CCIT Report #393

July 2002

QoS Provisioning and Routing with EDF Scheduling, Stochastic Burstiness and Stochastic Guarantees

Erez Biton and Ariel Orda
Department of Electrical Engineering
Technion - Israel Institute of Technology
Haifa 32000, Israel
E-mail: {berez@tx,ariel@ee}.technion.ac.il

July 24, 2002

Abstract

A scalable approach to end-to-end QoS provisioning requires to handle traffic aggregates. This translates into a looser characterization of traffic profiles. Scalability of QoS provisioning is further enhanced by the elasticity of typical (QoS demanding) applications, which translates into looser QoS requirements. Accordingly, this study considers QoS provision schemes for connections with *stochastic* traffic profiles and *stochastic* QoS requirements. We concentrate on the class of Rate-Controlled Earliest Deadline First (RC-EDF) scheduling disciplines, which have several well known advantages, in particular simplicity of implementation and flexibility. Assuming the Exponentially Bounded Burstiness (EBB) traffic model, we establish results that extend the deterministic study of RC-EDF, both for a single server in isolation and for networks of servers. For a single traffic shaper followed by an EDF scheduler, we establish stochastic bounds on the distribution of the delay for each session. In the general (multi-hop) setting, we first establish stochastic bounds on the distribution of the end-to-end delay for traffic shaper elements in series; then, we establish stochastic bounds for RC-EDF networks. Consequently, we formulate call admission control and routing schemes that identify feasible paths under various network optimization criteria. Here, we consider two settings: a deterministic setting, of Burstiness Constrained traffic (BC), and a stochastic setting of Exponentially Bounded Burstiness.

1 Introduction

Emerging Broadband high speed networks are expected to provide real time and multimedia applications with various Quality of Service (QoS) guarantees. However, providing such guarantees on a per-flow basis is typically unscalable. Thus, a scalable approach to end-to-end QoS provisioning, which would handle traffic aggregates, is called for. Such aggregates translate into a looser characterization of traffic profiles. Scalability of QoS provisioning is further enhanced by the elasticity of typical (QoS demanding) applications, which translates into looser (stochastic rather than deterministic) QoS requirements.

One of the major problems in the provision of QoS guarantees is identifying a feasible path that can meet the QoS requirements. Schedulability conditions, as well as worst case bounds on the delay, constitute valuable tools for quantifying the ability of a path to meet the QoS requirements.

The ability to support QoS requirements depends on the scheduling policies employed in the nodes. In this paper, we consider the Rate-Controlled Earliest-Deadline-First (RC-EDF) scheduling discipline [1]. In the rate-controlled class of service disciplines, the traffic of each connection is reshaped at every node to ensure that the traffic offered to the scheduler conforms to specific characteristics. In particular, it is typically used to enforce, at an internal network node, the same traffic parameters as at the network access point. Reshaping makes the traffic at each node more predictable and, therefore, simplifies the task of guaranteeing performance to individual connections. When used with a particular scheduling policy, it allows the specification of worst case delay bounds at each node. End-to-end bounds can then be computed as the sum of worst-case delay bounds at each node along the path. Some important advantages of a rate-controlled service discipline, especially when compared with Generalized Processor Sharing (GPS) [2], are simplicity of implementation and flexibility. Thus, a rate-controlled service discipline is often a better solution in terms of scalability.

The EDF scheduling policy associates a per-hop deadline with each packet and schedules packets in the order of their assigned deadlines. Exact schedulability conditions, which detect violations of delay guarantees in an EDF network switch, have been established in [3] under a *deterministic* setting. In [4], EDF has been proven to be an optimal scheduling discipline; that is, if a set of tasks is schedulable under any scheduling discipline, then this set is also schedulable under EDF. Also, RC-EDF was proven to outperform GPS in providing end-to-end delay guarantees in a network [5].

Under a deterministic setting, the input traffic bursts are assumed to be of bounded length. This is not the case in many commonly used input processes, in particular traffic aggregates. Hence, a setting that considers the stochastic nature of the traffic is desired. Under a stochastic setting, only stochastic QoS is guaranteed, *i.e.*, it is guaranteed that the end-to-end delay experienced by a high percentage of the packets does not significantly exceed the required delay. Such guarantees are appropriate for many applications, in particular multimedia applications, which can tolerate a certain amount of loss due to either late arrival or buffer overflow. Furthermore, schemes that guarantee no loss have a low connection-carrying capacity for bursty traffic. In other words, with stochastic guarantees, better network utilization can be achieved. Consequently, several studies have investigated the provision of *stochastic* QoS guarantees for *stochastic* traffic profiles, e.g., [6, 7, 8]; however, these studies were carried only in the context of the GPS scheduling discipline.

Due to the well known practical advantages of the RC-EDF service discipline, in this study we investigate its stochastic behavior using Exponentially Bounded Burstiness (EBB) processes [9] as source session traffic models. First, we study the single node case. We derive schedulability conditions for the EDF scheduling discipline as well as stochastic bounds on the delay experienced by a packet entering an EDF scheduling element. Furthermore, we introduce the concept of *EBB traffic shapers*, and derive stochastic bounds on the delay experienced by any packet entering such a traffic shaper. Next, we study the multiple node (*i.e.*, multi-hop path) case, and derive a stochastic bound on the end-to-end delay. This bound, for a network of EDF schedulers, packetized EBB traffic and stochastic QoS guarantees, is the main contribution of this study.

Some previous studies also considered QoS provisioning under a stochastic setting [10, 11, 12]. Our stochastic model and framework is different. In our case, the stochasticity of the end-to-end guarantees is (solely) due to the stochastic nature of the session input traffic, whereas in [10, 11, 12], the exploitation of statistical multiplexing results in a provision of (only) stochastic guarantees even for deterministically bounded input traffic.

Finally, with the stochastic end-to-end bounds at hand, we study the related call admission control and routing problems. Such problems, for EDF schedulers in a *deterministic* setting, have been studied in [13, 14] and several routing schemes have been proposed. However, the focus in those studies was on identifying feasible, but not necessarily (networkwide) efficient paths. Considering our framework of stochastic traffic profiles and stochastic QoS guarantees, we propose routing and call admission schemes, which aim at balancing loads and maximizing the ability to accommodate future calls.

The rest of the paper is structured as follows. In Section 2, we formulate the model. Next, in Section 3, we study the provisioning of QoS with EDF scheduling and stochastic guarantees: first, we consider a service element in isolation; then, we study the multiple node case and derive a stochastic bound on the end-to-end delay. In section 4 we discuss and present call admission and routing schemes; here, we consider both the deterministic and stochastic settings. In Section 5, we conclude the paper and discuss possible future work.

2 Model Formulation

We consider a store-and-forward network comprising of packet switches in which a packet scheduler is available at each output link. Packetized traffic from a particular connection entering the switch passes through a traffic shaper before being delivered to the scheduler. The traffic shaper regulates traffic, so that the output of the shaper satisfies certain pre-specified traffic characteristics. We focus on the Earliest Deadline First (EDF) scheduling discipline. The EDF scheduler associates a deadline $\bar{d}^j$ with each packet of a session j. The packets are served in the order of their assigned deadlines. Let L^j be the maximal packet size of a session j and $L_{\max}$ be the maximal packet size in the network.

The network is represented by a directed Graph $G(V, E)$, in which nodes represent switches and arcs represent links. V is the set of nodes and E is the set of links interconnecting them, and let $|V| = N$ and $|E| = M$. Each link $l \in E$ is characterized by a service rate r_l. We denote the number of sessions entering link l by $\mathcal{N}_l$. We assume that link propagation delays are negligible.

Following [15, 16, 17], we assume a source ("explicit") QoS routing framework, in which link state information is exchanged and maintained up-to-date among network nodes for path computation. Routing decisions are based on the image of the network at the source node.

A session j is routed through a path $\mathbf{p}^j$. Let $n\left(\mathbf{p}^j\right)$ be the total number of hops along $\mathbf{p}^j$. We denote by H the maximal possible number of hops along a path.

We consider stochastic traffic profiles, and adopt the traffic model introduced in [9], of Exponentially Bounded Burstiness (EBB) processes, defined as follows.

A stochastic process $A(t)$ is *Exponentially Bounded (EB)* with parameters (Λ, α) if, for any t and any $\sigma \geq 0$, the following bound applies:

$$\Pr\{A(t) \geq \sigma\} \leq \Lambda \cdot e^{-\alpha \cdot \sigma}. \tag{1}$$

Let $A(t)$ be the instantaneous traffic rate. $A(t)$ has *Exponentially Bounded Burstiness (EBB)* with parameters (ρ, Λ, α), if for any s,τ and any $\sigma \geq 0$, the following upper bound, on the tail distribution of the traffic arriving during the time interval $[s, \tau]$, holds:

$$\Pr\{A[s, \tau] \geq \rho(\tau - s) + \sigma\} \leq \Lambda \cdot e^{-\alpha \cdot \sigma}, \tag{2}$$

where $A[s, \tau] = \int_s^\tau A(t)\, dt$.

We assume a discrete time domain, in which the amount of information transmitted on a link with capacity $r = 1$ during one time slot is regarded as a unit of data. In this context, we have $A[s, \tau] = \sum_{n=s+1}^{\tau} A(n)$.

3 QoS Provisioning with Stochastic Guarantees

In order to handle stochastic traffic profiles, we need to formulate a stochastic version of the class of rate-controlled service disciplines introduced in [1]. We begin with a brief overview of the main results obtained there and in some subsequent studies, for the basic, deterministic setting. In [1], it is assumed that connections whose traffic satisfies certain (deterministic) burstiness constraints enter the network at various nodes. At each node along the path of a connection, traffic is reshaped to confirm to its original envelop before it enters the scheduler. Based on the traffic envelope of the connection, upper bounds on the scheduling delays at each node can be guaranteed. It is also shown in [1] that, for the traffic shapers considered there, reshaping the traffic to its original envelope does not introduce extra delays. Therefore, an upper bound on the end-to-end packet delay is simply the sum of the scheduling and propagation delays. A more general rate-controlled service discipline of arbitrary reshaping at each node was considered in [5]. There, it was shown that any end-to-end delay bounds that can be guaranteed by the GPS discipline, can also be achieved by a rate-controlled service discipline with "proper" reshaping.

In this study, we investigate the rate-controlled service discipline under a stochastic setting; more precisely, we adopt the traffic model introduced in [9], of Exponentially Bounded Burstiness (EBB) processes, to investigate the rate-controlled service discipline under a stochastic setting. With that at hand, we derive stochastic bound on the end-to-end delay in networks that employ the RC-EDF service discipline.

3.1 The single node case

Here, we consider a service element in isolation, which consists of an EBB traffic shaper followed by an EDF scheduler, as depicted in Fig. 1.

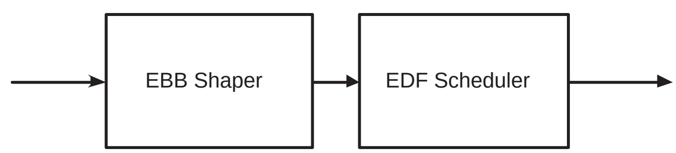

Figure 1: EDF scheduler

3.1.1 EDF Scheduler

Consider a set $\mathcal{N}$ of connections where each connection's traffic rate is $\left\{A^j(t)\right\}_{j\in\mathcal{N}}$. Each connection j requires stochastic delay guarantees as follows:

$$\Pr\left\{D^j(t) \geq \bar{d}^j + d\right\} \leq f_j(d).$$

We define the stochastic schedulability of a set of connections entering an EDF server as follows.

Definition 1 *Given are a scheduler and a set of $\mathcal{N}$ connections, where each connection $j \in \mathcal{N}$ is characterized by $\left(A^j, \bar{d}^j\right)$. The set of connections is said to be* EDF-schedulable *if, for all $t > 0$, stochastic delay guarantees are provided for each connection,* i.e., $\Pr\left\{D^j(t) \geq \bar{d}^j + d\right\} \leq f^j(d) \;\; \forall j \in \mathcal{N}$.

Proposition 1 *A set $\mathcal{N}$ of connections is EDF-schedulable if for all $k \in \mathcal{N}$ and for all t:*

$$\sum_{\hat{\tau}=0}^{t} \Pr\left\{ \sum_{j\in\mathcal{N}} A^j\left[t-\hat{\tau}, t+\bar{d}^k-\bar{d}^j\right] + \max_{\bar{d}^j > \hat{\tau}+\bar{d}^k} L^j > r\left(\hat{\tau}+\bar{d}^k+d\right)\right\} \leq f^k(d) \tag{3}$$

Proof: See Appendix A.

With this schedulability conditions at hand, we derive an upper-bound on the delay tail distribution for each session's traffic entering a single EDF server. We show that the delay experienced by a packet entering the EDF scheduler is exponentially bounded (*i.e.*, EB).

Proposition 2 *Suppose that $\left\{A^j\right\}_{j\in\mathcal{N}}$ are $|\mathcal{N}|$ independent $\left(\rho^j, \Lambda^j, \alpha^j\right)$-EBB processes sharing an EDF server with delay assignment $\left\{\bar{d}^j\right\}_{j\in\mathcal{N}}$. Then, at any time t, for any $d > \bar{d}^k$ and for all $k \in \mathcal{N}$,*

$$\Pr\left\{D^k(t) \geq d\right\} \leq \tilde{\Lambda} e^{-\tilde{\alpha} r d} \tag{4}$$

where $\frac{1}{\tilde{\alpha}} = \sum_{j\in\mathcal{N}} \frac{1}{\alpha^j}$, $\tilde{\Gamma} = \dfrac{\left(\sum_{j\in\mathcal{N}} \Lambda^j\right) e^{-\tilde{\alpha}\left(\sum_{j\in\mathcal{N}} \rho^j \bar{d}^j - L_{\max}\right)}}{1-e^{-\tilde{\alpha}\left(r-\sum_{j\in\mathcal{N}} \rho^j\right)}}$, *and* $\tilde{\Lambda} = \tilde{\Gamma} e^{\tilde{\alpha}\left(\sum_{j\in\mathcal{N}, j\neq k} \rho^j\right)\bar{d}^k}$.

Proof: From the proof of Proposition 1 (expression (A.33)), we have

$$\Pr\left\{D^k(t) \geq \bar{d}^k + d\right\} \leq \sum_{\hat{\tau}=0}^{t} \Pr\left\{\sum_{j\in\mathcal{N}} A^j\left[t-\hat{\tau}, t+\bar{d}^k-\bar{d}^j\right] + \max_{\bar{d}^j>\hat{\tau}+\bar{d}^k} L^j > r\left(\hat{\tau}+\bar{d}^k+d\right)\right\}. \tag{5}$$

Let $p_1, p_2, \ldots, p_{|\mathcal{N}|}$ be positive constants that sum to 1. Then,

$$\left\{\sum_{j\in\mathcal{N}} A^j\left[t-\hat{\tau}, t+\bar{d}^k-\bar{d}^j\right] + L_{\max} \geq r\left(\hat{\tau}+\bar{d}^k+d\right)\right\} \subset \bigcup_{j\in\mathcal{N}} \left\{A^j\left[t-\hat{\tau}, t+\bar{d}^k-\bar{d}^j\right] \geq p_j\left(r\left(\hat{\tau}+\bar{d}^k+d\right) - L_{\max}\right)\right\} \tag{6}$$

Recall that the session's traffic is EBB, *i.e.*,

$$\Pr\left\{A^j\left[t-\hat{\tau}, t+\bar{d}^k-\bar{d}^j\right] \geq \rho^j\left(\hat{\tau}+\bar{d}^k-\bar{d}^j\right) + \sigma^j\right\} \leq \Lambda_j e^{-\alpha^j\sigma^j}. \tag{7}$$

From (5), (6), (7) and the union bound, we get

$$\Pr\left\{D^k(t) \geq \bar{d}^k + d\right\} \leq \sum_{\hat{\tau}=0}^{t}\sum_{j\in\mathcal{N}} \Lambda^j e^{-\alpha^j\left(p_j\left(r\left(\hat{\tau}+\bar{d}^k+d\right)-L_{\max}\right)-\rho^j\left(\hat{\tau}+\bar{d}^k-\bar{d}^j\right)\right)}.$$

Choosing p_j such that

$$\alpha^j\left(p_j\left(r\left(\hat{\tau}+\bar{d}^k+d\right)-L_{\max}\right)-\rho^j\left(\hat{\tau}+\bar{d}^k-\bar{d}^j\right)\right) = C,$$

for a (any) constant C and $\sum_{j\in\mathcal{N}} p_j = 1$, we get

$$\Pr\left\{D^k(t) \geq \bar{d}^k + d\right\} \leq \left(\sum_{j\in\mathcal{N}}\Lambda^j\right)\sum_{\hat{\tau}=0}^{t} e^{-\tilde{\alpha}\left(rd+\sum_{j\in\mathcal{N}}\rho^j\bar{d}^j - L_{\max} + \left(r-\sum_{j\in\mathcal{N}}\rho^j\right)\left(\hat{\tau}+\bar{d}^k\right)\right)}$$

$$\leq \frac{\left(\sum_{j\in\mathcal{N}}\Lambda^j\right)}{1-e^{-\tilde{\alpha}\left(r-\sum_{j\in\mathcal{N}}\rho^j\right)}} e^{-\tilde{\alpha}\left(rd+\sum_{j\in\mathcal{N}}\rho^j\bar{d}^j - L_{\max} + \left(r-\sum_{j\in\mathcal{N}}\rho^j\right)\bar{d}^k\right)}.$$

The proposition follows by taking $d' = \bar{d}^k + d$.

3.1.2 EBB Traffic Shaper

Figure 2: EBB Traffic Shaper

The (ρ, Λ, α)-shaper has one input link and one output link with equal rate, as depicted in Fig. 2. The shaper receives an arbitrary stream on the input link and buffers data, if necessary, so that the output stream transmitted in the output link has EBB with parameters (ρ, Λ, α).

Suppose the rate of the traffic input to the shaper is represented by A_0. Let s_i be the time at which the ith packet starts to arrive on the input link, and L_i is the length in bits of the packet. Suppose that the packet exits on the output link at time t_i and let A_1 represent the rate of traffic exiting the shaper. The shaper transmits packets on the output link in an FCFS order, with the smallest possible delay, such that

$$\Pr\left\{W_\rho\left(A_1\right)(t) \geq \sigma\right\} \leq \Lambda e^{-\alpha\sigma} \tag{8}$$

for all packets.

Proposition 3 *Let the delay suffered by the ith packet, $D_i^{(A_0,A_1)} = t_i - s_i$, be as small as possible subject to the constraint (8). It holds that*

$$\Pr\left\{D_i^{(A_0,A_1)} \leq d\right\} \begin{cases} = 0 & d < t_{i-1} - s_i \\ \leq \Lambda e^{-\alpha(W_\rho(A_0)(s_i) - \rho d)} & t_{i-1} - s_i \leq d < \frac{W_\rho(A_0)(s_i)}{\rho} \\ = 1 & d \geq \frac{W_\rho(A_0)(s_i)}{\rho} \end{cases}, \tag{9}$$

where $W_\rho\left(A_0\right)(s)$ is defined as follows:

$$W_\rho\left(A_0\right)(s) = \max_{t \leq s}\left\{A_0(s,t) - \rho(t-s)\right\}. \tag{10}$$

Proof: In [18] it was shown that: (i) $D_i^{(A_0,A_1)} \geq \frac{1}{\rho}\left(W_\rho\left(A_0\right)\left(s_i\right)-\sigma\right)$ is necessary so that $W_\rho\left(A_1\right)\left(t_j\right) \leq \sigma$ is satisfied, (ii) $D_i^{(A_0,A_1)} = \frac{1}{\rho}\left(W_\rho\left(A_0\right)\left(s_i\right)-\sigma\right)$ implies that $W_\rho\left(A_1\right)\left(t_j\right) = \min\left\{W_\rho\left(A_0\right)\left(s_i\right), \sigma\right\}$. These results were proved in the context of Burstiness constrained traffic. However, it can be verified that the proof in [18] holds for the broader class of stochastic traffic as well.

¿From (ii) above it is easy to see that if $\left\{D_i^{(A_0,A_1)} \leq \frac{1}{\rho}\left(W_\rho\left(A_0\right)\left(s_i\right)-\sigma\right)\right\}$ then $\left\{W_\rho\left(A_1\right)\left(t_i\right) \geq \sigma\right\}$.

According to Theorem 1 in [9], A_1 has EBB with upper rate ρ if

$$\Pr\left\{W_\rho\left(A_1\right)\left(t_i\right) \geq \sigma\right\} \leq \Lambda e^{-\alpha\sigma}.$$

Thus, A_1 has EBB if

$$\Pr\left\{D_i^{(A_0,A_1)} \leq \frac{1}{\rho}\left(W_\rho\left(A_0\right)\left(s_i\right)-\sigma\right)\right\} \leq \Lambda e^{-\alpha\sigma},$$

Since the shaper transmits packets in an FCFS order we have

$$\Pr\left\{D_i^{(A_0,A_1)} \leq d\right\} = 0 \quad \forall\, d < t_{i-1} - s_i.$$

Note that $D_i^{(A_0,A_1)} = \frac{W_\rho(A_0)(s_i)}{\rho}$ implies that $W_\rho\left(A_1\right)\left(t_i\right) = 0$ and thus (8) holds. Since $D_i^{(A_0,A_1)}$ is as small as possible subject to the constraint (8) we have $\Pr\left\{D_i^{(A_0,A_1)} \leq d\right\} = 1\ \forall\, d \geq \frac{W_\rho(A_0)(s_i)}{\rho}$. Furthermore, one can see that $t_{i-1} - s_i < \frac{W_{\rho_0}(A_0)(s_i)}{\rho_0}$.

Consider an EBB stream A_0 with the parameters $(\rho, \Lambda_0, \alpha_0)$ entering a $(\rho, \Lambda_1, \alpha_1)$-shaper. We have

$$\Pr\left\{W_\rho\left(A_0\right)\left(s_i\right) \geq \sigma\right\} \leq \Lambda_0 e^{-\alpha_0\sigma}$$

and, from (9),

$$D_i^{(A_0,A_1)} \leq \frac{W_\rho\left(A_0\right)\left(s_i\right)}{\rho}.$$

Therefore,

$$\Pr\left\{D_i^{(A_0,A_1)} \geq d\right\} \leq \Lambda_0 e^{-\alpha_0\rho d}. \tag{11}$$

We note that the upper bound on the shaper delay tail distribution in (11) is very loose. The bound looseness results from taking $D_i^{(A_0,A_1)} =$

$\frac{W_\rho(A_0)(s_i)}{\rho}$. That is, assuming that the shaper delays all packet until the backlog $W_\rho\left(A_1\right)\left(t_i\right)$ is cleared.

The next proposition is necessary for the analysis in the next section where we consider (ρ, Λ, α)-shapers in series.

Proposition 4 *Let $\rho \geq 0$ be given. For all $i \geq 1$ it holds that*

$$W_\rho\left(A_1\right)\left(t_i\right) = \left[W_\rho\left(A_0\right)\left(s_i\right) - \rho \cdot D_i^{(A_0,A_1)}\right]^+ \qquad (12)$$

Proof: See Appendix B.

3.2 The multiple node case

We proceed to consider a sequence of service elements, which constitute a path in the network, as depicted in Fig. 3.

3.2.1 Traffic shaper elements in series

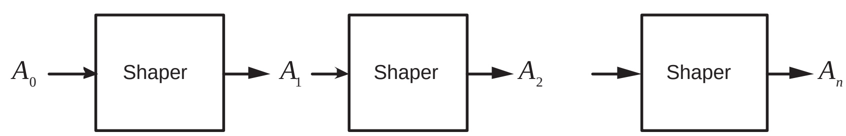

Figure 3: EBB Shapers in Series

First, we study the effect of connecting a sequence of some n (ρ, Λ, α)-shaper elements. Let $(\rho, \Lambda_k, \alpha_k)$ be the parameters of the kth shaper. A_k represents the rate of the traffic output of the kth shaper and A_0 represents the rate of the traffic input to the system. Let $D_i^{(A_0,A_k)}$ be the difference between the time at which the ith packet begins to exit the kth shaper (t_i^k) and the time at which it begins to arrive to the system (s_i).

Proposition 5 *For all $k = 1, 2, \ldots, n$ it holds that:*

$$\Pr\left\{D_i^{(A_0,A_k)} \le d\right\} \begin{cases} = 0 & d < t_{i-1}^k - s_i \\ \le \max_k \Lambda_k e^{-\alpha^k (W_\rho(A_0)(s_i) - \rho d)} & t_{i-1}^k - s_i \le d < \frac{W_\rho(A_0)(s_i)}{\rho} \\ = 1 & d \ge \frac{W_\rho(A_0)(s_i)}{\rho} \end{cases}, \tag{13}$$

Proof: We prove the proposition by induction on k. It holds for $k = 1$ by (9). Suppose it holds for some fixed k; we now show it still holds when k is replaced by $k+1$. To show that, we first apply Proposition 4 k times: in the mth application we use $A_{m-1}, A_m, s_i + D_i^{(A_0,A_{m-1})}, s_i + D_i^{(A_0,A_m)}, D_i^{(A_0,A_m)} - D_i^{(A_0,A_{m-1})}, \rho$ instead of $A_0, A_1, s_i, t_i, d_i, \rho$, respectively. This implies

$$W_\rho(A_k)\left(s_i + D_i^{(A_0,A_k)}\right) = \left[W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_k)}\right]^+. \tag{14}$$

$\Pr\left\{D_i^{(A_0,A_{k+1})} \le d\right\} = 0 \ \forall d < t_{i-1}^{k+1} - s_i$ holds since all the shapers transmit packets in an FCFS order.

Suppose that $D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} > 0$. Then, from (9),

$$\Pr\left\{D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} \le d \,\middle|\, D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} > 0\right\} \le \Lambda_{k+1} e^{-\alpha_{k+1}\left(W_\rho(A_k)\left(s_i + D_i^{(A_0,A_k)}\right) - \rho d\right)} \quad \forall\, t_{i-1}^{k+1} - s_i \le d < \frac{W_\rho(A_0)(s_i)}{\rho}.$$

Now if $D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} > 0$ then $W_\rho(A_k)\left(s_i + D_i^{(A_0,A_k)}\right) > 0$ and the '+' sign superscript in (14) can be removed. Hence,

$$\Pr\left\{D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} \le d \,\middle|\, D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} > 0\right\} \le \Lambda_{k+1} e^{-\alpha_{k+1}\left(W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_k)} - \rho d\right)} \quad \forall\, t_{i-1}^{k+1} - s_i \le d < \frac{W_\rho(A_0)(s_i)}{\rho}.$$

Therefore,

$$\Pr\left\{ D_i^{(A_0,A_{k+1})} \le d \middle| D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} > 0 \right\}$$
$$\le \Lambda_{k+1} e^{-\alpha_{k+1}(W_\rho(A_0)(s_i) - \rho d)}$$
$$\forall\, t_{i-1}^{k+1} - s_i \le d < \frac{W_\rho\left(A_0\right)\left(s_i\right)}{\rho}. \quad (15)$$

On the other hand, if $D_i^{(A_0,A_{k+1})} - D^{(A_0,A_k)} = 0$, it follows from the induction hypothesis that

$$\Pr\left\{ D_i^{(A_0,A_{k+1})} \le d \middle| D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} = 0 \right\} \le \max_k \Lambda_k e^{-\alpha^k(W_\rho(A_0)(s_i) - \rho d)}$$
$$\forall\, t_{i-1}^{k+1} - s_i \le d < \frac{W_\rho\left(A_0\right)\left(s_i\right)}{\rho}. \quad (16)$$

Thus, from (15) and (16) it follows that

$$\Pr\left\{ D_i^{(A_0,A_{k+1})} \le d \middle| D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} = 0 \right\} \le \max_{k+1} \Lambda_{k+1} e^{-\alpha_{k+1}(W_\rho(A_0)(s_i) - \rho d)}$$
$$\forall\, t_{i-1}^{k+1} - s_i \le d < \frac{W_\rho\left(A_0\right)\left(s_i\right)}{\rho}. \quad (17)$$

Note that $\Pr\left\{ D_i^{(A_0,A_{k+1})} \le d \right\} = 0 \;\forall\, d < t_{i-1}^{k+1} - s_i$ holds since all the shapers transmit packets in an FCFS order.

Finally, suppose that $D_i^{(A_0,A_k)} = \frac{W_\rho(A_0)(s_i)}{\rho}$; then, from (14), $W_\rho\left(A_k\right)\left(s_i + D_i^{(A_0,A_k)}\right) = 0$. Thus, $D_i^{(A_0,A_{k+1})} - D_i^{(A_0,A_k)} = 0$ and $D_i^{(A_0,A_{k+1})} = \frac{W_\rho(A_0)(s_i)}{\rho}$.

3.2.2 EDF-Scheduler and EBB-Shaper in Cascade

Consider two systems, S_1 and S_2, where system S_1 consists of a (ρ, Λ, α)-shaper and system S_2 consists of a "delay" subsystem and an identical

(ρ, Λ, α)-shaper connected in series as depicted in Fig. 4. The delay subsystem delays the ith arriving packet by $\theta_i \geq 0$ and then delivers it to the shaper. The following lemma relates the delays experienced by a packet in the two systems S_1 and S_2. More precisely, the lemma states that the delay distribution in system S_2 is upper-bounded by the same function that upper bounds the distribution of the delay in system S_1.

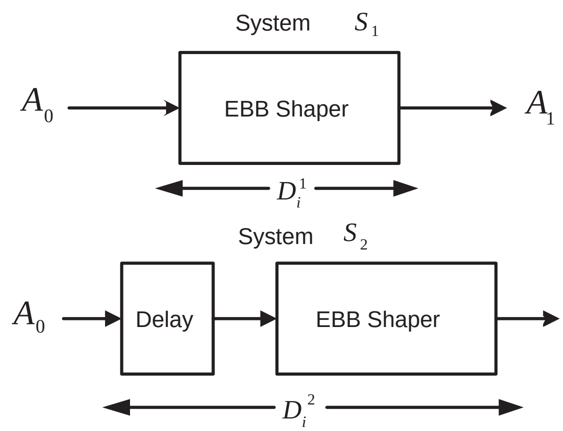

Figure 4: The Systems S_1 and S_2

Lemma 1 *Assume that packets arrive to systems S_1, S_2 according to the same arrival process A_0. If $D_i^{(1)}$ and $D_i^{(2)}$ are the delays of packet i in the traffic shaper in systems S_1 and S_2 respectively, then, for all $i = 1, 2, \ldots$,*

$$\Pr\left\{D_i^{(2)} + \theta_i \leq d\right\} \begin{cases} = 0 & d < t_{i-1}^{(1)} - s_i^{(1)} \\ \leq \Lambda e^{-\alpha\left(W_\rho(A_0)\left(s_i^{(1)}\right) - \rho d\right)} & t_{i-1}^{(1)} - s_i^{(1)} \leq d < \frac{W_\rho(A_0)\left(s_i^{(1)}\right)}{\rho} \\ = 1 & d \geq \frac{W_\rho(A_0)\left(s_i^{(1)}\right)}{\rho} \end{cases} . \tag{18}$$

Proof: The equalities in (18) are straightforward. Consider, then, $t_{i-1}^{(1)} - s_i^{(1)} \leq d < \frac{W_\rho(A_0)\left(s_i^{(1)}\right)}{\rho}$.

Let $s_i^{(1)}$ be the time the ith packet of length l_i arrives to either of the two systems. Let also $s_i^{(2)} = s_i^{(1)} + \theta_i$ be the time that packet i enters the shaper in system S_2. According to (9) we have, for $n = 1, 2$,

$$\Pr\left\{D_i^{(n)} \leq d\right\} \leq \Lambda e^{-\alpha\left(W_\rho(A_0)\left(s_i^{(n)}\right)-\rho d\right)} \tag{19}$$

Hence,

$$\Pr\left\{D_i^{(2)} + \theta_i \leq d\right\} \leq \Lambda e^{-\alpha\rho\left(\frac{1}{\rho}\left(W_\rho(A_0)\left(s_i^{(2)}\right)-\rho d\right)+\theta_i\right)}. \tag{20}$$

Set

$$\begin{aligned} g_i^{(n)}(\sigma) &= \frac{1}{\rho}\left(W_\rho(A_0)\left(s_i^{(n)}\right) - \sigma\right) \\ &= \max_{0 \leq s \leq s_i^{(n)}} \left\{ \frac{A_0\left(s, s_i^{(n)}\right) - \rho\left(s_i^{(n)} - s\right) - \sigma}{\rho} \right\} \quad \forall\, 0 \leq \sigma \leq W_\rho(A_0)\left(s_i^{(n)}\right). \end{aligned} \tag{21}$$

From the proof of Lemma 1 in [5] we have

$$g_i^{(1)}(\sigma) \leq g_i^{(2)}(\sigma) + \theta_i. \tag{22}$$

We note that, while (22) was proven in [5] in the context of deterministic traffic, it is easy to verify that the proof holds, with no modification, for stochastic traffic as well.

From (20) and (22), we have

$$\Pr\left\{D_i^{(2)} + \theta_i \leq d\right\} \leq \Lambda e^{-\alpha\rho\left(g_i^{(2)}(\rho d)+\theta_i\right)} \leq \Lambda e^{-\alpha\rho g_i^{(1)}(\rho d)},$$

therefore,

$$\Pr\left\{D_i^{(2)} + \theta_i \leq d\right\} \leq \Lambda e^{-\alpha\left(W_\rho(A_0)\left(s_i^{(1)}\right)-\rho d\right)}.$$

The following proposition considers a series of an EDF-scheduler and an EBB-shaper.

Proposition 6 *Assume that the output of a $(\rho, \Lambda_0, \alpha_0)$-shaper $\mathcal{A}_0$ enters a system $\mathcal{S}$, for which it is known that the delay experienced by a packet i is exponentially bounded as follows:*

$$\Pr\left\{D_i\left(t\right) \geq d\right\} \leq \tilde{\Lambda} e^{-\tilde{\alpha} r d} \quad \forall d \geq \bar{d}.$$

The output of system $\mathcal{S}$ enters $(\rho, \Lambda_1, \alpha_1)$-shaper $\mathcal{A}_1$.

The total delay, $\hat{D}_i$, experienced by a packet i, from the time it enters the scheduler till the time it exits $\mathcal{A}_1$, is exponentially bounded as follows:

$$\Pr\left\{\hat{D}_i \geq d\right\} \leq \left(\Lambda_0 + \tilde{\Lambda}\right) e^{-\frac{1}{\frac{1}{\alpha_0 \rho} + \frac{1}{\tilde{\alpha} r}} d} \quad \forall d \geq \bar{d}. \tag{23}$$

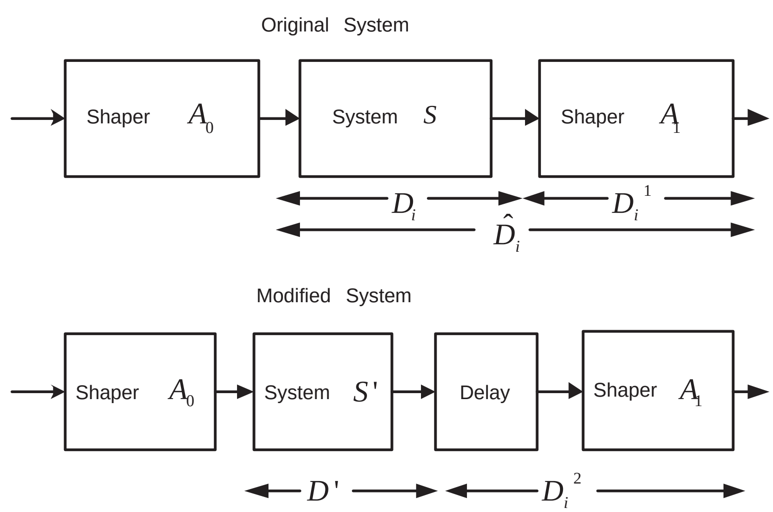

Figure 5: Original and Modified System

Proof: Let D_i be the delay of packet i in system $\mathcal{S}$, and let $D_i^{(1)}$ be its delay in $\mathcal{A}_1$. Therefore, $\hat{D}_i = D_i + D_i^{(1)}$. First, consider a modified system where a delay subsystem with $\theta_i = \left(\bar{d} - D_i\right)^+$ is inserted between $\mathcal{S}$ and $\mathcal{A}_1$, and let $D_i^{(2)}$ be the delay of packet i in $\mathcal{A}_1$ under this new arrangement. Next, replace S with a modified system $\mathcal{S}'$ in which all packets experience the same delay D' such that

$$\Pr\left\{D' \geq d\right\} \leq \tilde{\Lambda} e^{-\tilde{\alpha} r d} \quad \forall d \geq \bar{d}. \tag{24}$$

(See Fig. 5.) Observe now that, since the delay of every packet between its entrance to $\mathcal{S}'$ and its exit from the delay system is $\max\left\{D', \bar{d}\right\}$, the traffic entering the shaper $\mathcal{A}_1$ is a time-shifted version of the traffic exiting $\mathcal{A}_0$. Hence $D_i^{(2)} = D_i^{(A_0,A_1)}$ and from (11) we have

$$\Pr\left\{D_i^{(2)} \geq d\right\} \leq \Lambda_0 e^{-\alpha_0 \rho d}. \tag{25}$$

Applying sample path reasoning along with Lemma 1 we get

$$\Pr\left\{\hat{D}_i \geq d\right\} \leq \Pr\left\{D' + D^{(2)} \geq d\right\} \quad \forall d \geq \bar{d}. \tag{26}$$

(24) and (25) imply that both D' and $D_i^{(2)}$ are EB. Thus, we have to sum two EB processes. The proposition follows from the calculus of EB processes presented in [9].

Observe that the upper bound on the delay tail distribution presented in Proposition 6 is loose. This is due to the use of the upper bound given in (11) for an isolated shaper element. A tighter bound can be achieved for the special case where $\mathcal{A}_1$ is a $(\rho, \Lambda_0, \alpha_0)$-shaper. Consider the modified system presented in the proof of proposition 6. The traffic entering $\mathcal{A}_1$, which has EBB with the parameters $(\rho, \Lambda_0, \alpha_0)$, need not be delayed in $\mathcal{A}_1$; the $(\rho, \Lambda_0, \alpha_0)$-shaper, $\mathcal{A}_1$, should forward the input traffic immediately to the output with no further delay, since the input traffic has already the required output characteristics. Thus, we have:

$$\Pr\left\{\hat{D}_i \geq d\right\} \leq \tilde{\Lambda} e^{-\tilde{\alpha} r d} \quad \forall d \geq \bar{d}. \tag{27}$$

Note that even tighter upper bound may be achieved by using the more general (ρ, Λ, α)-shaper in $\mathcal{A}_1$, and by considering the actual parameters of the shaper rather than just those presented in (9).

3.2.3 End-to-end Delay

Finally, consider a connection k with $\left(\rho^k, \Lambda^k, \alpha^k\right)$-EBB traffic. The connection is routed through a path $\mathbf{p}$ in which all nodes employ the RC-EDF service discipline. We assume that the traffic shaper parameters for connection k at each node along its path are $\left(\rho^k, \Lambda^k, \alpha^k\right)$. Then, the stochastic bound on the end-to-end delay is given in the following theorem.

Theorem 1 *For any session k in an RC-EDF network with $\left(\rho^k, \Lambda^k, \alpha^k\right)$-EBB traffic and $\left(\rho^k, \Lambda^k, \alpha^k\right)$-traffic shapers, the end-to-end delay $D^k(\mathbf{p})$ is stochastically upper bounded as follows:*

$$\Pr\left\{D^k(\mathbf{p}) \geq d\right\} \leq \left(\sum_{l \in \mathbf{p}} \tilde{\Gamma}_l e^{\tilde{\alpha}_l \left(\sum_{\substack{j \in N_l \\ j \neq k}} \rho^j\right) \bar{d}_l^k}\right) e^{-\frac{1}{\sum_{l \in \mathbf{p}} \frac{1}{\tilde{\alpha}_l r_l}} d} \tag{28}$$

Proof: The end-to-end delay consists of the delays along a series of shapers and schedulers in cascade. We enforce, at internal nodes, the same traffic parameters as at the network access point. In this case, we showed in the previous section, namely expression (27), that the delay of a cascaded traffic shaper and EDF scheduler is exponentially bounded with the same bound as the delay of an EDF scheduler in isolation. Thus, the end-to-end delay can be calculated from the sum of the EB delays in the schedulers. The theorem follows from the calculus of EB processes in [9] and the the delay bound of a single EDF scheduler (4).

4 Call Admission Control and Routing algorithms

The related call admission control problem, namely, whether a new connection request can be accommodated, has been exclusively studied under deterministic settings and many schemes have been proposed. However, the corresponding routing problem, namely, the identification of a suitable path for establishing a new connection, has gained little attention. Thus, although the main focus of the paper is on EBB traffic and stochastic guarantees, the first part of this section considers routing algorithms with deterministic guarantees. These schemes, under the deterministic setting, provide some of the required basics for routing schemes under the stochastic settings, which we study next.

4.1 The Deterministic Setting

Under the deterministic setting we assume that the traffic arriving in a closed interval $[s, s+\tau]$ is deterministically bounded. More precisely, we assume that there is a nonnegative function $\bar{A}(\tau)$ called envelope of $A[s, s+\tau]$, such that for all nonnegative s and τ, $A[s, s+\tau] \leq \bar{A}(\tau)$. With such traffic characteristics, deterministic upper bound on the end-to-end delay can be guaranteed.

Under the deterministic setting, the routing problem has been studied in [13, 14] and several schemes have been proposed. However, the focus in those studies was on identifying feasible, but not necessarily (networkwide) efficient paths. For example, the proposed table-driven distributed route-selection scheme in [14] is guaranteed to find a "qualified" route (if such exists) that meets the performance requirement of the requested channel without compromising any of the existing guarantees. When a feasible path is found, the remainder delay, after subtracting the link delays from the end-to-end delay bound, is subdivided evenly among the links along the path. Such a path selection and a delay partition is not necessarily (networkwide) optimal. Thus, efficient schemes, which aim at balancing the load or, alternatively, at maximizing the ability to accommodate future calls, are called for. Obviously, such schemes should efficiently partition the delay.

First, we describe the deterministic model and the assumptions under which the routing schemes are suggested. We assume that the input traffic is Burstiness Constrained [18]. A traffic stream with a rate function $A(t)$, is *Burstiness Constrained (BC)* if

$$A^j[t, t+\tau] \leq \rho\tau + \sigma \ \ \forall\, \tau > 0,$$

where $A^j[t, t+\tau] = \int_t^{t+\tau} A(s)\, ds$, ρ is the long term upper rate of the arrival process and σ is the maximal burst size.

Based on more general traffic models, [3, 19] established schedulability conditions for an EDF scheduler. The authors of [3] considered bounded traffic with general envelop functions $\bar{A}^j(\tau)$ and established the following conditions:

a set of $\mathcal{N}$ connections is EDF-schedulable if and only if for all $t \geq 0$:

$$rt \geq \sum_{j \in \mathcal{N}} \bar{A}^j\left(t - \bar{d}^j\right)$$

and for all $\bar{d}^1 \leq t \leq \bar{d}^{|\mathcal{N}|}$:

$$rt \geq \sum_{j \in \mathcal{N}} \bar{A}^j\left(t - \bar{d}^j\right) + \max_{\bar{d}^j > t} L^j.$$

Taking $\bar{A}^j(\tau) = \left(\rho^j\tau + \sigma^j\right) 1(\tau)$, where $1(t) = \begin{cases} 1 & t \geq 0 \\ 0 & t < 0 \end{cases}$, we have:

a set of $\mathcal{N}$ connections, with Burstiness Constrained traffic, is EDF-schedulable if and only if for all $t \geq 0$:

$$rt \geq \sum_{j \in \mathcal{N}} \left[\rho^j\left(t - \bar{d}^j\right) + \sigma^j\right] 1\left(t - \bar{d}^j\right)$$

and for all $\bar{d}^1 \leq t \leq \bar{d}^{|\mathcal{N}|}$:

$$rt \geq \sum_{j \in \mathcal{N}} \left[\rho^j\left(t - \bar{d}^j\right) + \sigma^j\right] 1\left(t - \bar{d}^j\right) + \max_{\bar{d}^j > t} L^j.$$

Note that Burstiness Constrained traffic is a special case of the (C, σ, ρ) Token Bucket flows considered in [19] in which $C = \infty$.

4.1.1 Call Admission Control

To consider call admission control, we exploit the following terms: (i) the (work) availability function [19]

$$F(t) = rt - \sum_{j \in \mathcal{N}} \left[\rho^j \left(t - \bar{d}^j\right) + \sigma^j\right] 1 \left(t - \bar{d}^j\right) - \max_{\bar{d}^j > t} L_j,$$

and (ii) the minimum worst-case response time (MWRT) [14], which is the minimum delay that can be guaranteed to a new session while conforming to the schedulability conditions.

The call admission of a new session k can be depicted as follows:

Algorithm Admission Control - Minimum worst-case response time (AC-MWRT)

1. if $\sum_{j \in \mathcal{N}} \rho^j + \rho^k > r$ then exit "cannot accept session k"
2. calculate the minimum $\hat{d}^k$ such that

$$F(t) - \left(\rho^k \left(t - \hat{d}^k\right) + \sigma^k\right) u \left(t - \hat{d}^k\right) > 0$$

The establishment of efficient schemes to calculate $\hat{d}^k$ is not in the scope of this study. Such schemes can be found in [19].

4.1.2 Routing Algorithms

Here, we consider the more general problem, of optimizing the route choice in terms of balancing the loads and accommodating multiple connections.

First, we aim at balancing the loads by seeking a path for which the residual maximal rate (after establishing the new connection) of its bottleneck link is maximal. Algorithm MRB identifies such a path.

Let δ_l be the residual rate at link l, i.e., $\delta_l = r_l - \sum_{j \in \mathcal{N}_l} \rho^j$ and let the $\Delta^1 \geq \Delta^2 \geq \ldots \geq \Delta^U$ be the different values of δ_l, where $U \leq M$. Denote by D the session k required end-to-end delay.

Algorithm Maximum Residual Bottleneck (MRB)

1. for $u = 1$ to U

   (a) delete all links l for which $\delta_l < \Delta^u$

   (b) find the shortest path with respect to the metric $\left\{\hat{d}_l^k\right\}$.

   (c) if $\sum_{l \in \mathbf{p}^*} \hat{d}_l^k < D$ then

      i. $\mathbf{p}^*$ is the required path

      ii. subdivide the remainder delay $D - \sum_{l \in \mathbf{p}^*} \hat{d}_l^k$ evenly among the links along $\mathbf{p}^*$

2. if no path was identified, then the connection is not feasible

While the above algorithm aims at balancing the load, it does not efficiently partition the remainder delay along the identified path. The following algorithm, MWA, takes another approach to balance the load and to efficiently partition the delays along the path. MWA aims at maximizing the available work function (after accommodating the new connection) at the bottleneck link (with respect to the available work function). More precisely, Let $W_l\left(\bar{d}_l^k\right)$ be the minimal available work function at link l after accommodating a new connection k,

$$W_l^k\left(\bar{d}_l^k\right) = \min_{t>0}\left\{rt - \sum_{\substack{j \in \mathcal{N}_l \\ j \neq k}} \left[\left(\rho^j\left(t - \bar{d}_l^j\right) + \sigma^j\right) 1\left(t - \bar{d}_l^j\right)\right] - \left(\rho^k\left(t - \bar{d}_l^k\right) + \sigma^k\right) 1\left(t - \bar{d}_l^k\right) - \max_{\bar{d}^j > t} L_j\right\};$$

given a connection k and a feasible path $\mathbf{p}$ denote by $W_{\mathbf{p}}^k$ the minimal available work of its bottleneck link l^*, i.e.,

$$W_{\mathbf{p}}^k = \min_{l \in \mathbf{p}} W_l^k\left(\bar{d}_l^k\right).$$

The problem then is to find a feasible path $\mathbf{p}$ and a feasible partition $\bar{d}_l^k$ $\forall l \in$

$\mathbf{p}$ $(\bar{d}_l^k \geq \hat{d}_l^k, \sum_{l \in \mathbf{p}} \bar{d}_l^k \leq D)$ that maximizes $W_{\mathbf{p}}^k$.

Note that $W_l\left(\bar{d}_l^k\right)$ is a non-decreasing piecewise linear function of $\left(\bar{d}_l^k\right)$ and the function discontinuities occur at times $\bar{d}_l^k = \bar{d}_l^j \; \forall j \in \mathcal{N}_l$.

Let $W^1 \geq W^2 \geq \ldots \geq W^U$ be the different values of $W_l\left(\bar{d}_l^k\right) \quad \forall l \in E$ at all the discontinuities points. It hods that $U \leq M \cdot \mathcal{N}_{\max}$, where $\mathcal{N}_{\max} = \max_{l \in E} \mathcal{N}_l$.

Algorithm MWA

1. delete all links l which cannot accept session k
2. for $i = 1, 2, \ldots, U$
   (a) for all $l \in E$ compute $d_l^k(i)$ for which $W_l\left(d_l^k(i)\right) = W^i$
   (b) find a path $\mathbf{p}(i)$ that is the shortest path with respect to the metric $\left\{d_l^k(i)\right\}$
   (c) if $\sum_{l \in \mathbf{p}(i)} d_l^k(i) < D$ then
      i. $\mathbf{p}(i)$ is the required path
      ii. subdivide the remainder delay $D - \sum_{l \in \mathbf{p}(i)} d_l^k(i)$ evenly among the links along $\mathbf{p}(i)$
3. If no path was identified, then the connection is not feasible

Complexity For each possible value of the available work function at a discontinuity point W^i we compute the delay at each link in $O(1)$ steps and find a shortest path in $O(N \log N + M)$ steps (through dijkstra's algorithm). Thus the overall complexity is $O(M \cdot \mathcal{N}_{\max}(N \log N + M))$.

4.2 The Stochastic Setting

We now briefly discuss the problem in our framework, namely, RC-EDF schedulers under a stochastic setting.

4.2.1 Call Admission Control

Consider the stochastic end-to-end delay bound in (28). It is easy to see that the QoS guarantee of a session is affected by the traffic of other sessions at each service element along its path. Thus, each service element ν must meet certain performance bounds in order to provide QoS guarantees. These bounds, which characterize the achievable QoS at each service element should be advertised in the network for routing purposes. It can be verified that the following set of performance bounds is sufficient:

- $\sum_{j \in \mathcal{N}_\nu} \rho^j \leq \hat{\rho}_\nu$,
- $\frac{1}{\sum_{j \in \mathcal{N}_\nu} \frac{1}{\alpha^j}} \geq \hat{\alpha}_\nu$,
- $\frac{\left(\sum_{j \in \mathcal{N}} \Lambda^j\right)}{1 - e^{-\frac{1}{\sum_{j \in \mathcal{N}_\nu} \frac{1}{\alpha^j}}\left(r - \sum_{j \in \mathcal{N}_\nu} \rho^j\right)}} e^{-\frac{1}{\sum_{j \in \mathcal{N}_\nu} \frac{1}{\alpha^j}}\left(\sum_{j \in \mathcal{N}_\nu} \rho^j \bar{d}_\nu^j - L_{\max}\right)} \leq \hat{\Gamma}_\nu$.

Accordingly we propose the following algorithm for admission control of EBB sessions. The algorithm takes as inputs the characteristics and delay guarantees for the existing sessions in $\mathcal{N}_\nu$ and the characteristics for the new session k. It outputs the minimum (stochastically) guarantee-able deadline for session k.

Algorithm Admission Control - Minimum Deadline (AC-MD)
(**input:** $\left(\rho^j, \Lambda^j, \alpha^j, \bar{d}_\nu^j\right)_{j\in\mathcal{N}_\nu}$, $\left(\rho^k, \Lambda^k, \alpha^k\right)$; **output:** $\hat{d}_\nu^k$)

1. if $\sum_{j\in\mathcal{N}_\nu} \rho^j + \rho^k > \hat{\rho}_\nu$ then exit "cannot accept session k"
2. if $\frac{1}{\sum_{j\in\mathcal{N}_\nu} \frac{1}{\alpha^j} + \frac{1}{\alpha^k}} < \hat{\alpha}_\nu$ then exit "cannot accept session k"
3. let $\hat{d}_\nu^k$ be the minimal $\bar{d}_\nu^k$ such that

$$\frac{\left(\sum_{j\in\mathcal{N}_\nu} \Lambda^j + \Lambda^k\right) e^{-\frac{1}{\sum_{j\in\mathcal{N}_\nu} \frac{1}{\alpha^j} + \frac{1}{\alpha^k}}\left(\sum_{j\in\mathcal{N}_\nu} \rho^j \bar{d}_\nu^j + \rho^k \bar{d}_\nu^k - L_{\max}\right)}}{1 - e^{-\frac{1}{\sum_{j\in\mathcal{N}_\nu} \frac{1}{\alpha^j} + \frac{1}{\alpha^k}}\left(r - \sum_{j\in\mathcal{N}_\nu} \rho^j - \rho^k\right)}} \le \hat{\Gamma}_l$$

holds

4.2.2 Routing Algorithms

With the stochastic bound (28) on the end-to-end delay and the call admission control algorithm AC-MD at hand, we investigate the corresponding routing problem. We adopt the concept of *q-feasible* paths introduced in [8]. Namely, we assume that each session is associated with a certain probability q, which reflects its "sensitivity" to end-to-end fluctuations beyond the required delay D. Intuitively, a path is said to be feasible if the end-to-end delay fluctuations beyond the required value conforms with the session's "sensitivity". More precisely, a path $\mathbf{p}$ is q-feasible if $\Pr\{D(\mathbf{p}) \geq D\} \leq q$.

We begin by considering the basic problem of identifying q-feasible paths. If several q-feasible paths exist, we seek a path with the minimal end-to-end delay tail distribution. The problem, then, is to find a path that minimizes the delay distribution upper bound, i.e., solve the following problem:

$$\min_{\mathbf{p}} \left(\sum_{l\in\mathbf{p}} \hat{\Gamma}_l e^{\hat{\alpha}_l\left(\hat{\rho}_l - \rho^k\right)\bar{d}_l^k}\right) e^{-\frac{1}{\sum_{l\in\mathbf{p}} \frac{1}{\hat{\alpha}_l r_l}} D} \qquad D \ \geq \ \sum_{l\in p} \bar{d}_l^k, \ \underset{l\in\mathbf{p}}{\forall} \bar{d}_l^k \ \geq \ \hat{d}_l^k.$$

One can observe that the solution to this problem essentially involves high computational complexity. Thus, we consider the following problem, which follows by employing the trivial bound $\sum\limits_{l\in\mathbf{p}} \frac{1}{\hat{\alpha}_l r_l} \leq \frac{n(\mathbf{p})}{\min\limits_{l\in\mathbf{p}} \hat{\alpha}_l r_l}$:

$$\min_{\mathbf{p}} \quad \left(\sum_{l\in\mathbf{p}} \hat{\Gamma}_l e^{\hat{\alpha}_l \left(\hat{\rho}_l - \rho^k\right) \bar{d}_l^k} \right) e^{-\frac{\min\limits_{l\in\mathbf{p}} \hat{\alpha}_l r_l}{n(\mathbf{p})} D} \qquad D \;\geq\; \sum_{l\in p} \bar{d}_l^k, \;\; \forall_{l\in\mathbf{p}} \bar{d}_l^k \;\;\geq\;\; \hat{d}_l^k.$$

The following algorithm correctly solves the problem.

Let $\hat{\alpha r}^1 \geq \hat{\alpha r}^2 \geq \ldots \geq \hat{\alpha r}^U$ be the different values of $\hat{\alpha}_l r_l$, where $U \leq M$.

Algorithm Minimum Delay tail Distribution (MDD)

1. delete all links l which cannot accept session k
2. for $u \leftarrow 1$ to U
   (a) delete all links l with $\hat{\alpha}_l r_l < \hat{\alpha} r^u$
   (b) for $n \leftarrow 1$ to H
      i. find the shortest path with respect to the metric $\left\{\hat{\Gamma}_l e^{\hat{\alpha}_l\left(\hat{\rho}_l-\rho^k\right)\hat{d}_l^k}\right\}$ among n-hops paths under the constraint $\sum_{l\in\mathbf{p}} \hat{d}_l^k \leq D$ (through a restricted shortest path algorithm)
      ii. compute the stochastic bound
$$\left(\sum_{l\in\mathbf{p}(n,u)} \hat{\Gamma}_l e^{\hat{\alpha}_l\left(\hat{\rho}_l-\rho^k\right)\hat{d}_l^k}\right) e^{-\frac{\min_{l\in\mathbf{p}(n,u)} \hat{\alpha}_l r_l}{n(\mathbf{p}(n,u))}D}$$
3. among $H \cdot U$ paths choose the path $\mathbf{p}^*$ with the minimal bound
4. if $\left(\sum_{l\in\mathbf{p}^*} \hat{\Gamma}_l e^{\hat{\alpha}_l\left(\hat{\rho}_l-\rho^k\right)\hat{d}_l^k}\right) e^{-\frac{\min_{l\in\mathbf{p}^*} \hat{\alpha}_l r_l}{n(\mathbf{p}^*)}D} \leq q$ then
   (a) $\mathbf{p}^*$ is a q-feasible path
   (b) while $\left(\sum_{l\in\mathbf{p}^*} \hat{\Gamma}_l e^{\hat{\alpha}_l\left(\hat{\rho}_l-\rho^k\right)\bar{d}_l^k}\right) e^{-\frac{\min_{l\in\mathbf{p}^*} \hat{\alpha}_l r_l}{n(\mathbf{p}^*)}D} < q$ increase $\bar{d}_l^k$ to maximize
$$\min_{l\in\mathbf{p}^*}\left(\hat{\Gamma}_l - \frac{\left(\sum_{j\in\mathcal{N}_l}\Lambda^j\right) e^{-\frac{1}{\sum_{j\in\mathcal{N}_l}\frac{1}{\alpha^j}}\left(\sum_{j\in\mathcal{N}_l}\rho^j\bar{d}_l^j - L_{\max}\right)}}{1 - e^{-\frac{1}{\sum_{j\in\mathcal{N}_l}\frac{1}{\alpha^j}}\left(r-\sum_{j\in\mathcal{N}_l}\rho^j\right)}}\right)$$
5. else there is no q-feasible path

Notice that algorithm MDD balances the consumed resources along the identified path; however, it does not balance the load in the network. The following algorithm considers the more general problem, of optimizing the route choice in terms of balancing the loads and accommodating multiple connections.

Let δ_l be the residual rate at link l, i.e., $\delta_l = \hat{\rho}_l - \sum_{j \in \mathcal{N}_l} \rho^j$ and let the $\Delta^1 \geq \Delta^2 \geq \ldots \geq \Delta^U$ be the different values of δ_l, where $U \leq M$.

Algorithm MRB aims at balancing loads by seeking a q-feasible path for which the residual rate (after establishing the new connection) of its bottleneck link is maximal.

Algorithm Maximum Residual Bottleneck (MRB)

1. for $u = 1$ to U
   (a) delete all links l for which $\delta_l < \Delta^u$
   (b) execute algorithm MDD
   (c) if algorithm MDD identifies a q-feasible path $\mathbf{p}^*$, then $\mathbf{p}^*$ is the required path
2. if no q-feasible path was identified, then the connection is not feasible

The above algorithm focuses on the residual rate; however, it can be easily modified to consider each of the other bounds guaranteed by each service element, i.e., $\hat{\alpha}_l$ or $\hat{\Gamma}_l$.

5 Conclusion

This study establishes a scalable approach to end-to-end QoS provisioning of traffic aggregates. Such aggregates translate into looser (stochastic) characterization of traffic profiles. Previous studies (e.g., [6, 7, 8]) dealt with either the generalized processor sharing scheduling discipline (GPS), under both deterministic and stochastic settings, or with the rate-controlled earliest deadline first discipline (RC-EDF), under a deterministic setting (e.g., [3, 4, 5]). The present study is the first to provide end-to-end bounds for exponentially bounded burstiness (EBB) traffic and systems of RC-EDF schedulers.

These bounds provide the required foundations for several related network control problems. In this study, we demonstrated how the end-to-end delay bound can be applied in order to devise an adequate call admission control and routing schemes. Due to the complexity of the end-to-end bound under the stochastic setting, these schemes are quite complex. However, the complexity can be appropriately reduced by resorting to approximation techniques that employ scaling and rounding, hence obtaining sub-optimal (ϵ-optimal) solutions.

An important direction for future research is a performance comparison between the RC-EDF scheduling discipline and the GPS discipline under stochastic settings. In particular, under the deterministic setting, RC-EDF was proven to outperform GPS in providing end-to-end delay guarantees in a network [5]; whether a similar result can be obtained also under the stochastic setting is a challenging open question.

References

[1] H. Zhang and D. Ferrari, "Rate-controlled service disciplines," *Journal of High Speed Networks*, vol. 3, no. 4, pp. 389–412, 1994.

[2] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks - the single node case," *IEEE/ACM Transanctions on Networking*, vol. 1, 1993.

[3] J. Liebeherr, D. Wrege, and D. Ferrari, "Exact admission control for networks with a bounded delay service," *IEEE/ACM Transanctions on Networking*, vol. 4, no. 6, pp. 885–901, December 1996.

[4] Leonidas Georgiadis, Roch Guerin, and Abhay Parekh, "Optimal multiplexing on a single link: Delay and buffer requirements," *IEEE Trans. Information Theory*, vol. 43, no. 5, pp. 1518–1535, Septemer 1997.

[5] L. Georgiadis, R. Guerin, V. Peris, and K. N. Sivarajan, "Efficient network QoS provisioning based on per node traffic shaping," *IEEE Trans. Networking*, vol. 4, no. 4, pp. 482–501, August 1996.

[6] O. Yaron and M. Sidi, "Generalized processor sharing networks with exponentially bounded burstiness arrivals," *Journal of High Speed Networks*, vol. 3, pp. 375–387, 1994.

[7] Z.-L. Zhang, D. Towsley, and J. Kurose, "Statistical analysis of the generalized processor sharing scheduling discipline," *IEEE Journal of Selected Areas in Communications*, vol. 13, no. 6, pp. 1071–1080, August 1995.

[8] E. Biton and A. Orda, "QoS provision and routing with deterministic and stochastic guarantees," CCIT Report #318, Dept. of Electrical Engineering, Technion, Israel, June 2000. Available from http://comnet.technion.ac.il/~berez/bo00.pdf.

[9] O. Yaron and M. Sidi, "Performance and stability of communication networks via robust exponential bounds," *IEEE/ACM Transanctions on Networking*, vol. 1, no. 3, pp. 372–385, 1993.

[10] Vijay Sivaraman and Fabio M. Chiussi, "Providing end-to-end statistical delay guarantees with earliest deadline first scheduling and per-hop traffic shaping," in *PROCEEDINGS OF IEEE INFOCOM 2000*, March 2000, pp. 631–640.

[11] Vijay Sivaraman, Fabio M. Chiussi, and Mario Gerla, "End-to-end statistical delay service under gps and edf scheduling: A comparison study," in *PROCEEDINGS OF IEEE INFOCOM 2001*, April 2001, pp. 1113 – 1122.

[12] M. Andrews, "Probabilistic end-to-end delay bounds for earliest deadline first scheduling," in *PROCEEDINGS OF IEEE INFOCOM 2000*, March 2000.

[13] D. Ferrari and D. C. Verma, "A scheme for real-time channel establishment in wide-area networks," *IEEE Journal of Selected Areas in Communication*, vol. 8, pp. 368–379, April 1990.

[14] Chih-Che Chou and Kang G. Shin, "A distributed table-driven route selection scheme for establishing real-time video channels," in *International Conference on Distributed Computing Systems*, 1995, pp. 52–59.

[15] R. Guerin and A. Orda, "QoS-based routing in networks with inaccurate information: Theory and algorithms," *IEEE/ACM Transanctions on Networking*, vol. 7, no. 3, pp. 350–364, June 1999.

[16] Q. Ma and P. Steenkiste, "Quality of service routing for traffic with performance guarantees," in *IWQOS'97*, May 1997, pp. 115–126.

[17] A. Orda, "Routing with end to end QoS guarantees in broadband networks," *IEEE/ACM Transanctions on Networking*, vol. 7, no. 3, pp. 365–374, June 1999.

[18] R. L. Cruz, "A calculus for network delay, part i: Network elements in isolation," *IEEE Trans. Information Theory*, vol. 37, no. 1, pp. 114–131, January 1991.

[19] V. Firoiu, J. Kurose, and D. Towsley, "Efficient admission control of piecewise linear traffic envelopes at edf schedulers," *IEEE/ACM Transanctions on Networking*, vol. 6, no. 5, pp. 558–570, October 1998.

A Proof of proposition 1

Consider a tagged packet from connection $k \in \mathcal{N}$ that arrives at an EDF-scheduler at time t and is completely transmitted at time $t + D^k(t)$. Let

$W^{k,t}(t+\tau)$ be the workload in the scheduler at time $t+\tau$ that is served before the tagged packet (from connection k with arrival time t). The tagged packet has a deadline violation greater than d if for all τ, $0 \leq \tau \leq \bar{d}^k + d$, $W^{k,t}(t+\tau) > 0$, *i.e.*,

$$\Pr\left\{D^k(t) > \bar{d}^k + d\right\} = \Pr\left\{\min_{0 \leq \tau \leq \bar{d}^k + d}\left\{W^{k,t}(t+\tau)\right\} > 0\right\}.$$

Let $t - b(t)$ (where $b(t) \geq 0$) be the last time before t at which 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 $t - b(t)$ is guaranteed to exist. $b(t)$ is given by:

$$b(t) = \min\left\{z \middle| W^{\leq t+\bar{d}^k}(t-z) = 0, z \geq 0\right\},$$

where $W^{\leq x}(y)$ denotes the workload in the scheduler at time y due to packets with deadlines less than or equal to x. We have:

$$\begin{aligned}\Pr\left\{D^k(t) > \bar{d}^k + d\right\} &= \sum_{\hat{\tau}=0}^{t} \Pr\left(\left\{D^k(t) > \bar{d}^k + d\right\} \cap \{b(t) = \hat{\tau}\}\right) \\ &= \sum_{\hat{\tau}=0}^{t} \Pr\left(\left\{\min_{0 \leq \tau \leq \bar{d}^k + d}\left\{W^{k,t}(t+\tau)\right\} > 0\right\} \cap \{b(t) = \hat{\tau}\}\right) \qquad \text{(A.29)}\end{aligned}$$

Note that $\{b(t) = \hat{\tau}\}$ implies that, in each of the $\hat{\tau}$ time slots, r units of data have been transmitted. Also, note that $\left\{W^{k,t}(t+\tau) > 0\right\}$ $\forall \tau$ $\left(0 \leq \tau \leq \bar{d}^k + d\right)$ implies that some additional $r \cdot \tau$ data units have been transmitted from t to $t + \tau$. Moreover, $\{b(t) = \hat{\tau}\}$ also implies that the queue at time $t - \hat{\tau}$ did not contain any packets whose deadline was less than or equal to $t + \bar{d}^k$; these packets enter the scheduler only from $t - \hat{\tau}$ to $\min\left\{t + \tau, t + \bar{d}^k - \bar{d}^j\right\}$. Note that the deadlines of packets from session j that arrive after $t + \bar{d}^k - \bar{d}^j$ are greater than $t + \bar{d}^k$. Therefore,

$$\Pr\left(\left\{\min_{0\le\tau\le\bar{d}^k+d}\left\{W^{k,t}(t+\tau)\right\}>0\right\}\cap\{b(t)=\hat{\tau}\}\right)$$
$$=\Pr\left\{\min_{0\le\tau\le\bar{d}^k+d}\left\{\sum_{j\in\mathcal{N}}A^j\left[t-\hat{\tau},\min\left\{t+\tau,t+\bar{d}^k-\bar{d}^j\right\}\right]+R(t-\hat{\tau})-r(\hat{\tau}+\tau)\right\}>0\right\}, \tag{A.30}$$

where $R(t-b(t))$ denotes the amount of untransmitted data of a possible packet that is in transmission at time $t-b(t)$. We distinguish between two cases, namely whether at time $t-\hat{\tau}$ the scheduler is empty or transmitting a packet.

Case 1 $W(t-\hat{\tau})=0$. In this case, the scheduler is empty at time $t-\hat{\tau}$, *i.e.*, $R(t-\hat{\tau})=0$. We obtain from (A.30):

$$\Pr\left(\left\{D^k(t)>\bar{d}^k+d\right\}\cap\{b(t)=\hat{\tau}\}\right)$$
$$=\Pr\left\{\min_{0\le\tau\le\bar{d}^k+d}\left\{\sum_{j\in\mathcal{N}}A^j\left[t-\hat{\tau},\min\left\{t+\tau,t+\bar{d}^k-\bar{d}^j\right\}\right]-r(\hat{\tau}+\tau)\right\}>0\right\}$$
$$=\Pr\left\{\min_{0\le\tau\le\bar{d}^k+d}\left\{\sum_{j\in\mathcal{I}(\tau)}A^j[t-\hat{\tau},t+\tau]+\sum_{j\in\mathcal{N}-\mathcal{I}(\tau)}A^j\left[t-\hat{\tau},t+\bar{d}^k-\bar{d}^j\right]-r(\hat{\tau}+\tau)\right\}>0\right\} \tag{A.31}$$

where $j\in\mathcal{I}(\tau)$ if $t+\tau\le t+\bar{d}^k-\bar{d}^j$. Since the sessions traffic is stationary, it is easy to see that the minimum is achieved for $\mathcal{I}(\tau)=\emptyset$. Thus,

$$\Pr\left(\left\{D^k(t)>\bar{d}^k+d\right\}\cap\{b(t)=\hat{\tau}\}\right)$$
$$=\Pr\left\{\sum_{j\in\mathcal{N}}A^j\left[t-\hat{\tau},t+\bar{d}^k-\bar{d}^j\right]-r\left(\hat{\tau}+\bar{d}^k+d\right)>0\right\}.$$

Case 2 $W(t-b(t))>0$. The scheduler is transmitting traffic at time $t-\hat{\tau}$ from some connection j'. By the definition of $\hat{\tau}$ the traffic in transmission has a deadline greater that $t + \bar{d}^k$, that is, $\bar{d}^{j'} > \hat{\tau} + \bar{d}^k$. Without loss of generality, we assume that j' is such that $L_{j'} = \max_{\bar{d}^j > \hat{\tau} + \bar{d}^k} L^j$. From (A.29) and (A.30) and using similar arguments as in Case 1, we have:

$$
\begin{aligned}
&\Pr\left(\left\{D^k(t) > \bar{d}^k + d\right\} \cap \{b(t) = \hat{\tau}\}\right) \\
&= \Pr\left\{\min_{0 \le \tau \le \bar{d}^k + d} \left\{\sum_{j \in \mathcal{N}} A^j\left[t - \hat{\tau}, \min\left\{t + \tau, t + \bar{d}^k - \bar{d}^j\right\}\right] + \max_{\bar{d}^j > \hat{\tau} + \bar{d}^k} L^j - r(\hat{\tau} + \tau)\right\} > 0\right\} \\
&= \Pr\left\{\sum_{j \in \mathcal{N}} A^j\left[t - \hat{\tau}, t + \bar{d}^k - \bar{d}^j\right] + \max_{\bar{d}^j > \hat{\tau} + \bar{d}^k} L^j - r\left(\hat{\tau} + \bar{d}^k + d\right) > 0\right\}
\end{aligned}
\tag{A.32}
$$

From (A.29),(A.31) and (A.32) we get

$$
\begin{aligned}
&\Pr\left\{D^k(t) > \bar{d}^k + d\right\} \\
&\le \sum_{\hat{\tau}=0}^{t} \Pr\left\{\sum_{j \in \mathcal{N}} A^j\left[t - \hat{\tau}, t + \bar{d}^k - \bar{d}^j\right] + \max_{\bar{d}^j > \hat{\tau} + \bar{d}^k} L^j > r\left(\hat{\tau} + \bar{d}^k + d\right)\right\}.
\end{aligned}
\tag{A.33}
$$

■

B Proof of proposition 4

The proof of the proposition goes along similar lines to that of Lemma 5.1 in [18].

We show (12) by induction. It holds for the first packet ($i = 1$) since $W_{\hat{\rho}}(A_1)(t_1) = W_{\hat{\rho}}(A_0)(s_1) = D_1^{(A_0,A_1)} = 0$. We now assume that (12) holds for some fixed i and show that it also holds when i is replaced by $i+1$, *i.e.*, we show that

$$W_\rho(A_1)(t_{i+1}) = \left[W_\rho(A_0)(s_{i+1}) - \rho \cdot D_{i+1}^{(A_0,A_1)}\right]^+.$$

We have

$$\begin{aligned}
&W_\rho(A_1)(t_{i+1}) \\
&= [W_\rho(A_1)(t_i) + L_i - \rho(t_{i+1} - t_i)]^+ \\
&= \left[\left[W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)}\right]^+ + L_i - \rho(t_{i+1} - t_i)\right]^+.
\end{aligned} \tag{B.34}$$

Consider the following various cases:

Case 1: $W_\rho(A_1)(t_{i+1}) = 0$ In this case (B.34) implies

$$\begin{aligned}
0 &\geq W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)} + L_i - \rho(t_{i+1} - t_i) \\
&= W_\rho(A_0)(s_i) + L_i - \rho(s_{i+1} - s_i) - \rho D_{i+1}^{(A_0,A_1)}.
\end{aligned} \tag{B.35}$$

Since $D_{i+1}^{(A_0,A_1)} \geq 0$, (B.35) implies

$$[W_\rho(A_0)(s_i) + L_i - \rho(s_{i+1} - s_i)]^+ - \rho D_{i+1}^{(A_0,A_1)} \leq 0$$

Thus,

$$W_\rho(A_0)(s_{i+1}) - \rho D_{i+1}^{(A_0,A_1)} \leq 0$$

and

$$W_\rho(A_1)(t_{i+1}) = \left[W_\rho(A_0)(s_{i+1}) - \rho D_{i+1}^{(A_0,A_1)}\right]^+$$

in this case.

Case 2: $W_\rho(A_1)(t_{i+1}) > 0, W_\rho(A_1)(t_i) > 0$ In this case the plus sign superscripts can be removed from the right-hand side of (B.34). Thus

$$\begin{aligned}
&W_\rho(A_1)(t_{i+1}) \\
&= W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)} + L_i - \rho(t_{i+1} - t_i) \\
&= W_\rho(A_0)(s_i) + L_i - \rho(s_{i+1} - s_i) - \rho D_{i+1}^{(A_0,A_1)} \\
&= [W_\rho(A_0)(s_i) + L_i - \rho(s_{i+1} - s_i)]^+ - \rho D_{i+1}^{(A_0,A_1)} \\
&= W_\rho(A_0)(s_{i+1}) - \rho D_{i+1}^{(A_0,A_1)} \\
&= \left[W_\rho(A_0)(s_{i+1}) - \rho D_{i+1}^{(A_0,A_1)}\right]^+.
\end{aligned} \tag{B.36}$$

Note that (B.36) follows since $W_\rho(A_1)(t_{i+1}) > 0$ in this case.

Case 3: $W_\rho(A_1)(t_{i+1}) > 0, W_\rho(A_1)(t_i) = 0$ From Proposition 3 we have

$$\rho D_i^{(A_0,A_1)} \leq W_\rho(A_0)(s_i).$$

We will consider two subcases of this case.

Subcase 3A: $W_\rho(A_0)(s_i) = 0$ In this subcase it follows that $D_i^{(A_0,A_1)} = 0$. Hence it follows from (B.34) that

$$W_\rho(A_1)(t_{i+1}) = W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)} + L_i - \rho(t_{i+1} - t_i). \quad \text{(B.37)}$$

Subcase 3B: $W_\rho(A_0)(s_i) > 0$ From the definition of Case 3 ($W_\rho(A_1)(t_i) = 0$) and the induction hypothesis it follows that $W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)} \leq 0$. Since $\rho D_i^{(A_0,A_1)} \leq W_\rho(A_0)(s_i)$, we have $\rho D_i^{(A_0,A_1)} = W_\rho(A_0)(s_i)$. Hence it follows from (B.34) that

$$W_\rho(A_1)(t_{i+1}) = W_\rho(A_0)(s_i) - \rho D_i^{(A_0,A_1)} + L_i - \rho(t_{i+1} - t_i). \quad \text{(B.38)}$$

Now, (B.37) (B.38) are equivalent to (B.36). Thus, the identical reasoning used in Case 2 can be used to show that $W_\rho(A_1)(t_{i+1}) = \left[W_\rho(A_0)(s_{i+1}) - \rho D_{i+1}^{(A_0,A_1)}\right]^+$.
■