

Statistical Per-Flow Service Bounds in a Network with Aggregate Provisioning

Jörg Liebeherr, *Department of Computer Science, University of Virginia*

Stephen D. Patek, *Department of Systems and Information Engineering, University of Virginia*

Almut Burchard, *Department of Mathematics, University of Virginia*

Abstract—Scalability concerns of QoS implementations have stipulated service architectures where QoS is not provisioned separately to each flow, but instead to aggregates of flows. This paper determines stochastic bounds for the service experienced by a single flow when resources are managed for aggregates of flows and when the scheduling algorithms used in the network are not known. Using a recently developed statistical network calculus, per-flow bounds can be calculated for backlog, delay, and the burstiness of output traffic.

Index Terms—Statistical Multiplexing, Quality-of-Service, Admission Control, Network Calculus.

I. INTRODUCTION

Concerns about the scalability of Quality-of-Service architectures which offer service guarantees to individual network flows have led to the development of network services, in which service guarantees are provisioned to collections ('aggregates') of flows and where the network core does not perform any per-flow operations. The services defined in the Differentiated Services (DiffServ) architecture fall into this category. If at all, per-flow operations are performed only at the network entrance, for example, by regulating the amount of traffic from a flow that can enter the network.

Scenarios where core switches do not perform per-flow operations, and where the switch treats flows from an aggregate in a uniform fashion are sometimes referred to as 'aggregate scheduling'. Some deterministic bounds for aggregate scheduling at a single node are summarized in [5]. The general problem of stability with aggregate scheduling,

This work is supported in part by the National Science Foundation through grants ANI-9730103, ECS-9875688 (CA-REER), ANI-9903001, DMS-9971493, and ANI-0085955, and by an Alfred P. Sloan research fellowship.

that is, finiteness of delays and backlog, in a network where per-flow regulation is only performed at the network edge has been studied in [2], [17] and, under a broader set of assumptions, in [4]. The authors of [9] derive bounds on the utilization such that end-to-end delay bounds in a network with FIFO scheduling are finite, when the total number of flows in the network is known and when leaky-bucket regulators control the traffic from each flow at the network entrance. These bounds are extended to deadline-based scheduling algorithms in [20].

We are interested in computing statistical lower bounds on the service given to a flow in a network with aggregate provisioning, where the scheduling algorithms at the nodes are unknown. We assume that traffic is regulated at the network entrance by per-flow conditioning algorithms, and that there is no per-flow processing of traffic after it has entered the network. At network nodes, service is provisioned to aggregates of flows and we assume that no information is available on how the provisioned capacity is distributed to the flows. We assume that each node is workconserving, but do not require knowledge of the scheduling algorithms at the nodes. When calculating the service experienced by a single flow, treat that flow as having lower priority than any other flow in the network. Then, the service given to a single flow is computed from the service allocated to the aggregate that is left unused by the other flows. In a deterministic worst-case setting, the unused capacity that is available to a single flow will be small. However, by considering statistical multiplexing of flows, we observe that even under our pessimistic assumptions the lower bound on the service given to a single flow is considerable, especially if the number of flows is large. The statistical service seen by an individual

flow is much higher than the aggregate service divided by the number of flows.

The results in this paper can be used for verification of service level agreements with network service providers. If a network customer can measure its aggregate input to the network and the throughput of only a single flow, the customer can determine if the network service provider has provisioned the resources specified in the agreement. If the service seen by the single measured flow is worse than the statistical lower bound paper, the network service provider is likely to have under-provisioned network resources.

We derive our results within the context of a recently developed statistical network calculus. Statistically multiplexed arrivals from flows are presented in terms of *effective envelopes* [3], which are bounds on aggregate traffic that hold with high probability. Service guarantees to the flow aggregate are expressed in terms of *service curves*, which provide deterministic lower bounds on service guarantees. The service of a single flow is presented in terms of an *effective service curve*, which provides a lower bound on the service received by a single flow that holds with high probability. Specifically, we will be able to give a lower bound for the service experienced by a single flow in terms of an effective service curve of the form

$$S_j^{\varepsilon_1 + \varepsilon_2} = [S_C - \mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}]_+$$

where S_C is the service provisioned to the flow aggregate and $\mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}$ expresses a (strong) effective envelope of the statistically multiplexed arrivals from all flows in the aggregate. The parameters ε_1 and ε_2 are violation probabilities and are generally small, e.g., $\varepsilon_1, \varepsilon_2 = 10^{-9}$.

If one is interested in statistical bounds on delay, backlog, or loss to the aggregate as a whole (as opposed to bounds for a single flow), we refer to the rich literature on multiplexed regulated traffic e.g., [3], [8], [12], [13], [15], [18], [19]. Generally, these results focus on the analysis of a single node, and do not consider a multi-node network.

The remaining sections of this paper are structured as follows. In Section II, we review relevant results from the statistical network calculus in terms of effective service curves, as presented in [6]. In Section III we discuss the arrivals and service

provisioning of the flow aggregate. We extend the notion of effective envelopes from [3] to heterogeneous arrivals. In Section IV we present an effective service curve for a single flow at a node in which service is allocated to an aggregate of flows. In Section V, we discuss numerical examples for single node and feedforward networks and evaluate the service guarantees achievable with the constructed effective service curves.

II. NETWORK CALCULUS WITH STATISTICAL SERVICE GUARANTEES

Consider the traffic arrivals from a flow to a network node. The arrivals and departures, respectively, of a flow in the time interval $[0, t)$ are given in terms of stochastic processes $A(t)$ and $D(t)$ with non-negative increments, which are defined over a joint probability space and which satisfy $D(t) \leq A(t)$ for all $t \geq 0$. We assume a continuous time model with fluid left-continuous traffic arrival functions. Packetization delays and other effects of discrete-sized packets, such as the non-preemption of packet transmissions, are ignored. The backlog of a flow at time t , denoted by $B(t)$, is given by $B(t) = A(t) - D(t)$. The delay at time t , denoted as $W(t)$, is the delay experienced by an arrival which departs at time t , given by $W(t) = \inf\{d \geq 0 \mid A(t-d) \leq D(t)\}$. We will use $A(x, y)$ and $D(x, y)$ to denote the arrivals and departures in the time interval $[x, y)$, with $A(x, y) = A(y) - A(x)$ and $D(x, y) = D(y) - D(x)$. Moreover, we assume that $A(t) = D(t) = 0$ for $t < 0$. When analyzing delays in a network we ignore processing and propagation delays.

We make the following assumptions on the arrival functions.

- (A1) *Regulated arrivals.* The arrivals A of a flow are bounded by a function A^* , called the *arrival envelope*,¹ such that $A(t+\tau) - A(t) \leq A^*(\tau)$ for all $t, \tau \geq 0$. We assume that arrival envelopes are subadditive, that is, $A^*(x+y) \leq A^*(x) + A^*(y)$, for all $x, y \geq 0$.

The assumption of an arrival envelope translates into a requirement for per-flow traffic regulation at the network ingress. For example, a peak-rate

¹ A function E is called an *envelope* for a function f if $f(t+\tau) - f(\tau) \leq E(t)$ for all $t, \tau \geq 0$.

constrained leaky bucket regulator enforces that traffic from a flow adheres to the envelope $A^*(t) = \min(Pt, \sigma + \rho t)$ for a peak rate parameter P , an average rate parameter ρ , and a burstiness parameter σ .

Before we discuss the service guarantees we introduce some operators. The *convolution* $f * g$ of two functions f and g , is defined as $f * g(t) = \inf_{\tau \in [0, t]} \{f(t - \tau) + g(\tau)\}$, and the *deconvolution* $f \oslash g$ is defined as $f \oslash g(t) = \sup_{\tau \geq 0} \{f(t + \tau) - g(\tau)\}$.

Service guarantees to a flow at a network node are given in terms of service curves [10]. A *minimum service curve* is a function S which specifies a lower bound on the service given to a flow by $D(t) \geq A * S(t), \forall t \geq 0$. A *maximum service curve* for a flow is a function \bar{S} which specifies an upper bound on the service by $D(t) \leq A * \bar{S}(t), \forall t \geq 0$. In this paper, when service is allocated to flow aggregates we assume that the service curves are *strict* [5], in the sense that they guarantee the minimum deterministic service whenever a flow is backlogged, that is, $D(t_1, t_2) \geq S(t_2 - t_1)$ whenever $B(x) > 0$ for all $x \in (t_1, t_2)$.

To express probabilistic service guarantees, following [6], we define a (*minimum*) *effective service curve*² for an arrival process A as a nonnegative function \mathcal{S}^ε that satisfies for all $t > 0$,

$$Pr\{D(t) \geq A * \mathcal{S}^\varepsilon(t)\} \geq 1 - \varepsilon. \quad (1)$$

Given an effective service curve at a node, one can derive probabilistic bounds on backlog, delay, and the output process for effective service curves [6]. Specifically, the function $A^* \oslash \mathcal{S}^\varepsilon$ is a probabilistic bound for the departures on $[0, t]$, in the sense that, for all $t, \tau > 0$,

$$Pr\{D(t, t + \tau) \leq A^* \oslash \mathcal{S}^\varepsilon(\tau)\} \geq 1 - \varepsilon. \quad (2)$$

Similarly, \mathcal{S}^ε provides a backlog bound as

$$Pr\{B(t) \leq A^* \oslash \mathcal{S}^\varepsilon(0)\} \geq 1 - \varepsilon. \quad (3)$$

Finally,

$$d_{max} = \inf \{d \geq 0 \mid \forall t \geq 0 : A^*(t - d) \leq \mathcal{S}^\varepsilon(t)\} \quad (4)$$

²Henceforth, following the literature, the term ‘service curve’ refers to a minimum service curve, unless stated otherwise.

is a probabilistic delay bound that is violated with probability at most ε . By setting $\varepsilon = 0$, the above bounds correspond to the bounds of the deterministic calculus from [1].

If a flow passes through several nodes, each node guaranteeing an effective service curve, then the service offered by all nodes as a whole can be expressed as a convolution of the individual service curves.

Theorem 1: Effective Network Service Curve [6]. Consider a flow that passes through H network nodes in series. Let A^h and D^h denote the arrivals and departures at the h -th node, with $A^1 = A$, $A^h = D^{h-1}$ for $h = 2, \dots, H$ and $D^H = D$. Assume that effective service curves are given by nondecreasing functions $\mathcal{S}^{h, \varepsilon}$ at each node ($h = 1, \dots, H$). Further, assume that there exists a number $T \geq 0$ such that for all $t \geq 0$,

$$Pr\left\{D^h(t) \geq \inf_{x \in [0, T]} \{A^h(t - x) + \mathcal{S}^{h, \varepsilon}(x)\} \geq 1 - \varepsilon\right\} \quad (5)$$

for $h = 1, \dots, H$. Then, for any choice of $a > 0$,³

$$\mathcal{S}^{net, \varepsilon'} = \mathcal{S}^{1, \varepsilon} * \dots * \mathcal{S}^{H, \varepsilon} * \delta_{(H-1)a} \quad (6)$$

is an effective network service curve, with violation probability bounded by

$$\varepsilon' \leq H\varepsilon \left(1 + (H-1) \frac{T+a}{2a}\right). \quad (7)$$

Note that the theorem depends on a time scale T . The time scale can be established from a (deterministic or probabilistic) bound on the length of a busy period at a node, or from a priori backlog or delay bounds. Also note that the given effective network service curve deteriorates with the number of nodes H . We refer to [6] (cf. Theorem 5), which discusses when the assumption on the time scale T can be relaxed, and what the implications are of such a relaxation.

III. STATISTICAL MULTIPLEXING OF FLOWS

The fact that bursty or variable-rate traffic sources require less resources per flow when multiple flows are multiplexed is widely exploited for capacity provisioning of network traffic. In this paper,

³ δ_τ is the *impulse function* with $\delta_\tau(t) = \infty$ if $t > \tau$, and $\delta_\tau(t) = 0$ if $t \leq \tau$.

we express statistical multiplexing of flows using the notion of effective envelopes from [3]. Effective envelopes are functions that express probabilistic upper bounds for the traffic from an aggregate of flows. A desirable feature of effective envelopes as compared with other methods that express statistical multiplexing gain, e.g., effective bandwidth [7], [14], is that effective envelopes are easily related to the envelope functions used in the deterministic network calculus.

A. Effective Envelopes for a Flow Aggregate

Let us now consider a set \mathcal{C} of flows at a node, and let A_j and D_j , respectively, denote the arrival and departure processes for each flow $j \in \mathcal{C}$. We will refer to the set of flows as an aggregate of flows. Let $A_{\mathcal{C}}$ and $D_{\mathcal{C}}$ denote the aggregate arrivals and departures from the set \mathcal{C} at a network node, that is, $A_{\mathcal{C}}(t) = \sum_{j \in \mathcal{C}} A_j(t)$ and $D_{\mathcal{C}}(t) = \sum_{j \in \mathcal{C}} D_j(t)$.

A deterministic arrival envelope for the aggregate is simply given by $A_{\mathcal{C}}^*(t) = \sum_{j \in \mathcal{C}} A_j^*(t)$. However, such an envelope is pessimistic, and generally overestimates the bandwidth requirements of the aggregate. Therefore, we describe the traffic arrivals from an aggregate with a probabilistic bound, namely the *effective envelope* [3].

Given a set \mathcal{C} of flows with arrival process $A_{\mathcal{C}}$, an *effective envelope* for $A_{\mathcal{C}}$ is a function $\mathcal{G}_{\mathcal{C}}^{\varepsilon}$ such that for all $t, \tau \geq 0$,

$$\Pr\{A_{\mathcal{C}}(t, t + \tau) \leq \mathcal{G}_{\mathcal{C}}^{\varepsilon}(\tau)\} \geq 1 - \varepsilon. \quad (8)$$

A *strong effective envelope* for $A_{\mathcal{C}}$ for intervals of length ℓ is a subadditive function $\mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon}$ such that for each interval I_{ℓ} of length ℓ ,

$$\Pr\left\{\forall [t, t + \tau] \subseteq I_{\ell} : A_{\mathcal{C}}(t, t + \tau) \leq \mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon}(\tau)\right\} \geq 1 - \varepsilon. \quad (9)$$

Thus, an effective envelope provides a stationary bound for arrivals, which is violated with probability at most ε . A strong effective envelope is, in addition, a uniform bound for all subintervals in a larger interval.⁴ Strong effective envelopes are used

⁴In [3], the effective envelope is called *local effective envelope* and the strong effective envelope is called *global effective envelope*.

in our construction of effective service curves for a single flow with aggregate provisioning, and are constructed from an effective envelope.

B. Constructing Effective Envelopes for Heterogeneous Traffic

To construct effective envelopes for an aggregate of flows we consider an adversarial traffic model [12], where arrivals of flows to the network can individually exhibit a worst-case arrival pattern as allowed by assumption (A1), but sources do not conspire to construct a joint worst-case. In addition to assumption (A1) from Section II, we assume that the following hold for the arrival processes.

- (A2) *Stationarity*.⁵ For all $t \geq 0, t' \geq 0$ and for any $\tau > 0$ and any $x \geq 0$, $\Pr\{A_i(t, t + \tau) \leq x\} = \Pr\{A_i(t', t' + \tau) \leq x\}$.
- (A3) *Independence*. The arrivals from two flows $i, j \in \mathcal{C}$, A_i and A_j , are stochastically independent.

Since buffering and scheduling distort traffic and introduce correlations between flows, we only require assumptions (A2)–(A3) to hold for the arrivals at the first node of a flow's route through the network.

The following construction extends the derivations in [3] to an aggregate of flows with heterogeneous arrival envelopes, and is based on an application of the Chernoff bound [16]. By heterogeneity, we mean that the arrival envelopes A_i^* of flows can be different for each flow. The construction of effective envelopes $\mathcal{G}_{\mathcal{C}}^{\varepsilon}$ for a set \mathcal{C} of flows uses the moment generating function of A_j , denoted as $M_j(s, t) = E[e^{A_j(\tau, \tau + t)s}]$. As shown in [3], if assumptions (A1)–(A3) hold, we obtain $M_j(s, t) \leq \overline{M}_j(s, t)$, where

$$\overline{M}_j(s, t) = 1 + \frac{\rho_j t}{A_j^*(t)} \left(e^{s A_j^*(t)} - 1 \right). \quad (10)$$

With assumption (A3) and with the bound in Eqn. (10), we obtain from the Chernoff bound that

$$\Pr\{A_{\mathcal{C}}(t) \geq x\} \leq e^{-xs} \prod_{j \in \mathcal{C}} \overline{M}_j(s, t). \quad (11)$$

⁵We only need the stationary bound $E[A_j(t, t + \tau)] \leq \rho_j \tau$ where $\rho_j := \lim_{\tau \rightarrow \infty} A_j^*(\tau)/\tau$. This bound follows from (A1) and (A2).

Setting the right hand side equal to ε and solving for x gives

$$x = \frac{1}{s} \left(\sum_{j \in \mathcal{C}} \log \bar{M}_j(s, t) - \log \varepsilon \right). \quad (12)$$

For any choice of $s > 0$, Eqn. (12) is an effective envelope for the arrivals from \mathcal{C} . We select the value of the effective envelope at t to be

$$\mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) = \inf_{s > 0} \frac{1}{s} \left(\sum_{j \in \mathcal{C}} \log \bar{M}_j(s, t) - \log \varepsilon \right). \quad (13)$$

With this choice, $\mathcal{G}_{\mathcal{C}}^{\varepsilon}(t) \leq A_{\mathcal{C}}^*(t)$ is always satisfied. Since the derivative of the right hand side of Eqn. (13) is increasing in s , there is at most one minimum, which can be found by searching for the zero of the derivative. Note the similarity of Eqn. (13) to the effective bandwidth in [7], [14].

Given an effective envelope $\mathcal{G}_{\mathcal{C}}^{\varepsilon}$ for a set \mathcal{C} , we can construct a strong effective envelope $\mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon}$.

Lemma 1: Given an effective envelope $\mathcal{G}_{\mathcal{C}}^{\varepsilon}$ for a set \mathcal{C} of heterogeneous flows satisfying (A1) – (A3). There exists a strong effective envelope for the arrivals in \mathcal{C} satisfying

$$\mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon'}(t) \leq \mathcal{G}_{\mathcal{C}}^{\varepsilon}(\gamma t + a), \quad 0 \leq t \leq \ell,$$

where $\gamma > 1$ and $a \in (0, \ell)$ are arbitrary parameters, and

$$\varepsilon' \leq \varepsilon \cdot \frac{\ell \sqrt{\gamma} + 1}{a \sqrt{\gamma} - 1}.$$

The quality of the bound in the lemma for a given value of t depends on the selection of the two parameters a and γ . To get a close-to-optimal bound for a time scale near t^* , one should choose $a \approx \sqrt{\gamma}(\gamma - 1)t^*$. The largest subadditive function below $\mathcal{G}_{\mathcal{C}}^{\varepsilon}(\gamma t + a)$ is a strong effective envelope.

Proof: Fix $a_0 \in (0, \ell)$ and $\gamma_0 > 1$, and set

$$\begin{aligned} \tau_i &= a_0 \frac{\gamma_0^i - 1}{\gamma_0 - 1} \\ x_i &= \tau_{i+1} - \tau_i = a_0 \gamma_0^i \end{aligned}$$

where $i = 1, 2, \dots, n$ and n is the smallest number with $\tau_n \geq \ell$. Consider the intervals $I_{ij} = [jx_i, jx_i + \tau_{i+1}]$ for $j = 0, 1, \dots, \lceil \frac{\ell - \tau_{i+1}}{x_i} \rceil$. If $\tau \leq \tau_i$, then every interval $I_{\tau} \subseteq [0, \ell]$ of length τ is contained

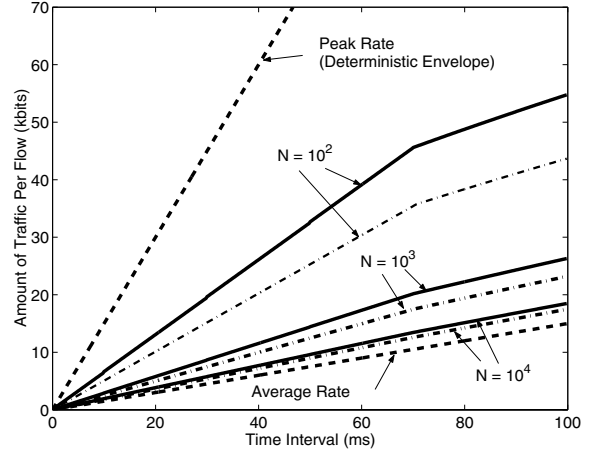


Fig. 1. Comparison of effective envelopes (dash-dotted lines) and strong effective envelopes (solid lines) for Type-1 flow from Section V for $t \leq 100$ ms, $\varepsilon = 10^{-9}$, and for number of flows $N = 100, 1000, 10000$. The curves show the normalized curves $\mathcal{G}_{\mathcal{C}}^{\varepsilon}/|\mathcal{C}|$ and $\mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon}/|\mathcal{C}|$ ($|\mathcal{C}| = N$). For constructing $\mathcal{H}_{\mathcal{C}}^{\ell, \varepsilon}$, we use $\ell = 2$ sec, $\gamma = 1.01$ and $t^* = 10$ ms.

in one of the I_{ij} . The total number N of intervals I_{ij} is bounded by

$$N \leq \sum_{i=1}^n \left\lceil \frac{\ell - \tau_{i+1}}{x_i} \right\rceil + 1 \quad (14)$$

$$\leq \sum_{i=1}^{\infty} \frac{\ell}{a_0 \gamma_0^i} = \frac{\ell}{a_0(\gamma_0 - 1)}. \quad (15)$$

It follows that

$$\begin{aligned} &Pr\{\exists [x, x + \tau] \subseteq [0, \ell] : \\ &A(x, x + \tau) > \mathcal{G}^{\varepsilon}(\gamma_0^2 + (1 + \gamma_0)a_0)\} \\ &\leq Pr\{\exists i, j : A(jx_i, jx_i + \tau_{i+1}) > \mathcal{G}^{\varepsilon}(\tau_{i+1})\} \\ &\leq N\varepsilon. \end{aligned} \quad (16)$$

Eqn. (16) holds since $\tau_{i-1} < \tau \leq \tau_i$ implies

$$\tau_{i+1} = \gamma_0^2 \tau_{i-1} + (\gamma_0 + 1)a_0 \quad (18)$$

$$\leq \gamma_0^2 \tau + (\gamma_0 + 1)a_0. \quad (19)$$

Setting $\gamma_0 = \sqrt{\gamma}$ and $a_0 = \frac{a}{\sqrt{\gamma} + 1}$ completes the proof. \square

In Figure 1 we illustrate the multiplexing gain captured by effective and strong effective envelopes. We plot the effective envelope for N identical flows with parameters as given in Section V, normalized by the number N of flows. The results show that, as N grows large, the effective envelope is close to the average traffic rate.

C. Service Provisioning and Busy Period Estimate

We assume that the service allocated to the aggregate arrivals is given by deterministic service curves. A lower bound of the service given to the aggregate is expressed in terms of a strict minimum service curve S_C and an upper bound for the service is given by a maximum service curve \bar{S}_C . We do not make assumptions on a specific scheduling algorithm at the node, as long as it can guarantee a strict minimum service curve. For example, any work-conserving scheduling algorithm that allocates a constant rate to the aggregate of flows delivers a strict minimum service guarantee. Generally, scheduling algorithms do not specify an explicit maximum service curve. In these cases, a maximum service curve of the aggregate can be given as $\bar{S}_C(t) = Ct$, where C is the output link capacity of the node.

With the traffic characterization A_C^* and the minimum service guarantee S_C we can derive bounds on the busy period at a node. We define the *busy period* for a given time t as the maximal time interval containing t during which the backlog from the flows in C remains positive. The beginning of the busy period for a time t is denoted by \underline{t} with $\underline{t} = \sup\{\tau \leq t : B(\tau) = 0\}$. For a strict service curve S_C , the deterministic network calculus yields that the number

$$T_0 = \inf\{\tau > 0 \mid A_C^*(\tau) \leq S_C(\tau)\} \quad (20)$$

provides an upper bound for the length of any busy period [5], i.e., for all $t \geq 0$,

$$t - \underline{t} \leq T_0. \quad (21)$$

Therefore, in any interval $[t - T_0, t]$, the backlog must be zero at least once. Since the bound T_0 can be conservative, the following lemma can be used to find a less conservative estimate.

Lemma 2: Consider an aggregate C of flows with given arrival and departure processes $A_C(t)$ and $D_C(t)$, a strict service curve $S_C(t)$, and T_0 as given in Eqn. (20), for all $t \geq 0$. Let $\mathcal{H}_C^{T_0, \varepsilon}$ be a strong effective envelope for the arrivals in time intervals of length T_0 , for some $\varepsilon < 1$. If there exists a number $T^\varepsilon < T_0$ such that

$$\mathcal{H}_C^{T_0, \varepsilon}(T^\varepsilon) \leq S_C(T^\varepsilon), \quad (22)$$

then T^ε is a probabilistic bound on the busy period, in the sense that for all $t \geq 0$,

$$\Pr\{t - \underline{t} \leq T^\varepsilon\} \geq 1 - \varepsilon. \quad (23)$$

Proof: Fix $t > 0$. By Eqn. (20) we have $t - \underline{t} \leq T_0$. We compute

$$\begin{aligned} & \Pr\{t - \underline{t} > T^\varepsilon\} \\ & \leq \Pr\{B_C(\underline{t} + T^\varepsilon) > 0 \text{ and} \\ & \quad D_C(\underline{t}, \underline{t} + T^\varepsilon) \geq S_C(T^\varepsilon)\} \end{aligned} \quad (24)$$

$$\leq \Pr\{A_C(\underline{t}, \underline{t} + T^\varepsilon) > \mathcal{H}_C^{T_0, \varepsilon}(T^\varepsilon)\} \quad (25)$$

$$\leq \varepsilon. \quad (26)$$

Eqn. (24) uses that the backlog is positive in the entire interval (\underline{t}, t) , and the definition of the strict service curve S_C . Eqn. (25) uses that $B_C(\underline{t} + T^\varepsilon) > 0$ implies $A_C(\underline{t}, \underline{t} + T^\varepsilon) > D_C(\underline{t}, \underline{t} + T^\varepsilon)$, and the assumption that $S_C(T^\varepsilon) \geq \mathcal{H}_C^{T_0, \varepsilon}(T^\varepsilon)$. Eqn. (26) follows by applying the definition of $\mathcal{H}_C^{T_0, \varepsilon}$ to the interval $I_{T_0} = [t - T_0, t]$, which contains the interval $[\underline{t}, \underline{t} + T^\varepsilon]$ by assumption. This proves the claim. \square

The lemma allows us to replace a deterministic bound on the busy period by a (possibly less pessimistic) probabilistic bound. A similar argument can be used to improve a given probabilistic bound. If, in all assumptions of the lemma, the deterministic bound T_0 is replaced by a probabilistic bound T_1 , which satisfies $\Pr\{t - \underline{t} \leq T_1\} \geq 1 - \varepsilon_1$, then we can construct another probabilistic bound T_2 , with $\Pr\{t - \underline{t} \leq T_2\} \geq 1 - (\varepsilon_1 + \varepsilon_2)$. We can thus recursively define a sequence of probabilistic bounds T_n on the busy period satisfying $\Pr\{t - \underline{t} \leq T_n\} \geq 1 - n\varepsilon$. Since the bound T_n decreases with n while the violation probability $n\varepsilon$ increases, one needs to pick a ‘good’ value for n .

In Figure 2, we illustrate the busy period estimates for a link with capacity $C = 100$ Mbps with Type-1 traffic as described in Section V. The figure shows the deterministic bound T_0 according to Eqn. (20) and the probabilistic bound $T_1 = T^\varepsilon$ from Lemma 2 with $\varepsilon = 10^{-9}$. The figure also includes a bound T_2 with probability $1 - 2\varepsilon$ obtained by a repeated application of Lemma 2. The busy periods are evaluated for $N = 250 - 500$ flows. The graph indicates that a single application of Lemma 2 significantly reduces the estimate on the

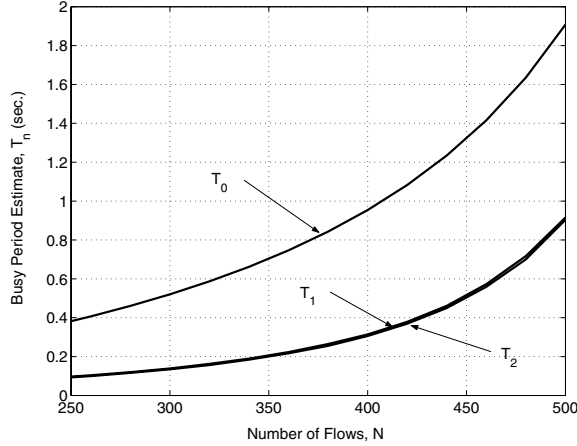


Fig. 2. Busy period estimates with Type-1 flow traffic from Section V for a link with $C = 100$ Mbps. The number T_n is a bound for the busy period with probability $1 - n\varepsilon$.

busy period, whereas successive iterations of the lemma do not result in noticeable improvements.

The busy period bound can be applied in Theorem 1 to give a bound on the range of the infimum in Eqn. (5). Note, however, that the above busy period bounds assume deterministic arrival envelopes, which are generally only available for the first node in the network.

For strict service curves S_C , we can directly provide busy period bounds when a priori bounds on backlog or delays are available [6]. Given a backlog bound b^* that satisfies $Pr\{B(t) > b^*\} \leq \varepsilon_1$ and T such that $\mathcal{G}_C^{\varepsilon_2}(T) + b^* \leq S_C(T)$, or given a delay bound d^* that satisfies $Pr\{D(t) > A_C(t - d^*)\} \leq \varepsilon_1$ and T such that $\mathcal{G}_C^{\varepsilon_2}(T + d^*) \leq S_C(T)$, then one can obtain from the definition of the strict service curve that $Pr\{t - \underline{t} > T\} \leq \varepsilon_1 + \varepsilon_2$.

IV. EFFECTIVE SERVICE FOR A FLOW WITH AGGREGATE SERVICE

Given the aggregate of flows and given the service provisioned to the aggregate as described in the previous section, we now address the problem of determining an effective service curve for a single flow from the aggregate. The effective service curve that we construct expresses the service as seen by a single flow in terms of a probabilistic bound. The basic idea is to construct the effective service curve for an individual flow from the leftover capacity that is not used by the aggregate. Roughly speaking, the main result is that a probabilistic lower bound for the service of a single flow from

the aggregate can be given by the effective service curve $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2} = [S_C - \mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}]_+$, where $\varepsilon_1 + \varepsilon_2$ is the probability that the effective service curve is violated, S_C is the service curve for the aggregate, $\mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}$ is a strong effective envelope for the aggregate, and T^{ε_1} is a probabilistic busy period bound. This effective service curve will be obtained from the following theorem.

Theorem 2: Effective service curve for a single flow in an aggregate of flows. Given a set \mathcal{C} of flows with a strict deterministic service curve S_C . Assume that the set $\mathcal{C} - \{j\}$ is allocated a deterministic maximum service curve \bar{S}_{-j} . Assume a probabilistic bound T^{ε_1} on the busy period as given in Lemma 2 and let $\mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}}$ denote a strong effective envelope for the arrivals from $\mathcal{C} - \{j\}$ for time intervals of length T^{ε_1} . Then the function defined on the interval $[0, T^{\varepsilon_1}]$ by

$$\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2} = [S_C - \mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}} * \bar{S}_{-j}]_+ \quad (27)$$

is an effective service curve for flow $j \in \mathcal{C}$, with violation probability $\varepsilon_1 + \varepsilon_2$.⁶ More precisely,

$$Pr \left\{ D_j(t) \geq \inf_{x \leq T^{\varepsilon_1}} \{A_j(t - x) + \mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}(x)\} \right\} \geq 1 - (\varepsilon_1 + \varepsilon_2). \quad (28)$$

Thus, a probabilistic service allocation for a single flow can be obtained from the service allocation of an aggregate by subtracting a probabilistic upper bound on the departures from all other flows.

The function $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}$ as defined in Eqn. (27) need not be monotonic in t . However, one can make $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}$ monotonic by replacing it with the largest nondecreasing function below $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}$. If S_C is convex and $\mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}}$ is concave, then $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}$ is always convex and nondecreasing.

Proof: Fix $t \geq 0$. Then

$$D_j(\underline{t}, t) = D_C(\underline{t}, t) - D_{-j}(\underline{t}, t) \quad (29)$$

$$\geq S_C(t - \underline{t}) - A_{-j} * \bar{S}_{-j}(t) + A_{-j}(\underline{t}) \quad (30)$$

$$\geq S_C(t - \underline{t}) - \inf_{x \leq t - \underline{t}} \{A_{-j}(\underline{t}, t - x) + \bar{S}_{-j}(x)\} \quad (31)$$

In Eqn. (30), we have used the definition of the service curves S_C and \bar{S}_{-j} , and the fact that $D_{-j}(\underline{t}) =$

⁶We use “[f]₊(t) = max{ $f(t)$, 0}”.

$A_{-j}(\underline{t})$ by definition of \underline{t} . In Eqn. (31), we have expanded the convolution operator, restricted the range of the infimum, and taken A_{-j} under the infimum. Since $D_j(\underline{t}) = A_j(\underline{t})$ and $D_j(\underline{t}, t) \geq 0$, this implies

$$D_j(t) \geq A_j(\underline{t}) + [S_C(t - \underline{t}) - \inf_{x \leq t - \underline{t}} \{A_{-j}(\underline{t}, t - x) + \bar{S}_{-j}(x)\}]_+ . \quad (32)$$

Let $\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2} = [S_C - (\mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}} * \bar{S}_{-j})]_+$ be the function defined in the statement of the theorem. Then

$$Pr\{D_j(t) \geq \inf_{x \leq T^{\varepsilon_1}} \{A_j(t - x) + \mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}(x)\}\} \quad (33)$$

$$\geq Pr \left\{ \begin{array}{l} D_j(t) \geq A_j(\underline{t}) + \mathcal{S}_j^{\varepsilon_1 + \varepsilon_2}(t - \underline{t}) \\ \text{and } t - \underline{t} \leq T^{\varepsilon_1} \end{array} \right\} \quad (34)$$

$$\geq Pr \left\{ \begin{array}{l} \forall t_0 \in [t - T^{\varepsilon_1}, t] : \\ A_{-j}(t_0, t) \leq \mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}}(t - t_0) \\ \text{and } t - \underline{t} \leq T^{\varepsilon_1} \end{array} \right\} \quad (35)$$

$$\geq 1 - (\varepsilon_1 + \varepsilon_2) , \quad (36)$$

as claimed. \square

Since the given effective service curve for the aggregate applies to all (workconserving) scheduling algorithms, it is pessimistic for most, particularly, FIFO. It is least conservative if flows in the set $\mathcal{C} - \{j\}$ are transmitted with higher priority than flow j .

There is a corresponding formulation of a per-flow service curve in the deterministic calculus, which can be found in [5], [11]. With the same notation as in Theorem 2, we can write a deterministic per-flow service guarantee as

$$S_j = [S_C - A_{-j}^* * \bar{S}_{-j}]_+ \quad (37)$$

However, since $A_{-j}^* = \sum_{k \in \mathcal{C}, k \neq j} A_k^*$ is large and does not exploit statistical multiplexing, the bounds with such a service curve may not have practical relevance.

The following corollary states the previously mentioned simpler bound on the minimum effective service to flow j . This bound does not assume that a maximum service curve is available and estimates $\mathcal{H}_{-j}^{T^{\varepsilon_1, \varepsilon_2}}$ by $\mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}$.

Corollary 1: The conclusions of Theorem 2 hold without change for

$$\mathcal{S}_j^{\varepsilon_1 + \varepsilon_2} = [S_C - \mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}]_+ .$$

Proof. The claim holds since $\mathcal{H}_C^{T^{\varepsilon_1, \varepsilon_2}}$ is a strong effective envelope for the flows in $\mathcal{C} - \{j\}$, and since $f \geq f * g$ follows by the definition of the convolution operator. \square

Thus, we have derived a lower bound on the effective service to a flow even when information is available only about the aggregate of flows. In the next section, we will see that even with the pessimistic lower bound of this section, we are able to extract a significant amount of multiplexing gain.

V. EVALUATION

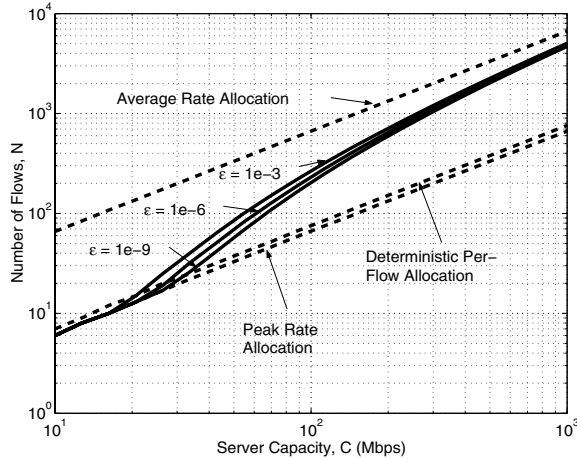
We now present numerical examples for the effective service curve. We assume that individual flows are regulated at the entrance to the network by using a peak rate limited leaky bucket with arrival envelope $A_j^*(\tau) = \min\{P_j\tau, \sigma_j + \rho_j\tau\}$ for flow j , where P_j is the peak rate, $\rho_j \leq P_j$ is the average rate, and σ_j is a burst size parameter. We consider two types of flows with parameters as given in the following table:

Type	Peak Rate P_j (Mbps)	Mean Rate ρ_j (Mbps)	Burst Size σ_j (bits)
Type 1	1.5	0.15	95400
Type 2	6.0	0.15	10345

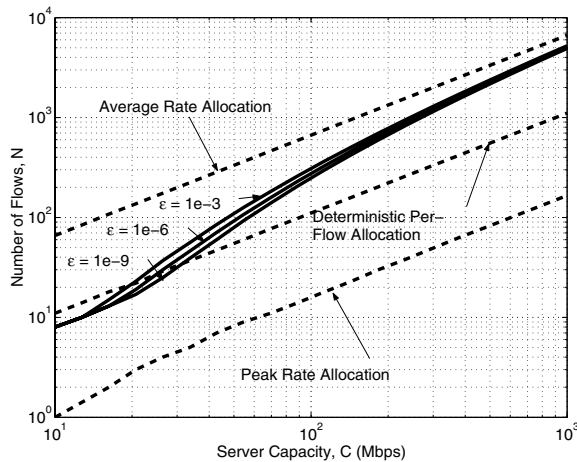
In the following we use A_1^* and A_2^* to denote the arrival envelope of a Type-1 and a Type-2 flow, respectively. We assume that the arrivals satisfy assumptions (A1)–(A3), and we construct effective envelopes as shown in Subsection III-B.

Service curves for the aggregate have a constant-rate form. We set $S_C(t) = Nc t$, where $c > 0$ is referred to as ‘per-flow capacity’ and $N = |\mathcal{C}|$ is the number of flows. We assume that the maximum service curve is given by $\bar{S}_C(t) = C t$, where C is the link capacity.

For the calculation of strong effective envelopes and effective service curves we use the busy period bounds from Subsection III-C. We use a busy period bound of $T = 2$ sec, which satisfies the deterministic busy period bound in the sense of Eqn. (21). By Lemma 1, we set $\gamma = 1.01$ and



(a) Type-1 Flows.



(b) Type-2 Flows.

Fig. 3. Example 1: Number of flows admitted on a link with capacity C to satisfy a delay bound of $d = 10$ ms with probability $1 - \varepsilon$.

$t^* = 10$ ms.⁷ For the construction of effective service curves $\mathcal{S}_j^\varepsilon$ we apply the simpler and more conservative bound from Corollary 1, rather than the bound from Theorem 2.

We compare the results obtained with effective service curves with the deterministic bound from Eqn. (37), as well as to the following non-statistical per-flow service provisioning schemes.

- A *peak rate allocation*, where each flow j has a service curve of $S_j(t) = P_j t$.
- An *average rate allocation*, where each flow j has a service curve of $S_j(t) = \rho_j t$.
- A *deterministic per-flow allocation* which de-

⁷In our examples, the numerical computations for effective envelopes are done in discrete intervals of length $\Delta = 0.2$ ms, and not in continuous time. This may introduce discretization errors.

livers worst-case delay guarantees to each flow. The resources allocated to a flow are determined by the smallest (deterministic) constant-rate service curve $S_j(t) = \hat{c}_j t$ that satisfies the delay bound d , i.e., $\hat{c}_j = \inf \{c \geq 0 \mid \forall t \geq 0 : A_j^*(t-d) \leq c t\}$. This allocation method assumes that the network nodes perform per-flow scheduling operations.

A. Example 1: Effective Service Curves at a Single Node

We consider an aggregate of N flows, where flows are either all Type-1 or all Type-2 flows. We use the effective service curve from Corollary 1 to determine how many flows can be provisioned on a link to meet a required service guarantee. Specifically, we determine how many flows can be put on a link with capacity C such that the probabilistic delay bound in Eqn. (4) does not exceed $d = 10$ ms. To do this, we find the largest N such that $\mathcal{S}_j^\varepsilon(t) = [Ct - \mathcal{H}_C^{T,\varepsilon}(t)]_+$ assures via Eqn. (4) a delay bound d with probability $1 - \varepsilon$.

In Figures 3(a) and 3(b), we show the number of flows that can receive the probabilistic delay bound of $d = 10$ ms as a function of the network capacity. The figures include plots using effective service curves with $\varepsilon = 10^{-3}, 10^{-6}, 10^{-9}$. We also include results for an average rate allocation (which does not satisfy the delay bound), the results for a peak rate allocation, and the results for a deterministic per flow allocation, which is $S_j(t) \approx 1.3140 t$ for Type-1 flows and $S_j(t) \approx 0.9016 t$ for Type-2 flows.

For small C , the number of flows is too small to extract multiplexing gain, and, consequently, the effective service curve constructed from Corollary 1 is inferior to a per-flow deterministic service curve allocation. On the other hand, when C grows large, the number of flows that can be admitted with the effective service curves are close to that of an average rate allocation.

B. Example 2: Multiple Nodes with Cross Traffic

We consider a network with two nodes, as shown in Figure 4, and determine the network service curve for a flow in this network. There are N_1 flows

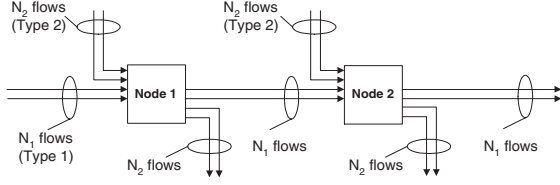


Fig. 4. Example 2: A network with two nodes and with cross traffic. The cross-traffic consists of N_2 Type-2 flows at the first node and at the second node.

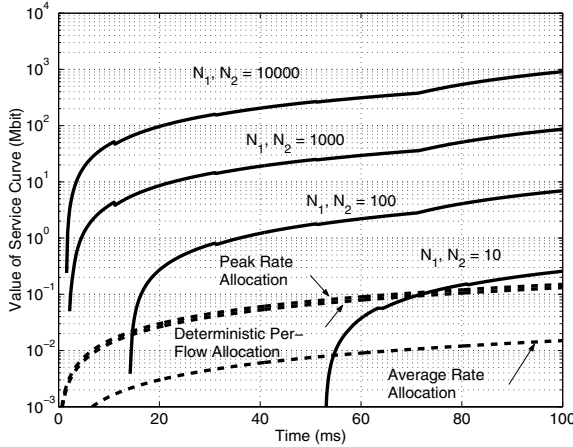


Fig. 5. Example 2: Network effective service curve for different values of N_1 , and are calculated for $\varepsilon = 10^{-9}$.

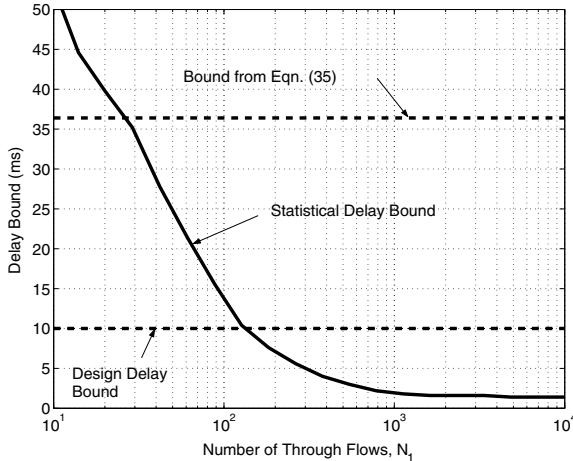


Fig. 6. Example 2: Achieved delay bounds with the effective service curve.

from Type-1 flows that pass through both nodes. We refer to these flows as ‘through flows’. At each node there are N_2 cross flows from Type-2. We set the number of through and cross flows to be equal at each node ($N_1 = N_2$).

We assume that the service guarantee for the aggregate of $N_1 + N_2$ flows at the first node is $S_C(t) = (N_1\hat{c}_1 + N_2\hat{c})t$, where $\hat{c}_1 = 1.3140$ and $\hat{c}_2 = 0.9016$. These service curves guarantee a deterministic delay bound of $d = 10$ ms for each

Type-1 or Type-2 flow, if each flow receives a rate allocation of \hat{c}_1 or \hat{c}_2 , respectively.

With this aggregate allocation we now construct effective service curves for different values of N_1 , N_2 , and for $\varepsilon = 10^{-9}$. The effective service curves for a Type-1 through flow at the first node and second node, respectively, are denoted by $\mathcal{S}_1^{1,\varepsilon}$ and $\mathcal{S}_1^{2,\varepsilon}$, and are given by $S^{h,\varepsilon}(t) = [(N_1\hat{c}_1 + N_2\hat{c})t - \mathcal{H}_C^{h,T,\varepsilon}(t)]_+$ ($h = 1, 2$) according to Corollary 1, where $\mathcal{H}_C^{h,T,\varepsilon}$ is the strong effective envelope of the aggregate of all flows at the h -th node. In the effective service curve, since T is a deterministic bound, the violation probability of the busy period is zero. Once we have the effective service curve for each node, we can calculate the network service curve $\mathcal{S}_1^{net,\varepsilon} = \mathcal{S}_1^{1,\varepsilon} * \mathcal{S}_1^{2,\varepsilon} * \delta_{(H-1)a}$, according to Theorem 1 with $H = 2$ and $a = 1$ ms, and determine probabilistic delay bounds from Eqn. (4).

The calculation of $\mathcal{H}_C^{1,T,\varepsilon}$ is straightforward as described in Subsection III-B. The calculation of $\mathcal{H}_C^{2,T,\varepsilon}$ at the second node, however, requires some thought. Note that, for the through flows, we cannot assume that the arrivals at the second node are independent. Therefore, we need to consider the entire set of N_1 through flows as a group. At the first node, the group of N_1 Type-1 flows obtains an effective service curve of $\mathcal{S}_{N_1}^{1,\varepsilon}(t) = [(N_1\hat{c}_1 + N_2\hat{c})t - \mathcal{H}_{N_2}^{1,T,\varepsilon}(t)]_+$, where the subscripts (in a slight abuse of notation) indicate the set of flows, and where $\mathcal{H}_{N_2}^{1,T,\varepsilon}$ is a strong effective envelope for the Type-2 flows at the first node, which is constructed as described in Subsection III-B. With the statistical bound in Eqn. (2) we can construct an effective envelope for the Type-1 departures from the first node by $A_{N_1}^* \circ \mathcal{S}_{N_1}^{1,\varepsilon}$. This effective envelope can be turned into a strong envelope via Lemma 1. Let us call this envelope $\mathcal{H}_{N_1}^{2,T,\varepsilon}$. The strong effective envelope for the Type-2 flows at the second node, denoted by $\mathcal{H}_{N_2}^{2,T,\varepsilon}$, is same as $\mathcal{H}_{N_2}^{1,T,\varepsilon}$. Finally, the sum $\mathcal{H}_C^{2,T,2\varepsilon} = \mathcal{H}_{N_1}^{2,T,\varepsilon} + \mathcal{H}_{N_2}^{2,T,\varepsilon}$ gives us a strong effective envelope for all arrivals at the second node. As a last issue, we need to ensure that T is selected so that it is greater than the busy period at each node. Here, we use a worst-case estimate (using the deterministic network calculus) to provide us with such a bound.

In Figure 5, we plot the resulting effective net-

work service curves for $N_1, N_2 = 100, 1000, 10000$ flows. A comparison shows that for $N_1 > 100$, the statistical lower bound of the service to a single through flow offers more service than a deterministic per-flow allocation. Thus, when the number of flows is large, the aggregate provisioning will result in lower, albeit statistical, delay bounds. This point is emphasized in Figure 6, where we plot the statistical delay bounds achieved with the network service curve by applying Eqn. (4) as a function of the number of flows N_1 . If $N_1 > 100$, the delay bounds offered by our statistical lower bound on the service are better than a deterministic per-flow allocation.

VI. CONCLUSION

We have presented a method to compute statistical lower bounds on the service given to a single flow in a network in which service is provisioned to aggregates of flows. The lower bounds assume that the service allocated to the aggregate workconserving, but does not assume that the scheduling algorithms in the network are known. The derivations were done in the context of a statistical network calculus that expresses the lower bound of the service given to a flow in terms of an effective service curve. By describing the service at each node with an effective service curve, the service given by the network as a whole is simply expressed as a concatenation of the per-node service curves. Thus, we were able to derive probabilistic end-to-end guarantees. The calculus requires that bounds on the busy period, or a priori bounds on backlog or delay are available. We have shown how such bounds can be derived.

REFERENCES

- [1] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan. Performance bounds for flow control protocols. *IEEE/ACM Transactions on Networking*, 7(3):310–323, June 1999.
- [2] M. Andrews. Instability of FIFO in session-oriented networks. In *Proceedings of the Eleventh Annual ACM-SIAM Symposium on Discrete Algorithms (SODA 2000)*, pages 440–447, October 1996.
- [3] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Statistical service assurances for traffic scheduling algorithms. *IEEE Journal on Selected Areas in Communications. Special Issue on Internet QoS*, 18(12):2651–2664, December 2000.
- [4] A. Borodin, J. M. Kleinberg, P. Raghavan, J. Sudan, and D. P. Williamson. Adversarial queuing theory. *Journal of the ACM*, 48(1):13–38, 2001.
- [5] J. Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [6] A. Burchard, J. Liebeherr, and S. D. Patek. A calculus for end-to-end statistical service guarantees. Technical Report CS-2001-19, University of Virginia, Computer Science Department, May 2002.
- [7] C. S. Chang. Stability, queue length, and delay of deterministic and stochastic queueing networks. *IEEE Transactions on Automatic Control*, 39(5):913–931, May 1994.
- [8] C. S. Chang, Y. Chiu, and W. T. Song. On the performance of multiplexing independent regulated inputs. In *Proceedings of ACM SIGMETRICS 2001*, pages 184–193, 2001.
- [9] A. Charny and J. Y. Le Boudec. Delay bounds in a network with aggregate scheduling. In *Proceedings of QOFIS*, October 2000.
- [10] R. L. Cruz. Quality of service guarantees in virtual circuit switched networks. *IEEE Journal on Selected Areas in Communications*, 13(6):1048–1056, August 1995.
- [11] R. L. Cruz. SCED+: Efficient management of quality of service guarantees. In *IEEE Infocom*, pages 625–634, 1998.
- [12] A. Elwalid, D. Mitra, and R. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous, regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.
- [13] V. Firoiu, J.-Y. Le Boudec, D. Towsley, and Z.-L. Zhang. Theories and models for internet quality of service. *Proceedings of the IEEE*, 2002. To appear.
- [14] F. Kelly. Notes on effective bandwidths. In *Stochastic Networks: Theory and Applications*. (Editors: F.P. Kelly, S. Zachary and I.B. Ziedins) *Royal Statistical Society Lecture Notes Series*, 4, pages 141–168. Oxford University Press, 1996.
- [15] G. Kesidis and T. Konstantopoulos. Extremal traffic and worst-case performance for queues with shaped arrivals. In *Proceedings of Workshop on Analysis and Simulation of Communication Networks*, Toronto, November 1998.
- [16] A. Papoulis. *Probability, Random Variables, and Stochastic Processes*. 3rd edition. McGraw Hill, 1991.
- [17] L. Tassiulas and L. Georgiadis. Any work-conserving policy stabilizes the ring with spatial re-use. *IEEE/ACM Transactions on Networking*, 4(2):205–208, April 1996.
- [18] M. Vojnovic and J.-Y. Le Boudec. Bounds for independent regulated inputs multiplexed in a service curve network element. In *Proceedings of IEEE Globecom 2001*, San Antonio, Texas, November 2001.
- [19] M. Vojnovic and J.-Y. Le Boudec. Elements of a probabilistic network calculus for packet scale rate guarantee nodes. In *Proceedings of Fifteenth International Symposium on Mathematical Theory of Networks and Systems, University of Notre Dame*, 2002.
- [20] Z. Zhang, Z. Duan, and T. Hou. Fundamental trade-offs in aggregate packet scheduling (extended abstract). In *Proceedings of SPIE Vol. 4526*, August 2001.