On the Impacts of Traffic Shaping on End-to-End Delay Bounds in Aggregate Scheduling Networks

Markus Fidler

Department of Computer Science, Aachen University
Ahornstr. 55, 52074 Aachen, Germany
fidler@i4.informatik.rwth-aachen.de

Abstract. The Differentiated Services architecture allows for the provision of scalable Quality of Service by means of aggregating flows to a small number of traffic classes. Among these classes a Premium Service is defined, for which end-to-end delay guarantees are of particular interest. However, in aggregate scheduling networks the derivation of such worst case delays is significantly complicated and the derived bounds are weakened by the multiplexing of flows to aggregates.
A means to minimize the impacts of interfering flows is to shape incoming traffic, so that bursts are smoothed. Doing so reduces the queuing delay within the core of the domain, whereas an additional shaping delay at the edge is introduced. In this paper we address the issue of traffic shaping analytically. We derive a form that allows to quantify the impacts of shaping and we show simulation results on the derivation of end-to-end delay bounds under different shaping options.

1 Introduction

Differentiated Services (DS) [2] addresses the scalability problems of the former Integrated Services approach by an aggregation of flows to a small number of traffic classes. Packets are identified by simple markings that indicate the respective class. In the core of the network, routers do not need to determine to which flow a packet belongs, only which aggregate behavior has to be applied. Edge routers mark packets and indicate whether they are within profile or, if they are out of profile, in which case they might even be discarded at the edge router. A particular marking on a packet indicates a so-called Per Hop Behavior (PHB) that has to be applied for forwarding of the packet. The Expedited Forwarding (EF) PHB [8] is intended for building a service that offers low loss and low delay, namely a Premium Service. For this purpose delay bounds are derived for a general topology and a defined load in [4]. However, these bounds can be improved, when additional information concerning the current load and the special topology of a certain DS domain is available.

In [12] a resource manager for DS domains called Bandwidth Broker (BB) is conceptualized. The task of a BB in a DS domain is to perform a careful admission control and to set up the appropriate configuration of the domain's edge routers. While doing so, the BB knows about all requests for resources of certain Quality of Service (QoS) classes. Besides it can learn about the domain's topology by implementing a routing protocol listener. Thus, a BB can have access to all information that is required, to base the admission control on delay bounds that are derived for individual flows, for the current load, and for the actual mapping of flows onto the topology of the administrated domain, applying the mathematical methodology of Network Calculus [15].

In this paper we investigate the impacts of traffic shaping on end-to-end delay bounds. The remainder of this paper is organized as follows: In section 2 the required background on Network Calculus is given. Section 3 addresses the impacts of traffic shaping. Related simulation results are given in section 4. Section 5 concludes the paper. Proofs are given in the appendix.

2 Network Calculus

Network Calculus is a theory of deterministic queuing systems based on the calculus for network delay presented in [5, 6] and on Generalized Processor Sharing in [13, 14]. Extensions, and a comprehensive overview on current Network Calculus are given in [3, 11]. Here only a few concepts are covered briefly. In the sequel upper indices j indicate links and lower indices i indicate flows.

Flows can be described by arrival functions $F(t)$ that are given as the cumulated number of bits seen in an interval $[0, t]$. Arrival curves $\alpha(t)$ are defined to give an upper bound on the arrival functions with $\alpha(t_2 - t_1) \geq F(t_2) - F(t_1)$ for all $t_2 \geq t_1 \geq 0$. In DS networks, a typical constraint for incoming flows can be given by the affine arrival curve $\alpha_{r,b}(t) = b + r \cdot t$. Usually the ingress router of a DS domain meters incoming flows against a leaky bucket algorithm, which allows for bursts b and a rate r. Non-conforming traffic is either shaped or dropped.

The service that is offered by the scheduler on an outgoing link can be characterized by a minimum service curve, denoted by $\beta(t)$. A special characteristic of a service curve is the rate-latency type that is given by $\beta_{R,T}(t) = R \cdot [t - T]^+$, with a rate R and a latency T. The term $[x]^+$ is equal to x, if $x \geq 0$, and zero otherwise. Service curves of the rate-latency type are implemented for example by Priority Queuing (PQ) or Weighted Fair Queuing (WFQ).

Based on the above concepts, bounds for the backlog, the delay, and for the output flow can be derived. If a flow i that is constrained by α_i^j is input to a link j that offers a service curve β^j, the output arrival curve α_i^{j+1} of flow i can be given by (1).

$$\alpha_i^{j+1}(t) = \sup_{s \geq 0}[\alpha_i^j(t+s) - \beta^j(s)] \tag{1}$$

Multiplexing of flows can simply be described by addition of the belonging arrival functions, respective arrival curves. For aggregate scheduling networks with FIFO service elements, families of per-flow service curves $\beta_\theta^j(t)$ according to (2), with an arbitrary parameter $\theta \geq 0$ are derived in [7, 11]. $\beta_\theta^j(t)$ gives a family of service curves for a flow 1 that is scheduled in an aggregate manner with a flow, or a sub-aggregate 2 on a link j. The term $1_{t>\theta}$ is zero for $t \leq \theta$.

$$\beta_\theta^j(t) = [\beta^j(t) - \alpha_2(t - \theta)]^+ 1_{t>\theta} \tag{2}$$

3 Traffic Shaping

A means to reduce the impacts of interfering bursts on network performance is to shape incoming traffic at the edge of a domain [12]. Queuing of the initial bursts is in this case performed at the edge, with the aim to minimize the delay within the core. Especially, if heterogeneous aggregates have to be scheduled, shaping allows to reduce impacts of different types of flows on each other [15]. However, to our knowledge the work on shapers in [10] has not been extended to aggregate scheduling networks and an analysis of the effects of shaping on the derivation of end-to-end delay bounds is missing in current literature.

Theorem 1 (Bound for Output, General Case) *Consider two flows* 1*, and* 2 *that are* α_1^j*, respective* α_2^j *upper constrained. Assume these flows are served in FIFO order and in an aggregate manner by a node* j *that is characterized by a minimum service curve of the rate-latency type* $\beta_{R,T}^j$*. Then, the output of flow* 1 *is* α_1^{j+1} *upper constrained according to (3), where* θ *is a function of* t *and has to comply with (4).*

$$\alpha_1^{j+1}(t) = \alpha_1^j(t+\theta) \tag{3}$$

$$\theta(t) = \frac{\sup_{v>0}[\alpha_1^j(v+t+\theta(t)) - \alpha_1^j(t+\theta(t)) + \alpha_2^j(v) - R^j \cdot v]}{R^j} + T^j \tag{4}$$

Corollary 1 (Bound for Output, Single Leaky Bucket) *In case of a leaky bucket constrained flow* 1*, with rate* r_1 *and burst size* b_1^j*, (4) can be simplified applying* $\alpha_1^j(v+t+\theta(t)) - \alpha_1^j(t+\theta(t)) = r_1 \cdot v$*. As an immediate consequence,* θ *becomes independent of* t*. With (3) we find that the output flow* 1 *is leaky bucket constrained with* r_1 *and* $b_1^{j+1} = \alpha_1(\theta)$*. Further on, if the flow or sub-aggregate* 2 *is leaky bucket constrained with rate* r_2 *and burst size* b_2^j*, the* $\sup[\dots]$ *in (4) is found for* $v \to 0$ *resulting in* $\theta = b_2^j/R^j + T^j$ *and* $b_1^{j+1} = b_1^j + r_1 \cdot (T^j + b_2^j/R^j)$*.*

Definition 1 (Sustainable Rate Shaping) *Assume a flow that is leaky bucket constrained with the parameters* $(r_1, \bar{b}_1)$*, where* r_1 *is called the sustainable rate. If this flow is input to a traffic shaper that consists of a bit-by-bit system with a shaping rate* r_1*, and a packetizer with a maximum packet size* $l_{\max}$*, the output flow is constrained by* $(r_1, b_1 = l_{\max})$ *[10]. Further on, the shaper adds a worst-case delay of* $\bar{b}_1/r_1$*.*

In [11] it is shown that shaping at the sustainable rate does not worsen the end-to-end delay bounds in Integrated Services networks, if the reserved rate matches the shaping rate, and in turn the rate of the flow. However, this assumption does not hold true for DS domains. DS Premium resources are intended to be reserved statically and PQ is a likely means of implementation. Thus, the allocated Premium capacity usually exceeds the capacity that is requested by Premium traffic sources. In this scenario shaping at the sustainable rate can significantly worsen delay bounds, whereas Premium bursts that are not shaped can result in unwanted interference and increase queuing delays in the core of the domain. Hence, adaptivity when setting the shaping rates is required.

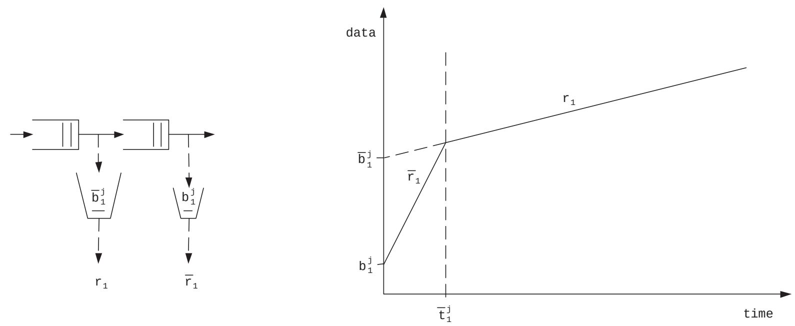

Fig. 1. Two Leaky Bucket Constraint

Definition 2 (Two Leaky Bucket Constraint) *Consider a two leaky bucket configuration according to figure 1. Define* $(r_1, \overline{b}_1^j)$*, and* $(\overline{r}_1, b_1^j)$ *to be the parameters of the first, respective second leaky bucket, with* $\overline{r}_1 > r_1$ *and* $\overline{b}_1^j > b_1^j$*. The resulting arrival curve is defined in (5). It allows for bursts of size* b_1^j*, then it ascends by* $\overline{r}_1$ *until* $\overline{t}_1^j = (\overline{b}_1^j - b_1^j)/(\overline{r}_1 - r_1)$*, and finally it increases with rate* r_1*.*

$$\alpha_1^j(t) = \min[b_1^j + \overline{r}_1 \cdot t, \overline{b}_1^j + r_1 \cdot t] = \begin{cases} b_1^j + \overline{r}_1 \cdot t & \text{, if } t \leq \overline{t}_1^j = \frac{\overline{b}_1^j - b_1^j}{\overline{r}_1 - r_1} \\ \overline{b}_1^j + r_1 \cdot t & \text{, else} \end{cases} \quad (5)$$

An arrival curve of the type in (5) can be given, if a leaky bucket constrained flow with the arrival curve $\alpha_1^j(t) = \overline{b}_1^j + r_1 \cdot t$ traverses a combination of a bit-by-bit traffic shaper with rate $\overline{r}_1$ and a packetizer with a maximum packet size $l_{\max}$ [10]. Then, the output arrival curve is two leaky bucket constrained with the parameters $(r_1, \overline{b}_1^j)$ and $(\overline{r}_1, l_{\max})$. The shaper adds a worst-case delay of $\overline{b}_1^j/\overline{r}_1$.

Theorem 2 (Bound for Output, Two Leaky Bucket) *Consider two flows 1 and 2 that are* α_1^j*, respective* α_2^j *upper constrained. Assume that these flows are served in FIFO order and in an aggregate manner by a node* j *that is characterized by a minimum service curve of the rate-latency type* β_{R^j,T^j}^j*. If the input flow 1 is constrained by two leaky buckets with* $(r_1, \overline{b}_1^j)$*,* $(\overline{r}_1, b_1^j)$*, and* $\overline{t}_1^j = (\overline{b}_1^j - b_1^j)/(\overline{r}_1 - r_1)$*, the output flow is two leaky bucket constrained with* $(r_1, \overline{b}_1^{j+1})$*, and* $(\overline{r}_1, b_1^{j+1})$*, where* $b_1^{j+1} = \alpha_1^j(\theta(0))$ *and* $\overline{b}_1^{j+1} = \overline{b}_1^j + r_1 \cdot \theta(\overline{t}_1^j)$*.*

Definition 3 (Minimum Interference Shaping) *We define minimum interference shaping to be a configuration, where all flows* i *with* $(r_i, \overline{b}_i^j)$ *of the set of flows* $\mathbb{I}$ *that form an aggregate are shaped with a rate* $\overline{r}_i$*, so that* $\sum_{i=\mathbb{I}_j} \overline{r}_i \leq R^j$ *holds on all links* j *of the set of links* $\mathbb{J}$ *of the domain, where* $\mathbb{I}_j$ *is the set of flows* i *that traverse a link* j*. Thus, flows are constrained by* $(r_i, \overline{b}_i)$ *and* $(\overline{r}_i, b_i = l_{\max})$*.*

For the settings given in definition 3 the sup[...] in (4) is found on all links $j \in \mathbb{J}$ for any $t \geq 0$ with $v \to 0$, whereby θ is constant over time with $\theta = b_2^j/R^j + T^j$. Hence, the impact of interfering flows is reduced to the impact of their effective burst size after shaping. The output constraint of the flow of interest 1 that is scheduled in an aggregate manner with a flow, or a sub-aggregate 2 on a link j is given by the parameters $(\overline{r}_1, b_1^{j+1} = b_1^j + \overline{r}_1 \cdot (b_2^j/R^j + T^j))$, and $(r_1, \overline{b}_1^{j+1} = \overline{b}_1^j + r_1 \cdot (b_2^j/R^j + T^j))$, with $\overline{t}_1^{j+1} = (\overline{b}_1^j - b_1^j)/(\overline{r}_1 - r_1) - (b_2^j/R^j + T^j)$. If $\overline{t}_1^{j+1} \leq 0$, the output constraint is reduced to a single leaky bucket constraint.

For over-provisioned links minimum interference shaping allows for a variety of settings of the per-flow shaping rates $\overline{r}_i$. However, the use of high priority traffic classes, such as a PQ-based Premium class, can lead to starvation of other services including the Best-Effort (BE) Service. Thus, it is reasonable to limit the Premium burst size by shaping and to restrict the overall Premium rate, as is supported by various router implementations. Here, we define a parameter d_q to give an upper bound on the allowed queuing delay of Premium traffic within the core of the network, from which an upper bound of the Premium burst size can be derived. To set up corresponding shaping rates $\overline{r}_i$, we apply a two step approach. The maximum allowed shaping delay d_{s_i} is computed as $d_{s_i} = d_{r_i} - (d_{t_i} + d_q)$ with d_{r_i} denoting the requested maximum delay for flow i, and d_{t_i} giving the end-to-end propagation delay on the corresponding path. If $d_{s_i} > 0$, the corresponding shaping rate follows as $\overline{r}_i = \max[r_i, \overline{b}_i/d_{s_i}]$, otherwise the target delay cannot be guaranteed. Then, if still all of the conditions in definition 3 hold, and, if the queuing delay within the core can be derived by Network Calculus to be smaller than d_q for all flows $i \in \mathbb{I}$, a solution is found. Note that the configuration of the shapers does not have to be updated during the lifetime of the corresponding flows, since all shaping rates $\overline{r}_i$ are set to account for queuing delay of at most d_q. Thus, the approach scales similar to sustainable rate shaping. If the rate of the Premium traffic shall be restricted in addition, the conditions in definition 3 have to be replaced by stricter ones.

4 Evaluation Results

We implemented an admission control for an application within the framework of a Bandwidth Broker [15]. The admission control currently knows about the topology of the domain statically, whereas a routing protocol listener can be added. Requests for Premium capacity are signalled in a Remote Procedure Call (RPC) style. The requests consist of a Committed Information Rate (CIR), a Committed Burst Size (CBS), and a target maximum delay. Whenever the admission control receives a new request, it computes the end-to-end delay for all requests that are active concurrently. If none of the target maximum per-flow delays is violated, the new request is accepted, which otherwise is rejected. Note that requests are usually made for aggregated traffic flows that use the same path from the ingress to the egress router to allow for scalability.

For performance evaluation a simulator that generates such Premium resource requests is used. Sources and sinks are chosen uniformly from a predefined set. Start and end times are modelled as negative exponentially distributed with a mean inter-arrival time $1/\lambda$ and a mean service time $1/\mu$. Thus, $\rho = \lambda/\mu$ can be defined as the network load, that is the mean number of concurrently active requests. The target delay, CIR, and CBS are either used as parameters or as uniformly distributed random variables for the following simulations.

Different topologies have been used [9], whereby the results that are included in this paper have been obtained for the G-WiN topology of the German research network (DFN) [1]. The level one nodes of this topology are core nodes. End systems are connected to the level two nodes that are edge nodes. Links are either Synchronous Transfer Mode (STM) 4, STM 16, or STM 64 connections. The link propagation delay is assumed to be 2 ms. Shortest Path First (SPF) routing is applied to minimize the number of hops along the paths. Further on, Turn Prohibition (TP) is used to ensure the feed-forward property of the network that is required for a direct application of Network Calculus [16]. For the G-WiN topology the combination of SPF and TP increases the length of only one path by one hop compared to SPF routing. Simulation results of a Guaranteed Rate Service, which only considers capacity constraints, have shown that the impacts of TP on SPF routing are comparably small [9]. Further on, the TP algorithm can be configured to prohibit turns that include links with a comparably low capacity with priority [16].

The emulated Premium Service is assumed to be based on PQ. Thus, service curves are of the rate-latency type. The latency is set to the time it takes to transmit 4 Maximum Transmission Units (MTU) of 9.6 kB, to account for non-preemptive scheduling, packetization, and a router internal buffer for 2 MTU.

Simulation results that compare the different shaping options are shown in figure 2. The performance measure that we apply is the ratio of accepted requests divided by the overall number of requests. Simulations have been run, until the 0.95 confidence interval of the acceptance ratio was smaller than 0.01. Requests for Premium capacity are generated by the simulator with random parameters. The CIR is chosen uniformly from 10 Mb/s to 80 Mb/s, the CBS from 1 Mb to 8 Mb, and the target worst case delay from 40 ms to 80 ms. The CIR and CBS are comparably large to model service requests for aggregated traffic trunks. In case of sustainable rate shaping, we find that the acceptance ratio drops to less than 0.2, independent of the actual load ρ with $0 \leq \rho \leq 400$. This is due to the static shaping configuration, which can result in comparably large shaping delays, independent of the requested delay bound. We address this shortcoming by minimum interference shaping, where shaping rates are adaptive. An end-to-end queuing delay of $d_q = 4$ ms respective $d_q = 8$ ms has been applied, to quantify the influence of the setting of d_q. However, we find only minor impacts of d_q in the investigated scenario. Minimum interference shaping allows to increase the acceptance ratio significantly compared to the option without shaping as well as compared to sustainable rate shaping, as can be seen from figure 2.

For illustrative purposes the cumulative density functions of the respective delay bounds are shown in figure 3 for a load of $\rho = 50$. The requested delay bounds are set to infinity to achieve an acceptance ratio of 1.0 for all of the

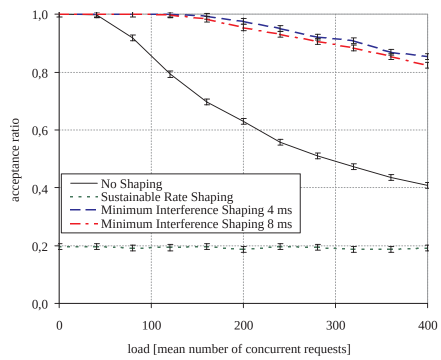

Fig. 2. Acceptance Ratio

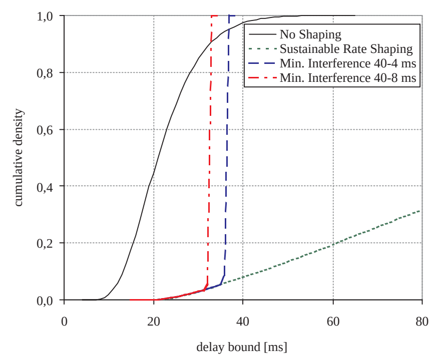

Fig. 3. Delay Bounds

investigated shaping options, to allow for comparison. Here we find the reason for the bad performance of sustainable rate shaping. The delays that are introduced by shaping frequently exceed the range of 40 to 80 ms that is applied for figure 2. For a delay bound of infinity, minimum interference shaping applies the smallest possible shaping rate and becomes the same as sustainable rate shaping. Therefore, results are added for minimum interference shaping for a target delay of 40 ms. Figure 3 clearly shows the impacts of the parameter d_q. In case of $d_q = 4$ ms at most 4 ms of queuing delay are allowed to occur in the core of the network. Thus, shapers are configured so that the propagation delay and the shaping delay sum up to 36 ms for a target delay bound of 40 ms. In case of $d_q = 8$ ms, the propagation delay and shaping delay sum up to 32 ms, leaving room for up to 8 ms of queuing delay in the core of the network which, however, are not required for a load of $\rho = 50$.

Apart from the measured performance increase, traffic shaping is of particular interest, if a Guaranteed Rate Service and a Premium Service are implemented based on the same PHB. In this case a traffic mix with significantly heterogeneous traffic properties and service requirements results. For example Guaranteed Rate Transmission Control Protocol (TCP) streams that can have a large burstiness but no strict delay requirements can share the same PHB with extremely time critical Premium voice or video traffic. In this scenario traffic shaping can be applied to control the impacts of bursty Guaranteed Rate traffic on the Premium Service [15].

As a further benefit, traffic shaping reduces the impacts of EF on the BE class. The starvation of the BE class that can be due to priority scheduling of EF traffic is bound by the maximum EF burst size at the respective outgoing interface. Shaping EF bursts at the network's ingress nodes, reduces this burst size significantly, resulting in less impacts on the BE class. For the experiment in figure 3 and a load of $\rho = 50$ we find that the aggregated Premium burst size within the core of the network stays below 1 Mbit on all links, resulting in less than 0.5 ms BE starvation on a 2.4 Gb/s STM 16 link.

5 Conclusions

In this paper we have addressed the impacts of traffic shaping in aggregate scheduling networks. For this purpose the notation of two leaky bucket constrained arrival curves was introduced. A general per-flow-based service curve has been derived for a FIFO aggregate scheduling rate-latency service element. This form has been solved for the special case of a two leaky bucket constrained flow of interest and a two leaky bucket output constraint has been obtained.

We found that the shaping rate has to be chosen carefully in aggregate scheduling networks, wherefore we evolved an adaptive shaping scheme. Our scheme allows to configure shaping rates individually for a wide variety of heterogenous flows. It minimizes the interference within an aggregate scheduling domain, while it allows to support individual application-specific delay bounds. Thus, it can be applied to adapt end-to-end delay bounds to support heterogenous aggregates, while still allowing for scalability.

References

1. Adler, H.-M., *10 Gigabit/s Plattform für das G-WiN betriebsbereit*, DFN Mitteilungen, 60, 2002.
2. Blake, S., et al., *An Architecture for Differentiated Services*, RFC 2475, 1998.
3. Chang, C.-S., *Performance Guarantees in Communication Networks*, Springer, TNCS, 2000.
4. Charny, A., and Le Boudec, J.-Y., *Delay Bounds in a Network with Aggregate Scheduling*, Proceedings of QofIS, Springer, LNCS 1922, 2000.
5. Cruz, R. L., *A Calculus for Network Delay, Part I: Network Elements in Isolation*, IEEE Transactions on Information Theory, vol. 37, no. 1, pp 114-131, 1991.
6. Cruz, R. L., *A Calculus for Network Delay, Part II: Network Analysis*, IEEE Transactions on Information Theory, vol. 37, no. 1, pp 132-141, 1991.
7. Cruz, R. L., *SCED+: Efficient Management of Quality of Service Guarantees*, Proceedings of IEEE Infocom, 1998.
8. Davie, B., et al., *An Expedited Forwarding PHB*, RFC 3246, 2002.
9. Einhoff, G., *Quality of Service Routing for an IP Premium Service based on MPLS Traffic Engineering*, Master's Thesis, Aachen University, 2003.
10. Le Boudec, J.-Y., *Some Properties Of Variable Length Packet Shapers*, Proceedings of ACM Sigmetrics, 2001.
11. Le Boudec, J.-Y., and Thiran, P., *Network Calculus A Theory of Deterministic Queuing Systems for the Internet*, Springer, LNCS 2050, 2002.
12. Nichols, K., Jacobson, V., and Zhang, L., *A Two-bit Differentiated Services Architecture for the Internet*, RFC 2638, 1999.
13. Pareck, A. K., and Gallager, R. G., *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Single-Node Case*, IEEE/ACM Transactions on Networking, vol. 1, no. 3, pp. 344-357, 1993.
14. Pareck, A. K., and Gallager, R. G., *A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks: The Multiple-Node Case*, IEEE/ACM Transactions on Networking, vol. 2, no. 2, pp. 137-150, 1994.
15. Sander, V., *Design and Evaluation of a Bandwidth Broker that Provides Network Quality of Service for Grid Applications*, PhD Thesis, Aachen University, 2002.
16. Starobinski, D., Karpovsky, M., and Zakrevski, L., *Application of Network Calculus to General Topologies using Turn-Prohibition*, Proceedings of IEEE Infocom, 2002.

Appendix

Proof 1 (Proof of Theorem 1) Substitution of (2) in (1) and application of $\inf_{\theta\geq 0}[\ldots]$ yields (6) for the output arrival curve α_1^{j+1} of flow 1.

$$\alpha_1^{j+1}(t) = \inf_{\theta\geq 0}\Big[\sup_{s\geq 0}\big[\alpha_1^j(t+s) - [\beta^j(s) - \alpha_2^j(s-\theta)]^+ 1_{s>\theta}\big]\Big] \quad (6)$$

With $\sup_{0\leq s\leq\theta}[\alpha_1^j(t+s) - [\beta^j(s) - \alpha_2^j(s-\theta)]^+ 1_{s>\theta}] = \alpha_1^j(t+\theta)$, (7) follows.

$$\alpha_1^{j+1}(t) = \inf_{\theta\geq 0}\Big[\sup\big[\alpha_1^j(t+\theta), \sup_{s>\theta}[\alpha_1^j(t+s) - [\beta^j(s) - \alpha_2^j(s-\theta)]^+]\big]\Big] \quad (7)$$

Then, service curves of the rate-latency type $\beta^j_{R^j,T^j} = R^j\cdot[t-T^j]^+$ are assumed. The condition $R^j\cdot(s-T^j) - \alpha_2^j(s-\theta)\geq 0$ can be found to hold for $\theta\geq\theta'$ with $\theta' = \sup_{s>0}[\alpha_2^j(s) - R^j\cdot s]/R^j + T^j$ [11], whereby $\theta'\geq T^j$. For $\theta\geq\theta'$ (8) and (9) follow.

$$\alpha_1^{j+1}(t) = \inf_{\theta\geq\theta'}\Big[\sup\big[\alpha_1^j(t+\theta), \sup_{s>\theta}[\alpha_1^j(t+s) - R^j\cdot(s-T^j) + \alpha_2^j(s-\theta)]\big]\Big] \quad (8)$$

$$\alpha_1^{j+1}(t) = \inf_{\theta\geq\theta'}\Big[\sup\big[\alpha_1^j(t+\theta), \sup_{v>0}[\alpha_1^j(t+v+\theta) - R^j\cdot(v+\theta-T^j) + \alpha_2^j(v)]\big]\Big] \quad (9)$$

For different settings of θ a θ^* is defined as a function of $(t+\theta)$ in (10). With $\theta^*\geq\theta'$ (11) can be given.

$$\theta^*(t+\theta) = \frac{\sup_{v>0}[\alpha_1^j(t+v+\theta) - \alpha_1^j(t+\theta) + \alpha_2^j(v) - R^j\cdot v]}{R^j} + T^j \quad (10)$$

$$\sup_{v>0}[\alpha_1^j(t+v+\theta) - \alpha_1^j(t+\theta) + \alpha_2(v) - R^j\cdot v] - R^j\cdot(\theta-T^j) \lesseqqgtr 0,\ \text{if } \theta \gtreqqless \theta^* \quad (11)$$

With (10), and (11) the outer $\sup[\ldots]$ in (9) is solved in (12).

$$\alpha_1^{j+1}(t) = \inf\Big[\inf_{\theta>\theta^*}\big[\alpha_1^j(t+\theta)\big], \inf_{\theta'\leq\theta\leq\theta^*}\big[\alpha_1^j(t+\theta) + \sup_{v>0}[\alpha_1^j(t+v+\theta) - \alpha_1^j(t+\theta) + \alpha_2^j(v) - R^j\cdot v] - R^j\cdot(\theta-T^j)\big]\Big] \quad (12)$$

The $\inf[\ldots]$ in (12) is found for $\theta = \theta^*$, which proofs theorem 1. Here, $\theta < \theta'$ is not investigated. Instead, it can be shown that the bound in theorem 1 is attained in the same way as for the special case of a single leaky bucket constrained flow 1 in [11]. Thus, we cannot find a better form for $\theta < \theta'$. □

Proof 2 (Proof of Theorem 2) Based on (4), $\theta(t)$ is derived here for a two leaky bucket constrained flow 1. For the flow 2 arrival curve sub-additivity is assumed without loss of generality.

Case 1 $(t = 0)$ With $\alpha_1^{j+1}(t) = \alpha_1^j(t+\theta(t))$ according to (3) we find the output burst size $b_1^{j+1} = \alpha_1^j(\theta(0))$. Equation (4) is applied at $t = 0$ to find $\theta(0)$.

Case 2 $(0 < t < \bar{t}_1^j - v(t) - \theta(t))$ For this case (13) can be derived from (4), where $v(t)$ is the v for which the $\sup_v[\dots]$ in (4) is found.

$$\theta = \frac{\sup_{0<v\leq \bar{t}_1^j - t - \theta}[\bar{r}_1 \cdot v + \alpha_2^j(v) - R^j \cdot v]}{R^j} + T^j \tag{13}$$

Thus, θ is independent of t for $0 < t < \bar{t}_1^j - v - \theta$. With $\alpha_1^{j+1}(t) = \alpha_1^j(t+\theta)$ according to (3) the output arrival curve of flow 1 increases with $\bar{r}_1$.

Case 3 $(\bar{t}_1^j - v(t) - \theta(t) \leq t < \bar{t}_1^j - \theta(t))$ Equation (4) yields (14) for this case. Note that $b_1^j + \bar{r}_1 \cdot \bar{t}_1^j = \bar{b}_1^j + r_1 \cdot \bar{t}_1^j$.

$$\theta(t) = \frac{(\bar{r}_1 - r_1)\cdot(\bar{t}_1^j - t - \theta) + \sup_{v\geq \bar{t}_1^j - t - \theta}[r_1 \cdot v + \alpha_2^j(v) - R^j \cdot v]}{R^j} + T^j \tag{14}$$

$$= \frac{(\bar{r}_1 - r_1)\cdot(\bar{t}_1^j - t) + \sup_{v\geq \bar{t}_1^j - t - \theta}[r_1 \cdot v + \alpha_2^j(v) - R^j \cdot v] + T^j \cdot R^j}{R^j + \bar{r}_1 - r_1} \tag{15}$$

For $t > \bar{t}_1^j - v(t) - \theta(t)$ it can be immediately seen from (15) that any increase of t results in a corresponding decrease of θ by $(\bar{r}_1 - r_1)/(R^j + \bar{r}_1 - r_1)$. With $\alpha_1^{j+1}(t) = \alpha_1^j(t+\theta)$ according to (3) the output arrival curve of flow 1 increases with less than $\bar{r}_1$. Applying the leaky bucket parameters $(\bar{r}_1, b_1^{j+1})$ in theorem 2 overestimates the output arrival curve, which is allowed, since arrival curves are defined to give an upper bound on the respective arrival functions. However, as long as an increase of t results in a comparably smaller decrease of θ, smaller values v that fulfill $t \geq \bar{t}_1^j - v(t) - \theta(t)$ can be applied in (15). As a consequence, if the $\sup[\dots]$ in (15) is found for $t = \bar{t}_1^j - v(t) - \theta(t)$, it can occur that $t = \bar{t}_1^j - v(t) - \theta(t)$ also holds if t is increased by an infinitesimal Δt, resulting in a dependance of the $\sup[\dots]$ in (15) on t. For sub-additive flow 2 arrival curves, it can be shown that if t is increased, θ decreases slower than t increases. Here, for simplicity differentiable flow 2 arrival curves are assumed. Then, $\partial\alpha_2(t)/\partial t \geq R^j - \bar{r}_1$ at $t = \bar{t}_1^j - v(t) - \theta(t)$ holds, because otherwise case 2 would apply. By substitution of this condition in (15) it follows that θ decreases, if t increases. Further on, $\partial\alpha_2(t)/\partial t \leq R^j - r_1$ at $t = \bar{t}_1^j - v(t) - \theta(t)$ holds, wherefrom it can be found that θ decreases slower than t increases. Following the same argumentation as above, the leaky bucket parameters $(\bar{r}_1, b_1^{j+1})$ are applied.

Case 4 $(t \geq \bar{t}_1^j - \theta(t))$ In this case, (16) can immediately be derived from (4).

$$\theta = \frac{\sup_{v>0}[r_1 \cdot v + \alpha_2^j(v) - R^j \cdot v]}{R^j} + T^j \tag{16}$$

Note that $\theta(t)$ according to (16) is constant for $t \geq \bar{t}_1^j - \theta(t)$. With (3), the output arrival curve of flow 1 is given as $\alpha_1^{j+1}(t) = \alpha_1^j(t + \theta(t))$. The conditions $t + \theta(t) \geq \bar{t}_1^j$, and thus $\alpha_1^j(t+\theta(t)) = \bar{b}_1^j + r_1 \cdot (t + \theta(t))$ hold for $t \geq \bar{t}_1^j - \theta(t)$. Resulting, the output arrival curve of flow 1 increases with rate r_1 for $t \geq \bar{t}_1^j - \theta(t)$. The output burst size can be derived as $\bar{b}_i^{j+1} = \alpha_1^j(t+\theta(t)) - r_1 \cdot t = \bar{b}_1^j + r_1 \cdot \theta(t)$ for any $t \geq \bar{t}_1^j - \theta(t)$, so that $\bar{b}_i^{j+1} = \bar{b}_1^j + r_1 \cdot \theta(\bar{t}_1^j)$ holds. □