

Pay Bursts Only Once Holds for (Some) Non-FIFO Systems

Jens B. Schmitt, Nicos Gollan, Steffen Bondorf, Ivan Martinovic
Distributed Computer Systems (DISCO) Lab, University of Kaiserslautern, Germany

Abstract—Non-FIFO processing of flows by network nodes is not a rare phenomenon. Unfortunately, the state-of-the-art analytical tool for the computation of performance bounds in packet-switched networks, network calculus, cannot deal well with non-FIFO systems. The problem lies in its conventional service curve definitions. Either the definition is too strict to allow for a concatenation and consequent beneficial end-to-end analysis, or it is too loose and results in infinite delay bounds. Hence, in this paper, we propose a new service curve definition and demonstrate its strength with respect to achieving both finite delay bounds and a concatenation of systems resulting in a favorable end-to-end delay analysis. In particular, we show that the celebrated pay bursts only once phenomenon is retained even without any assumptions on the processing order of packets. This seems to contradict previous work [15]; the reasons for this are discussed.

I. INTRODUCTION

A. Motivation

In the recent past, network calculus [6], [13] has shown promise as an alternative methodology, besides classical queueing theory, for the performance analysis of packet-switched networks. It is being used as a basic tool for attacking several important network engineering problems: most prominently in the Internet’s Quality of Service proposals IntServ and DiffServ, but recently also in other environments like, e.g., wireless sensor networks [10], switched Ethernets [18], Systems-on-Chip (SoC) [5], or even to speed-up simulations [11], to name a few. Unfortunately, it comes up short in a basic aspect: its delay bound calculation depends on a strict FIFO processing of work units within the flow under analysis. This issue about how to deal with a possible *non-FIFO processing* within a flow is largely unexplored (and sometimes may be even overseen in applications). In the next subsection, we provide an overview of the rather little previous work we could find on this. In this paper, we are concerned with the scheduling within a single flow under analysis, so we deal with the questions *when* work units are processed by a node and *in which order*. The first question about how a node provides its capacity to a flow is flexibly answered by the network calculus concept of service curves, the second question has so far always been answered rigidly by assuming a FIFO processing order. The goal of this paper is to provide a more flexible answer to this. In the following, we provide some arguments why besides being of theoretical interest this issue should be addressed.

Assuming that the work units of a flow under analysis are processed in FIFO order constitutes a logical break for a worst-

case methodology as we discuss now. Assume a flow traverses a system. The *real delay*¹ for a work unit input at time t and output at time t' is simply defined as

$$rd(t) = t' - t.$$

Any processing order other than FIFO results in an increase of the worst-case real delay. This can be seen by the following argument: Assume at time t_0 a work unit which experiences the worst-case real delay is input to a FIFO system. Now assume we can change the processing order of work units. If the work unit is further delayed by scheduling work units that arrived later, then certainly the real delay of that work unit under the new processing order will be worse. Processing that work unit earlier will make its new real delay $rd'(t_0)$ smaller, yet, the work unit which was processed just ahead of the above work unit is now leaving the system when the above work unit would have left; however, that work unit has arrived at time $t_1 \leq t_0$ such that the new real delay of that work unit $rd'(t_1)$ is higher than or equal to the one from the FIFO worst-case work unit, i.e., $rd(t_0) \leq rd'(t_1)$. So, in this sense FIFO can be viewed as a best-case assumption (with respect to the worst-case delay) on the processing order of the flow under analysis. Therefore, it can be seen as consistent with a worst-case methodology to release the FIFO processing assumption.

Additional motivation to deal with non-FIFO processing comes from real-world scenarios like, e.g., packet reordering by Internet routers, link-level retransmissions in wireless networks, or content-dependent packet scheduling in wireless sensor networks. For example, in several studies of Internet traffic it has been shown that packet reordering is a frequent event (see for example [2]). According to these studies this is due to a growing amount of parallelism on a global (use of multiple paths) as well as on a local (device) level. In particular, for scalability reasons high-speed routers often contain a complex multi-stage switching fabric which cannot ensure to preserve the order of arrivals at its output. This is due to a common design trade-off where FIFO service at the input queues is relaxed in order to avoid head-of-line blocking by choosing from a set of packets in the input queue (often some window-based scheme is used).

So, from a methodological as well as an application perspective there is a clear motivation for an investigation on how network calculus can be extended towards an analysis without any FIFO assumptions.

¹The word *real* is chosen for the purpose of contrasting it to the *virtual delay*, later on defined as delay under FIFO processing of a flow.

B. Related Work

There is surprisingly little existing work on the treatment of non-FIFO systems in the context of network calculus. Remarkably, in his pioneering paper [7], Cruz briefly showed how to derive a delay bound for a single work-conserving server under a general scheduling assumption (comprising any non-FIFO processing order) based on the observation that the maximum backlogged period can be bounded given that traffic is regulated. Similar results can also be found in [6]. Yet, the multiple node case as well as more general server models are not treated therein.

In [12], Le Boudec and Charny investigate a non-FIFO version of the Packet Scale Rate Guarantee (*PSRG*) node model as used in DiffServ’s Expedited Forwarding. They show that the delay bound from the FIFO case still applies in the single node case while it does not in a two node case. They leave more general concatenation scenarios for further study.

In [17], the problem of computing tight delay bounds for a network of arbitrary (non-FIFO) *aggregate multiplexers* is addressed. The tightness of the bounding method is shown by sample path arguments. Yet, in contrast to the problem setting in this paper, a FIFO assumption on the processing order within a flow is made and non-FIFO behavior is only allowed between flows. Bouillard et al. recently provided more advanced and general results for the same setting in [4], still basing on FIFO processing per flow.

To the best of our knowledge, the only previous work that also tries to derive end-to-end delay bounds without any FIFO processing within the flow under analysis assumption was done by Rizzo and Le Boudec [15]. They investigate delay bounds for a special server model, non-FIFO guaranteed rate (*GR*) nodes, and show that a previously derived delay bound for *GR* nodes [8] is not valid in the non-FIFO case (against common belief). Furthermore, they derive a new delay bound based on network calculus results. Their delay bound does not exhibit the nice pay bursts only once phenomenon any more. Based on sample path arguments they argue that their bound is tight and thus conclude “pay bursts only once does not hold for non-FIFO guaranteed rate nodes”. In contrast, we show that non-FIFO systems may still possess a concatenation property. This puzzling contradiction is discussed in detail in Section V, since it is a subtle issue that needs to be resolved thoroughly.

II. PRELIMINARIES ON NETWORK CALCULUS

Network calculus is a min-plus system theory for deterministic queuing systems. A detailed treatment of network calculus can be found in [6] and [13].

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

$$\mathcal{F} = \{f : \mathbb{R}^+ \rightarrow \mathbb{N}_0, \forall t \geq s : f(t) \geq f(s), f(0) = 0\}.$$

In particular, the input function $F(t)$ and the output function $F'(t)$, which cumulatively count the number of work units that are input to, respectively output from, a system \mathcal{S} , are in \mathcal{F} .

Throughout the paper, we assume in- and output functions to be continuous in time and *discrete* in space unless specified otherwise. We assume work units of equal size, which is convenient for some of the discussions with respect to the processing order of work units. Arbitrarily, we assume in- and output functions to be left-continuous. There are two important min-plus algebraic operators:

Definition 1: (Min-plus Convolution and Deconvolution) The min-plus convolution and deconvolution of two functions $f, g \in \mathcal{F}$ are defined to be

$$(f \otimes g)(t) = \inf_{0 \leq s \leq t} \{f(t-s) + g(s)\},$$

$$(f \oslash g)(t) = \sup_{u \geq 0} \{f(t+u) - g(u)\}.$$

With these definitions, let us now turn to the performance characteristics of flows:

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

$$b(t) = F(t) - F'(t).$$

The *virtual delay* for a work unit input at time t is defined as

$$vd(t) = \inf \{\tau \geq 0 : F(t) \leq F'(t + \tau)\}.$$

So, this is the point where a FIFO assumption “sneaks” in network calculus as far as delay is concerned, because $rd(t) = vd(t)$ for all t only under FIFO processing of the flow. We use the usual network calculus terminology of the so-called *virtual* delay in contrast to the *real* delay as defined above (see Section I-A). Next, the central network calculus concepts of arrival and service curves are introduced:

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

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

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

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

$$F' \geq F \otimes \beta.$$

If F is left-continuous and β is continuous, this is equivalent to the following condition

$$\forall t \geq 0 : \exists s \leq t : F'(t) \geq F(s) + \beta(t-s).$$

A typical example of a service curve is given by a so-called rate-latency function $\beta_{R,T}(t) = R(t-T) \cdot 1_{\{t>T\}}$, where $1_{\{cond\}}$ is 1 if the condition *cond* is satisfied and 0 otherwise. Also, nodes operating under a delay-based scheduler and guaranteeing that a work unit arriving at any time t will leave the node at time $t' \leq t + T$ for some fixed $T > 0$, i.e. $\forall t \geq 0 : rd(t) \leq T$, are known to provide a service curve $\delta_T = \infty \cdot 1_{\{t>T\}}$. We also call these bounded latency nodes.

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

Theorem 1: (Performance Bounds) Consider a system \mathcal{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:

$$\begin{aligned} \text{backlog: } & \forall t : b(t) \leq (\alpha \otimes \beta)(0) =: v(\alpha, \beta), \\ \text{virtual delay: } & \forall t : vd(t) \leq \inf \{t \geq 0 : (\alpha \otimes \beta)(-t) \leq 0\} \\ & =: h(\alpha, \beta), \\ \text{output (arrival curve } \alpha' \text{ for } F') : & \alpha' = \alpha \otimes \beta. \end{aligned}$$

One of the strongest results of network calculus 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 \mathcal{S}_1 and \mathcal{S}_2 . Assume that \mathcal{S}_i offers a service curve β_i 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 end-to-end analysis of a tandem of servers achieves tight performance bounds, which in general is not the case for an iterative per-node application of Theorem 1. In particular, a linear scaling of performance bounds in the number of traversed queues is achieved and also referred to as pay bursts only once phenomenon in [13].

III. CONVENTIONAL NETWORK CALCULUS AND NON-FIFO SYSTEMS

In this section, we investigate how existing network calculus can cope with non-FIFO systems. The crucial aspect is the node model. We start with the typical service curve model as defined in the previous section and turn to other node models like strict and adaptive service curves, only to find out that all of them encounter problems under non-FIFO processing.

A. Using Service Curves (SC) for Non-FIFO Systems

As the *SC* definition bears the advantages that many systems belong to that class and that it possesses a concatenation property, it is tempting to apply it also in the case of non-FIFO systems. Yet, the following example shows that it is impossible to bound the real delay in non-FIFO systems based solely on the *SC* definition:

Example 1: (*SC* Cannot Bound the Real Delay) Assume a single node system \mathcal{S} which offers a rate-latency service curve $\beta = \beta_{2,1}$ to a flow F which is constrained by an affine arrival curve $\alpha = \gamma_{1,0}$. Now assume the flow to be greedy, that means $F = \alpha$ and the server to be lazy, that means $F' = F \otimes \beta$. Thus, we obtain

$$\begin{aligned} F' &= \alpha \otimes \beta = \gamma_{1,0} \otimes \beta_{2,1} = \gamma_{1,0} \otimes \gamma_{2,0} \otimes \delta_1 \\ &= (\gamma_{1,0} \wedge \gamma_{2,0}) \otimes \delta_1 \leq \gamma_{1,0} \otimes \delta_1 < \gamma_{1,0} = F. \end{aligned}$$

Hence, $\forall t \geq 0 : F'(t) < F(t)$, or equivalently, $\forall t \geq 0 : b(t) > 0$, which means the system remains backlogged at all times. Therefore, without any assumptions on the processing order

a certain work unit can be kept forever in the system under these circumstances. Thus, the real delay of that work unit is unbounded. Note that using the standard FIFO processing assumption, we can of course bound the real delay of the system by $\forall t \geq 0 : rd(t) = vd(t) \leq \frac{3}{2}$.

From this example, we see that the *SC* property is too weak as a node model for analyzing non-FIFO systems. Therefore, it is sensible to look for more stringent node models as is done in the following subsection.

B. Using Strict Service Curves (S^2C) for Non-FIFO Systems

A number of systems provide more stringent service guarantees than captured by *SC*, fulfilling a so-called strict service curve [13] (also known as strong service curve [1] and related to the universal service curve concept in [14]).

Definition 5: (Strict Service Curve – S^2C) Let $\beta \in \mathcal{F}$. System \mathcal{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 $\beta(u)$. A backlogged period of duration u at time t is defined by the fact that $\forall s \in (t - u, t] : b(s) > 0$. More formally, \mathcal{S} offers a *strict* service curve β to a flow iff

$$\forall t \geq 0 \wedge \forall u \geq 0 \wedge (\forall s \in (t - u, t] : b(s) > 0) :$$

$$F'(t - u) \geq F'(u) + \beta(t - u).$$

Note that any node satisfying S^2C also satisfies *SC*, but not vice versa. For example, a bounded latency node does not provide its service curve δ_T as a *strict* service curve. In fact, in a continuous time and space model for in- and output functions, it does not provide any S^2C apart from the trivial case $\beta = 0$ [13] (for a discrete space model with unit packets, a (very low) strict service curve can be derived for a bounded latency node as $\beta_{1/T,T}$). On the other hand, there are many schedulers that offer strict service curves; for example, most of the generalized processor sharing-emulating schedulers (e.g., PGPS [14], WF²Q [3], or round robin schedulers like SRR [9], to name a few), offer a strict service curve of the rate-latency type. More generally, all nodes providing some form of capacity guarantee can be captured by S^2C .

Now for bounding the real delay under S^2C : In fact, as was already shown by Cruz [7] (and can also be found in [6] (Lemma 1.3.2)), the *intersection* point between an arrival and a *strict* service curve constitutes a bound on the length of the maximum backlogged period and thus also a bound on the real delay for such a system:

Theorem 3: (Real Delay Bound for Single S^2C Node) Consider a system \mathcal{S} that offers a strict service curve β . Assume a flow F traversing the system has an arrival curve α . Then we obtain the following bound on the real delay:

$$rd(t) \leq \sup\{s \geq 0 : \alpha(s) \geq \beta(s)\} =: i(\alpha, \beta).$$

So, the situation has improved in comparison to the *SC* case: Based on the single node result one can conceive, for the multiple node case, an iterative application of Theorem 3 together with the output bound from Theorem 1 (note that for the latter we only require *SC*). More specifically, if a tandem

of n S^2C non-FIFO nodes, each providing a strict service curve $\beta_j, j = 1, \dots, n$, is to be traversed by an α -constrained flow then a bound on the real delay can be calculated as

$$rd(t) \leq \sum_{j=1}^n i(\alpha \otimes \bigotimes_{k=1}^{j-1} \beta_k, \beta_j).$$

Setting for example $\beta_j = \beta_{R,T}, j = 1, \dots, n$ and $\alpha = \gamma_{r,b}$ this results in

$$rd(t) \leq \frac{n(b + RT) + \frac{n}{2}(n-1)rT}{R-r} \quad (1)$$

Here, we see the typical drawback of additive bounding methods, with the burst of the traffic being paid n times as well as a quadratic scaling of the bound in the number of nodes [13]. The key to avoid this behavior is to perform an end-to-end analysis based on a concatenation theorem. Yet, as is known S^2C does not possess such a concatenation property (an example can be found in [16]). The problem is due to the fact that there may be a system backlogged period but some of the subsystems are not backlogged. Another way to view this, is that the backlogged period of a composite system cannot be bounded based on the individual systems providing a strict service curve.

So, from this discussion we can conclude that S^2C is too strict as a concept in order to allow for tight bounds under non-FIFO processing, since it possesses no concatenation property for the multiple node case. Now, since SC is too loose and S^2C too strict, it is sensible to look for some intermediate node model.

C. Using Adaptive Service Guarantees for Non-FIFO Systems

The so-called adaptive service guarantee, introduced in the context of window flow control systems [1], is a node model that is known to be falling in between SC and S^2C in a certain sense.

Definition 6: (Adaptive Service Guarantee – ASG) If the service provided by a system \mathcal{S} for a given input function F results in an output function F' we say that \mathcal{S} offers an adaptive service guarantee $(\tilde{\beta}, \beta)$ iff $\forall t \geq 0$:

$$F'(t) \geq \sup_{s < t} \left\{ \left(F'(s) + \tilde{\beta}(t-s) \right) \wedge \inf_{u \in [s,t]} \{ F(u) + \beta(t-u) \} \right\} \quad (2)$$

ASG enables to bound the *virtual* delay based on the backlog of a flow (delay-from-backlog bound), something that can also be done under S^2C but not under SC . In contrast to S^2C , however, ASG has a concatenation property (albeit a bit more complex). So, this makes it interesting because in this sense it lies in between SC and S^2C . Yet, unfortunately, it cannot bound the real delay without assuming FIFO processing of the flow under analysis (an example showing this weakness is provided in [16]). Hence, ASG suffers from the same deficiency as SC with respect to bounding the real delay, though it is a stronger guarantee but in a sense that does not help here.

D. Discussion and Some Remarks

To sum up, we started with the service curve (SC) definition as the standard network calculus node model, yet found it too weak to bound the real delay under non-FIFO processing. Next, we tried a more stringent node model, the strict service curve (S^2C). Actually, under S^2C a single-node bound on the real delay can be given. However, S^2C has no concatenation property, thus impeding a true end-to-end delay analysis. Hence, we were looking for a node model falling in between SC and S^2C . Adaptive service guarantees (ASG) are lying in this intermediate area. They have some properties similar to strict service curves but still possess a concatenation property. Unfortunately, as we demonstrated they do not provide a bound on the real delay. So they are too weak, again.

All these variants of node models are based on min-plus algebraic definitions, which effectively means nodes are characterized by the amount of work they guarantee to do in certain intervals (the intervals for which this is specified differ). In fact, there is a “dual world” of max-plus algebraic definitions for node models. Known models of this category are so-called guaranteed rate (GR) nodes [8] and packet-scale rate guarantees ($PSRG$) [13]. These work by guaranteeing certain delivery deadlines to work units, which they are assigned when arriving at a node. Since these deadlines are in absolute time, these node models can always provide a bound on the real delay (given a finite input at each point in time). Their concatenation properties are derived based on equivalence theorems with min-plus algebraic counterparts. In particular, GR is equivalent with SC (modulo packetization) and $PSRG$ is equivalent with ASG (modulo packetization) (see chapter 2 and 7 of [13] for details). However, the crucial point here is that these equivalences are only applicable for FIFO systems and are actually shown not to be applicable for concatenation of non-FIFO systems in [13], [12]. Thus, GR and $PSRG$ have no known concatenation properties under non-FIFO processing. In fact, the delay bound derived for non-FIFO GR nodes in [15] is an additive bound based on iterative application of per-node delay bounds. This is discussed in more detail in Section V.

The bottom line of this discussion is that we need a new node model: it should

- 1) allow for calculating a bound on the real delay without FIFO processing order assumptions, and, yet, also
- 2) have a concatenation property, and thus be min-plus based, in order to avoid loose additive bounds.

IV. SUFFICIENTLY STRICT SERVICE CURVES

In this section, we introduce a new node model called sufficiently strict service curves (S^3C). This node model allows to efficiently bound the real delay over a tandem of non-FIFO systems using a concatenation result. We show that there are actually systems conforming to this new node model and provide numerical examples for the superiority of S^3C -based delay bounds over conventionally derived delay bounds (based on S^2C).

A. The New Service Curve Definition

Central to the new service curve definition is the notion of a *maximum achievable dwell period* (MADP).

Definition 7: (Maximum Achievable Dwell Period) Given a system \mathcal{S} , the maximum achievable dwell period at time t , denoted as $D(t)$, is the length of the interval $[t_o(t), t]$, i.e., $D(t) = t - t_o(t)$, where $t_o(t)$ is the arrival time of the oldest work unit in the system at time t under *all possible processing orders*. If the system is empty at time t , then by definition $t_o(t) = t$ and $D(t) = 0$.

Some comments on this definition of the MADP are appropriate:

- A tacit assumption is that changing the processing order of work units does not affect the output function F' of the system. This is true under identically sized work units (as assumed in this paper) and a continuous data model with infinitesimally small work units. The definition would require more elaboration if discrete, but variable size work units were considered. We leave this for further study.
- The term “all possible processing orders” in the definition may warrant some further clarification. For min-plus algebraic node models, which only provide guarantees on the amount of work done over specific intervals, possible processing orders amount to all permutations of work units that ever share a buffer.

An example for an MADP should be illustrative: For a work-conserving constant rate server with a single buffer, the MADP at time t equals the backlogged period at time t ; consequently, $t_o(t)$ equals the start of the last backlogged period. The processing order that achieves the MADP for such a simple buffer is LIFO. However, the MADP of a tandem of work-conserving servers can be shorter than the backlogged period of the overall system, which is anyway generally unbounded.

Now using the notion of the MADP, the new service curve definition can be introduced:

Definition 8: (Sufficiently Strict Service Curve - S^3C) Given a system \mathcal{S} with input function F and output function F' , $\beta \in \mathcal{F}$ is a *sufficiently strict service curve* (S^3C) iff

$$\forall t \geq 0 : F'(t) \geq F(t - D(t)) + \beta(D(t)).$$

Note that S^3C implies SC .

B. Properties of S^3C

Now we show that the S^3C definition achieves for non-FIFO systems the two attractive features known from conventional network calculus with FIFO systems: a bound on the delay, yet now on the real instead of the virtual delay, and a concatenation property. These two properties of S^3C allow to recover the celebrated pay bursts only once phenomenon in the non-FIFO case.

1) S^3C Can Bound the Real Delay: First we show that S^3C allows to bound the real delay:

Theorem 4: (Real Delay Bound for an S^3C System) Consider a system \mathcal{S} that offers a sufficiently strict service curve

β . Assume a flow F traversing the system has an arrival curve α . Then we obtain the following bound on the real delay:

$$rd(t) \leq \sup\{s \geq 0 : \alpha(s) \geq \beta(s)\} = i(\alpha, \beta).$$

Proof: As above we denote the MADP at time t by $D(t)$ and the arrival time of the oldest possible work unit (corresponding to $D(t)$) by $t_o(t)$. We can make the following observation for the backlog of the system at time t :

$$\begin{aligned} b(t) &= F(t) - F'(t) \\ &= F(t) - F(t_o(t)) - (F'(t) - F'(t_o(t))) \\ &\leq \alpha(D(t)) - \beta(D(t)). \end{aligned}$$

Here, we used the arrival curve as well as the S^3C property from the assumptions. This relation implies that

$$\alpha(D(t)) \geq \beta(D(t)) + b(t).$$

This now allows to bound the MADP at time t , by the following observation:

$$\begin{aligned} D(t) &\leq \sup\{0 \leq s \leq t : \alpha(s) \geq \beta(s) + b(t)\} \\ &\leq \sup\{0 \leq s \leq t : \alpha(s) \geq \beta(s)\} \\ &\leq \sup\{s \geq 0 : \alpha(s) \geq \beta(s)\} \\ &= i(\alpha, \beta). \end{aligned}$$

Since the bound is independent of t , it applies for all t and, furthermore, since a bound over all MADPs is also a bound for all real delays, the proof is completed. ■

So, like S^2C , S^3C is strong enough to enable a finite bound on the real delay, in contrast to SC and ASG .

2) S^3C Has a Concatenation Property: Now we show that S^3C also has a concatenation property. The concatenation property is not as simple to achieve as for SC and some restrictions apply. To make these restrictions clearly visible and demonstrate where they are required, we develop the concatenation result in a modular fashion, from a set of lemmas, which eventually result in a concatenation theorem for S^3C .

We use the following notation, assuming a tandem \mathcal{T} of n systems $\mathcal{S}_1, \dots, \mathcal{S}_n$:

- $t_o^{(i)}(t)$: the (local) arrival time of the oldest possible work unit at time t for system \mathcal{S}_i ,
- $t_o^{(1, \dots, n)}(t)$: the arrival time (at \mathcal{S}_1) of the oldest possible work unit at time t for the tandem \mathcal{T} ,
- $T_o(t) := t_o^{(1)}\left(t_o^{(2)}\left(\dots t_o^{(n)}(t)\right)\right)$.

The first lemma states a basic concatenation rule for a tandem of S^3C nodes:

Lemma 1: Consider a flow with input function F that traverses the tandem \mathcal{T} . Assume that each system \mathcal{S}_i offers a sufficiently strict service curve β_i , $i = 1, \dots, n$ to the flow. Denoting the output function of \mathcal{T} as F' , we obtain

$$F'(t) \geq F(T_o(t)) + \left(\bigotimes_{i=1}^n \beta_i\right)(t - T_o(t)).$$

The proof of Lemma 1 is relatively straightforward, though notationally somewhat cumbersome and can be found in [16].

Note that Lemma 1 is close to the concatenation property we desire. What is further required is how the arrival times of the oldest possible work units in each of the systems are related with the arrival time of the oldest possible work unit in the overall system. More technically, we need to show that $t_o^{(1,\dots,n)}(t) = T_o(t)$. This is done in the next two lemmas.

Lemma 2: Assuming a tandem \mathcal{T} of n systems $\mathcal{S}_1, \dots, \mathcal{S}_n$, it applies that

$$t_o^{(1,\dots,n)}(t) \geq T_o(t).$$

The proof of Lemma 2 can be done via an induction over the individual systems and can be found in [16].

So far, we made no specific assumptions about the systems \mathcal{S}_i . Yet, for the last step we now need to make a further assumption: We assume each of the systems to have a *common waiting room* for the work units that await service at that system. A common waiting room (c.w.r.) in this context implies that the scheduling of a node is free to choose *any* of the backlogged work units for being processed next. This is, for example, the case if a single buffer is used at a node.

Lemma 3: Assuming a tandem \mathcal{T} of n systems $\mathcal{S}_1, \dots, \mathcal{S}_n$, each with c.w.r., it applies that

$$t_o^{(1,\dots,n)}(t) \geq T_o(t).$$

The proof depends on the c.w.r. assumption as it constructs a sequence of scheduling decisions that result in a work unit that arrived at $T_o(t)$. Again, the proof can be found in [16].

The following theorem sums up what is a consequence of Lemma 1, 2, and 3:

Theorem 5: (Concatenation Theorem for a Tandem of n S^3C Systems) Consider a flow with input function F that traverses a tandem \mathcal{T} of n systems $\mathcal{S}_1, \dots, \mathcal{S}_n$. Assume that each \mathcal{S}_i has a c.w.r. and offers a sufficiently strict service curve β_i to the flow. Then the tandem \mathcal{T} offers a sufficiently strict service curve $\bigotimes_{i=1}^n \beta_i$ to the flow F .

Proof: Simply combine Lemma 2 and 3 to find that $t_o^{(1,\dots,n)}(t) = T_o(t)$. Enter this into Lemma 1's statement and we obtain (denoting the output function after \mathcal{T} as F' again)

$$\begin{aligned} F'(t) &\geq F\left(t_o^{(1,\dots,n)}(t)\right) + \left(\bigotimes_{i=1}^n \beta_i\right)\left(t - t_o^{(1,\dots,n)}(t)\right) \\ &= F\left(t - D^{(1,\dots,n)}(t)\right) + \left(\bigotimes_{i=1}^n \beta_i\right)\left(D^{(1,\dots,n)}(t)\right), \end{aligned}$$

where $D^{(1,\dots,n)}(t)$ denotes the MADP of the tandem \mathcal{T} . ■

We point out that the concatenation needs to be performed in one step because after the concatenation we generally have no c.w.r. for the overall system and thus Lemma 3 is not applicable any more. This is also why we have to prove all these statements in generality for n systems, so we can apply them together in one step. Nevertheless, we now have both properties we strived for: S^3C provides a real delay bound and has a concatenation property under the further assumption of a c.w.r. at the individual systems. The c.w.r. assumptions should not be a strong limitation in practice as the modeling starts with simple components for which this assumption is

usually met by using a single buffer for the backlogged work units at such components. Based on the concatenation theorem a complex system can be constructed. For such a complex system, the c.w.r. assumption will not and does not need to hold any more since the bound on the real delay from Theorem 4 does not depend on that assumption.

3) *Pay Bursts Only Once Holds for Non-FIFO S^3C Tandem:* By combining the results from Theorem 4 and 5, we can, for the case of a tandem of n non-FIFO nodes, each providing an S^3C $\beta_j, j = 1, \dots, n$, which is traversed by an α -constrained flow, derive a bound on the real delay as

$$\forall t \geq 0 : rd(t) \leq i \left(\alpha, \bigotimes_{j=1}^n \beta_j \right).$$

Looking at the same special case as in Section III-B, i.e., $\beta_j = \beta_{R,T}$ and $\alpha = \gamma_{r,b}$, we obtain the following

$$\forall t \geq 0 : rd(t) \leq \frac{b + nrT}{R - r} + nT,$$

which generally improves considerably on the additive bound based on S^2C from Section III-B. We can perceive again the pay bursts only once phenomenon as the burst term appears only once as well as a linear scaling in the number of nodes. We provide some more quantitative observations in Subsection IV-D. Yet, before doing so, we first want to answer an important question: ...

C. Are There Any Systems Providing S^3C ?

Admittedly, the definitions of the MADP and consequently S^3C are somewhat peculiar, raising the question whether there are any systems actually satisfying S^3C . In fact, it is hard to *directly* verify that a system provides S^3C . Instead, we show that S^2C nodes with c.w.r. provide S^3C .

Proposition 1: A node with c.w.r. providing an S^2C β also provides β as S^3C .

Proof: First, we show that the MADP is identical to the length of the backlogged period for S^2C with c.w.r. Let us define the backlogged period at time t as $B(t) = \sup_{s \leq t} \{t - s : \forall u \in (s, t] : b(u) > 0\}$. That $D(t) \leq B(t)$ is clear from the renewal character of an empty system just before $t - B(t)$, so there can be no more work units in the system with an arrival time earlier than $t - B(t)$. Further, under the assumption of c.w.r., it is clear that LIFO is a possible processing order and results in an arrival time for the oldest work unit of $t - B(t)$. Thus $D(t) \geq B(t)$, and altogether $D(t) = B(t)$. Now as the S^2C guarantee applies to any backlogged periods, it also applies for $(t - B(t), t]$. Further using that $t - B(t)$ is the start of the last backlogged period before t , i.e. $F(t - B(t)) = F'(t - B(t))$, we obtain

$$F'(t) \geq F'(t - B(t)) + \beta(B(t)) = F(t - D(t)) + \beta(D(t)),$$

which constitutes β as S^3C . ■

So, the question whether there are any systems providing S^3C can be answered positively, because many schedulers are known to provide S^2C and thus also provide S^3C if c.w.r. is satisfied.




Fig. 1. Possible maximum utilizations for a target delay bound under S^2C and S^3C .

D. Numerical Examples

To give some feeling for the improvements achievable by using the S^3C -based end-to-end delay analysis compared to a node-by-node bounding based on S^2C we provide a numerical example when applying the results for admission control purposes. In addition, we demonstrate what cost is incurred for releasing the FIFO assumption. For these numerical experiments we use simple settings: As arrival curve for the flow to be analyzed we assume a token bucket $\gamma_{r,b}$, where we set $b = 5[Mb]$ and vary the rate r to achieve a certain utilization; for the service curves of the nodes to be traversed we use a rate-latency function $\beta_{R,T}$ with $R = 20[Mbps]$ and $T = 0.01[s]$. We assume $n = 10$ nodes to be traversed by the flow under analysis.

1) *Worst-Case Admission Control:* A comparison between S^3C and S^2C -based bounding methods when used in admission control is provided in Figure 1. Here, the acceptable utilizations for a given delay bound are shown for both methods. This information can be used to decide how much traffic can be admitted. As can be clearly seen, the S^3C -based method outperforms the S^2C -based method by far, especially for lower delay bounds.

2) *FIFO vs. Non-FIFO Delay Bounds:* In another numerical example, we investigate the cost of releasing the FIFO assumption in terms of delay bounds. For that purpose, we vary the utilization by increasing the sustained rate of the flow under analysis (while at the same time scaling the bucket depth accordingly). As can be observed from Figure 2, only for higher utilizations there is a significant difference between the FIFO and non-FIFO delay bounds (at least for the S^3C -based bound). The bottom line is that only for highly utilized systems it is necessary to enforce a FIFO behavior, as far as delay bounds are concerned. For systems with lower utilizations, optimizations such as for example link aggregation or multi-stage switching fabrics do not incur a high cost in terms of worst-case delay bounds.

E. Simulations

We have built a queueing model for discrete-event simulations using OMNeT++, in order to test some actual system




Fig. 2. FIFO vs. non-FIFO delay bounds depending on the utilization of the system.

behavior against our bounds. To that end, we implemented several queue management strategies: FIFO, LIFO, Shortest-In-System (SIS), and windowed random (RND(w)). SIS always selects the packet with the latest arrival time to the overall system. RND(w) chooses randomly, but restricted to the first w packets from a FIFO buffer, mimicking typical optimizations for core Internet routers with high-speed switching fabrics to avoid head-of-line blocking.

1) *Experimental Design:* We simulate a tandem of n servers, each equipped with a single buffer; each server is TDMA-like with On and Off periods of duration $t_{on}^{(i)}$ and $t_{off}^{(i)}$, respectively; during an On period each server sends at rate $C^{(i)}$; each server applies one of the queue management strategies mentioned above to select the next packet eligible for service. For the analytical derivation of the bounds, we can abstract server i as S^3C with $\beta_{C^{(i)} / (1 + t_{off}^{(i)} / t_{on}^{(i)}), t_{off}^{(i)}}$.

We use a Markov-Modulated On-Off source with rates λ for the On-Off and μ for the Off-On transition, as well as a data rate P while in state On. Furthermore, the source is regulated by a token bucket $\gamma_{r,b}$ before entering the network.

The factor values corresponding to the simulation results in the next subsection can be found in Table I. The primary factor in these experiments is the utilization, which is varied by using different token bucket rates. For all factor combinations and all queue management strategies 100 simulation runs were performed; the results below report the average of the observed worst-case delays over these 100 replications.

2) *Results:* In Figure 3, the observed worst-case delays for all queue management strategies as well as the analytical worst-case delay bounds under FIFO and non-FIFO processing are displayed (for a utilization of 50%), from left to right in ascending order. We can observe that the non-FIFO bound is never violated. In contrast, the FIFO bound is actually violated by all queue management strategies (apart from FIFO, of course, which is slightly smaller); in particular, it is already violated by RND(2). We also see that SIS is closest to the non-FIFO bound and performs worse than LIFO. Clearly, the RND(w) delays deteriorate with increasing window size, and for $w = 8$ we are roughly at two third of the non-FIFO bound.

Factor	Value
n	10
$C^{(i)}$	$200 \frac{p}{s}$
$t_{on}^{(i)}$	$0.02 s$
$t_{off}^{(i)}$	$i = n : 0.02 s, \text{ else } t_{off}^{(i)} = t_{off}^{(i+1)} - 0.002 s$
λ	$20 \frac{1}{s}$
μ	$20/3 \frac{1}{s}$
P	$100 \frac{p}{s}$
r	$10, 20, \dots, 90 \frac{p}{s}$
b	$20 p$

TABLE I
FACTOR VALUES FOR THE SIMULATION EXPERIMENTS.




Fig. 3. Worst-case delays for different queue management strategies.

The effect of varying the utilization of the system is shown in Figure 4. Here, we only show the results for SIS and RND(2) as well as the analytical bounds. As we can observe SIS exhibits the same growth behavior as the non-FIFO bound and stays close up to high utilizations, though the gap tends to become larger. RND(2) is pretty insensitive to the utilization, but violates the FIFO bound for all utilizations.

V. A NOTE ON “PAY BURSTS ONLY ONCE DOES NOT HOLD FOR NON-FIFO GUARANTEED RATE NODES” [15]

In Section IV, we established that the “pay bursts only once” phenomenon also holds for non-FIFO systems. This seems to be in sharp contrast to a literature result by Rizzo and Le Boudec [15]. They claim (in their article’s title) that “pay bursts only once does not hold for non-FIFO guaranteed rate nodes”. While they use a different server model, guaranteed rate (*GR*) nodes, some schedulers can be abstracted as both, *GR* or *S³C*, and thus it is important to resolve this contradiction, as for such nodes only one statement can be true.

A. The Result from [15]

Some preliminaries are necessary: A *guaranteed rate (GR) node* is a server that guarantees an upper bound on the departure time of any given packet of a flow by limiting the deviation of its departure time from a hypothetical *FIFO* constant rate server. To that end, it uses the concept of a *guaranteed rate clock (GRC)* for any given packet p^j of length l_j , where p^j is the j^{th} packet arriving at the node, $A(p^j)$ its arrival time, and R the server’s allocated rate for the flow under analysis:

$$GRC(p^j) = \begin{cases} 0, & j = 0 \\ \frac{l_j}{R} + \max\{A(p^j), GRC(p^{j-1})\}, & j > 0 \end{cases}$$

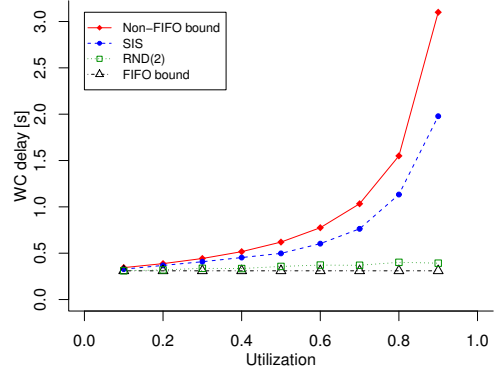


Fig. 4. The effect of utilization on worst-case delays.

A *GR* node guarantees that any packet p^j is delivered by the time $GRC(p^j) + e$, where the error term e depends on the actual packet scheduling algorithm.

Note that a *GR* node does not need to be *FIFO*, though its definition restricts its non-*FIFO* behavior. Possible reorderings of packets are restricted to packets leaving *early* with respect to their deadline. For this reason, a delay bound can be derived for a *GR* node even if it is not assumed to be *FIFO*. However, in order to concatenate several *GR* nodes one draws upon their equivalence with *SC* in the *FIFO* case [13]. Therefore, in [15], an additive delay bound based on the per-node delay bound for non-*FIFO GR* nodes is derived. More specifically, under a token-bucket arrival curve $\gamma_{r,b}$ and assuming a propagation delay τ^i between nodes i and $i+1$ as well as a maximum packet size l_{max} , the bound in [15] for a tandem of n *GR* nodes with equal service rate R is given as

$$d^{[15]} = n \frac{b}{R} + \frac{r l_{max} n (n-1)}{2R^2} + \frac{r}{R} \sum_{k=1}^n \sum_{i=1}^{k-1} e^i + \sum_{i=1}^n (e^i + \tau^i).$$

As can be seen, the burst is paid n times and a quadratic scaling in n is incurred, as it is typical for additive bounding methods. [15] claims that the bound is *tight* in general and thus concludes that “pay bursts only once does not hold for non-*FIFO GR* nodes”.

B. The Bound From [15] Is Not Tight

A counter-example on the tightness of the bound in [15] can be constructed by assuming constant rate servers and choosing the following parameters: $n = 2$, $R = 1$, $e^i = 0$, $r = 0$, $b = 2$, $\tau^i = 0$, and all packets being of unit length. Then $d^{[15]} = 4$. However, it can be seen that $d^{tight} = 3$ is a valid delay bound for that scenario: In the worst case, the two packets of the flow arrive at the same time and one is served within one time unit and immediately transferred to the second node; the second packet is served in another time unit while the first is served by the second node; now the second node is free to serve the second packet in another time unit, amounting for the second packet in 3 time units delay altogether. Changing the order of the packets cannot matter here, as they are arriving at the same




Fig. 5. Comparison of delay bounds from S^3C and [15]. The parameter settings are: $R = 20$, $T = e^i = 0.01$, $\tau^i = 0$, $r = \rho R$, $b = \frac{r}{5}$, unit packets for a tandem of $n = 10$ servers (i.e., $l_{max} = 1$), with ρ being the utilization.

time anyway. Any spreading of the two packets can only result in a lower delay. This example can be generalized to a series of $n \geq 2$ nodes, where the tight delay bound $d^{tight} = n + 1$, but $d^{[15]} = n \frac{b}{R} = 2n$. This shows that the delay bound from [15] is not tight in the general case. In their tightness proof, Rizzo and Le Boudec make the following assumption: “In order to simplify the notation, we assume that $r = R \dots$ ” Under this assumption their delay bound is actually tight, yet for any scenario with $r < R$, it is not. A more detailed discussion about this issue is provided in [16].

In contrast, using S^3C allows to calculate the tight delay bound: We use $\beta^{stair}(t) := \lfloor t \rfloor$ as S^3C ; for the concatenation of n nodes we have $\bigotimes_{i=1}^n \beta^{stair} = \beta^{stair} \otimes \delta_{n-1}$; thus we obtain $d^{S^3C} = i(\gamma_{0,2}, \beta^{stair} \otimes \delta_{n-1}) = n + 1$.

C. Comparison with S^3C -based Bounds

In principle, comparing the bound from [15] with an S^3C -based one is like a comparison between apples and oranges. On the one hand, a GR node is less stringent with respect to giving a guarantee about the work being done in a specific time interval, on the other hand, a GR node allows for much less reordering between work units. Nevertheless, we provide a comparison between the two bounding methods for the following sample scenario: Assume we are given a tandem of PGPS servers (each with a single buffer) guaranteeing a rate R with latency $T = e_i$; such nodes can be abstracted as GR as well as S^3C nodes (because they are S^2C with c.w.r.). Then Figure 5 provides a numerical example comparing the respective delay bounds. Up to a very high utilization, S^3C yields less conservative bounds than [15]. Only at a utilization of 95% the increased potential for reordering due to large backlogs outweighs the benefits of the concatenation-based S^3C bound. Whether a concatenation-based bound can be found under the non-FIFO GR model remains open.

VI. CONCLUSION

It was our goal to extend the scope of network calculus towards non-FIFO systems, as it is theoretically interesting, methodologically consistent, and useful in real-world applications. It turned out that existing service curve definitions are

not satisfying under non-FIFO processing of a flow: they are either too loose to enable any bounding or too strict to allow for an efficient end-to-end analysis. Therefore, we introduced a new service curve definition, S^3C , which allows to bound the delay and at the same time enables an efficient end-to-end analysis. By some numerical examples, we showed that the new analysis based on S^3C is clearly superior to existing methods; by simulations, we validated the bounds against actual system behavior. Remarkably, S^3C allows to recover the pay bursts only once phenomenon for non-FIFO systems.

ACKNOWLEDGEMENTS: We are very grateful to Markus Fidler for insightful comments on an early version of this paper. This work is supported by an Individual grant from the German Research Foundation (DFG).

REFERENCES

- [1] Rajeev Agrawal, R. L. Cruz, Clayton M. Okino, and Rajendran Rajan. A framework for adaptive service guarantees. In *Proc. of Allerton Conf.*, pages 693–702, 1998.
- [2] J. C. R. Bennett, C. Partridge, and N. Shectman. Packet reordering is not pathological network behavior. *IEEE/ACM Trans. Netw.*, 7(6):789–798, 1999.
- [3] J. C. R. Bennett and H. Zhang. WF²Q: Worst-case fair weighted fair queueing. In *Proc. IEEE INFOCOM*, pages 120–128, March 1996.
- [4] A. Bouillard, L. Jouhet, and E. Thierry. Tight performance bounds in the worst-case analysis of feed-forward networks. In *Proc. IEEE INFOCOM*, pages 1–9, March 2010.
- [5] S. Chakraborty, S. Kuenzli, L. Thiele, A. Herkersdorf, and P. Sagmeister. Performance evaluation of network processor architectures: Combining simulation with analytical estimation. *Computer Networks*, 42(5):641–665, 2003.
- [6] C.-S. Chang. *Performance Guarantees in Communication Networks*. Telecommunication Networks and Computer Systems. Springer-Verlag, 2000.
- [7] R. L. Cruz. A calculus for network delay, Part I: Network elements in isolation. *IEEE Trans. Inf. Th.*, 37(1):114–131, January 1991.
- [8] P. Goyal, S. S. Lam, and H. M. Vin. Determining end-to-end delay bounds in heterogeneous networks. *Multimedia Syst.*, 5(3):157–163, 1997.
- [9] C. Guo. SRR: An O(1) time complexity packet scheduler for flows in multi-service packet networks. *IEEE/ACM Trans. Netw.*, 12(6):1144–1155, December 2004.
- [10] Petr Jurcik, Anis Koubãa, Ricardo Severino, Mário Alves, and Eduardo Tovar. Dimensioning and worst-case analysis of cluster-tree sensor networks. *ACM Trans. Sen. Netw.*, 7:14:1–14:47, September 2010.
- [11] H. Kim and J.C. Hou. Network calculus based simulation: theorems, implementation, and evaluation. In *Proc. IEEE INFOCOM*, March 2004.
- [12] J.-Y. Le Boudec and A. Charny. Packet scale rate guarantee for non-fifo nodes. In *Proc. IEEE INFOCOM*, pages 23–26, June 2002.
- [13] J.-Y. Le Boudec and P. Thiran. *Network Calculus A Theory of Deterministic Queuing Systems for the Internet*. Number 2050 in Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 2001.
- [14] A. K. Parekh and R. G. Gallager. A generalized processor sharing approach to flow control in integrated services networks: The single-node case. *IEEE/ACM Trans. Netw.*, 1(3):344–357, June 1993.
- [15] G. Rizzo and J.-Y. Le Boudec. Pay bursts only once does not hold for non-fifo guaranteed rate nodes. *Performance Evaluation*, 62(1-4):366–381, 2005.
- [16] J. Schmitt, N. Gollan, S. Bondorf, and I. Martinovic. Pay bursts only once holds for (some) non-fifo systems. Technical Report 378/10, University of Kaiserslautern, Germany, July 2010. <http://disco.informatik.uni-kl.de/publications/SGBM10-1.pdf>.
- [17] J. Schmitt, F. Zdarsky, and M. Fidler. Delay bounds under arbitrary aggregate multiplexing: When network calculus leaves you in the lurch... In *Proc. IEEE INFOCOM*, April 2008.
- [18] T. Skeie, S. Johannessen, and O. Holmeide. Timeliness of real-time IP communication in switched industrial ethernet networks. *IEEE Trans. Ind. Inf.*, 2(1):25–39, February 2006.