

Evaluation of Bandwidth Broker Signaling

Manuel Günter, Torsten Braun
Institute of Computer Science and Applied Mathematics
University of Berne
Neubrückestrasse 10, CH-3012 Bern, Switzerland
<http://www.iam.unibe.ch/~rvs/>
[mgunter|braun]@iam.unibe.ch

Abstract

The Differentiated Services (DiffServ) architecture for the Internet implements a scalable mechanism for quality-of-service (QoS) provisioning. Bandwidth brokers represent the instances of the architecture, that automate the provisioning of a DiffServ service between network domains. Although several bandwidth broker implementations (e.g. [1]) have been proposed, the alternatives and trade-offs of the different viable approaches of inter-broker communication were not studied up to now.

This paper presents the broker signaling trade-offs considered in the context of a DiffServ scenario used by the Swiss National Science Foundation project CATI [8], and it presents results gathered by simulations.

1 Introduction

The DiffServ architecture (RFC 2475) uses automated bandwidth brokers [7] to negotiate service level agreements (SLA) between different autonomous systems. These agreements describe the volume of DiffServ traffic that can be exchanged between two domains and the price for that traffic. If all the domains between two end users have engineered their networks properly and have established SLAs for the DiffServ volume expected, the DiffServ architecture can guarantee end-to-end QoS. However, it is obvious that the local traffic volumes produced by end-users show a dynamic behaviour which has to be reflected also in the SLAs between core networks. Concrete numbers do not exist since no DiffServ service is established yet, but the University of Berkeley, California currently processes initial studies on what QoS service the Internet users value [9].

One option to cope with the changing user requirements is to signal each change in flow activities through the bandwidth brokers to the core networks. However, this is not

desirable, since this would lead to the equivalent signaling scaling problem that the Integrated Services architecture (IntServ) faces, thus undermining the main advantage of DiffServ. Therefore, the signaling between the bandwidth brokers must reflect aggregated changes and should be decoupled to some degree from user flow forwarding. The simulation presented in this paper is used to describe and evaluate pit-falls and trade-offs of such aggregation and decoupling.

This work was performed in the context of the Swiss National Science Foundation project *Charging and Accounting for the Internet* CATI. The project was originally based on an IntServ (RFC 1633) Internet architecture but diverted into the DiffServ research because of the well-known scalability problems of IntServ (RFC 2208).

In chapter 3 we identify the main trade-off of broker signaling as the trade-off between scalability and cost on one hand and end-to-end QoS guarantees on the other hand. We also describe the terminology and assumptions of the simulation which is used to find a solution to the trade-off. Chapter 4 presents an overview of the simulation. Chapter 5 presents the results of the simulation and chapter 6 concludes. The next chapter describes the DiffServ architecture.

2 The Differentiated Services architecture

The DiffServ architecture uses the IP packet's DiffServ Code Point (DSCP) to specify a per-hop behaviour (PHB) for that packet. Once the DSCP is set, all traffic with the same DSCP code is treated in the same way, regardless of its other characteristics (e.g. source/destination address, port). Thus, traffic of many different flows is classified into a small number of traffic classes. This *aggregation* mechanism then easily scales to large core networks, that forward huge numbers of flows.

Two proposed differentiated services PHB are *Expedited Forwarding* (RFC 2598) and *Assured Forwarding* (RFC

2597). The expedited forwarding (EF) service is used to provide the characteristics of a virtual leased line (constant bit-rate). The assured forwarding (AF) service offers less end-to-end guarantees, but allows bursty traffic. Both services are similar in the sense, that the service is expressed in terms of a maximum bit rate. Basically, host networks set up SLAs with their Internet Service Provider (ISP) where they agree on such a rate of DiffServ traffic that the host network can inject into the Internet. ISPs will forward DiffServ packets according to the DSCP marking. ISP networks will queue and schedule DiffServ packets separate from normal ('best-effort') IP packets. Based on the SLAs, the ISPs will engineer their network so that pure DiffServ traffic cannot congest it. They will setup SLAs with adjacent networks, thus enabling end-to-end QoS for DiffServ traffic. Bandwidth brokers are software agents that automate the SLA negotiation. Upon SLA negotiation for new incoming DiffServ traffic, they have to check if their network is able to support it without congestion and they have to (re-)negotiate SLAs.

The planing and provisioning of a working DiffServ network causes additional work and expenses for an ISP. Therefore, and also for the implicit resource reservation, customers will be charged by the ISPs for DiffServ traffic. As mentioned before, the ISPs need the customer to commit to a maximum bit-rate in order to provide the service. This is described in the SLAs. If a customer sends more DiffServ packets than agreed upon in the SLA, the ISP will police that traffic. In the case of the EF service, it will simply drop the exceeding packets. In the case of exceeding AF service traffic, the ISP will increase the packets' drop precedence.

Figure 1 shows an example of a working DiffServ scenario. Here, two host networks (H1 and H2) have established an SLA with an ISP A for 500 kilo-bits per second (Kbps) assured traffic each. They inject that amount of DiffServ traffic plus a large amount of best-effort traffic through fast access links. ISP A forwards all traffic to ISP B. The brokers of the two ISPs have already established a sufficient SLA (1Mbps) between them, thus the DiffServ traffic can continue on its path to the destinations. The link between ISP A and B is only of limited size, thus congestion can occur. However, this congestion only affects the best-effort traffic.

Many questions are left open in this scenario. How and when is the SLA between ISP A and B established? How is the bit-rate of this SLA determined? What happens when the DiffServ traffic produced by the hosts splits to reach many destinations and travels through many ISPs? These questions fall back to the question of the signaling between the bandwidth brokers, which is the subject of this paper. The next section will present trade-offs of a broker signaling design.

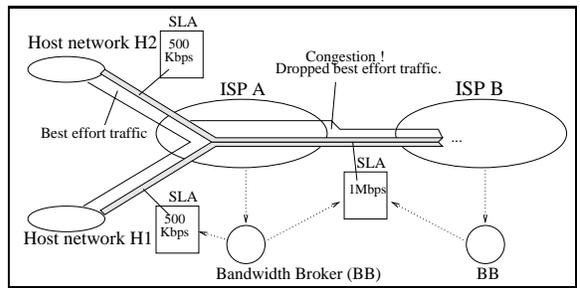


Figure 1. The ideal DiffServ scenario.

3 The design space of the broker signaling

The Differentiated Services architecture consists of a data transport level and a control level. The data transport level of DiffServ includes the different kind of DSCP codes and their corresponding per hop behaviour. The work in this area is far progressed within the Internet Engineering Task Force (IETF). However, when it comes to evaluate the global behaviour of DiffServ, the control level must be specified. Bandwidth brokers play the main role at the control level of the DiffServ architecture. As mentioned before, the design space of the broker signaling is not explored up to know.

3.1 The granularity of the DiffServ control

We consider the granularity of the broker signaling as a key issue of the overall performance evaluation of DiffServ. Therefore, this work evaluates different signaling aggregation levels by their impact on the scalability, the end-to-end behaviour and the cost of the DiffServ control architecture. On one hand, a fine grained control structure (per-flow and using end-to-end signaling) as provided by IntServ supports end-to-end QoS. Unfortunately, it does not scale to large networks such as backbone networks in the ever growing Internet. If, on the other hand, the per-hop nature of DiffServ is combined with a coarse grained per-hop control structure, statements about end-to-end QoS are limited to statistical evidence. Nevertheless, it is end-to-end QoS guarantees that the end-users want, and what they are willing to pay for. Furthermore, a fine grained control level might add a significant management burden to the providers. Additional costs must be paid by the users. We can conclude, that the ratio of possible end-to-end QoS value compared to the costs will define the competitiveness of the DiffServ architecture in the data transport market. The ratio depends directly on the granularity of the involved inter-broker signaling.

The simulations we describe in this paper were performed to identify the ideal level of granularity for DiffServ signaling in respect of the tradeoff between end-to-end QoS

on one hand and scalability & costs on the other hand.

3.2 Reservation and notification strategies.

The bandwidth brokers buy and sell reservations of DiffServ bandwidth. Each such purchase is expressed in an SLA. Before a reservation is granted, a bandwidth broker must check if it really can grant that request. This includes a check of the capacity of the broker's own network, but it can also include signaling to neighbour brokers to either just inform them that the DiffServ traffic volume will change (notification) or to reserve additional bandwidth from them. Both reservation and notification must be handled by a broker signaling protocol. For clarity, we use the term *reservation* for the negotiation of an SLA describing the conditions under which the reserved amount of DiffServ traffic is forwarded with the expected per-hop behaviour. Such a reservation may be triggered by an incoming *notification* of a DiffServ reservation request submitted elsewhere, and it may be blocked until a further notification has been issued to the involved neighbours. The following list summarises the granularity levels that we want to evaluate. The three basic classes include issuing *no notification*, issuing an *end-to-end notification* or issuing a *limited notification*.

No notification. The most simple (absolutely coarse grained) DiffServ control structure could foresee only reservations between brokers, but no notifications propagated further. This adds no notification costs or scalability problem to the DiffServ architecture, but it would not allow end-to-end guarantees. Especially, when DiffServ traffic has to be shaped¹ inside of the core networks, new SLAs need to be established. This is situation shown in figure 2. Note, that if the host network H1 has paid money to the ISP A for injecting DiffServ traffic to the Internet, its users will certainly complain when their traffic is shaped and subsequently dropped or congested in the core network C. In a reasonable scenario, the ISP networks would *measure* the DiffServ traffic and *overprovision* their networks and their SLAs with the adjacent networks. Such measurement based reservation together with significant overprovisioning is what we call the *adaptive reservation* scenario. It can enable cooperating ISPs to give some statistical end-to-end QoS guarantees. A particular problem in this scenario is the loose cooperation between the ISPs. Overprovisioning causes cost for an ISP. The end-to-end QoS is lost when one ISP is not overprovisioning sufficiently.

End-to-end notifications. Before the establishment of a new SLA, the broker notifies the involved neighbour bro-

¹Shaping in the context of the simulation means the 'degradation' of DiffServ traffic due to local non SLA conformity. For EF it means dropping of packets, for AF it means increasing of the drop precedence.

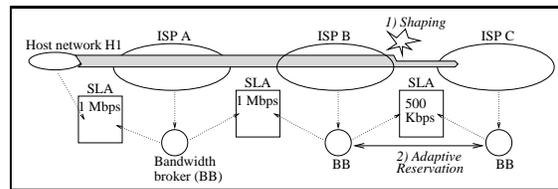


Figure 2. Adaptive reservation triggered by shaping.

ker(s), and sets up a new SLA with them if necessary, to accommodate the new DiffServ traffic. The neighbour brokers acts likewise, notifying upstream. Thus, when the originally requested SLA is accepted, all SLAs from traffic source to sink have already been updated. Obviously, this fine-grained approach allows for end-to-end QoS. Figure 3 depicts this situation. However, if an SLA would be set up or modified for each flow, this DiffServ signaling would be of the same complexity as IntServ signaling and thus suffer from the same scalability problems. Even when only aggregated flows trigger notifications, the number of notifications that a backbone provider has to handle grows with the square of the number of host networks. Furthermore, each new aggregated flow (even if it is relatively small) would face the delay of the end-to-end notification between the brokers. The bandwidth brokers in the backbone would become a bottleneck with negative impact on the DiffServ traffic performance. In the view of the authors, it is clearly undesirable that such end-to-end broker signaling burdens the DiffServ architecture with a scalability problem.

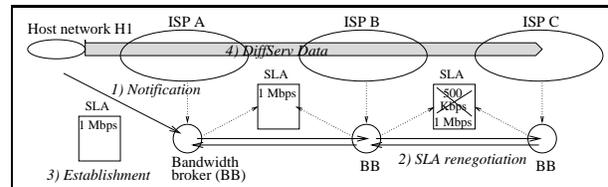


Figure 3. End-to-end notification with end-to-end QoS guarantee.

Limited notification. A simple approach to address the scalability problem is to decrease the granularity of the notifications, so that not each flow or change in an aggregated flow triggers notifications and that not each notification is propagated to further brokers. The bandwidth broker has to judge the relevance of the notifications. The second option is depicted in figure 4. The obvious problem with such an approach is, that it may loose the end-to-end QoS property. Another problem is that the notification process needs flow destination information. If one notification covers different flow aggregations it is not possible in advance to tell which ISPs will experience an impact.

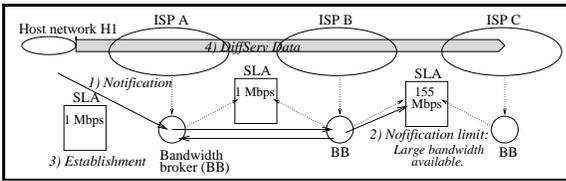


Figure 4. Limited notification.

As mentioned before, in the view of the authors, a fine grained end-to-end notification is not suitable for the DiffServ control level. Therefore, only the adaptive reservation scenario and the limited notification scenario are evaluated by our simulator.

4 The simulator

The simulator is written in Java, for the sole purpose of evaluating bandwidth broker signaling. Due to space limitations we can only briefly describe its underlying assumptions and the network types we used.

4.1 Terminology and assumptions of the simulation

Our simulation uses a coarse grained model of the Internet. The inter-network is modelled as interconnected autonomous systems. Some of these systems are host networks, which act as traffic sources and sinks, the rest are ISP networks² which act as pure transport networks.

Business Assumptions. Each bandwidth broker represents a business party, namely the ISP of the network that it controls. Business models for traffic forwarding may be complex. We have chosen the following three basic assumptions: (1) ISPs demand money from other networks that want to reserve for the injection of DiffServ traffic into their networks. (2) Host networks do not demand money for incoming DiffServ traffic. (3) ISPs avoid breaking SLAs.

The assumptions (1) and (2) enable the simulator to simulate the exchange of money between the brokers. However, this is not subject of this paper. Assumption (3) is highly important in the context of this paper. This is because the desired end-to-end QoS can only be achieved if the ISPs are collaborative. Host network do not need to stick to their SLAs, since their traffic is policed by the ISPs.

Traffic generation. The DiffServ traffic is modelled as aggregated flows. All flows between two distinct host networks are modelled as one aggregated flow. In an inter-network with n host networks, each host network generates

²For notation convenience we will often refer to such networks simply as ‘ISPs’.

$n - 1$ aggregated flows which add up to a total of $n(n - 1)$ aggregated flows. The simulation allows the flow generation to be parameterised in two ways: (1) The total amount of traffic a single host network generates can randomly vary between a minimum and a maximum value. (2) The percentage of traffic assigned to an aggregated flow can change randomly with a parameterised speed which we call the *fluctuation* of the traffic distribution.

Simulation rounds The simulator runs a given number of simulation rounds. A single round has five different phases, namely the traffic calculation, the notification (and reservation), the traffic forwarding, the usage based charging and the adaptive reservation. For each simulation round, the phases are executed in the given order.

4.2 Networks types

Currently, the simulator features two kinds of networks: the Dumbbell- and the Slalom networks. Both have a parametrised size.

Dumbbell. This network has two interconnected backbone networks. As shown in figure 5 there is an equal number of n host networks attached to each of the two core network. Thus, the channel between the two core networks is a possible bottleneck.

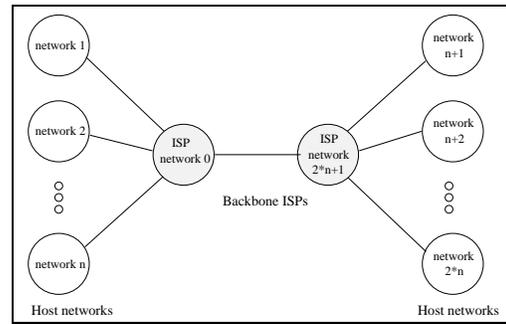


Figure 5. The Dumbbell network.

Slalom. This network is shown in figure 6. The number of backbone networks is customisable. The purpose of this network is to evaluate the end-to-end QoS behaviour, when the DiffServ traffic crosses several autonomous systems. Furthermore, some host networks have more than one ISP.

5 Simulation results

We start by evaluating the adaptive reservation scenario by measuring the end-to-end QoS. Then, we compare the results with a limited notification scenario. We identify the

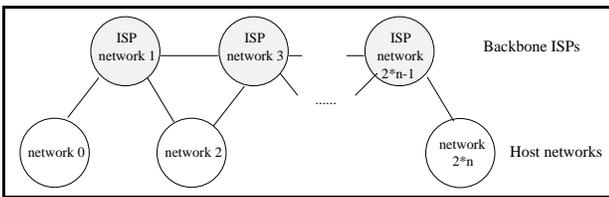


Figure 6. The Slalom network.

'dumbbell' problem when using a naive limited notification approach, and propose an improved solution.

5.1 The adaptive reservation scenario

A valuable service in the adaptive reservation scenario is usually achieved using massive overprovisioning. We used concrete numbers for Frame Relay overprovisioning from [4]. There, a Frame Relay provider would conduct network capacity management on a weekly basis. They provision new trunks between Frame Relay switches when trunk utilisation exceeds 50 percent. The provider will reimburse a user if the delivery success rate is below 99.8 percent. This maps nicely to a DiffServ simulation where the corresponding overprovisioning is 100 percent. Thus, if a broker measures, that outgoing DiffServ traffic exceeds 50 percent of the agreed value in the appropriate SLA, it will renegotiate the SLA. Using only a medium traffic fluctuation our simulation showed that 99.87 percent of the injected DiffServ traffic reached the destination. This seems to be an encouraging result because (1) it shows that the coarse grained nature of the simulator can still produce appropriate results, and (2) because the end-to-end QoS in this scenario is apparently economically interesting. However, it cannot be assumed, that all ISPs will want to set up such a high overprovisioning. Furthermore, measurements with larger traffic fluctuation and with more intermediate ISPs showed a poorer end-to-end behaviour.

Figure 7 shows a simulation of 100 rounds on the Slalom network with 9 backbone ISPs and 10 host networks. There are therefore 90 different aggregated flows. A total amount of 200 traffic units³ is injected into the network at each simulation round. The fluctuations of the flows is set to be high here. This means that between two rounds, some aggregated flows will shrink massively, while others will grow. The brokers arrange for an overprovisioning of 20 percent.

At the beginning of the simulation, no SLAs were set up, thus there is no reservation. All DiffServ traffic generated from the host networks is therefore not policy-conform and is shaped. After the 10th round, the content of the SLAs is adapted reasonably and the shaping reaches a stable level.

³Given the coarse grained structure of the simulator, it would be misleading to use concrete traffic units. Furthermore, the units used here allow a nice integration into the figures.

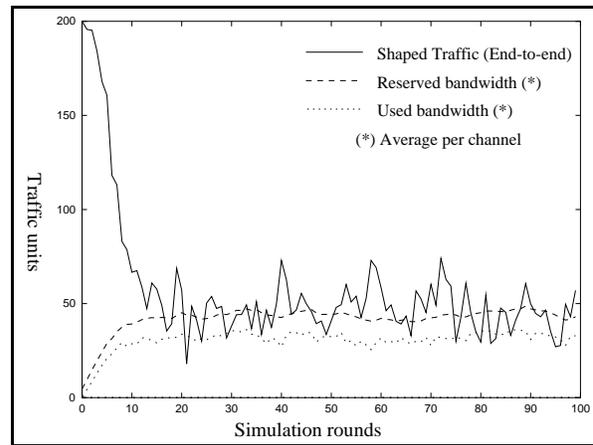


Figure 7. Adaptive reservation with strong fluctuations.

Figure 7 shows the total amount of shaped traffic. Furthermore, the reservation and usage is shown as *average per channel*. Here, they nicely show the 20 percent overprovisioning. As we see in this example, there is a massive loss (shaping) of DiffServ traffic (about 20 percent) in the adaptive reservation scenario, because of an insufficient overprovisioning, heavy traffic fluctuations and a large number of intermediate ISPs.

5.2 Limited notification scenario

In the limited notification scenario, a broker only notifies and reserves upon notifications of significant changes. There are two kinds of problems here. The first is the 'dumbbell' problem, named after the network type that reveals this problem. The other problem is that of the missing destination information in notifications. The next sections describe the problems and propose particular solutions.

The dumbbell problem. The first approach for limited notification was to see the notification and reservation as one process. Thus, a broker reacts upon reservation requests by checking its outgoing SLAs and propagating reservation requests, if necessary. In this approach, the broker includes a reservation threshold. If a new inbound reservation causes the reservation on an outbound SLA to exceed this threshold, the broker would issue a new reservation there, before accepting the inbound request. The threshold effectively limits the number of notifications. However, it can have severe impact on the end-to-end QoS as the following simulation run indicates:

In the dumbbell network of the simulator (presented in figure 5), the host networks have only one channel to an access ISP. Using the naive limited notification approach, the

host networks reserve a constant amount of DiffServ traffic which suffices all their future needs. Although the amount of the traffic sent for the different destinations changes during the simulation, the total amount of the traffic a host network presents to its backbone ISP stays within the SLA. However, since the traffic distribution scheme of each host network changes, the traffic going through the bottleneck channel between the backbone ISPs may also change. Unfortunately, since the host networks don't reserve new bandwidth, no notification is sent, and thus no renegotiation of the SLA between the backbone ISP takes place. Consequently traffic is shaped at the bottleneck channel. Figure 8 shows the situation for the `Dumbbell` network with four host networks on each side. Only in the first round, when no reservation is set up at all, notifications are exchanged. Then, no notification is sent at all for the reason mentioned above. Therefore, as reflected in figure 8, the reservation stays constant. Subsequently, traffic is shaped without hope for the better.

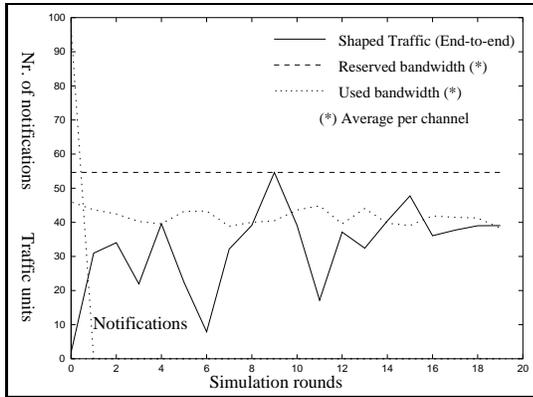


Figure 8. The dumbbell problem.

Lack of destination information in notifications. One approach to limit the notifications is to use one notification to cover several subsequent aggregated flows. Usually, when host networks set up SLAs these SLAs should last some time, thus covering several subsequent flows. However, in that case the notification of such an SLA cannot (in general) include the information of the destination of these flows. There are some special cases however, such as virtual private networks (VPN)[3]. If a host network wants to establish a QoS enabled VPN [5] it could set up an SLA describing the VPN requested. Usually, the VPN peers are known in advance, such as a company's head-quarters and its branch offices. Therefore, the notification of a new QoS VPN can lead to SLAs that cover several aggregated flows and can include their destination information.

Proposed solutions. For the two presented problems with limited notification we propose several solutions and show their viability by simulation. The dumbbell problem can be addressed by decoupling notification from reservation. The dumbbell problem occurs, because necessary notifications are not propagated. The notification chain was interrupted, because it did not lead to a reservation in some place. For the problem concerning the lack of destination information we propose the use of exponential estimation based on measurements.

Decoupled notification limitation mechanism. The decoupled notification limitation mechanism is only a small extension to the presented reservation threshold mechanism. Here, the notification is not directly coupled to a reservation. Upon the reception of a notification, that announces DiffServ traffic on an incoming channel, the bandwidth broker reacts according to the following scheme:

- Estimate the impact on the local network.
- Estimate the impact on the outgoing channels. Use destination information if provided.
- Use the estimation and a *reservation threshold* to determine whether to reserve bandwidth (renegotiate the SLA).
- Use the estimation and a *notification threshold* to determine whether to notify other bandwidth brokers. Typically, this threshold is lower than the reservation threshold. Furthermore, the ISPs should all agree on the value of this threshold.
- Use a *minimal notification size* threshold that stops the propagation of notifications concerning only small changes of DiffServ traffic. Such small notifications might occur when estimating the impact of incoming notifications in absence of destination information (see next paragraph).

We also propose to use adaptive reservation and over-provisioning to smooth out the coarse grained nature of the limited notification approach. For the estimation of the size of the needed reservation and notification in case of missing destination information we propose to use the measurements described in the next section.

Exponential Estimation. An ISP with n channels ($n > 1$) can use a distribution matrix D ($n \times n$ matrix). The entry d_{ij} of the matrix D contains the probability that DiffServ traffic coming in on channel i will leave on channel j . Initially, D contains equal probabilities. However, under the assumption that no routing loops occur, no traffic will leave

the ISP the same way it entered it. Furthermore, as mentioned before, the ISPs do not act as traffic sinks. Thus the initial D is:

$$d_{ij} = \begin{cases} 0 & : i = j \\ \frac{1}{n-1} & : i \neq j \end{cases}$$

Periodically, the ISP can compile measurements of DiffServ traffic into the matrix M , where m_{ij} contains the amount of traffic measured, that entered the network from channel i and left it through channel j . The matrix M can be used to update the matrix D in the following way:

$$D_{\text{new}} = \alpha D_{\text{old}} + (1 - \alpha) \text{normRows}(M)$$

Here $\alpha \in [0..1]$ expresses, to what extend the old estimation is still valid after new measurements. In the simulations, α was set to 0.5. The $\text{normRows}()$ function normalises the absolute traffic measurements to relative values:

$$\text{normRows}(m_{ij}) = \frac{m_{ij}}{\sum_{k=1}^n m_{ik}}$$

To estimate the impact p on an outbound channel j of a notification about DiffServ traffic of the amount a coming from channel i we can simply calculate $p = ad_{ij}$.

The next section shows, how using such exponential estimation together with the extended limited notification mechanism improved the DiffServ performance in the simulation.

Improved simulation results Without having the destination information of aggregated flows, there are more unknown factors, and there need to be more notifications. However, this more realistic scenario is feasible and reasonable as the following example will show. Figure 9 shows the performance under the same conditions as the example for the adaptive reservation scenario (figure 7). Even though we have up to nine intermediate ISPs for a flow, high traffic fluctuation, little overprovisioning (20 percent), and the destination information is not included in the notifications, the performance is reasonable. The percentage of DiffServ traffic that is shaped is only 11 percent of the total amount of DiffServ traffic presented to the network.

In the first rounds of the simulation, many notifications are necessary to set up the SLAs, but soon the notification limitations restrict the number of notifications to a reasonable level.

If we assume the special case, when the destination information is included in the notifications (e.g. for VPN flows) the result is even improving. Figure 10 depicts the simulation results in this case, using the same harsh network conditions as in the previous example. The shaping decreases to 8 percent of the total DiffServ traffic and there are also less notifications necessary.

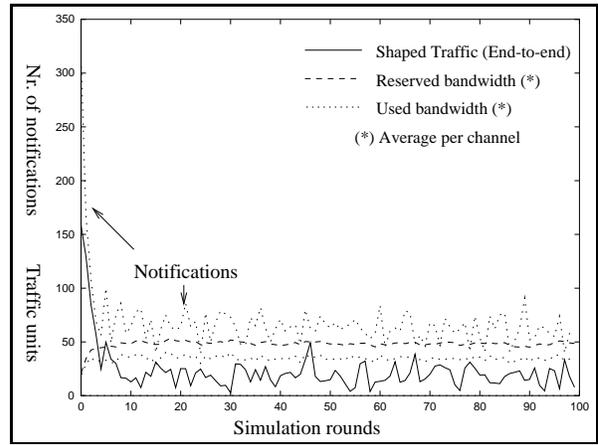


Figure 9. Performance of proposed solution.

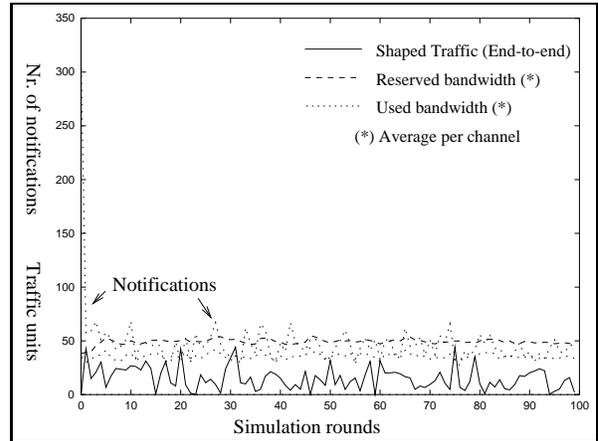


Figure 10. Proposed solution using destination information.

6 Conclusions

Besides of the definition of per-hop behaviour of differentiated services in the Internet, there is a need to investigate in the control level of the DiffServ architecture. The control level consists mainly of bandwidth brokers that use signaling between each other, ideally to establish end-to-end quality-of-service. The simulations presented are a step towards finding the ideal signaling granularity, taking the trade-off of between scalability and end-to-end QoS into account. We worked out the following conclusions:

- An adaptive reservation mechanism based on measurements is a light weight solution, but cannot be used to provide reasonable end-to-end guarantees.
- A limited notification scenario can encounter two major pitfalls: (1) The 'dumbbell' problem, where the

notifications do not reach the bottleneck channel, and thus cannot trigger the needed reservations. (2) The 'missing destination information' problem, where an SLA covers several future flow aggregations in advance.

- Nevertheless, limited notification is a viable way to reduce the number of notifications thus being scalable, but keeping reasonable end-to-end behaviour. The decoupling of notifications and reservations, a set of few thresholds as well as a traffic estimation mechanism produce encouraging results in the simulation.
- Services such as a virtual private network service, that allow for the setup of SLAs describing large flow aggregations and that include destination information are beneficial for a limited notification DiffServ control mechanism.

This paper showed, that the limited notification approach to the bandwidth broker signaling is a favourable option. However, the results were only produced by a coarse grained simulation.

Related Work. Xiao and Ni [10] compare different QoS architectures for the Internet. They present bandwidth broker signaling in the DiffServ control architecture and propose to extend the RSVP protocol for that purpose. This is equivalent to the end-to-end notification scenario against which we argued in section 3.2.

The Internet2 QoS working group [6](QBone) is about to implement DiffServ provisioning using bandwidth brokers. The group proposes a measurement architecture for DiffServ performance evaluation. That approach resembles the adaptive reservation scenario. They are starting to investigate in the inter-broker signaling to automate the establishment of a premium service; but currently the brokers are configured manually. Another adaptive-reservation based approach for QoS provisioning is described in Duffield et al. [2]. They use real traffic traces to show significant capacity savings when comparing an elaborate adaptive scheme to static provisioning. The work optimises the provisioning for VPNs. However, the work only considers a single provider scenario.

Future Work. The control level of the DiffServ architecture bears many subjects to current research. These topics include the business model, the security architecture necessary, a design for the monetary transactions involved and SLA routing. We want to deploy the results of this paper in a more general service broker architecture that we proposed in [5]. An implementation of bandwidth brokers using the limited notification approach is planned. The simulation results of the proposed solution are encouraging but far from

perfect. Nevertheless, they seem to indicate that by putting more intelligence to the bandwidth brokers, the end-to-end quality can be further improved without decreasing the scalability of the architecture. By finding an optimum, a deployable quality-of-service architecture for the Internet can finally become a reality.

Acknowledgements. This work has been performed in the framework of the project Charging and Accounting Technology for the Internet - CATI (CAPIV 5003-054559/1 and MEDeB 5003-054560/1) which is funded by the Swiss National Science Foundation (SNF). The authors would like to acknowledge all contributions from their project and working colleagues.

References

- [1] British Columbia Institute of Technology. Bandwidth broker high level design, 1998. <http://www.merit.edu/working.groups/i2-qbone-bb/>.
- [2] N. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. Ramakrishnan, and J. van der Merwe. A flexible model for resource management in virtual private networks. In *SIGCOMM*. ACM, 1999.
- [3] P. Ferguson and G. Huston. What is a VPN - part I. *The Internet Protocol Journal*, 1(1), 1998.
- [4] P. Ferguson and G. Huston. What is a VPN - part II. *The Internet Protocol Journal*, 1(2), 1998.
- [5] M. Günter, T. Braun, and I. Khalil. An architecture for managing QoS-enabled VPNs over the Internet. In *IEEE Conference on Local Computer Networks*, 1999.
- [6] Internet2 QoS Working Group. Qbone. <http://www.internet2.edu/qos/qbone/>, 1999.
- [7] F. Reichmeyer, L. Ong, A. Terzis, L. Zhang, and R. Yavatka. A two-tier resource management model for differentiated services networks, 1998. IETF work in progress.
- [8] B. Stiller, T. Braun, M. Günter, and B. Plattner. Charging and accounting technology for the Internet. In *Multimedia Applications, Services, and Techniques*, May 1999.
- [9] P. Varaiya, R. Edell, and B. Rupp. Measuring what Internet users value: the INDEX project. <http://www.INDEX.Berkeley.EDU/public/index.phtml>, 1998.
- [10] X. Xiao and L. M. Ni. Internet QoS: A big picture. *IEEE Network*, 13(2), March/April 1999.