Differentiated Services: A New Approach for Quality of Service in the Internet

Florian Baumgartner, Torsten Braun and Pascal Habegger
University of Berne
Institute of Computer Science and Applied Mathematics
Neubrückstrasse 10
CH-3012 Berne
Switzerland
baumgart@iam.unibe.ch
braun@iam.unibe.ch
habegger@iam.unibe.ch

Abstract

With the grown popularity of the Internet and the increasing use of business and multimedia applications the users' demand for higher and more predictable quality of service has risen. A first improvement to offer better than best-effort services was made by the development of the integrated services architecture and the RSVP protocol. But this approach proved only suitable for smaller IP networks and not for Internet backbone networks. In order to solve this problem the concept of differentiated services has been discussed in the IETF, setting up a working group in 1997. While RSVP classifies packets according to application flow properties, differentiated services are based on the idea that the user negotiates a service profile with his Internet service provider (ISP) for specially marked packets and then transmits marked packets over the ISP network. A further significant difference to RSVP consists in the fact that for scaling reasons the service profile is only negotiated and policed for a set of aggregated flows. This article gives an overview of the activities of the IETF with regards to differentiated services and presents several proposals for the implementation of differentiated services.

Keywords

Differentiated Services, Internet, Quality of Service

1 INTRODUCTION

A central problem of today's Internet exists in the mostly unpredictable service and the often very low quality of transmission. At present there does not exist any satisfactory solution to this problem.

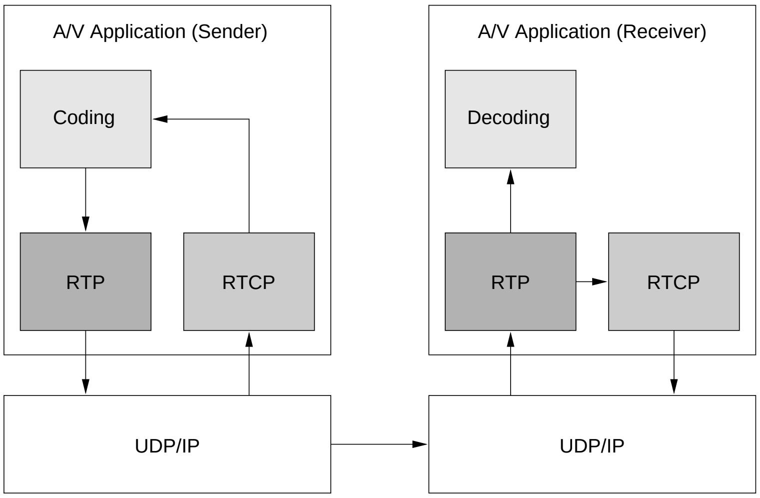

Figure 1 Monitoring of QoS using RTCP and RTP

A pragmatic approach to achieve good quality of service (QoS) is an adaptive design of the applications to react to changes of the network characteristics (e.g. congestion). Immediately after detecting a congestion situation the transmission rate may be reduced by increasing the compression ratio or by modifying the A/V coding algorithm. For this purpose functions to monitor quality of service are needed. For example, such functions are provided by the Real-Time Transport Protocol (RTP) (Schulzrinne *et al.* 1996) and the Real-Time Control Protocol (RTCP). The receiver in Figure 1 measures the delay and the rate of the packets received. This information is transmitted to the sender via RTCP. With this information the sender can detect if there is congestion in the network and adjust the transmission rate accordingly. This may affect the coding of the audio or video data. If only a low data rate is achieved, a coding algorithm with lower quality has to be chosen. Without adaptation the packet loss would increase, making the transmission completely useless.

1.1 Integrated Services and RSVP

Adaptive methods have their limitations when an application requires a certain minimum bandwidth to achieve a reasonable QoS. In these cases a minimal QoS has to be guaranteed by resource reservation. Special applications with real-time requirements depend on resource and bandwidth reservation. This is the reason why the Integrated Services (IntServ) working group defined several services which extend the simple best-effort service: the Controlled Load Service and the Guaranteed Service.

These services are provided for flows i.e. application data streams between end systems. For example three flows exist in Figure 2, two from sender *S* to the receiver *R1* and one flow from *S* to *R2*. Between the sender and *R1* several applications

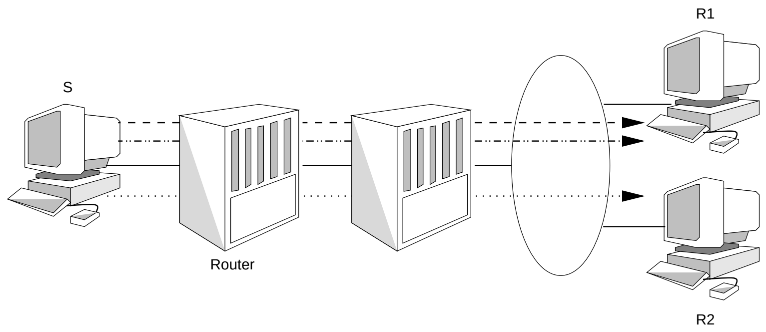

Figure 2 Application flows

may be active (e.g. a data transmission via FTP and a terminal emulation), or an application may support several flows at the same time (a WWW browser opens TCP-connections to a server).

The resources for a flow are reserved within the end systems and the routers using signaling protocols. For this reason, network elements like routers, nodes and even the operating systems within the end systems have to check whether sufficient CPU time, memory and network bandwidth are available in order to provide a certain service (admission control). Resources then have to be reserved and assigned to the packets of the respective flow (scheduling). Finally, the compliance to the negotiated traffic characteristics has to be monitored (policing).

The Resource Reservation Setup Protocol (RSVP) has been developed as a signaling protocol for resource reservation (Braden *et al.* 1997). RSVP extends the IP protocol stack, i.e. data is transmitted unchanged using IP. It exchanges only signaling information describing the QoS to be given to the TCP/IP or UDP/IP flows. The RSVP resource reservation is receiver-oriented. The receiver generates the reservation message containing the desired service parameters for the received application data flow.

RSVP has been criticized mainly for its limited applicability in large IP networks. The RSVP working group of the IETF has evaluated the applicability of the current version. The flow-based approach is considered as the main problem of RSVP since resources are reserved for every single flow. This cannot be realized with conventional routers if large networks with millions of users and possibly several flows per user have to be supported. Routers are not able to store such a huge number of flow states because of limited memory resources. Secondly, the amount of flows will increase the complexity of packet scheduling in the routers. Scheduling is essential for guaranteed services. A further disadvantage is the lack of standards for accounting and billing, making resource reservation as a result quite unrealistic. For these reasons it is recommended to use RSVP only in small confined networks (Mankin *et al.* 1997).

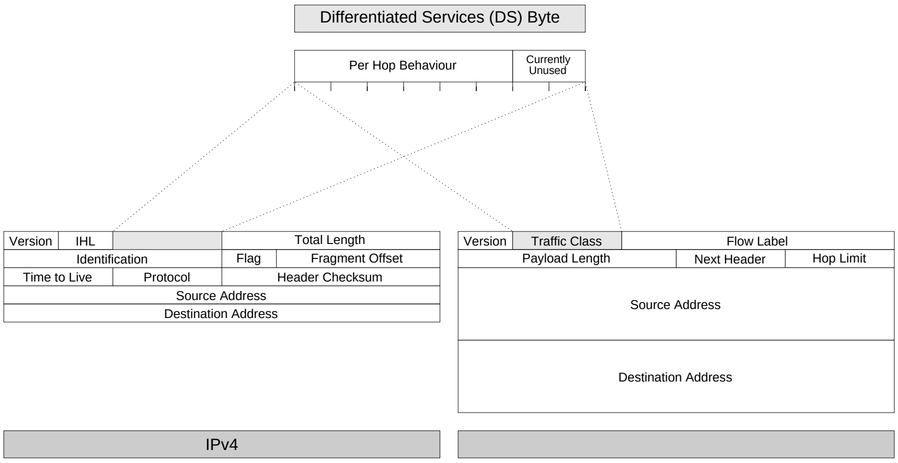

Figure 3 DS byte in IPv4 and IPv6

2 DIFFERENTIATED SERVICES: BASICS AND TERMINOLOGY

A demand for higher-level services apart from best-effort has been recognized, but these services cannot be realized using the integrated services approach, particularly in large IP networks. The differentiated services model tries to avoid the disadvantages of best-effort networks and the integrated services approach.

The idea of differentiated services is based on the aggregation of flows, i.e. reservations have to be made for a set of related flows (e.g. for all flows between two subnets). Furthermore, these reservations are rather static since no dynamic reservations for a single connection are possible. Therefore, one reservation may exist for several, possibly consecutive connections.

IP packets are marked with different priorities by the user (either in an end system or at a router) or by the service provider. According to the different priority classes the routers reserve corresponding shares of resources, in particular bandwidth. This concept enables a service provider to offer different classes of QoS at different costs to his customers.

The differentiated services approach allows customers to set a fixed rate or a relative share of packets which have to be transmitted by the ISP with high priority. The probability of providing the requested quality of service depends essentially on the dimensions and configuration of the network and its links, i.e. whether individual links or routers can be overloaded by high priority data traffic. Though this concept cannot guarantee any QoS parameters as a rule it is more straightforward to implement than continuous resource reservations and it offers a better QoS than mere best-effort services.

For packet marking the so-called DS byte (for differentiated services) in the header of each IP packet is mapped to the IPv4 Type Of Service octet (TOS) or to the IPv6 Traffic Class octet (Figure 3). Six bits of this byte are used to define the per-hop behavior (PHB) that a packet experiences in each router. The remaining two bits correspond to the currently unused (CU) field which is reserved for purposes not yet specified and may be assigned later.

The meaning of the individual bits in the PHB field are not yet standardized and are part of ongoing discussions in the Differentiated Services working group (DiffServ) of the IETF. The proposal in (Baker *et al.* 1998) suggests to use one bit for tagging in and out of profile packets and to distinguish service classes with different priorities using the other five bits. Thus, a minimal backward compatibility to the TOS field in IPv4 can be kept. (Nichols *et al.* May 1998) suggests to standardize two different services Default (DE) and Expedited Forwarding (EF) by using two code points in the PHB field. Since the value of the PHB field for a certain service may change at the edge of different ISP networks because of missing standards, it might be necessary to change the value of the PHB field at the border of two networks.

It has to be pointed out that size, meaning and name of the bit fields in the DS byte are subject of further discussions within the DiffServ working group and might change again in the near future. Therefore, the explanations presented here are merely a representation of the status quo of the DiffServ working group. Several sites on the WWW, which are referenced at the end of this text, contain up-to-date information of the exact DS byte definition and should be consulted first of all.

3 SERVICES OF THE DIFFERENTIATED SERVICES APPROACH

At present, several proposals exist for the realization of differentiated services. The approach allowing the combination of different services like Premium and Assured Service seems to be very promising. In both approaches absolute bandwidth is allocated for aggregated flows. They are based on packet tagging indicating the service to be provided for a packet.

A similar idea is pursued by the Scalable Resource Reservation Protocol (SRP). Flows are aggregated automatically at each link, so that the network does not have to know every single flow. No particular signaling protocol is deployed. Only three different packet types (RESERVED, REQUEST, BEST-EFFORT) are introduced, which differ by the tag in the packet header.

An alternative approach (user-share differentiation, USD) assigns bandwidth proportionally to aggregated flows in the routers (for example all flows from or to an IP address or a set of addresses). A similar service is provided by the Olympic service. Here, three priority levels are distinguished assigning different fractions of bandwidth to the three priority levels gold, silver and bronze, for example 60% for gold, 30% for silver and 10% for bronze.

In the following these services are described in more detail.

3.1 Premium Service

With Premium Service the user negotiates with his ISP a maximum bandwidth for sending packets through the ISP network. Furthermore, the aggregated flow is de-

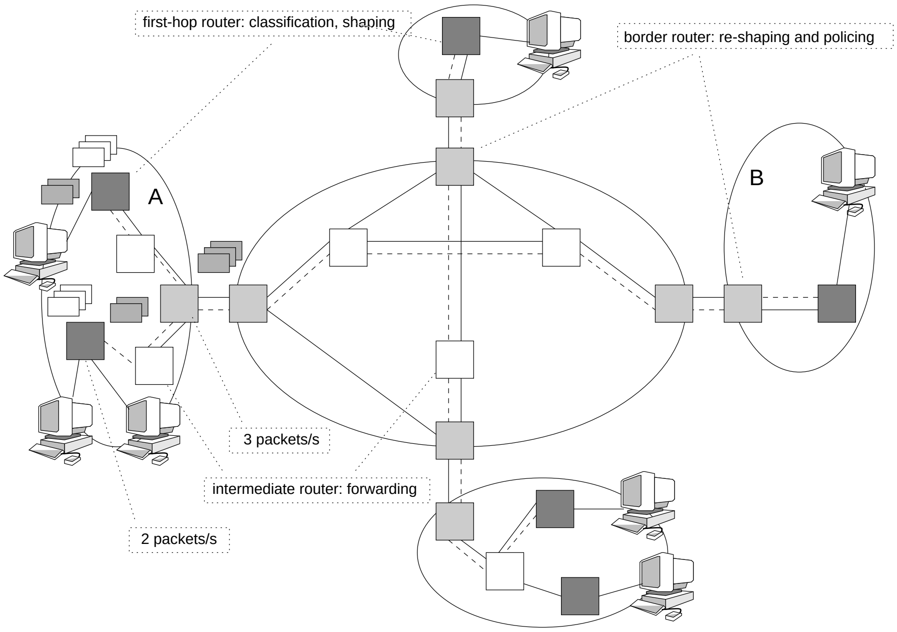

Figure 4 Premium Service

scribed by the packets' source and destination addresses. In Figure 4 users and ISPs have agreed on a rate of three packets/s for traffic from A to B. The user configures the first-hop router in the individual subnet accordingly. In the example above a packet rate of two packets/s is allowed in every first-hop router as it can be expected that no two end systems will use the full bandwidth of two packets/s at the same time.

First-hop routers have the task to classify the packets received from the end systems, i.e. to analyze if the Premium Service shall be provided to the packets or not. If yes, the packets are tagged as Premium Service (P-bit) and the data stream is shaped according to the maximum bandwidth. The user's border router re-shapes the stream (e.g. three packets per second) and transmits the packets to the ISP's border router, which performs policing functions, i.e. it checks whether the user's border router remains below the negotiated bandwidth of three packets/s. If each of the two first-hop routers allows two packets/s, one packet per second will be dropped by shaping or policing at the border routers. All first-hop and border-routers own two queues, one for packets with the P-bit set and one for all other (see Figure 4). If the P-queue contains packets these are transmitted prior to others. The implementation of two queues in every router of the network (ISP and user network) equals to the realization of a virtual network for Premium Service traffic.

Premium Service offers a service corresponding to a private leased line, with the advantage of making free network capacities available to other tasks, resulting in lower fees for the users.

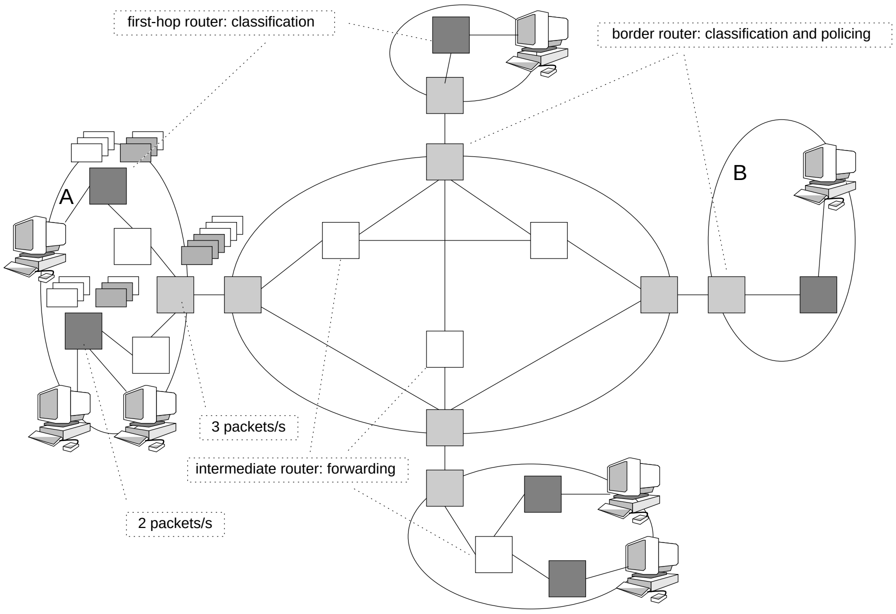

Figure 5 Assured Service

3.2 Assured Service

A potential disadvantage of Premium Service is the weak support for bursts and the fact that a user has to pay even if he is not using the whole bandwidth. The Assured Service tries to offer a service which cannot guarantee bandwidth but provides a high probability that the ISP transfers high-priority-tagged packets reliably. The definition of concrete services has not yet happened, but it is obvious to offer services corresponding to the controlled load service. The probability for packets to be transported reliably depends on the network capacity. An ISP may choose the sum of all bandwidths for Assured Service to remain below the bandwidth of the weakest link. In this case, only a small portion of the available capacity may be allocated in the ISP network. An advantage of the Assured Service is that users do not have to establish a reservation for a relative long time. With ISDN or ATM, users might be unable to use the reserved bandwidth because of the burstiness of their traffic, whereas Assured Service allows the transmission of short time bursts.

With the Assured Service the user negotiates a service profile with his service provider, e.g. the maximum amount or rate of high priority, i.e. Assured Service, packets. The user may then tag his packets as high priority within the end system or the first-hop router, i.e. tag them with an A-bit (see Figure 5). To avoid modifications in the end systems the first-hop router may analyze the packets with respect to their IP addresses and UDP-/TCP-Port and then assign them the according priority, i.e. set the A-bit for conforming Assured Service packets. The maximum rate of high-priority (A-bit) packets must not be exceeded. This is done by (re-)classification in

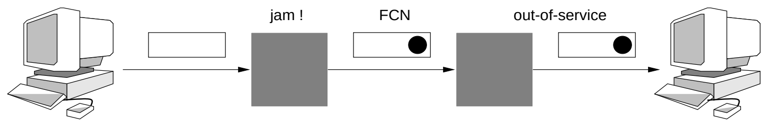

Figure 6 Receiver-oriented realization of Assured Service

the first-hop routers and in the user's border routers at the border to the ISP network. Nevertheless, the service provider has to check if the user remains below the maximum rate for high priority packets and apply corrective actions such as policing if necessary.

For example, the border router at the network entrance will tag the non conforming packet as low priority (out of service, out of profile). An alternative would be to charge higher fees for non conforming packets by the ISP. The tagging of low and high priority packets is done by use of the DS byte.

Bursts are supported by making buffer capacity available for storing bursty traffic. Inside the network, especially in backbone networks bursts can be expected to be compensated statistically.

(a) Receiver-oriented scenarios

One problem of the Assured Service is the negotiation of the service profile between the sender and the ISP. If an Internet user connects to a WWW server, the receiver should be able to determine the transmission quality and take over the costs. Therefore, the receiver should be able to set up a user profile with the ISP. At the border between the ISP and the receiver's network a border router knows the profile agreement (see Figure 6). This router checks whether the received data flow conforms to the service profile. Otherwise, the ISP's border router sets the forward congestion notification (FCN) bit.

This bit might also be set by routers in the network to indicate a congestion situation. If the packet conforms to the profile the border router resets the bit. For a set FCN-bit the receiver has to slow down the sender's data flow, e.g. by delaying TCP-acknowledgments or by the setting of flow control information. If the receiver does not react, the border router may drop future packets.

(b) Adaptation of applications

The Assured Service can be combined with the concept of application adaptation. An application can monitor via RTP/RTCP the throughput respectively the loss rate. According to this, more or less packets might be tagged as high priority. If the network is idle the application might transmit best-effort instead of high-priority packets and save costs. On the other hand the application has to increase the number of high priority packets, if a high loss of low priority packets is detected.

The maximum rate of high-priority packets has to be re-negotiated with the service provider, requiring the support of dynamic reconfiguration or signaling.

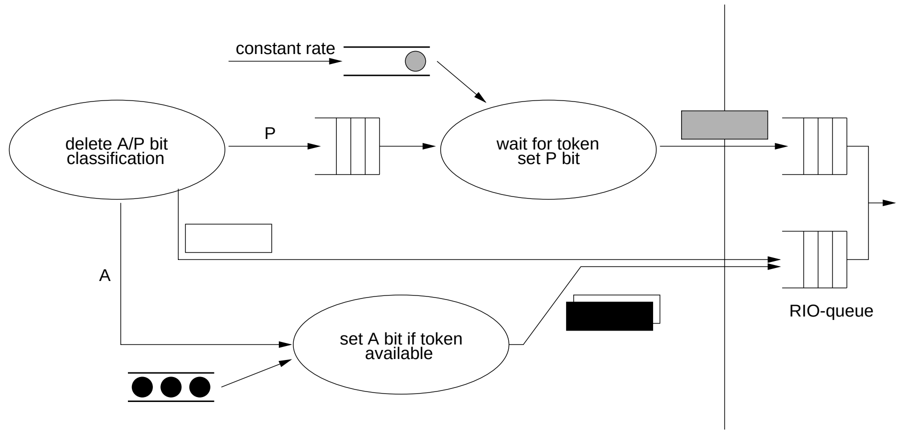

Figure 7 First-hop router for Premium and Assured Service

3.3 Router implementation for Assured and Premium Service

The implementation of Assured and Premium Service requires several modifications of the routers. Mainly classification, shaping, and policing functions have to be performed to the router. These functions are necessary at the border between two networks, for example at the transition of the customer network to the ISP or between the ISPs. Service profiles have to be negotiated between the ISPs similar to the transition to the user.

(a) First-hop router

Figure 7 shows the first-hop router function for Premium and Assured Service. Received packets are classified and according to this the A- or P-bit is set if the packet should be supported with Assured or Premium Service. As parameter for the classification source and destination addresses or information of higher protocols (e.g. port numbers) may be used. A pure best-effort packet will directly be forwarded to the so-called RIO-queue. Also, the Assured Service packets get to this queue. The Assured Service packets are checked whether they conform to the service profile. The A-bit will only be kept if the Assured Service bucket contains a token. Otherwise the A-bit will be deleted and the packets are handled as best-effort packets. The RIO-queuing shall guarantee that best-effort packets are dropped prior to Assured Service packets, if the capacity is exceeded.

(b) Border router

Similar to the first-hop router an intermediate router has to perform shaping functions in order to guarantee that not more than the allowed packet rate is transmitted to the ISP. This is important since the ISP will check whether the user remains within the negotiated service profile. The border router in Figure 8 will therefore drop non conforming Premium service and reset the A-bit of non conforming Assured Service

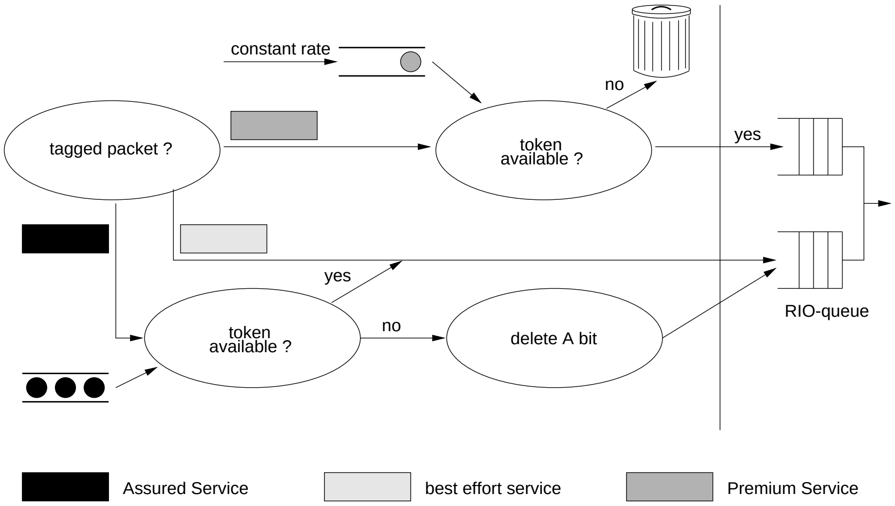

Figure 8 Policing in a border router

packets. Assured Service and best-effort packets share the same queue since both types of packets may belong to the same source. A common queue avoids re-ordering of packets. This is especially important for TCP performance reasons.

(c) Queuing

An important element in the implementation of Premium and especially Assured Service is a proper procedure for dropping packets in overload conditions. To distribute the available bandwidth fairly among the flows in congestion situations, it is recommended to identify and to drop packets of aggressive data flows.

The fundamental mechanism suggested therefore is the Random Early Detection (RED) mechanism. RED is a new technique for router queue management and is supposed to eliminate disadvantages of traditional queuing mechanisms.

With traditional queuing every supported queue accepts packets as long as possible. If there is no space left in a queue arriving packets are dropped, i.e. the packets at the end of the queue are discarded. This method has two significant disadvantages:

- If bursts arrive at nearly full queues, the likelihood for packets of the burst to get lost is high. But queues are also intended for buffering packets in the case of bursts. Therefore, it is recommended to provide space for those bursts.
- Full queues cause higher delays than queues with lower utilization. Especially for real time or interactive applications higher delays are not desired.

RED (Braden *et al.* 1997) is a mechanism trying to keep the queue length below a certain limit in order to provide space for bursts. This is achieved by dropping packets even if the queue length is relatively small (see Figure 9).

Below the lower threshold no packets are dropped. The more the queue length exceeds the lower threshold, the higher is the likelihood for dropping a received packet. The dropping is done randomly to prevent dropping the packets of a certain application data flow.

If the queue length reaches the upper threshold, all packets are dropped. With this mechanism the following advantages shall be achieved:

- Bursts can be supported better since there is always a certain queue capacity reserved for incoming bursts.
- By the lower average queue length the delays are reduced, providing better support for real time applications.

RED is especially capable of dividing the available bandwidth fairly among TCP data flows, as packet loss automatically leads to a reduction of an TCP data flows packet rate. The situation with non TCP conforming data as for example real-time applications based on UDP or multicast applications without an adaptation or flow control mechanism is more problematic. They have to be treated special to prevent them from overloading the network.

The queuing algorithm RIO (RED with In and Out) (Clark *et al.* 1997) has been suggested for Assured Service implementation. RIO is an extension of the RED mechanism. A common queue is provided for in-profile and out-of-profile packets, but different dropping procedures (dropper) are applied. The dropper for out-of-profile packets (out-dropper) drops discards packets earlier i.e. at a substantially lower queue length, than the dropper for in-profile packets, i.e for packets with set A bit. Moreover, the dropping probability of the out-dropper increases more rapidly than the probability of the in-dropper (see Figure 9). This tries to keep the dropping probability of in-profile packets low.

For the implementation of different service types routers have to support several queues, e.g. a queue for Assured or Premium Service. Special bits, e.g. in the TOS field or in the traffic class field of IPv4 respectively IPv6 indicate which service shall be provided to the packet.

3.4 User-Share Differentiation

Based upon packet tagging Premium and Assured Service models can fulfill the stipulated service parameters like bit rates with a high degree of probability only if the ISP network is dimensioned appropriately and non best-effort traffic is transmitted between certain known networks only.

If for instance two users have contracted a bit rate of 1 Mbps for Assured Service packets with an ISP and both wish to receive data simultaneously at a rate of 1 Mbps each from a WWW server which is connected to the network with a 1.5 Mbps link, the requested quality of service cannot be provided.

The User-Share Differentiation approach (Wang 1997) avoids this problem by contracting not absolute bandwidth parameters but relative bandwidth shares. A user

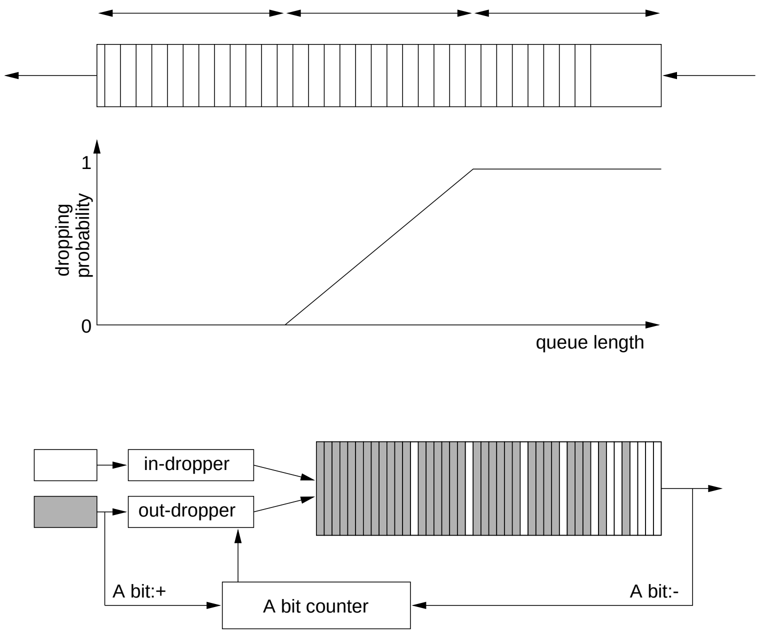

Figure 9 The queuing algorithm RIO

will be guaranteed only a certain relative amount of the available bandwidth in an ISP network. In practice, the size of this share will be in direct relation to the charged costs.

In Figure 10, user A has allocated only half of the bandwidth of user B and one third of the bandwidth of user C. If A and B access the network on low bandwidth links with a capacity of 30 kbps at the same time, e.g. user B will receive a bandwidth of 20 kbps but user C will get merely 10 kbps. If B and C access the same or possibly a different network via a common high bandwidth link with a capacity of 25 Mbps, B will receive 10 Mbps and C only 15 Mbps.

Simpler router configuration is an important advantage of the USD approach. However, absolute bandwidth guarantees cannot supported.

3.5 Olympic Service

The Olympic Service (Nichols *et al.* February 1998) specifies an appropriate service to be deployed within an ISP or a domain. Deployment of this service requires the implementation of a rate-based link share scheduler behavior at each hop. Three service levels are distinguished: gold, silver and bronze. In case of a congested link packets with ”Olympic gold” service will get a larger share of the link than packets sent using the ”Olympic silver” service which in turn get a larger share than packets

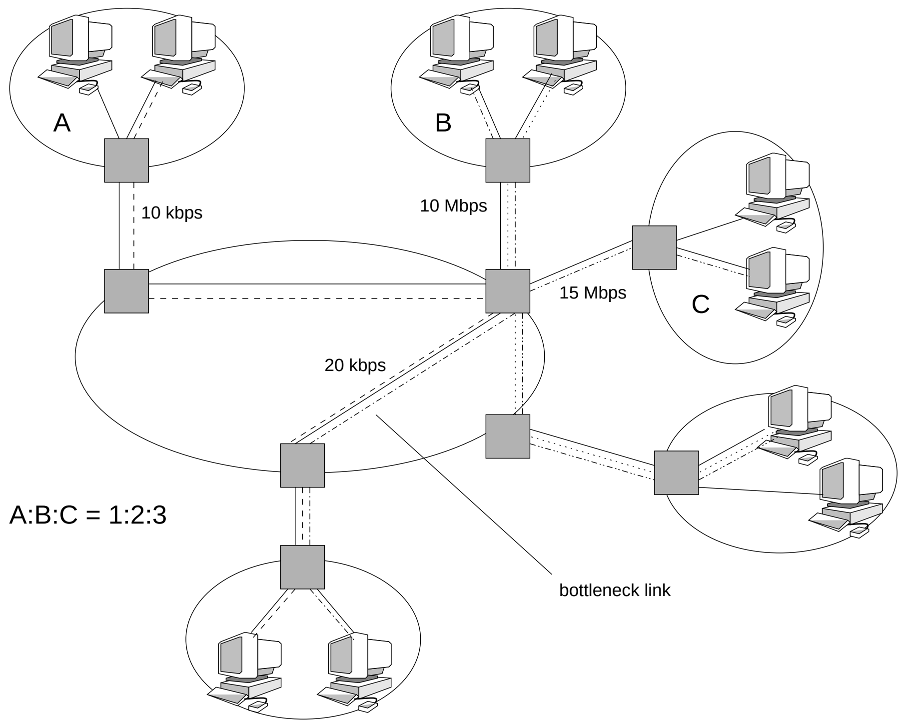

Figure 10 User Share Differentiation (USD)

with "Olympic bronze" service. When there are no packet flows of gold or silver, packets with "Olympic bronze" service may utilize the entire output link.

By marking packets for a link share flows are classified at a boundary. The exact method of service discrimination is not specified but should be selected in a way that it makes a perceptible difference to customers. A possible configuration of the link sharing could be to allocate 60% for gold, 30% for silver and 10% for bronze, although different configurations could be thought of. Customers do not specify a particular traffic profile for the Olympic Service nor is there admission control, shaping and policing of flows in any way.

3.6 Scalable Reservation Protocol

The Scalable Resource Reservation Protocol (SRP) developed at the *Institute for Computer Communications and Applications* (ICA) of ETH Lausanne represents yet another proposal in addition to Assured and Premium Service for a possible implementation of differentiated services in the Internet (Almesberger *et al.* 1997). As indicated by its name much effort has been spent on making the protocol well scalable even for large numbers of packet flows. End systems (i.e. sender and receiver) play an active part in resource reservation while additional control of the sender's behavior is done at the affected routers. Each router aggregates all incoming data flows and monitors this aggregated data stream in order to estimate the necessary resources (now and in future) at that node.

The so-called estimators play an important role in the process of resource reservation. It is their job to estimate the amount of resources needed for reservation. Estimators are deployed in the sender, the receiver and the routers in between. At the sender it helps to make an (optimistic) prediction on the required reservation of network resources for the data to be transported. The estimator of the receiver computes a (conservative) estimation of the resources actually reserved by the network and periodically sends this information back to the sender.

Without requiring explicit signaling of flow parameters the reservation mechanism consists of a reservation protocol and a feedback protocol which will be discussed in the following.

(a) The Reservation Protocol

The reservation protocol is deployed from sender to receiver requiring that sender, receiver and all routers in between have implemented this protocol. Three different packet types are distinguished by a tag to be defined in the packet headers.

REQUEST Packets marked as REQUEST belong to flows wishing to reserve network resources. If a router forwards such a packet he agrees to accept packets tagged as RESERVED in the future at the same transmission rate. Thus, an implicit reservation at the router takes place.

RESERVED If there exists already a reservation at the router and if packets marked as RESERVED arrive at a rate agreed-upon in an earlier stage, the router has to forward them and must not discard them.

BEST-EFFORT No reservations exist in the nodes for these packets, and the packets may be deleted by the routers in case of congestion. This service corresponds to today's best-effort service of the Internet.

A sender wishing to make a reservation begins with the transmission of data packets marked by him as so-called REQUEST packets, which already contain the application data. On arriving at a router they are inspected by admission control functions. They monitor the arriving aggregated flow of packets tagged as RESERVED and estimate the amount of local resources needed to maintain a ”good” quality of service. These resources consist of the available bandwidth, the buffers' sizes and further local resources of the router. When the router receives a packet tagged as REQUEST for forwarding it has to decide whether the QoS will deteriorate by adding the packet to the existing RESERVED-flow. If this is not the case, the packet, which continues to be marked as REQUEST, can be forwarded, and the estimator of the router has to be accordingly updated.

If the necessary additional resources are not available, the packet is degraded to best-effort service by appropriate tagging before being forwarded. In particular, no reservations are performed at the router. Packets marked as BEST-EFFORT or REQUEST may be deleted by a router in case of congestion. An end-to-end reservation is only achieved if packets arriving at the receiver are still marked as REQUEST, i.e. resources are allocated at each router on the transport path. By degrading packets marked as RESERVED in case of insufficient resources at a router, a sender cannot get a better QoS by sending only RESERVED-packets.

Reserved resources need not to be released explicitly by the sender. The estimators in the routers will observe an over-allocation of resources after some time after the end of the flow. They will adjust the estimated share of reserved resources in the routers.

(b) The Feedback Protocol

Periodically, the receiver sends back feedback information to the sender containing the arriving rates of REQUEST- and RESERVED-packets measured at the receiver. To this end a special feedback protocol needs to be implemented, e.g. RTP/RTCP (Schulzrinne *et al.* 1996), in order to notify the sender about the current transmission quality. On receiving this feedback information the sender may begin to send packets tagged as RESERVED while observing a transmission rate based on the received feedback from the receiver. If the sender wishes more resources to be allocated for his flow he can keep on sending packets tagged as REQUEST.

SRP has been tested using simulations, although some topics need further investigation. Policing at network borders and multicasting are not covered in working drafts currently available and are subject of on-going research. The use of SRP for Virtual Private Networks (VPN) is not advised since at each router individual packet flows are aggregated to one large flow which is then treated uniformly. For packet tagging the PHB-field in the DS byte could be used. The necessary code points will be applied for at a future meeting of the DiffServ working group.

4 COEXISTENCE OF DIFFERENTIATED AND INTEGRATED SERVICES

Integrated and differentiated services do not necessarily have to be considered as competing concepts. It is rather advisable to combine both approaches. While differentiated services are recommended for rather large IP networks, the approach chosen for integrated services can be appropriate for limited-size networks, e.g. corporate networks or virtual private networks (VPN).

Both services will be integrated if e.g. VPNs extend over a large IP backbone network. Such a VPN might consist of a client subnet and a server subnet interconnected by a large ISP network, possibly by use of a tunnel. By means of differentiated services techniques both subnets can be linked allocating bandwidth for the traffic between the two subnets.

In such a case it is necessary to map integrated services parameters to differentiated services parameters, similar to the mapping of integrated services parameter

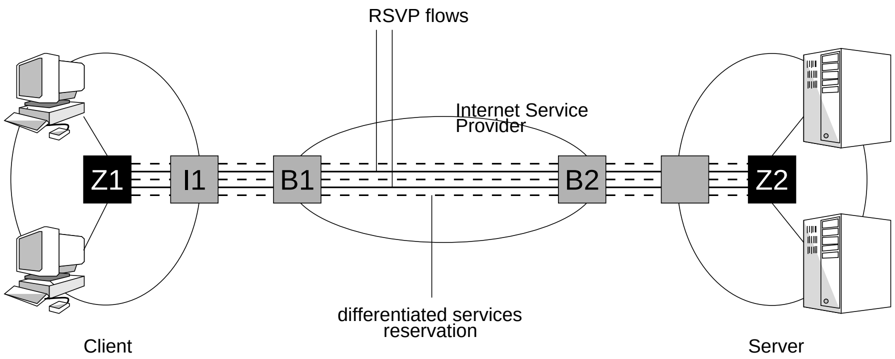

Figure 11 Reservations with RSVP

to ATM parameters or IEEE 802.1p priorities. In the past such mappings have been defined by the Integrated Services Over Specific Link Layers (issll) working group. In the same way a transformation from integrated services to differentiated services has to be made. In this respect Ford *et al.* (1998) suggest to map guaranteed service to Premium Service and controlled load to Assured Service.

A general framework for the integration of differentiated services and integrated services is proposed in (Bernet *et al.* 1998). The order of events for making an RSVP-reservation in a scenario illustrated in Figure 4 is as follows. First, a sender (server) generates PATH messages. In the server network these messages are processed according to the RSVP protocol by the border router *Z2* and other RSVP-routers lying between *Z2* and the sender. In the example a reservation for differentiated services has been made between *Z2* and *Z1*, e.g. a Premium Service with a bandwidth of 1 Mbps. In the network between *Z2* and *Z1* RSVP-messages are transparently forwarded for routers not knowing RSVP. Only the router *Z1* processes the PATH-message again. The message arrives at one of the receivers (clients) who can then make a RSVP-reservation using a RESV-message.

This message is processed again by *Z1* and *Z2*. *Z2* has to check whether the requested reservation (e.g. 600 kbps) is covered by differentiated services reservation. This is for instance the case if no RSVP-reservation over the differentiated services network has been made yet. If there exists already an RSVP-reservation of 500 kbps between the two subnets, the new reservation of 600 kbps cannot be realized and will therefore be rejected by *Z2*. Finally, the RESV-message reaches the corresponding sender of the PATH-packet.

When the sender begins to send the real data, *Z2* has to do the appropriate mapping on a Differentiated Service. For instance, the DS byte in a packet has to be set to the correct PHB value corresponding to Premium Service if a guaranteed service was requested. *Z1* will then reset the DS byte again.

5 SUMMARY AND OUTLOOK

At first the differentiated services model seems to be a highly promising concept to provide qualitatively improved services for the Internet since it avoids the obvious drawbacks of the integrated services architecture. In general, however, guaranteed services for application flows following this approach are not possible. It is questionable whether customers will be satisfied with these kinds of services. It seems to be rather interesting to integrate the two concepts of integrated and differentiated services. An important aspect for the success of differentiated services will be if it will be possible to perform appropriate dimensioning of an IP network in a manner that the available bandwidth on all links will be sufficient to forward all differentiated services packets. This presents a very demanding challenge to network planners.

The tasks in the IETF working groups concerning the standardization consist of defining the precise syntax of the DS byte. Moreover, a definition of the management information bases (MIBs) is needed to create a common basis for the configuration of differentiated services parameters in a router. Finally, all the queuing algorithms based on various differentiated services have to be defined in order to implement these services in heterogeneous router environments.

Up-to-date information on the development in the DiffServ working group and related Internet drafts are available on the official homepage of the working group in the WWW. The mailing list dealing with many aspects of differentiated services and the corresponding mail archive offer a close view at on-going discussions and decisions made by the working group in recent time. The URLs to the mentioned resources on the Internet can be found at the end of this text.

Further investigation has to be done on the support of dynamically changing service requirements. Usually, a customer has to negotiate a service contract with the ISP before making use of a service, e.g. by phone, fax, email or WWW form. The agreed-upon parameters then have to be used by the network operators to configure the routers accordingly. Approaches based on *active networking* could be used for this task, e.g. allowing the customer to run configuration scripts on the routers of the ISP. A different approach would be to use a signaling protocol of the requested service parameters, possibly an adapted version of RSVP.

So far, the deployment of differentiated services for multicast services has been hardly investigated. SRP is one of the few approaches where researchers are considering multicasting explicitly. The difficulties essentially lie in the fact that the total need of bandwidth for an IP multicast flow does not only depend on the transmission rate but also on the size of the multicast group and how the individual group members are spread. The latter two criteria however are very difficult to determine in advance. These parameters may dynamically change because of the receiver-oriented IP multicast concept.

For obvious reasons differentiated services could be implemented using networks with QoS capabilities (e.g. ATM). This of course requires a suitable mapping of differentiated services to ATM services. Especially in the area of ATM different concepts of IP switching are going to establish themselves. However, IP switching tries to bypass IP routers and to forward the packets using switching as often as possible. This in turn may lead to switched packets bypassing shaping and policing functions in the routers, which is inconsistent with the differentiated services architecture. For these scenarios it has to be ensured that either all packets always pass routers with shaping and policing functions or that these functions are realized at so-called ingress and egress routers.

REFERENCES

W. Almesberger, T. Ferrari, J.Y. Le Boudec (November 1997) Scalable Resource Reservation for the Internet (work in progress), *Internet Draft*, draft-almesberger-srp-00.txt.

W. Almesberger, J.Y. Le Boudec, T. Ferrari (March 1998) Encoding of SRP packet types in the DS byte (work in progress), *Internet Draft*, draft-watfjyl-srp-ds-00.txt.

F. Baker, S. Brim, T. Li, F. Kastenholz, S. Jagannath, J. Renwick (April 1998) IP Precedence in Differentiated Services Using the Assured Service (work in progress), *Internet Draft*, draft-ietf-diffserv-precedence-00.txt.

Y. Bernet, R. Yavatkar, P. Ford, F. Baker, L. Zhang (April 1998) A Framework for End-to-End QoS Combining RSVP/Intserv and Differentiated Services (work in progress), *Internet Draft*, draft-bernet-intdiff-00.txt.

B. Braden, D. Clark, J. Crowcroft, B. Davie, S. Deering, D. Estrin, S. Floyd, V. Jacobson, G. Minshall, C. Partridge, L. Peterson, K. Ramakrishnan, S. Shenker, J. Wroclawski, L. Zhang (March 1997) Recommendations on Queue Management and Congestion Avoidance in the Internet (work in progress), *Internet Draft*, draft-irtf-e2e-queue-mgt-00.txt.

R. Braden, L. Zhang, S. Berson, S. Herzog, S. Jamin (September 1997) Resource ReSerVation Protocol (RSVP) -Version 1 Functional Specification, *Request for Comments*, 2205.

D. Clark, J. Wroclawski (July 1997) An Approach to Service Allocation in the Internet (work in progress), *Internet Draft*, draft-clark-diff-svc-alloc-00.txt.

P. Ford, Y. Bernet (March 1998) Integrated Services Over Differentiated Services (work in progress), *Internet Draft*, draft-ford-issll-diff-svc-00.txt.

A. Mankin, F. Baker, B. Braden, S. Bradner, M. O'Dell, A. Romanow, A. Weinrib, L. Zhang (September 1997) Resource ReSerVation Protocol (RSVP) Version 1 Applicability Statement, *Request for Comments*, 2208.

K. Nichols, S. Blake (February 1998) Differentiated Services Operational Model and Definitions (work in progress), *Internet Draft*, draft-nichols-dsopdef-00.txt.

K. Nichols, S. Blake (May 1998) Definition of the Differentiated Services Field (DS Byte) in the IPv4 and IPv6 Headers (work in progress), *Internet Draft*, draft-ietf-diffserv-header-00.txt.

H. Schulzrinne, S. Casner, R. Frederick, V. Jacobson (January 1996) RTP: A Transport Protocol for Real-Time Applications, *Request for Comments*, 1889.

Z. Wang (November 1997) User-Share Differentiation (USD) Scalable bandwidth allocation for differentiated services (work in progress), *Internet Draft*, draft-wang-diff-serv-usd-00.txt.

6 INTERESTING WWW-PAGES OF THE IETF

- Official homepage of the DiffServ working group: http://www.ietf.org/html.charters/diffserv-charter.html
- Mailing list of the DiffServ working group: diff-serv@baynetworks.com
- Information and Internet drafts of the DiffServ working group: http://diffserv.lcs.mit.edu

7 BIOGRAPHY

Florian Baumgartner M.S. (1997) in information science and physics at the University of Regensburg (Germany). Since 1998 research assistant at the Institute of Computer Science and applied Mathematics with research interests in the simulation and optimization of networks, differentiated services and active networks.

Torsten Braun Diploma (1990) and Ph.D. (1993) degree in Computer Science at the University of Karlsruhe. 1990–1994 research assistant at the Institute of Telematics of the University of Karlsruhe. 1994–1995 visiting scientist at the Institut National de Recherche en Informatique et en Automatique (INRIA) in Sophia-Antipolis (France). 1995–1997 guest scientist and senior consultant at the IBM European Networking Center at Heidelberg (Germany). Since 1998 professor of Computer Science at the University of Berne (Switzerland).

Pascal Habegger M.S. (1996) in Computer Science at the University of Berne. Since 1996 Ph.D. candidate at the Institute of Computer Science and Applied Mathematics with research interests in network simulations, differentiated services, QoS, ATM.