

A Case for Decomposition of FIFO Networks

Florin Ciucu and Jörg Liebeherr
Department of Electrical and Computer Engineering
University of Toronto

Abstract—Recent studies showing that the output of traffic flows at packet switches has similar characteristics as the corresponding input enables a *decomposition* analysis of a network where nodes can be studied in isolation, thus simplifying an end-to-end analysis of networks. However, network decomposition results available today are mostly many-sources asymptotics. In this paper we explore the viability of network decomposition in a non-asymptotic regime with a finite number of flows. For traffic with Exponentially Bounded Burstiness (EBB) we derive statistical bounds for the output traffic at a FIFO buffer and compare them with bounds on the input. Evaluating the accuracy of the output bounds with exact results available for special cases and by numerical illustrations we find that conditions for network decomposition appear favorable even if the number of flows is relatively small.

I. INTRODUCTION

A fundamental problem encountered in analyzing network performance is that factors such as finite node capacity or scheduling distort the pattern of arriving traffic. As a direct consequence, product-form formulas which exactly express the performance of queueing networks are generally restricted to BCMP [1] and Kelly [16] networks.

For more general networks, recent works [11], [26] have shown that it is possible to analyze a network node by disregarding the distortion effects experienced by the node's arrivals at previously traversed nodes. In this way networks can be *decomposed* [11] and a node can be analyzed in isolation of the others.

The common idea behind network decomposition is that the output of traffic flows has similar characteristics as the input. For general scheduling algorithms the decoupling of effective bandwidths of flows was proven in large-buffer [25] and many-sources [26] asymptotic regimes. Under the many-sources asymptotic regime it was shown that a downstream queue behaves similarly as if the upstream queue capable of serving many flows was removed [11]. Using simulations it was observed in [14] that, unlike more complex scheduling mechanisms such as round-robin (RR) or weighted-RR, FIFO is adequate for networks with CBR service and that traffic reshaping is not necessary. In [2] was conjectured that flows in FIFO networks do not become burstier when replaced by Poisson input processes, i.e., the flows are said to be better-than-Poisson. The burstiness of flows was shown to be preserved across multiple nodes with FIFO scheduling in a many-sources asymptotic regime, even if the flows are not statistically independent [28].

In this paper we address the question whether network decomposition holds in non-asymptotic regimes. Concretely

we derive bounds in terms of statistical envelopes on the output at FIFO nodes with finite number of flows. We consider both Exponentially Bounded Burstiness (EBB) arrivals [27] with and without independent increments. Numerical evaluations of the obtained bounds show that the output envelopes closely match the input envelopes even for relatively small number of flows, thus indicating that network decomposition may be a viable analysis approach in finite regimes.

We show for the many-sources asymptotic regime that the derived output envelopes converge exponentially fast to the arrival envelopes, if a peak rate on the arrivals exists. This type of fast convergence was also shown to hold in the context of the decomposition results from [11]. Otherwise, we show for arrivals described with a compound Poisson process with infinite peak rate that the output envelopes is bounded by the input envelope shifted by a small offset averaging the size of two packets.

The traffic description from this paper with envelope functions is common to the framework of the network calculus [4], [5]. Envelope functions provide bounds limiting the amount of traffic entering or leaving a network node. Envelopes can be either deterministic (strict and never violated) [10], or have some statistical structure allowing for small violation probabilities [21], [27]. Statistical envelopes capture the multiplexing gain characteristic to aggregates of many independent arrivals [3], and are available for many traffic types, e.g., deterministically regulated, Markov-modulated processes, or fractional Brownian motion [23]. For many priority-based scheduling algorithms, such as Static Priority (SP) or Earliest Deadline First (EDF), per-flow statistical envelopes can be constructed at the output of nodes [23].

So far, the study of envelopes in the context of FIFO scheduling has been mainly carried out in the context of the deterministic network calculus. For instance, it was shown that per-flow (deterministic) output envelopes at a FIFO node are tighter than those obtained when the flow receives the lowest priority at a node with SP scheduling [4], [10]. Per-flow output envelopes for FIFO scheduling were further improved in [7]. End-to-end delay bounds for aggregate FIFO scheduling were derived in networks with general topologies [6], [15], [29], and networks with restricted topologies (e.g. acyclic or sink-tree) [12], [22].

In a probabilistic framework, we show that FIFO output envelopes can be much tighter than SP output envelopes, especially at high utilizations, thus extending similar results available in the deterministic network calculus. Moreover, the FIFO output envelope constructions can be iterated for

a probabilistic end-to-end network analysis using a statistical network service curve approach [5], [9], [13]. Consider for instance a through flow traversing H nodes in a feed-forward FIFO network. By iteratively constructing statistical output envelopes from input envelopes, an envelope description of the cross traffic at each of the H nodes can be eventually obtained. From these descriptions, a lower bound on the service received by the through flow at each node can be expressed using service curves which can be further convolved. The resulting network service curve describes the service received by the through flow across the entire network, such that end-to-end backlog and delay bounds for the through flow become available by applying single-node results. As recently shown, the advantage of using network service curves in the end-to-end analysis is that they can yield much tighter end-to-end performance bounds those those derived by adding per-node bounds [9], [13].

The rest of this paper is organized as follows. In Section II we describe the model for a network node and provide a review of some FIFO results in the context of deterministic envelopes. In Section III we derive backlog bounds and construct output envelopes for SP and FIFO scheduling for flows with EBB arrivals and independent increments, and apply these results to M/M/1 and M/D/1 queues. In Section IV we consider general EBB arrivals. In Section V we present numerical results where we consider M/M/1 and M/D/1 queueing scenarios that compare our FIFO bounds with exact results, and then consider Markov-modulated On-Off flows for which we illustrate the convergence of output envelopes to input envelopes. We also test network decomposition in a scenario with two nodes. We conclude the paper in Section VI.

II. SYSTEM MODEL AND REVIEW OF DETERMINISTIC ANALYSIS

We consider a network node with capacity Nc as in Figure 1, where N is a scale parameter representing the number of flows. The node is traversed by N flows of which n are said to be through and $N - n$ are cross flows. The fraction of the through flows is denoted by $\phi = \frac{n}{N}$. We use a continuous time model. All arrivals and departures at the node are modelled by non-decreasing, left-continuous processes. The arrivals of the flows are denoted by $A_i(t)$, $i = 1, \dots, N$, and the arrivals of all through flows are denoted by $A(t)$. For each process $A_i(t)$ we assume that $A_i(0) = 0$ and the departure process $D_i(t)$ satisfies $D_i(t) \leq A_i(t)$ for all $t \geq 0$. We write for convenience $A_i(s, t) = A_i(t) - A_i(s)$. The node serves the arrivals in a fluid-flow manner and stores the backlog in an infinite sized buffer. The backlog process of the through arrivals is denoted by $B(t) = A(t) - D(t)$.

We are interested in the through arrivals $A(t)$ which are described in the paper using statistical envelopes. A statistical envelope is a function $\mathcal{G}(t)$ such that for all $0 \leq s \leq t$ and $\sigma \geq 0$ [9]

$$\Pr\left(A(s, t) > \mathcal{G}(t - s) + \sigma\right) \leq \varepsilon(\sigma), \quad (1)$$

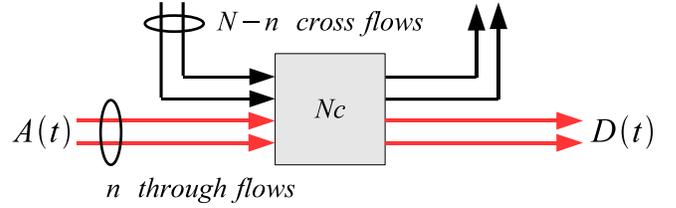


Fig. 1. A network node with normalized capacity Nc traversed by n through and $N - n$ cross flows. The fraction of through flows is $\phi = \frac{n}{N}$.

where $\varepsilon(\sigma)$ is an error functions specifying the probabilities that the arrivals violate the envelope. For the arrivals $A(t)$ we will derive backlog bounds and construct statistical output envelopes for the departures $D(t)$ for both SP and FIFO. While this paper is mainly concerned with FIFO, the results derived for SP serve as worst-case benchmarks, assuming that the through flows receive the lowest priorities.

For simplicity we only consider the case of homogeneous flows. If the flows are heterogeneous then the subsequent analysis can be carried out similarly at the expense of increasing notation.

We review a scenario where the arrivals of flows are bounded by deterministic envelopes which are statistical envelopes with zero error functions. Then, in the remaining part of the paper, we will consider the general case of arrivals bounded by statistical envelopes.

Suppose that each flow $A_i(t)$, $i = 1, \dots, N$, is bounded by a leaky-bucket deterministic envelope $\rho t + \sigma$, where ρ is the arrival rate and σ is the burst (amount of arrivals in a very short interval of time), i.e., [10]

$$A_i(s, t) \leq \rho(t - s) + \sigma,$$

for all $0 \leq s \leq t$.

To compare the output (deterministic) envelopes of the departures $D(t)$ for SP and FIFO scheduling we consider two asymptotic regimes: (a) a small fraction of through flows (i.e. $\phi \rightarrow 0$), and (b) a high utilization $u = \frac{\rho}{c}$ (i.e. $u \rightarrow 1$).

If the through flows are given the lowest priority at a SP scheduler, then the worst-case burst in $D(t)$ occurs when the maximum backlog of $A(t)$ is reached and served instantaneously. This happens at the end of the maximum time period during which the node exclusively serves the $N - n$ cross flows, i.e., after a duration t^* satisfying

$$N(1 - \phi)\rho t^* + N(1 - \phi)\sigma = Nct^*,$$

which yields $t^* = \frac{(1 - \phi)\sigma}{c - (1 - \phi)\rho}$. The envelope for $D(t)$ is then

$$\mathcal{G}_{SP}(t) = n\rho t + \frac{n\sigma}{1 - u(1 - \phi)}.$$

Keeping n fixed and letting $\phi \rightarrow 0$, i.e., $N \rightarrow \infty$, the envelope becomes

$$\mathcal{G}_{SP}(t) = n \left(\rho t + \frac{\sigma}{1 - u} \right). \quad (2)$$

Subject to the condition

$$\frac{n\sigma}{1-u} \leq c$$

this output envelope is tight in the sense that there exists an arrival pattern which yields a burst in the departures of the order predicted by the output envelope.

In turn, if the flows are served in a FIFO manner then the envelope of $D(t)$ can be constructed by applying Theorem 6.2.2 from [4], i.e.,

$$\mathcal{G}_{FIFO}(t) = n \left(\rho t + (1+u)\sigma \right),$$

after letting $\phi \rightarrow 0$. According to Theorem 6.2.2 in [4] this output envelope is also tight.

For the second asymptotic regime, i.e., by letting the utilization $u \rightarrow 1$, the burst predicted by $\mathcal{G}_{SP}(t)$ converges to $N\sigma$ (the same burst predicted by the output envelope constructed with Proposition 1.3.4 and Theorem 1.4.3 from [4]). On the other hand, with FIFO scheduling, the burst predicted by $\mathcal{G}_{FIFO}(t)$ converges to $2n\sigma$. This shows that the difference between the (tight) bursts predicted for SP and FIFO can be substantial, especially at high utilizations.

Finally we emphasize that even with FIFO scheduling, the (deterministic) output envelope gains a non-negligible burst $nu\sigma$ in addition to the burst $n\sigma$ of the arrival envelope. The rest of the paper will provide evidence that in a statistical framework the output envelopes converge to the input envelope, i.e., only negligible bursts are gained at the output of FIFO nodes.

III. OUTPUT BOUNDS: THE CASE OF INDEPENDENT INCREMENTS ARRIVALS

In this section we consider the case when all the flows in Figure 1 have statistically independent increments. Compound Poisson processes or arrivals at GI/GI/1 queues with finite moment generating functions are examples of arrival processes covered in this section. In the next section we will relax the independent increments assumption.

We assume that for each flow $A_i(t)$ the moment generating function satisfies for all $t \geq 0$ and $\theta > 0$

$$\sup_{s \geq 0} E \left[e^{\theta A_i(s, s+t)} \right] = e^{\theta r(\theta)t}, \quad (3)$$

where $r(\theta)$ depends on θ but not on t . Here $r(\theta)$ is a special case of an effective bandwidth [18] which holds for many Markov-modulated processes but excludes self-similar processes where $r(\theta)$ also depends on t . Using Chernoff bound Eq. (3) implies a statistical envelope characterization of the arrivals for all $0 \leq s \leq t$ and $\sigma \geq 0$

$$Pr \left(A_i(s, t) > r(\theta)(t-s) + \sigma \right) \leq e^{-\theta\sigma}. \quad (4)$$

This is a special case of EBB statistical envelopes which have error functions in a more general form, i.e., $\varepsilon(\sigma) = Me^{-\theta\sigma}$, where $M > 0$ is a parameter [27].

The next two theorems construct backlog and output bounds for the through process $A(t)$. We first present the theorem for SP which will be then used to derive the FIFO bounds.

Operations in the usual algebra	Operations in the $(\min, +)$ algebra
$a + b$	$\min \{a, b\}$
$a \cdot b$	$a + b$

TABLE I

In the proof of the theorem we will use the notion of a statistical service curve (referred in [5] as a dynamic F-server) which is a random process $\mathcal{S}(s, t)$ describing probabilistic lower bounds on the service received by $A(t)$, in the sense that for all $t \geq 0$

$$D(t) \geq A * \mathcal{S}(t).$$

The inequality is assumed to hold almost surely, and ‘*’ is the $(\min, +)$ convolution operator, defined as $A * \mathcal{S}(t) = \inf_{0 \leq s \leq t} \{A(s) + \mathcal{S}(s, t)\}$.

Theorem 1: (SP: INDEPENDENT INCREMENTS) Consider the node from Figure 1, where the flows of $A(t)$ receive the lowest priority at a SP scheduler. Each flow $A_i(t)$, $i = 1, \dots, N$, has an effective bandwidth $r(\theta)$ satisfying Eq. (3) for all $\theta > 0$. All flows are independent and have independent increments. Let θ be the maximum positive value satisfying $c \geq r(\theta)$. Then we have the following probabilistic bounds on the through traffic:

- 1) BACKLOG BOUND: For all $t, \sigma \geq 0$

$$Pr \left(B(t) > \sigma \right) \leq e^{-\theta\sigma}. \quad (5)$$

- 2) OUTPUT ENVELOPE: For all $0 \leq s \leq t$ and $\sigma \geq 0$

$$Pr \left(D(s, t) > nr(\theta)(t-s) + \sigma \right) \leq e^{-\theta\sigma}. \quad (6)$$

Since the backlog bound is invariant to N the same bound holds for the entire backlog process at the node, here denoted by $B_N(t)$. The proof idea is to apply a maximal inequality to a suitable constructed supermartingale. More exactly, we adapt the proof for the derivation of backlog bounds in GI/GI/1 queues from [19] to a network calculus setting, in a similar manner as in [8].

PROOF. Let J be the set of cross flows, and let the entire arrival process at the node $A_N(t) = A(t) + \sum_{j \in J} A_j(t)$. Following [13], the process

$$\mathcal{S}(s, t) = Nc(t-s) - \sum_{j \in J} A_j(s, t)$$

is a statistical (leftover) service curve for $A(t)$ such that for all $t \geq 0$

$$D(t) \geq A * \mathcal{S}(t).$$

Then the backlog process $B(t)$ satisfies

$$\begin{aligned} B(t) &= A(t) - D(t) \\ &\leq A(t) - A * \mathcal{S}(t) \\ &\leq \sup_{0 \leq s \leq t} \{A_N(s, t) - Nc(t-s)\}. \end{aligned}$$

Let us now construct the process

$$Z(s) = A_N(t - s, t) - Ncs ,$$

for $0 \leq s \leq t$, and let $\theta > 0$ as in the theorem. Defining a filtration of σ -algebras

$$\mathcal{F}_s = \sigma \{A_N(t - s, t)\}$$

it follows from the independent increments property of the arrivals that for all $s, u \geq 0$

$$\begin{aligned} E \left[e^{\theta Z(s+u)} \mid \mathcal{F}_s \right] &= e^{\theta Z(s)} e^{\theta N u (r(\theta) - c)} \\ &\leq e^{\theta Z(s)} , \end{aligned}$$

thus showing that $e^{\theta Z(s)}$ is a supermartingale. Consequently we can write for all $t, \sigma \geq 0$

$$\begin{aligned} Pr(B(t) > \sigma) &\leq Pr \left(\sup_{0 \leq s \leq t} \{A_N(s, t) - Nc(t - s)\} > \sigma \right) \\ &\leq Pr \left(\sup_{0 \leq s \leq t} e^{\theta Z(s)} > e^{\theta \sigma} \right) \\ &\leq E \left[e^{\theta Z(0)} \right] e^{-\theta \sigma} \\ &= e^{-\theta \sigma} . \end{aligned}$$

In the penultimate equation we applied Doob's inequality to the supermartingale $e^{\theta Z(s)}$. This completes the proof for the backlog. The proof for the output envelope is similar and is omitted here. \square

Theorem 2: (FIFO: INDEPENDENT INCREMENTS) Consider the scenario from Theorem 1, but assume now FIFO scheduling for all flows. For $\theta_1, \theta_2 > 0$ satisfying $c \geq r(\theta_2)$ and $\theta_2 c > \theta_1 \phi r(\theta_1)$, let us denote

$$K(\theta_1, \theta_2) = \frac{c\theta_2}{\theta_2 c - \theta_1 \phi r(\theta_1)} .$$

Then we have the following probabilistic bounds on the through traffic:

1) BACKLOG BOUND: For all $t, \sigma \geq 0$

$$P(B(t) > \sigma) \leq \inf_{\theta_1, \theta_2} K(\theta_1, \theta_2) e^{-\theta_1 \sigma} . \quad (7)$$

2) OUTPUT ENVELOPE: For all $0 \leq s \leq t$, $\sigma \geq 0$, and $\theta < \theta_1$

$$\begin{aligned} P(D(s, t) > nr(\theta)(t - s) + \sigma) \\ \leq \left(1 + K(\theta_1, \theta_2) \frac{\theta}{\theta_1 - \theta} \right) e^{-\theta \sigma} . \end{aligned} \quad (8)$$

PROOF. Let $t, \sigma \geq 0$, and let us denote by $B_N(t)$ the entire backlog process at the node. Critical to the proof is the following property of FIFO:

$$B(t) > \sigma \implies \left\{ \begin{array}{l} \exists s \leq t : A(s, t) > \sigma \\ \text{and } B_N(s) \geq Nc(t - s) \end{array} \right\} . \quad (9)$$

Here s denotes the arrival time of the earliest data unit, call it x , in $B(t)$. To have that $B(t) > \sigma$, it is necessary to

have at least σ arrivals in $A(s, t)$, i.e., the first clause in the righthand side of Eq. (9). Also, to have that x belongs to $B(t)$, it is necessary to have the entire backlog $B_N(s)$ exceed the maximal amount of data that can be processed in the time interval $(s, t]$, i.e., the second clause in the righthand side of Eq. (9).

Define a time discretization parameter $\tau_0 > 0$. For $0 \leq s \leq t$, denote by $j = \lfloor \frac{t-s}{\tau_0} \rfloor$ the integer part of $\frac{t-s}{\tau_0}$. Fix $\theta_1, \theta_2 > 0$ satisfying the condition from the theorem. Then we can write for the backlog process $B(t)$

$$\begin{aligned} Pr(B(t) > \sigma) &\leq Pr \left(\bigcup_{s \leq t} \{A(s, t) > \sigma, B_N(s) \geq Nc(t - s)\} \right) \\ &\leq Pr \left(\bigcup_{j \geq 0} \left\{ \begin{array}{l} A(t - (j+1)\tau_0, t) > \sigma, \\ B_N(t - j\tau_0) \geq Ncj\tau_0 \end{array} \right\} \right) \\ &\leq \sum_{j \geq 1} e^{\theta_1(\phi N r(\theta_1)j\tau_0 - \sigma)} e^{-\theta_2 Nc(j-1)\tau_0} \\ &\leq \frac{e^{\theta_2 Nc\tau_0}}{N(\theta_2 c - \theta_1 \phi r(\theta_1))\tau_0} e^{-\theta_1 \sigma} . \end{aligned} \quad (10)$$

In the second line we applied Eq. (9). In the third line we applied the monotonicity of $A(t)$ and that the total backlog process is stochastically increasing. In the fourth line we used the backlog bound from Eq. (5) in Theorem 1, which corresponds to the total backlog process as well. Finally we bounded the sum using the inequality $\sum_{j \geq 1} e^{-aj} \leq \int_0^\infty e^{-ax} dx$, for all $a > 0$. The proof on the backlog is complete after optimizing $\tau_0 = \frac{1}{\theta_2 Nc}$, and minimizing over θ_1 and θ_2 .

For the output envelope, let us fix $s \leq t$ and $\theta < \theta_1$. We can write for the output process

$$\begin{aligned} Pr(D(s, t) > nr(\theta)(t - s) + \sigma) \\ \leq Pr(A(s, t) + B(s) > nr(\theta)(t - s) + \sigma) \\ \leq E \left[e^{\theta B(s)} \right] e^{-\theta \sigma} . \end{aligned} \quad (11)$$

In the last line we used the independent increments property and the Chernoff bound. With the backlog bound from Eq. (7), we can bound the moment generating function of $B(s)$ by

$$\begin{aligned} E \left[e^{\theta B(s)} \right] &= 1 + \int_1^\infty Pr \left(B(s) > \frac{1}{\theta} \log x \right) dx \\ &\leq 1 + K(\theta_1, \theta_2) \frac{\theta}{\theta - \theta_1} \int_1^\infty dx \frac{\theta - \theta_1}{\theta} \\ &= 1 + K(\theta_1, \theta_2) \frac{\theta}{\theta_1 - \theta} , \end{aligned}$$

thus completing the proof. \square

The next two subsections apply Theorem 2 to M/M/1 and M/D/1 queueing systems. The goal is to analyze the tightness of the FIFO output bounds from a stochastic ordering point of view and compare the bounds with exact results. For two random variables X, Y we say that X is stochastically smaller than Y , and denote this relationship by

$$X \leq_{st} Y ,$$

if $Pr(X > \sigma) \leq Pr(Y > \sigma)$ for all σ [24].

A. Example 1: The M/M/1 queue

Consider that packets of each flow in Figure 1 arrive as a Poisson process with rate λ_0 and with exponentially distributed sizes with mean $\frac{1}{\mu}$. The total arrival rate is denoted by $\lambda = N\lambda_0$. Assume that the N arrival processes are statistically independent, and that the queue is stable, i.e., the utilization factor $u = \frac{\lambda_0}{\mu c} < 1$. For all $\theta < \mu$, the effective bandwidth of a flow is

$$r(\theta) = \frac{\lambda_0}{\mu - \theta}.$$

First we consider the asymptotic regime $\phi \rightarrow 0$. For fixed θ_1 satisfying the conditions from Theorem 2 we have that $K(\theta_1, \theta_2) \rightarrow e$. Therefore, letting $\theta_1 \rightarrow \mu$ we get from Eq. (7) that the FIFO backlog $B(t)$ satisfies

$$B(t) \leq_{st} Y + \frac{1}{\mu},$$

where Y is exponentially distributed with mean $1/\mu$. From the convexity of the exponential function we have $E[e^{\theta B(t)}] \leq E[e^{\theta(Y + \frac{1}{\mu})}]$. Using this in Eq. (11) yields for all s, t

$$D(s, t) \leq_{st} A(s, t) + Y + \frac{1}{\mu}. \quad (12)$$

This indicates that as $\phi \rightarrow 0$, the output $D(s, t)$ gains a burst of roughly two packet sizes relative to the arrivals $A(s, t)$.

In turn, since all the flows are independent and have the same service rate, the queue is quasi-reversible such that the departure process $D(t)$ is Poisson [17]. Note however that this exact result holds in steady state, whereas the stochastic bound from Eq. (12) holds for all values of s, t .

Now we consider a non-asymptotic regime where ϕ can take any value. To derive the FIFO output envelope for any ϕ we first optimize $\theta_2 = \mu(1 - u)$ in Theorem 2. This yields the output envelope for all $0 \leq s \leq t$ and $\sigma \geq 0$

$$Pr\left(D(s, t) > \frac{\phi u \mu N c}{\mu - \theta}(t - s) + \sigma\right) \leq \left(1 + e^{\frac{1-u}{1-u-\phi u} \frac{\theta}{\mu - \theta_1} \frac{\theta}{\theta_1 - \theta}}\right) e^{-\theta \sigma}, \quad (13)$$

where $0 < \theta < \theta_1 < \mu \frac{1-u}{1-u+\phi u}$.

In the steady-state, the quasi-reversibility of the queue yields the output envelope for large values of s, t and all $\sigma \geq 0$

$$Pr\left(D(s, t) > \frac{\phi u \mu N c}{\mu - \theta}(t - s) + \sigma\right) \leq e^{-\theta \sigma}. \quad (14)$$

This bound is clearly tighter than the bound from Eq. (13). Numerical comparisons between them will be shown in Section V.

B. Example 2: The M/D/1 queue

We consider a similar system as before with the exception that all packet sizes are equal to $\frac{1}{\mu}$. For all positive θ the effective bandwidth of a flow is now given by

$$r(\theta) = \frac{\lambda_0}{\mu} \left(e^{\frac{\theta}{\mu}} - 1\right).$$

The arrival envelope of $A(t)$ is then given by

$$Pr\left(D(s, t) > \frac{\phi u \mu N c}{\theta} \left(e^{\frac{\theta}{\mu}} - 1\right) (t - s) + \sigma\right) \leq e^{-\theta \sigma}. \quad (15)$$

The output envelope of $A(t)$ is given as in Eq. (15), but with the violation probability from Eq. (8). The conditions on θ_1 and θ_2 are in this case

$$\frac{\theta_2}{\mu} > \max\left\{u \left(e^{\frac{\theta_2}{\mu}} - 1\right), \phi u \left(e^{\frac{\theta_1}{\mu}} - 1\right)\right\}.$$

IV. OUTPUT BOUNDS: GENERAL CASE

In this section we extend the bounds for SP and FIFO scheduling to arrival processes without independent increments. We maintain the requirement of finite and time-invariant effective bandwidths for the flows. The class of arrival processes covered in this section extends to some Markov-modulated processes (e.g. Markov-modulated On-Off) and regulated processes.

As in the previous section, we first derive SP results and then use these results for FIFO scheduling.

Theorem 3: (SP) Consider the scenario from Theorem 1, without the requirement of independent increments for the flows. Let us denote

$$K(\theta) = \frac{ce}{c - r(\theta)}.$$

whenever $c > r(\theta)$. Then we have the following probabilistic bounds on the through traffic:

1) BACKLOG BOUND: For all $t, \sigma \geq 0$

$$Pr(B(t) > \sigma) \leq \inf_{\theta} K(\theta) e^{-\theta \sigma}. \quad (16)$$

2) OUTPUT ENVELOPE: For all $0 \leq s \leq t$ and $\sigma \geq 0$

$$Pr(D(s, t) > nr(\theta)(t - s) + \sigma) \leq K(\theta) e^{-\theta \sigma}. \quad (17)$$

The theorem extends results from [13] to a continuous time setting. Since the backlog bound from Eq. (16) is invariant to N , the same bound holds for the entire backlog process $B_N(t)$ at the node. Both bounds from the theorem are weaker than the corresponding bounds from Theorem 1, especially at high utilizations. The reason is that $K(\theta) > 1$ in Theorem 3. Note also that $K(\theta)$ rapidly increases with the utilization.

PROOF. We recall that in the proof of Theorem 1 we proved for the backlog process $B(t)$

$$B(t) \leq \sup_{0 \leq s \leq t} \{A_N(s, t) - Nc(t - s)\},$$

where $A_N(t)$ is the total arrival process at the switch. For $0 \leq s \leq t$ and a discretization parameter τ_0 , let us denote

$j = \lfloor \frac{t-s}{\tau_0} \rfloor$ the integer part of $\frac{t-s}{\tau_0}$. Using the monotonicity of the arrival processes, we can write for all $\sigma \geq 0$

$$\begin{aligned} Pr(B(t) > \sigma) &\leq Pr\left(\sup_{0 \leq s \leq t} \{A_N(s, t) - Nc(t-s)\} > \sigma\right) \\ &\leq Pr\left(\sup_{j \geq 1} \{A_N(t - j\tau_0, t) - Nc(j-1)\tau_0\} > \sigma\right) \\ &\leq \sum_{j \geq 1} e^{\theta Nr(\theta)j\tau_0} e^{-\theta Ncj\tau_0} e^{\theta Nc\tau_0} e^{-\theta\sigma} \\ &\leq \frac{e^{\theta Nc\tau_0}}{\theta N(c-r(\theta))\tau_0} e^{-\theta\sigma}. \end{aligned}$$

In the last line we bounded the sum using the integral inequality from the proof of Theorem 1. The proof for the backlog is complete after optimizing $\tau_0 = \frac{1}{\theta Nc}$ and minimizing over θ . The proof for the output envelope is similar and is omitted here. \square

For FIFO scheduling we need to make two statistical independence assumptions between the through process $A(t)$ and past backlogs:

(A1) $A(s, t)$ is statistically independent of the entire backlog process $B_N(s)$ at time s , for all $s \leq t$.

(A2) $A(s, t)$ is statistically independent of A 's backlog process $B(s)$ at time s , for all $s \leq t$.

These are technical assumptions which will be used in the proof of the theorem. Note that both assumptions hold when the n through flows have statistically independent increments (the cross flows do not need to have independent increments as well). For large number of flows N and a small number n of through flows, both assumptions are justified since the two backlogs are mainly determined by a large fraction of cross flows which are independent of $A(t)$. On the other hand, a large fraction of through flows at the node can create correlations between the arrivals $A(s, t)$ and the two backlogs at time s . For example, large values of $B(s)$ can imply a large number of arrivals $A(s, s+u)$ for u relatively small.

We point out that other independence assumptions have been made in the literature to avoid technical complications due to statistical dependencies. For instance, randomization effects are invoked in M/M/1 (dense) queueing networks to justify the assumption that each packet has its size independently regenerated at each traversed node [20]. With this assumption, inter-departure times and packet sizes are not correlated in M/M/1 *packet* queueing networks, and product-form formulas become available.

Theorem 4: (FIFO) Consider the scenario from Theorem 2, without the requirement of independent increments for the flows. For all $\theta_1, \theta_2 > 0$ satisfying $c > r(\theta_2)$ and $\theta_2c > \theta_1\phi r(\theta_1)$, let us denote

$$K(\theta_1, \theta_2) = \frac{(ce)^2\theta_2}{(c-r(\theta_2))(\theta_2c - \theta_1\phi r(\theta_1))}.$$

Then we have the following probabilistic bounds on the through traffic:

1) BACKLOG BOUND: For all $t, \sigma \geq 0$

$$P(B(t) > \sigma) \leq \inf_{\theta_1, \theta_2} K(\theta_1, \theta_2) e^{-\theta_1\sigma}. \quad (18)$$

2) OUTPUT ENVELOPE: For all $0 \leq s \leq t, \sigma \geq 0$, and $\theta < \theta_1$

$$\begin{aligned} P(D(s, t) > nr(\theta)(t-s) + \sigma) \\ \leq \left(1 + K(\theta_1, \theta_2) \frac{\theta}{\theta_1 - \theta}\right) e^{-\theta\sigma}. \end{aligned} \quad (19)$$

The two bounds in the theorem are weaker than the bounds from Theorem 2. The reason is that without accounting for independent increments, the term $K(\theta_1, \theta_2)$ in the theorem gains a factor $\frac{ce}{c-r(\theta_2)} > 1$ relative to $K(\theta_1, \theta_2)$ from Theorem 2.

PROOF. The proof is analogous to the proof of Theorem 2. The difference is that in the fourth line of Eq. (10) we now use the bound on the entire backlog process $B_N(t)$ from Theorem 3 instead of the bound from Theorem 1. \square

For the rest of the section we discuss the convergence of the FIFO backlog and output envelopes from Theorem 4 by fixing the number of through flows n and letting $N \rightarrow \infty$.

Assuming that each flow $A_i(t)$ is bounded by some finite peak rate P , i.e.,

$$\sup_{s \leq t} \left\{ \frac{\sup \{\sigma : Pr(A_i(s, t) > \sigma) > 0\}}{t-s} \right\} \leq P < \infty,$$

we will next show that the backlog process $B(t)$ from Theorem 4 satisfies for all $\sigma \geq 0$

$$P(B(t) > \sigma) \leq Ke^{-N\theta\sigma}, \quad (20)$$

for some constants $K, \theta > 0$. This indicates that $B(t)$ converges to zero exponentially fast in the number of flows N , and, consequently, that the output $D(t)$ converges to the input $A(t)$.

Since $r(\theta) \rightarrow 0$ as $\theta \rightarrow 0$ [18] we can choose θ_2 sufficiently small such that $c > r(\theta_2)$ in Theorem 4. We next select $\theta_1 = \frac{N}{2n} \frac{\theta_2c}{P}$ such that $\theta_2c - \theta_1\phi r(\theta_1) > \frac{\theta_2c}{P}$ since $r(\theta_1) < P$. The conditions of Theorem 4 are then met. It then follows that Eq. (18) implies Eq. (20) with the constants

$$\theta = \theta_1, \quad K = \frac{2ec^2}{c-r(\theta_2)}.$$

For the same values of θ_1 and θ_2 , the prefactor of the exponential in the output envelope bound from Eq. (19) converges to 1 as $N \rightarrow \infty$. Therefore the output envelope from Eq. (19) converges exponentially fast to the input envelope of $A(t)$ obtained from Eq. (4) by properly scaling the rate $r(\theta)$ with the number n of through flows

$$Pr(A(s, t) > nr(\theta)(t-s) + \sigma) \leq e^{-\theta\sigma}.$$

We point out that the critical condition used to prove the convergence of the backlog and output envelope is that the

flows have a finite peak rate. This condition appears also in [11] in order to prove decomposition results. Arrivals modelled with Markov-modulated On-Off process (see Subsection V-B) satisfy this condition, whereas the compound Poisson process from Subsection III-A does not.

V. NUMERICAL EXAMPLES

This section presents numerical examples that illustrate the accuracy of the output envelopes derived in this paper. Subsection V-A presents the cases of M/M/1 and M/D/1 queues which permit a comparison of FIFO output envelopes with envelopes obtained from exact results. Subsection V-B considers a single-node scenario with Markov-Modulated On-Off arrival processes where we show the convergence of the output envelopes to the arrival envelopes. Finally, Subsection V-C illustrates the viability of network decomposition in a two-node network scenario with Markov-modulated On-Off arrivals.

All numerical examples are shown by measuring time in milliseconds, nodes capacity $Nc = 100Mbps$, and letting the violation probability be 10^{-6} .

A. M/M/1 and M/D/1 Queues

Assume that the node from Figure 1 is modelled either with an M/M/1 or M/D/1 queue as in Subsections III-A and III-B. For the numerical setting we let the mean packet size and utilization factor $\frac{1}{\mu} = 400 Bytes$ and $u = 75\%$, respectively.

Figure 2(a) shows statistical arrival and output envelopes for two fractions $\phi = 0.5$ or $\phi = 0.95$ of flows at the M/M/1 queue. The output envelope is given in Eq. (13), and the arrival envelope is given in Eq. (14); the latter is also an output envelope by using quasi-reversibility. The two envelopes closely match at all time scales, even if the envelope from Eq. (14) is constructed from an exact result which holds in the steady-state. For M/D/1, Figure 2(b) shows an even closer match between the two envelopes which may indicate that M/D/1 queues yield less bursty traffic at the output than M/M/1 queues.

B. Markov-Modulated On-Off Processes

We assume that all flows at the node are statistically independent Markov-modulated On-Off processes. Such a process is illustrated in Figure 4, and can be defined as follows. There is a two-state homogenous and continuous-time Markov chain $X(t)$ with the generator matrix

$$Q = \begin{pmatrix} -\mu & \mu \\ \lambda & -\lambda \end{pmatrix}.$$

The states of $X(t)$ are denoted ‘On’ and ‘Off’, and μ and λ represent the transition rates from the ‘On’ state to the ‘Off’ state, and vice-versa, respectively.

We say that $A(t)$ is a Markov-modulated On-Off process with modulated process $X(t)$ if the arrival rate of the process is either P or zero, depending whether $X(t)$ is in the ‘On’ and ‘Off’ states, respectively. In the steady-state, the average

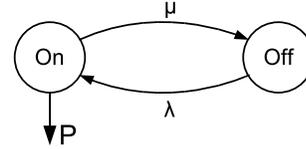


Fig. 4. A Markov-modulated On-Off traffic model

time in the ‘On’ state is $\frac{1}{\mu}$ and the average time in the ‘Off’ state is $\frac{1}{\lambda}$. The value $T = \frac{1}{\mu} + \frac{1}{\lambda}$ is the average time for a Markov process to change state twice.

The effective bandwidth of a single Markov-modulated On-Off process satisfies for all $\theta > 0$ [18]

$$r(\theta) \leq \frac{1}{2\theta} \left(P\theta - \mu - \lambda + \sqrt{(P\theta - \mu + \lambda)^2 + 4\lambda\mu} \right). \quad (21)$$

We consider the following numerical settings. The peak rate of a flow is $P = 1.5 Mbps$ and the average rate is $R = 0.15 Mbps$, such that the maximum number of flows is $N = 667$. We let $T = 10 ms$ such that μ and λ can be computed from P , r , and T .

In Figure 3 we show two sets of three statistical envelopes for the aggregate $A(t)$, depending on the number of through flows n . The first envelope is the arrival envelope from Eq. (4), where $r(\theta)$ is scaled by n . The second is the FIFO output envelope of $A(t)$ from Eq. (19). The third is the output envelope of $A(t)$ from Eq. (17), when its flows receive the lowest priorities at a SP scheduler. The parameters $\theta, \theta_1, \theta_2$ in the equations are optimized numerically, according to the conditions from Theorems 3 and 4.

In Figure 3(a) we consider a moderate utilization factor $u = 0.75$, such that there are $N = 500$ flows at the node. In Figure 3(b) we consider a high utilization factor $u = 0.95$, implying that $N = 633$. We consider $n = 1$ (i.e. $\phi = 0.2\%$ in (a) and $\phi = 0.1\%$ in (b)) and $n = 10$ (i.e. $\phi = 2\%$ in (a) and $\phi = 1\%$ in (b)).

Both figures illustrate that the FIFO output envelopes converge to the arrival envelopes as the number n of through flows decreases, or as the ratio $\phi = \frac{n}{N}$ decreases, agreeing with the convergence result from the end of Section IV. Moreover, we note that the SP output envelopes can be very pessimistic, especially at high utilizations.

C. The case for decomposition

We test the viability of network decomposition in the network scenario with two nodes from Figure 5. The network is traversed by n through flows and each node is also traversed by $N - c$ cross flows. We assume that all flows are independent Markov-modulated On-Off processes with the same numerical values as in Subsection V-B.

We compare the total backlog at the second node with the total backlog at the second node obtained in a simplified scenario where the first node is removed. The simplified bound is thus obtained by assuming that the arrivals of $A^2(t)$ are

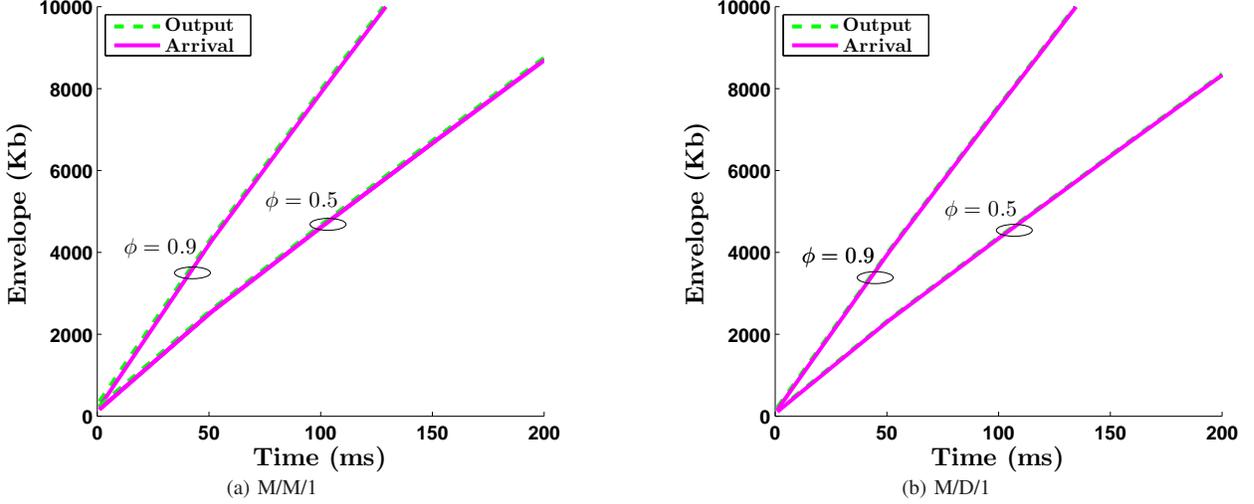


Fig. 2. Statistical envelopes for a fraction of flows ($\phi = 0.5$ and $\phi = 0.9$) at a queue as a function of time ($Nc = 100$ Mbps, utilization factor $u = 0.75$, average packet size $\frac{1}{\mu} = 400$ Bytes, violation probability 10^{-6})

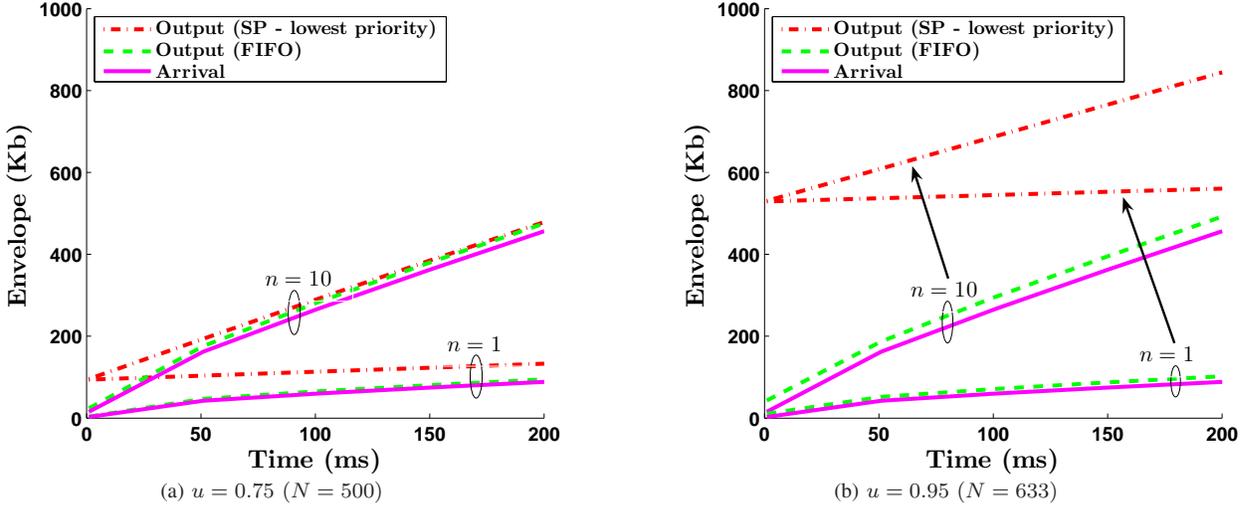


Fig. 3. Statistical arrival and output envelopes for $n = 1, 10$ Markov Modulated On-Off flows as a function of time ($Nc = 100$ Mbps, per-flow average rate $R = 0.15$ Mbps and peak rate $P = 1.5$ Mbps, $T = 10$ ms, violation probability 10^{-6})

exactly those of $A^1(t)$, and the expression for this bound is given in Eq. (16).

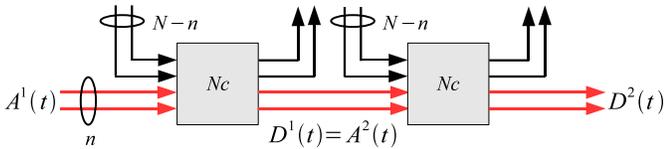


Fig. 5. A network with two nodes and cross traffic

A bound on the entire backlog process $B_N^2(t)$ at the second node can be derived by first constructing a bound on the moment generating function of $A^2(t)$ in a similar manner as

the output envelope construction from Theorem 4. We obtain for $\theta > 0$

$$\sup_{s \geq 0} E \left[e^{\theta A^2(s, s+t)} \right] \leq \left(1 + K(\theta_1, \theta_2) \frac{\theta}{\theta_1 - \theta} \right) e^{\theta nr(\theta)t},$$

where $K(\theta_1, \theta_2)$ and the conditions on θ_1, θ_2 are given in Theorem 4. Then we can slightly modify the proof of Theorem 3 to account for the prefactor of the exponential above, and obtain for all $\sigma \geq 0$

$$Pr \left(B_N^2(t) > \sigma \right) \leq \inf_{\theta_1, \theta_2, \theta} K(\theta) \left(1 + K(\theta_1, \theta_2) \frac{\theta}{\theta_1 - \theta} \right) e^{-\theta \sigma}, \quad (22)$$

where $K(\theta)$ is given in Theorem 3. The constraints on $\theta_1, \theta_2, \theta$ are as in Theorems 3 and 4.

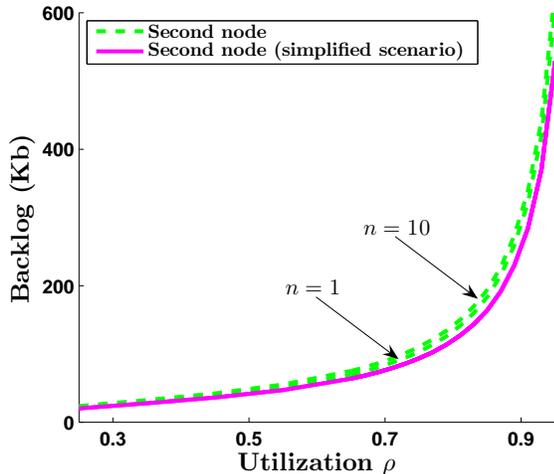


Fig. 6. Backlog bounds at the second node as a function of utilization ($Nc = 100$ Mbps, through flows $n = 1$ and $n = 10$, per-flow average rate $R = 0.15$ Mbps and peak rate $P = 1.5$ Mbps, $T = 10$ ms, violation probability 10^{-6})

In Figure 6 we plot the backlog bounds from Eq. (22) and Eq. (16), the latter corresponding to the simplified scenario, as a function of the utilization factor u . The bound from Eq. (22) is shown for two different values of the number of through flows (i.e. $n = 1$ and $n = 10$). Note that the bound in the simplified scenario is invariant to n . The plots illustrate that by decreasing n , or equivalently ϕ , the backlog bounds at the second node from Eq. (22) tend to closely match the backlog bounds derived in the simplified scenario, thus justifying network decomposition.

VI. CONCLUSIONS

In this paper we have investigated the viability of network decomposition in networks with FIFO scheduling and a finite number of flows. We described traffic with statistical EBB envelopes and provided envelope construction at the output of FIFO nodes. We showed that the derived output envelopes converge exponentially fast to the input envelopes. Our numerical results illustrate that the convergence becomes visible for relatively small number of flows. The statistical output envelopes derived for FIFO were shown to be much tighter than those obtained for SP, thus extending a similar result from the deterministic network calculus. The explicit expressions for statistical output envelopes can be applied to a probabilistic end-to-end network analysis of feed-forward networks with FIFO scheduling.

REFERENCES

- [1] F. Baskett, K. M. Chandy, R. R. Muntz, and F. G. Palacios. Open, closed and mixed networks of queues with different classes of customers. *Journal of the ACM*, 22(2):248–260, April 1975.
- [2] T. Bonald, A. Proutière, and J. W. Roberts. Statistical guarantees for streaming flows using expedited forwarding. In *INFOCOM*, pages 1104–1112, 2001.

- [3] R. Boorstyn, A. Burchard, J. Liebeherr, and C. Oottamakorn. Effective envelopes: Statistical bounds on multiplexed traffic in packet networks. In *Proceedings of IEEE Infocom 2000*, pages 1223–1232, Tel Aviv, Mar. 2000.
- [4] J. Y. Le Boudec and P. Thiran. *Network Calculus*. Springer Verlag, Lecture Notes in Computer Science, LNCS 2050, 2001.
- [5] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [6] A. Charny and J.-Y. Le Boudec. Delay bounds in a network with aggregate scheduling. In *Proceedings of the First International Workshop on Quality of Future Internet Services*, pages 1–13, 2000.
- [7] V. Chlovi, J. Echagüe, and J.-Y. Le Boudec. Worst case burstiness increase due to FIFO multiplexing. *Performance Evaluation*, 49(1-4):491–506, 2002.
- [8] F. Ciucu. Network calculus delay bounds in queueing networks with exact solutions. In *20th International Teletraffic Congress (ITC)*, pages 495–506, June 2007.
- [9] F. Ciucu, A. Burchard, and J. Liebeherr. Scaling properties of statistical end-to-end bounds in the network calculus. *IEEE Transactions on Information Theory*, 52(6):2300–2312, June 2006.
- [10] R. Cruz. A calculus for network delay, parts I and II. *IEEE Transactions on Information Theory*, 37(1):114–141, Jan. 1991.
- [11] D. Y. Eun and N. B. Shroff. Network decomposition: theory and practice. *IEEE/ACM Transactions on Networking*, 13(3):526–539, 2005.
- [12] M. Fidler. Extending the network calculus pay bursts only once principle to aggregate scheduling. In *International Workshop on Quality of Service in Multiservice IP Networks (QoS-IP)*, pages 19–34, 2003.
- [13] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *IEEE 14th International Workshop on Quality of Service (IWQoS)*, pages 261–270, June 2006.
- [14] M. Grossglauser and S. Keshav. On cbr service. In *INFOCOM*, pages 129–137, 1996.
- [15] Y. Jiang. Delay bounds for a network of guaranteed rate servers with FIFO aggregation. *Computer Networks*, 40(6):683–694, 2002.
- [16] F. P. Kelly. Networks of queues with customers of different types. *Journal of Applied Probability*, 3(12):542–554, Sept. 1975.
- [17] F. P. Kelly. Networks of quasi-reversible queues. *Applied Probability - Computer Science: The Interface*, 1:3–26, 1982.
- [18] F. P. 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.
- [19] J. F. C. Kingman. A martingale inequality in the theory of queues. *Cambridge Philosophical Society*, 59:359–361, 1964.
- [20] L. Kleinrock. *Queueing Systems*, volume 1. John Wiley and Sons, 1975.
- [21] J. Kurose. On computing per-session performance bounds in high-speed multi-hop computer networks. In *Proceedings of ACM Sigmetrics*, pages 128–139, June 1992.
- [22] L. Lenzi, L. Martorini, E. Mingozzi, and G. Stea. Tight end-to-end per-flow delay bounds in FIFO multiplexing sink-tree networks. *Performance Evaluation*, 63(9):956–987, 2006.
- [23] C. Li, A. Burchard, and J. Liebeherr. A network calculus with effective bandwidth. *IEEE/ACM Transactions on Networking*, 15(6):1442–1453, Feb. 2007.
- [24] D. Stoyan. *Comparison Methods for Queues and Other Stochastic Models*. John Wiley & Sons Inc., 1983.
- [25] G. de Veciana, C. Courcoubetis, and J. C. Walrand. Decoupling bandwidths for networks: A decomposition approach to resource management. In *Proceedings of IEEE Infocom*, pages 466–473, 1994.
- [26] D. Wischik. The output of a switch, or, effective bandwidths for networks. *Queueing Systems*, 32:383–396, 1999.
- [27] O. Yaron and M. Sidi. Performance and stability of communication networks via robust exponential bounds. *IEEE/ACM Transactions on Networking*, 1(3):372–385, June 1993.
- [28] Y. Ying, R. Mazumdar, C. Rosenberg, and F. Guillemin. The burstiness behavior of regulated flows in networks. In *IEEE Networking*, pages 918–929, 2005.
- [29] Z.-L. Zhang, Z. Duan, and Y. T. Hou. Fundamental trade-offs in aggregate packet scheduling. In *International Conference on Network Protocols (ICNP)*, pages 129–137, 2001.