A Transform for Network Calculus and its Application to Multimedia Networking

Krishna Pandit[a], Jens Schmitt[b], Claus Kirchner[b] and Ralf Steinmetz[a]

[a]Technische Universität Darmstadt, Germany
[b]Technische Universität Kaiserslautern, Germany

ABSTRACT

The rapid increase of multimedia traffic has to be accounted for when designing IP networks. A key characteristic of multimedia traffic is that it has strict Quality of Service (QoS) requirements in a heterogeneous manner. There are many different traffic types which have different throughput and delay requirements. In such a setting, scheduling by service curves is a useful method as it allows for assigning each flow exactly the service it requires. When hosting heterogeneous multimedia traffic, the utilization of packet-switched networks can be increased by using bandwidth/delay decoupled scheduling disciplines. It has been shown in previous work[1] how optimal network service curves are obtained with them, where optimal means that each multimedia flow receives the required service with the least possible consumption of resources. A basic result from Network Calculus is that the network service curve is obtained by the min–plus convolution of the node service curves. We state a theorem on the min–plus convolution in this work, which simplifies the computation of the min–plus convolution of service curves of bandwidth/delay decoupled schedulers. The rather complex min–plus convolution simplifies to merely shifting the functions and taking the minimum. The theorem follows from the continuous $\Gamma\Delta$-transform, which we develop. With this theorem, we derive the optimal service curves for the nodes along a path. Further, we show how the admission control can be improved when networks are designed based on service curves. Considering one node, reallocating the service curves leads to admitting more flows. Then we point out scenarios where sub-optimal allocation of service curves in a node can increase the number of admitted flows to the network. The key results are accompanied by numerical examples. On a broader scale, this paper advances the research in analytically modeling packet-switched networks by pointing out novel properties and a new application of Network Calculus.

Keywords: Quality of Service, Min-plus Algebra, Service Curve, Bandwidth-Delay Decoupled Scheduler

1. INTRODUCTION

1.1. Motivation

Future IP networks will carry heterogeneous traffic from a vast variety of application types, which includes several forms of multimedia. Some of the traffic will stem from multimedia applications which have hard Quality of Service (QoS) requirements. The typical QoS parameters of multimedia flows are throughput, delay and loss. Deterministic QoS guarantees can only be given, if the following three ingredients are enforced: traffic regulation, scheduling and admission control. Regulated traffic is best described with the arrival curve concept, which gives an upper bound of the traffic. Schedulers can be modeled with service curves. This is useful when considering heterogeneous multimedia traffic as it is possible to ensure each multimedia flow receives exactly the service it requires. The relationship between service curves and the actual scheduler is well studied. With the schedulers SCED from Sariowan and Cruz[2] as well as HFSC from Stoica[3] general service curves can be scheduled with sufficient efficiency. We are especially interested in service curves for bandwidth/delay decoupled scheduling

Please send correspondence to Krishna Pandit (Krishna.Pandit@kom.tu-darmstadt.de)

disciplines, as Schmitt[1] has shown that with those *optimal* network service curves are obtained. Optimal implies that the QoS requirement is met with the least possible consumption of resources. The challenge of admission control is to admit as many flows as possible while assuring that each flow still receives the contracted service. Optimal network service curves maximize the number of admissible multimedia flows requiring QoS. This is in the interest of the network service provider if he can charge multimedia flows separately, but also for marketing purposes. Network Calculus provides methods to compute performance bounds based on arrival and service curves. The knowledge of performance bounds is essential when offering deterministic QoS guarantees. A basic result of Network Calculus is that the network service curve is given by the min–plus convolution of the node service curves along the path. In this paper, we will show how to optimally design service curves for each node so that the network service curve is met. Therefore, this work generalizes the previously cited work of Schmitt,[1] which we review in the following subsection.

Beyond that, the remainder of this paper is organized as follows. In Section 2 we develop a novel transform for Network Calculus. This transform leads to our theorem for calculating the min–plus convolution of a certain kind of functions typical for Network Calculus. We introduce it and analyse its properties in Section 3. We then apply this theorem to admission control, first in Section 4 to optimally allocate service curves along a path. Then, in Section 5, we take the node perspective in order to perform admission optimization. Finally we conclude and give an outlook.

1.2. Background

For the following we assume the reader to be familiar with the *arrival curve* and *service curve* notions of Network Calculus, especially the TSpec, which is the Internet Engineering Task Force (IETF) arrival curve for multimedia flows, and the Latency-Rate (LR) scheduler. Further, we use the min–plus convolution as well as the theorems on concatenation of service curves and virtual delay bound. All of the above notions and results are excellently described in the textbook by Le Boudec and Thiran,[4] as well as introduced in our technical report.[5]

As mentioned above, the inspiration for this work is a result by Schmitt,[1] where he develops the optimal service curve for bandwidth/delay-decoupled scheduling disciplines. Such scheduling disciplines are well suited for multimedia networks, as these comprise traffic with diverse throughput and delay requirements. Each multimedia flow can then be assigned a service which offers exactly the requirements. The bandwidth/delay-decoupled scheduler is characterized by a non-linear service curve. The optimal network service curve for an arrival curve is derived, i.e., the network service curve which meets the QoS requirements while minimizing the resource consumption of a single flow. For a TSpec, which is the most widely spread arrival curve, the optimal service curve β^{opt} is that of a so-called *L2R scheduler*. The L2R scheduler has a latency and two rates. 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 equals the sustained rate of the arrival curve. Therefore, β^{opt} is originally 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. The inflection point I is chosen such that the delay bound is met at the point where the TSpec changes from the peak rate to the sustained rate. Further, we define $U = I - L$, which is the peak rate interval, i.e., the time that the flow is served at peak rate. In our considerations the peak rate interval is more important than the absolute inflection point, hence we will denote L2R service curves by the 4-tuple (L, U, R, r) hereafter. In Figure 1, a TSpec and its according optimal service curve β^{opt} are depicted. Note that in all figures Δ indicates that the following value belongs to the underlying slope. Schmitt further gives numerical examples which point out the benefits of this approach. In this work we extend those results to a wider perspective. 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 obtain optimal service curves.

1.3. Related Work

Much of the related work evolves around scheduling disciplines. This is well discussed by Schmitt.[1] Fidler and Recker[6] discuss the Fenchel transform in the context of Network Calculus. This is subject of Section 2.

2. A TRANSFORM FOR NETWORK CALCULUS

The min–plus convolution is 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. In conventional

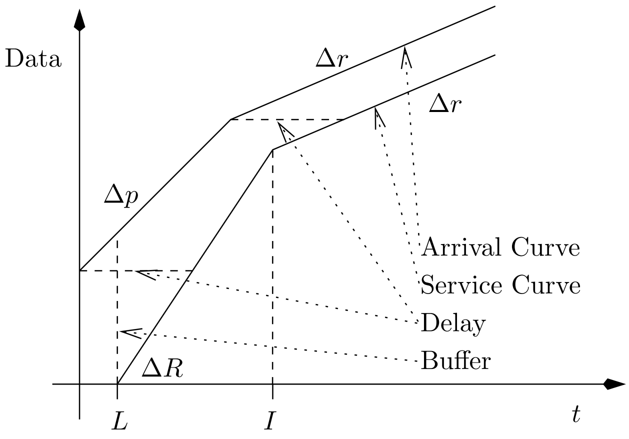

Figure 1. L2R scheduler

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 the legacy book by Press et al.,[7] 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.[8] it is briefly pointed out that this transform carries over the min–plus convolution in one domain to an addition in another domain. In the book by Hiriart-Urruty and Lemarechal[9] one finds the following definition for the convex conjugate function. The Fenchel transform, or convex conjugate function, is given by

$$f^*(s) = \sup\{sx - f(x) \mid x \in dom\, f\}. \tag{1}$$

Further, the bi-conjugate function is given by

$$f^{**}(x) = \sup\{sx - f^*(s) \mid s \in dom\, f^*\}. \tag{2}$$

I.e., the supremum of the term is taken for all x, where $f(x)$ is defined. 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 convolution is outlined in Chapter 3 in the book by Le Boudec and Thiran,[4] where it is pointed out that the min–plus convolution of 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 we have $f = f^{**}$. But in general, i.e., for non-convex functions, we cannot 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,[10] as well as the work by Fidler and Recker.[6] The latter further introduce the concave conjugate function, which is the equivalent of the Fenchel transform for concave functions, and present a graphical interpretation of some Network Calculus results. The next transform we consider is the $\Gamma\Delta$-transform. The $\Gamma\Delta$-transform can easily be represented graphically by so-called *point clouds*, This transform is explained in the book by Bacelli et al.,[8] and thoroughly described in the dissertation of Jäkel (in German),[11] whose notation we adopt. 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,

$$b_i = \sum_{k,t \in \mathbb{Z}} s_{x_i}(k,t)\delta^k\gamma^t = \sum_{k=-\infty}^{\infty} \sum_{t=-\infty}^{\infty} s_{x_i}(k,t)\delta^k\gamma^t. \tag{3}$$

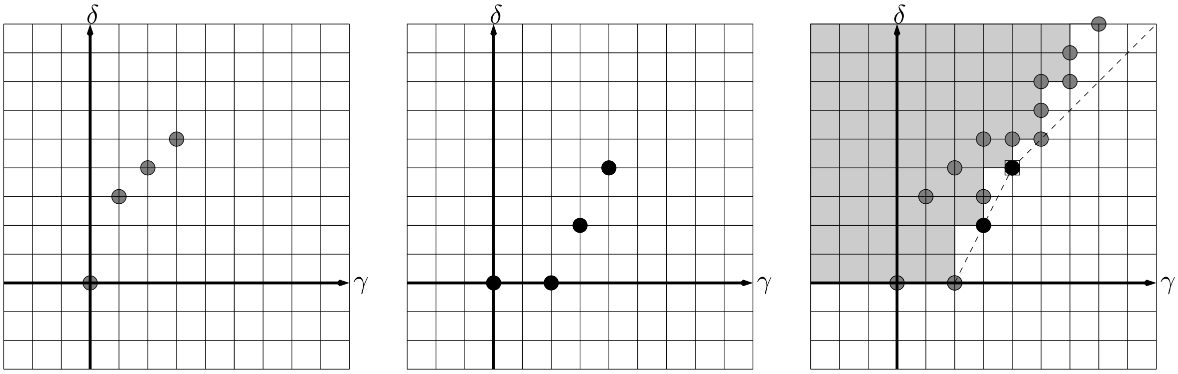

a) $s_{x_1}(k,t)$ b) $s_{x_2}(k,t)$ c) $(s_{x_1} \otimes s_{x_2})(k,t)$

Figure 2. Discrete $\Gamma\Delta$-transform

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

$$s_{x_i}(k,t) = \begin{cases} 0 & x_i(k) = t \\ +\infty & x_i(k) \neq t \end{cases} . \tag{4}$$

This operation is called the $\Gamma\Delta$-*transform*. For two elements $b_1, b_2 \in \mathbb{B}[[\gamma,\delta]]$ we define their sum (in the min-plus sense) component-wise as the usual minimum operation of their coefficients. Denote by $s_{x_i}(k,t), i = 1,2$ the coefficients of $b_i, i = 1,2$. Then the coefficients $(s_{x_1} \oplus s_{x_2})(k,t)$ of $b_1 \oplus b_2$ are given by

$$(s_{x_1} \oplus s_{x_2})(k,t) = \min_{k,t} [s_{x_1}(k,t), s_{x_2}(k,t)]. \tag{5}$$

The multiplication of two elements $b_1, b_2 \in \mathbb{B}[[\gamma,\delta]]$ is more involved. Again, let $s_{x_i}(k,t), i = 1,2$ denote the coefficients of $b_i, i = 1,2$. Then the coefficients $(s_{x_1} \otimes s_{x_2})(k,t)$ of $b_1 \otimes b_2$ are given by

$$(s_{x_1} \otimes s_{x_2})(k,t) = \min_{i_1+i_2=k, j_1+j_2=t} [s_{x_1}(i_1,j_1) + s_{x_2}(i_2,j_2)]. \tag{6}$$

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 convolution of two functions x_1 and x_2, we sample them (i.e., we obtain a sequence $x_1(k), x_2(k), k \in \mathbb{Z}$), apply the $\Gamma\Delta$-transform, add the resulting elements in $\mathbb{B}[[\gamma,\delta]]$ and apply the inverse $\Gamma\Delta$-transform. 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.

$$x_1(k) = \begin{cases} 0 & k = 0 \\ 2+k & 1 \le k \le 3 \\ +\infty & \text{otherwise} \end{cases} \tag{7}$$

The coefficients of the $\Gamma\Delta$-transform of $x_1(k)$ are 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.

$$x_2(k) = \begin{cases} 0 & k = 0 \\ 2(k-2) & 2 \le k \le 4 \\ +\infty & \text{otherwise} \end{cases} \tag{8}$$

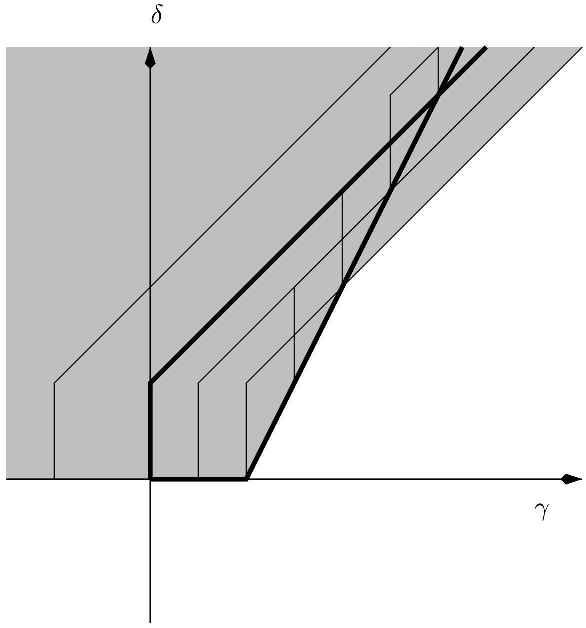

Figure 3. Continuous transform

The coefficients of the $\Gamma\Delta$-transform of $x_2(k)$ are depicted in Figure 2 b). In Figure 2 c) the multiplication operation $(s_{x_1} \otimes s_{x_2})(k,t)$ 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 convolution of two such service curves would have the switch from the slope from the service curve, to the slope from the arrival curve. For reference, the resulting min–plus convolution of the sampled curves is depicted by the dashed line. The reverse $\Gamma\Delta$-transform is obtained by taking the hull 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 convolution of piecewise linear functions.[10] Therefore, the $\Gamma\Delta$-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 approaches the min–plus convolution. Taking this to the limit, we end up at the *continuous* $\Gamma\Delta$-*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 $\Gamma\Delta$-transform. This is depicted in Figure 3. Underlying are an arrival curve, and a service curve. They are obtained when all their points are summed with each other, and it can be easily verified that the border of this denotes the min–plus convolution. With this in mind, we bring the attention back to Equation (6) 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.

3. MIN-PLUS CONVOLUTION UNDER NETWORK CALCULUS CONSTRAINTS

In this section we develop theorems on the min–plus convolution under Network Calculus constraints, which are beneficial for allocating resources by service curves. Note that all functions f, g in this section belong to $\mathcal{F}$, where we set

$$\mathcal{F} := \quad \{ f : \mathbb{R} \to \mathbb{R}_0^+ \mid \quad f(t) = 0 \text{ for } t < 0, f(t_1) \le f(t_2) \text{ for } t_1 \le t_2 \} \ . \tag{9}$$

i.e., $\mathcal{F}$ is the set of nonnegative wide-sense increasing functions. 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, an inflection point $Q \in \mathbb{R}$ is simply a point with $f'(Q^-) \neq f'(Q^+)$; i.e., the slope of f changes at Q.

Definition 3.1 (Convex and concave inflection points). *We define a convex inflection point Q as an inflection point, where the slope to its right is greater than the slope to its left. More mathematically, a convex inflection point Q has the property that $f'(Q^+) > f'(Q^-)$. 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 Q has the property that $f'(Q^+) < f'(Q^-)$*

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. Firstly, all functions are wide-sense increasing. Secondly, all functions are 0 for $t \leq T$ and *concave* for $t > T$. This is reasonable, as all underlying functions are arrival curves or 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 curves of bandwidth/delay decoupled schedulers from Schmitt,[1] which is the work we build upon, satisfy the second property.

Therefore, we are interested in the min–plus–convolution of special functions f and g. Let $i = 0, \ldots, m$, $j = 0, \ldots, n$, $A_i, B_j \in \mathbb{R}^+$ with $A_k > A_l$ and $B_k > B_l$ if $k > l$. Now we can state the functions f and g to be investigated in the rest of this section.

$$f(t) = \begin{cases} 0 & t \leq A_0 \\ a_0(t - A_0) & A_0 < t \leq A_1 \\ a_0(A_1 - A_0) + a_1(t - A_1) & A_1 < t \leq A_2 \\ \vdots & \\ \sum\limits_{i=0}^{m-2} a_i(A_{i+1} - A_i) + a_{m-1}(t - A_{m-1}) & A_{m-1} < t \leq A_m \\ \sum\limits_{i=0}^{m-1} a_i(A_{i+1} - A_i) + a_m(t - A_m) & t > A_m \end{cases} \tag{10}$$

$$g(t) = \begin{cases} 0 & t \leq B_0 \\ b_0(t - B_0) & B_0 < t \leq B_1 \\ b_0(B_1 - B_0) + b_1(t - B_1) & B_1 < t \leq B_2 \\ \vdots & \\ \sum\limits_{i=0}^{n-2} b_i(B_{i+1} - B_i) + b_{n-1}(t - B_{n-1}) & B_{n-1} < t \leq B_n \\ \sum\limits_{i=0}^{n-1} b_i(B_{i+1} - B_i) + b_n(t - B_n) & t > B_n \end{cases} \tag{11}$$

The functions f and g are to have exactly one convex inflection point each at A_0 and B_0, respectively, and only concave inflection points thereafter. I.e., we have $a_k < a_l$ and $b_k < b_l$ for $l < k$. In other words, we assume $a_m < a_{m-1} < \ldots < a_0$ and $b_n < b_{n-1} < \ldots < b_0$. The shape of f and g is given in Figure 4. Again, Δ indicates value of the slope.

We are now ready to formulate the theorem, which is the foundation of this work. It states, that 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 depicted in Figure 5.

Theorem 3.2. *The min–plus–convolution of f and g is given by*

$$(f * g)(t) = \inf_{0 \leq s \leq t} \{f(t - s) + g(s)\} = \min\{f(t - B_0), g(t - A_0)\}. \tag{12}$$

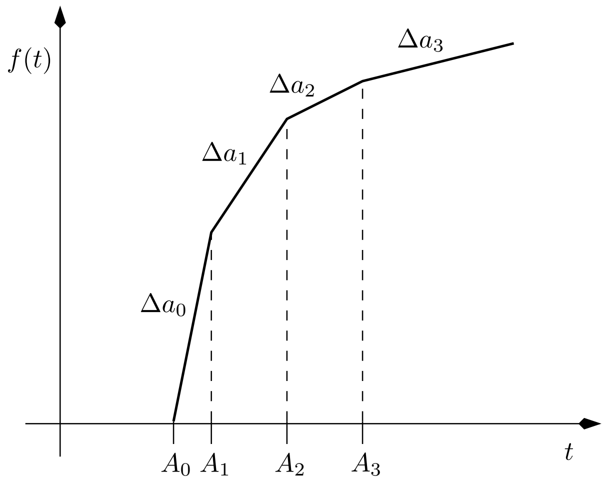

Figure 4. Shape of functions f and g

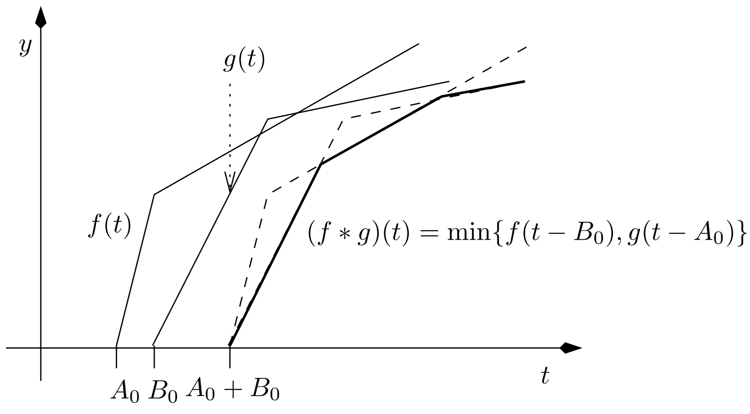

Figure 5. Application of Theorem 3.2

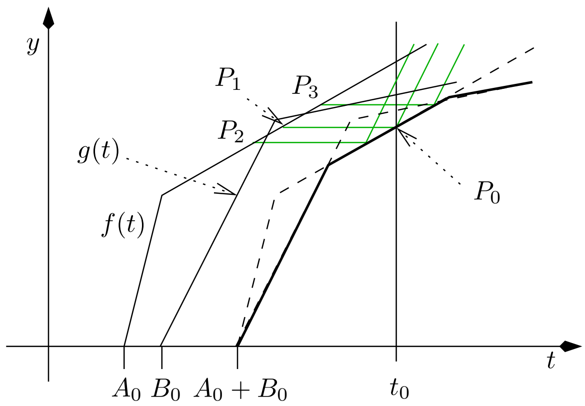

Figure 6. Proof of Theorem 3.2

Proof. While the complete proof is given in the technical report[5] we limit ourselves here to a graphical proof. As pointed out in the previous section, the min–plus convolution can be obtained by using the continuous $\Gamma\Delta$-transform. The continuous $\Gamma\Delta$-transform of two functions is given by adding all points of one function to all points of the other function. The reverse transform, i.e., taking the minimum over all t, then yields the min–plus convolution. We show now that the minimum of the shifted functions is always the overall minimum.

In Figure 6, all points of $g(t)$ (only the first few are shown) are added to different points of $f(t)$, namely P_1, P_2 and P_3. At time t_0, the lowest point is P_0, which is obtained when all points of $g(t)$ are added to the point P_1. The value at P_0 is

$$f(t_0 - B_0) + g(B_0) = f(t_0 - B_0). \tag{13}$$

It can be seen, that $f(t_0 - B_0 - \epsilon) + g(B_0 + \epsilon)$ yields a higher value at t_0 for any $\epsilon > 0$. The reason is that the slope of $g(t)$ between B_0 and $B_0 + \epsilon$ is greater than the slope of $f(t)$ between $t_0 - B_0$ and $t_0 - B_0 - \epsilon$. This is depicted where the points of $g(t)$ are added to point P_2. Similarly, $f(t_0 - B_0 + \epsilon) + g(B_0 - \epsilon)$ always yields a higher value at t_0 for any $\epsilon > 0$. The value for $f(t_0 - B_0 + \epsilon)$ is greater than $f(t_0 - B_0)$ and $g(B_0 - \epsilon)$ is always 0. This is depicted where the points of $g(t)$ are added to point P_3. Therefore, moving in both directions yields a higher value at t_0 which concludes the proof. $\square$

In this case, an alternative proof is to dissect the function f and g into a burst-delay component and a concave component and compute the convolution then. The advantage of the proof given here is that it is more general and easily extensible to functions with mixed concave and convex inflection points. We skip this since it would exceed the scope of this paper and is not required for our applications to multimedia networks, which we will show in the remainder of this paper. 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. However, we only consider piecewise linear functions as all functions in contemporary Network Calculus are piecewise linear. The slopes always denote some kind of rate. We make following two observations. The convolution of f and g is also piecewise linear and concave from $A_0 + B_0$ onwards. Further, the number of concave inflection points is well defined. 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. These two observations lead to the following theorem.

Theorem 3.3. *Given n functions $f_1, f_2, \ldots, f_n$, each with the properties of f, the min–plus convolution of all of them is given by*

$$(f_1 * f_2 * \ldots * f_n)(t) = \min_{j=1,2,\ldots,n} (f_j(t + A_{0,j} - \sum_{i=1}^{n} A_{0,i})). \tag{14}$$

Proof. Utilizing the distributivity,

$$(f_1 * f_2 * \ldots * f_n)(t) = (((f_1 * f_2) * f_3) * \ldots * f_n)(t). \tag{15}$$

This can be computed using Theorem 3.2 recursively, which yields the above term. This is equivalent to shifting all functions f_i by the sum of all A_0's other than the own one and taking the minimum. ◻

In the remainder of this paper we apply the results from this section to resource allocation in networks.

4. ALLOCATION OF SERVICE CURVES ALONG A PATH

This section deals with the relationship of the network service curve to the node service curves. The goal is to meet the QoS requirements, in particular the delay bound, of a multimedia flow while allocating as little resources to it as possible.

4.1. Determining optimal node service curves

With the theorems of the previous section we have established a relationship between functions and their min–plus convolution. Recall that the network service curve is the concatenation, i.e., the min–plus convolution, of the node service curves. Therefore, we can use our theorem to answer the question, which service curves the nodes on the path must have, so that the concatenation of these service curves yields the network service curve desired for the multimedia flow.

Assume a network service curve β^{opt} has been obtained as shown in Section 1.2. Further, assume k nodes, where each node has a service curve β_j, with $j = 1 \ldots k$. Recalling Theorem 3.3 we have

$$\beta^{opt} = \min_{j=1,2,\ldots,k}(\beta_j(t + L_j - \sum_{i=1}^{k} L_i)). \tag{16}$$

It can be seen that the latency L of the network service curve is the sum of all L_j's of all node service curves. To the right of L, the slope of the network service curve will first be the minimum of all R_j's of all node service curves. In other words, if any node has an R_j higher than another one, it is wasted as only the minimum defines the network service curve. This property is depicted in Figure 5. Imagine $f(t)$ and $g(t)$ to be service curves. The area between the dashed lines and the solid line denoting the convolution are wasted resources. The same rationale applies for all times, therefore it is beneficial that all node service curves are equal after their convex inflection point. Therefore, optimal node service curves are given by 4-tuples (L_j, U_j, R_j, r_j), where all R_j's, U_j's and r_j's are equal. Setting the latencies L_j is a question of delay distribution along a path, which we will not discuss further. Trivially, if one node offers less resources the delay bound cannot 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 peak rate interval U, this would not have any effect on the network service curve and therefore be wasted resources, i.e., prevent more multimedia flows to be admitted.

4.2. Numerical Example

The following numerical examples illustrate these findings. Assume a low bandwidth, short delay flow that satisfies a TSpec with peak rate 9000 bytes/s, sustained rate 1000 bytes/s and buffer 2000 bytes. Without loss of generality we assume a fluid model, i.e., $M = 0$. Hence we have

$$TSpec_1(b, r, p, M) = (2000 \text{ bytes}, 1000 \text{ bytes/s}, 9000 \text{ bytes/s}, 0). \tag{17}$$

In other words, a flow can send at its peak rate for 250 ms and then has to reduce to sustained rate. The flow traverses a path consisting of 5 nodes and has a delay bound of 500 ms. Using the method from Schmitt[1] the optimal network service curve is the 4-tuple

$$\beta^{opt}_{net} = (L, U, R, r) = (500 \text{ ms}, 250 \text{ ms}, 9000 \text{ bytes/s}, 1000 \text{ bytes/s}). \tag{18}$$

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

$$\beta^{opt}_{node} = (L, U, R, r) = (100 \text{ ms}, 250 \text{ ms}, 9000 \text{ bytes/s}, 1000 \text{ bytes/s}). \tag{19}$$

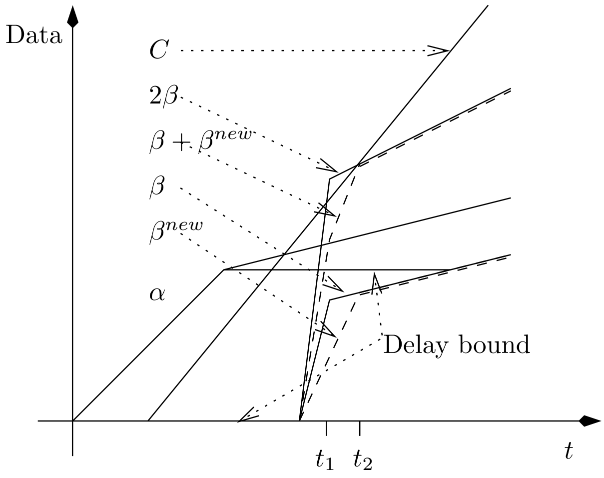

Figure 7. Local Reallocation

5. REALLOCATION OF SERVICE CURVES IN NODES

5.1. Local Reallocation

In this subsection we take a perspective of a single node. There are instances where altering the service curves within a node can lead to increasing the number of multimedia flows admitted. The rationale is explained by a numerical example and depicted in Figure 7.

There exists a maximum service curve C which denotes the capacity of a node. Without loss of generality, we assume this to be an LR service curve, which we call the capacity service curve C. Assume a node has two requests of flows, each with an arrival curve α in form of the following TSpec.

$$TSpec_2(b, r, p, M) = (1500 \text{ bytes}, 2500 \text{ bytes/s}, 10000 \text{ bytes/s}, 0).$$

A service curve β is chosen according to the method of Schmitt[1] such that the delay bound is met at the two crucial points, namely the kinks of the arrival curve.

$$\beta = (L, U, R, r) = (300 \text{ ms}, 40 \text{ ms}, 40000 \text{ bytes/s}, 2500 \text{ bytes/s}) \tag{20}$$

Per se they can not be accepted as the sum of the service curves, 2β, exceeds the C between t_1 and t_2. However, if a service curve β^{new} can be found such that the delay bound remains met at the two crucial points and the sum $\beta + \beta^{new}$ never exceeds the capacity service curve C, both flows can be accommodated. For this example, a possible β^{new} is given by

$$\beta^{new} = (L, U, R, r) = (300 \text{ ms}, 80 \text{ ms}, 21250 \text{ bytes/s}, 2500 \text{ bytes/s}). \tag{21}$$

From Theorem 3.2 we can be sure that the network service curve is not affected if one or more nodes apply this method.

5.2. Global Reallocation

We next discuss how deficiencies of one node can be compensated by other nodes. We require the following definition.

Definition 5.1 (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*

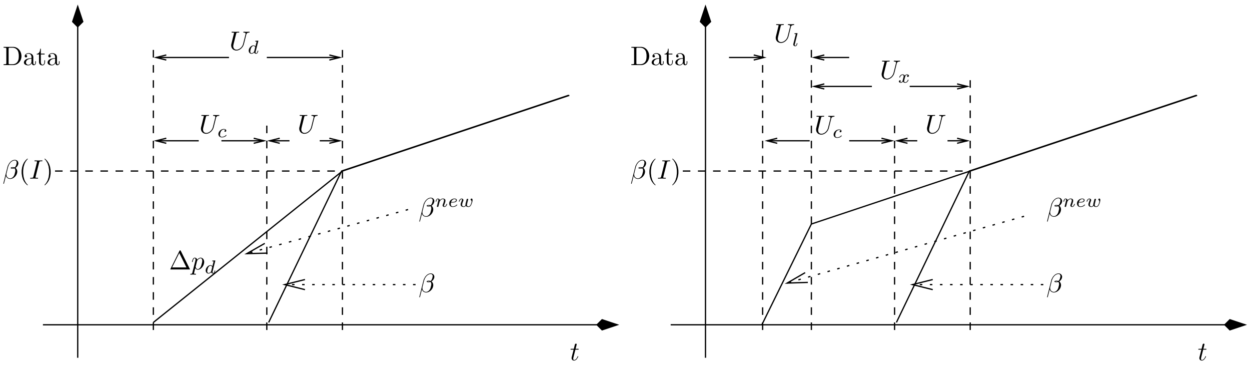

Figure 8. Compensated latency with deficient peak rate (left) and with deficient peak rate serving time (right)

context, the reduction is done either by the shortcoming node itself, or one other node along the path, or a combination of nodes along the path.

There are 3 ways in which a node can fail to offer the demanded service and therefore be forced to reject a request: latency, peak rate and sustained rate. If a node cannot offer the sustained rate, then the flow cannot be accepted. There is no way that the other nodes can compensate this, as 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. This corresponds directly to the shift of Theorem 3.3. To illustrate this with the example of Section 4.2, if one node can only offer a latency of 140 ms, then the network service curve can still be 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 of 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 it is not able to offer the rate for the required time. A crucial quantity is $\beta(I)$, which we define as the amount of data that is served at the inflection point. It follows from Theorem 3.2 that having less data served at the inflection point would decrease the minimum of the shifted functions. This would destroy the network service curve as the inflection point is crucial for the delay bound.

Consider a node that 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.

$$\beta(I) = p_d U_d = pU \tag{22}$$

$$U_c = U_d - U = \beta(I)\left(\frac{1}{p_d} - \frac{1}{p}\right) \tag{23}$$

Again, it is arbitrary which nodes make up the latency. This case is depicted in the left graph of Figure 8, where Δ again labels the slope.

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 $\beta(I)$ after the peak rate stopped. The latency that has to be compensated is then $U_l + U_x - U$.

$$\beta(I) = pU_l + rU_x = pU \tag{24}$$

$$U_c = U_l + U_x - U = \frac{(p-r)(\beta(I) - pU_l)}{pr} \tag{25}$$

This case is depicted in the right graph of Figure 8. It can be seen that if the ratio between peak rate p and sustained rate r is large, the compensated latency grows rapidly. Note that the service curves β^{new} in this section are not optimal, in the sense that they guarantee a lower delay for the first packets than required. This does not improve the overall performance as the worst-case delay, which is at the inflection point, remains untouched. Therefore, some resources are wasted. The bottom line is that no shortcoming of a node can be compensated by any other offering only a higher rate. All compensations require other nodes guaranteeing a lower latency.

6. CONCLUSION AND OUTLOOK

In this paper we developed the continuous $\Gamma\Delta$-transform and devised a theorem on the computation of the min–plus convolution under Network Calculus constraints. These results advance the research in Network Calculus, which is a well-suited tool for modeling multimedia networks. Beyond that our results were applied to assist admission control decisions in networks with heterogeneous multimedia traffic. With the theorem, optimal service curves of single nodes along a path in order to meet a network service curve were determined. Further, we showed how shortcomings of nodes can be compensated locally and globally to improve admission control. Subject of future work will be to bring this theory closer to practice, i.e., to develop algorithms that actually conduct the admission control. Finally, such a system is to be implemented in a testbed under ALTQ.[12]

REFERENCES

1. 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.
2. H. Sariowan, R. Cruz, and G. Polyzos, "SCED: A Generalized Scheduling Policy for Guaranteed Quality of Service," *IEEE/ACM Transactions on Networking* **7**, pp. 669–684, October 1999.
3. 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**, Oct. 1997. Proceedings of SIGCOMM'97 Conference.
4. J.-Y. L. Boudec and P. Thiran, *Network Calculus*, no. 2050 in Lecture Notes in Computer Science, Springer-Verlag, 2001.
5. K. Pandit, J. Schmitt, C. Kirchner, and R. Steinmetz, "Optimal Allocation of Service Curves by Exploiting Properties of the Min-plus Convolution," Tech. Rep. TR-KOM-2004-08, Technische Universiät Darmstadt, Nov. 2004.
6. M. Fidler and S. Recker, "Conjugate Network Calculus: A Dual Approach Applying the Legendre Transform," *Elsevier Computer Networks Journal* , 2005. Accepted for publication.
7. W. Press, S. Teukolsky, W. Vellerling, and B. Flannery, *Numerical Recipes in C*, Cambridge University Press, second ed., 1992.
8. 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.
9. J.-B. Hiriart-Urruty and C. Lemarechal, *Fundamentals of Convex Analysis*, Springer-Verlag, 2001.
10. K. Pandit, C. Kirchner, J. Schmitt, and R. Steinmetz, "Optimization of the Min-plus Convolution Computation under Network Calculus Constraints," Tech. Rep. TR-KOM-2004-04, Technische Universität Darmstadt, June 2004.
11. H. Jaekel, *Ein Beitrag zur algebraischen Analyse verteilter und ereignisdiskreter Systeme*, Logos-Verlag, 2003. ISBN 3-8325-0238-6.
12. 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.