

A SURVEY OF INTERNET QoS SIGNALING

DIMITRA VALI, HELLENIC TELECOMMUNICATIONS ORGANIZATION — OTE S.A.

SARANTIS PASKALIS AND LAZAROS MERAKOS, UNIVERSITY OF ATHENS

ALEXANDROS KALOXYLOS, UNIVERSITY OF PELOPONNISOS

ABSTRACT

The need for real-time services support over the Internet drives research efforts toward the provisioning of quality of service (QoS) guarantees in IP networks. Among the various aspects of QoS provisioning, QoS state establishment and maintenance in intermediate routers is a major factor, responsible for dynamic resource allocation. The dynamic manipulation of QoS state is possible through the utilization of appropriate QoS signaling that triggers the respective resource allocation in QoS-capable network elements. At first, protocol design assumed a homogeneous underlying network and resulted in end-to-end QoS protocols that applied a specific QoS configuration in all routers along the path. However, the need for accommodating network heterogeneity and flexibility led to a two-tier resource management model that utilizes separate signaling for intra- and inter-domain reservations and requires different signaling processing in domain interior and border routers. This article gives an overview of the QoS signaling protocols designed for the Internet and describes their characteristics. Moreover, the identified protocols are classified depending on their applicability for intra- or inter-domain usage. A comparison of the various protocols based on some common signaling elements is also provided, and future trends in the Internet QoS signaling area are identified.

The Internet was designed as a best-effort network, offering no QoS assurances for the supported services. However, the imminent dominance of the Internet Protocol (IP) as a de-facto telecommunication standard is the leading factor in an increasing demand for the efficient support of real-time services over the Internet with QoS assurance. The most prominent methods for offering QoS in the Internet are network over-provisioning, traffic engineering, and differentiated packet treatment inside routers.

The advent of fiber and the latest advances in optical networks and router technology that allow for excess network and router capacity made over-provisioning a possible approach for QoS support over the Internet. Excess resources, however, can prove to be rather expensive, especially in access networks (wired or wireless) where bandwidth is a scarce resource. On the other hand, bandwidth availability and router capacity in general cannot be considered as an infinite resource, considering that the increasing penetration of broadband access technologies to the users (e.g., xDSL, Fiber to the Home, or Fiber to the Curb technologies) is expected to increase significantly the resources required in the core networks.

For the aforementioned reasons, efficient mechanisms for supporting end-to-end QoS through the Internet should be

developed and implemented. One way of supporting QoS in the Internet is through traffic engineering, where routing of QoS packets does not follow the traditional IP routing protocols (i.e. OSPF and BGP), but instead takes into account available resources and expected traffic on the various network links. Following this approach, some paths can be over-provisioned and used for the most demanding packet flows (marked with the appropriate label), whereas others could be left for the best-effort traffic. A traffic engineering technique that can provide service differentiation is MPLS. MPLS is a forwarding scheme that uses a fixed-length label inside the packet's header to decide packet handling. MPLS-capable routers are configured to forward packets with the same label to the same outgoing interface. In this manner label-switched paths that correspond to different QoS classes can be built across the Internet, where all packets marked with the same label follow the same route.

In this article, however, we focus on QoS mechanisms that maintain the standard routing functionality, and rely on differentiated packet treatment inside routers to support end-to-end QoS provision. Following this approach, QoS is assured via adequate packet classification, queuing, and scheduling inside routers that reflect the QoS characteristics of the packet. The QoS-related configuration of routers along the data

path can take place via the end-to-end exchange of messages, i.e., QoS signaling. Signaling messages are processed in the routers along the path and enable the reservation of resources for the requested QoS service. Moreover, signaling exchange allows for the negotiation of requested and received QoS characteristics between the initiating and the receiving user and between the initiating user and the network.

An independent but important aspect in QoS provision is the existence and operation of policy and admission control inside the routers. Policy control determines whether the requesting user is entitled to make the requested reservation, while admission control determines whether the node has sufficient resources to facilitate the reservation. If both checks succeed, the requested reservation can be established by configuring the respective router parameters. Policy and admission control can also be centralized, where only minimal enforcement stubs operate in each node. However, policy/admission control interaction with QoS signaling is minimal and indirect (related to error cases).

The result of fertile research work in the area of Internet QoS signaling for more than a decade has resulted in a plethora of proposed mechanisms for allocating resources either inside a domain (intra-domain) or in between domains (inter-domain). This article gives an overview of the various QoS mechanisms, describes the major QoS protocols, classifies them into broader signaling categories, and attempts a comparison between them based on their individual characteristics. Major trends and principles in the QoS signaling area are also discussed.

The rest of the article is organized as follows. We first present QoS background issues, and then we identify and describe the single-tier resource management architecture and the respective end-to-end QoS signaling protocols. Then we present the two-tier signaling architecture and its protocol representatives for both intra- and inter-domain QoS signaling. We summarize the presented approaches and discuss their merits and drawbacks, and we present other efforts to give an overview of the Internet QoS area. Finally, we conclude the article.

BACKGROUND

The Internet community soon realized the vision of end-to-end QoS services and introduced the Integrated Services (IntServ) architecture [1] to implement this vision into specifications. IntServ supports end-to-end signaling, QoS state establishment, and management for per-flow differentiated treatment in intermediate routers along the data path. The signaling protocol that emerged to meet the integrated services requirements is RSVP (Resource reSerVation Protocol) [2]. The IntServ architecture was designed to facilitate every QoS element (router functionality, signaling, and accounting) in a fine-grained manner. To achieve this goal, IntServ was founded on the underlying assumption that a homogeneous Internet environment equipped with IntServ-enabled routers and end hosts would be the common case.

The IntServ architecture in general and the RSVP protocol in particular received criticism, mainly due to the scalability issues raised by the state maintenance for every data flow in intermediate routers across the end-to-end path. Therefore, the Internet community considered other alternatives to the QoS provision problem. This time the target was a lightweight QoS architecture putting as little burden on the routers as possible and providing coarse-grained traffic prioritization based on the statically contracted service level agreements (SLAs) between users and the network. SLAs specify the

amount and types of traffic each side has agreed to send and receive. The outcome was DiffServ (Differentiated Architecture) [3].

DiffServ networks are statically configured to support a small set of QoS levels (per hop behaviors (PHBs)) and do not use any QoS signaling for state establishment and maintenance in routers. DiffServ routers prioritize the data packets according to a six-bit field in the IP packet header (DiffServ code point (DSCP)) that reflects the requested QoS level. This procedure results in aggregating reservations for different users sharing the same QoS level. Appropriate packet marking takes place either at end-hosts or at DiffServ edge routers before the traffic enters the DiffServ network. In addition, DiffServ edge routers perform traffic classification and traffic conditioning procedures (including metering, marking, shaping, and policing) based on the contracted SLAs.

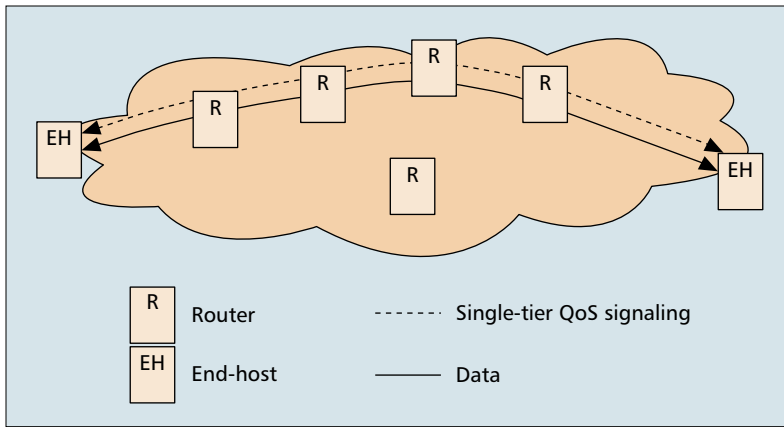
The DiffServ model is simple and avoids processing complexity and signaling overhead in network routers (especially in interior routers). Furthermore, it does not mandate any specific QoS protocol implementation to the end-user application. DiffServ, however, is rather rigid since the users cannot dynamically change the amount of reserved resources according to their current traffic requirements. Moreover, the aggregation of different flows belonging to the same QoS level can result in an unfair distribution of resources among the flows inside the same aggregate (e.g., due to the aggressiveness of some bursty flows). For all the aforementioned reasons, statically configured QoS domains, unless being highly over-provisioned, may fail to offer the agreed QoS to the users.

Furthermore, the realization of the fact that the Internet is a concatenation of technologically and administratively different domains (autonomous systems (ASs)) led to the identification of separate QoS techniques for the efficient support of intra- and inter-domain QoS. Thus, a two-tier resource management model was proposed in [4], with the intra-domain QoS signaling performing resource management inside a domain, and the inter-domain signaling managing resource allocation between domains. The two tiers must be closely coordinated to enable provision of the necessary end-to-end QoS support. The two-tier model increases the degrees of freedom regarding end-to-end QoS support, since each domain is free to choose any QoS support mechanism for allocating resources internally, as long as proper co-operation takes place with the respective inter-domain signaling protocol.

SINGLE-TIER QoS SIGNALING

Single-tier signaling offers end-to-end QoS guarantees assuming a homogeneous QoS architecture for the Internet, where all routers support the same QoS mechanisms. Single-tier signaling traverses the end-to-end path between the communicating users following standard IP routing, and establishes and maintains a QoS state in every intermediate IP router. The subsequent data packets follow the same route with the QoS signaling and take advantage of the specific QoS behavior that the signaling has configured in routers along the path (Fig. 1).

A major representative of single-tier signaling is RSVP, designed for application in a homogeneous end-to-end IntServ architecture. RSVP provides uni-directional reservation of resources for each application data flow, and adopts a receiver-based resource reservation approach for QoS sessions, suited for handling reservations in a multicast environment. Resources are reserved and respective soft states are installed for each data flow on RSVP-aware routers along the path, using a two-pass signaling scheme (Fig. 2). The sender initi-



■ Figure 1. Single-tier QoS signaling.

ates the signaling procedure by sending an end-to-end PATH message (message 1 in Fig. 2) toward the receiver, in order to inform the receiver about the traffic shape of the data flow and the available resources of routers along the path. The traffic characteristics of the source and the requested QoS service type are represented by specific objects inside the RSVP messages. The PATH message installs route-specific soft states in all traversing routers, in order to enable the response messages to follow the reverse path back to the sender. The receiver responds with a RESV message (message 2 in Fig. 2) that performs the actual reservation of resources along the path. Admission and policy control takes place in each traversed RSVP router. The RESV message is sent hop-by-hop between routers following the reverse path, according to the route-specific state established by the PATH message. Merging of the various reservation messages heading for the same sender takes place along the route, allowing for efficient multicast reservations. As a result of the aforementioned signaling procedure, reservation and signaling soft states are installed in all RSVP-enabled routers along the path, which are maintained by periodic end-to-end refresh messages.

While RSVP succeeds in assuring QoS separately for each data flow, it has received criticism regarding the complexity and the processing overhead it implies in routers. Since the amount of states stored in each router increases linearly with the number of traversing QoS flows, the number of states stored inside routers can raise scalability issues, especially in backbone routers that handle a large volume of traffic. Moreover, the use of refresh messages transmitted periodically for each flow has been accounted for contributing to signaling and processing overhead inside the network. The latter has been addressed in [5] where, among other optimizations, the use of a Bundle message is proposed, which consists of a bundle header and a body carrying multiple RSVP messages dealing with separate sessions.


As an alternative to RSVP, YESSIR (YET another Sender Session Internet Reservation) [6] is proposed for uni-directional per-flow reservation within the IntServ QoS model. YESSIR is built as an extension of the RTCP (Real Time Control Protocol) protocol, an in-band control protocol for RTP flows. The YESSIR design is based on the assumption that a large fraction of applications requiring guaranteed quality of service are real-time applications, and as such will use the RTP (Real-time Transport Protocol) to deliver their data. RTCP is used for the transportation of the YESSIR reservation messages. Soft reservation states are installed on each

router along the path, as the RTCP/YESSIR message travels from the sender to the receiver (Fig. 3). Reservations are installed along the path following a one-pass sender-based reservation mechanism.

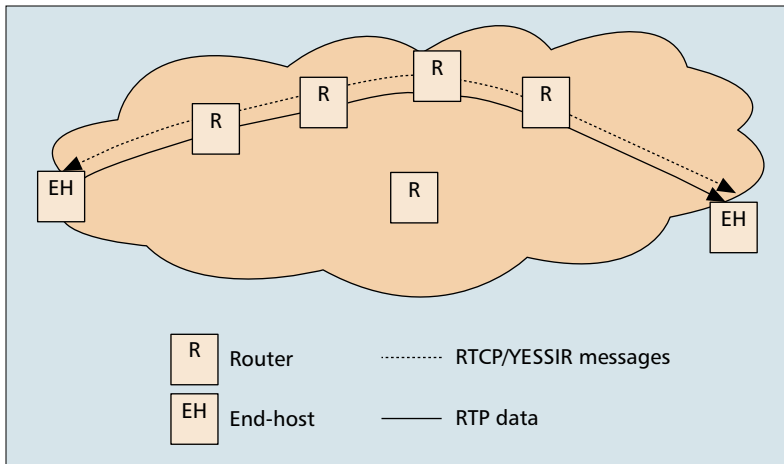
YESSIR seems to perform better than RSVP in terms of reservation set-up time, signaling processing, and message overhead. Measurements [6] indicate that YESSIR reservation setup time is three times faster than that of RSVP, and that the YESSIR refresh message processing overhead is approximately half of that of RSVP refresh. Moreover, YESSIR is designed to take advantage of the RTP and RTCP protocol features and avoids the implementation of an independent reservation protocol.

Boomerang [7] is another reservation protocol designed with simplicity in mind. Boomerang establishes per-flow bi-directional reservations, using a single protocol message. The Boomerang reservation message is sent by the initiating node (which may be the sender or the receiver) to the far-end node, where it is echoed back to the initiating node. Resources are allocated hop-by-hop (using soft states) in all routers traversed by the reservation message in both directions (Fig. 4). When the initiating node receives the Boomerang message, it verifies the success of the reservation by examining the appropriate message flags set in the message by the routers along the path. Reservation messages are then sent periodically from the initiating to the far-end node to keep the reservation alive along the upstream and downstream paths.

The Boomerang message is routed according to standard routing procedures. This ensures that the reservation will be made along the correct path for both upstream and downstream traffic and eliminates the need for storage of route-specific states along the path. Signaling states are also eliminated from the routers, due to the use of a single reservation message. However, per-flow reservation states are still installed in each router along the path. Boomerang's short messages, along with the elimination of routing and signaling states, contribute to an enhanced protocol performance, especially in terms of processing overhead. Based on implementation results [7], Boomerang signaling overhead (in terms of bytes/sec) is about four to six times lower than that of RSVP. Boomerang also results in lower memory and processing requirements in routers when compared to RSVP. The Boomerang memory gain is around 400 bytes/flow and the processing gain ranges between two and four, depending on the specific type of message processed.



■ Figure 2. RSVP message exchange.



■ Figure 3. YESSIR message exchange.

Similar to RSVP, both YESSIR and Boomerang propose the storage of per-flow QoS states in all routers along the path, which is a major point of criticism due to scalability reasons. Regarding Boomerang, however, a protocol extension has been proposed in [7] for aggregating per-flow states in routers, when necessary. In addition, both YESSIR and Boomerang follow the soft state approach introduced by RSVP. As a result, they use refresh mechanisms for keeping the reservations alive, which adds to the signaling overhead. Moreover, YESSIR caters to RTP traffic only.

INSIGNIA [8] is a QoS protocol proposed for the support of adaptive services in mobile ad-hoc networks. INSIGNIA is an in-band signaling system that supports restoration and adaptation of the reserved resources to the continuously changing conditions of the ad-hoc networks. The INSIGNIA control messages are transported inside IP data packets (using IP header options), which follow the dynamic routing protocols proposed for ad-hoc networks, and install per-flow soft states in traversed routers. In general, in-band signaling systems are considered to be well suited to the rapidly changing environment of mobile ad-hoc networks, enabling fast resource reservation and restoration after topology changes. Simulation results study the effect of mobility on the INSIGNIA network performance (i.e., percentage of packets delivered as QoS packets, as best-effort packets, or being lost) and show that INSIGNIA supports relatively constant QoS delivery (approximately 80 percent of packets receive QoS) under slow and moderate mobility conditions (between three/six and 18 km/h) [8].

Overall, single-tier signaling assumes a homogeneous QoS architecture being applied end-to-end in the Internet, which is not usually the case. This simplified assumption, however, eliminates any interworking issues between different QoS protocols in the domain boundaries, adding to the simplicity of the QoS protocol implementation. The simplicity of this signaling category is further emphasized by the use of standard routing procedures throughout the network for routing the QoS signaling messages.

Concluding this section, we would also like to present an alternative method for supporting QoS that does not, however, directly involve the reservation of resources. This proposal suggests the use of the Session Initiation Protocol (SIP) [9]. SIP is a request-response protocol for initiating and managing communications sessions in the Internet. SIP intentionally does not involve itself with the reservation of resources and can,

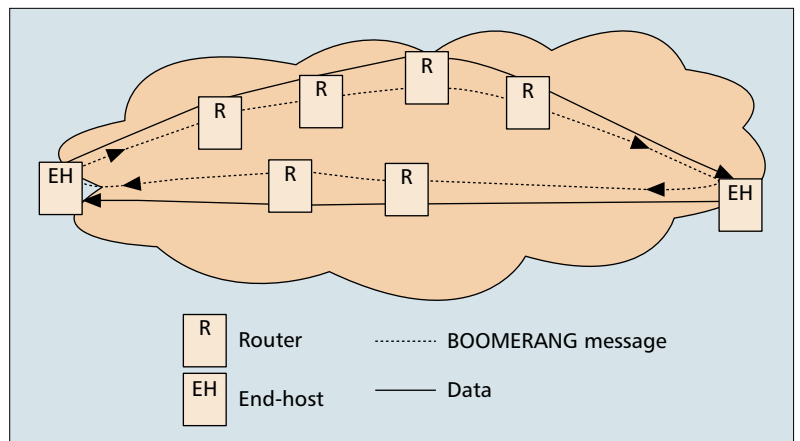
in principle, work independently from the resource reservation protocol that may be used along the data path. Regarding QoS, SIP could provide the transport mechanism enabling QoS negotiation, policy, and AAA (authentication, authorization, and accounting) enforcement. QoS negotiation is enabled via the transport inside the SIP messages of Session Description Protocol [10] objects that describe the capabilities of the end-hosts and the characteristics of the connection (e.g., media, bandwidth codec, etc.). Regarding AAA/policy enforcement, SIP provides a universal transport means to carry AAA/policy requests to the local AAA/policy server, which has the authority to grant or deny access to local (edge router) or remote resources (through brokering with peer servers). SIP does not perform the actual reservation of

resources, which is assured by the underlying reservation protocol (e.g., RSVP). It is important, however, that appropriate coordination exists between the two protocols so that requested resources are reserved before the data communication between the end systems begins [11]. Moreover, to alleviate the user terminal from QoS resource functionality implementation, SIP extensions are proposed in [12] (Q-SIP protocol). These extensions allow for terminal QoS request information to be transported inside Q-SIP messages. Q-SIP messages are addressed to specially designed proxies inside the network that are responsible for performing the actual reservation of resources.

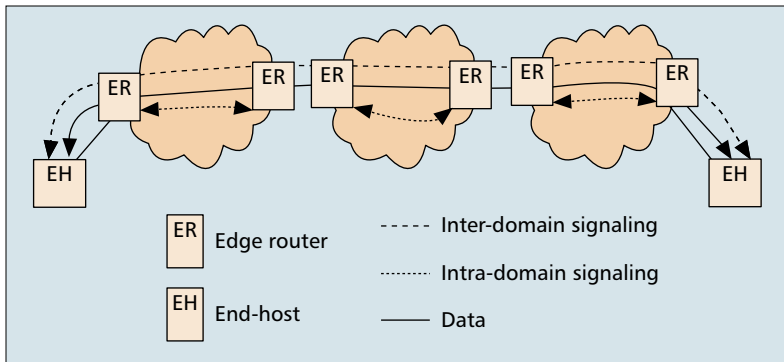
TWO-TIER QOS SIGNALING

The Internet is a concatenation of different autonomous systems that are administratively and technologically independent. Taking this into account, the two-tier resource management model is proposed [4], where the need for two separate resource management categories is identified: one inside the administrative domain (intra-domain signaling) and one between domains (inter-domain signaling). Appropriate combination of the intra- and the inter-domain signaling results in efficient end-to-end management of resources.

The two-tier signaling architecture, illustrated in Fig. 5, implies that each domain is allowed to use its own QoS mechanism or protocol internally, allowing for concatenation of the various heterogeneous domains. The provision, however, of end-to-end QoS requires that appropriate interworking



■ Figure 4. Boomerang message exchange.



■ **Figure 5.** Two-tier QoS signaling.

between the intra- and the inter-domain QoS protocols take place at the domain boundaries.

In the following paragraphs, the major intra- and inter-domain signaling protocols are presented.

INTRA-DOMAIN QoS SIGNALING

Intra-domain QoS signaling is the signaling used for performing resource allocation inside a domain. The purpose of intra-domain resource allocation is to determine whether sufficient resources are available for traffic flowing through each domain and, if so, to allocate resources for this traffic. Intra-domain signaling performs resource reservation based on the user's QoS requirements that it receives via interaction with the end-to-end inter-domain QoS signaling. Appropriate QoS parameter mapping takes place at the domain boundaries or at specific resource management entities inside the domain.

The simplest way to allocate resources inside a domain is through static configuration of resources inside routers for a small number of QoS levels according to the DiffServ QoS model. Following this approach, no signaling is required for allocating resources inside the domain. However, as already mentioned, statically configured QoS domains may fail to offer the agreed QoS to the user, unless highly over-provisioned. Over-provisioning may be particularly expensive, especially in access domains, where the use of QoS signaling is advised for the reliable offer of the agreed QoS and the efficient management of resources.

As an alternative to static router configuration for QoS, the IntServ architecture can be deployed in the domain interior, relying on RSVP¹ for dynamic resource reservation inside the domain [4]. However, per-flow reservation signaling results in a number of QoS states stored inside routers that is analogous to the number of traversing flows. Therefore, it seems that IntServ would be appropriate for access domains where the number of QoS flows is manageable but the resources are scarce, while DiffServ would be a good candidate for core domains where over-provisioning is possible and the number of traversing flows is rather large.

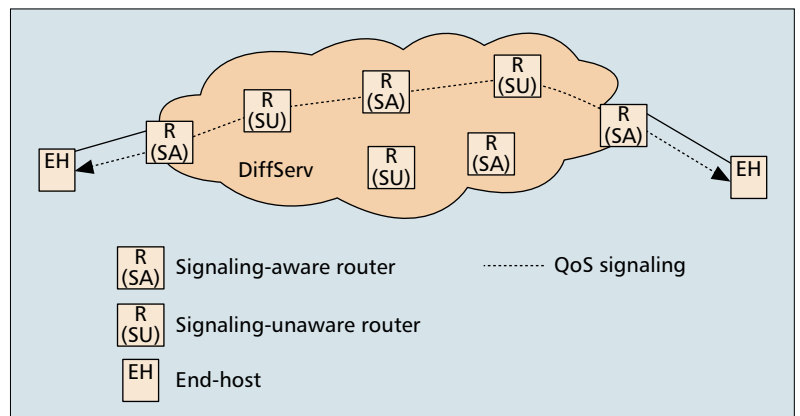
In an attempt to combine the simplicity of DiffServ with the dynamic resource management enabled by QoS signaling, a number of signaling-aware routers can be introduced inside a DiffServ

¹ In principle, there does not seem to exist any technical barrier for deploying any single-tier protocol (e.g. YESSIR or Boomerang) for performing intra-domain reservation. However, this scenario as well as the respective requirements for interworking with inter-domain protocols has not been addressed in previous research work and, therefore, is not considered in the current study.

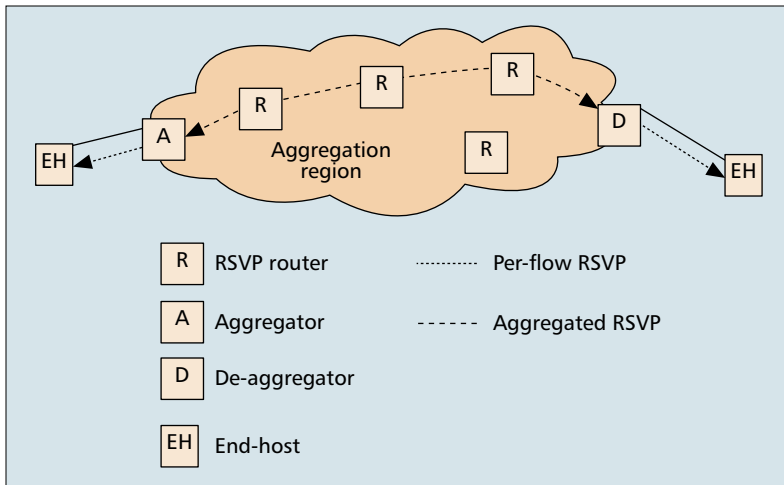
domain. While DiffServ domains are in general signaling-unaware, the DiffServ signaling-aware routers could participate in QoS signaling and perform admission control and aggregate (PHB-based) resource allocation for the DiffServ domain (Fig. 6). The number of signaling-aware routers inside the DiffServ domain can vary from only edge routers to all interior routers [13], depending on the specific characteristics of the domain, as well as the desired degree of reliability in offering the agreed QoS. This approach enables per-flow admission control and dynamic allocation of resources to the various DiffServ QoS levels.

A first approach for enabling dynamic resource allocation inside a signaling-aware DiffServ domain is through the use of RSVP [13]. RSVP messages install per-QoS-level reservation states in RSVP-aware routers while they cross transparently the RSVP-unaware routers. Although RSVP-aware nodes in the DiffServ region participate in RSVP signaling, they perform classification and scheduling of the subsequent data traffic based on the packets' respective DSCPs, allowing for data and reservation state aggregation based on the DiffServ principles. Aggregation, however, does not take place in the control plane, since per-flow RSVP signaling states are stored inside routers. The relative number of RSVP-aware routers in the DiffServ region should, therefore, be carefully selected to keep the state maintenance in the DiffServ network relatively simple.

Moreover, to enable aggregation in the control plane of a dynamically managed DiffServ domain, RSVP aggregation can be used between domain edges [13]. Aggregated RSVP [14] is an RSVP extension designed for handling aggregate reservations that cross an aggregation region (e.g., a DiffServ region) and share common ingress and egress routers (aggregator and de-aggregator). Aggregated RSVP assumes end-to-end RSVP signaling exchange between users and the existence of a number of RSVP-aware routers inside the aggregation region. Aggregate RSVP messages are exchanged between the aggregators and de-aggregators located at the edges of the aggregation region (Fig. 7), as a response to the per-flow RSVP messages. Aggregate RSVP messages result in storing per-aggregate signaling and reservation states in routers inside the aggregation region. Initial reservation levels for each aggregate can be established between edge routers, based on anticipated traffic patterns. The end-to-end RSVP messages should cross transparently the aggregation region, so that RSVP routers inside the aggregation region do not process them. For this reason, the IP protocol number



■ **Figure 6.** Dynamic resource management inside a DiffServ domain.



■ Figure 7. RSVP aggregation.

of the end-to-end RSVP messages is changed upon entering the aggregation region (at the aggregator) and is restored at the de-aggregator.

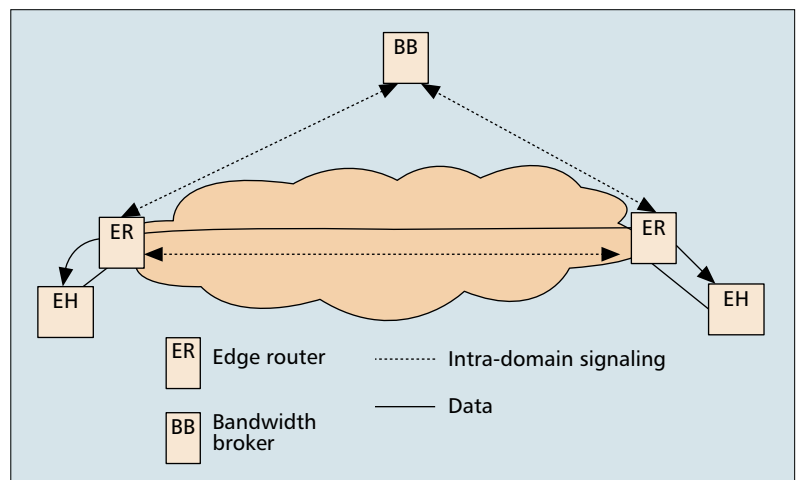
Another framework designed for edge-to-edge dynamic resource allocation inside a DiffServ domain is the Resource Management in DiffServ (RMD) framework [15]. The RMD framework specifies the Per Domain Reservation (PDR) and Per Hop Reservation (PHR) protocols that are both triggered at the ingress DiffServ edge router by the end-to-end inter-domain QoS protocol (e.g. RSVP). At DiffServ edge routers the inter-domain protocol QoS parameters are mapped to the DiffServ DSCP. The PDR protocol operates edge-to-edge within the DiffServ domain and establishes dynamic per-flow states at the edge routers, enabling per-flow admission control on behalf of the DiffServ domain. The PDR protocol is only processed by the DiffServ edge nodes. The PHR protocol operates hop-by-hop between interior DiffServ routers and reserves requested resources inside the respective DiffServ QoS levels. DiffServ interior routers store per-QoS level reservation states. PDR messages are usually carried encapsulated inside the PHR messages. The RMD framework message sequence is based on a request-response procedure, where reservation requests are sender-initiated and Acknowledgments are sent back by the receiver. RMD reservation states can be either soft or hard states. A refresh mechanism is used for the soft state maintenance. Simulation results show that the mean processing delay of PHR reservation messages are more than 1300 times smaller than that of the RSVP reservation message [15]. This is due to the RMD aggregate state maintenance that significantly reduces the look-up time in the reservation state table, compared to RSVP per-flow state maintenance.

A similar effort for offering QoS guarantees inside a DiffServ domain is the DiffRes reservation protocol [16]. DiffRes uses a combination of end-to-end (or edge-to-edge) and hop-by-hop messages to dynamically perform aggregate reservations inside a DiffServ domain. For each flow, after collecting resource availability information along the route, the sender (or the ingress router) issues a reservation request message that performs the actual resource reservation in the DiffServ routers along the route. Per-flow soft states are stored in the DiffServ edge routers while per-QoS level states are installed in the core routers. DiffServ core routers add the requested amount of resources to the respective QoS level. In addition,

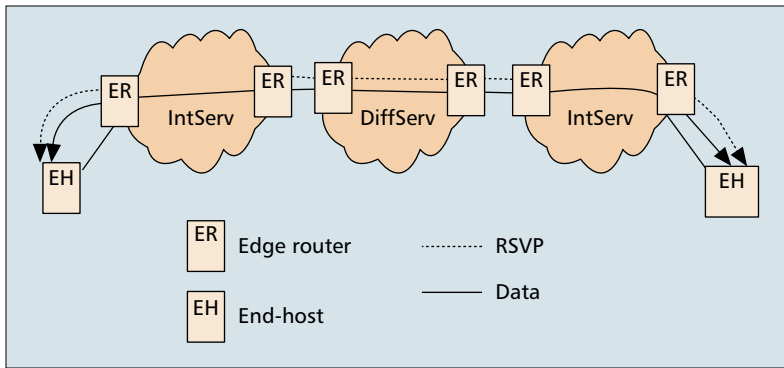
DiffRes employs a mechanism of hop-by-hop Acknowledgments for avoiding duplication of reservations, in case an end-to-end reservation message is lost somewhere along the route. However, this mechanism requires the storage of short-lived per-flow states in DiffServ interior routers (for the time period of the reservation establishment) and a per-hop router Acknowledgment overhead.

The RSVP aggregation, RMD, and DiffRes protocols are generally triggered by domain edge routers, and enable dynamic per-flow admission control inside the DiffServ network region, while avoiding per-flow QoS state storage in the signaling-aware DiffServ core routers. The described mechanisms generally follow the soft state approach, i.e., they rely on refresh message exchange for maintenance of stored states. In addition, DiffRes provides a mechanism for avoiding duplication of reservations in case of lost reservation messages, requiring additional storage and processing overhead.

A centralized approach to intra-domain signaling is offered by the bandwidth broker (BB) architecture (Fig. 8) [17]. The BB is a logical domain entity aware of the domain's policies and available resources and charged with two major responsibilities. The first responsibility is to manage network resources on behalf of the domain by setting the parameters of the domain routers; the second responsibility is to manage the inter-domain link resources and signaling to other BBs. More specifically, the BB performs QoS parameter mapping, admission control, and resource management for the domain as a response to the end-to-end inter-domain (inter-BB) QoS signaling protocol. For intra-domain resource allocation, the BB could use custom protocols (e.g., SNMP, LDAP) to directly allocate resources to each domain router [17, 18], or it can trigger and delegate the allocation of resources to an intra-domain resource management protocol (e.g. RSVP) [4]. Resource reservation inside BB-based domains can take place either on a per-flow (access domains) or a per-aggregate (core domains) basis. An intra-domain aggregation scheme is also proposed, where aggregate allocations using RSVP take place between all ingress/egress router pairs of a core domain based on aggregate traffic measurements [4]. Simulation results show that this aggregation scheme offers preferential treatment to QoS packets without starving the best-effort traffic.



■ Figure 8. Intra-domain QoS signaling in BB architectures.



■ **Figure 9.** End-to-end inter-domain RSVP usage.

In general, BB-based architectures leverage routers from admission control procedures and storage burden. However, dynamic configuration of domain routers by the central BB may prove to be complex. Moreover, BB-based architectures suffer from the weaknesses that characterize all centralized architectures, including single point of failure, high processing load in centralized nodes, etc.

INTER-DOMAIN QoS SIGNALING

Inter-domain signaling is used end-to-end in a two-tier resource management architecture for handling the reservations between neighboring domains. Moreover, inter-domain signaling is used for sending the end-user's QoS requirements to the internal QoS mechanisms of the transit domains.

The end-to-end (triggered by end-hosts) inter-domain usage of RSVP for resource reservation in mixed IntServ/DiffServ networks is proposed in [13]. According to this proposal, IntServ can be applied in access network domains, whereas DiffServ can support aggregation of traffic in core domains. At the boundaries between domains, special RSVP-aware routers (that may be part of either the IntServ or the DiffServ network) perform per-flow admission control for the DiffServ domain, based on resource availability inside the DiffServ network and on the customer-defined policy. If the DiffServ domain is RSVP-unaware, the RSVP messages will cross transparently the DiffServ domain (Fig. 9). To facilitate the RSVP usage over DiffServ networks, the introduction of the DCLASS object is proposed to carry the agreed DSCP values within RSVP messages [19], enabling the sender's data packet marking with the appropriate DiffServ QoS level.

However, to avoid RSVP per-flow processing in core routers, various protocols have been designed to aggregate reservations on backbone inter-domain links. The aggregation of reservations is designed to reduce storage and processing cost in core routers. This cost is related to the number of aggregate states maintained in the routers and the number of signaling messages that are necessary for the establishment and maintenance of these states (signaling load).

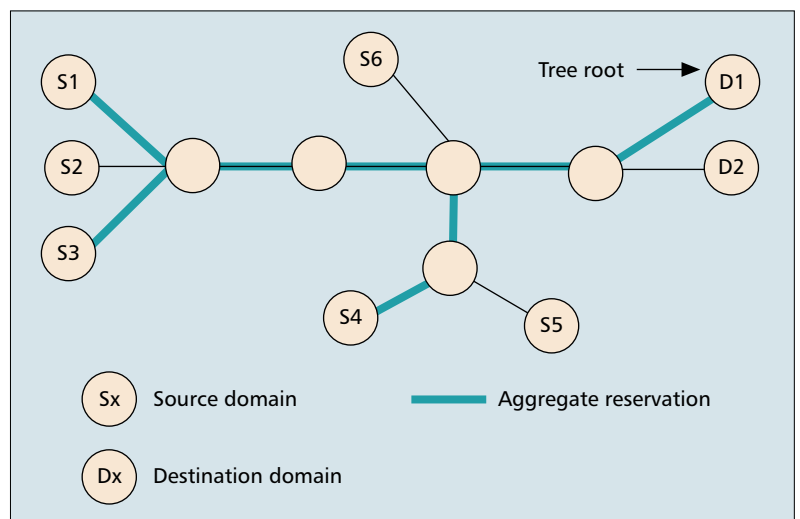
The Border Gateway Reservation Protocol (BGRP) [20] is a QoS signaling protocol designed for aggregate inter-domain usage between heterogeneous domains (autonomous systems). BGRP operates end-to-end only between domain border routers and is designed to aggregate reservations between domains improving scalability. BGRP uses the sink-tree aggregation approach and performs reservation aggregation by building a sink tree for each destination domain (Fig. 10). Reservations from different source domains that are

destined toward the same destination domain are aggregated along the path, forming a sink-tree rooted at the destination domain edge router. Basic functionality of BGRP includes a PROBE message sent by the source domain toward the receiver for determining reservation resource availability and recording the reservation path. The destination domain edge router terminates the PROBE and responds with a GRAFT message along the reverse path, which performs the actual inter-domain reservation and triggers the intra-domain QoS mechanisms in transit domains. The destination domain edge router is the de-aggregation point for the reservation

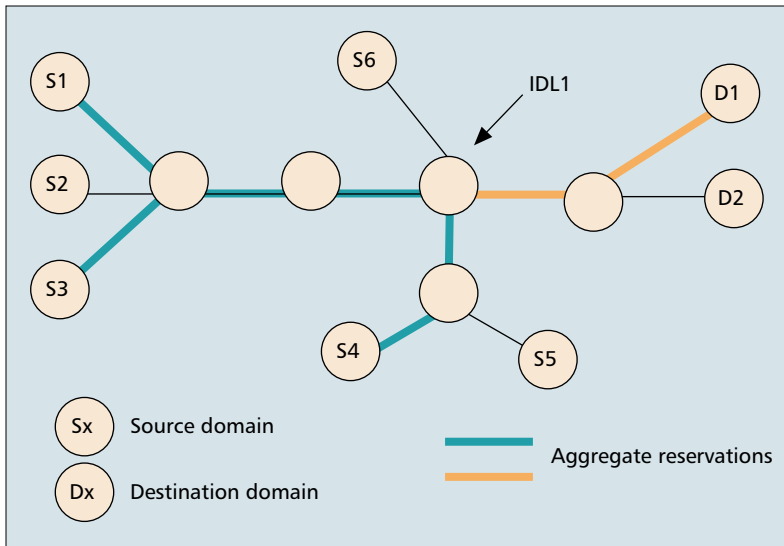
aggregate, while the source domain edge routers are the aggregation points. By performing sink tree-based aggregation of reservations toward each destination domain, BGRP results in storing per-destination domain QoS states in border routers. The latter can be considered to be a significant contribution to scalability when compared to the per-flow RSVP, if we consider that the number of domains (ASs) in the Internet is approximately 160,000 [21], whereas the number of possible host pairs, i.e., the maximum number of flows, is $(230,000,000)^2$ [22].

BGRP offers a significant gain in the number of maintained states in each router compared to RSVP. While the number of states in RSVP is equal to the number of QoS flows passing through a router, the respective number of states in BGRP corresponds to the number of destination domains that correspond to these flows. The gain in maintained states can reach a couple of orders of magnitude, depending on the network load and the diversification of the destination domains. Regarding the signaling load, BGRP reduces the message rate of the required refresh messages. The gain in the refresh message rate follows a similar pattern to the state gain, with an extra dependence on the refresh period (the smaller the refresh period, the greater the gain) [20].

The Shared-segment based Inter-domain Control Aggregation Protocol (SICAP) [23] is another approach for supporting aggregate inter-domain reservations between autonomous systems. SICAP combines the shared-segment and the tree-based aggregation approaches to create tree-based reservation



■ **Figure 10.** BGRP example: individual reservations from S1, S3, and S4 toward destination D1 are aggregated into an aggregate reservation tree with root D1.



■ **Figure 11.** SICAP example: individual reservations from S1, S3, and S4 toward destination D1 are aggregated into an aggregate reservation tree with root IDL1, while another reservation aggregate is created for the path segment between IDL1 and D1.

aggregates that do not necessarily extend to the destination domains. Apart from the destination domain de-aggregator point, intermediate de-aggregation locations (IDLs) are elected along the path so that reservation requests that share a common path segment but do not end-up at the same destination domain are aggregated up to one of their common routers along the path. Similar to BGRP, SICAP performs receiver-based reservations and uses a two-phase setup mechanism for their establishment. Upon the reception of the first QoS request, the respective source domain edge router initiates a REQ message toward the destination domain that collects information along the route regarding resource availability and network topology. The destination domain edge router uses the collected topology info and uses a specific choice algorithm to determine the appropriate IDL(s) along the route. By sending a RESV message along the reverse path, the destination domain edge router establishes consecutive aggregate segments between itself, the IDL(s), and the source domain edge router (Fig. 11). Future reservation requests that cross the same IDL(s) are dynamically incorporated into the existing aggregate(s), forming an aggregate reservation tree(s) rooted at the IDL(s). At each IDL, the aggregated reservations are de-aggregated and re-aggregated toward the next IDL or the destination domain. SICAP results in storing per-IDL states in each router along the path, further improving scalability as compared to BGRP (or RSVP). Simulation results [23] show that SICAP has consistently lower state requirements than BGRP. SICAP can result in a BGRP state reduction that varies between 0.5 and 0.8 (or more), depending on the specific load conditions in the network. SICAP signaling message load is similar to that of BGRP.

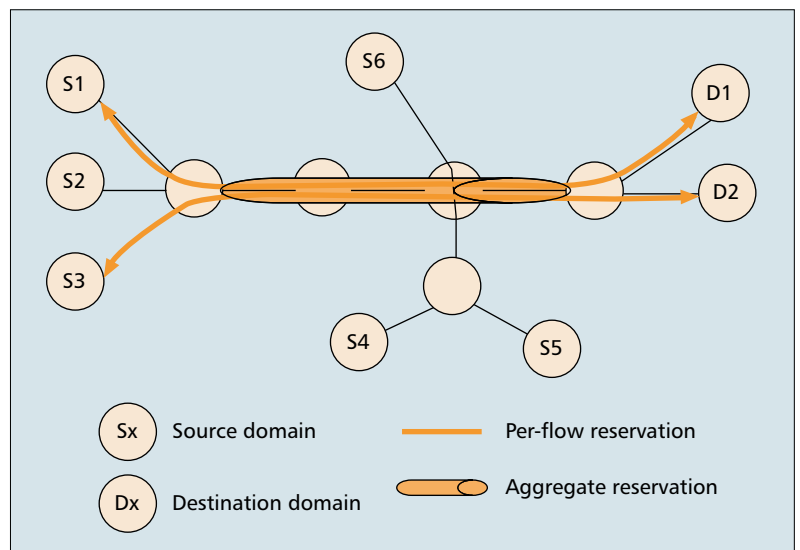
An architecture for dynamically aggregating reservations sharing a common path segment is proposed in [24]. The DARIS (Dynamic Aggregation of Reservations for Internet Services) architecture assumes the existence of a central resource management entity (similar to the BB approach) inside each DiffServ domain that has complete knowledge and control of the resources

inside the domain and also employs the inter-domain BGP routing table. DARIS enables the creation of an aggregate between two arbitrary domains as soon as a threshold k of active common reservations between the two domains is exceeded. In this case, all intermediate edge routers can substitute the respective per-flow states with a single aggregate state (Fig. 12). The reservation protocol that manages the aggregate reservations is the Domain Manager Signaling Protocol (DMSP). When a new resource request is received at a domain, a DMSP request message is sent toward the destination domain to check whether there exist other reservations along the path. If the threshold k of existing reservations is exceeded for a path segment, a new aggregate is established for these flows via the DMSP response message. The aggregation initiator becomes the aggregator domain, and the last domain of the common path segment is the de-aggregator domain. New reservations that traverse the same path segment can be incorporated into the existing reservation aggregate via appropriate signaling

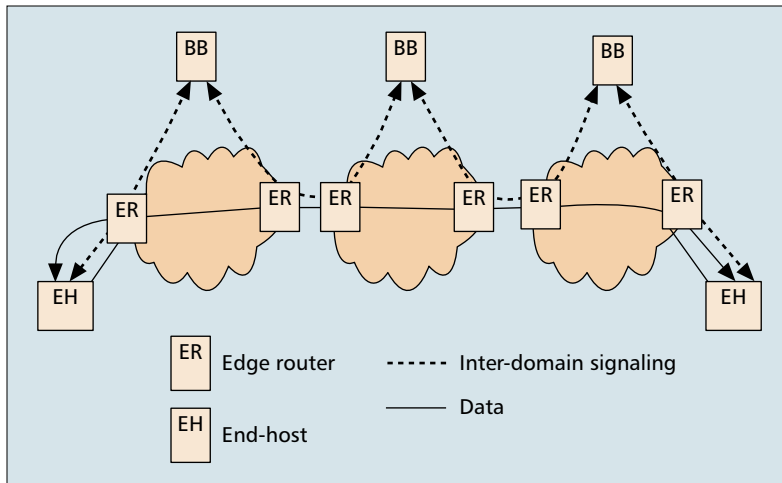
exchange. Simulation results [24] show that the DARIS state savings in the average lie in the range of one to two orders of magnitude when compared to a per-flow protocol that does not perform aggregation.

BGRP, SICAP, and DMSP are designed to aggregate individual reservations on inter-domain links into larger aggregates. These aggregation protocols result in storing fewer QoS states in domain border routers as compared to RSVP, thus enhancing scalability. Moreover, BGRP, SICAP, and DMSP are generally triggered by domain edge routers, unlike RSVP, which is triggered by end-hosts. However, DMSP assumes the underlying functionality of an end-to-end reservation protocol triggered by end-hosts. Instead, BGRP and SICAP assume that there is an established form of communication of QoS requests from end-hosts to domain edge routers, and vice versa.

BGRP aggregates reservations along the destination domain rooted sink-trees, which results in the storage of per-



■ **Figure 12.** DARIS example: individual reservations from S1 and S3 toward destinations D1 and D2 respectively are aggregated into an aggregate reservation along their common path segment.



■ **Figure 13.** *Inter-domain QoS signaling in BB architectures.*

destination domain states in domain border routers. SICAP further reduces the stored states by introducing intermediate de-aggregation points using topology-related information. DMSP takes a different approach, aggregating in a non-tree like fashion, introducing dynamic discovery of aggregators and de-aggregators that takes into consideration the current network load. An advantageous feature of DMSP resulting from its pure shared-segment design principle is that it allows for the creation of hierarchical aggregates (smaller aggregates can be aggregated into a larger one), which is not possible in tree-related protocols such as BGRP and SICAP. Moreover, a pure shared-segment approach that allows for dependence between the discovery of the aggregate end-points and the current traffic load may result in fewer aggregate stored states than tree-based approaches. The DARIS architecture, however, is centralized and relies on the BGP information for the de-aggregator discovery.

In BB-based architectures, inter-domain QoS signaling is triggered by the respective BBs when the received traffic exceeds the a priori bilateral resource agreements that exist between adjacent domains [4]. Resource agreements between domains are usually on a per-aggregate basis to enhance the resource allocation mechanism scalability. When a specific reservation aggregate needs to be modified, the requesting domain's BB sends an aggregate modification request to its neighboring BB along the path, which in turn may terminate or forward the request downstream (depending on its own inter-domain agreements). During propagation of the request between BBs, the intra-domain reservation mechanisms should be triggered, enabling the necessary adaptation of resources inside the transit domains (Fig. 13).

Following the basic concepts described above, an RSVP extension for establishing and modifying inter-domain resource aggregates is proposed in [4]. In addition, the QBone Signaling Workgroup proposed SIBBS (Simple Inter-domain Bandwidth Broker Signaling) [25], an inter-domain resource reservation protocol for deployment in BB-based architectures. SIBBS specifies a new set of reservation request and response messages, as well as a set of globally well-known services (GWS) that are globally identifiable throughout the network, each of which

requires different treatment regarding QoS. Irrespective of the specific signaling protocol used for inter-BB communication, special care should be given to the granularity level of the protocol, taking into account that there is a trade-off between accurate end-to-end QoS provision on one hand and scalability on the other [26].

SUMMARY AND DISCUSSION

QoS signaling is among the major elements constituting the QoS support mechanism. This article has focused on the identification of the most representative QoS signaling protocols for the Internet and identified their applicability to the basic Internet QoS architectures.

The simplified assumption for a homogeneous end-to-end Internet architecture (single-tier) gave birth to the IntServ QoS architecture. IntServ assumes identical QoS functionality in all intermediate routers along the data path, which is based on the support of per-flow differentiated treatment. Several single-tier end-to-end protocols have been proposed to implement the respective QoS signaling protocol, with RSVP being the major representative. While end-to-end per-flow protocols succeed in offering accurate QoS guarantees, high processing and storage overhead in routers is implied, since per-flow QoS states are stored in all traversing routers. Per-flow state maintenance in each intermediate router is a point of great dispute, mainly due to the scalability issues raised in backbone networks.

The necessity for a scalable and simplified alternative to IntServ led to the design of the DiffServ QoS architecture. DiffServ relies on packet differentiation based on a small number of per-hop behaviors statically configured inside routers. Static router configuration eliminates the need for QoS signaling, while packet classification in a small number of DiffServ QoS classes reduces significantly the number of states stored. However, DiffServ provides only an approximation of the requested QoS to the various flows, treating similar flows the same way.

The inherent heterogeneity of Internet domains gave birth to a two-tier QoS signaling architecture, where the intra-domain signaling performs reservations inside a domain and the inter-domain signaling performs reservations between domains. Separation of intra- and inter-domain signaling allows for greater flexibility and efficiency in resource allocation, where aggregation of resources can take place in core domains and border links, while per-flow QoS can be offered inside access domains. Further flexibility can be achieved by extending the architecture to a multi-tier signaling paradigm.

The major representatives of both intra- and inter-domain signaling protocols are summarized in Table 1, while Table 2 summarizes the basic characteristics of each of the protocols reviewed. One can observe that the soft-state approach is supported by all considered protocols, since soft-states offer flexibility in cases of node failure. Single-tier protocols usually assume per-flow processing in each router and are initiated by end-hosts. On the other hand, inter-domain protocols usually

Single-tier		RSVP YESSIR Boomerang INSIGNIA
Two-tier	<i>Intra-domain</i>	RSVP Aggr. RSVP RMD DiffRes
	<i>Inter-domain</i>	RSVP BGRP SICAP DMSP SIBBS

■ **Table 1.** *QoS signaling protocols classification.*

	Reservation granularity	Reservation initiation	Triggering point	Architecture	Soft or hard states
RSVP	Per-flow	Receiver-based	End-host	Distributed	Soft
YESSIR	Per-flow	Sender-based	End-host	Distributed	Soft
Boomerang	Per-flow	Sender-based or receiver-based	End-host	Distributed	Soft
INSIGNIA	Per-flow	Sender-based	End-host	Distributed/in-band	Soft
Aggr. RSVP	Per-aggregate	Receiver-based	Edge router	Distributed	Soft
RMD	Per-aggregate	Sender-based	Edge router	Distributed	Soft or hard
DiffRes	Per-aggregate	Sender-based	Edge router	Distributed	Soft
BGRP	Per-aggregate	Receiver-based	Edge router	Distributed	Soft
SICAP	Per-aggregate	Receiver-based	Edge router	Distributed	Soft
DMSP	Per-aggregate	Sender-based	Domain agent	Centralized	Soft
SIBBS	Per-aggregate	Receiver-based	Bandwidth broker	Centralized	Soft

■ **Table 2.** Basic characteristics of QoS signaling protocols.

operate on a per-aggregate basis and are initiated by edge routers or domain agents. Intra-domain protocols can either perform per-flow or per-aggregate reservations, depending on the positioning of the domain inside the network (core or access). In general, distributed architectures seem to be preferable to agent-based architectures, basically due to scalability problems arising from centralized functionality. The sender- or receiver-based initiation of the reservation has been for a long time a point of dispute between protocol designers. The supporters of sender-initiated reservations call on the avoidance of storage of backward routing information in the nodes situated in the forwarding path. This information is necessary in receiver-initiated schemes so that reservation messages sent by the receiver to the sender follow the same backward path as the forward one. On the other hand, sender-initiated reservation schemes may result in unnecessary reservation of resources inside routers, since reservation in a router takes place without knowing if the other routers along the path can allocate the requested resources or not.

Despite the great research advances in the Internet QoS signaling area, there is no large-scale deployment of the QoS supporting protocols. This can be attributed to various reasons, not necessarily related to technical issues. Several core providers seem to find over-provisioning a simple alternative to QoS support, attributed to the great advances in fiber optics communications. However, considering that bandwidth will always be a finite resource, the support of service differentiation will eventually be necessary.

Moreover, the technical maturity of the Internet QoS has not yet been followed by the respective development of the necessary business models, and charging and policy frameworks. The business models of the Internet players are still not adapted to dynamic negotiations of QoS agreements among users and access and core network providers. Innovative QoS-aware service models assume the deployment of dynamic charging techniques as well as efficient AAA methods.

RELATED WORK

As depicted in the previous sections, several QoS standards and proposals have been produced during the last decade. The need to cast a global view on the Internet QoS area has driven various research efforts, each of which is focusing on different QoS aspects.

More specifically, a first effort to give a global overview of Internet QoS is presented in [18]. The authors in [18] focus on the various mechanisms that can be used for offering QoS in the Internet and describe the IntServ/DiffServ architectures, MPLS, traffic engineering, and constrained-based routing. The advantages and drawbacks of each method are discussed, as well as the motives behind the introduction of each method and the relations between them. In addition, a comparison between ATM networks and IP router networks using DiffServ and MPLS is attempted. The RSVP protocol is generally used throughout the article to cover the signaling needs of the described QoS architectures.

The two-tier QoS signaling approach for the Internet was first presented in [4], which described a realization of this model based on bandwidth brokers. The focus of this work is on the presentation of intra-domain and inter-domain resource allocation procedures. Moreover, a measurement-based approach for the reservation of aggregate resources between domains is proposed. Apart from RSVP, the authors in [4] do not consider the applicability of other QoS protocols for intra- or inter-domain usage.

The IETF NSIS group (Next Steps in Signaling) [27] is currently very actively working on the design of a generic signaling protocol managing general-purpose states. Since QoS provision is basically an issue of establishing and maintaining reservation states inside network routers, QoS signaling is considered as a special case of the NSIS generic signaling protocol.

NSIS is trying to identify which aspects of the existing QoS protocols could constitute useful elements for the future Internet signaling protocols have. They have released two Internet Drafts that review the current QoS architectures and protocols for the Internet. An analysis of the existing QoS solutions, trying to point out open issues in QoS signaling, is provided in [28]. The QoS solutions considered are the end-to-end usage of RSVP, the IntServ over DiffServ architecture, the static assignment of trunk reservations based on DiffServ, and the Aggregated RSVP. Advantages and disadvantages of each QoS solution are further discussed.

In an effort to gain useful lessons for the design of the future generic Internet signaling protocol, the authors in [29] present a review of some of the major QoS signaling protocols and analyze their features. Special focus is given to the RSVP protocol, which is directly examined for compatibility with the NSIS requirements. The RSVP protocol's characteristics and

extensions are studied in detail, including RSVP reliability, performance, security and mobility issues, RSVP usage with MPLS, ATM, GMPLS, as well as mechanisms for refresh reduction and reservation across tunnels. A short description and evaluation is given for other QoS signaling proposals. Useful conclusions are drawn regarding the design of the NSIS generic signaling protocol.

NSIS is expected to give valuable input in the Internet signaling area. The group has been very active during the past few years and has already specified the requirements for the generic signaling protocol [30] and the requirements for a QoS solution for mobile terminals (using Mobile IP) [31]. Moreover, the NSIS group has designed a general framework for the NSIS Internet signaling protocol [32], where the overall signaling protocol suite is split into a generic lower NSIS transport layer and several upper signaling layers, one for each specific signaling application (e.g., QoS signaling layer). The interactions between the two layers are also described in [32]. A protocol for the NSIS transport layer that uses the existing IP transport layer protocols has been further proposed [33]. The NSIS signaling layer protocol for the QoS has been specifically addressed in [34].

The NSIS framework supports the multi-tier paradigm addressed in our survey by specifically examining scenarios that include signaling end-to-end, end-to-edge (intra-domain), and edge-to-edge (inter-domain) [32]. To support this infrastructure within a single signaling framework, NSIS supports local objects that are used only within the boundaries of a domain.

CONCLUSIONS

The increasing demand for real-time services support over the Internet necessitates the evolution of the Internet from a best-effort network to a network that can reliably offer QoS guarantees. Various efforts have been made toward this goal, comprising the design of QoS architectures and signaling protocols for the Internet.

In this article the major QoS signaling protocols have been identified and classified under the general prism of the two-tier resource management model for the Internet. Moreover, the individual characteristics of each protocol are presented and relatively compared. The general trends and principles behind Internet QoS signaling are also discussed.

Major research is still being carried out regarding QoS provision in the Internet, since none of the proposed protocols seems to offer the ideal compromise between protocol complexity, network scalability, and accuracy of offered quality of service. Among the interesting conclusions derived through previous research is that efficient QoS provisioning in the Internet would not be offered by the universal adoption of a single end-to-end QoS architecture. Instead, a two-tier signaling architecture, allowing for aggregation of resources in backbone networks and per-flow QoS offering inside access domains, serves relatively well the purpose of network scalability and offered QoS accuracy. This can be achieved by the use of separate QoS signaling inside and in-between network domains that would further allow for aggregation of allocated resources inside and in-between core domains. Aggregation in the QoS data and control plane allows for scalability in core routers and domains, while assuring end-to-end QoS.

The generic approach of a multi-tier signaling architecture for the Internet is also supported by the NSIS (Next Steps in Signaling) IETF Working Group, which is currently defining a unified signaling framework for the Internet.

- [1] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: An Overview," IETF RFC 1633, June 1994.
- [2] R. Braden et al., "Resource Reservation Protocol (RSVP)," IETF RFC 2205, Sept. 1997.
- [3] S. Blake et al., "An Architecture for Differentiated Services," IETF RFC 2475, Dec. 1998.
- [4] L. Terzis et al., "A Two-Tier Resource Management Model for the Internet," *Global Internet*, Dec. 1999.
- [5] L. Berger et al., "RSVP Refresh Overhead Reduction Extensions," IETF RFC 2961, Apr. 2001.
- [6] P. Pan and H. Schulzrinne, "YESSIR: A Simple Reservation Mechanism for the Internet," *Int'l. Wksp. Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Cambridge, England, July 1998.
- [7] G. Feher et al., "Boomerang — A Simple Protocol for Resource Reservation in IP Networks," *IEEE Wksp. QoS Support for Real-Time Internet Applications*, Vancouver, Canada, June 1999.
- [8] S. Lee et al., "INSIGNIA: An IP-Based Quality of Service Framework for Mobile Ad Hoc Networks," *J. Parallel and Distributed Computing (Academic Press)*, Special issue on Wireless and Mobile Computing and Commun., Apr. 2000.
- [9] J. Rosenberg et al., "SIP: Session Initiation Protocol," IETF RFC 3261, June 2002.
- [10] Handley and Jacobson, "SDP: Session Description Protocol," IETF RFC 2327, Apr. 1998.
- [11] G. Camarillo, W. Marshall, and J. Rosenberg, "Integration of Resource Management and Session Initiation Protocol (SIP)," IETF RFC 3312, Oct. 2002.
- [12] S. Salsano and L. Veltri, "QoS Control by means of COPS to support SIP-based applications," *IEEE Network*, Mar./Apr. 2002.
- [13] Y. Bernet et al., "A Framework for Integrated Services Operation over DiffServ Networks," IETF RFC 2998, Nov. 2000.
- [14] F. Baker et al., "Aggregation of RSVP for Ipv4 and Ipv6 reservations," IETF RFC 3175, Sept. 2001.
- [15] L. Westberg et al., "Resource Management in DiffServ (RMD): A Functionality and Performance Behavior Overview," *Protocols for High-Speed Networks*, 2002.
- [16] D. Sisalem, S. Krishnamurthy, and S. Dao, "DiffRes: A Light Weight Reservation Protocol for the Differentiated Services Environment," Tech. Rep., HRL Laboratories, Malibu, USA, Dec. 1999.
- [17] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," IETF RFC2638, July 1999.
- [18] Xiao and L. M. Ni, "Internet QoS: A Big Picture," *IEEE Network*, vol. 13, no. 2, Mar./Apr. 1999.
- [19] Y. Bernet, "Format of the RSVP DCLASS Object," IETF RFC 2996, Nov. 2000.
- [20] P. Pan, E. Hahne, and H. Schulzrinne, "BGRP: A Tree-Based Aggregation Protocol for Inter-Domain Reservations," *J. Commun. and Networks*, vol. 2, no. 2, June 2000.
- [21] CAIDA project, <http://www.caida.org/analysis/routing/atypes/>
- [22] ISC Internet Domain Survey, <http://www.isc.org/index.pl?/ops/ds/reports/2004-01/>
- [23] R. Sofia, R. Guerin, and P. Veiga, "SICAP: A Shared-segment Inter-Domain Control Aggregation Protocol," Technical Report, ESE, University of Pennsylvania, Oct. 2002.
- [24] R. Bless, "Dynamic Aggregation of Reservations for Internet Services," *ICTSM 10*, Oct. 2002.
- [25] QBone Signaling Design Team, Final Report, <http://qbone.internet2.edu/bb/>
- [26] M. Gunter and T. Braun, "Evaluation of Bandwidth Broker Signaling," *Proc. ICNP '99*, Nov. 1999.
- [27] IETF Next Steps In Signaling Working Group (NSIS), <http://www.ietf.org/html.charters/nsis-charter.html>
- [28] H. de Meer et al., "Analysis of Existing QoS Solutions," June 2002.
- [29] J. Manner, X. Fu, and P. Pan, "Analysis of Existing Quality of Service Signaling Protocols," Internet draft, May 2004.
- [30] M. Brunner, "Requirements for Signaling Protocols," IETF RFC 3726, Apr. 2004.

-
- [31] H. Chaskar, "Requirements of a Quality of Service (QoS) Solution for Mobile IP," IETF RFC 3583, Sept. 2003.
- [32] R. Hancock, "Next Step in Signaling: Framework," Internet draft, Oct. 2003.
- [33] H. Schulzrinne and R. Hancock, "GIMPS: General Internet Messaging Protocol for Signaling," Internet draft, May 2004.
- [34] S. Van den Bosch, G. Karagiannis, and A. McDonald, "NSLP for Quality-of-Service Signaling," Internet draft, May 2004.

BIOGRAPHIES

ALEXANDROS KALOXYLOS (kaloxyl@uop.gr) received his B.Sc. from the department of computer science at the University of Crete in 1993, his M.Phil. from the department of computing and electrical engineering at the Heriot-Watt University in 1994, and his Ph.D. degree from the department of informatics at the University of Athens in 1999. Since 1995 he has been a senior researcher at the Communications Network Laboratory of the University of Athens. Currently he is an assistant professor at the University of Peloponnesus. He has participated in numerous projects in the context of EU programs, as well as in national research initiatives. His research interests include mobile and wireless networks, QoS provisioning, and formal description techniques.

LAZAROS MERAKOS (merakos@di.uoa.gr) received his diploma in electrical and mechanical engineering from the National Technical University of Athens in 1978, and M.S. and Ph.D. degrees in electrical engineering from the State University of New York, Buffalo, in 1981 and 1984, respectively. He has been a profes-

sor in the department of informatics and telecommunications, University of Athens, since 1994. His research interests are in the design and performance analysis of broadband networks, and wireless/mobile communication systems and services. He has led the research activities of the Communications Network Laboratory of the University of Athens, and he is the president of the board of the Greek Universities Network and the Greek Schools Network.

SARANTIS PASKALIS (paskalis@di.uoa.gr) received his B.Sc. degree (with honors) in computer science from the department of informatics at the University of Athens in 1995 and his M.Sc. (with honors) in computer science (communication systems and networks) from the same department in 1998. He received his Ph.D. from the same department in 2004. He is a staff member of the Communication Networks Laboratory (CNL) of the University of Athens. His research interests include performance analysis, mobility management, and QoS protocol design and evaluation.

DIMITRA VALI (dvali@otereseach.gr) received her M.Sc. degree in communication systems and networks in 1998 and her B.Sc. degree in computer science in 1995 from the department of informatics at the University of Athens. In 1995 she became a staff member of the Communication Networks Laboratory of the University of Athens. Since 1999 she has been with the Hellenic Telecommunications Organization of Greece (OTE), working for the Mobile/Wireless Communications Laboratories of OTE Research. Her interests particularly include QoS protocol design, evaluation, and performance analysis.