuit-switching approach. However, to support the stringent requirements for high quality continuous media transmissions, a packet-switching network must offer services that provide guarantees on delays, delay variations, throughput, and error rate.

Packet-switching networks which support service guarantees to connections must tightly control the use of network resources by limiting both the number of connections as well as the traffic transmitted by each connection. The control of network traffic is performed by relying on admission control functions and policing functions as follows [1, 3, 7, 10, 11]:

1. *Admission Control:* When a network client requests the establishment of a connection, it submits a specification of its maximum traffic together with the desired service guarantees. Admission control functions in the network verify if guarantees can be given without violating any previously given guarantees. If the new connection may result in violations of service guarantees, the connection will not be established. Otherwise the network commits to support the service guarantees for the entire lifetime of the connection.

2. *Traffic Policing:* To ensure that all established connections adhere to the traffic specification given to the network during connection establishment, the network monitors all traffic that enters the network. Traffic from a connection that exceeds its specification is not allowed to enter the network.

In this study, we consider connection-oriented networks that offer a service with deterministic bounds on the network delay for all packets from a connection. We refer to such a service as a *bounded delay service*. The main design goal for a bounded delay service is to maximize the efficiency of the network, that is, to maximize the number of connections that can be supported without violating any delay bound guarantees. The efficiency of a bounded delay service is largely influenced by three factors: (1) the specification which describes the worst case traffic from a connection, (2) the scheduling discipline at the network switches, and (3) the accuracy of the admission control functions.

The traffic specification of a connection describes the maximum traffic that is generated by this connection. Admission control functions use this specification to determine whether to accept or to reject a new connection. Also, traffic monitoring by the policing functions is based on the traffic specification. If the traffic specification for a connection does not precisely describe its actual traffic, the admission control functions will overestimate the resource requirements for a connection. The literature contains several proposals for both discrete traffic specifications [7, 9], which describe traffic as a sequence of packet arrivals with nonzero length, and continuous traffic specifications [4, 14], where traffic is regarded as a continuous stream of data. Since actual network traffic is discrete in nature, discrete traffic specifications enable a more precise description of network traffic.

The scheduling discipline at a network switch determines the order of packet transmission, and thus, controls the variable multiplexing delays of packets. In the presence of admission control and policing, which limit the number of connections and the traffic on the connections, many packet scheduling disciplines can provide bounds on multiplexing delays [6]; however, most scheduling disciplines are not suited for use in a network with bounded delay services. As an example, with FCFS scheduling a bound on the multiplexing delay of a packet is given by the maximum backlog of untransmitted packets. Thus, a network switch based on FCFS scheduling cannot offer a diverse set of delay bounds to connections. Recently, a considerable research effort has resulted in the development of various scheduling techniques for bounded delay services [5, 7, 8, 12, 14, 18, 20].

The admission control functions must have conditions available that can detect if the delay at the network switches may result in a delay bound violation. We refer to conditions that must be satisfied to guarantee that no delay bounds are violated at a network switch as *schedulability conditions*. If exact schedulability conditions are not available, the admission control tests will unnecessarily limit the number of connections in the network, resulting in an inefficient use of network resources.

The three components of a network with bounded delay services, i.e., traffic specification, packet scheduling, and admission control, are highly interdependent in the way they impact the efficiency of a bounded delay service. For example, consider a packet scheduler which can support a large number of connections with delay bound constraints. If the schedulability conditions for this scheduler are inaccurate, the admission control will overestimate the maximum delays of packets, which in turn will decrease the acceptance rate of connection requests. Even if admission control is based on accurate schedulability conditions, an inappropriate traffic specification will overestimate the network traffic, which again will result in a low acceptance rate for connections.

Note, however, that the time cost of high accuracy for any of these three components can be expensive. Accurate traffic specifications increase the complexity of the policing and admission control functions. Also, the complexity of packet scheduling at network switches cannot be arbitrarily high, otherwise scheduling cannot be performed at the speed of the transmission links. Finally, testing accurate schedulability conditions may be time-consuming, resulting in long connection establishment delays. Thus, any implementation of a bounded delay service constitutes a tradeoff between high efficiency and low complexity of the above network components.

In this study, we provide a framework that enables us to quantify the tradeoffs of policing, scheduling, and admission control. We consider two packet scheduling disciplines, Earliest-Deadline-First (EDF) and Static-Priority (SP), both of which have been considered for implementations of bounded delay services [7, 19]. We derive exact schedulability conditions for both scheduling disciplines.[1] Our conditions apply to most traffic specifications, both continuous and discrete, considered in the literature. Knowledge of the exact conditions enables us to compare the efficiency gain obtained by using exact conditions for admission control rather than less accurate, yet computationally less demanding, conditions. Additionally, we will be able to illustrate the efficiency/complexity tradeoff of continuous and discrete traffic specifications.

EDF-schedulers assign each packet a *deadline*, computed as the sum of the arrival time and the delay bound of a packet. The EDF scheduling algorithm always selects the packet with the earliest deadline for transmission. With EDF scheduling, a switch can support a variable set of delay bounds for a large number of connections. However, EDF requires that queued packets be sorted according to their deadlines. Ferrari and Verma presented sufficient schedulability conditions for EDF scheduling for a bounded delay service in [7]. Using a traffic specification which neglects the burstiness of network traffic, Zheng and Shin have derived necessary and sufficient schedulability conditions [21].

A Static-Priority (SP) scheduler provides multiple priority levels, and each connection is assigned to one priority level. SP-schedulers always select the highest-priority packet with the earliest arrival time for transmission. Since SP-schedulers can be implemented with a fixed number of FCFS queues, i.e., one FCFS queue for each priority level, the complexity of scheduling is very low. However, at most one delay bound can be associated with each priority level, thus, limiting the flexibility of SP-schedulers for providing different delay bounds to connections. Using a fluid flow traffic specification, necessary and sufficient schedulability conditions for SP-schedulers are presented in [4]. However, the conditions are not exact for more realistic discrete traffic scenarios. For a particular discrete traffic specification [7], Zhang and Ferrari [20], and Zhang [19] have derived several sufficient schedulability conditions.

The remainder of this study is structured as follows. In Sections 2–4 we present our assumptions on the network and discuss the formal framework used in this study. In Section 2 we give a description of a general traffic model that can express most existing traffic specifications. In Section 3 we discuss assumptions for the network switches, and in Section 4 we give a formal definition of schedulability. In Sections 5 and 6 we provide the necessary and sufficient schedulability conditions for EDF and SP packet scheduling, respectively. In Section 6, we also present several sufficient schedulability conditions for SP packet schedulers which are derived from the exact conditions. In Section 7, we discuss examples that compare the efficiency of EDF and SP scheduling for schedulability conditions with a different degree of accuracy, as well as for different traffic specifications. The conclusions of this study are given in Section 8.

[1] Note that accurate schedulability conditions for EDF and SP scheduling are available in the context of real-time computer systems to determine the delay bound violations of so-called real-time tasks [13, 15]. However, the results obtained for real-time tasks, if applied in a network context, cannot express burstiness of network traffic, arbitrary delay bounds, and nonpreemption of packet transmission.

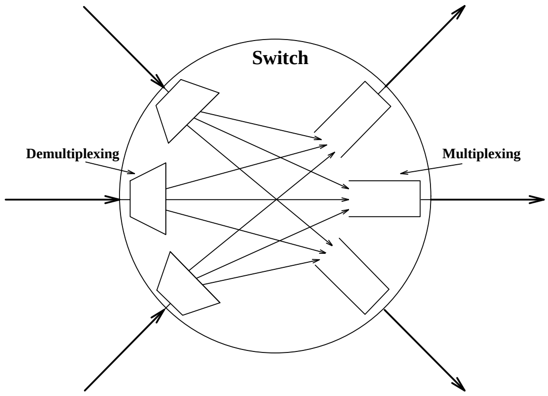

Figure 1: Switch Architecture.

2 Traffic Model

In Figure 1 we show a simplified architecture of a network switch as considered in this study. Traffic that enters the switch through an incoming link is demultiplexed and then routed to the transmission queue of an outgoing link. Each transmission queue has one scheduler that determines the order in which packets in the queue are transmitted. For our purposes it is sufficient to consider a single scheduler for a transmission queue at an arbitrary network switch.

We assume that the scheduler experiences variable-length packet arrivals from a set of connections denoted by $\mathcal{N}$, where $\mathcal{N} = \{1, 2, \ldots, |\mathcal{N}|\}$. To characterize the traffic that arrives at the scheduler of a network switch, we use a general traffic model that allows us to express both continuous and discrete traffic specifications. Continuous traffic specifications regard all traffic as a continuous stream of packets with infinitesimally small size. Discrete traffic specifications consider packets with finite length that arrive at discrete time instants. For discrete traffic specifications we assume that the arrival of a packet occurs instantaneously, that is, a packet arrival is considered complete if the last bit of the packet is received.

For all traffic specifications, we use a right-continuous function A_j to describe the *(actual) traffic arrival* from connection j, where $A_j[t, t+\tau]$ provides the actual arrivals from connection j in time interval[2] $[t, t+\tau]$. (We assume that traffic is measured in terms of the transmission time at the scheduler.)

The maximum traffic from a connection $j \in \mathcal{N}$ is characterized by a *rate-controlling function* A_j^*. We assume that A_j^* is right-continuous, and for continuous traffic specifications, we also assume that A_j^* is concave. The relation between actual and maximum traffic is such that for all times $t > 0$ and for all $\tau \geq 0$, A_j is bounded by A_j^* in the following way [2, 4]:

$$A_j[t, t+\tau] \leq A_j^*[0, \tau] \tag{1}$$

[2] We use [a,b] to denote the set of all x with $a \leq x \leq b$, (a,b] to denote the set of all x with $a < x \leq b$, [a,b) to denote the set of all x with $a \leq x < b$, and (a,b) to denote the set of all x with $a < x < b$.

Note that for continuous traffic specifications, equation (1) follows from the concavity of A_j^*. If equation (1) holds, we say that A_j is *rate-controlled by* A_j^*, denoted by $A_j \prec A_j^*$. In the following we will use $A_j^*(t)$ and $A_j^*(t^-)$ as short-hand notations for $A_j^*[0,t]$ and $A_j^*[0,t)$, respectively, and we set $A_j^*(t) = 0$ and $A_j(t) = 0$ for all $t < 0$.

For discrete traffic specifications we assume that the maximum transmission time of a packet from a connection j is limited by a parameter s_j. In continuous traffic specifications which do not explicitly express packet boundaries, s_j is interpreted as the maximum time interval during which the transmission of traffic from connection j cannot be interrupted.

With the above traffic characterization by rate-controlling functions $A_j^*(t)$, we can express a large class of policing functions for traffic monitoring. As an example, we present a simple traffic specification that is derived from a variation of the leaky bucket traffic policing mechanism [17]. The traffic specification uses three parameters to characterize the traffic from a connection j: the period T_j, the burst size b_j, and the maximum transmission time of a packet s_j. For each connection j there exists a counter with initial value b_j. Each time the connection transmits a packet with transmission time $s \leq s_j$ to the scheduler, the counter is decremented by one. Packets cannot be sent to the scheduler if the counter is zero. The counter is incremented by one after each T_j time units, if its value is less than b_j, and not incremented otherwise. With this characterization, we obtain a discrete traffic specification with the following rate-controlling function $\hat{A}_j^*(t)$ for a connection j:

$$\hat{A}_j^*(t) = b_j s_j + \left\lfloor \frac{t}{T_j} \right\rfloor s_j \tag{2}$$

For the above discrete traffic specification we can provide a corresponding continuous traffic specification which can also be interpreted in terms of a leaky bucket. Here, the counter takes on a continuous range of values with initial value set to $b_j s_j$. The counter is continuously decremented by the amount of traffic that is sent to the switch. If the counter reaches zero, no traffic can be transmitted. If the value of the counter is less than $b_j s_j$, the counter is continuously incremented with rate s_j/T_j, i.e., in a time interval of length $\Delta t > 0$ the increment is given by $(s_j/T_j)\Delta t$. The counter is not incremented if its value reaches $b_j s_j$. The given continuous specification is equivalent to the specification in [4] and results in the following rate-controlling function:

$$\tilde{A}_j^*(t) = b_j s_j + t\frac{s_j}{T_j} \tag{3}$$

Note that the rate-controlling functions in (2) and (3) satisfy $\hat{A}_j^*(t) \prec \tilde{A}_j^*(t)$, i.e., $\tilde{A}_j^*(t)$ is a rate-controlling function for $\hat{A}_j^*(t)$. Since continuous traffic specifications do not account for the discrete nature of packetized network traffic, discrete traffic characterizations more precisely characterize traffic on a connection. In Figure 2 we illustrate the rate-controlling functions, as well as (actual) arrival functions that conform to the rate-controlling function in the sense of the '$\prec$' relation. In Figure 2(a) we show the functions for the discrete traffic specification with $\hat{A}_j^*$ as given in equation (2), and in Figure 2(b) for the continuous traffic specification with $\tilde{A}_j^*$ as given in equation (3).

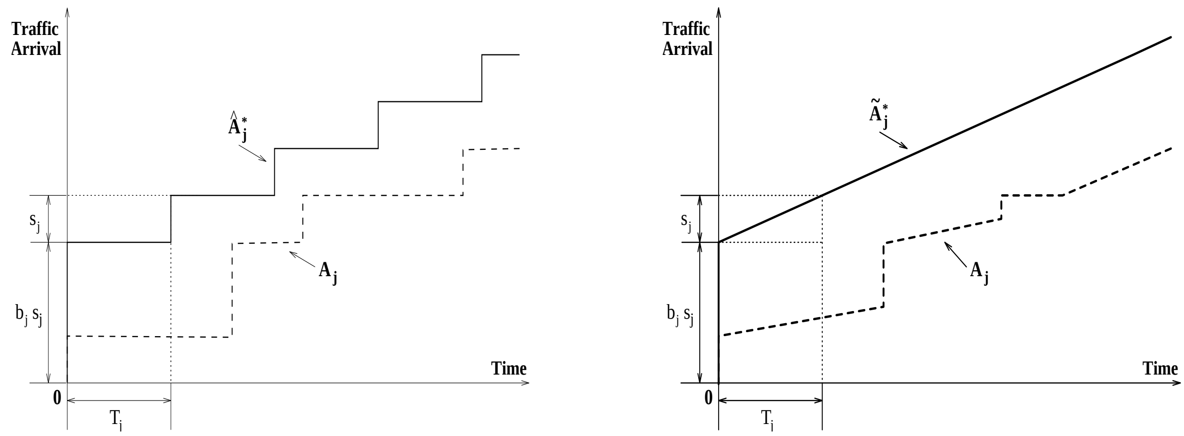

(a) Discrete Traffic Specification.

(b) Continuous Traffic Specification.

Figure 2: Traffic Characterizations.

3 Packet Scheduling

Packet transmission at a network switch is managed by the packet scheduler. Both the EDF-scheduler and the SP-scheduler considered in this study are work-conserving, that is, the scheduler always transmits traffic if its queue is not empty. We assume that schedulers are nonpreemptive. For discrete traffic specifications, nonpreemption implies that the only time instants when the scheduler selects a packet for transmission are (a) upon completion of a packet transmission if additional packets are waiting for transmission, and (b) upon arrival of a packet at an empty scheduler. In continuous traffic specifications we use s_j to define the longest time that the transmission of traffic from connection j cannot be preempted.

We use $W(t)$ to denote the *workload* (or backlog) of traffic at time $t > 0$ waiting to be transmitted by the packet scheduler. By assuming $W(t) = 0$ if $t < 0$, the workload in the scheduler at time $t \geq 0$ due to a set $\mathcal{N}$ of connections with arrival functions $\{A_j\}_{j \in \mathcal{N}}$ is given by [16]:

$$W(t) = \sup_{0 \leq u \leq t} \left\{ \sum_{j \in \mathcal{N}} A_j[u,t] - (t-u) \right\} \tag{4}$$

We denote by $W(t^-)$ the workload at time t excluding the arrivals at time t, that is $W(t^-) = \lim_{h \to 0} W(t-h)$.

A *busy period* of a packet scheduler is a time interval where the scheduler queue is nonempty. Thus, a time interval $[t_1, t_2]$ is a busy period if $W(t_1^-) = 0$, $W(t_2^-) = 0$, and $W(t) > 0$ for all $t_1 \leq t < t_2$. Our definition of a busy period does not coincide with conventional definitions, since we allow a busy period to end at a time t even if the backlog is nonzero during any finite time interval which includes t. This is illustrated in Figure 3.

We say that a packet scheduler is *stable* if all its busy periods are finite. Note that stability of

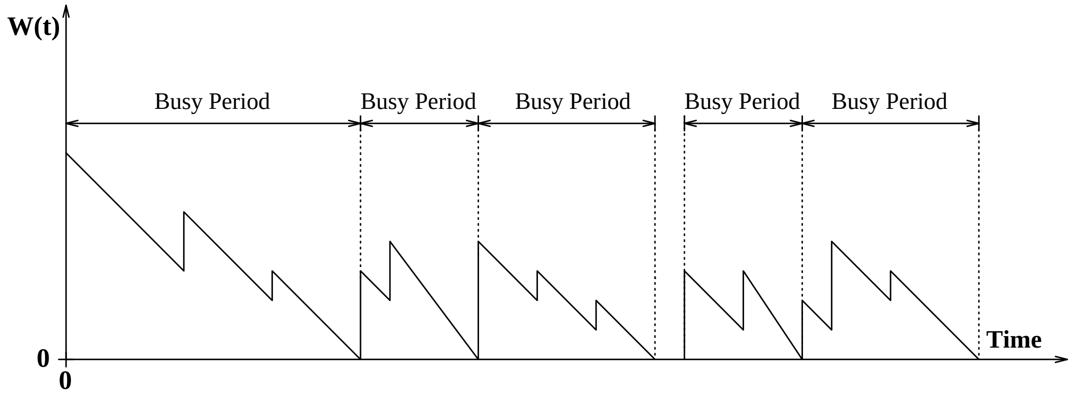

Figure 3: Busy Periods of a Packet Scheduler.

a packet scheduler also implies that the delays in the scheduler queue are finite. The condition for stability of a work-conserving packet schedulers is given by [4]:

$$\lim_{t\to\infty} \frac{\sum_{j=1}^{N} A_j^*(t)}{t} < 1 \qquad (5)$$

For discrete traffic specifications, we allow equality in (5).

4 Schedulability Conditions

The maximum tolerable delay[3] of any packet from connection j in the packet scheduler of a network switch is referred to as the *delay bound* and denoted by d_j. A packet from connection j with a delay bound d_j that arrives at the scheduler at time t is assigned a *deadline* of $t + d_j$. If a packet is not transmitted by its deadline then a *deadline violation* has occurred. We say that a set of connections is *schedulable* if deadline violations never occur. Schedulability is formally defined as follows:

Given a scheduler and a set $\mathcal{N}$ of connections where each connection $j \in \mathcal{N}$ is characterized by (A_j^*, d_j). The set of connections is said to be schedulable *if for all $t > 0$ and for all arrival functions $\{A_j\}_{j\in\mathcal{N}}$ with $A_j \prec A_j^*$ no deadline violation occurs for any connection.*

The conditions which determine if a set of connections is schedulable are referred to as *schedulability conditions*. The efficiency of a bounded delay service is largely influenced by the choice of the schedulability conditions. In Sections 5 and 6 we present the best possible, that is necessary and sufficient, schedulability conditions for an Earliest-Deadline-First scheduler (*EDF-scheduler*) and a Static-Priority scheduler (*SP-scheduler*) for connections with rate-controlled arrival functions.

[3]The delay includes queueing and transmission delays.

5 Earliest-Deadline-First Packet Schedulers

An Earliest-Deadline-First scheduler (EDF-scheduler) assigns each arriving packet a timestamp corresponding to its deadline, i.e., a packet from connection j with a delay bound d_j that arrives at the scheduler at time t is assigned a timestamp of $t + d_j$. The EDF-scheduler maintains a single queue of untransmitted packets, and the queue is sorted in increasing order of packet deadlines. The scheduler always selects the packet in the first position of the queue, that is, the packet with the lowest deadline, for transmission; however, the transmission of a packet is not interrupted by the arrival of a packet with a lower deadline. Since the scheduler queue of an EDF-scheduler must be sorted according to deadlines, each packet arrival involves a search operation to find the correct position of the newly arrived packet in the scheduler queue.

Next we present the necessary and sufficient conditions for schedulability in an EDF-scheduler for traffic specifications that are rate-controlled in the sense of equation (1). We assume without loss of generality that connections are ordered so that $i < j$ whenever $d_i < d_j$. We use B_1 to denote the end of the first busy period for an arrival scenario where all connections transmit according to the rate-controlling function A_j^*, that is,

$$B_1 = \min_{t>0} \left\{ \sum_{j \in \mathcal{N}} A_j^*(t) - t = 0 \right\} \tag{6}$$

Then the schedulability conditions are given as follows:

Theorem 1 *A set $\mathcal{N}$ of connections where each connection $j \in \mathcal{N}$ is characterized by (A_j^*, d_j), is EDF-schedulable for all $A_j \prec A_j^*$ if and only if for all $t \leq B_1$:*

$$t \geq \sum_{j \in \mathcal{N}} A_j^*(t - d_j) \tag{7}$$

and for all t with $d_1 \leq t < d_{|\mathcal{N}|}$:

$$t \geq \sum_{j \in \mathcal{N}} A_j^*(t - d_j) + \max_{d_k > t} s_k \tag{8}$$

The first condition in Theorem 1 is the schedulability condition for a preemptive EDF-scheduler, and the second condition considers that packet transmissions are non-preemptive. A formal proof of Theorem 1 is given in Appendix A.

Example:

Consider a set of connections that conform to the two traffic specifications discussed in Section 2 with rate-controlling functions as given in equations (2) and (3). For the discrete traffic specification in (2) it is not feasible to obtain a closed form expression for equations (7) and (8). However, for the continuous specification in (3). We can modify the conditions in equations (7) and (8) to:

$$\begin{cases} t \geq \sum_{i=1}^{|\mathcal{N}|} s_i \left(b_i + \frac{t-d_i}{T_i} \right) & \text{for } 0 \leq t \leq B_1 \\ t \geq \sum_{i=1}^{j} s_i \left(b_i + \frac{t-d_i}{T_i} \right) + \max_{k>j} s_k & \text{for } d_j \leq t < d_{j+1}, 1 \leq j < |\mathcal{N}| \end{cases} \tag{9}$$

If $\sum_{j=1}^{|\mathcal{N}|} s_j/T_j < 1$, i.e., the stability condition in equation (5) is satisfied, then we obtain for the schedulability conditions that the following must hold:

$$d_j \geq \frac{\sum_{i=1}^{j} s_i \left(b_i - \frac{d_i}{T_i} \right) + \max_{k>j} s_k}{1 - \sum_{i=1}^{j} \frac{s_i}{T_i}} \qquad \text{for all } j \in \mathcal{N} \tag{10}$$

Recall that the rate-controlling function $\tilde{A}_j^*$ for the continuous traffic specification from equation (3) is a rate-controlling function for the discrete specification $\hat{A}_j^*$ from equation (2), i.e., $\hat{A}_j{}^*(t) \prec \tilde{A}_j^*(t)$. Hence, the condition in equation (10) is a sufficient schedulability condition for the discrete traffic specification from equation (9).

6 Static-Priority Packet Schedulers

In this section, we consider that the transmission of traffic is handled by a Static-Priority scheduler (SP-scheduler). An SP-scheduler distinguishes P priority levels and maintains one FIFO queue for each priority. Each connection is assigned a priority p with $1 \leq p \leq P$, and packets arriving on a connection are inserted into the FIFO queue for this connection's priority. At the beginning of a busy period, or after completing the transmission of a packet, the SP-scheduler always selects the first packet in the nonempty FIFO queue with the highest priority for transmission.

An SP-scheduler can only support one delay bound for all connections from the same priority level. Thus, the SP-scheduler is less flexible than an EDF-scheduler in offering different delay bounds to connections. Since the SP-scheduler does not maintain a sorted list of untransmitted traffic as the EDF-scheduler does, the scheduling operations of an SP-scheduler involve less overhead than scheduling with EDF. Due to its simplicity which enables packet scheduling at very high data rates, SP-schedulers are attractive for bounded delay services.

We assume that each connection j is assigned a priority p with $1 \leq p \leq P$. We use $\mathcal{C}_p$ to denote the set of connections with priority p, where a lower priority index indicates a higher priority. All connections in $\mathcal{C}_p$ have the same delay bound d_p, with $d_p < d_q$ for $p < q$. Thus, the priority of a connection is high if its delay bound is short.

Since with continuous traffic specifications the packets on a connection j are infinitesimally small, we define the maximum transmission time of packets by s_j^* for all $j \in \mathcal{N}$ as follows:

$$s_j^* = \begin{cases} s_j & \text{for discrete traffic specifications} \\ 0 & \text{for continuous traffic specifications} \end{cases}$$

Recall that s_j in continuous traffic specifications gives the maximum time interval during which transmissions of traffic from connection j cannot be interrupted. For each priority level p, we define $s_p = \max_{j \in \mathcal{C}_p} s_j$, and $s_p^* = \max_{j \in \mathcal{C}_p} s_j^*$.

We use the term *priority-p busy period* to denote a busy period that is generated by connections with priority equal to or higher than p, and we denote by B_1^p the first priority-p busy period where all connections $j \in \bigcup_{q<p} \mathcal{C}_q$ transmit according to their rate-controlling function A_j^*, that is,

$$B_1^p = \min_{t>0} \left\{ \sum_{q=1}^{p} \sum_{j \in \mathcal{C}_q} A_j^*(t) - t = 0 \right\} \tag{11}$$

With these definitions, we now give the necessary and sufficient schedulability conditions for SP-schedulers.

Theorem 2 *A set $\mathcal{N}$ of connections, where each connection $j \in \mathcal{N}$ is characterized by (A_j^*, d_p), is SP-schedulable for all $A_j \prec A_j^*$ if and only if for all priorities p and for all $0 \le t \le B_1^p - d_p$ there exists a τ with $\tau \le d_p - s_p^*$ such that:*

$$t + \tau \ge \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t + \tau^-) - s_p^* + \max_{r>p} s_r \tag{12}$$

A complete proof of Theorem 2 is presented in Appendix B. Comparing Theorems 1 and 2, we see that testing (exact) schedulability for SP-schedulers requires significantly more effort than for EDF-schedulers. First, condition (12) must be tested for each priority level. Second, for a fixed priority p and fixed value of t, condition (12) must possibly be tested for the entire range of values of τ.

An equivalent formulation of Theorem 2 can be given in terms of the maximum delay of a priority-p packet in the scheduler, denoted by D_p^{max}:

$$D_p^{max} = \max_{t \le B_1^p - d_p} \left\{ \min \left\{ \tau \mid t + \tau \ge \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t + \tau^-) - s_p^* + \max_{r>p} s_r, \tau \ge 0 \right\} \right\} \tag{13}$$

This notation is similar to the one used by Cruz in [4] for sufficient schedulability conditions that apply to continuous traffic specifications. However, the sufficient conditions in [4] can be extended to also consider discrete traffic specifications. Then, the results in [4] provide the following expression for the maximum delay $\overline{D}_p^{max}$:

$$\overline{D}_p^{max} \leq \max_{t+d_p \leq B_1^p} \left\{ \max \left\{ \tau \mid t+\tau \leq \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t+\tau^-) - s_p^* + \max_{r>p} s_r, \tau \geq 0 \right\} \right\} \tag{14}$$

Obviously, all maximum delays that satisfy the condition in (13) also satisfy the condition in (14). The difference between (13) and (14) may appear subtle, and in fact, both expression are identical for most continuous traffic specifications. However, the difference between the two expressions can be large with discrete traffic specifications.

In the following theorems we present two sufficient schedulability condition for SP-schedulers with a lower computational complexity. Both conditions follow directly from Theorem 2. Theorem 4 was first presented by Zhang and Ferrari in [20].

Theorem 3 *A set of $\mathcal{N}$ rate-controlled connections that is characterized by (A_j^*, d_p) is SP-schedulable for all $A_j \prec A_j^*$ if for all priorities p and for all $d_p \leq t \leq B_1^p$ the following holds:*

$$t \geq \sum_{j \in \mathcal{C}_p} A_j^*(t-d_p) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t^-) + \max_{r>p} s_r \tag{15}$$

Theorem 4 (Zhang/Ferrari 1993 [20]) *A set $\mathcal{N}$ of rate-controlled connections characterized by (A_j^*, d_p) is SP-schedulable for all $A_j \prec A_j^*$ if for all priorities p the following holds:*

$$d_p \geq \sum_{q=1}^{p} \sum_{j \in \mathcal{C}_q} A_j^*(d_p) + \max_{r>p} s_r \tag{16}$$

Example:

For discrete traffic models, it is typically not feasible to simplify the conditions in Theorems 2 – 4. However, for the continuous traffic specification in equation (3), the conditions can be simplified as follows. Let us assume that there is only one connection p in each priority set $\mathcal{C}_p$. In this case, we can rewrite the condition in (12) as:

$$t \left(1 - \sum_{j=1}^{p} \frac{s_j}{T_j}\right) + \tau \left(1 - \sum_{j=1}^{p-1} \frac{s_j}{T_j}\right) \geq \sum_{j=1}^{p} b_j s_j + \max_{r>p} s_r \qquad \text{for all } p = 1, 2, \ldots, P \tag{17}$$

Clearly, for fixed τ the condition is satisfied for all $t \geq 0$ if it is satisfied for $t = 0$. Thus, for $\sum_{j=1}^{p} s_j/T_j < 1$, the connections are schedulable if d_p is set to:

$$d_p \geq \frac{\displaystyle\sum_{j=1}^{p} b_j s_j + \max_{r>p} s_r}{1 - \displaystyle\sum_{j=1}^{p-1} \frac{s_j}{T_j}} \qquad \text{for all } p = 1, 2, \ldots, P \tag{18}$$

Since for this traffic specification, equations (13) and (14) coincide, the above conditions are equivalent with the condition given in [4].

Now we turn to Theorem 3. With the continuous traffic specification, the conditions for priority p in Theorem 3 are clearly satisfied for all $t \geq d_p$ if they are satisfied for $t = d_p$. Therefore, we obtain:

$$d_p \geq \frac{\displaystyle\sum_{j=1}^{p} b_j s_j - \frac{s_p d_p}{T_p} + \max_{r>p} s_r}{1 - \displaystyle\sum_{j=1}^{p} \frac{s_j}{T_j}} \qquad \text{for all } p = 1, 2, \ldots, P \tag{19}$$

By simple manipulations we obtain the same condition as in equation (18). Thus, for the continuous traffic specification from equation (3), the conditions given in Theorem 2 and Theorem 3 are identical, and hence, Theorem 3 gives the necessary and sufficient condition.

Finally, Theorem 4 yields the following simplified conditions for the continuous traffic specification given by equation (2).

$$d_p \geq \frac{\displaystyle\sum_{j=1}^{p} b_j s_j + \max_{r>p} s_r}{1 - \displaystyle\sum_{j=1}^{p} \frac{s_j}{T_j}} \qquad \text{for all } p = 1, 2, \ldots, P \tag{20}$$

As in the example discussed in Section 5, the schedulability conditions in (18) and (20) can be used as sufficient schedulability conditions of a discrete traffic specification with a rate-controlling function as in equation (2).

	Group Index j	Delay Bound d_j	Transmission Time per Packet (Max.) s_j	Burst Size b_j	Period T_j
Low Delay Group	1	2 ms	200 μs	8 packets	0.5 – 2 ms
Medium Delay Group	2	4 ms	200 μs	9 packets	0.3 – 2.5 ms
High Delay Group	3	8 ms	200 μs	9 packets	2.5 – 10 ms

Table 1: Parameter Set for Schedulers with 50 Mbps Transmission Rate.

7 Numerical Examples

In Sections 5 and 6 we presented several schedulability conditions for EDF and SP schedulers. Here we use the results to empirically show how the choice of packet schedulers, traffic specifications, and schedulability conditions impact the efficiency of bounded delay services. We consider the necessary and sufficient conditions for EDF schedulers from Theorem 1 and SP-schedulers from Theorem 2 for both continuous and discrete traffic specifications. For SP schedulers, we also use the sufficient schedulability conditions described in Section 6. With these conditions, we present examples that demonstrate the inherent tradeoff between complexity and efficiency in providing bounded delay services.

For the sake of presentation, we consider groups of connections rather than individual connections. By focusing our study upon three groups of connections with similar characteristics, we are able to give an intuitive graphical representation of efficiency.

To characterize the traffic that arrives at a scheduler, we use the discrete and continuous traffic specifications derived from a variation of the leaky bucket traffic policing mechanism as discussed in Section 2. The specifications from equations (2) and (3) are as follows:

$$\begin{aligned} \hat{A}_j^*(t) &= b_j s_j + \left\lfloor \frac{t}{T_j} \right\rfloor s_j && \text{(discrete traffic)} \\ \tilde{A}_j^*(t) &= b_j s_j + t\,\frac{s_j}{T_j} && \text{(continuous traffic)} \end{aligned}$$

We investigate schedulability for a set of three connection groups at a scheduler that operates at 50 Mbps. Table 1 shows the parameters for the connection groups, which we refer to as low delay group, medium delay group, and high delay group. The maximum transmission time of a packet[4] s_j is kept constant at 200 μs for all connection groups corresponding to packets with a maximum size of 1250 Bytes. The maximum burst sizes b_j are 8–9 packets in each connection group. The

[4] For the continuous specification, recall that s_j denotes the longest transmission time that cannot be interrupted.

delay bounds of packets are given by $d_1 = 2$ ms for the low delay group, $d_2 = 4$ ms for the medium delay group, and $d_3 = 8$ ms for the high delay group. The periods T_j of the connection groups are such that the maximum average data rate varies between 4–16 Mbps for the low delay group, 4–26 Mbps for the medium delay group, and 0.8–3.2 Mbps for the high delay group.

The results of the efficiency comparison for the given parameter set are graphically illustrated in Figures 4 and 5. By considering different transmission periods for the connection groups from Table 1, we graph the range of values for which the connection groups are schedulable as described in Section 4. These graphs, referred to as *schedulability graphs*, and interpret them as follows. The volume below the surface in each graph depicts the period values at which no deadline violation occurs for any feasible traffic arrival sequence $\{A_j\}_{j=1,2,3}$ that conforms to a rate-controlling functions $\{A_j^*\}_{j=1,2,3}$ with $A_j \prec A_j^*$. All period values that are not schedulable in the worst case lie in the region above the surface. With these schedulability graphs, we can directly compare the efficiency of two schedulability conditions as follows. If the surface obtained for one schedulability condition lies completely above the surface for a second condition, then the first schedulability condition has a higher efficiency than the second condition.

7.1 Example 1

In the first example, we compare the efficiency of the exact schedulability conditions for EDF schedulers from Theorem 1 and the exact schedulability conditions for SP schedulers from Theorem 2. Figures 4(a) and 4(b) show the schedulability graphs for the EDF scheduler with the chosen discrete and continuous traffic specifications, respectively. The graphs clearly show that the discrete traffic specification yields a better utilization than the continuous traffic specification. Note that the schedulability graph in Figure 4(b) differs significantly from Figure 4(a) in the range where the period of the low delay group is small, that is, the data rate of the low delay group is large.

Figures 4(c) and 4(d) show the corresponding schedulability graphs for the SP scheduler[5] with both the discrete and continuous traffic specifications. Similar to EDF, the discrete traffic specification for SP schedulers enables a more efficient utilization of the scheduler than the continuous traffic specification for SP. Comparing Figure 4(a) with Figure 4(c), and Figure 4(b) with Figure 4(d), we see that EDF is significantly more efficient than SP for both the discrete and continuous traffic specification for this choice of parameters.

It is noteworthy to compare the schedulability graph for the discrete traffic SP scheduler in Figure 4(c) with the graph for the continuous traffic EDF scheduler in Figure 4(d). For this example, these two graphs are similar with respect to efficiency. Thus, we obtain similar schedulability graphs by selecting a complex scheduler with a simple condition, i.e., EDF with a continuous traffic specification, or a simple scheduler with a complex condition, i.e., SP with a discrete traffic

[5]If two connection groups have different delay bounds, the SP scheduler assigns a higher priority to the connection group with the smaller delay bound.

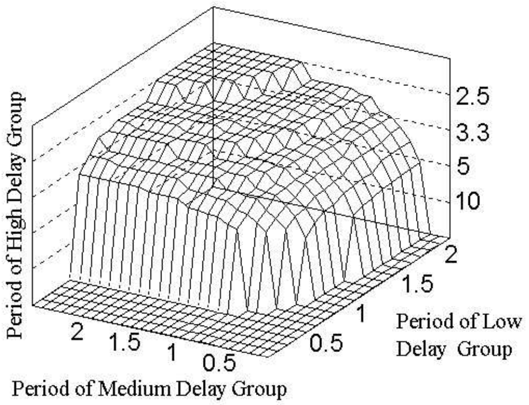

(a) EDF Scheduler (discrete traffic).

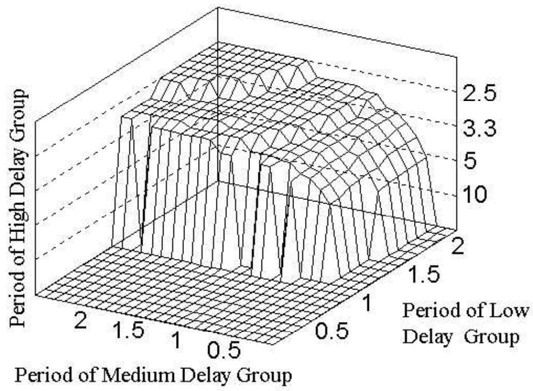

(b) EDF Scheduler (continuous traffic).

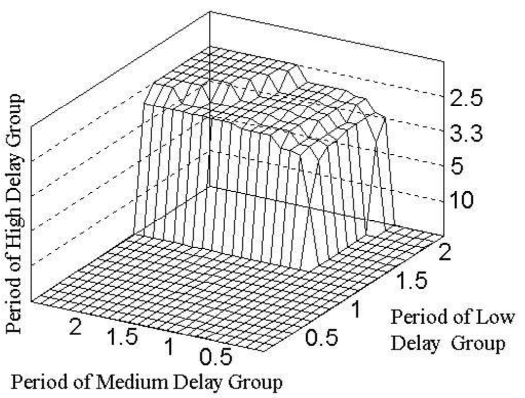

(c) SP Scheduler (discrete traffic).

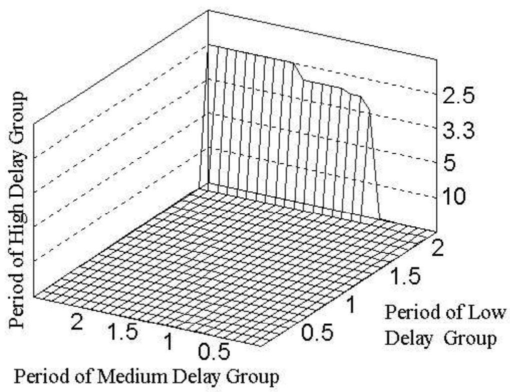

(d) SP Scheduler (continuous traffic).

Figure 4: Exact conditions for EDF and SP Schedulers (time values in milliseconds).

specification. These results indicate that the selection of packet scheduler and traffic specification are independent in their impact on the efficiency of the bounded delay service.

7.2 Example 2

Here we compare the efficiency of the sufficient schedulability conditions for SP schedulers given in equation (14) and Theorems 3–4 with the exact conditions for SP from Theorem 2. We consider the following schedulability conditions for SP schedulers:

- (EC) Exact schedulability conditions given by Theorem 2.
- (SC1) Sufficient schedulability conditions given by Theorem 3.
- (SC2) Sufficient schedulability conditions given by equation (14).
- (SC3) Sufficient schedulability conditions given by Theorem 4.

Since we showed that condition (SC1) and condition (SC2) are identical to condition (EC) for the continuous traffic specification, we consider only the discrete traffic specification in this example. We consider the same set of parameters as in Example 1, except that we restrict the period T_1 of the low delay group such that it varies between 1–1.5 ms, i.e. the maximum average data rate varies between 6.7–10 Mbps.

In Figures 5(a)–5(d) we show the schedulability graphs obtained for the sufficient conditions of the SP scheduler. To better illustrate the tradeoff between the schedulability conditions, the period T_1 is altered so that the maximum average data rate of the low delay group varies between 6.7 and 10 Mbps. Figure 5(a) depicts the graph for the necessary and sufficient conditions (EC). Observe that the schedulability graph corresponding to the sufficient condition (SC1) in Figure 5(b) is similar to the schedulability graph of the exact condition in Figure 5(a). For our example the more computationally intensive exact conditions do not provide a significant increase in efficiency over that of (SC1).

Figure 5(c) and 5(d) illustrate the schedulability regions for (SC2) and (SC3), respectively. For our choice of parameters, these two conditions are clearly less efficient than the tight condition presented in Figure 5(a). Note that the efficiency corresponding to (SC2) is low as compared to the tight condition (EC) even though condition (SC2) is exact for continuous traffic specifications.

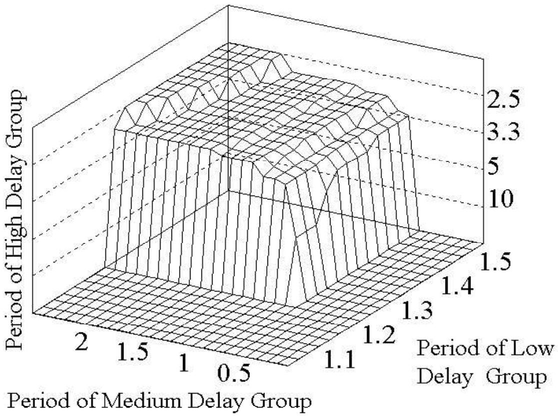

(a) Schedulability Condition (EC).

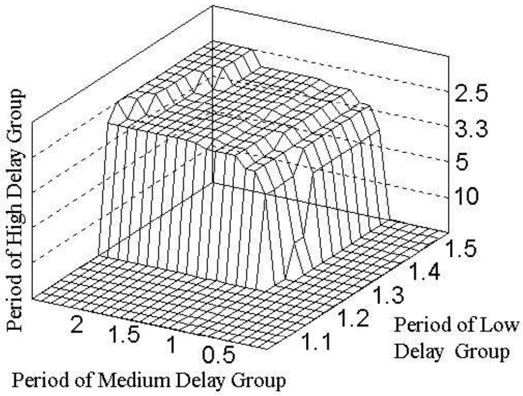

(b) Schedulability Condition (SC1).

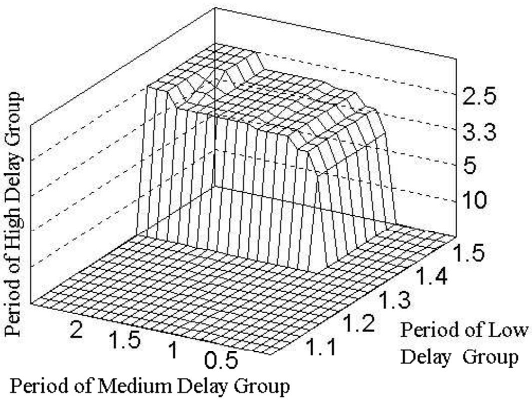

(c) Schedulability Condition (SC2).

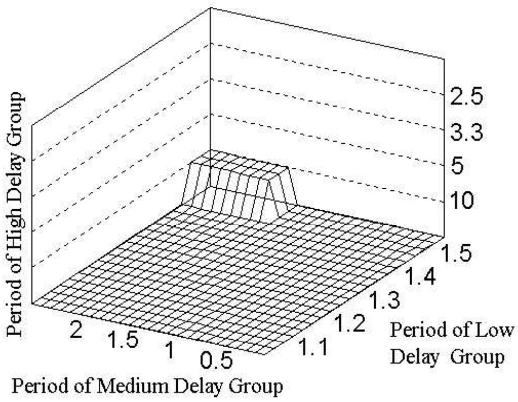

(d) Schedulability Condition (SC3).

Figure 5: Sufficient conditions for an SP Scheduler (time values in milliseconds).

8 Conclusions

We have studied admission control functions in connection-oriented packet-switching networks that offer a bounded delay service, that is, a service that provides deterministically bounded network delays for all packets on a connection. Admission control for a bounded delay service requires the knowledge of the so-called schedulability conditions, i.e., conditions which detect delay bound violations at network switches. Before a new connection can be established, the admission control functions must verify that the schedulability conditions are satisfied at all network switches. The number of connections that are accepted by the admission control functions was found to heavily depend on the traffic specifications used to describe the maximum traffic of a connection, the scheduling disciplines at the network switches, and the accuracy of the schedulability conditions.

We considered networks where packet transmission is based on the Earliest-Deadline-First (EDF) or the Static-Priority (SP) algorithm, and proved exact schedulability conditions using a general traffic model that can represent a large class of both continuous and discrete traffic specifications. We also presented sufficient schedulability conditions for SP schedulers. Using the schedulability conditions we presented examples that illustrated the tradeoffs in a network with a bounded delay service. We showed that the efficiency achieved by relatively complex EDF packet schedulers is significantly superior to the efficiency of low-complexity SP schedulers. By employing various schedulability conditions with a variable degree of complexity and accuracy, we quantified the loss of efficiency in a bounded delay service if sufficient schedulability conditions are used. We also showed that by selecting continuous traffic specifications, which reduce the complexity of the admission control functions, and hence reduce the time cost of connection establishment, the efficiency of a bounded delay services can be severely decreased.

References

[1] K. Bala, I. Cidon, and K. Sohraby. Congestion Control for High Speed Packet Switched Networks. In *Proc. IEEE Infocom '90*, pages 520–526, May 1990.

[2] C.-S. Chang. Stability, Queue Length and Delay, Part I: Deterministic Queueing Networks. Technical Report RC 17708, IBM Research Division, Yorktown Heights, February 1992.

[3] D. D. Clark, S. Shenker, and L. Zhang. Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanisms. In *Proc. Sigcomm '92*, pages 14–26, August 1992.

[4] R. L. Cruz. A Calculus for Network Delay, Part I: Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[5] A. Demers, S. Keshav, and S. Shenker. Analysis and Simulation of a Fair Queueing Algorithm. In *Proc. Sigcomm '89*, pages 1–12, 1989.

[6] D. Ferrari. Real-Time Communication in an Internetwork. *Journal of High-Speed Networks*, 1(1):79–103, 1992.

[7] D. Ferrari and D. C. Verma. A Scheme for Real-Time Channel Establishment in Wide-Area Networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.

[8] S. J. Golestani. A Stop-and-Go Queueing Framework for Congestion Management. In *ACM Sigcomm '90*, pages 8–18, September 1990.

[9] S. J. Golestani. A Framing Strategy for Congestion Management. *IEEE Journal on Selected Areas In Communications*, 9(7):1064–1077, September 1991.

[10] D.P. Hong and T. Suda. Congestion Control and Prevention in ATM Networks. *IEEE Network*, 5(7):10–16, July 1991.

[11] J. M. Hyman, A. A. Lazar, and G. Pacifici. Joint Scheduling and Admission Control for ATS-based Switching Nodes. In *Proc. Sigcomm '92*, pages 223–234, August 1992.

[12] C. R. Kalmanek, H. Kanakia, and S. Keshav. Rate Controlled Servers for Very High-Speed Networks. In *Proc. Globecom '90*, pages 300.3.1–300.3.9, December 1990.

[13] C. L. Liu and J. W. Layland. Scheduling Algorithms for Multiprogramming in a Hard Real Time Environment. *Journal of the ACM*, 20(1):46–61, January 1973.

[14] A. K. J. Parekh. *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks*. PhD thesis, Massachusetts Institute of Technology, February 1992.

[15] J. A. Stankovic and K. Ramamritham (editors). *Hard Real-Time Systems*. IEEE Computer Society Press, 1988.

[16] L. Takács. *Introduction to the Theory of Queues*. Oxford University Press, 1962.

[17] J. S. Turner. New Directions in Communications (or Which Way to the Information Age?). *IEEE Communications Magazine*, 25(8):8–15, October 1986.

[18] D. Verma, H. Zhang, and D. Ferrari. Guaranteeing Delay Jitter Bounds in Packet Switching Networks. In *Proc. Tricomm '91*, Chapel Hill, North Carolina, April 1991.

[19] H. Zhang. *Service Disciplines for Packet-Switching Integrated-Services Networks*. PhD thesis, University of California - Berkeley, November 1993.

[20] H. Zhang and D. Ferrari. Rate-Controlled Static-Priority Queueing. In *Proc. IEEE Infocom '93*, pages 227–236, April 1993.

[21] Q. Z. Zheng and K. G. Shin. On the Ability of Establishing Real-Time Channels in Point-to-Point Packet Switched Networks. to appear: IEEE Transactions on Communications.

A Proof of Theorem 1

The proof of Theorem 1 proceeds in four steps. First we derive an expression for the traffic that is transmitted before an arbitrary packet. Using this expression, we will show sufficiency and necessity of the conditions in (7) and (8) in the second and the third step, respectively. Then we will show that condition (7) is satisfied for all $t \geq 0$ if the condition holds for all $t \leq B_1$.

(a) Workload served before an arbitrary packet

We will derive the workload transmitted before a tagged packet from connection $k \in \mathcal{N}$ that arrives at the EDF-scheduler at time t and is completely transmitted at time $t+\delta$. We use $W^{\leq x}(y)$ to denote the workload in the scheduler at time y due to packets with deadlines less than or equal to x, and $W^{k,t}(t+\tau)$ $(0 \leq \tau \leq \delta)$ to denote the workload in the scheduler at time $t+\tau$ that is served before the tagged packet from connection k with arrival time t.

Let $t-\hat{\tau}$ $(\hat{\tau} \geq 0)$ be the last time before t when the scheduler does not contain traffic with a deadline less than or equal to the deadline of the tagged packet. Since the scheduler is empty before time 0, the time $t-\hat{\tau}$ is guaranteed to exist. Using the definition of $W^{\leq x}(y)$, $\hat{\tau}$ is given by:

$$\hat{\tau} = \min\{z \mid W^{\leq t+d_k}(t-z) = 0, z \geq 0\} \tag{21}$$

Hence, in time interval $[t-\hat{\tau}, t+\delta)$ the scheduler always contains work with a deadline less than or equal to $t+\delta$. With $\hat{\tau}$ we can determine $W^{k,t}(t+\tau)$, the workload that is transmitted before the tagged packet. $W^{k,t}(t+\tau)$ is composed of:

- The remaining transmission time of the packet that is in transmission at time $t-\hat{\tau}$, denoted by $R(t-\hat{\tau})$. With equation (21), this packet has a deadline greater than $t+d_k$. (For continuous traffic specifications $R(t-\hat{\tau})$ denotes the remaining time until the current transmission of traffic can be interrupted).
- $A_j^{\leq t+d_k}[t-\hat{\tau}, t+\tau]$, that is, all arrivals from connection j in time interval $[t-\hat{\tau}, t+\tau]$ with deadlines less than or equal to $t+d_k$. Note that $A_k^{\leq t+d_k}[t-\hat{\tau}, t+\tau]$ includes the tagged packet.

From equation (21) we obtain that in time interval $[t-\hat{\tau}+R(t-\hat{\tau}), t+\tau]$, the EDF-scheduler only transmits traffic with a deadline less than or equal to $t+d_k$. Therefore, we obtain the following expression for $W^{k,t}(t+\tau)$ with $0 \leq \tau \leq \delta$:

$$W^{k,t}(t+\tau) = \sum_{j \in \mathcal{N}} A_j^{\leq t+d_k}[t-\hat{\tau}, t+\tau] + R(t-\hat{\tau}) - (\hat{\tau}+\tau) \tag{22}$$

Since all traffic from a connection j that arrives after time $t+d_k-d_j$ has a deadline greater than $t+d_j$, we can rewrite (22) as:

$$W^{k,t}(t+\tau) = \sum_{j \in \mathcal{N}} A_j[t-\hat{\tau}, \min\{t+\tau, t+d_k-d_j\}] + R(t-\hat{\tau}) - (\hat{\tau}+\tau) \tag{23}$$

(b) Proof of Sufficiency

Consider the tagged packet from connection k that arrives at time t. The packet does not have a deadline violation if there exists a τ $(0 \le \tau \le d_k)$ such that

$$W^{k,t}(t+\tau) = 0 \tag{24}$$

where $W^{k,t}(t+\tau)$ is as given in equation (23). For the delay of the tagged packet we must distinguish whether at time $t - \hat{\tau}$, the scheduler is empty or is transmitting a packet.

<u>Case 1:</u> $W(t - \hat{\tau}) = 0$
In this case, the scheduler is empty at time $t - \hat{\tau}$. Consequently, $R(t - \hat{\tau}) = 0$. For $\tau = d_k$ we obtain from equation (23):

$$\begin{aligned} W^{k,t}(t + d_k) &= \sum_{j \in \mathcal{N}} A_j[t - \hat{\tau}, t + d_k - d_j] - (\hat{\tau} + d_k) && (25) \\ &\le \sum_{j \in \mathcal{N}} A_j^*(\hat{\tau} + d_k - d_j) - (\hat{\tau} + d_k) && (26) \end{aligned}$$

Equation (26) follows from equation (25) with the property of A_j^* from equation (1). With equation (7) we have:

$$W^{k,t}(t + d_k) \le 0 \tag{27}$$

Hence, there exists a $\tau \le t + d_k$ such that $W^{k,t}(t+\tau) = 0$.

<u>Case 2:</u> $W(t - \tau) > 0$
The scheduler is transmitting traffic at time $t - \hat{\tau}$ from some connection i. Due to equation (21), the traffic in transmission has a deadline greater than $t + d_k$, that is,

$$d_i > \hat{\tau} + d_k \tag{28}$$

Without loss of generality we assume that i is such that $s_i = \max_{d_j > \hat{\tau} + d_k} s_j$. For $\tau = d_k$ we obtain for equation (23):

$$W^{k,t}(t + d_k) = \sum_{j \in \mathcal{N}} A_j[t - \hat{\tau}, t + d_k - d_j] + R(t - \hat{\tau}) - (\hat{\tau} + d_k) \tag{29}$$

Since the remaining transmission time at time $t - \hat{\tau}$ is at most s_i, we obtain with the definition of A_j^*:

$$W^{k,t}(t + d_k) \le \sum_{j \in \mathcal{N}} A_j^*(\hat{\tau} + d_k - d_j) + \max_{d_j > \hat{\tau} + d_k} s_j - (\hat{\tau} + d_k) \tag{30}$$

With equation (28), we obtain for equation (30) that:

$$W^{k,t}(t + d_k) \le 0 \tag{31}$$

With the definition of $W^{k,t}$ from equation (23), there exists a $\tau \le d_k$ such that $W^{k,t}(t+\tau) = 0$.

(c) Proof of Necessity

Assume that the inequality in (7) is violated at time $t > 0$, that is:

$$t < \sum_{j \in \mathcal{N}} A_j^*(t - d_j) \tag{32}$$

Now consider the following scenario. The scheduler is empty for all time less than 0, and starting at time 0 all connections j submit traffic at their maximum rate, that is, according to $\{A_j^*\}_{j \in \mathcal{N}}$. With $A_j^{\leq t}(t) = A_j(t - d_j)$, we obtain that the workload in the scheduler at time t due to traffic with deadlines less than or equal to t, $W^{\leq t}(t)$, is given by:

$$W^{\leq t}(t) = \sum_{j \in \mathcal{N}} A_j^*(t - d_j) - t \tag{33}$$

With the assumption from (32) we obtain $W^{\leq t}(t) > 0$, that is, at time t the scheduler contains traffic that has deadlines less than or equal to t. Therefore, there is at least one packet in the scheduler at time t with a deadline violation.

Suppose that inequality (8) is violated, that is, there exists a connection $i \in \mathcal{N}$ such that at a time t with $d_1 \leq t \leq d_{|\mathcal{N}|}$ the following holds:

$$t < \sum_{j \in \mathcal{N}} A_j^*(t - d_j) + \max_{d_j > t} s_j \tag{34}$$

Now assume a scenario where the scheduler is empty for all times less than[6] 0^-, and at time 0^- a packet from connection i with $s_i = \max_{d_j > t} s_j$ arrives with a transmission time of s_i, and at time 0 packets from connections j ($j < i$) arrive according to $\{A_j^*\}_{j<i}$. Since the EDF-scheduler is non-preemptive, the packet from connection i is transmitted before the packets from connections j ($j < i$) with a smaller deadline. Then $W^{\leq t}(t)$ is given by:

$$W^{\leq t}(t) = \sum_{j \in \mathcal{N}} A_j^*(t - d_j) - t + \max_{d_j > t} s_j \tag{35}$$

With the assumption from (34) we have $W^{\leq t}(t) > 0$. Thus, there must be a packet in the scheduler at time t with a deadline violation.

(d) If condition (7) holds for all $t \leq B_1$ then it holds for all $t \geq 0$.

Clearly, condition (7) cannot be violated outside a busy period. Now, consider a busy period $[t_1, t_2]$ with $t_1 \geq B_1$. Due to the property of A_j^* given in equation (1) we have that:

$$t_2 - t_1 \leq B_1 \tag{36}$$

[6] Time t^- denotes the time instant immediately before time t.

Assume that condition (7) is violated at time t with $t_1 \leq t \leq t_2$. Thus,

$$t < \sum_{j\in\mathcal{N}} A_j^*(t-d_j) < \sum_{j\in\mathcal{N}} A_j^*(t_1) + \sum_{j\in\mathcal{N}} A_j^*(t_1, t-d_j] \tag{37}$$

Since t_1 is the beginning of a busy period we have that $t_1 > \sum_{j\in\mathcal{N}} A_j^*(t_1)$. With equation (1) we obtain from (37) that:

$$t - t_1 < \sum_{j\in\mathcal{N}} A_j^*(t_1, t-d_j] \leq \sum_{j\in\mathcal{N}} A_j^*(t-d_j-t_1) \tag{38}$$

From the definition of A_j^* we have that $t - t_1 \leq B_1$, and we have found a deadline violation in $[0, B_1]$. □

B Proof of Theorem 2

Similar to the proof of Theorem 1, we proceed in four steps. First we obtain an expression for $W^{p,t}(t+\tau)$, the workload in the scheduler at time $t+\tau$ that is served before a packet from priority p that arrived at time t. Then we prove sufficiency and necessity of the schedulability condition. Finally, we show that for each priority p condition (12) holds for all $t \geq 0$ if condition (12) holds for $t \leq B_1^p - d_p$.

(a) Workload served before an arbitrary packet

Assume that a packet (the *tagged packet*) from a connection $k \in \mathcal{C}_p$ arrives at the scheduler at time t with a transmission time of $s \leq s_p^*$, and that its transmission begins at time $t+\delta$. The arrival time t falls into a priority-p busy period which started at time $t - \hat{\tau}$, i.e.,

$$\hat{\tau} = \min\{z \mid \sum_{q=1}^{p} \sum_{j\in\mathcal{C}_q} W_j(t-z) = 0, z \geq 0\} \tag{39}$$

Denoting by $W^{p,t}(t+\tau)$ $(0 \leq \tau \leq \delta)$ the workload in the SP-scheduler at time $t+\tau$ that is served *before* the tagged packet, $W^{p,t}(t+\tau)$ is determined for $t \leq t+\tau \leq t+\delta$ by:

- $R(t-\hat{\tau})$, the remaining transmission time of a priority-r packet with $r > p$ that is in transmission at time $t - \hat{\tau}$ (For continuous traffic specifications $R(t-\hat{\tau})$ denotes the remaining time until the current transmission of traffic can be interrupted.)

- Traffic from priority-p connections that arrives in time interval $[t-\hat{\tau}, t]$, i.e., before or with the arrival of the tagged packet, however, not including the tagged packet. This traffic is given by $A_j[t-\hat{\tau}, t+\tau] - s$ for $j \in \mathcal{C}_p$.

- Traffic from higher priority connections that arrives in time interval $[t-\hat{\tau}, t+\tau)$, given by $A_j[t-\hat{\tau}, t+\tau)$ for for $j \in \mathcal{C}_q$ and $q < p$.

- the length of time interval $[t-\hat{\tau}, t+\tau]$.

Formally, $W^{p,t}(t+\tau)$ is given for all $0 \leq \tau \leq \delta$ by:

$$W^{p,t}(t+\tau) = \sum_{j \in \mathcal{C}_p} A_j[t-\hat{\tau}, t] + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j[t-\hat{\tau}, t+\tau) - s + R(t-\hat{\tau}) - (\hat{\tau}+\tau) \tag{40}$$

Since the transmission of the tagged priority-p packet begins at time $t+\delta$, there cannot be any workload left that is to be served before the tagged packet. More precisely:

$$\delta = \min\{z \mid W^{p,t}(t+z) = s, z \geq 0\} \tag{41}$$

(b) Proof of Sufficiency

We will show that, for an arbitrary packet on connection k (with $k \in \mathcal{C}_p$) with transmission time $s \leq s_k^*$ that arrives at time t, condition (12) guarantees that the packet will depart before $t+d_p$. From the definition of A_j^* in (1) we have:

$$\sum_{j \in \mathcal{C}_p} A_j[t-\hat{\tau}, t] \leq \sum_{j \in \mathcal{C}_p} A_j^*(\tau) \tag{42}$$

$$\sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j[t-\hat{\tau}, t+\tau) \leq \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(\hat{\tau}+\tau^-) \tag{43}$$

Since the remaining nonpreemptable transmission time of priority-r traffic ($r > p$) at time $t-\hat{\tau}$ is maximal if low priority traffic with maximum transmission time[7] starts transmission at $t-\tau^-$, we obtain:

$$R(t-\hat{\tau}) \leq \max_{r>p} s_r \tag{44}$$

With equations (42)–(44), we can give the following bound for $W^{p,t}(t+\tau)$ in equation (40):

$$W^{p,t}(t+\tau) \leq \sum_{j \in \mathcal{C}_p} A_j^*(\tau) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(\tau+\tau^-) - s + \max_{r>p} s_r - (\hat{\tau}+\tau) \tag{45}$$

Since the transmission time of the tagged packet is given by $s \leq s_p^*$, we obtain with condition (12) that there exists a τ' with $0 \leq \tau' \leq d_p - s_k^*$ such that $W^{p,t}(t+\tau') \leq 0$. Hence, there exists a $\tau'' \leq d_p - s_k^*$ such that $W^{p,t}(t+\tau'') = s_p^*$. Therefore, with (41) the tagged packet begins transmission at time $t+\tau''$. Since the packet has at most a transmission time of s_k^*, the tagged packet does not cause a deadline violation.

[7] For discrete traffic specifications, the traffic will result from a single packet.

(c) Proof of Necessity

Let us assume that the condition in equation (12) does not hold, that is, there exists a priority p and a time interval $[t, t+d_p - s_p^*]$ within a priority-p busy period such that, for all $0 \le \tau \le d_p - s_p^*$:

$$t + \tau < \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t+\tau^-) - s_p^* + \max_{r>p} s_r \tag{46}$$

Now assume a scenario where the SP-scheduler is empty before time 0^-, and at time 0^- traffic from connection $i \in \mathcal{C}_r$ with $s_i = \max_{r>p} s_r$ arrives. Suppose that, starting at time 0, all connections j with priorities p or higher transmit the maximum traffic permitted by their rate-controlling functions A_j^*, with one exception: the last packet arrival before t from a connection k with $k \in \mathcal{C}_p$ and $s_k^* = s_p^*$ is delayed until time t. In other words, if the last packet arrival from connection $k \in \mathcal{C}_p$ before t occurs at time $t - z$ where

$$z = \min\{z' \mid A_k^*(t - z') < A_k^*(t), z' \ge 0\}, \tag{47}$$

then the packet arrival is delayed until time t. The delayed packet is assumed to have a transmission time of s_k^*.

If the delayed packet from connection $k \in \mathcal{C}_p$ with arrival time t has not started transmission at time $t+\tau$, then the traffic that arrives in time interval $[0^-, t+\tau]$ and is transmitted before the delayed packet consists at least of:

- $s_i = \max_{r>p} s_r$, the transmission time of traffic that arrived at time 0^-
- $A_j^*(t) - s_p^*$ with $j \in \mathcal{C}_p$, the traffic from priority p that arrived in time interval $[0, t]$ excluding the packet with arrival time t.
- $A_j^*(t+\tau)$ with $j \in \mathcal{C}_q$ and $q < p$, the high priority traffic which arrives in time interval $[0, t+\tau)$.

Here we do not assume that $t \le B_1^p - d_p$, that is, in time interval $[0, t]$, the SP-scheduler could be empty or transmit traffic from priorities $q > p$. However, in the best case, the SP-scheduler is always transmitting traffic in time interval $[0^-, t+\tau]$. Hence, we obtain the following lower bound for $W^{p,t}(t+\tau)$, the workload that is transmitted before the delayed packet:

$$W^{p,t}(t+\tau) > \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t+\tau^-) + \max_{r>p} s_r - (t+\tau) \tag{48}$$

With our assumption in (46) we obtain that $W^{p,t}(t) > 0$ in the entire time interval $[t, t+d_p - s_p^*]$. Thus, if the packet from connection k that arrives at time t has a transmission time of s_p^* a deadline violation occurs for this packet at $[t, t+d_p - s_p^*]$.

(d) If condition (12) holds for all $t \leq B_1^p - d_p$ then it holds for all $t \geq 0$.

The proof is similar to the corresponding proof for EDF-schedulers. We consider a priority-p busy period $[t_1, t_2]$ with $t_1 \geq B_1^p$. Assume that condition (12) is violated for priority p, with fixed values for t and τ such that $t_1 \leq t \leq t + \tau \leq t_2 - d_p$ and $0 \leq \tau \leq d_p - s_p^*$. Then the following holds:

$$t + \tau < \sum_{j \in \mathcal{C}_p} A_j^*(t) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t + \tau^-) - s_p^* + \max_{r>p} s_j \tag{49}$$

Equation (49) is equivalent to:

$$t + \tau < \sum_{j \in \mathcal{C}_p} A_j^*(t_1^-) + \sum_{j \in \mathcal{C}_p} A_j^*(t_1, t] + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t_1^-) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t_1, t + \tau^-] - s_p^* + \max_{r>p} s_j \tag{50}$$

Since t_1 is the beginning of a priority-p busy period, we have that:

$$t_1 > \sum_{j \in \mathcal{C}_p} A_j^*(t_1^-) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t_1^-) \tag{51}$$

With (50) and the properties of A_j^* from equation (1) we can rewrite the inequality in (50) as:

$$t + \tau - t_1 < \sum_{j \in \mathcal{C}_p} A_j^*(t - t_1) + \sum_{q=1}^{p-1} \sum_{j \in \mathcal{C}_q} A_j^*(t + \tau - t_1^-) - s_p^* + \max_{r>p} s_r \tag{52}$$

Using again the properties of A_j^*, we obtain that $t - t_1 < B_1^p$. Hence, we have obtained a violation of condition (12) in time interval $[0, B_1^p - d_p]$. $\square$