Multimedia Communications Lab
Department of Electrical Engineering and Information Technology
Technische Universität Darmstadt
Merckstr. 25
64283 Darmstadt, Germany

Technische Universität Darmstadt

Technical Report

Optimal Allocation of Service Curves by Exploiting Properties of the Min–plus Convolution

Krishna Pandit, Jens Schmitt, Claus Kirchner, and Ralf Steinmetz

Version 1.1
First published: November 29th, 2004
Last major revision: November 29th, 2004
Last minor revision: December 22nd, 2004

TR-KOM-2004-08

Version history:

Version 1.0 published on November 29th 2004
Version 1.1 published on December 22nd 2004

Abstract

Providing deterministic Quality of Service (QoS) in packet-switched networks like the Internet, which allows non-critical traffic as well, remains an open research issue. It has been shown [17] how optimal network service curves for bandwidth/delay decoupled scheduling disciplines are found. In this paper, we show how to optimally allocate service curves to nodes along the path to obtain the optimal network service curve. We further show how the nodes can locally optimize their resources to accommodate as many service curve requests as possible. These results are achieved by exploiting novel properties of the convolution operation in the Min-plus algebra, which are derived in this work. Therefore, as a second contribution, this paper gives insight on the min-plus convolution, as well as transforms in the min–plus algebra. Particularly, an efficient method to compute the min-plus convolution under Network Calculus constraints is developed. The key results of this paper are accompanied by numerical examples.

Optimal Allocation of Service Curves by Exploiting Properties of the Min–plus Convolution

Krishna Pandit*, Jens Schmitt†, Claus Kirchner‡, Ralf Steinmetz

1 Introduction

1.1 Motivation

Quality of Service (QoS) in packet-switched networks remains a much debated research issue. There have been recent doubts and frustrations, especially regarding QoS in the Internet. However, it is unlikely that the status quo is the final peal of wisdom. In the future, there will be a multi-service IP network which will carry many different types of traffic, which are today carried over legacy networks. The applications range from time-uncritical FTP connections, over semi-critical Voice-over-IP conversations and video conferences to mission critical connections such as air traffic control or emergency rescue calls. The latter require deterministic guarantees and even though it will certainly not make up a large proportion of the bandwidth, it is worth optimizing due to its expensive implementation and high significance. The mission critical applications usually are low bandwidth, short delay flows, which we focus on in this paper. Further there will be closed, specialized networks, optionally based on IP. An example of such a network would be in a car, where steer-by-wire, the multimedia entertainment system as well as the log book share the communcation channel. This work builds upon a result by Schmitt [17], where the optimal network service curves for such flows were deduced. Here we discuss how to set the parameters of the service curves in the nodes along the path in order to obtain the desired network service curve. To date, most research on QoS has gone into developing and

*K. Pandit and R. Steinmetz are with the Multimedia Communications Lab, Department of Electrical Engineering and Information Technology, Technische Universität Darmstadt, Darmstadt, Germany, {Krishna.Pandit, Ralf.Steinmetz}@kom.tu-darmstadt.de

†J. Schmitt is with the Distributed Computer Systems Lab, Department of Computer Science, University of Kaiserslautern, Kaiserslautern, Germany, jschmitt@informatik.uni-kl.de

‡C. Kirchner is with the Technomathematics Group, Department of Mathematics, Technische Universität Kaiserslautern, Kaiserslautern, Germany, kirchner@mathematik.uni-kl.de

optimizing new architectures. The development of a mathematical framework has fallen short. Our long term vision is a system theory for packet-switched networks, which plays a role similar to classical system theory for the analysis of electronic circuits, wireless channels, etc. Such a system theory would be attractive as development costs could be saved. Currently, simulations have to be conducted to evaluate new ideas regarding algorithms for packet-switched networks. If they prove worthwhile, a prototype implementation is next, and at the end stands a product. In opposition, in many fields mathematical frameworks exist that allow a theoretical analysis prior to simulation. An example is the field of developing new codes for wireless channels. E.g., in the paper by Sandhu et al. [15] this is demonstrated by using the existing theoretical framework to analyze the gain of a non-linear space/time block code. Besides the cost factor an advantage is that some negative side effects of simulations can be overcome (e.g. random number generation, parameter estimation) and radically different concepts, which usually do not fit into existing simulation environments, can be evaluated. Last but not least, mathematical properties from the analysis can lead to improvements of an implementation, as demonstrated in this paper. It will be shown how the properties of the min–plus convolutioncan be used for optimal admission control.

Network Calculus, which was invented by Cruz [6] [7], has the potential of becoming the tool that fulfills this vision. Therefore, a significant part of this paper deals with enhancing Network Calculus. Explicitly, we shed light onto transforms in the min–plus algebra and develop new properties of the min–plus convolution, which is the second contribution of this paper.

This paper is organized as follows. In the next section we review the necessary results from Network Calculus. We then give an overview of related work, especially the paper from Schmitt [17] that this work builds upon. In Section 5 we derive a theorem for calculating the min–plus convolutionof a certain kind of functions typical for Network Calculus, and analyse its properties. We then apply this theorem, first in Section 7 to optimally allocate service curves along a path, and then in Section 8 to optimally utilize the resources of one node. Finally we conclude and give an outlook.

2 Background

2.1 Selected Results from Network Calculus

Network Calculus [6] [7] is a theory for deterministic queueing systems. The underlying idea is that service guarantees can be achieved by regulating the traffic and deterministic scheduling. Analogous to conventional system theory, a system consists of an input, a transfer function and an output. The input, mostly referred to as *arrival curve*, is an abstraction of the traffic regulation, and the transfer function, mostly referred to as *service curve*, is an abstraction of the scheduling. A *dioid* is a mathematical structure with addition (often denoted by $\oplus$) and multiplication (often denoted by $\otimes$) as two inner operations. The difference to conventional system theory is that the dioid $\{\mathbb{R} \cup \infty, \min, +\}$ is used, i.e., that addition and multiplication are replaced by minimum and addition, respectively. This is often referred to as *Min-plus Algebra*. The reason to switch to Min-plus Algebra is that this way linearity is preserved. In the following, we recapitulate the results from Network Calculus which are relevant for this paper. They can all be found in the excellent text of Le Boudec and Thiran [2]. As in conventional system theory, a key operation in Network Calculus is the min-plus convolution. Note that the infimum (inf) of a set A is similar to the minimum (min) of this set, with the sole difference that the former does not have to be in the set. The same applies for the supremum (sup) and maximum (max) of a set A. We set

$$\mathcal{F} \quad := \quad \{f : \mathbb{R} \to \mathbb{R}_0^+ \mid f(t) = 0 \text{ for } t < 0, f(t_1) \leq f(t_2) \text{ for } t_1 \leq t_2\}$$

i.e. $\mathcal{F}$ is the set of nonnegative wide-increasing functions.

Definition 1 [Min-plus convolution] Let f and g be two functions or sequences of $\mathcal{F}$. The min-plus convolution of f and g is the function

$$(f \otimes g)(t) = \inf_{0 \leq s \leq t} \{f(t-s) + g(s)\} \tag{1}$$

The traffic bound is given by an *arrival curve*, which denotes the largest amount of traffic allowed to be sent in a given time interval.

Definition 2 [Arrival Curve] Given a wide-sense increasing function α defined for $t \geq 0$, we say that a flow R is constrained by α iff for all $s \leq t$

$$R(t) - R(s) \leq \alpha(t-s)$$

We say that R has α as an arrival curve, or also that R is α-smooth.

The arrival curve can be viewed as an abstraction of the regulation algorithm. The most prominent example for a traffic regulation algorithm is the Leaky Bucket [19], which is often also referred to as Token Bucket. Its arrival curve is given by the following equation.

$$\alpha(t) = b + rt \text{ for } t > 0 \tag{2}$$

Therefore, no more than b data units can be sent at once and the long-term rate is r.

From this the Traffic Specification (TSpec) of Integrated Services evolved. It consists of two concatenated Token Buckets and has the following arrival curve.

$$\alpha(t) = \inf(M + pt, b + rt) \text{ for } t > 0 \tag{3}$$

A *greedy shaper* with the shaping curve σ optimally delays packets, so that the output has σ as an arrival curve, and sends all bits as soon as possible.

Theorem 1 [Greedy Shaper] Consider a greedy shaper with shaping curve σ, which is sub-additive and $\sigma(0) = 0$. Assume that the shaper buffer is empty at time 0, and that it is large enough so that there is no data loss. For an input flow R, the output R^o is given by

$$R^o = R \otimes \sigma \tag{4}$$

We omit the proof as it can be found in [2]. The service curve is an abstraction of the scheduling.

Definition 3 [Service Curve] Consider a system S and a flow through S with input and output functions R and R^o, respectively. We say that S offers to the flow a service curve β if and only if $\beta \in \mathcal{F}$ and $R^o \geq R \otimes \beta$.

Theorem 2 [Concatenation of service surves] Consider n nodes, each offering a service curve β_i. A flow traversing these nodes in sequences is then offered the service curve β.

$$\beta = \beta_1 \otimes \beta_2 \otimes \ldots \otimes \beta_n \tag{5}$$

We omit the proof as it can be found in [2].

Due to its application in the Integrated Services context, a prominent service curve is the rate-latency function.

Definition 4 [Rate-latency functions $\beta_{R,T}$]

$$\beta_{R,T} = R[t-T]^+ = \begin{cases} R(t-T) & \text{if } t > T \\ 0 & \text{otherwise} \end{cases} \tag{6}$$

for some $R \geq 0$ (the 'rate') and $T \geq 0$ (the 'delay').

We refer to this hereafter as the LR scheduler or LR service curve. It is also the basis of the L2R scheduler, which will be introduced in Section 3 and used throughout the paper.

2.2 Conventional system theory

In order to not cause confusion with max/min-plus system theory we use the term conventional system theory. In this section we give a brief review of conventional system theory. For a detailed description we refer the reader to the excellent textbook by Oppenheim et al. [11].

The Z-Transforms plays an important role in the context of conventional system theory; in particular, the Fourier transform is a special case. We denote by $x(n)$ the input to a linear, time-invariant system. Then the (discrete) Z-Transform $X(z)$ of the signal $x(n)$ is defined as:

$$X(z) = \sum_{n=-\infty}^{\infty} x(n) z^{-n} \tag{7}$$

The (continuous) inverse Z-Transform is then given by:

$$x(n) = \frac{1}{2\pi} \oint X(z) z^{n-1} dz \tag{8}$$

The output $y(n)$ of an input $x(n)$ of a linear, time-invariant system with the impulse response $h(n)$ is given by the convolution sum:

$$y(n) = \sum_{k=-\infty}^{\infty} x(k) h(n-k) \tag{9}$$

Further, it can be shown that

$$Y(z) = X(z) H(z) \tag{10}$$

$Y(z)$, $X(z)$ and $H(z)$ are the z-transforms of $y(n)$, $x(n)$ and $h(n)$, respectively. Eq. 10 implies that the convolution from Eq. 9 can be carried over to a multiplication in the z-domain. Similar results exist for the continuous case, e.g., the Fourier transform or, more generally, the Laplace transform. An alternative way to compute the convolution in conventional system theory is to apply the Fourier transform to the functions, multiply the result and apply the inverse Fourier transform. In certain cases this procedure can be favourable in terms of computational complexity compared to the convolution operation itself. Therefore, it is worth looking into transforms for the min-plus algebra in order to compute the Min-plus convolution, which we will do in section 4.

3 Related Work

As mentioned before, the inspiration for this work is a result by Schmitt [17], where he develops the optimal service curve for a bandwidth/delay-decoupled scheduler. The optimal service curve is denoted by s^{opt}. The

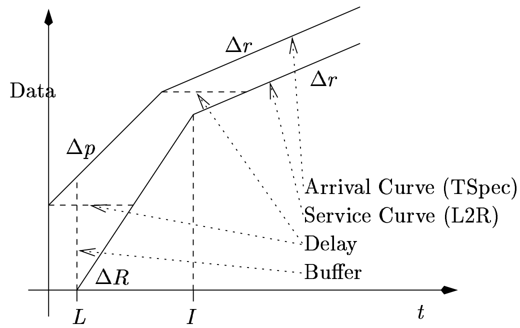

Figure 1: L2R scheduler

bandwidth/delay-decoupled scheduler is characterized by a L2R service curve, i.e., a service curve that has a latency, and two rates. The latency is determined by factors such as the hardware and the scheduling discipline and is therefore not a design parameter for us. After the latency the flow is served with a rate R up to a certain time which we refer to as the inflection point I. After that it switches down to rate r, which is the sustained rate of the arrival curve. As depicted in Figure 1, s^{opt} is given by the 4-tuple (L, I, R, r), where L is the latency and I the inflection point where the rate switches from R to r. Further, we define $U = I - L$, which is the time that the flow is served at peak rate.

In this case, optimality is defined on a per-flow perspective, i.e. minimizing the resource consumption of a single flow. The optimal network service curve for an arrival curve is derived. In this work we extend those results to a wider perspective, in particular we concentrate on two aspects. The implications of an optimal network service curve to the service curves of the nodes along the path are studied as well as the allocation of resources within one node to accommodate as many flows as possible. It turns out that the inflection point is chosen such that the delay bound is met at the point where the arrival curve makes its last rate change to the sustained rate. For the Tspec, this is the point where it switches from peak rate to sustained rate, as shown in Figure 1. Besides deriving the optimal service curve, numerical examples are given which show the benefit of this approach. Schmitt further discusses related work, which obviously covers most of our related work as well. Most interesting are the deterministic schedulers, as they are the basis for allocating resources with the service curve approach. Sariowan et al.[16] and Stoica et al. [18] describe algorithms which schedule according to service curves.

Looking at the big picture, our work can be seen as a step towards developing a system theory for packet-switched networks on the basis of network calculus. There are several researchers around the world who share this goal, each approaching it from a different angle. What they have in common is that the focus is on developing new min–plus based methods for Network Calculus, which solves an issue in contemporary networking research. We exemplary point out two approaches. Bearing in mind that the common goal is rather vague, they are far apart from our work as well as each other. Chang et al. [4] extend the min–plus system theory to describe constrained traffic regulation and dynamic service guarantees. Constrained traffic regulation is achieved by concatenation of the so-called *g-clipper* and *maximum service regulator*. To obtain dynamic service guarantees results from filtering theory with time-varying impulse responses are mapped to min–plus system theory. Another big battlefield is enhancing Network Calculus with probabilistic elements. Leading is the group around Liebeherr, who [10] introduce the concept of Statistical Network Calculus. In the first incarnation this is based on the assumption that an arrival curve does not deterministically bound the incoming traffic but bounds it only with a certain probability. Similarly, a statistical service curve [3] is a service curve that only offers the service with a certain probability. Targeting the same goal from a different angle are Pandit et al. [13], who conduct a simulation based analysis of a Token Bucket bounded queue.

4 A Transform for Network Calculus

The min–plus convolutionis a key operation in Network Calculus. It has a similar role as the conventional convolution in conventional system theory. Therefore, it is worth analysing this operation to find efficient ways to compute it. In conventional system theory, computation algorithms for the convolution often rely on transforms, since the convolution in the time domain corresponds to the multiplication in the frequency domain. According to [14], an efficient numerical algorithm to compute the conventional convolution is via the Fast Fourier Transform (FFT). Hence, we turn our attention to transforms. When looking for a transform for Network Calculus, the first thing that comes to mind is the *Fenchel transform*, also known as the convex conjugate function. In the book by Bacelli et al.[1] it is briefly pointed out that this transform carries over the min–plus convolutionin one domain to an addition in another domain. In the book by Hiriart-Urruty and Lemarechal [8] one finds the following definition for the convex conjugate function:

Definition 5 The Fenchel transform, or convex conjugate function, is given by

$$f^*(s) = \sup\{sx - f(x) \mid x \in dom\, f\}$$

Further, the bi-conjugate function is given by

$$f^{**}(x) = \sup\{sx - f^*(s) \mid s \in dom\, f^*\}$$

Unfortunately, a shortcoming of the Fenchel transform is that it only works well for convex functions. For piecewise linear, convex functions (and slightly more general ones) an efficient algorithm to obtain the min–plus convolutionis outlined in Chapter 3 in the book by Le Boudec [2], where it is pointed out that the min–plus convolutionof two piecewise linear, convex functions is obtained by simply sorting the slopes of the individual functions. In general, the bi-conjugate of a function yields the convex closure of the function. For closed convex functions f (e.g. the piecewise linear, convex functions in [2] fall into this category) we have $f = f^{**}$. But in gerenal, i.e. for non-convex functions, we can not hope for equality. Therefore, the Fenchel transform is not suited in the context of Network Calculus. Applications and problems of the Fenchel transform in the context of Network Calculus are described in the technical report by Pandit et al.[12]. Furthermore example computations are provided in this work. The next transform we consider is the Γ-transform. The Γ-transform can easily be represented graphcically by so-called *point clouds*, This transform is explained in the book by Bacelli et al.[1], and thoroughly described in the dissertation of Holger Jäkel (in German) [9]. Here we limit ourselves to a descriptive discussion on this transform, the mathematically inclined reader is referred to the above two texts. An element $b \in \mathbb{B}[[\gamma, \delta]]$ is a formal power series in two variables (γ, δ) with Boolean coefficients (the subscript x can be ignored for now):

$$b = \sum_{k,t \in \mathbb{Z}} s_x(k,t)\delta^k \gamma^t = \sum_{k=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} s_x(k,t)\delta^k \gamma^t \tag{11}$$

The possible "value" for the coefficients $s_x(k,t), k, t \in \mathbb{Z}$ in the min–plus algebra are $s_x(k,t) = e = 0$ and $s_x(k,t) = \epsilon = +\infty$. A sequence $x(k), k \in \mathbb{Z}$ can be transformed into $\mathbb{B}[[\gamma, \delta]]$ via

$$s_x(k,t) = \begin{cases} 0 & x(k) = t \\ +\infty & x(k) \neq t \end{cases} \tag{12}$$

This operation is called the Γ-*transform*.

For two elements $b_1, b_2 \in \mathbb{B}[[\gamma, \delta]]$ we define their sum (in the min-plus sense) componentwise as the usual minimum operation of their coefficients. Denote by $s_i(k,t), i = 1, 2$ the coefficients of $b_i, i = 1, 2$. Then the coefficients $(s_1 \oplus s_2)(k,t)$ of $b_1 \oplus b_2$ are given by

$$(s_1 \oplus s_2)(k,t) = \min_{k,t} s_1(k,t), s_2(k,t) \tag{13}$$

The multiplication of two elements $b_1, b_2 \in \mathbb{B}[[\gamma, \delta]]$ is more involved. Again, let $s_i(k,t), i = 1, 2$ denote the coefficients of $b_i, i = 1, 2$. Then the coefficients $(s_1 \otimes s_2)(k,t)$ of $b_1 \otimes b_2$ are given by

$$(s_1 \otimes s_2)(k,t) = \min_{i_1+i_2=k, j_1+j_2=t} [s_1(i_1,j_1) + s_2(i_2,j_2)] \tag{14}$$

While the addition operation is trivial, the multiplication operation is not as intuitive. The graphical interpretation of the addition operation is simply to take all points of the underlying point clouds, choose the minimum of those and fix it as a new point for the point cloud of the sum. The multiplication operation is explained below using a back-of-the-envelope example.
In order to compute the min–plus convolutionof two functions f and h, we sample them (i.e. we obtain a sequence $f(k), h(k), k \in \mathbb{Z}$), apply the Γ-Transform, add the resulting elements in $\mathbb{B}[[\gamma, \delta]]$ and apply the inverse Γ-Transform. However, it will also show that this is merely a different representation of obtaining the min–plus convolutionwithout any gain on the computing complexity. We have to limit ourselves to a discrete excerpt (referred to as "sample" in the following) of the actually continuous functions, as with the number of points the complexity rises exponentially.

Let $x_1(k)$ be a sample of a token bucket arrival curve with bucket depth 2 and slope 1.

$$x_1(k) = \begin{cases} 0 & k = 0 \\ 2+k & 1 \leq k \leq 3 \\ +\infty & \text{otherwise} \end{cases}$$

The Γ-transform Γ_1 of $x_1(k)$ is depicted in Figure 2 a). Note that we always assume the function to be 0 at $k = 0$. Similarly, $x_2(k)$ is a sample of a LR service curve with latency 2 and slope 2.

$$x_2(k) = \begin{cases} 0 & k = 0 \\ 2(k-2) & 2 \leq k \leq 4 \\ +\infty & \text{otherwise} \end{cases}$$

The Γ-transform Γ_2 of $x_2(k)$ is depicted in Figure 2 b). In Figure 2 c) the multiplication operation of Γ_1 and Γ_2 is depicted. Every point is added to every other one. Therefore, since we have 4 points in each cloud, the result should be 16 points. The avid reader counts only 15, as one point, the one with the square around it, falls twice on the same spot. This is also the point where the resulting min–plus convolutionof two such service curves would have the switch from the slope from the service curve, to the slope from the arrival curve. The resulting min–plus convolutionof the curves is depicted by the dashed line. The reverse Γ-transform is obtained by taking the corners of the area above or left of each point. This area is shaded in Figure 2 c).

It can be clearly seen that the number of points we picked is not sufficient to obtain the correct result even for the depicted range. It is also not possible to find some characteristic points which describe the min–plus convolutionof

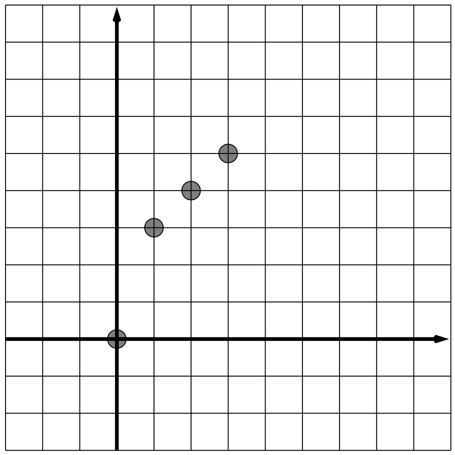

a) $s_1(k,t)$

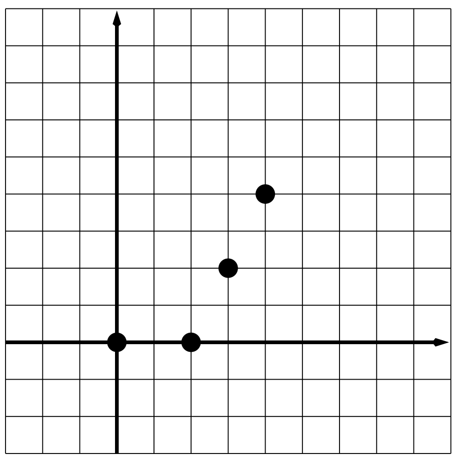

b) $s_2(k,t)$

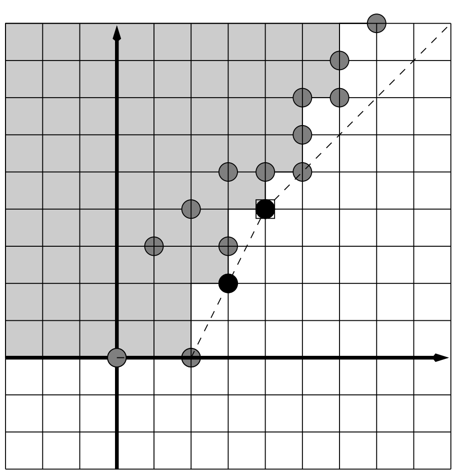

c) $s_1(k,t) \otimes s_1(k,t)$

Figure 2: Discrete transform

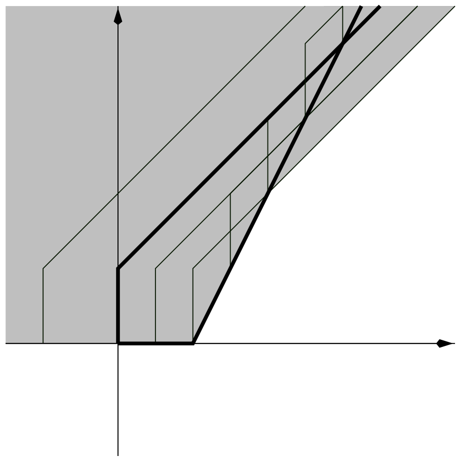

Figure 3: Continuous transform

piecewise linear functions. Therefore, the Γ-transform is not very helpful to simplify the general computation of the min–plus convolution.

The main problem is the discretization or sampling procedure. The accuracy of the computation depends strongly on the amount of discrete points we choose. However, the more points we take into account, the closer the result becomes to the min–plus convolution. Taking this considerations to the limit, we end up at the *continuous* Γ*-transform*. A function, such as an arrival curve or service curve, can be viewed as an infinite set of points. We proceed with them as with the points in the Γ-transform. This is depicted in Figure 3. Underlying are an arrival curve, and a service curve. The are obtained when all their points are summed with each other, and it can be easily verified that the border of this are denotes the min–plus convolution. With this in mind, we bring the attention back to Eq. 14 and realize that this merely is another representation of the min–plus convolution. Taking a close look at this graphical representation, we obtain a theorem which is given in the following section.

5 Min-plus Convolution under Network Calculus constraints

In this section we derive an efficient method to calculate the min–plus convolutionunder Network Calculus constraints.

Remark 1 All functions f, g in this section belong to $\mathcal{F}$. This is important, since if the range of a function f is not a subset of the real numbers, one can not speak of "the slope of f or g" that easily. Accordingly, the following definition would not make sense any more.

First we define the convex inflection point, as it will play a key role in the course of this paper. Recall that for a function $f \in \mathcal{F}$, f piecewise linear, a reflection point $R \in \mathbb{R}$ is simply a point with $f'(R^-) \neq f'(R^+)$; i.e. the slope of f changes at R.

Definition 6 [Convex and concave inflexion points] We define a *convex inflection point* R as an inflection point, where the slope to its right is greater than the slope to its left. More mathematically, a convex inflection point R has the property that $f'(R^+) > f'(R^-)$. Analogously, we call an inflection point where the slope to its right is less than the slope its left a *concave inflection point*. Thus, a concave inflection point R has the property that $f'(R^+) < f'(R^-)$

Next we discuss what we mean by Network Calculus constraints. This refers to the shape of the functions. We argue that the relevant functions have two properties

1. All functions are monotonously (wide-sense) increasing

2. All functions are 0 for $t \leq T$ and *concave* for $t > T$.

As pointed out in section 2.1, the underlying functions are the arrival curves as well as the service curves. The first property is straight-forward, as we are dealing solely with cumulative functions. The second property holds for arrival curves, as they are are sub-additive. Further, the service curve from [17], which is the work we build upon, satisfies the second property.

Therefore, we are interested in the min–plus–convolution of special functions f and g. Let $i, j = 0, \ldots, n$, $A_i, A_j, B_i, B_j \in \mathbb{R}^+$ with $A_i > A_j$ and $B_i > B_j$ if $i > j$. Then define $I_j^A = [A_{j-1}, A_j]$ with $A_{-1} = -\infty$ and $A_{n+1} = +\infty$. Now we can state the functions f and g to be investigated in the rest of this section.

$$f(t) = \begin{cases} 0 & t \in I_0^A \\ a_0(t - A_0) & t \in I_1^A \\ a_0(A_1 - A_0) + a_1(t - A_1) & t \in I_2^A \\ \vdots & \\ \sum\limits_{i=0}^{n-2} a_i(A_{i+1} - A_i) + a_{n-1}(t - A_{n-1}) & t \in I_n^A \\ \sum\limits_{i=0}^{n-1} a_i(A_{i+1} - A_i) + a_n(t - A_n) & t \in I_{n+1}^A \end{cases}$$

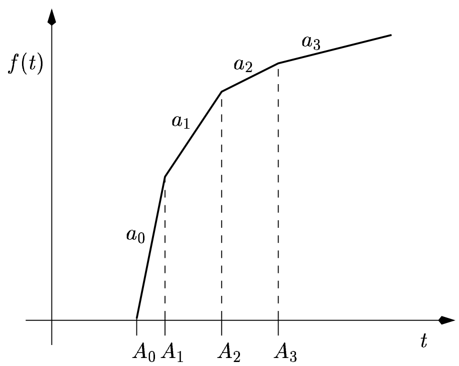

Figure 4: Shape of functions f and g

$$g(t) = \begin{cases} 0 & t \in I_0^B \\ b_0(t - B_0) & t \in I_1^B \\ b_0(B_1 - B_0) + b_1(t - B_1) & t \in I_2^B \\ \vdots & \\ \sum_{i=0}^{n-2} b_i(B_{i+1} - B_i) + b_{n-1}(t - B_{n-1}) & t \in I_n^B \\ \sum_{i=0}^{n-1} b_i(B_{i+1} - B_i) + b_n(t - B_n) & t \in I_{n+1}^B \end{cases}$$

The functions f and g have exactly one convex inflection point each at A_0 and B_0, respectively, and only concave inflection points thereafter. I.e., we have $a_i < a_j$ and $b_i < b_j$ for $j < i$. In other words, we assume $a_n < a_{n-1} < \ldots < a_1$ and $b_n < b_{n-1} < \ldots < b_1$. In the style of L2R functions, we call them *LnR functions*. The shape of f and g is given in Figure 4.

We are now ready to formulate the theorem, which is the foundation of this work.

Theorem 3 The min–plus–convolution of f and g is given by

$$(f * g)(t) = \inf_{0 \le s \le t} \{f(t - s) + g(s)\} = \min\{f(t - B_0), g(t - A_0)\}$$

Proof We assume $A_0 \le B_0$ without loss of generality. For $t \le A_0$ let $s = 0$ and we have

$$(f * g)(t) = f(t) + g(0) = 0$$

For $A_0 < t \le A_0 + B_0$ let $s = A_0$ which yields

$$(f * g)(t) = f(A_0) + g(t - A_0) = 0$$

Let $t > A_0 + B_0$ and assume that

$$f(t - B_0) \ge g(t - A_0)$$

Let x be such that $A_0+B_0+x<t$. Then we have the following inequalities:

$$\frac{f(t-B_0)-f(A_0)}{t-B_0-A_0} \leq \frac{f(A_0+x)-f(A_0)}{x} \tag{15}$$
$$\frac{g(t-A_0)-g(t-A_0-x)}{x} \leq \frac{g(t-A_0)-g(B_0)}{t-B_0-A_0} \tag{16}$$
$$\frac{f(t-B_0)-f(A_0)}{t-B_0-A_0} \geq \frac{g(t-A_0)-g(B_0)}{t-B_0-A_0} \tag{17}$$

Inequality (17) simply restates our assumption $f(t-B_0) \geq g(t-A_0)$. Inequality (15) and (16) hold, since f respectively g are concave. We will give a proof below.
We can write down a series of inequalities:

$$\begin{aligned}\frac{g(t-A_0)-g(t-A_0-x)}{x} &\leq \frac{g(t-A_0)-g(B_0)}{t-B_0-A_0}\\ &\leq \frac{f(t-B_0)-f(A_0)}{t-B_0-A_0}\\ &\leq \frac{f(A_0+x)-f(A_0)}{x}\end{aligned}$$

Multiplying by $x>0$ and rearranging we obtain

$$g(t-A_0)+f(A_0) \leq f(A_0+x)+g(t-A_0-x)$$

Setting $y := A_0+x$ and recalling $f(A_0)=0$ we have

$$(\dagger) \qquad g(t-A_0) \leq f(y)+g(t-y)$$

Proof of inequality (15)

We have $t > A_0+B_0$. For $x \geq 0$ with $t \geq A_0+B_0+x$ we have $A_0 \leq A_0+x \leq t-B_0$, hence A_0+x can be written as unique convex combination of the points A_0 and $t-B_0$. Therefore we need to solve the following:

$$A_0+x = \alpha \cdot A_0 + (1-\alpha)\cdot(t-B_0)$$

A simple computation yields

$$\alpha = \frac{A_0+x-t+B_0}{A_0-t+B_0} = 1-\frac{x}{t-A_0-B_0} < 1$$

Hence we have

$$1-\alpha = \frac{x}{t-A_0-B_0}$$

Now we use that f is concave. Recall, that a function η is concave on $[a, b]$, if for every $\alpha \in [0, 1]$ the following inequality holds:

$$\alpha \cdot \eta(a) + (1-\alpha) \cdot \eta(b) \leq \eta(\alpha \cdot a + (1-\alpha) \cdot b)$$

Setting $a = A_0$ and $b = t - B_0$, we obtain

$$\begin{aligned} f(A_0 + x) &\geq \alpha f(A_0) + (1-\alpha) f(t - B_0) \\ = f(A_0) - \frac{x}{t - A_0 - B_0} f(A_0) &+ \frac{x}{t - A_0 - B_0} f(t - B_0) \qquad (18) \\ = f(A_0) &+ \frac{x}{t - A_0 - B_0} f(t - B_0) - f(A_0)) \qquad (19) \end{aligned}$$

Rearranging yields inequality (15).

Proof of inequality (16)

We have $t > A_0 + B_0$. For $x \geq 0$ with $t \geq A_0 + B_0 + x$ we have $B_0 \leq t - A_0 - x \leq t - A_0$, hence $t - A_0 - x$ can be written as unique convex combination of the points B_0 and $t - A_0$. Therefore we need to solve the following:

$$t - A_0 - x = \alpha \cdot B_0 + (1-\alpha) \cdot (t - A_0)$$

A simple computation yields

$$\alpha = \frac{-x}{B_0 + A_0 - t} = \frac{x}{t - A_0 - B_0} \leq 1$$

Now we use that g is concave. Recall, that a function η is concave on $[a, b]$, if for every $\alpha \in [0, 1]$ the following inequality holds:

$$\alpha \cdot \eta(a) + (1-\alpha) \cdot \eta(b) \leq \eta(\alpha \cdot a + (1-\alpha) \cdot b)$$

Setting $a = B_0$ and $b = t - A_0$, we obtain

$$\begin{aligned} g(t - A_0 - x) &\geq \alpha g(B_0) + (1-\alpha) g(t - A_0) \\ = \frac{x}{t - A_0 - B_0} g(B_0) &+ \left(1 - \frac{x}{t - A_0 - B_0}\right) g(t - A_0) \qquad (20) \\ = g(t - A_0) &+ \frac{x}{t - A_0 - B_0} (g(B_0) - g(t - A_0)) \qquad (21) \end{aligned}$$

So we have

$$\frac{g(t - A_0 - x) - g(t - A_0)}{x} \geq \frac{g(B_0) - g(t - A_0)}{t - A_0 - B_0}$$

Multiplying by -1 yields inequality (16), which concludes the proof.

In other words, the convolution is obtained by taking the minimum of the two functions shifted such that their convex inflection points are at $t = A_0 + B_0$. This is also depicted in Figure 5.

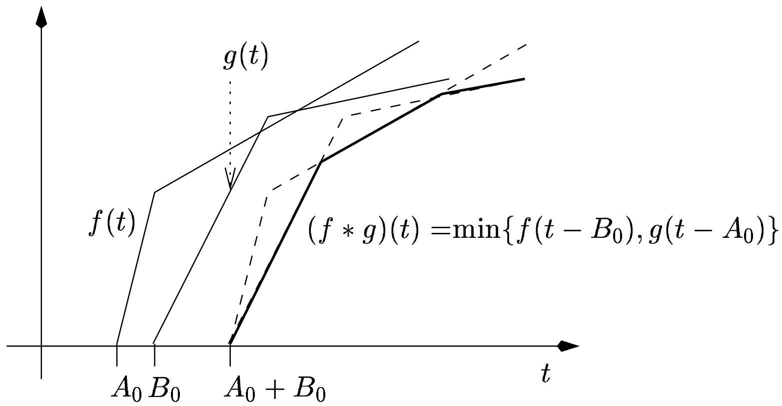

Figure 5: Application of Theorem 3

Note that the proof does not require f and g $(f, g \in \mathcal{F})$ to be piecewise linear and works for any function that has only one convex inflection point. Note that the convex inflection point has to be the first inflection point, but that is given for all functions out of $\mathcal{F}$. However, we only consider piecewise linear functions as all function in contemporary Network Calculus are piecewise linear. The slopes always denote some kind of rate. The piecewise linearity is also subject of the next remark.

Remark 2 The convolution of f and g is also piecewise linear and concave from $A_0 + B_0$ onwards.

Proof of Remark 2

Since f, g are piecewise linear and concave, $f(t - B_0)$ and $g(t - A_0)$ also posess these properties, since they evolve from f and g, respectively, by shifting along the x-axis. That the convolution is again piecewise linear is clear, since f and g are piecewise linear and *continuous.* It remains to show that the convolution is a concave function.
Let f, g be concave from a point R onwards. Set $p(t) := \min\{f(t), g(t)\}$. Then $p(t)$ is concave.
We need to show that for any given $x, y \in \mathbb{R}$ and $\lambda \in [0, 1]$ we have

$$p(\lambda x + (1 - \lambda)y) \quad \geq \quad \lambda p(x) + (1 - \lambda)p(y) \tag{22}$$

We assume w.l.o.g. $f(\lambda x + (1-\lambda)y) \leq g(\lambda x + (1-\lambda)y)$. Then we have

$$\begin{aligned}
p(\lambda x + (1-\lambda)y) &= \min\{f(\lambda x + (1-\lambda)y), g(\lambda x + (1-\lambda)y)\} \\
&= f(\lambda x + (1-\lambda)y) && (23) \\
&\geq \lambda f(x) + (1-\lambda) f(y) && (24) \\
&\geq \lambda \min\{f(x), g(x)\} + (1-\lambda)\min\{f(y), g(y)\} \\
&= \lambda p(x) + (1-\lambda) p(y)
\end{aligned}$$

Note that (23) holds because of our assumption and (24) uses the concavity of f.

Remark 3 Let f and g have m and n concave inflection points, respectively, i.e., $m+1$ and $n+1$ slopes greater than 0, respectively, The convolution of f and g then has $m+n$ concave inflection points and $m+n+1$ slopes greater than zero.

This leads to the following theorem.

Theorem 4 Given n functions $f_1, f_2, \ldots, f_n$, each with the properties of f. The min–plus convolutionof all of them is given by

$$(f_1 * f_2 * \ldots * f_n)(t) = \min_{j=1,2,\ldots,n} (f_j(t + A_{0,j} - \sum_{i=1}^{n} A_{0,i}))$$

Proof Utilizing the distributivity,

$$(f_1 * f_2 * \ldots * f_n)(t) = (((f_1 * f_2) * f3) * \ldots * f_n)(t)$$

This can be computed using Theorem 3 recursively, which yields the above term. Note that a function f_i is shifted by the sum of all A_0's other than the own one.

6 System Model

In the remainder of this paper, we apply the results obtained so far to prevailing network problems. We assume to be dealing with networks where a typical path is 6 hops and the number of nodes is approximately 50. Mission-critical, as well as semi-critical and uncritical data traffic traverse this network. Therefore, this network could be either a domain in the future Internet, where the mission-critical traffic would be some air traffic control, the semi-critical traffic is a Internet radio stream and the uncritical an Email. However, it could as well be a network in an automobile. There the mission-critical data flow is the brake-by-wire signal, the semi-critical a Video-on-Demand from the entertainment system and the uncritical again an Email. For critical flows the network service curve is given to reserve resources along the path. The scheduler is based on service curves. As mentioned in the related work, SCED [16] and HFSC [18] are the front runners.

We understand that this does not seem viable today, however, we do believe that if this theory is advanced further, the system theory concept will find its way into network design. We further do not believe that strict priority queueing is the ultimate answer to such a setup, even though it seems to be destined to be used when mission-critical flows are in the network.

All nodes are controlled by one central entity, which we call the Network Calculator. It has knowledge about the sum of all service curves each node is serving as well as the path of each source-destination pair. Note that it does not have to know the state of every flow. New requests are sent to this entity which checks whether it can be accommodated. If the request is not directly admissible, it checks whether it can be admitted by altering the service curves in the nodes preserving the delay guarantees. There are two turning knobs here, path optimal allocation and node optimal allocation of service curves. These two are discussed in the following two sections.

Making End-to-end guarantees requires the interconnection of domains. We neglect this and related issues as they would be out of scope of this paper.

It will turn out that shortcomings of nodes with respect to providing service can only be compensated by other nodes reducing their latency. Therefore we require the following definition.

Definition 7 [Compensated latency] We define *compensated latency* U_c as the time that has to be reduced from the initially allocated latency in order to meet the service curve requirement. Depending on the context, the reduction is done either the shortcoming node itself, or other node along the path, or a combination of nodes along the path.

Also depending on the context, the service curve requirement can be the local service curve of a node or the network service curve. To accommodate as many flows as possible in a node, some flows will not be allocated their initial service curve. Of course, this only works, if the new service curve gives better guarantees than the initial one. Therefore, we require the notion of a *dominating service curve.*

Definition 8 [Dominating Service Curve] We call $\beta_d(t)$ a *dominating service curve* over $\beta(t)$, if $\beta_d(t) \geq \beta(t)$ for all t.

In other words over any time $\beta_d(t)$ offers always at least the service that $\beta(t)$ does.

7 Path Optimal Allocation of Service Curves

This section deals with the relationship of the network service curve to the service curves of the individual nodes. We assume, a network service curve

s^{opt} has been obtained as shown in Section 3. Therefore, it is given by a 4-tuple (L, I, R, r). Now the question arises, which service curves the nodes on the path have to have so that the concatenation of these service curves yields the desired network service curve. We assume k nodes, where each node has a service curve s_i, with $i = 1 \ldots k$. The service curves s_i are also given by 4-tuples, namely (L_i, I_i, R_i, r_i). By using Theorem 4 we have

$$s^{opt} = \min_{j=1,2,\ldots,k} (s_j(t + L_j - \sum_{i=1}^{k} L_i))$$

First, we notice that the latency L of the network service curve is the sum of all L_i's of each node service curve along the path. Right of L, the slope of the network service curve will first be the minimum of all R_i's of all node service curves. In other words, if any node has an R_i higher than another one, it is wasted as only the minimum defines the network service curve. Therefore, it makes sense if all R_i's, U_i's and r_i's are equal. The setting of L_i's is a question of delay distribution along a path. As pointed out in Section 3, much work exists on this. An popular assumption is to distribute the delay equally among all nodes. Therefore, the optimal parameters for each node service curve are $(L/k, L/k + U, R, r)$. Trivially, if one node offers less resources the delay bound can not be met anymore without the other nodes raising their resources. However, if one node offered a higher rate R_i or r_i, or a longer period ΔU of sending at rate R, this would not have any effect on the network service curve and therefore be wasted resources.

The following numerical examples illustrate these findings. Assume a low bandwidth, short delay flow that suffices a TSpec with peak rate 9000 bytes/s, sustained rate 1000 bytes/s and buffer 2000 bytes. Without loss of generality we assume the maximum packet size is 0. Therefore we have, $TSpec(b, r, p, M) = (2000, 1000, 9000, 0)$. The flow traverses a path consisting of 5 nodes and has a delay bound of 500 ms. In other words, a flow can send at its peak rate for 250 ms and then has to reduce to sustained rate. Using the method from [17] the optimal network service curve is

$$s^{opt}_{netw} = \begin{cases} 0 & t \leq 500 \\ 9000 & 500 < t \leq 750 \\ 1000 & t > 750 \end{cases}$$

Assuming an equal delay distribution over all nodes, the optimal node service curves are for each node

$$s^{opt}_{node} = \begin{cases} 0 & t \leq 100 \\ 9000 & 100 < t \leq 350 \\ 1000 & t > 350 \end{cases}$$

We now discuss the consequences of a node not being able to offer the demanded service. There are 3 ways in which a node can fail to offer the

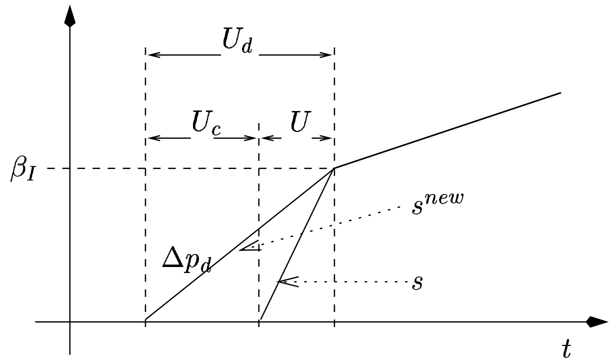

Figure 6: Compensated latency with deficient peak rate

demanded service: latency, peak rate and sustained rate. If a node can not offer the sustained rate, then the flow can not be accepted, as there is no way that the other nodes can make this up. The sustained rate of the network service curve is determined by the smallest sustained rate of all nodes. A node not making the delay requirement can be compensated rather easily. The time that it exceeds the delay requirement has to be made up by one or many of the other nodes guaranteeing a lower latency, such that the sum of the latencies of all nodes remains the required latency. To illustrate this with the example above, if one node can only offer a latency of 140 ms, then the network service curve can still achieved if the other nodes along the path save 40 ms. This can either be done by reducing the latency of one node to 60 ms or by the 4 other nodes having a latency or 90 ms each, or any other combination for which the sum of all latencies is 500 ms. The last possible failure is when a node is not able to offer the demanded peak rate. Here we distinguish whether the node is not able to offer the rate itself or whether is not able to offer the rate for the required time. A crucial quantity is β_I, which we defined as the amount of data that is served at the inflection point. Consider a node can only offer a deficient peak rate p_d. If it can offer it for a time U_d, such that $p_d U_d = \beta_I$, then the delay bound can be met by reducing the latency. The amount by which the latency has to be reduced, $U_d - U$, is obtained by the following consideration.

$$\beta_I = p_d U_d = pU \tag{25}$$

$$U_c = U_d - U = \beta_I (\frac{1}{p_d} - \frac{1}{p}) \tag{26}$$

Again, it is arbitrary which nodes make up the latency. This case is depicted in Figure 6.

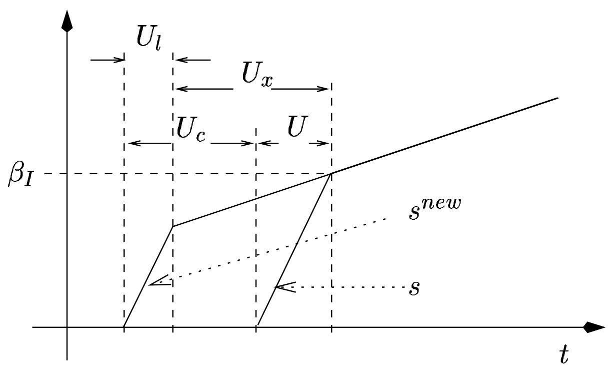

Figure 7: Compensated latency with deficient peak rate serving time

The next case is a node that can offer the peak rate only for a time $U_l < U$. The term U_x is the time needed to reach β_I after the peak rate stopped. The latency that has to be compensated is then $U_l + U_x - U$.

$$\beta_I = pU_l + rU_x = pU \tag{27}$$

$$U_c = U_l + U_x - U = \frac{(p-r)(\beta_I - pU_l)}{pr} \tag{28}$$

This case is depicted in Figure 7. It can be seen that if the ratio between peak rate p and sustained rate r is large, the compensated latency grows rapidly.

The bottom line is that no shortcoming of a node can be compensated by any other offering only a higher rate. All compensations rely on one or many other nodes guaranteeing a lower latency.

8 Node Optimal Allocation of Service Curves

We shift our attention now to one node, and apply our results to the optimization of the service curves. As mentioned in Section 6, new requests come in and it has to be decided whether the node can admit the flow. It is not relevant for the following concept whether this decision is taken by the Network Calculator or outsourced to the node itself. There exists a maximum service curve which denotes the capacity of a link. Without loss of generality, we assume this to be an LR curve, which we call the capacity service curve C. The sum of all prevailing service curves as well as the capacity service curve C are known. Figure 8 shows a capacity service curve,

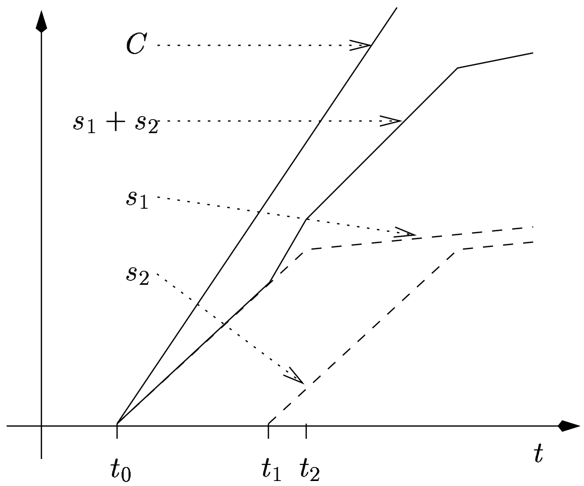

Figure 8: Not achievable service curves I

two individual service curves and their summed service curve. The first criterion for a flow to be admitted, is that the sum of all service curves does not exceed the capacity service curve. Remember that the service curve denotes how many packets have to be serviced by time t. In the case of a maximum service curve, it denotes how many packets can maximally be serviced. Obviously, if the number of packets that have to be service exceed the number that maximally can be serviced, the delay guarantee can not be made. The second criterion is that the rate in the summed service curve at any time t must never be greater than the rate at that time t in the capacity service curve. In Figure 8 the first criterion is met, while the second is violated between the times t_1 and t_2. This is much clearer to see in the derivatives of the service curves, which are given in Figure 9.

Consider the following numerical example. A node receives two requests. Both are low bandwidth short delay flows and suffice the same TSpec with peak rate 9000 bytes/s, sustained rate 1000 bytes/s and buffer 2000 bytes. Without loss of generality we assume the maximum packet size is 0. Therefore we have, $TSpec(b, r, p, M) = (2000, 1000, 9000, 0)$. In other words, a flow can send at its peak rate for 250 ms and then has to reduce to sustained rate. Flow 1 requires a delay bound from the node $d_{max} = 100$ ms. Flow 2 requires $d_{max} = 300$ ms. Let the available capacity of the router at the time the requests are made be the rate-latency curve with $L = 100$ ms and $R = 15000$ bytes/s. According to [17] the optimal service curve for Flow 1 is

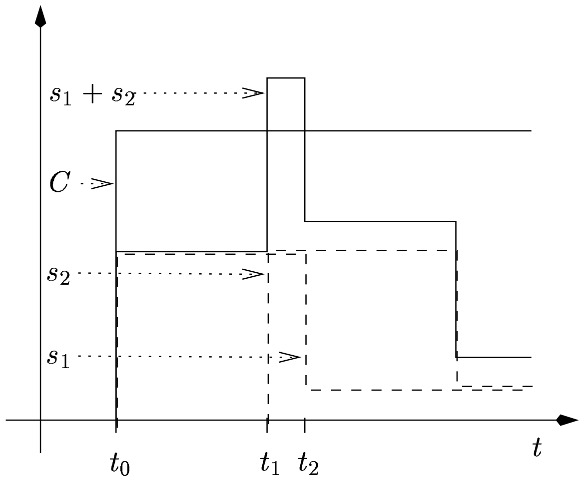

Figure 9: Not achievable service curves II

$$s_1^{opt} = \begin{cases} 0 & t \leq 100 \\ 9000 & 100 < t \leq 350 \\ 1000 & t > 350 \end{cases}$$

Similarly, the optimal service curve for Flow 2 is

$$s_2^{opt} = \begin{cases} 0 & t \leq 300 \\ 9000 & 300 < t \leq 550 \\ 1000 & t > 550 \end{cases}$$

Clearly, offhand both flows can not be accepted, as together they require a rate of $(9000{+}9000) bytes/s = 18000$ bytes/s in the interval $t \in [300ms, 350ms]$. But, using Theorem 3 Flow 1 can be assigned the service curve

$$s_1^{new} = \begin{cases} 0 & t \leq 100 \\ 11000 & 100 < t \leq 300 \\ 1000 & t > 300 \end{cases}$$

This allows both flows to be accommodated.

9 Conclusion and Outlook

In this paper we introduce some novel properties of the min–plus convolutionunder Network Calculus constraints. All functions relevant for Network Calculus known to us are piecewise linear and wide-sense increasing functions with only one convex kink. Explicitly, we show that the min–plus convolutionof functions with these characteristics can by efficiently computed by taking the minimum of the shifted functions. We apply this result to determine the optimal service curves of single nodes along a path in order to meet a network service curve. Therefore, our theorem brings insight on the relationship between the service curve of the single nodes and the network service curve. A key result is that no shortcoming of a node can be compensated by the other nodes merely increasing their rate. Rather, it is the latency of the other nodes that is crucial for compensation. With this work we move a step closer to allowing networks to be designed by service curves. Subject of future work will be to bring this theory closer to practice. I.e., derive an optimization algorithm that actually conducts the admission control. Further, we are working on implementing the system in a testbed under ALTQ [5]. On a final note, it remains a long way to bring a system theory for networking to the level of conventional system theory.

References

[1] F. Baccelli, G. Cohen, G.-J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity: An Algebra for Discrete Event Systems.* John Wiley and Sons, 1992.

[2] J.-Y. L. Boudec and P. Thiran. *Network Calculus.* Number 2050 in Lecture Notes in Computer Science. Springer-Verlag, 2001.

[3] A. Burchard, J. Liebeherr, and S. Patek. A Calculus for End-to-end Statistical Service Guarantees. Technical Report CS-2001-19, University of Virginia, Department of Computer Science, 2002.

[4] C.-S. Chang. *Performance Guarantees in Communication Networks.* Springer-Verlag, 2000.

[5] K. Cho. A Framework for Alternate Queueing: Towards Traffic Management by PC-UNIX Based Routers. In *Proceedings of USENIX 1998 Annual Technical Conference, New Orleans, USA*, June 1998.

[6] R. Cruz. A Calculus for Network Delay. I. Network Elements in Isolation. *IEEE Transactions on Information Theory*, 37(1):132–141, January 1991.

[7] R. Cruz. A Calculus for Network Delay. II. Network Analysis. *IEEE Transactions on Information Theory*, 37(1):114–131, January 1991.

[8] J.-B. Hiriart-Urruty and C. Lemarechal. *Fundamentals of Convex Analysis.* Springer-Verlag, 2001.

[9] H. Jaekel. *Ein Beitrag zur algebraischen Analyse verteilter und ereignisdiskreter Systeme.* Logos-Verlag, 2003. ISBN 3-8325-0238-6.

[10] J. Liebeherr, S. D. Patek, and E. Yilmaz. Tradeoffs in Designing Networks with End-to-End Statistical QoS Guarantees. In *Proceedings of the IEEE/IFIP Eighth International Workshop on Quality of Service (IWQoS '2000)*, pages 221–230, 2000.

[11] A. Oppenheim, A. Willsky, and S. Nawab. *Signals and Systems.* Prentice Hall, second edition, 1995. ISBN 0138147574.

[12] K. Pandit, C. Kirchner, J. Schmitt, and R. Steinmetz. Optimization of the Min-plus Convolution Computation under Network Calculus Constraints. Technical Report TR-KOM-2004-04, Technische Universit¨t Darmstadt, June 2004.

[13] K. Pandit, J. Schmitt, and R. Steinmetz. Network Calculus Meets Queueing Theory - A Simulation Based Approach to Bounded Queues. In *2004 Twelfth IEEE International Workshop on Quality of Service (IWQoS'04), Montreal, Canada*, pages 114–118, June 2004.

[14] W. Press, S. Teukolsky, W. Vellerling, and B. Flannery. *Numerical Recipes in C.* Cambridge University Press, second edition, 1992.

[15] S. Sandhu, K. Pandit, and A. Paulraj. On Non-Linear Space-Time Block Codes. In *Proceedings of 2002 International Conference on Acoustics, Speech, and Signal Processing, Orlando, USA*, volume 3, pages 2417 – 2420, May 2002.

[16] H. Sariowan, R. Cruz, and G. Polyzos. SCED: A Generalized Scheduling Policy for Guaranteed Quality of Service. *IEEE/ACM Transactions on Networking*, 7(5):669–684, October 1999.

[17] J. Schmitt. On the Allocation of Network Service Curves for Bandwidth/Delay-Decoupled Scheduling Disciplines. In *Proceedings of IEEE Global Telecommunications Conference 2002 (GLOBECOM 2002), Taipei, Taiwan.* IEEE, Nov. 2002.

[18] I. Stoica, H. Zhang, and T. S. E. Ng. Hierarchical Fair Service Curve Algorithm for Link-Sharing, Real-Time and Priority Service. *ACM Computer Communication Review*, 27(4), Oct. 1997. Proceedings of SIGCOMM'97 Conference.

[19] J. Turner. New Directions in Communications (or Which Way to the Information Age?). *IEEE Communications Magazine*, 24(10):8–15, Oct 1986.