[SZF08]

Jens Schmitt, Frank Zdarsky, Markus Fidler; Delay Bounds under Arbitrary Multiplexing: When Network Calculus Leaves You in the Lurch.In: Proceedings of IEEE INFOCOM, p. 1669-1677, IEEE Press, April 2008.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

Delay Bounds under Arbitrary Multiplexing: When Network Calculus Leaves You in the Lurch...

Jens B. Schmitt Frank A. Zdarsky Distributed Computer Systems Lab (DISCO) University of Kaiserslautern, Germany Markus Fidler Multimedia Communications Lab (KOM) Darmstadt University of Technology, Germany

Abstract-Network calculus has proven as a valuable and versatile methodology for worst-case analysis of communication networks. One issue in which it is still lacking is the treatment of aggregate multiplexing, in particular if the FIFO property cannot be assumed when flows are merged. In this paper, we address the problem of bounding the delay of individual traffic flows in feed-forward networks under arbitrary multiplexing. Somewhat surprisingly, we find that direct application of network calculus results in loose bounds even in seemingly simple scenarios. The reasons for this "failure" of network calculus are discussed in detail and a method to arrive at tight delay bounds for arbitrary (aggregate) multiplexing is presented. This method is based on the solution of an optimization problem. For the special case of sink-tree networks this optimization problem is solved explicitly, thus arriving at a closed-form expression for the delay bound. Numerical experiments illustrate that in sink-tree networks the improvement over bounds based on direct application of network calculus can be considerable.

Index Terms—Performance Evaluation, Performance Bounds, Network Calculus, Feed-forward Networks, Arbitrary Multiplexing.

I. INTRODUCTION

Network calculus is a min-plus system theory for deterministic queuing systems which builds on the calculus for network delay in [1], [2]. The important concept of *service curve* was introduced in [3]–[7]. The service curve based approach facilitates the efficient analysis of tandem queues where a linear scaling of performance bounds in the number of traversed queues is achieved as elaborated in [8] and also referred to as pay bursts only once phenomenon in [9]. A detailed treatment of min-plus algebra and of network calculus can be found in [10] and [9], [11], respectively.

Network calculus has found numerous applications, most prominently in the Internet's Quality of Service (QoS) proposals IntServ and DiffServ, but also in other scenarios as, for example, wireless sensor networks [12], [13], switched Ethernets [14], Systems-on-Chip (SoC) [15], or even to speedup simulations [16]. Hence, besides queueing theory it has established as a valuable methodology.

However, as a relatively young theory, compared to e.g. traditional queueing theory, there is also a number of challenges network calculus still has to master. To name a few: recently there has been much interest and progress with respect to stochastic extensions (see [8], [17], [18] for recent advances); tool support for network calculus has been addressed by [19], [20] and brings about new interesting perspectives. A very tough challenge is also found in the treatment of nontandem topologies with aggregate multiplexing of multiple flows. While this has been addressed from the beginning [2], there are still many open issues. For aggregate multiplexing in general network topologies there is a very fundamental issue about the circumstances under which a finite delay bound exists at all [21], [22]. In [23] a sufficient condition for stability in general network topologies and an explicit delay bound are given. Extensions of this approach are provided by [24] and [25]. Yet, for larger networks this puts a heavy constraint on the utilization of the network since the maximum allowable utilization is inversely proportional to the network diameter. The problems in the analysis of general topologies arise due to cyclic dependencies between flows and the consequent difficulties in bounding their network-internal burstiness. Implementing special (fair) aggregation schemes wherever flows are merging can alleviate the utilization problem but requires special actions on the data path [26], [27]. A class of topologies which avoids those problems are feed-forward networks, which are known to be stable for all utilizations ≤ 1 [2]. In this paper, we focus on this class of networks. While many networks are obviously not feed-forward, many important instances like switched networks, wireless sensor networks, or MPLS networks with multipoint-to-point label switched paths are, or can be made feed-forward by using, e.g., the turnprohibition algorithm [28].

In feed-forward networks, there has been some work on aggregate multiplexing recently: [29] treats the case of feedforward networks under FIFO multiplexing for token-bucket constrained flows and rate-latency servers. The left-over service curve for a flow of interest is derived. It is again of the rate-latency type with minimally possible latency. [30] shows that this does not result in a tight delay bound, and derives tight delay bounds under knowledge about the arrival curve of the flow of interest for the special case of sink-trees and, again, under token bucket constrained flows and rate-latency servers. Another work [31] also investigates sink-tree networks, but now under dual token-bucket constrained flows and constant rate servers, for which delay bounds are derived by summing per-node bounds, expectedly not arriving at tight bounds but reported as being at least close under practical conditions.

Besides being very specific with respect to traffic and server models, all of the above work assumes FIFO aggregate multiplexing. However in practice, as argued in [32], many devices cannot be accurately described by FIFO because packets arriving at the output queue from different input ports may experience different delays when traversing a node. This is due to the fact that many networking devices like routers are implemented using input-output buffered crossbars and/or multistage interconnections between input and output ports. Hence, packet reordering on the aggregate level is a frequent event (not so on the flow level) and should not be neglected in modelling. Therefore, in this work we drop the FIFO multiplexing assumption and make essentially no assumptions on the way aggregates are multiplexed at servers, i.e. we assume arbitrary multiplexing a.k.a general or blind multiplexing [1], [9]. On the level of a single flow, however, we still assume FIFO. This assumption is sometimes called FIFO-per-microflow [33] or locally FCFS multiplexing [1].

There is actually little work on arbitrary multiplexing: Some results are reported in [9] (see Section II), and there is some work on the burstiness increase due to arbitrary multiplexing at a single node [34]. In previous work [19] related to network calculus tool support, we have proposed and implemented a number of network calculus analysis methods for arbitrary multiplexing feed-forward networks. Some of these are presented in Section III, but as will be demonstrated, they were not the "final word on this topic". While not addressing aggregate multiplexing, we also mention [33] here, because it demonstrates that releasing FIFO assumptions can lead to interesting and somewhat unexpected phenomena, similar to what we will be dealing with later on.

The goal of our work is to calculate tight delay bounds in feed-forward networks of arbitrary multiplexers. With respect to traffic and server models we address a more general case than previous work on FIFO multiplexing, in particular we assume piecewise linear concave arrival curves and convex service curves, which encompass the majority of practical traffic and server models.

In essence, the main contributions of this paper are

- exposition of a fundamental problem of network calculus to achieve tight delay bounds in any non-FIFO aggregate multiplexing (→Section III);
- a novel method to achieve tight delay bounds in feedforward networks under arbitrary multiplexing (→Section IV);
- closed-form expressions for tight delay bounds in sinktree networks (→Section V).

II. NETWORK CALCULUS BASICS

As network calculus is built around the notion of cumulative functions for input and output flows of data, the set of real-valued, non-negative, and wide-sense increasing functions passing through the origin plays a major role:

$$\mathcal{F} = \{ f : \mathbb{R}^+ \to \mathbb{R}^+, \forall t \ge s : f(t) \ge f(s), f(0) = 0 \}$$

In particular, the input function F(t) and the output function F'(t), which cumulatively count the number of bits that are input to, respectively output from, a system S, are in \mathcal{F} . Throughout the paper, we assume in- and output functions to be continuous in time and space. Note that this is not a general limitation as there exist transformations between discrete and continuous time models [9].

Definition 1: (Min-plus Convolution and Deconvolution) The min-plus convolution respectively deconvolution of two functions f and g are defined to be

$$(f \otimes g)(t) = \inf_{0 \le s \le t} \left\{ f(t-s) + g(s) \right\}$$

$$\left(f \oslash g\right)(t) = \sup_{u \ge 0} \left\{f(t+u) - g(u)\right\}$$

It can be shown that the triple $(\mathcal{F}, \wedge, \otimes)$, where \wedge denotes the minimum operator (which ought to be taken pointwise for functions), constitutes a dioid [9]. Also, the min-plus convolution is a linear operator on the dioid $(\mathbb{R} \cup \{+\infty\}, \wedge, +)$, whereas the min-plus deconvolution is not. These algebraic characteristics result in a number of rules that apply to those operators, many of which can be found in [9], [11].

Let us now turn to the performance characteristics of flows which can be bounded by network calculus means:

Definition 2: (Backlog and Delay) Assume a flow with input function F that traverses a system S resulting in the output function F'. The backlog of the flow at time t is defined as

$$x(t) = F(t) - F'(t)$$

Assuming FIFO delivery, the virtual delay for a bit input at time t is defined as

$$d(t) = \inf \{ \tau \ge 0 : F(t) \le F'(t+\tau) \}$$

Next, the arrival and departure processes specified by input and output functions are bounded based on the central network calculus concepts of arrival and service curves:

Definition 3: (Arrival Curve) Given a flow with input function F a function $\alpha \in \mathcal{F}$ is an arrival curve for F iff

$$\forall t,s \geq 0, s \leq t : F(t) - F(t-s) \leq lpha(s) \Leftrightarrow F = F \otimes lpha$$

A typical example of an arrival curve is given by an affine arrival curve $\gamma_{r,b}(t) = b + rt$, t > 0 and $\gamma_{r,b}(t) = 0$, $t \le 0$ which corresponds to token-bucket traffic regulation.

Definition 4: (Service Curve) If the service provided by a system S for a given input function F results in an output function F' we say that S offers a service curve β iff

$$F' \ge F \otimes \beta$$

A typical example of a service curve is given by a socalled rate-latency function $\beta_{R,T}(t) = R[t-T]^+$, where the [.]⁺ operator is defined as $[x]^+ = x \lor 0$, and \lor denotes the maximum operator. A number of systems fulfill, however, a stricter definition of service curve [9], which is particularly useful as it permits certain derivations that are not feasible under the more general minimum service curve model.

Definition 5: (Strict Service Curve) Let $\beta \in \mathcal{F}$. System S offers a *strict* service curve β to a flow if, during any backlogged period of duration u the output of the flow is at least equal to $\beta(u)$.

Note that any strict service curve is also a service curve, but not the other way around. Many schedulers offer strict service curves, for example most of the generalized processor sharingemulating schedulers offer a strict service curve of the ratelatency type. Strict service curves will play a crucial role in this paper, since they, in contrast to service curves, allow to bound the maximum backlogged period of a system. More specifically, that bound \overline{d} is given as the non-zero intersection point between arrival and service curve, i.e. $\alpha(\overline{d}) = \beta(\overline{d})$.

Using those concepts it is possible to derive *tight* performance bounds on backlog, (virtual) delay and output:

Theorem 1: (Performance Bounds) Consider a system S that offers a service curve β . Assume a flow F traversing the

system has an arrival curve α . Then we obtain the following performance bounds:

Backlog: $\forall t : x(t) \leq (\alpha \oslash \beta) (0) =: v(\alpha, \beta)$ Delay: $\forall t : d(t) \leq \inf \{t \geq 0 : (\alpha \oslash \beta) (-t) \leq 0\} =: h(\alpha, \beta)$ Output (arrival curve α' for F'): $\alpha' = \alpha \oslash \beta$

Note that, if FIFO cannot be assumed, the bound on the maximum backlogged period under the assumption of a strict service curve can be used as an alternative delay bound instead of the horizontal deviation $h(\alpha, \beta)$.

One of the strongest results of network calculus (albeit being a simple consequence of the associativity of \otimes) is the concatenation theorem that enables us to investigate tandems of systems as if they were single systems:

Theorem 2: (Concatenation Theorem for Tandem Systems) Consider a flow that traverses a tandem of systems S_1 and S_2 . Assume that S_i offers a service curve β_i , i = 1, 2 to the flow. Then the concatenation of the two systems offers a service curve $\beta_1 \otimes \beta_2$ to the flow.

Using the concatenation theorem, it is ensured that an endto-end analysis of a tandem of servers still achieves tight performance bounds, which in general is not the case for an iterative per-node application of Theorem 1.

So far we have only covered the single flow case, the next result factors in the existence of other interfering flows. In particular, it states the minimum service curve available to a flow at a single node under cross-traffic from other flows at that node.

Theorem 3: (Left-over Service Curve under Arbitrary Multiplexing) Consider a node multiplexing two flows 1 and 2 in arbitrary order. Assume that the node guarantees a *strict* minimum service curve β to the aggregate of the two flows. Assume that flow 2 has α_2 as an arrival curve. Then

$$\beta^1 = \left[\beta - \alpha_2\right]^+$$

is a service curve for flow 1 if $\beta^1 \in \mathcal{F}$, often also called the *left-over* service curve for the *flow of interest*. Note that we require the service curve to be strict. In [9], an example is given showing that the theorem otherwise would not hold.

III. NETWORK CALCULUS IN FEED-FORWARD NETWORKS

In this section, several methods of computing delay bounds in feed-forward networks of arbitrary multiplexers are presented. All of them are based on the direct application of well-known network calculus results.

A. Network Calculus Based Bounding Methods

The methods we present in the following compute delay bounds for a certain flow of interest. In order to do so all of them require to compute the network-internal flow constraints of each flow that is interfering with the flow of interest. This can be easily done in a feed-forward network by the application of the output bound from Theorem 1 and noting that the multiplexing of flows is performed by their addition.

For ease of exposition we now present the different alternatives in probably the most simple conceivable example scenario as illustrated in Figure 1. Despite its simplicity this scenario will already exhibit that any of the presented network calculus based methods runs into problems with respect to achieving tight delay bounds. 1) Total Flow Analysis (TFA): The first bounding method is probably the most direct application of basic network calculus results and is already mentioned in [2] to show that feedforward networks are stable. The idea of this method is to compute per-node delay bounds for the total traffic offered to the respective node, which is why it has been called total flow analysis (TFA), and then sum those up for the end-to-end delay bound of the flow of interest. Under the assumption of arbitrary multiplexing this means the per-node delay bounds have to be computed as the maximum busy periods at the nodes because the total flow must be considered to be non-FIFO.

For the example scenario in Figure 1, the TFA delay bound for flow 1 can be computed as follows

$$d^{TFA} = d_1 + d_2$$
 with

 $(\alpha_1 + \alpha_2)(d_1) = \beta_1(d_1), ((\alpha_1 + \alpha_2) \oslash \beta_1)(d_2) = \beta_2(d_2)$

2) Separated Flow Analysis (SFA): An obvious weakness of the TFA is that it makes no use of the concatenation theorem, which is known to provide a clear advantage over additive per-node bounds. To this end, instead of treating the total flow, the next method first separates the service provided to the flow of interest at each node, which is why it is called separated flow analysis (SFA), before applying the concatenation theorem to the tandem of left-over service curves. The separation is based on Theorem 3 for the left-over service at a single node under arbitrary multiplexing. Note that in contrast to the TFA, the horizontal deviation can be used for the calculation of the end-to-end delay bound because each flow is assumed to be served in FIFO order (FIFO-permicroflow).

For the example scenario in Figure 1, the SFA delay bound for flow 1 can be computed as follows

$$d^{SFA} = h\left(\alpha_1, \left[\beta_1 - \alpha_2\right]^+ \otimes \left[\beta_2 - \left(\alpha_2 \oslash \beta_1\right)\right]^+\right) \quad (1)$$

3) TFA vs. SFA: In comparison between TFA and SFA, SFA is a clear winner due to the application of the concatenation theorem. Let us illustrate this fact by setting $\alpha_i = \gamma_{r_i,b_i}$ and $\beta_i = \beta_{R_i,T_i}$ for i = 1, 2 in the example scenario of Figure 1. We obtain the following delay bounds

$$d^{TFA} = T_1 + T_2 + \frac{b_1 + b_2 + (r_1 + r_2)T_1}{R_1 - r_1 - r_2} + \frac{b_1 + b_2 + (r_1 + r_2)(T_1 + T_2)}{R_2 - r_1 - r_2}$$
$$d^{SFA} = T_1 + T_2 + \frac{b_1}{(R_1 \wedge R_2) - r_2} + \frac{b_2 + r_2T_1}{R_1 - r_2} + \frac{b_2 + r_2(T_1 + T_2)}{R_2 - r_2}$$

It can be easily checked that $d^{TFA} \ge d^{SFA}$ is always true and thus the use of SFA is strictly superior to TFA. However, taking a closer look at Equation (1) it becomes clear that d^{SFA} cannot be a tight delay bound since the service at node 1 would be required to be lazy and infinitely fast after the latency at the same time, which can obviously not be the case for any sample path of the system. This observation is also exhibited



Figure 1. Simple two nodes, two flows scenario.

in the fact that the burst term b_2 appears twice in d^{SFA} , i.e. multiplexing with the interfering flow is *paid twice*. However, the flow of interest cannot be overtaken by the interfering flow's burst twice.

4) Pay Multiplexing Only Once SFA (PMOO-SFA): A problem with the SFA is the order in which it applies the arbitrary multiplexing theorem and the min-plus convolution. In fact, there is some choice: in the example scenario from Figure 1, we could also first take the convolution of the two nodal service curves and afterwards apply the arbitrary multiplexing to the resulting single node system. In general, the idea is to concatenate single systems first, in order to be able to pay for the multiplexing with interfering flows only once, which is why this method is called pay multiplexing only once SFA (PMOO-SFA).

Specifically in the example scenario from Figure 1, a delay bound can be calculated as¹

$$d^{PMOO} = h\left(\alpha_1, \left[\left(\beta_1 \otimes \beta_2 - \alpha_2\right)\right]^+\right)$$

Instantiating $\alpha_i = \gamma_{r_i,b_i}$ and $\beta_i = \beta_{R_i,T_i}$ for i = 1, 2 in the example scenario, we obtain

$$d^{PMOO} = T_1 + T_2 + \frac{b_1 + b_2 + r_2(T_1 + T_2)}{(R_1 \wedge R_2) - r_2}$$
(2)

So, as can be observed, based on a clever application order of convolution and the arbitrary multiplexing result, the multiplexing with flow 2 is paid only once (b_2 appears only once). Actually, d^{PMOO} has an intuitive form, with the sum of the nodal latencies and each burst term as well as the burstiness increases at each of the nodes paid at the minimum residual rate of the nodes. So, PMOO-SFA seems to constitute a nice application of the basic network calculus results. In fact, we provided a generalization of this method in [19], where the main issue is dealing with overlapping interference scenarios as depicted for the canonical example in Figure 2 (on page 4).

B. Network Calculus Crisis: Convolution Considered Harmful

So far, everything seems to turn out well. However, taking a closer look at the delay bounds for the example scenario for SFA and PMOO-SFA reveals a startling phenomenon: SFA can outperform PMOO-SFA! In particular, let $b_2 = 0, T_1 = 0$ then we obtain

$$d^{SFA} = T_2 + \frac{b_1}{(R_1 \wedge R_2) - r_2} + \frac{r_2 T_2}{R_2 - r_2}$$
$$d^{PMOO} = T_2 + \frac{b_1 + r_2 T_2}{(R_1 \wedge R_2) - r_2}$$

which means that the difference

$$d^{PMOO} - d^{SFA} = r_2 T_2 \left(\frac{1}{(R_1 \wedge R_2) - r_2} - \frac{1}{R_2 - r_2} \right) \ge 0,$$

i.e. the improvement of SFA over PMOO-SFA, in this special case, can be made arbitrarily large by e.g. increasing the rate at the second server R_2 . While it is only a special case and in many cases the PMOO-SFA is superior to the SFA (set for example $T_2 = 0$ and the PMOO-SFA always outperforms the SFA), it nevertheless shows that the PMOO-SFA cannot give a tight delay bound in all scenarios. So the question is what goes wrong here, why can we not find a tight bound for this

¹Note that since $\beta_1 \otimes \beta_2$ is not necessarily strict, this calculation may not safely be done at this stage; however it will be justified by our later findings.

simple scenario by the application of basic network calculus results?

With the SFA, it is clear as discussed above that it does not give a tight bound because it does not correspond to a realizable sample path of the system. Yet, the PMOO-SFA seems like a perfect application of network calculus. In fact, we encounter in this seemingly simple and innocent looking scenario a case where the application of the min-plus convolution is detrimental with respect to finding tight bounds. This can be explained physically when carefully examining the PMOO-SFA delay bound in Equation (2): the burstiness increase of flow 2 due to the latency of node 2, r_2T_2 , is paid at the minimum of the residual rates of node 1 and 2. However, that burstiness increase can only be experienced at node 2 and never at node 1, so the PMOO-SFA delay bound cannot capture this physical reality. The convolution effectively "swallows" the topological information that the burstiness increase due to node 2 can only be paid at node 2 (or subsequent nodes in larger scenarios), but not at node 1. The min-plus convolution necessarily is blind for such topological details because it is by definition commutative and thus the order of nodes cannot matter to it when concatenating them.

In this simple example, a very fundamental problem with the direct application of basic network calculus results is exhibited, because the algebraic structure of network calculus being a dioid breaks down here since commutativity is lost. Any aggregate multiplexing that is non-FIFO, even if more knowledge than the arbitrary multiplexing assumption is given, runs into this problem and can thus not be dealt with in the conventional framework of network calculus. The physical reason for this is the following: Without FIFO, there is the possibility at any given node that data from interfering flows which the flow of interest has not yet encountered on previous nodes will overtake the flow of interest, and will thus result in a burstiness increase of the interfering traffic. This burstiness increase must be accounted for at servers downstream starting from the point where it overtook the flow of interest but not on servers upstream from this point.

IV. OPTIMIZATION-BASED BOUNDING METHOD

In this section, we present an alternative method to compute delay bounds in feed-forward networks of arbitrary multiplexing nodes. This method consists of formulating an optimization problem based on the knowledge about arrival constraints of interfering traffic flows and service guarantees provided by each node. Since the formulation of the general feedforward network case is notationally heavy, we present the method for the canonical example of overlapping interference displayed in Figure 2. In fact, this example captures all of the main difficulties for the application of the optimizationbased bounding method (its generalization can be found in [35]). The optimization problem is set up under general arrival



Figure 2. Overlapping interference scenario.

 $\langle 0 \rangle$

 $\langle 0 \rangle$

and strict service curves and solved for the case of tokenbucket constrained flows and rate-latency servers. Next we demonstrate that the solution to the optimization problem arrives at a tight bound by providing a sample path in which the delay bound is actually experienced. At the end of this section, we prove a result on how to compute a left-over service curve under piece-wise linear concave arrival and convex service curves, a very relevant case in practice.

A. The Case of Overlapping Interference

Let us assume the canonical example scenario of overlapping interference from Figure 2 with the respective denotations. Furthermore, assume that the nodes provide strict service curves $\beta_k, k = 1, 2, 3$.

Assume $0 \le t_0 \le t_1 \le t_2 \le t_3$ such that t_{k-1} is the start of the last backlogged period at node k before t_k . From the strict service curve property at each of the nodes it is ensured that for k = 1, 2, 3

$$F^{(k)}(t_k) - F^{(k-1)}(t_{k-1}) \ge \beta_k (t_k - t_{k-1})$$

where $F^{(k)}$ denotes the total flow entering node k+1. Due to the wide-sense increasing nature of the input and output functions and since $F_i^{(k)}(t_k) \ge F_i^{(k)}(t_{k-1}) = F_i^{(k-1)}(t_{k-1})$ for the selected t_k such that $F_i^{(k)}(t_k) - F_i^{(k-1)}(t_{k-1}) \ge 0$ this can be rewritten as

$$\begin{split} F_{1}^{(3)}\left(t_{3}\right) &- F_{1}^{(2)}\left(t_{2}\right) \\ &\geq \left[\beta_{3}\left(t_{3}-t_{2}\right) - \left(F_{3}^{(3)}\left(t_{3}\right) - F_{3}^{(2)}\left(t_{2}\right)\right)\right]^{+} \\ F_{1}^{(2)}\left(t_{2}\right) &- F_{1}^{(1)}\left(t_{1}\right) \\ &\geq \left[\beta_{2}\left(t_{2}-t_{1}\right) - \left(F_{3}^{(2)}\left(t_{2}\right) - F_{3}^{(1)}\left(t_{1}\right)\right)\right] \\ &- \left(F_{2}^{(2)}\left(t_{2}\right) - F_{2}^{(1)}\left(t_{1}\right)\right)\right]^{+} \\ F_{1}^{(1)}\left(t_{1}\right) &- F_{1}^{(0)}\left(t_{0}\right) \\ &\geq \left[\beta_{1}\left(t_{1}-t_{0}\right) - \left(F_{2}^{(1)}\left(t_{1}\right) - F_{2}^{(0)}\left(t_{0}\right)\right)\right]^{+} \end{split}$$

These can be added up so that

$$F_{1}^{(3)}(t_{3}) - F_{1}^{(0)}(t_{0}) \\\geq \left[\beta_{3}(t_{3} - t_{2}) - \left(F_{3}^{(3)}(t_{3}) - F_{3}^{(2)}(t_{2})\right)\right]^{+} \\+ \left[\beta_{2}(t_{2} - t_{1}) - \left(F_{3}^{(2)}(t_{2}) - F_{3}^{(1)}(t_{1})\right) \\- \left(F_{2}^{(2)}(t_{2}) - F_{2}^{(1)}(t_{1})\right)\right]^{+} \\+ \left[\beta_{1}(t_{1} - t_{0}) - \left(F_{2}^{(1)}(t_{1}) - F_{2}^{(0)}(t_{0})\right)\right]^{+}$$

From the arrival constraints we know that

$$F_{2}^{(1)}(t_{1}) - F_{2}^{(0)}(t_{0}) \leq F_{2}^{(0)}(t_{1}) - F_{2}^{(0)}(t_{0}) \leq \alpha_{2}(t_{1} - t_{0})$$

$$F_{2}^{(2)}(t_{2}) - F_{2}^{(0)}(t_{0}) \leq F_{2}^{(0)}(t_{2}) - F_{2}^{(0)}(t_{0}) \leq \alpha_{2}(t_{2} - t_{0})$$

$$F_{3}^{(2)}(t_{2}) - F_{3}^{(1)}(t_{1}) \leq F_{3}^{(1)}(t_{2}) - F_{3}^{(1)}(t_{1}) \leq \alpha_{3}(t_{2} - t_{1})$$

$$F_{3}^{(3)}(t_{3}) - F_{3}^{(1)}(t_{1}) \leq F_{3}^{(1)}(t_{3}) - F_{3}^{(1)}(t_{1}) \leq \alpha_{3}(t_{3} - t_{1})$$

Now comes an important step: we introduce slack variables

in order to rewrite the previous inequations as

$$\begin{array}{rcl} F_2^{(1)}\left(t_1\right) - F_2^{(0)}\left(t_0\right) &=& \alpha_2\left(t_1 - t_0\right) - s_2^{(1)} \\ F_2^{(2)}\left(t_2\right) - F_2^{(1)}\left(t_1\right) &\leq& \alpha_2\left(t_2 - t_0\right) - \alpha_2\left(t_1 - t_0\right) + s_2^{(1)} \\ F_3^{(2)}\left(t_2\right) - F_3^{(1)}\left(t_1\right) &=& \alpha_3\left(t_2 - t_1\right) - s_3^{(2)} \\ F_3^{(3)}\left(t_3\right) - F_3^{(2)}\left(t_2\right) &\leq& \alpha_3\left(t_3 - t_1\right) - \alpha_3\left(t_2 - t_1\right) + s_3^{(2)} \end{array}$$

where $0 \le s_2^{(1)} \le \alpha_2 (t_1 - t_0)$ and $0 \le s_3^{(2)} \le \alpha_3 (t_2 - t_1)$. The slack variables $s_i^{(k)}$ have the following interpretation: they account for how much of the accumulated burstiness of the respective flow i is transferred from node k to the next node. As such they will be the decision variables for the optimization problem to be solved, because the whole bounding problem is about where each flow's burstiness is to be paid. Taking it all together and using the logical condition $C_1 = (0 \le t_0 \le t_1 \le t_2 \le t_3)$ we arrive at

$$\begin{aligned} F_1^{(3)}(t_3) &- F_1^{(0)}(t_0) \geq \beta^1 \left(t_3 - t_0 \right) \\ &= & \inf_{\mathcal{C}_1} \left\{ \begin{bmatrix} \beta_3 \left(t_3 - t_2 \right) - \left(\alpha_3 \left(t_3 - t_1 \right) - \alpha_3 \left(t_2 - t_1 \right) + s_3^{(2)} \right) \end{bmatrix}^+ \\ &+ \begin{bmatrix} \beta_2 \left(t_2 - t_1 \right) - \left(\alpha_3 \left(t_2 - t_1 \right) - s_3^{(2)} \right) \\ &- \left(\alpha_2 \left(t_2 - t_0 \right) - \alpha_2 \left(t_1 - t_0 \right) + s_2^{(1)} \right) \end{bmatrix}^+ \\ &+ \begin{bmatrix} \beta_1 \left(t_1 - t_0 \right) - \left(\alpha_2 \left(t_1 - t_0 \right) - s_2^{(1)} \right) \end{bmatrix}^+ \end{aligned} \end{aligned}$$

Thus we have to solve a constrained optimization problem in order to find the left-over service curve of flow 1 with which we can then compute the delay bound according to Theorem 1 in the conventional manner. The optimization problem cannot be solved in the general setting assumed so far. When setting $\beta_i = \beta_{R_i,T_i}$ and $\alpha_i = \gamma_{r_i,b_i}$, i = 1, 2, 3, we proceed under the logical condition

$$\mathcal{L}_2 = \left(s_2^{(1)} \le b_2 + r_2 T_1, \ s_3^{(2)} \le b_3 + r_3 \left(T_2 + \frac{s_2^{(1)} + r_2 T_2}{R_2 - r_2} \right) \right)$$

(for details on this condition see the appendix) as follows: $\beta^1 (t_3 - t_0)$

$$\begin{split} &= \inf_{C_1, C_2} \left\{ \left[R_3 \left[t_3 - t_2 - T_3 \right]^+ - \left(r_3 \left(t_3 - t_2 \right) + s_3^{(2)} \right) \right]^+ \right. \\ &+ \left[R_2 \left[t_2 - t_1 - T_2 \right]^+ - \left(b_3 + r_3 \left(t_2 - t_1 \right) - s_3^{(2)} \right) \right. \\ &- \left(r_2 \left(t_2 - t_1 \right) + s_2^{(1)} \right) \right]^+ \\ &+ \left[R_1 \left[t_1 - t_0 - T_1 \right]^+ - \left(b_2 + r_2 \left(t_1 - t_0 \right) - s_2^{(1)} \right) \right]^+ \right\} \\ &\geq \inf_{C_1, C_2} \left\{ \left(R_3 - r_3 \right) \left[t_3 - t_2 - T_3 - \frac{r_3 T_3 + s_3^{(2)}}{R_3 - r_3} \right]^+ \\ &+ \left(R_2 - r_2 - r_3 \right) \left[t_2 - t_1 - T_2 - \frac{b_3 + \left(r_2 + r_3 \right) T_2 - s_3^{(2)} + s_2^{(1)}}{R_2 - r_2 - r_3} \right]^+ \\ &+ \left(R_1 - r_2 \right) \left[t_1 - t_0 - T_1 - \frac{b_2 + r_2 T_1 - s_2^{(1)}}{R_1 - r_2} \right]^+ \\ &= \left(\left(R_1 - r_2 \right) \wedge \left(R_2 - r_2 - r_3 \right) \wedge \left(R_3 - r_3 \right) \right) \times \\ &\qquad \inf_{C_2} \left\{ \left[t_3 - t_0 - \left(T_1 + T_2 + T_3 \right) - \frac{r_3 T_3 + s_3^{(2)}}{R_3 - r_3} \right]^+ \\ \end{split}$$

$$-\frac{b_3 + (r_2 + r_3)T_2 - s_3^{(2)} + s_2^{(1)}}{R_2 - r_2 - r_3} - \frac{b_2 + r_2T_1 - s_2^{(1)}}{R_1 - r_2} \right]^+ \left\{$$

Hence, to find a closed form for the left-over service curve is equivalent to solving the following optimization problem:

$$\min \left(\frac{1}{R_1 - r_2} - \frac{1}{R_2 - r_2 - r_3}\right) s_2^{(1)} + \left(\frac{1}{R_2 - r_2 - r_3} - \frac{1}{R_3 - r_3}\right) s_3^{(2)}$$

subject to C_2 . This constitutes a simple linear program that can be solved easily when the parameters are given. The general solution of this problem depends on the relations between the residual rates at the nodes. With $A := \frac{1}{R_1 - r_2} - \frac{1}{R_2 - r_2 - r_3}$, $B := \frac{1}{R_2 - r_2 - r_3} - \frac{1}{R_3 - r_3}$, and $C := \frac{r_3}{R_2 - r_2}$ the solution vector for the slack variables $\left(s_2^{(1)}, s_3^{(2)}\right)$ is obtained as

$$\begin{array}{l} A \leq 0, B > 0 : (b_{2} + r_{2}T_{1}, 0) \\ A \leq 0, B \leq 0 : \left(b_{2} + r_{2}T_{1}, b_{3} + r_{3}\left(T_{2} + \frac{b_{2} + r_{2}T_{1} + r_{2}T_{2}}{R_{2} - r_{2}}\right)\right) \\ A > 0, B > 0 : (0, 0) \\ A > 0, B \leq 0, \\ -CB \leq A \\ A > 0, B \leq 0, \\ -CB > A \end{array} : \left(0, b_{3} + r_{3}\left(T_{2} + \frac{r_{2}T_{2}}{R_{2} - r_{2}}\right)\right) \\ A > 0, B \leq 0, \\ -CB > A \\ \end{array} : \left(b_{2} + r_{2}T_{1}, b_{3} + r_{3}\left(T_{2} + \frac{b_{2} + r_{2}T_{1} + r_{2}T_{2}}{R_{2} - r_{2}}\right)\right) \end{array}$$

From this the left-over (l.o.) service curve for flow 1 can be given. We state it for the last case $(A > 0, B \le 0, -CB > A)$:

$$\beta^1 = \beta_{R^{l.o.}, T^{l.o.}} \tag{3}$$

with

$$\begin{array}{lll} R^{l.o.} & = & \left(R_1 - r_2\right) \wedge \left(R_2 - r_2 - r_3\right) \wedge \left(R_3 - r_3\right) \\ T^{l.o.} & = & T_1 + T_2 + T_3 + \frac{b_3 + r_3 (T_2 + T_3)}{R_3 - r_3} \\ & & + \frac{b_2 + r_2 (T_1 + T_2)}{R_2 - r_2 - r_3} + \frac{(R_2 - r_2 - R_3) r_3 \frac{b_2 + r_2 (T_1 + T_2)}{R_2 - r_2}}{(R_2 - r_2 - r_3)(R_3 - r_3)} \end{array}$$

This case corresponds to a situation where $R_1 - r_2 < R_2 - r_2 - r_3$ (A > 0) and $R_2 - r_2 - r_3 \ge R_3 - r_3$ ($B \le 0$) and $r_3 > \frac{(R_2 - R_1)R_3}{R_2 - r_2}$ (-CB > A). It is interesting because it constitues a case where the burst of flow 2 is not paid at node 1, although node 1's residual service is slower than that for node 2. Instead it is paid at node 2, because its effect of increasing the burstiness of flow 3, which is then paid at node 3, is stronger than if it had been paid at node 1. This shows that for some scenarios the delay bound can depend on fairly contrived conditions on the parameters of the scenario.

Setting $F_3 = 0$ and $\beta_3 = +\infty$ in the overlapping interference scenario results in the simple example scenario from Section III. Here, the left-over service curve for flow 1 and the corresponding delay bound are

$$\beta^{1} = \beta_{(R_{1} \wedge R_{2}) - r_{2}, T_{1} + T_{2} + \frac{b_{2} + r_{2}T_{1}}{(R_{1} \wedge R_{2}) - r_{2}} + \frac{r_{2}T_{2}}{R_{2} - r_{2}}}$$
$$d^{OPT} = h\left(\alpha_{1}, \beta^{1}\right) = T_{1} + T_{2} + \frac{b_{1} + b_{2} + r_{2}T_{1}}{(R_{1} \wedge R_{2}) - r_{2}} + \frac{r_{2}T_{2}}{R_{2} - r_{2}}$$

The burstiness increase due to node 2 (r_2T_2) is now correctly accounted for at node 2 instead of node 1 as for the PMOO-SFA. At the same time the burst terms are only accounted for once, thus the PMOO gain is still retained.

B. Tightness of Delay Bound

Using the optimization-based method to find the left-over service curve for a flow of interest actually achieves a tight delay bound. Let us illustrate this for the interesting case treated in the previous subsection $(A > 0, B \le 0, -CB > A)$. The other cases can be treated similarly, but are simpler. Assume the left-over service curve in Equation (3) is tight, i.e. it is the largest service curve (for all interval lengths) that can be given to the flow of interest. Then the delay bound can be easily seen to be tight from the tightness of the delay bound in a simple single node system as mentioned in Section II for Theorem I. Hence, it has to be shown that the left-over service curve in Equation (3) is tight. First we show the tightness of the latency, i.e. there is actually a sample path such that the system offers no service to the flow of interest for the duration of the latency:

We track a specific bit of flow 1 (the flow of interest), called the bit-under-observation (b-u-o), which arrives at node 1 at time t_0 , and create a worst-case sample path for the time until this bit leaves node 3. This constitues the latency for the flow of interest.

All interfering flows are assumed to be constantly claiming their sustained rates r_i , however when they burst will be detailed below. Similarly, when the servers take their latencies is given below. Note that not all the servers are lazy, i.e. some provide more than their service curves ensure (see 2) below). In fact, otherwise the worst-case sample path with respect to the latency would not be attained.

- 1) At time t_0 : node 1 takes its latency T_1 ; flow 2 bursts.
- 2) At time $t_1 = t_0 + T_1$: node 1 starts serving infinitely fast, thus the backlog of flow 2, $b_2 + r_2T_1$, and the bu-o are passed on to node 2; node 2 immediately takes its latency; flow 3 bursts.
- 3) At time t₂ = t₁ + T₂: node 2 starts serving flow 2 at the minimal rate R₂ with strict priority over flow 3; flow 2 has a backlog at node 2 of x₂⁽²⁾ (t₂) = b₂+r₂ (T₁ + T₂).
 4) At time t₃ = t₂ + x₂⁽ⁱ⁾(t₂) / R₂-r₂: flow 3 starts to be served by node 2 with strict priority over flow 1; flow 3 has a (2)
- 4) At time $t_3 = t_2 + \frac{x_2^{(2)}(t_2)}{R_2 r_2}$: flow 3 starts to be served by node 2 with strict priority over flow 1; flow 3 has a backlog at node 2 of $x_3^{(2)}(t_3) = b_3 + r_3T_2 + r_3\frac{x_2^{(2)}(t_2)}{R_2 - r_2}$; at the same time bits from flow 2 arrive at rate r_2 and are still served with priority over flow 3, hence reducing the service rate for flow 3 to $R_2 - r_2$.
- the service rate for flow 3 to $R_2 r_2$. 5) At time $t_4 = t_3 + \frac{x_3^{(2)}(t_3)}{R_2 - r_2 - r_3}$: the b-u-o leaves node 2 and is passed on to node 3; node 3 takes its latency; flow 3 has a backlog at node 3 of $x_3^{(3)}(t_4) = (R_2 - r_2 - R_3) \frac{x_3^{(2)}(t_3)}{R_2 - r_2 - r_3}$. Note that $B \leq 0$ implies $R_2 - r_2 \geq R_3$ and thus a backlog for flow 3 builds up at node 3.
- 6) At time t₅ = t₄ + T₃: flow 3 is continued to be served by node 3 with strict priority over flow 1; flow 3 has a backlog of x₃⁽³⁾ (t₅) = r₃T₃ + x₃⁽³⁾ (t₄).
 7) At time t₆ = t₅ + x₃⁽³⁾(t₅): the b-u-o leaves node 3.
- 7) At time $t_6 = t_5 + \frac{x_3 t_5}{R_3 r_3}$: the b-u-o leaves node 3. This sample path is also illustrated in Figure 3. The latency for the b-u-o can thus be computed as

$$t_6 - t_0 = T_1 + T_2 + T_3 + \frac{b_2 + r_2(T_1 + T_2)}{R_2 - r_2} + \frac{b_3 + r_3 T_2 + r_3 \frac{b_2 + r_2(T_1 + T_2)}{R_2 - r_2 - r_3}}{R_2 - r_2 - r_3}$$

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings.

+
$$\frac{r_3T_3 + (R_2 - r_2 - R_3)\frac{b_3 + r_3T_2 + r_3\frac{b_2 + r_2(T_1 + T_2)}{R_2 - r_2}}{R_2 - r_2 - r_3}}{R_3 - r_3}$$

This expression can be checked to be equal to $T^{l.o.}$ and thus the latency of the left-over service curve is shown to be tight because it is attained for the above described sample path.

For a time interval longer than the latency, $t > T^{l.o.}$, the service curve can be seen to be tight by considering the following sample path: During a period of length $t - T^{l.o.}$ each interfering flow creates traffic according to its sustained rate. This is followed by a period of length $T^{l.o.}$ for which the sample path is the one described above to illustrate the tightness of the latency. Note that it is important to take the outage time of the service for the flow of interest (the latency of the left-over service curve) at the end of the interval. Otherwise, if for example the last server is not the bottleneck server, this server could play out buffered data faster than at the minimum residual rate over all nodes, resulting in a seemingly larger service curve. Yet, this would not constitute the worstcase sample path.

C. Generalization to Piecewise Linear Curves

Based on the method presented in the previous subsections, we can compute delay bounds for the case of token-bucket constrained interfering flows and rate-latency servers. Now we provide a result that allows to generalize that method to the frequent case of piecewise linear concave arrival and convex service curves.

Proposition 1: Given piecewise linear concave arrival curves $\alpha_i = \bigwedge_{k_i=1}^{m_i} \gamma_{r_{k_i},b_{k_i}}$ for each interfering flow $i = 1, \ldots, n$ and piecewise linear convex strict service curves $\beta_j = \bigvee_{l_j=1}^{m_j} \beta_{R_{l_j},T_{l_j}}$ for each node $j = 1, \ldots, m$ on the path of the flow of interest, the left-over service curve for the flow of interest is given by

$$\beta^{l.o.} \ge \bigvee_{i=1}^{n} \bigvee_{j=1}^{m} \bigvee_{k_i=1}^{n_i} \bigvee_{l_j=1}^{m_j} \beta^{l.o.}_{\{k_i\},\{l_i\}}$$
(4)

where $\beta_{\{k_i\},\{l_j\}}^{l.o.}$ are end-to-end left-over service curves for a specific combination of a single token bucket per interfering flow and a single rate-latency curve per node.

Proof: Note that we can also view left-over service curves as functions $\beta^{l.o.}(\vec{\alpha},\vec{\beta})$ of the arrival curves of the interfering flows as well as the service curves of the nodes on the path of the flow of interest. It should be clear that left-over service curves are wide-sense decreasing in their arrival curve arguments and wide-sense increasing in their service curve arguments. That means, if any of the arrival curves α_i is substitued by a (strictly) larger one, then the left-over



Figure 3. Worst-case sample path illustration.

service curve can only become smaller. For the service curve arguments β_i , if any becomes larger, then the left-over service curve can only become larger, too.

Now we prove by a combination of structural induction on the given scenario and contradiction:

Induction start: For $n_i = m_j = 1; i = 1, ..., n; j = 1, ..., m$ the statement in the proposition is tautological.

Induction step: Assume the proposition to be correct for given n_i and m_j . Now suppose, at first for $n_i + 1$ Equation (4) is wrong, i.e. $\exists t \geq 0$ such that

$$\begin{split} \beta^{l.o.} \left(\left(\alpha_1, ..., \alpha_i \wedge \gamma_{r_{n_{i+1}}, b_{n_{i+1}}}, ..., \alpha_n \right), \vec{\beta} \right)(t) \\ < \beta^{l.o.} \left(\vec{\alpha}, \vec{\beta} \right)(t) \end{split}$$

or

$$\beta^{l.o.} \left(\left(\alpha_1, ..., \alpha_i \land \gamma_{r_{n_i+1}, b_{n_{i+1}}}, ..., \alpha_n \right), \vec{\beta} \right) (t)$$

$$< \beta^{l.o.} \left(\left(\alpha_1, ..., \gamma_{r_{n_i+1}, b_{n_{i+1}}}, ..., \alpha_n \right), \vec{\beta} \right) (t)$$

None of the two can be the case because they would violate the wide-sense decreasing nature of the left-over service curve for the arrival curve arguments. Thus the proposition is also true for $n_i + 1$. The induction step over the service curve arguments follows along similar lines. \blacksquare So, under the assumption of piecewise linear curves we now have to solve a *set* of linear programs and then compute the pointwise maximum of the service curves in this set. This can become pretty compute-intensive if the amount of linear segments used to model arrival and service curves grows. In particular, the number of left-over service curves to be computed and thus the number of linear programs to be solved is $\prod_{i=1}^{n} \prod_{j=1}^{m} n_i m_j$.

V. EXPLICIT DELAY BOUNDS FOR SINK TREES

Using the optimization-based bounding method presented in the previous section, we derive explicit delay bounds for sinktree networks in this section. Sink trees are a frequent special case of feed-forward networks. For example, MPLS networks use sink-tree aggregation of label switched paths from one edge of a network to the other. Another example are wireless sensor networks, which also often organize their topology as a sink-tree towards a base station collecting the data.

While applying the optimization based method to general feed-forward networks is well possible if the parameter values of arrival and service curves are given, it is hard to conceive that a closed-form solution for delay bounds under general arrival and service curves can be given. For sink-tree networks this is different: The optimization problem of finding the leftover service curve for a flow of interest has a form that allows to solve it under a general setting of piecewise linear concave arrival and convex service curves (without instantiation of their parameters). For a comprehensive discussion of how the optimization problem is solved see [35], here we just state the result for the left-over service curve in the following proposition. The closed-form expression of the delay bound for a flow of interest then follows simply from the application of Theorem 1 by calculating the horizontal deviation between the flow's arrival curve and its left-over service curve.

Proposition 2: Assume a sink-tree scenario as illustrated in Figure 4, note that the arrival curves α_i for the interfering



Figure 4. Sink-tree scenario.

flows can be calculated using the output bound from Theorem 1. The left-over service curve for the flow of interest under piecewise linear concave arrival curves $\alpha_i = \bigwedge_{k_i=1}^{n_i} \gamma_{r_{k_i},b_{k_i}}$ and convex strict service curves $\beta_i = \bigvee_{l_i=1}^{m_i} \beta_{R_{l_i},T_{l_i}}$ is given by $\beta^{l.o.} = \bigvee_{i=1}^{n} \bigvee_{k_i=1}^{m_i} \bigcup_{l_i=1}^{m_i} \beta_{R_{l_i},\{l_i\}} \sum_{i=1}^{T_{l.o.}^{l.o.}} \beta_{k_i,\{l_i\}} \sum_{l_i=1}^{T_{l.o.}^{l.o.}} \beta_{k_i,\{l_i\}$

with

$$R_{\{k_i\},\{l_i\}}^{l.o.} = \bigwedge_{i=1}^n \left(R_{l_i} - \sum_{j=1}^i r_{k_j} \right)$$
$$T_{\{k_i\},\{l_i\}}^{l.o.} = \sum_{i=1}^n \left(T_i + \frac{\mathbf{b}_{k_i}}{\Lambda_{j=i}^n (R_{l_j} - \sum_{s=1}^j r_{k_s})} + \sum_{i=i}^n \frac{r_{k_i} T_{l_j}}{\Lambda_{s=j}^n (R_{l_s} - \sum_{v=1}^s r_{k_v})} \right)$$

The latency terms of the left-over service curve actually have an intuitive form: Each interfering flow's burst and its burstiness increases at subsequent nodes are paid at the minimum residual rate server they actually traverse. This insight also means that if residual rates are monotonically decreasing on the path of the flow of interest then PMOO-SFA actually becomes tight. Furthermore, note that the latency terms, and therefore also the delay bound, scale *quadratically* in the number of nodes on the path of the flow of interest.

A. Numerical Experiments

We choose a simple experimental setup with fully occupied binary trees, where each node acts as a source of data. The sources are token-bucket constrained with a sustained rate of 10 Mbps and a bucket depth of 1 Mb. Each node offers a strict rate-latency service curve with a latency of 0.1 ms and the service rate dimensioned such that a certain target utilization is achieved. The server utilizations are varied over the experiments. All calculations have been performed using the DISCO Network Calculator [19].

In the first experiment we compare, under varying tree depths, all of the methods described in the paper: TFA, SFA, PMOO-SFA, and the optimization-based bounding method (denoted as TIGHT in the figures). The delay bounds for a leaf in the sink-tree under a utilization of 20% at each of the nodes are displayed in Figure 5. Clearly, we can observe a considerable improvement of the delay bounds with increasing tree depths. Interestingly, it can also be observed, that the PMOO-SFA, while being second best for tree depths up to 12, is outperformed by SFA and even TFA for larger tree depths. This stresses again the insights gained in Section III.

In a second experiment, the effect of the utilization on the delay bounds are investigated. Both, the tight bound and the PMOO-SFA are computed for utilizations of 20, 50, and



Figure 5. All bounding methods compared at a utilization of 20%.



Figure 6. PMOO-SFA vs. tight bound at several utilizations.

90%, again for varying tree depths. The results are shown in Figure 6. As can be observed, the tight bound reacts pretty mildly on higher utilizations, with an approximately linear increase of the delay bound when increasing the tree depth (the quadratic regime mentioned above only shows for even larger tree depths), while the PMOO-SFA suffers badly from higher utilizations and exhibits a steep super-linear increase in its bounds for increasing tree depths.

VI. CONCLUSIONS

In this paper, we have demonstrated that direct application of network calculus in aggregate multiplexing scenarios for which FIFO cannot be assumed is problematic with respect to the tightness of delay bounds. The nice algebraic characteristics of network calculus cannot be preserved under arbitrary multiplexing. Based on this insight, we have proposed an alternative method to arrive at tight bounds for arbitrary multiplexing in feed-forward networks. The method consists of formulating an optimization problem. For the case of piecewise linear concave arrival and convex service curves (a very frequent case), we obtain a set of linear programming problems, which can be solved by standard methods, and from which the solution for a tight delay bound can be composed. For the special case of sink-tree networks we have provided a closed-form expression of the delay bound based on this method. Numerical experiments have shown a clear edge for the novel bounding method over methods based on direct application of network calculus.

This full text paper was peer reviewed at the direction of IEEE Communications Society subject matter experts for publication in the IEEE INFOCOM 2008 proceedings

APPENDIX

In Section IV we derived the service curve β^1 under the precondition that the slack variables $s_2^{(1)}$ and $s_3^{(2)}$ have upper bounds given by C_2 . Here, we show that this assumption does not restrict the solution space of the minimization problem.

If $s_2^{(1)} > b_2 + r_2 T_1$ we can conclude from

$$s_2^{(1)} \le \alpha_2 (t_1 - t_0) = b_2 + r_2 (t_1 - t_0)$$

that $t_1 - t_0 > T_1$ in this case. We write $s_2^{(1)} = b_2 + r_2 T_1 + r_2 \tau$ where $0 < \tau \le t_1 - t_0 - T_1$ and derive $\beta^{1}(t_{3}-t_{0})$

$$\geq \inf_{C_{1}} \left\{ (R_{3} - r_{3}) \left[t_{3} - t_{2} - T_{3} - \frac{r_{3}T_{3} + s_{3}^{(2)}}{R_{3} - r_{3}} \right]^{+} \\ + (R_{2} - r_{2} - r_{3}) \left[t_{2} - t_{1} - T_{2} - \frac{b_{3} + (r_{2} + r_{3})T_{2} - s_{3}^{(2)} + b_{2} + r_{2}T_{1} + r_{2}\tau}{R_{2} - r_{2} - r_{3}} \right]^{+} \\ + (R_{1} - r_{2}) \left[t_{1} - t_{0} - T_{1} + \frac{r_{2}\tau}{R_{1} - r_{2}} \right]^{+} \right\} \\ \geq \inf_{C_{1}} \left\{ (R_{3} - r_{3}) \left[t_{3} - t_{2} - T_{3} - \frac{r_{3}T_{3} + s_{3}^{(2)}}{R_{3} - r_{3}} \right]^{+} \\ + (R_{2} - r_{2} - r_{3}) \left[t_{2} - t_{1} - T_{2} - \frac{b_{3} + (r_{2} + r_{3})T_{2} - s_{3}^{(2)} + b_{2} + r_{2}T_{1}}{R_{2} - r_{2} - r_{3}} \right]^{+} r_{2}\tau \\ + (R_{1} - r_{2})(t_{1} - t_{0} - T_{1}) + r_{2}\tau \}$$

Since
$$t_1 - t_0 - T_1 > 0$$
 we can conclude that
 $\beta^1 (t_3 - t_0) \ge ((R_1 - r_2) \land (R_2 - r_2 - r_3) \land (R_3 - r_3)) \times \left[t_3 - t_0 - T_1 - T_2 - T_3 - \frac{r_3 T_3 + s_3^{(2)}}{R_3 - r_3} - \frac{b_3 + (r_2 + r_3) T_2 - s_3^{(2)} + b_2 + r_2 T_1}{R_2 - r_2 - r_3}\right]^{+1}$

The lower bound that we derived here for $s_2^{(1)} > b_2 + r_2 T_1$ coincides with the previously derived bound instantiated at connected with the previously derived bound instantated at $s_2^{(1)} = b_2 + r_2 T_1$. Hence, we cannot find a smaller service curve for $s_2^{(1)} > b_2 + r_2 T_1$. In the same way it can be shown, that no smaller service curve can be derived for $s_3^{(2)} > b_3 + r_3 \left(T_2 + \frac{s_2^{(1)} + r_2 T_2}{R_2 - r_2}\right)$.

REFERENCES

- [1] R. L. Cruz, "A calculus for network delay, Part I: Network elements in isolation," IEEE Transactions on Information Theory, vol. 37, pp. 114-131, Jan. 1991.
- [2] R. L. Cruz, "A calculus for network delay, Part II: Network analysis," IEEE Transactions on Information Theory, vol. 37, pp. 132-141, Jan. 1991.
- [3] R. L. Cruz, "Quality of service guarantees in virtual circuit switched networks," IEEE Journal on Selected Areas in Communications, vol. 13, pp. 1048-1056, Aug. 1995.
- [4] H. Sariowan, R. L. Cruz, and G. C. Polyzos, "Scheduling for quality of service guarantees via service curves," in Proc. IEEE ICCCN, pp. 512-520, Sept. 1995.
- C.-S. Chang, "On deterministic traffic regulation and service guarantees: [5] A systematic approach by filtering," IEEE Transactions on Information Theory, vol. 44, pp. 1097–1110, May 1998. J.-Y. Le Boudec, "Application of network calculus to guaranteed service
- [6] networks," IEEE Transactions on Information Theory, vol. 44, pp. 1087-1096, May 1998
- R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan, "Performance bounds for flow control protocols," *IEEE/ACM Transactions on Networking*, 171 vol. 7, pp. 310–323, June 1999.
- F. Ciucu, A. Burchard, and J. Liebeherr, "A network service curve [8] approach for the stochastic analysis of networks," in Proc. ACM SIG-METRICS, pp. 279-290, June 2005.

- [9] J.-Y. Le Boudec and P. Thiran, Network Calculus A Theory of Deterministic Queuing Systems for the Internet. No. 2050 in Lecture Notes in Computer Science, Berlin, Germany: Springer-Verlag, 2001.
- [10] F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat, Synchronization and Linearity: An Algebra for Discrete Event Systems. Probability and Mathematical Statistics, John Wiley & Sons Ltd., 1992.
- [11] C.-S. Chang, Performance Guarantees in Communication Networks. Telecommunication Networks and Computer Systems, Springer-Verlag, 2000.
- [12] J. Schmitt and U. Roedig, "Sensor network calculus a framework for worst case analysis," in Proc. Distributed Computing on Sensor Systems (DCOSS), pp. 141-154, June 2005.
- A. Koubaa, M. Alves, and E. Tovar, "Modeling and worst-case dimen-sioning of cluster-tree wireless sensor networks," in *Proc. IEEE RTSS*, [13] pp. 412-421, 2006. [14] T. Skeie, S. Johannessen, and O. Holmeide, "Timeliness of real-time
- ip communication in switched industrial ethernet networks," IEEE Transactions on Industrial Informatics, vol. 2, pp. 25-39, Feb. 2006.
- [15] S. Chakraborty, S. Kuenzli, L. Thiele, A. Herkersdorf, and P. Sagmeister, "Performance evaluation of network processor architectures: Combining simulation with analytical estimation." Computer Networks. vol. 42. no. 5, pp. 641–665, 2003.
- [16] H. Kim and J. Hou, "Network calculus based simulation: theorems, implementation, and evaluation." in Proc. IEEE INFOCOM. Mar. 2004.
- [17] M. Fidler, "An end-to-end probabilistic network calculus with moment
- [17] M. Pidei, An end-to-end probabilistic network calculus with moment generating functions," in *Proc. of IEEE IWQoS*, pp. 261–270, Jun 2006.
 [18] Y. Jiang, "A basic stochastic network calculus," in *Proc. ACM SIG-COMM*, pp. 123–134, Sept. 2006.
 [19] J. Schmitt and F. Zdarsky, "The DISCO Network Calculator a toolbox for worst case analysis," in *Proc. of VALUETOOLS*, ACM, Nov. 2006.
 [20] A. Bouillard and E. Thierry, "An algorithmic toolbox for network for network and the USE of the U
- calculus," Tech. Rep. RR-6094, Unité de recherche INRIA Rennes, 2007. [21] C.-S. Chang, "Stability, queue length and delay of deterministic and stochastic queueing networks," *IEEE Transactions on Automatic Con-*
- trol, vol. 39, pp. 913-931, May 1994. + [22] M. Andrews, "Instability of fifo in session-oriented networks," in Proc.
 - SODA, pp. 440-447, Mar. 2000.
 - [23] A. Charny and J.-Y. L. Boudec, "Delay bounds in a network with aggregate scheduling," in Proc. QofIS, pp. 1-13, Sept. 2000.
 - Y. Jiang, "Delay bounds for a network of guaranteed rate servers with [24] fifo aggregation." *Computer Networks*, vol. 40, no. 6, pp. 683–694, 2002. [25] Z.-L. Zhang, Z. Duan, and Y. T. Hou, "Fundamental trade-offs in
 - aggregate packet scheduling," in Proc. ICNP, pp. 129-137, Nov. 2001. [26] J. A. Cobb. "Preserving quality of service guarantees in spite of flow
 - aggregation," IEEE/ACM Transactions on Networking, vol. 10, pp. 43-53, Feb. 2002.
 - [27] Y. Jiang, "Per-domain packet scale rate guarantee for expedited forwarding," IEEE/ACM Transactions on Networking, vol. 14, pp. 630-643, June 2006
 - [28] D. Starobinski, M. Karpovsky, and L. A. Zakrevski, "Application of network calculus to general topologies using turn-prohibition," IEEE/ACM
 - Transactions on Networking, vol. 11, no. 3, pp. 411-421, 2003. [29] M. Fidler and V. Sander, "A parameter based admission control for differentiated services networks," Computer Networks, vol. 44, no. 4, pp. 463-479, 2004.
 - [30] L. Lenzini, L. Martorini, E. Mingozzi, and G. Stea, "Tight end-toend per-flow delay bounds in fifo multiplexing sink-tree networks."
 - Performance Evaluation, vol. 63, no. 9, pp. 956–987, 2006.
 [31] G. Urvoy-Keller, G. Hébuterne, and Y. Dallery. "Traffic engineering in a multipoint-to-point network," *IEEE Journal on Selected Areas in* Communications, vol. 20, pp. 834–849, May 2002. [32] J.-Y. Le Boudec and A. Charny, "Packet scale rate guarantee for non-fifo
 - nodes," in Proc. IEEE INFOCOM, pp. 23-26, June 2002.
 - [33] J.-Y. Le Boudec and G. Rizzo, "Pay bursts only once does not hold for non-fifo guaranteed rate nodes." Performance Evaluation, vol. 62, no. 1-4, pp. 366-381, 2005.
 - [34] J. Echagüe and V. Cholvi, "Worst case burstiness increase due to arbitrary aggregate multiplexing," in Proc. of VALUETOOLS, ACM, Nov. 2006.
 - J. Schmitt, F. Zdarsky, and M. Fidler, "Delay bounds under arbitrary multiplexing," Technical Report 360/07, University of Kaiserslautern, Germany, July 2007. [35]