

Window Flow Control in Stochastic Network Calculus – The General Service Case

Michael Beck
University of Kaiserslautern
Distributed Computer Systems Lab (DISCO)
Germany
beck@cs.uni-kl.de

Jens Schmitt
University of Kaiserslautern
Distributed Computer Systems Lab (DISCO)
Germany
jschmitt@cs.uni-kl.de

ABSTRACT

The problem of Window Flow Controlled (WFC) systems has been analysed successfully in Deterministic Network Calculus (DNC). While many results of DNC have been carried over to its stochastic extension, the problem of WFC systems has not been solved so far. This paper presents the first approach to analyse WFC in the context of Stochastic Network Calculus (SNC) for a general service inside the feedback loop of the controller. The key idea is to keep track of how much the service deviates from being subadditive. The new method is illustrated in numerical examples and its properties are discussed.

Categories and Subject Descriptors

C.4 [Computer Systems Organization]: Performance of Systems—*Modeling techniques*

General Terms

Theory

Keywords

Stochastic Network Calculus, Feedback, Flow Control

1. INTRODUCTION

Stochastic Network Calculus (SNC) has matured in recent years to provide an alternative method for performance analysis of stochastic queueing systems (see e.g. [15, 13, 9]). Many results from the deterministic network calculus (DNC) have been transferred into the stochastic domain, some have been rather immediate some have required considerable effort (e.g., deriving the end-to-end service [7]). One major remaining open issue is the stochastic analysis of feedback-based systems, such as Window Flow Controlled (WFC) transport protocols, e.g. TCP. While there are very elegant solutions for WFC in the deterministic setting [1, 5, 16], WFC in SNC has been identified for some time already as a very challenging open research question [15, 12, 13, 8]. Moreover, being able to analyse WFC systems in SNC would

be very relevant to open up new application areas for SNC such as modelling smart grid systems [14], for instance.

In this paper, we present an approach to analyse WFC in SNC under very general assumptions. As we demonstrate below, as long as the service in the feedback loop is subadditive the analysis is a rather direct transfer from DNC (see [3] for a detailed discussion, where especially the continuous time case needs care, though). Yet, the assumption of a subadditive service inside the feedback loop is restrictive and there are several scenarios where subadditivity cannot be assumed; most prominently, this is the case if a tandem of servers has to be traversed. This is why we tackle the hard case of general service in the feedback loop in this work.

The key idea of our approach is to stochastically control how far the service deviates from being subadditive and cast this into the setting of MGF-calculus [6, 11], a branch of SNC, as further contribution to the violation probabilities of the performance bounds. We demonstrate the method by providing some numerical results in case of a server in the feedback loop that is not subadditive and would thus not be analysable by direct methods (compare [3]).

2. NOTATION AND SOME BASIC RESULTS

First, we provide the notation and some basic results of (deterministic) network calculus. For further details one can also refer to [4, 6]. For ease of presentation we assume a discrete time model throughout this work, see [3] for results on continuous time.

We start by defining (arrival) flows, which represent a fluid stream of data. These flows enter and depart service elements.

DEFINITION 1. *We denote an arrival flow by its cumulatives A ; i.e., $A(t)$ units of data arrive in the interval $[0, t]$. The bivariate extension of A is defined by:*

$$A(s, t) := A(t) - A(s).$$

Note that flows are always additive, i.e.

$$A(s, t) = A(s, r) + A(r, t)$$

for all r, s, t .

We introduce two service descriptions here, one for the uni-

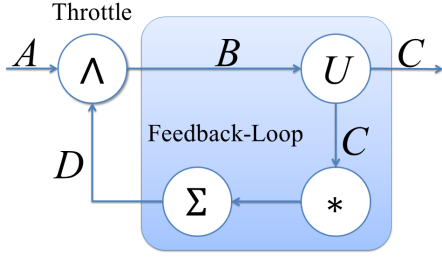


Figure 1: A window flow controller: the input A is throttled at the \wedge -element. The departures of the system are C and the feedback-loop consists of the dynamic U -server, an unspecified service element and a window-element. The star is a placeholder for zero, one, or more elements, like delay elements, scalars or dynamic servers.

variate calculus (usually used in DNC and the tail-bound-branch of SNC) and the bivariate calculus (used in the MGF-branch of SNC):

DEFINITION 2. We say a service element offers a service curve U , if for any input-output pair A, B and time t

$$B(t) \geq A \otimes U(t) := \min_{0 \leq s \leq t} \{A(s) + U(t-s)\}.$$

Let U be a bivariate function with $U(s, t) \leq U(s, t')$ for all $t \leq t'$. We say a service element is a dynamic U -server, if for any input-output pair A, B and time t :

$$B(t) \geq A \otimes U(0, t) := \min_{0 \leq s \leq t} \{A(0, s) + U(s, t)\}.$$

The operators \otimes are called (univariate or bivariate) min-plus convolution, as they resemble the ordinary convolution operator under standard algebra. Note that the bivariate U is in general not additive, i.e., there may exist r, s, t with $U(s, t) \neq U(s, r) + U(r, t)$. The service U can, however, be *subadditive*, i.e., for all r, s, t :

$$U(t) \leq U(s) + U(t-s) \quad (1)$$

$$U(s, t) \leq U(s, r) + U(r, t) \quad (2)$$

in the univariate and bivariate case, respectively.

The following is crucial for the end-to-end analysis of networks consisting of several service elements.

THEOREM 3. Consider two service elements, such that the output of the first service element is the input to the second. If both service elements have a service curve U_i (are a dynamic U_i -server, $i = 1, 2$), the system also has a service curve, given by $U_1 \otimes U_2$ (is a dynamic $U_1 \otimes U_2$ -server).

The second network-operation we give here is central to this work. It describes how a WFC system as presented in Figure 1 is handled. In this feedback system, the original input A is fed to a throttle-element first, which governs how much data is admitted to the inside of the system (the feedback-loop).

This is realized by taking the minimum of the input A and the output of the service element at any time, such that:

$$B(t) = A(t) \wedge D(t).$$

Such systems are studied for example in [1, 5, 16, 6]. We formulate from there the following theorem:

THEOREM 4. Assume the whole feedback-loop in Figure 1 is described by a service curve U_{fb} (is a dynamic U_{fb} -server). The throttle element \wedge has a service curve U_\wedge (is a dynamic U_\wedge -server), with:

$$U_\wedge(t) = \bigwedge_{k=0}^{\infty} U_{fb}^{(k)}(t), \quad \left(U_\wedge(s, t) = \bigwedge_{k=0}^{\infty} U_{fb}^{(k)}(s, t) \right).$$

Here the notation $U_{fb}^{(k)}$ stands for the k -fold self-convolution of U_{fb} . Further, for any U we define $U^{(0)}$ to be the neutral element of the convolution: $U^{(0)}(t) = \mathbf{1}(t) = \infty$ for all $t > 0$ and $\mathbf{1}(0) = 0$ (for the bivariate case we have $\mathbf{1}(s, t) = \infty$ for all $s < t$ and $\mathbf{1}(t, t) = 0$ for all t , respectively). The expression $\bigwedge_{k=0}^{\infty} U^{(k)} =: \bar{U}$ is known as the *subadditive closure* of service U . As the name suggests the subadditive closure of a service description is subadditive. Further, we have for any subadditive U :

$$\bigwedge_{k=0}^{\infty} U^{(k)} = \mathbf{1} \wedge U.$$

To complete the analysis of the feedback system in Figure 1 we apply Theorem 4 and Theorem 3 and get a service description for the whole system by:

$$U_{sys} = U_\wedge \otimes U.$$

This service description can then be used to derive performance bounds, like end-to-end-delay, as detailed later.

In network calculus one typically assumes systems to be empty at time zero, i.e., for all participating flows holds $A(0) = 0$. The window element in the feedback-loop serves two purposes: 1) it kick-starts our system, by initially admitting a certain amount of data to U and 2) it controls how much data is present inside the feedback-loop, thus guaranteeing a maximum backlog, which may occur on any element inside of it.

DEFINITION 5. A window element Σ is a function from \mathbb{N}_0 to \mathbb{R}_0^+ , such that for any input-flow A it produces an output B

$$B(t) = A(t) + \Sigma(t) \quad (3)$$

$$\Sigma(s, t) \geq -A(s, t) \quad \forall s \leq t. \quad (4)$$

We further define the bivariates $\Sigma(s, t)$ as for flows.

The window elements relays its input A without delay and adds the window size Σ to it. The constraint (4) enforces, that the output $B(s, t)$ of a window element cannot be negative (the window element can negate the whole input, though, by reducing its window size). Note that a window-element

is not a flow: we do not have the property that $\Sigma(t) \leq \Sigma(t')$ for all $t \leq t'$.

One can rewrite the properties of a window element in the form of service curves (if Σ is constant) or dynamic servers (for time-varying window sizes):

$$B(t) = A \otimes W_\Sigma(0, t)$$

with

$$W_\Sigma(t) = \begin{cases} \Sigma & \text{if } t = 0 \\ \infty & \text{else} \end{cases}$$

or

$$W_\Sigma(s, t) = \begin{cases} \Sigma(t) & \text{if } s = t \\ \infty & \text{if } s < t \end{cases}$$

Note, however, as Σ is missing monotonicity, it is not a service curve (dynamic server) in the sense of Definition 2.

Using the above calculus one can derive probabilistic performance bounds. To this end we define two performance measures of a system with input-output pair A, B :

DEFINITION 6. *The backlog q of the system at time t is defined by:*

$$q(t) := A(t) - B(t).$$

The virtual delay d at time t is defined by:

$$d(t) := \min\{T : A(t) \leq B(t+T)\}.$$

So far we have expressed the theory of network calculus in terms of its flows and service processes. To calculate stochastic bounds on q and d we need to add a probabilistic component on top of that. There are different methods [13] to describe the stochastic nature of flows and service descriptions each leading to their own slightly different performance guarantees. The one we use in this paper is provided in the following theorem [6, 11]:

THEOREM 7. *Let time be discrete and let the arrival flow A and the service process U be stochastically independent. Let $\theta > 0$, such that there exist bounds:*

$$\mathbb{E}(e^{\theta A(s,t)}) \leq e^{\theta \rho_A(\theta)(t-s) + \theta \sigma_A(\theta)} \quad (5)$$

$$\mathbb{E}(e^{-\theta U(s,t)}) \leq e^{\theta \rho_U(\theta)(t-s) + \theta \sigma_U(\theta)} \quad (6)$$

for all $s \leq t$ and $\rho_A(\theta) + \rho_U(\theta) =: \rho(\theta) < 0$. Then:

$$\mathbb{P}(q(t) > x) \leq e^{\theta(\sigma_A(\theta) + \sigma_U(\theta) - x)} (1 - e^{\theta \rho(\theta)})^{-1} \quad \forall t \geq 0$$

$$\mathbb{P}(d(t) > T) \leq e^{\theta(\sigma_A(\theta) + \sigma_U(\theta) + \rho_U(\theta)T)} (1 - e^{\theta \rho(\theta)})^{-1} \quad \forall t \geq 0.$$

The condition $\rho(\theta) < 0$ is a stability condition for the system. There are versions of the above theorem for continuous time, as well as for the case in which A and U are not stochastically independent [2]. The above bounds use Moment-Generating Functions (MGFs) to describe the behaviour of A and U , which is why it is called MGF-calculus for short.

We point out here that none of the state-of-the-art methods is able to describe the service process resulting from Theorem 4 efficiently. This is why the solution of WFC systems eluded stochastic network calculus, so far.

3. PROBLEM EXPOSITION

Subadditivity plays an important role when analysing WFC systems, as the throttle's service is just the subadditive closure of the elements inside the feedback loop. Hence, feedback-loops containing subadditive elements are much easier to analyse. We present this situation now in detail.

We look at a feedback loop which contains a subadditive service element U (lying on the path from the sender to the receiver) and a window-element (lying on the return path). Such a system can be easily analysed, by applying the most general result concerning WFC systems as found in the textbook of Chang [6]. To that end, denote the space of bivariate functions which are monotonically increasing in their second variable as $\tilde{\mathcal{F}}$. We see immediately that all flows and all dynamic-server descriptions lie in $\tilde{\mathcal{F}}$, as they fulfill $F(s, t) \leq F(s, t')$ for all $t \leq t'$.

DEFINITION 8. *An operator $\pi : \tilde{\mathcal{F}} \rightarrow \tilde{\mathcal{F}}$ is called σ -additive if:*

$$\pi \left(\bigwedge_{n=1}^{\infty} F_n \right) = \bigwedge_{n=1}^{\infty} \pi(F_n), \quad (7)$$

where F_n is any sequence in $\tilde{\mathcal{F}}$ and the infima are understood pointwisely.

One can easily verify that the space of σ -additive operators is closed under taking countable minima and compositions. Further they distribute over countable minima:

$$\left(\bigwedge_{n=1}^{\infty} \pi_n \right) \circ \pi = \bigwedge_{n=1}^{\infty} (\pi_n \circ \pi), \quad \pi \circ \left(\bigwedge_{n=1}^{\infty} \pi_n \right) = \bigwedge_{n=1}^{\infty} (\pi \circ \pi_n)$$

An example of a σ -additive operator is the convolution with some bivariate $U \in \tilde{\mathcal{F}}$, defined by $\pi_U(A) := A \otimes U$. Other examples needed for our specific feedback system are

$$\begin{aligned} \pi_{\bar{e}}(A) &:= A \\ \pi_{+w}(A)(s, t) &:= \begin{cases} A(s, t) & \text{if } s \neq 0 \\ A(0, t) + w & \text{if } s = 0 \end{cases} \end{aligned}$$

where $w \in \mathbb{R}$. The above operators represent the identity operator and a window-element, respectively; their σ -additivity as well as the following properties are easy to verify:

- π_{+w} commutes with π_U , i.e. $\pi_U \circ \pi_{+w} = \pi_{+w} \circ \pi_U$.
- $\pi_{+w}^2 = \pi_{+2w}$, and U is subadditive with $U(t, t) = 0$ for all t iff π_U is idempotent, i.e. $\pi_U^2 = \pi_U$.

The whole feedback loop can be expressed by successively applying σ -additive operators and it is a σ -additive operator

itself:

$$\pi_{fb} := \pi_{+w} \circ \pi_U,$$

i.e. $D = \pi_{fb}(B)$. The relations between the flows in Figure 1 are hence given by:

$$B(t) = A(t) \wedge \pi_{fb}(B)(0, t). \quad (8)$$

Lemma 5.7.2. in [6] can be applied resulting in $B \geq \pi_{fb}^*(A)$ for any B fulfilling $B \geq A \wedge \pi_{fb}(B)$, where $\pi_{fb}^* = \pi_{\bar{\varepsilon}} \wedge \bigwedge_{n=1}^{\infty} \pi_{fb}^n$ is the closure of π_{fb} .

We show now that π_{fb}^* is tractable when considering a sub-additive U together with a window-element.

EXAMPLE 9. *Applying the above properties to π_{fb} , we observe that*

$$\pi_{fb}^n \geq \pi_{+nw} \circ \pi_U.$$

And for the closure of π_{fb}

$$\pi_{fb}^* \geq \pi_{\bar{\varepsilon}} \wedge \bigwedge_{n=1}^{\infty} \pi_{+nw} \circ \pi_U = \pi_{\bar{\varepsilon}} \wedge \pi_U.$$

Applying the result of Chang we have for the departures C of the feedback system:

$$\begin{aligned} C(t) &\geq \pi_U(B)(0, t) \geq \pi_U \circ \pi_{fb}^*(A)(0, t) \\ &= \bigwedge_{n=0}^{\infty} \pi_U \circ \pi_{fb}^n(A)(0, t) \\ &\geq \bigwedge_{n=0}^{\infty} \pi_{+nw} \circ \pi_U(A)(0, t) = \pi_U(A)(0, t) \end{aligned}$$

We use the distributivity of σ -additive operators in the second line. We see that the whole system behaves, just like the unthrottled one. With given MGF-bounds on A and U we can use this bound on C to produce stochastic performance bounds, as in Theorem 7.

The above example makes the role of subadditivity clear: without U being subadditive, we would not have $\pi_U^2 = \pi_U$ and the description of π_{fb}^* would include π_U^n . The repeated application of π_U prevents deriving stochastic performance bounds as above. Note further, that the window-element causes only minor difficulties, as its operator commutes with π_U .

The most important example for a non-subadditive U is a service, which results from applying Theorem 3. How one can preserve subadditivity instead for a concatenation of service elements is discussed in [3]. However, preserving subadditivity in that way comes at the cost of a decreased service.

We now leave the notations of σ -additive operators, and take a step back to have a closer look at subadditivity in the most simple scenario. For this discussion we retreat to the univariate case and let the placeholder in Figure (1) be empty. Further, we assume a window-element of fixed size, such that $U_{fb} = U + \Sigma$.

As the subadditive closure of the feedback-loop appears in the solution of Theorem 4, subadditivity of the service curve

U and how much is ‘‘missing’’ to it are crucial questions in WFC systems. Assume we can give the following bound, describing by what quantity a service curve U fails subadditivity:

$$U(t) - U \otimes U(t) \leq b \quad \forall t \in \mathbb{R}_0^+. \quad (9)$$

Or, written differently:

$$U^{(2)}(t) \geq U(t) - b \quad \forall t \in \mathbb{R}_0^+.$$

If we consider some U and b to fulfill Equation (9), we achieve for $\Sigma \geq b$ the subadditivity of $U + \Sigma = U_{fb}$, since

$$\begin{aligned} (U + \Sigma) \otimes (U + \Sigma)(t) &= U^{(2)}(t) + 2\Sigma \\ &\geq U(t) - b + 2\Sigma \\ &\geq U(t) + \Sigma = U_{fb}(t). \end{aligned}$$

A subadditive U_{fb} , however means that¹

$$U_{sys}(t) = \bar{U}_{fb} \otimes U(t) = (\mathbf{1} \wedge U + \Sigma) \otimes U(t) = U(t).$$

This effect, that a system with a large enough window can be analysed as unthrottled, is well known and can for example be found in [1]. One can think of the window-element to compensate for how the service deviates from subadditivity.

4. WFC WITH GENERAL SERVICE

We revisit the solution of Theorem 4. Without further assumptions, we need to find an MGF-bound in the sense of (6) for the end-to-end service:

$$U_{sys} = \left(\bigwedge_{n=0}^{\infty} (U_{fb})^{(n)} \right) \otimes U(s, t),$$

where U_{fb} represents the whole feedback-loop. Just applying the definition of the MGF we obtain

$$\begin{aligned} \Phi_{U_{sys}(s,t)}(-\theta) &= \mathbb{E}(e^{-\theta \wedge (U_{fb})^{(n)} \otimes U(s,t)}) = \mathbb{E}\left(\bigvee_{n=0}^{\infty} e^{-\theta U_{fb}^{(n)} \otimes U(s,t)}\right). \end{aligned}$$

A naive approach would be to use $\mathbb{E}(X \vee Y) \leq \mathbb{E}(X) + \mathbb{E}(Y)$ for some positive random variables X and Y , resulting in:

$$\Phi_{U_{sys}(s,t)}(-\theta) \leq \sum_{n=0}^{\infty} \mathbb{E}(e^{-\theta U_{fb}^{(n)} \otimes U(s,t)})$$

This however is problematic, since finite representations, as in (6) for $U_{fb}^{(n)}$ with $n \rightarrow \infty$, are hard to achieve, even if U_{fb} consists of a single server U and the window element Σ only.

Remembering condition (9) we choose another path here. To fix notations, insert for the placeholder in Figure (1) a dynamic V -server, such that $U_{fb}(s, t) = U \otimes V(s, t) + \Sigma(t)$. We saw in the univariate case for $b < \Sigma$ an easy solution to the feedback inequality. The same holds for the bivariate setting with a dynamic window; assume it holds for all $s \leq t$:

$$U \otimes V(s, t) - (U \otimes V)^{(2)}(s, t) \leq b \leq \Sigma_{min}(t) \quad (10)$$

¹As they are σ -additive operators the convolution with a function is distributive over minima, i.e.: $(U_1 \wedge U_2) \otimes V = U_1 \otimes V \wedge U_2 \otimes V$

with $\Sigma_{min}(t) := \min_{s \leq t} \{\Sigma(s)\}$. Then we have:

$$\begin{aligned} & U_{fb} \otimes U_{fb}(s, t) \\ &= \min_{s \leq r \leq t} \{U \otimes V(s, r) + \Sigma(r) + U \otimes V(r, t) + \Sigma(t)\} \\ &\geq \min_{s \leq r \leq t} \{U \otimes V(s, r) + U \otimes V(r, t)\} + \Sigma_{min}(t) + \Sigma(t) \\ &\geq U \otimes V(s, t) + (\Sigma_{min}(t) - b) + \Sigma(t) \\ &\geq U_{fb}(s, t) \end{aligned}$$

and U_{fb} is subadditive.

We use this property to achieve a service description for the whole system by:

$$\begin{aligned} & U_{sys}(s, t) \\ &= \overline{U_{fb}} \otimes U(s, t) = (\mathbf{1} \wedge U_{fb}) \otimes U(s, t) \end{aligned} \quad (11)$$

$$\begin{aligned} &= U(s, t) \wedge \min_{s \leq r \leq t} \{(U \otimes V)(s, r) + \Sigma(r) + U(r, t)\} \\ &\geq U(s, t) \wedge (U \otimes V) \otimes (U \otimes V)(s, t) + \Sigma_{min}(t) \end{aligned} \quad (12)$$

$$\begin{aligned} &\stackrel{(10)}{\geq} U(s, t) \wedge U \otimes V(s, t) - b + \Sigma_{min}(t) \\ &\geq U(s, t) \wedge U \otimes V(s, t) \\ &= U \otimes V(s, t). \end{aligned} \quad (13)$$

In (12) and we used the monotonicity of min-plus convolution: $U(s, t) \geq U \otimes V(s, t)$ for any V with $V(t, t) = 0$ for all t . So, under the assumption of (10) we obtain at least the same service for the throttled system as for an unthrottled one where the servers U and V would have to be traversed. Hence, we are interested in the probability of (10) happening and call that event E . With this information at hand we can analyse the whole system by:

$$\begin{aligned} & \mathbb{P}(d_{sys}(t) > x) \\ &= \mathbb{P}(d_{sys}(t) > x | E)\mathbb{P}(E) + \mathbb{P}(d_{sys}(t) > x | \neg E)\mathbb{P}(\neg E) \\ &\leq \mathbb{P}(d_{U \otimes V}(t) > x) + \mathbb{P}(\neg E) \end{aligned} \quad (14)$$

where $\mathbb{P}(d_{U \otimes V} > x)$ can be calculated by applying Theorem 3 and Theorem 7.

We now discuss condition (10) and its corresponding probability. For ease of presentation, we leave the placeholder blank, i.e., $U \otimes V = U$.

One can rewrite (10) by:

$$\max_{0 \leq s \leq t} \{U(s, t) - U \otimes U(s, t)\} \leq b,$$

which is just the expression we arrive at when bounding the buffer in bivariate deterministic network calculus, if feeding a ‘‘flow’’ U in a dynamic $U \otimes U$ -server. Such stochastically dependent systems can be analysed in MGF-Calculus by using Hölder’s inequality. For this assume some MGF-bounds of the form (5) and (6) for U (we denote the σ and ρ corresponding to bound (5) by $\bar{\sigma}$ and $\bar{\rho}$ to distinguish them from the ones used in (6)). The probabilistic backlog bound in

such a system is $(\frac{1}{p} + \frac{1}{q} = \frac{1}{p'} + \frac{1}{q'} = 1)$:

$$\begin{aligned} & \mathbb{P}(q(t) > b) \leq \mathbb{P}(\max_{0 \leq s \leq t} \{U(s, t) - U \otimes U(s, t)\} > b) \\ &\leq e^{-\theta b} \sum_{s=0}^t \mathbb{E}(e^{\theta(U(s, t) - U \otimes U(s, t))}) \\ &\leq e^{-\theta b} \sum_{s=0}^t (\mathbb{E}(e^{p\theta U(s, t)}))^{1/p} (\mathbb{E}(e^{-\theta q U \otimes U(s, t)}))^{1/q} \\ &\leq e^{-\theta b} \sum_{s=0}^t e^{\theta \bar{\rho}(p\theta)(t-s) + \theta \bar{\sigma}(p\theta)} (\mathbb{E}(e^{-\theta q U \otimes U(s, t)}))^{1/q} \\ &\leq e^{-\theta b + \theta \bar{\sigma}(p\theta)} \sum_{s=0}^t e^{\theta \bar{\rho}(p\theta)(t-s)} \\ &\quad \cdot \left(\sum_{r=s}^t e^{\theta q \rho(qp'\theta)(r-s) + \theta q \sigma(qp'\theta)} e^{\theta q \rho(qq'\theta)(t-r) + \theta q \sigma(qq'\theta)} \right)^{1/q} \\ &\leq e^{-\theta b + \theta \bar{\sigma}(p\theta) + \theta \sigma(qp'\theta) + \theta \sigma(qq'\theta)} \\ &\quad \cdot \sum_{s=0}^t e^{\theta \bar{\rho}(p\theta)(t-s)} e^{\theta \rho(qp'\theta)(t-s)} \left(\sum_{r'=0}^{t-s} e^{\theta q(\rho(qq'\theta) - \rho(qp'\theta))r'} \right)^{1/q} \end{aligned}$$

We have used Chernoff’s inequality in the second line and Hölder’s inequality in line 3 and 5. If we assume, w.l.o.g. that $q' > p'$ holds, we have that the last sum is convergent in t . We denote the limit of that series just by B and can proceed with:

$$\begin{aligned} & \mathbb{P}(q(t) > b) \\ &\leq e^{-\theta b + \theta(\bar{\sigma}(p\theta) + \sigma(qp'\theta) + \sigma(qq'\theta))} B \sum_{s=0}^t e^{\theta \bar{\rho}(p\theta)(t-s) + \theta \rho(qp'\theta)(t-s)} \\ &\leq e^{-\theta b + \theta(\bar{\sigma}(p\theta) + \sigma(qp'\theta) + \sigma(qq'\theta))} B \sum_{s'=0}^t e^{\theta(\bar{\rho}(p\theta) + \theta \rho(qp'\theta))s'}. \end{aligned}$$

The above sum, however, does in general not converge. This is due to the three inequalities:

$$\begin{aligned} & \bar{\rho}(p\theta) \geq \bar{\rho}(\theta) \quad \forall p \geq 1, \\ & \bar{\rho}(\theta) \geq -\rho(\theta) \quad \forall \theta \geq 0, \\ & -\rho(\theta) \geq -\rho(q\theta) \quad \forall q \geq 1. \end{aligned}$$

For the case $q' = p' = 2$, the exponents in the last sum vanish and the expression reduces to $t - s + 1 \leq e^{\theta(t-s)\frac{1}{\theta}}$, such that the bound becomes:

$$\begin{aligned} & \mathbb{P}(q(t) > b) \\ &\leq e^{-\theta b + \theta(\bar{\sigma}(p\theta) + 2\sigma(2q\theta))} \sum_{s'=0}^t e^{\theta(\bar{\rho}(p\theta) + \theta \rho(2q\theta) + \frac{1}{q\theta})s'}. \end{aligned}$$

Note that the system $\xrightarrow{U} U \rightarrow U$ is unstable (the concatenation of the two servers with service $U \otimes U$ is at most as large as the arrivals U), which is why the above bound is valid only for any finite time t .

Overall, we can summarize our findings in the main result of this paper:

THEOREM 10. *Consider a WFC system as in Figure 1 with the placeholder being a dynamic V -server.*

Let $\frac{1}{p} + \frac{1}{q} = \frac{1}{p'} + \frac{1}{q'} = 1$ and $q' > p'$. Assume the following MGF-bounds on $U \otimes V$:

$$\begin{aligned}\mathbb{E}(e^{-\theta U \otimes V(s,t)}) &\leq e^{\theta \rho(\theta)(t-s) + \sigma(\theta)} \\ \mathbb{E}(e^{\theta U \otimes V(s,t)}) &\leq e^{\theta \bar{\rho}(\theta)(t-s) + \bar{\sigma}(\theta)}.\end{aligned}$$

The whole system fulfills the probabilistic delay-bound:

$$\begin{aligned}\mathbb{P}(d_{sys}(t) > x) &\leq \mathbb{P}(d_{U \otimes V}(t) > x) \\ &+ e^{-\theta \Sigma_{min}(t+T) + \sigma_E(\theta, p, p')} B \sum_{s=0}^{t+T} e^{\theta \rho_E(\theta, p, p')s}\end{aligned}$$

with

$$\begin{aligned}\sigma_E(\theta, p, p') &= \bar{\sigma}(p\theta) + \sigma(qp'\theta) + \sigma(qq'\theta) \\ \rho_E(\theta, p, p') &= \begin{cases} \bar{\rho}(p\theta) + \rho(qp'\theta) & \text{if } p' \neq 2 \\ \bar{\rho}(p\theta) + \rho(2q\theta) + \frac{1}{q\theta} & \text{if } p' = 2 \end{cases} \\ B &= \begin{cases} \left(\frac{1}{1 - e^{\theta q(\rho(qq'\theta) - \rho(qp'\theta))}} \right)^{1/q} & \text{if } p' \neq 2 \\ 1 & \text{if } p' = 2 \end{cases}\end{aligned}$$

and $d_{U \otimes V}$ being the delay of an unthrottled tandem consisting of a dynamic U - and V -server.

We want to emphasize that with this theorem it is for first time possible to analyse a general WFC system in the context of SNC. Our solution does not rely on the subadditivity of U or V directly, in contrast to what was presented in Section 3 and [3]. Instead we ask for the probability of failing the subadditivity by at least the window size Σ , which allows an analysis of general service elements inside the feedback loop. Note that U and V do not need to be single service elements themselves, they could instead result from Theorem 3 or include further elements (like fixed delay-elements or scaling elements [10]).

REMARK 11. *It is interesting to note, that, if $U \otimes V$ is subadditive already, we have $\mathbb{P}(\neg E) = 0$ and the whole system's service reduces immediately to*

$$U_{sys}(s, t) = U(s, t) \wedge U \otimes V(s, t) + \Sigma_{min}(t).$$

One can directly apply Theorem 7 on U_{sys} to achieve an end-to-end delay bound in this case. Further, we observe that step (13) is not a necessary one. We could, for example, shift b below $\Sigma_{min}(t)$ and continue directly with:

$$U_{sys}(s, t) \geq U(s, t) \wedge U \otimes V(s, t) + \Sigma_{min}(t) - b. \quad (15)$$

We can then apply Theorem 10 on this U_{sys} (which is at least as large as $U \otimes V$). One can view this as a shift in the violation probabilities towards the subadditive part (event E) of the bound. We investigate this tradeoff in the following section.

5. NUMERICAL EVALUATION

In this section, we investigate how the bound derived in Theorem 10 evolves in its parameters. Further, we quantify the impact of WFC on the delay of the system, by comparing it to a similar unthrottled system. For this section we assume U to be a constant rate server $U_o(s, t) = u(t - s)$, which also serves a crossflow $A_U(s, t)$ at higher priority than B .

A well-known result in SNC states that B receives a service $U(s, t) = u(t - s) - A_U(s, t)$. Similarly, we insert for the placeholder in Figure 1 a service element V , which also offers a constant rate $V_o(s, t) = v(t - s)$ shared with a higher priority crossflow $A_V(s, t)$, such that $V(s, t) = v(t - s) - A_V(s, t)$. Note that both service descriptions are subadditive by themselves, but when applying Theorem 3 this property is lost.

To account for the typically smaller size of acknowledgements flowing back to the throttle, we assume $v > u$. The crossflows in this example consist of i.i.d. exponentially distributed increments $a_U(t)$ and $a_V(t)$, respectively. The arrivals to the WFC system, denoted by A also consist of i.i.d. exponentially distributed increments. All flows are stochastically independent of each other. More sophisticated crossflows or arrivals are possible to analyse, as well as dropping the independency assumption, yet this is not the focus of our evaluation. We further assume a constant window-size Σ , for all times $t \geq 0$.

A corresponding unthrottled system would just consist of the flow A being fed into the service element U , thus Theorem 7 could be applied directly.

To achieve reasonable values in the bounds of Theorem 7 and Theorem 10, we numerically optimized the parameter θ and the Hölder-pairs p, q, p', q' . If not specified otherwise we used the following set of parameters in our calculations: the bound is taken at time $t = 5$ and asks for a delay $T = 10$, i.e., we consider the probability $\mathbb{P}(d_{sys}(5) > 10)$. The parameter of the exponential distributions for the arrivals and crossflows is given by $\lambda = 4$ (we assume all three flows to have the same rate λ for simplicity), while the server-rates are $u = 1$ and $v = 2$. This corresponds to a utilization of 50% and 25%, respectively. We present the results for a window size of $\Sigma = 15$.

5.1 Throttled vs. Unthrottled System

First we want to compare the system to its unthrottled counterpart. To that end, we alter the arrival rates λ , resulting in utilizations from 30% to 80%. We plot the corresponding violation probabilities for the performance bounds on a logarithmic scale for the throttled, as well as the unthrottled system. We did this for different window-sizes $\Sigma = 10, 15, 20$. The results are displayed in Figure 2 as black and blue lines for the throttled and unthrottled system, respectively. As expected, the throttled system behaves better, the larger Σ is; for $\Sigma = 20$ the throttled system behaves almost identically to the unthrottled one.

5.2 Dependence on Delay

A major difference between the unthrottled and the throttled analysis lies in the dependency on the delay T . While for the unthrottled system an increase in T leads to a decrease in the violation probability, we see in the bound of Theorem 10, that the term $\mathbb{P}(\neg E)$ increases in T . In Figure 3, one sees for the black line how the bound evolves for an increasing T . The two red lines show how the bound differs when choosing $b = \frac{\Sigma}{2}$, $\frac{9\Sigma}{10}$ as suggested in Remark 11 (Equation (15)). The trend here for larger T is, that the bound becomes worse, the larger the difference between b and Σ is. However, for small values of T there is a very slight improvement for $b = \frac{9\Sigma}{10}$ and even for $b = \frac{\Sigma}{2}$. In this




Figure 2: A graph showing the violation probabilities depending on the utilization of throttled (black) and unthrottled (blue) systems for different window sizes. The red lines are equal to 1 , 10^{-3} and 10^{-6} .




Figure 3: A graph showing how the bound evolves when increasing the delay T . The red colored lines represent a shift towards the violation probability of event E .

scenario, trading a higher violation probability for the event E is not worthwhile the gain from a better service description U_{sys} .

To investigate the composition of the delay bound further, we separated the two parts in Figure 4. The blue circles correspond to the delay-part of the violation probability $\mathbb{P}(d_{U \otimes V}(t) > x)$ and the red circles correspond to the violation probability of event E , while the solid black circles are the sum of both. It can be clearly observed that from a certain point onwards the probability $\mathbb{P}(\neg E)$ dominates the overall violation probability. The additional lines drawn into the graph show how the different parts of the bound are affected when we use $b = \frac{\Sigma}{2}$, $\frac{3\Sigma}{4}$, $\frac{9\Sigma}{10}$ in Equation (15). It can be seen that the delay-part (blue) of the probability experiences no considerable change, while the probability of violating event E (red) increases significantly, when $b < \Sigma$.




Figure 4: A graph showing the different components of the violation probability: the blue circles are the delay-part, while the red circles represent the subadditivity-part. The black circles are the sum of both parts. The lines show the same for different shifts towards the violation probability $\mathbb{P}(\neg E)$.

5.3 Convergence to Unthrottled System

In Figure 5, we consider the convergence of the throttled system towards the unthrottled one when increasing the window size. Clearly, from Theorem 10 the violation probability $\mathbb{P}(\neg E)$ vanishes for increasing window sizes. However, the throttled system does not fully converge to the unthrottled one, since the delay-part of the bound still differs:

$$\mathbb{P}(d_U(t) > T) \leq \mathbb{P}(d_{U \otimes V}(t) > T).$$

The size of the gap, which cannot be closed by increasing the window size further is completely dependent on the service descriptions U and V . We present in the graph the same system as before, but vary $v = 2, 1.5, 1.1$. One can see clearly how the gap to the delay of the unthrottled system (red) increases, when reducing the rate of v .

6. CONCLUSION AND OUTLOOK

In this paper, we have dealt with the long-standing problem of analysing WFC systems in SNC. While such feedback loops had been solved in deterministic network calculus more than a decade ago, its counterpart in the stochastic setting has been a well-known open problem [15, 12, 13, 8]. We presented how far subadditive service carries DNC solutions for WFC systems into stochastic network calculus (see also [3]). In that discussion, we encountered the very general notion of σ -additive operators and saw as a tractable example a feedback-loop containing a single subadditive server. Unfortunately, this method reaches the end of the road as soon as operators appear which no longer commute, or are not idempotent. This is not untypical in applications, for example if tandems of servers are involved.

Therefore, we approached subadditivity in a different way and hence, for the first time, successfully analysed general WFC systems in the context of SNC. Instead of assuming subadditive service elements, we leverage the stochastic nature of the problem and ask for the probability of the feedback loop not being subadditive. This effectively allows the

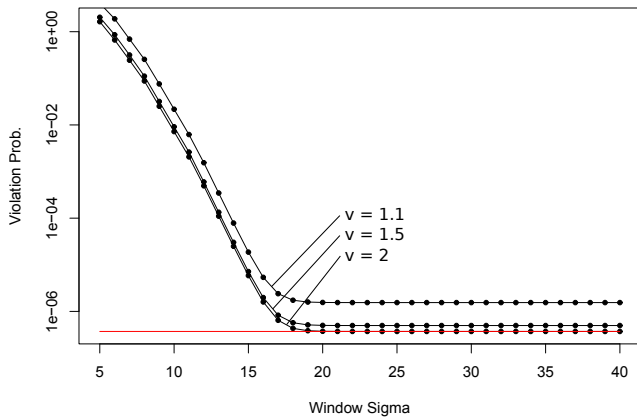


Figure 5: A graph showing the convergence of the throttled system towards the unthrottled one, when increasing the window Σ . The red line is the bound for the unthrottled system. The black lines show the throttled system, for different rates of v .

analysis of WFC systems in MGF-based network calculus. The resulting bounds consist of two parts: first, a delay-bound of a conventional unthrottled system, containing the feedback loop as service; second, a probability of violating the subadditivity, by more than the window size. The structure of our result makes a direct comparison between throttled and unthrottled systems possible.

The analysis of WFC systems in stochastic network calculus is not completed yet, but has rather just begun. While the now available methods can handle varying window sizes Σ , they can take only limited advantage of their variations. The presented method uses a backlog bound for the system $\xrightarrow{U} U \otimes U \rightarrow$. The “arrivals” and “service” in this scenario are strongly correlated. While using Hölder’s inequality deals correctly with that dependence, it also neglects its possible advantages. As the arrivals and the service in this system are positively correlated one can hope to improve the bounds significantly, when taking the dependencies into account.

Besides improving bounds in the above sense, one can extend and build upon this work: one direction is to break the “end-to-end” feedback-loop into several hops, resulting in a tandem of WFC systems. Another interesting question would be how to effectively handle stochastic dependencies between the “upstream”-service U , the “down-stream”-elements and the window-process Σ . Answering this will push the applicability of SNC even further. By better grasping the occurring dependencies one can eventually aim at analysing systems like the window-controlled TCP in SNC.

7. REFERENCES

- [1] R. Agrawal, R. L. Cruz, C. Okino, and R. Rajan. Performance bounds for flow control protocols. *IEEE/ACM Transactions on Networking*, 7(3):310–323, June 1999.
- [2] M. A. Beck. A first course in stochastic network calculus. Online lecture notes, October 2013. <http://goo.gl/qSIIwr>.
- [3] M. A. Beck and J. Schmitt. Window flow control in stochastic network calculus. Technical Report 391/15, University of Kaiserslautern, September 2015. <http://goo.gl/K4sWhQ>.
- [4] J.-Y. Le Boudec and P. Thiran. *Network Calculus - A Theory of Deterministic Queuing Systems for the Internet*. Springer, 2001.
- [5] C.-S. Chang. On deterministic traffic regulation and service guarantees: A systematic approach by filtering. *IEEE Transactions on Information Theory*, 44(3):1097–1110, May 1998.
- [6] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer, 2000.
- [7] F. Ciucu, A. Burchard, and J. Liebeherr. A network service curve approach for the stochastic analysis of networks. In *Proc. of ACM SIGMETRICS*, June 2005.
- [8] F. Ciucu, M. Fidler, J. Liebeherr, and J. Schmitt. Network calculus. *Dagstuhl Reports*, pages 63–83, March 2015.
- [9] F. Ciucu and J. Schmitt. Perspectives on network calculus – no free lunch, but still good value. In *Proc. of ACM SIGCOMM*, August 2012.
- [10] F. Ciucu, J. Schmitt, and H. Wang. On expressing networks with flow transformations in convolution-form. In *Proc. of IEEE INFOCOM*, April 2011.
- [11] M. Fidler. An end-to-end probabilistic network calculus with moment generating functions. In *Proc. of IWQoS*, June 2006.
- [12] M. Fidler. Survey of deterministic and stochastic service curve models in the network calculus. *IEEE Communications Surveys Tutorials*, 12(1):59–86, 2010.
- [13] M. Fidler and A. Rizk. A guide to the stochastic network calculus. *IEEE Communications Surveys and Tutorials*, 17(1):92–105, 2015.
- [14] Y. Ghiassi-Farrokhfal, S. Keshav, and C. Rosenberg. Toward a realistic performance analysis of storage systems in smart grids. *IEEE Transactions on Smart Grid*, 6(1):402–410, Jan 2015.
- [15] Y. Jiang and Y. Liu. *Stochastic Network Calculus*. Springer, 2008.
- [16] J.-Y. Le Boudec and P. Thiran. A note on time and space methods in network calculus. In *International Zurich Seminar on Broadband Communications*, pages 267–272, Feb 1998.