Fairness of Assured Service

Florian Baumgartner[†], Torsten Braun[†] and Christian Siebel[‡]
baumgart|braun@iam.unibe.ch
[†]*Institute of Computer Science and Applied Mathematics*
University of Berne
Neubrückstrasse 10, CH-3012 Berne, Switzerland
siebel@tzd.telekom.de
[‡]*Deutsche Telekom AG, Technology Center*
Am Kavalleriesand 3, D-64295 Darmstadt, Germany

Abstract

The Assured Service was proposed as a possible service at IETF's Differentiated Service (DiffServ) working group. The basic idea of this service is the negotiation of a certain service profile between the ISP (Internet Service Provider) and the customer. The user tags packets leaving his end system (network) according to this profile as in- or out of profile. The ISP then forwards the in profile packets with higher priority and especially favours them in overload situations by discarding out of profile packets prior to in profile packets. One of the central questions concerning Assured Service is whether the available bandwidth is shared fairly among several flows. Some initial research in this field is done in [IN98] and [BB]. In contrast to [IN98] especially the interaction of different kind of flows (UDP and TCP) is investigated.

1 Introduction

In addition to the existing best effort Internet services, there is currently a great demand for high level services capable to provide QoS (Quality of Service). On the other hand it has been recognized, that in larger networks such services cannot be provided by an RSVP [BZB+97] based integrated service approach, because of the missing scalability [MBB+97].As an alternative, especially for large IP networks (e.g. backbones), the concept of Differentiated Services has been developed [BBC+98]. The initial approach [CW97] proposed one Assured Service class and one bit for the in and out of profile marking. Newer proposals [HBWW98] suggest the introduction of four classes for Assured Forwarding with three priorities per class. In this evaluation we use the model of [CW97] with one Assured Service class and one bit indicating the priority of the packet.

2 Differentiated Services

In contrast to the Integrated Service architecture, the Differentiated Service approach [BBC+98] is based on the aggregation of application data flows, i.e. reservations are done for several related flows, e.g. for all flows between two subnetworks. These reservations are more static i.e. no dynamic reservation for every communication is done. So reservations last for several sequential communication flows.

According to the Differentiated Services concept IP packets are marked with different priorities. This can be done within the user's end-system or router or by the ISP. Every router reserves a certain amount of resources (especially bandwidth) for every service class. An ISP is then able to provide several differently priced service-classes to his customers.

Differentiated Services allow the users to define a certain rate or share of packets to be forwarded by the ISP with high priority. This concept cannot guarantee QoS as a rule, but is easier to implement as flow-based resource reservation and offers a better quality of service (comparable to the controlled load service of the integrated services architecture) than best effort traffic.

The probability that the desired quality of service is provided depends mostly on the provisioning of the network, i.e. on the probability that routers are overloaded by high priority data.

Marking of packets is supported by the so called DS-field (Differentiated Services field) in the IP header, that is mapped to the Type of Service Byte at IPv4 (see 1) and to the Traffic Class Byte of the IPv6 header. The DS-Field contains a six bit wide field called DSCP (Differentiated Services Code Point), specifying the desired handling in the router. The last two bits (currently unused, cu) are reserved for future use.

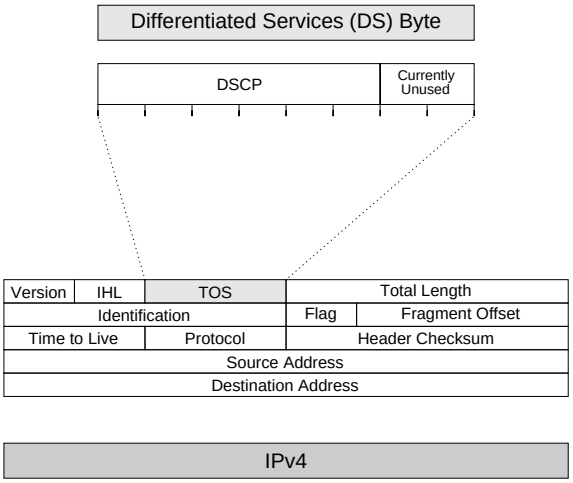

Figure 1: DSCP in IPv4 and IPv6

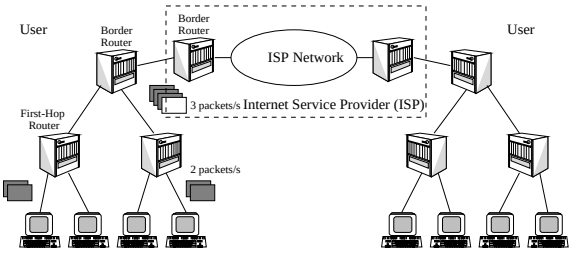

Figure 2: Common simulation scenario

Figure 2 shows the principle of Assured Service by an example. The user's networks and the ISP networks represent so called differentiated service domains being connected to other DS domains by border routers. In the user's network the first hop router represents a special border router at the link to the end-systems. To avoid changes in the end-systems the first hop router may analyse the forwarded packets by their IP address and UDP/TCP ports and then assign a certain priority, i.e. setting the in profile bit for service conformant packets. Of course, the maximum rate of assured service packets has to be regarded. This is assured by shaping in the first hop routers and reshaping functionality in the users border router at the link to the ISP network. Nevertheless, the ISP has to check whether the user obeys the maximum allowed rate of high priority packets and if necessary correct it. To achieve this, not conforming packets are tagged as out of profile at the the entrance to the ISP's network.

2.1 Simulation Model

For the simulation a common situation for Assured Service was chosen. Two user networks are connected via an ISP (see figure 2). Each user respectively the first hop router tags a certain part (dependent on the negotiated profile) of his packets and forwards them to the ISP. In our simulations we evaluate different numbers of connections, so the ISP is the communications bottleneck.

2.2 Topology in ns

The two main components of the Differentiated Service Concept are the RIO (RED with in and out) Queue [CW97] and the mechanism for tagging packets as in or out of profile. Each one is realised in `ns` [ns] as a link between network nodes. Figure 3 shows an example with three connections between S_i to C_i and the bottleneck link between I_1 and I_2. The solid line represents the RIO queue, the dashed one the "tagger" components.
In our simulation all links including the link between I_1 and I_2 are capable to transmit 1 Mbit/s with a delay of one millisecond.

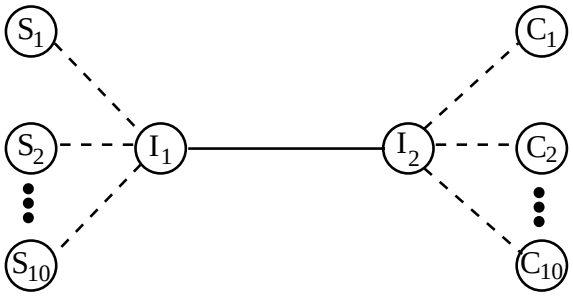

Figure 3: Simulation scenario in *ns*

2.3 Traffic

A connection source is attached to the node S_i communicating with the sink at C_i. Several traffic types have been simulated. We will examine the fairness of several constant bit rate UDP flows, several TCP flows, and the interaction of both. Ten client/server pairs sending at full bandwidth cause a tenfold overload on the bottleneck link. Another assured bandwidth is assigned to every Server S_i, causing each "tagger" queue S_i-I_1 to mark another percentage of forwarded packets as in profile. Every source sends with the same maximum bit rate of 1 Mbit/s. The single sources start delayed to each other. Also there is always overload on the bottleneck links, we will change the amount of assured bandwidth varying the congestion of

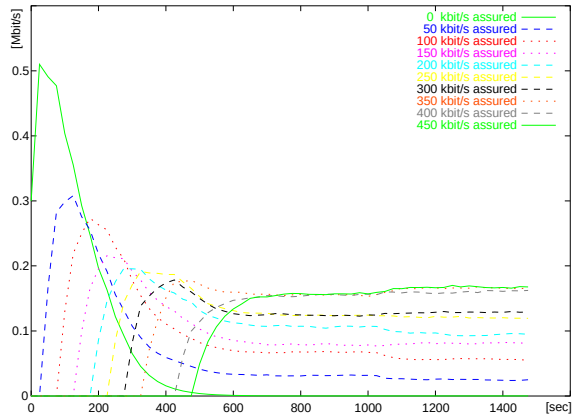

Figure 4: bandwidths of ten cbr flows causing heavy congestion

in-profile traffic.

3 Results

This section presents various simulation results. Two main cases have been simulated. First of all, it is investigated, how the flows interact, if the sum of allocated in profile traffic exceeds the capacity of the bottleneck. In a second step not more in profile capacity is allocated as the bottleneck link is able to transport.

3.1 Heavy Overload

First evaluations have been done to simulate heavy overload of in profile traffic on the bottleneck link between I_1 and I_2. The single flows get assured bandwidths from 0 to 450 kbit/s and start to send at different times. The flow starting to send at last gets the highest assured bandwidth. The total amount of assured bandwidth is 2.250 Mbit/s which means a more then the double in profile overload.
Figure 4 shows results using ten constant bit rate UDP-type connections. The first flow (the one with no assured bandwidth) gets no bandwidth, whereas the flow started sending at last and with the highest assured bit rate reaches the highest throughput. Between these two extrema, all the other flows share the bandwidth in a relatively fair way. The second column of table 1 shows the percentage of assured bandwidth really achieved.

ass. bw	UDP	TCP	UDP+TCP
0	-	-	-
50	49 %	94 % †	88 %
100	55 %	86 % †	4 % †
150	54 %	86 % †	85 %
200	48 %	57 % †	<1 % †
250	48 %	47 % †	80 %
300	49 %	39 % †	<1 % †
350	47 %	37 % †	80 %
400	40 %	27 % †	<1 % †
450	37 %	28 % †	73 %

Table 1: percentage of assured bandwidth reached. A †marks the TCP flows.

It can be seen, that each UDP flow can achieve about 40 to 55 percent of his assured bandwidth. Now we use exactly the same topology and assured bandwidths but TCP instead of the constant bit rate UDP traffic.
Figure 5 shows the resulting bandwidth values. The flow with no assured bandwidth is not able to transmit any data in the congestion situation, whereas the other perform according more or less to their assured bandwidth. On the other hand, the single bandwidths do not differ as nice as this was the case with the UDP flows. The reason for this is the TCP congestion control, causing a source to reduce the sending bandwidth in an overload situation. The third column in table 1 shows - like before with constant bit rate

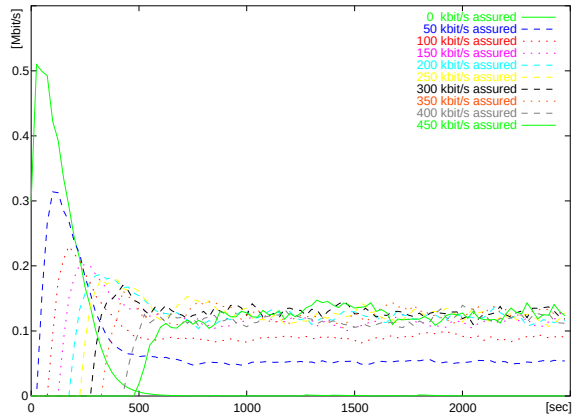

Figure 5: bandwidth of ten TCP flows causing heavy congestion

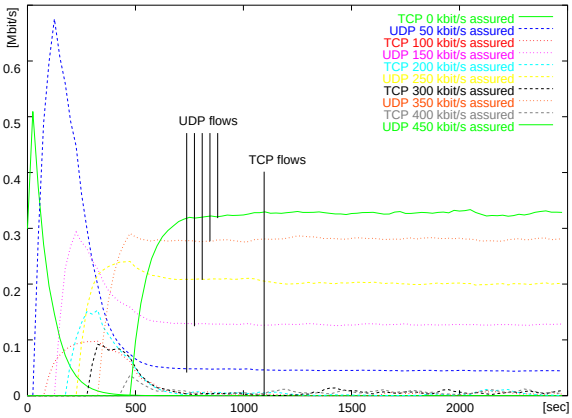

Figure 6: bandwidth of mixed cbr TCP flows causing heavy congestion

UDP - the reached percentage of assured bandwidth. It is obvious, that there is a decrease of achieved bandwidth dependent on the amount of assured bit rate. The lower the assured bit rate the higher is the probability, that a flow gets this bandwidth.

As a final aspect of the evaluation of congestion situations with assured service the interaction of both types of traffic has been examined. The constant bit rate UDP traffic is supposed to suppress the TCP flows during overload. Figure 6 shows the graphs confirming this expectation. The total amount of assured bit rates allocated by UDP flows is about 1.250 Mbit/s. So the very aggressive UDP flows alone lead to a congestion on the bottleneck link, leaving no bandwidth for the TCP connections. (see column four on table 1). The suppression of TCP by aggressive UDP flows is not a special problem of Assured Service, but a general problem in the Internet. So there is a demand for mechanisms being able to detect and police aggressive flows.

3.2 No Congestion

So far the simulation showed the sharing of bandwidth during an extreme overload situation. As mentioned before the most important issue for the success of Assured Service is the proper dimensioning of the network. In the previous section simulations have been done with a more than double load than the bottleneck link is capable to transmit. In this section, we will show the interaction of several flows using assured bandwidths that the network is capable to provide.

Figure 7 shows the graphs. Similar to the previous scenario, different assured bandwidth values have been allocated for the connections. The sum of assured bandwidths is 675 kbit/s. So the bottleneck can forward all in profile packets. Table 2 shows the percentage of the assured bandwidth a flow was capable to reach in the last 100 seconds of the simulation. For flows with no assured bandwidth the bit rate achieved is given. As above flows with small assured bandwidth perform in general better. So, the constant

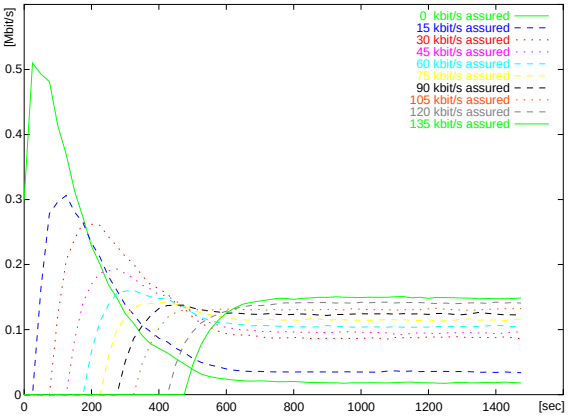

Figure 7: bandwidths of ten cbr flows

ass. bw	UDP	TCP	UDP+TCP
0	18 kbit/s	30 kbit/s	0 kbit/s †
15	220 %	379 % †	690 %
30	286 %	230 % †	100 % †
45	219 %	166 % †	273 %
60	175 %	146 % †	100 % †
75	154 %	152 % †	184 %
90	134 %	132 % †	99 % †
105	125 %	128 % †	149 %
120	117 %	125 % †	99 % †
135	113 %	129 % †	131 %

Table 2: bandwidth reached in percent of the assured bandwidth. The †marks the TCP flows.

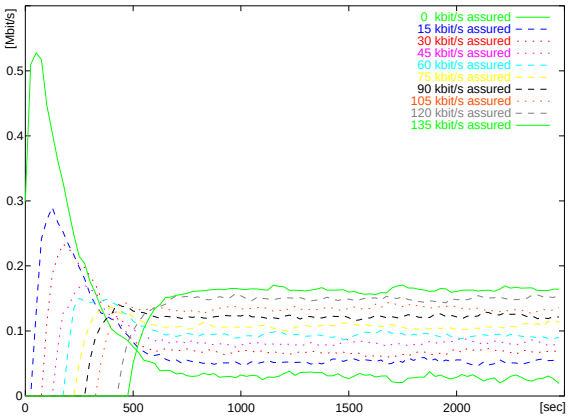

Figure 8: bandwidths of ten TCP flows

bit rate UDP flow with only 15 kbit/s assured bandwidth exceeds this bit rate by 220 percent, whereas the flow with 135 kbit/s assured bandwidth only gains 113 percent (152 kbit/s).
As before in the case with heavy congestion figure 8 shows the behaviour of ten TCP flows, figure 9 depicts the interaction of mixed UDP and TCP traffic. For the bit rates achieved in the last 100 seconds see table 2. In contrast to the situation with the sum of in profile traffic exceeding the capacity of the bottleneck link, now every flow gets as least the assured bandwidth. Of course the very aggressive constant bit rate UDP sources use nearly the whole bandwidth not allocated by assured traffic. But Assured Service is capable to protect the TCP flows in a way they can meet their profile.

4 Assured TCP Flows only

As final aspect now the capability of Assured Service to protect TCP flows against aggressive UDP flows shall be examined. For this purpose we use a similar simulation scenario as before, but assign assured bit rates to the TCP flows

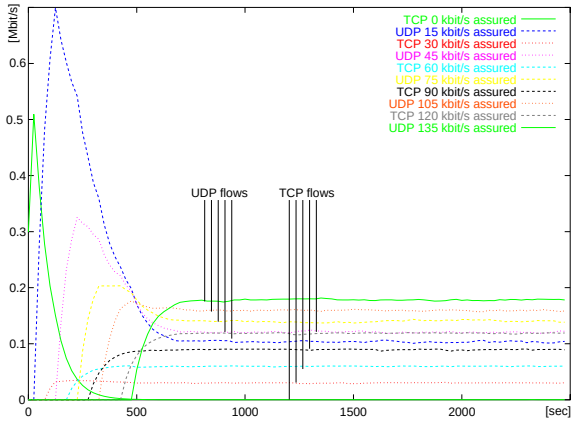

Figure 9: bandwidths of ten mixed UDP and TCP flows

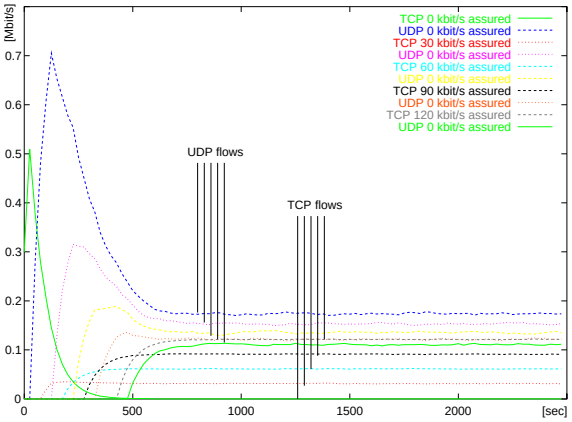

Figure 10: assigning assured bandwidth to tcp flows only

only, while the UDP traffic is transported with best effort. In analogy to the former simulations each traffic source starts sending timely delayed. Figure 10 shows the progression of the simulation, table 3 the bandwidths each flow achieved in the last 100 seconds of the simulation.
Every TCP flows is able to transmit data at least at the corresponding assured bandwidth, while the very aggressive constant bit rate UDP flows occupy the rest. As can be seen on table 3 it is almost impossible for the TCP flows to transmit data with more than their assured bandwidth.
This shows, that Assured Service is surely capable to guarantee high priority TCP flows at least their assured bandwidth, while aggressive best effort traffic blocks each out of profile TCP transmission.

ass. bw	Traffic Type	bw in kbit/s
0	TCP	0
0	UDP	172
30	TCP	30.9
0	UDP	152.6
60	TCP	61.1
0	UDP	136.6
90	TCP	91.0
0	UDP	121.9
120	TCP	121.3
0	UDP	112.1

Table 3: bandwidth reached

5 Conclusion and Outlook

The simulation results presented in this contribution show, that the Assured Service represents an promising approach to support Quality of Service in large IP networks. It has been shown,

that a proper dimensioning of the network is crucial for the probability to achieve the assured bandwidth. The development of methods and tools for this purpose will be a task for future research.
The fairness which can be achieved by Assured Service face the same problems best effort TCP traffic does: Traffic with congestion control mechanisms can be blocked by misbehaving, aggressive flows. As an advantage of Assured Service could been showed, that Assured Service now is able to protect TCP flows against aggressive traffic and can guarantee a certain minimum bitrate. Nevertheless the detection and policing of aggressive flows will one of future tasks in the Internet.

References

[BB] F. Baumgartner and T. Braun. Evaluierung von Assured Service für das Internet. to be published in KiVS'99.

[BBC+98] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weis. An architecture for differentiated services. Internet Draft draft-ietf-diffserv-arch-01.txt, August 1998. work in progress.

[BZB+97] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin. Resource reservation protocol (rsvp) -version 1 functional specification. Request for Comments 2205, September 1997.

[CW97] D. Clark and J. Wroclawski. An approach to service allocation in the internet, work in progress. Internet Draft draft-clark-diff-svc-alloc-00.txt, Juli 1997.

[HBWW98] Juha Heinanen, Fred Baker, Walter Weiss, and John Wroclawski. Assured forwarding phb group. Internet Draft draft-ietf-diffserv-af-02.txt, October 1998. work in progress.

[IN98] J. Ibanez and K. Nichols. Preliminary simulation evaluation of assured service. Internet Draft draft-ibanez-diffserv-assured-evald-00.txt, August 1998. work in progress.

[JNP98] Van Jacobson, K. Nichols, and K. Poduri. An expedited forwarding phb. Internet Draft draft-ietf-diffserv-af-02.txt, October 1998. work in progress.

[MBB+97] A. Mankin, F. Baker, B. Braden, S. Bradner, M. O'Dell, A. Romanow, A. Weinrib, and L. Zhang. Resource reservation protocol (rsvp) version 1 applicability statement. Request for Comments 2208, September 1997.

[ns] Ucb/lbnl/vint network simulator - ns (version 2). URL: http://www-mash.CS.Berkeley.EDU/ns/.