# The EuQoS System: A Solution for QoS Routing in Heterogeneous Networks

*Xavi Masip-Bruin, Marcelo Yannuzzi, René Serral-Gracià, and Jordi Domingo-Pascual, Technical University of Catalunya (UPC), Spain*

*José Enríquez-Gabeiras and María Ángeles Callejo, Telefónica I+D (TID), Madrid, Spain*

*Michel Diaz and Florin Racaru, Analysis and Architecture of Systems Laboratory (LAAS), France*

*Giovanni Stea and Enzo Mingozzi, University of Pisa, Italy*

*Andrzej Beben and Wojciech Burakowski, Warsaw University of Technology (WUT), Poland*

*Edmundo Monteiro and Luís Cordeiro, University of Coimbra (UoC), Portugal*

## ABSTRACT

EuQoS is the acronym for "end-to-end quality of service support over heterogeneous networks," which is a European research project aimed at building an entire QoS framework, addressing all the relevant network layers, protocols, and technologies. This framework, which includes the most common access networks (xDSL, UMTS, WiFi, and LAN) is being prototyped and tested in a multidomain scenario throughout Europe, composing what we call the EuQoS system. In this article we present the novel QoS routing mechanisms that are being developed and evaluated in the framework of this project. The preliminary performance results validate the design choices of the EuQoS system, and confirm the potential impact this project is likely to have in the near future.

## INTRODUCTION

New demands for using multimedia applications over the Internet, such as IP telephony, video, telemedicine, tele-engineering, and tele-education, have spurred the emergence of several research topics aimed at providing customers with the required quality of service (QoS) at different network layers. One of these research topics deals with the problem of finding a feasible path between a source and destination node satisfying one or more QoS constraints. This is precisely the main function of QoS routing (QoSR).

Despite the fact that routing decisively contributes to the provision of QoS, two main factors prevent QoSR from being widely deployed. First, the problem of QoSR with multiple constraints is NP-hard. This means that while numerous heuristics have been proposed, only a few exact solutions exist [1]. Second, delivering end-to-end QoS to users connected to the Internet through different access networks requires several other building blocks to be properly engineered and interconnected, which is still a big challenge for the research and industry communities [2]. Several hot topics, such as admission control, signaling, traffic engineering (TE), and network management, need further research efforts to find solutions appealing enough to challenge the usual overprovisioning strategies.

The EuQoS project [3] gathers research centers, universities, telcos, and consultants working in all the abovementioned hot topics. The main goal of this project is to design and implement an architectural network model (the EuQoS system) capable of guaranteeing end-to-end QoS across heterogeneous networks. EuQoS subscribers will be able to use both EuQoS-enabled applications to communicate with guaranteed QoS as well as legacy applications. This requires coordinated QoS mechanisms to be placed in both the applications and the network. In particular, the EuQoS system includes:

- Authentication, authorization, and accounting (AAA)
- Charging
- Connection admission control (CAC)
- Signaling and service negotiation
- Monitoring and measurements
- QoS routing (QoSR)
- Network management
- TE and resource optimization

At this stage, the EuQoS team has already designed and developed a first prototype of the EuQoS system. The prototype is starting to be deployed in the setting shown in Fig. 1. This multidomain testbed is built upon a core network composed by GEANT (the European research network) and the National Research and Education Networks (NRENs) of the partners involved.

In this article we report an overview of the experiences gained in developing and testing such a complete and heterogeneous QoS network architecture, with specific focus on QoSR issues.

The rest of the article is organized as follows. In the next section the architecture and QoS model are overviewed. Then, the EuQoS approach to QoSR is addressed in detail. After that, a preliminary evaluation of the proposed solutions is shown. Finally, we conclude highlighting the open issues and the directions for future work.

## THE EuQoS ARCHITECTURE AND QoS MODEL

The EuQoS architecture has been defined according to the following rules:
- The customers' applications should be able to negotiate the content and quality of each communication.
- Network administrators should have the freedom to use any of the existing network technologies, and the EuQoS system should be deployable on top of them.
- The proposed mechanisms should be incremental, in the sense that they should coexist with the existing installed base.

Figure 2 illustrates the main building blocks of the EuQoS architecture. In the rest of this section we outline the key components in Fig. 2 and the QoS model.

### MAIN BUILDING BLOCKS

In order to supply the desired freedom to network administrators, a *virtual network layer* has been defined, which decouples network decisions from network technologies. To achieve this goal, the virtual network layer is split in two sublayers, a technology-independent (TI) one and a technology-dependent (TD) one. As shown in Fig. 2a, the TI sublayer consists of a logical entity, called the resource manager (RM), which is in charge of managing QoS for each domain. The RM coordinates domain-wide CAC decisions, stores and manages peering agreements with neighboring domains, and controls the interdomain routing process. Whenever necessary, the RM decisions are enforced in specific device configurations by means of resource allocators (RAs), which are located at the TD sublayer (see the bottom of Fig. 2a and b).

In EuQoS, QoS resource management is handled on a per-session basis. At first, the remote communicating applications agree on the content and quality of their communication. This negotiation at the application layer requires a signaling protocol between end users. To this end, we developed *EQ-SIP*, which is an extension of SIP [4] that includes new mechanisms for negotiating particular QoS parameters. The EuQoS support to the applications includes (see Fig. 2b):
- A QoS control module (QCM) that links the QoS requests of the users to the network connection
- Application signaling (ASIG) that implements the EQ-SIP protocol in the users' terminals
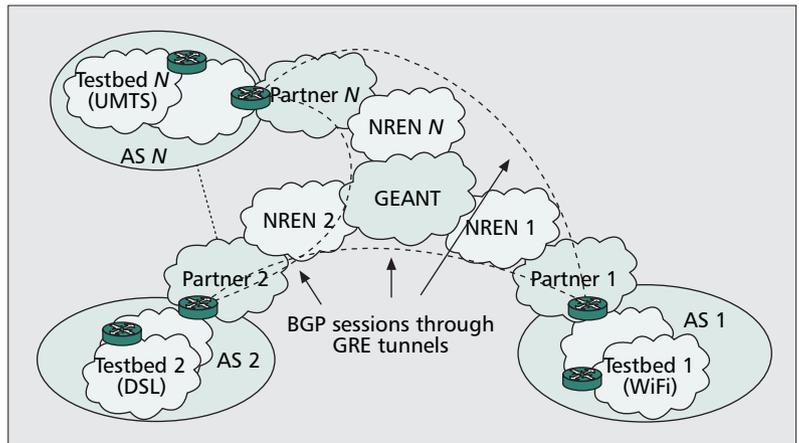


■ **Figure 1.** *Testbed architecture: access networks, NRENs, and GEANT.*

- An enhanced transport protocol that provides the QoS adaptations needed to handle the different QoS classes in the network layer

Each application request reaches the virtual network layer through the application programming interface (API) of the access network where the caller is located. Upon receiving the requests, the TI sublayer checks the feasibility of an end-to-end path (i.e., the capability of all the networks involved to provide the requested QoS). As a result, an end-to-end path fulfilling the QoS demands needs to be computed. The EuQoS system supports two different models to compute this path: a "loose" model where the data path is determined by a QoSR protocol on a per-domain basis, and a "hard" model where the data path, or part of it, is established using a TE mechanism (e.g., multiprotocol label switching with TE, MPLS-TE). Whereas the hard model is still in the design stage, the loose model is currently implemented in the prototype, and hence is the focus of the description in this article. Inside each domain the virtual network layer basically includes the following functionalities (Fig. 2b).

**Signaling and service negotiation** (SSN) encompasses support for application signaling (EQ-SIP), horizontal signaling between the RMs, and vertical signaling between the RMs and RAs. In order to check the availability of resources during the call setup phase, the horizontal signaling messages must reach all the RMs along the path. However, the customers' traffic is never routed through the RMs, so the signaling messages have to be forwarded out of the normal data paths, as depicted in Fig. 3. To achieve this goal, an extension of the Next Steps in Signaling (NSIS) protocol [5] has been designed. We call this extension *EQ-NSIS*. This pioneering implementation of NSIS is used for signaling and exchanging QoS requirements between RMs across different domains (Figs. 2a and 2b). This extension is necessary because an approach that needs to redirect the end-to-end signaling messages from some routers toward the RMs, such as the one shown in Fig. 3, is not fully solved by the NSIS protocol. To tackle this problem, the EuQoS team designed a middle layer between the NSIS Transport Layer Protocol (NTLP) [6] and the NSIS Signaling Layer

Protocol (NSLP) [7]. This middle layer is called the Hybrid Path (HyPath) [8], and is transparent to the NSIS layers, since the interface between the NTLP and the NSLP remains unchanged.

The operation of EQ-NSIS with the additional HyPath layer in the border routers and RMs in the different domains is illustrated in Fig. 3. When a user makes a QoS request to the EuQoS system, EQ-NSIS signaling starts and must reach all the RMs along the chosen path. In the first domain the HyPath in the local RM uses the RM's routing module to discover the local border (egress) router for the data path. After that, the HyPath asks the NTLP to send an NSIS message to the corresponding border router. This message contains the NSLP payload and some additional HyPath information. Once in the border router, the EQ-NSIS signaling message is sent toward the end user's domain. In this scenario all border routers are HyPath aware. In each downstream domain the EQ-NSIS signaling message is intercepted by the ingress border router and redirected to the local RM (Fig. 3).

After processing the message, each RM resumes signaling with a message back to the ingress border router. Signaling is restarted in the ingress border router, and the NSIS message continues toward the next domain. This process continues along all downstream domains until the last domain is reached. With this architecture all the requirements to achieve end-to-end network signaling are met, and no changes are needed in the definitions of the NTLP and NSLP layers.

For vertical signaling between RMs and RAs, the EuQoS team developed *EQ-COPS*, which is an extension of the Common Open Policy Service (COPS) [9]. EQ-COPS provides a scheme to map high-level QoS domain policies into specific low-level network device configurations, coping with both the required autonomy of QoS management inside each domain and the need to establish a TI sublayer composed by the RMs.

**CAC in each domain**. The CAC module in the RM checks for availability of resources both inside the domain (intradomain CAC) and in the links between peering domains (interdomain link CAC). CAC at the RA level is also enforced.

**The monitoring and measurement system (MMS)** provides a dedicated system in order to evaluate the real values of the QoS metrics provided by the network.

**Traffic engineering and resource optimization (TERO)** is in charge of interdomain QoSR configuration and resource provisioning.

**The security (SAAA) and charging (CHAR)** modules are also included.

As the EuQoS system is targeted for guaranteed QoS, BGP-4 cannot be used as the interdomain routing protocol. Thus, an Enhanced QoS Border Gateway Protocol (EQ-BGP) has been developed, building on a former extension of BGP-4 called qBGP [10]. EQ-BGP is the protocol in charge of determining the QoSR paths between end users, and is described in detail in the next section. In summary, the end-to-end QoS paths are built using the following key components:
• The RMs
• The RAs
• EQ-BGP
• EQ-NSIS
• EQ-COPS

### THE EuQoS QoS MODEL

QoS in EuQoS is achieved by implementing a set of five end-to-end classes of service (CoSs), specifically:
• IP telephony
• Real-time interactive
• Multimedia streaming
• High-throughput data
• Best effort

These CoSs are known and visible to the applications of end users. The traffic generated by a given application is submitted to the appropriate end-to-end CoS once the connection setup process has been successfully completed.

All the functions in the RM, RA, and EQ-BGP routers are CoS-specific. For instance, different routing tables, routing decision processes, provisioning strategies, traffic control mechanisms, and CAC policies exist for the different CoSs. Each domain is free to provide its own
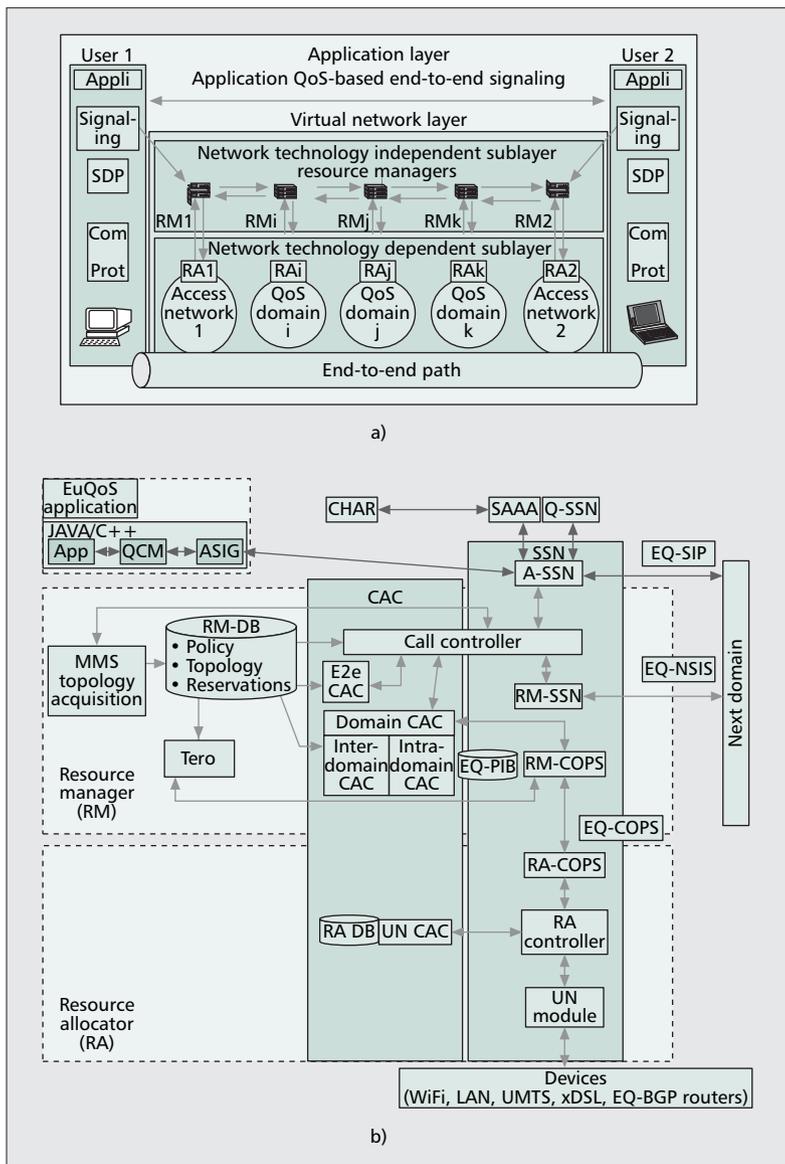


■ **Figure 2**. *The EuQoS model: a) the high-level EuQoS architecture; b) the main building blocks within a domain.*

implementation of a CoS, as far as it is compliant with its specifications, which in the framework of the EuQoS project are based on International Telecommunicatioin Union (ITU) recommendations. Neighboring domains establish per-CoS peering agreements, called peering service level specifications (p-SLSs), in order to regulate the transit of traffic for the different CoSs through their interdomain links.

## QoSR IN THE EUQOS SYSTEM

The EuQoS system targets to provide end-to-end QoS across heterogeneous networks. This motivates encompassing QoSR issues at the access networks, and at the intradomain and interdomain levels. So far, the major research efforts in the project have been devoted to interdomain QoSR, which is considered the most important issue by the telecom operators participating in the project. Furthermore, in practical settings the network terminals (e.g., Universal Mobile Telecommunications System [UMTS] mobile phones, WiFi notebooks, digital subscriber line [DSL] modems/routers) are typically connected through stub networks. Routing in these terminals is usually handled by means of default routing, so in practice the QoSR decisions need to be made *between* the source and destination access networks, but not *within* these latter. Even though QoSR decisions are not necessarily needed inside the access networks, QoS still needs to be delivered in these networks. In the EuQoS system this is managed by means of CoS subscription, QoS policies (e.g., CAC on a per-CoS basis and traffic shaping), resource reservations during an EuQoS session, and QoS monitoring.

Several new mechanisms related to interdomain QoSR were developed within the EuQoS project. Hereafter, we describe the most important ones: the p-SLSs, the EQ-BGP protocol, and the main tasks of the TERO module.

***p-SLSs*** — In the EuQoS system, an autonomous system (AS) negotiates p-SLSs with its neighbors. The p-SLSs regulate the transit and QoS guarantees of traffic belonging to a given CoS at an interdomain link in one direction. Thus, two ASs negotiating a p-SLS are called the *customer* and *provider* of that p-SLS, meaning that the traffic flows from the former to the latter. More specifically, an AS sends routing advertisements for a given CoS only along interdomain links at which it has previously negotiated a p-SLS as a provider. The p-SLSs formally specify:
• The amount of traffic a customer can inject at the interdomain link, and the actions the provider will take against nonconforming traffic
• The QoS the provider guarantees to the admitted traffic
Packets of a given CoS can leave an AS through an interdomain link only if a p-SLS exists for that CoS at that interdomain link. Thus, interdomain QoSR is constrained by the p-SLSs, which are controlled by the TERO module.

***EQ-BGP*** — EQ-BGP is the interdomain QoSR protocol proposed and developed within the EuQoS project. Its purpose is to select and adver-
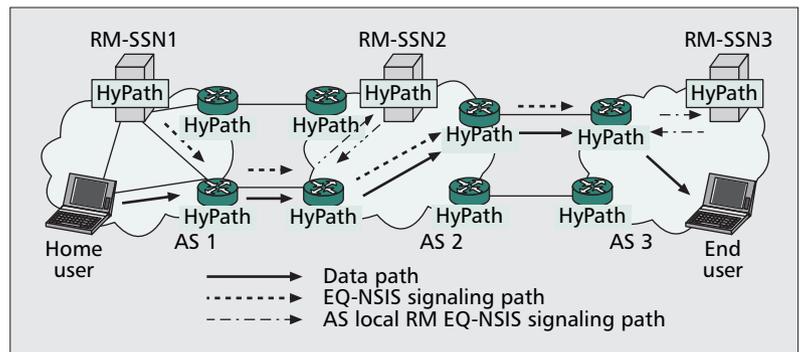


■ **Figure 3.** *EQ-NSIS signaling with HyPath.*

tise the QoS paths for the different CoSs [11]. EQ-BGP extends the BGP-4 routing protocol in several ways. First, it includes an optional path attribute, QoS_NLRI, that conveys information about the QoS capabilities of a path. Second, it includes a QoS assembling function for computing aggregated values of the QoS parameters for *entire* routing paths. For example, this assembling function can supply the sum of the delays for each segment of a path. Third, EQ-BGP has a QoS-aware decision process for selecting the best end-to-end path for the different CoSs. Fourth, EQ-BGP handles multiple routing tables in order to store the available paths for each end-to-end CoS.

EQ-BGP performs QoSR in multidomain networks by taking into account both intra- and interdomain QoS information about the available end-to-end CoSs. Thus, EQ-BGP sets the road map for the available QoS paths between each pair of source and destination networks. These paths are called *end-to-end QoS paths*, and are computed and advertised by EQ-BGP routers for each CoS separately.

In Fig. 4 we show an example of how the QoSR information is computed and advertised across different domains using EQ-BGP. For the sake of simplicity, we assume a simple network consisting of three domains, A, B, and C, that support the same end-to-end CoSs. We assume that each EQ-BGP router is aware of the *nominal* values of the QoS parameters that are assured both inside its particular domain ($Q_A$, $Q_B$, or $Q_C$, depending on the domain) as well as on its corresponding interdomain links ($Q_{A->B}$, $Q_{B->A}$, $Q_{B->C}$, or $Q_{C->B}$, also depending on the domain). All these nominal QoS values are computed by the TERO module during the network provisioning process and correspond to the *maximum admissible load* controlled by the intra- and interdomain CAC functions. These nominal QoS parameters typically change at provisioning timescales (e.g., on the order of days or weeks). The reason for this is to avoid routing advertisements conveying only slight variations of the QoS values. This provides a scalable EQ-BGP routing protocol, but clearly the success of the approach requires adaptive provisioning and strict admission control policies. During our simulations, the approach of advertising aggregated nominal QoS parameters in conjunction with adaptive provisioning has proven to achieve excellent results.

Now, let us consider the case when domain C advertises a new prefix, say NLRI_C. The routing information is propagated toward domain A
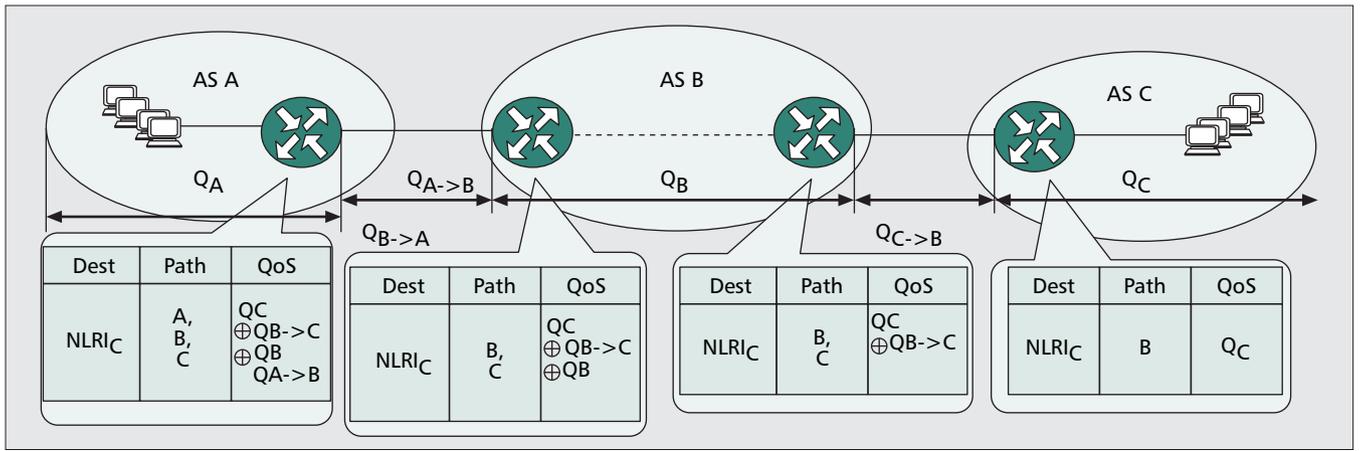
**■ Figure 4.** *Example of EQ-BGP operation.*

through domain B. Figure 4 shows how the QoSR tables of the border EQ-BGP routers become populated. During this process EQ-BGP routers aggregate the nominal values of the QoS parameters along the path, taking into account the nominal QoS contributions of the intradomain segments as well as those of the interdomain segments of the path. For example, domain A learns an end-to-end QoS path toward the destination $NLRI_C$ with quality $[Q_C \oplus Q_{B->C} \oplus Q_B \oplus Q_{A->B}]$ for a particular CoS, wherein the operator $\oplus$ denotes the appropriate QoS assembling function.

*TERO* — TERO is in charge of interdomain routing configuration and resource provisioning. More specifically, it controls the interdomain routing process based on QoS requirements so as to steer the traffic through the ASs in the most effective way, optimizing interdomain resources such as bandwidth and buffer space on interdomain links. Furthermore, it configures queues and policers at interdomain links so as to provision the necessary resources to support QoS traffic across neighboring domains. TERO actions regarding routing and provisioning can be taken either as a reaction to the variation of the network topology, or periodically for maintenance and optimization reasons. Thus, TERO works on a *network provisioning* timescale, (e.g., days or weeks), so it operates on a much longer timescale than an EuQoS session lifetime.

TERO interacts with the border routers through EQ-COPS (Fig. 2b) to configure the EQ-BGP protocol. More specifically, TERO configures EQ-BGP routers so that:

• EQ-BGP update messages are allowed to flow through an interdomain link whenever a new p-SLS is negotiated.
• The QoS-NLRI information is properly updated before EQ-BGP messages are advertised to neighboring ASs as well as inside the domain. In fact, the QoS-NLRI information advertised to upstream domains assembles ($\oplus$) the *nominal* QoS-NLRI information included in the update messages received from downstream domains, with the nominal values of the QoS parameters that are assured in both intra- and interdomain links of the domain (Fig. 4).
• When an EQ-BGP router receives multiple

candidate paths for the same destination, it runs the EQ-BGP decision process. This process includes an additional step to BGP-4 decision process, which takes into account a new parameter called the *degree of preference* (DoP). The latter is computed by EQ-BGP routers based on the *QoS preference parameters* provided by TERO. All these parameters are described in the following subsection.

### DEGREE OF PREFERENCE (DoP)

When a border router receives an EQ-BGP update from another AS, it associates a DoP with that update. The latter is exploited in the EQ-BGP decision process to select the best route among different updates advertising the same destination for a given CoS. The DoP parameter has a *local*, AS-wide meaning and is never advertised to other ASs. Since in EuQoS there is one decision process for each CoS, the computation of the DoP can — in principle — be different from one CoS to another. In the first prototype, the following formula has been used for all CoSs, except the best-effort class:

$$DoP = \sum_{i \in \{IPTD, IPDV, IPLR\}} \frac{f_i}{\max\{0, \left[M_i - (Q_i \oplus Q_i')\right]\}}$$

(1)

The DoP computed by an EQ-BGP border router is the sum of three terms, each associated with a different QoS parameter carried in the EQ-BGP updates: IP packet transmission delay (IPTD), IP packet delay variation (IPDV), and IP packet loss ratio (IPLR). Each term consists of:

• A *QoS preference* $f_i$, which accounts for the relative importance of the QoS parameter $i$ with respect to the others.
• *A parameter value* $Q_i$, the assembled value of the QoS parameter $i$ carried in the incoming update.
• *A parameter value* $Q_i'$, the actual (real) value of the QoS parameter $i$ in the interdomain link from which the EQ-BGP router receives the incoming update. This parameter basically takes into account the current load on the interdomain links and is used to *locally* compute Eq. 1, but is never included in the
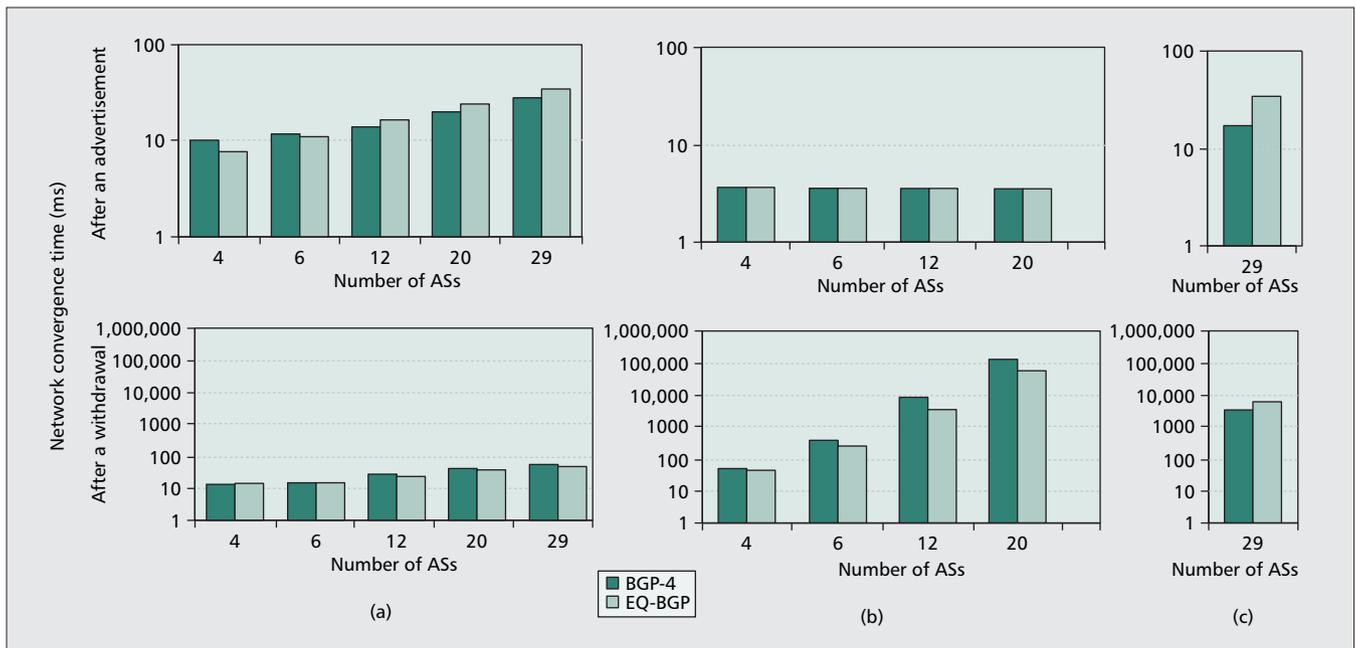
**■ Figure 5.** *Comparison of EQ-BGP and BGP-4 convergence time after a route advertisement or a route withdrawal, in the case of: a) Ring topology; b) full mesh topology; c) Internet like topology.*

QoS-NLRI information advertised to upstream domains (for scalability reasons only nominal QoS values are assembled and advertised to other domains).

• A *maximum value $M_i$* allowed for the QoS parameter $i$, taken from ITU Recommendations.

If a border router receives more than one update for the same destination, it selects the one with the *lowest* DoP. In fact, the DoP increases with the value of the QoS parameters, and goes to infinite (forcing the decision process not to select a specific route) if the value of one of the assembled QoS parameters exceeds the maximum $M_i$. The EQ-BGP decision process can be summarized as follows:

1 Choose the route with the highest BGP-4 local preference.
2 If the BGP-4 local preferences are equal, choose the route with the lowest DoP.
3 If the DoPs are equal, choose the route with the shortest AS path.
4 If the AS path lengths are equal, choose the route with the lowest BGP-4 MED.
5 If the BGP-4 MEDs are equal, prefer external routes over internal routes (eBGP over iBGP).
6 If the routes are still equal, prefer the one with the lowest Interior Gateway Protocol (IGP) metric to the next-hop router.
7 If more than one route is still available, run BGP-4 tie-breaking rules.

Different QoS preferences $f_i$ are assigned by TERO to different CoSs. For instance, for the IP telephony and real-time interactive CoSs, the IPTD and IPDV are equally important and more important than IPLR. However, for multimedia streaming, IPTD is more important than IPDV, whereas for high-throughput data, IPLR is the most important of all. While this mechanism provides a sufficient degree of flexibility, fine tuning of the QoS preference parameters requires extensive simulations and tests of the prototype.

## EVALUATION METHODOLOGY

In this section we evaluate the performance of EQ-BGP. Our objective is to analyze the impact of the new components in EQ-BGP on its scalability. The evaluation is performed by comparing the performance of EQ-BGP against BGP-4 for different topologies. To this end, we assess two different metrics:

• **The network convergence time (NCT)** defined as the total amount of time that elapses between the advertisement of a new prefix (or the withdrawal of a known one) and the time instant when the last update message caused by this event is processed.

• **The total number of messages** exchanged during the network convergence time.

In our experiments we consider three types of network topologies with different numbers of ASs: full mesh, ring, and a representative topology for the Internet. The full mesh topology was selected because it allows the maximum number of alternative paths, thus representing a "worst case" scenario. On the other hand, the ring network (or b-clique) is commonly used to analyze the routing decision algorithm, as there are exactly two disjoint paths between each pair of domains. To complete the evaluation, we analyze the performance of EQ-BGP in topologies derived from routers operating in the Internet [12].

For the sake of simplicity, we assume that each AS is represented by a single EQ-BGP router, connected to its neighbors through links of 1 Mb/s of capacity and constant delay of 1 ms. Although the link parameters were arbitrarily chosen, the results obtained for the NCT can easily be scaled taking into account actual link characteristics. In addition, we assume that all
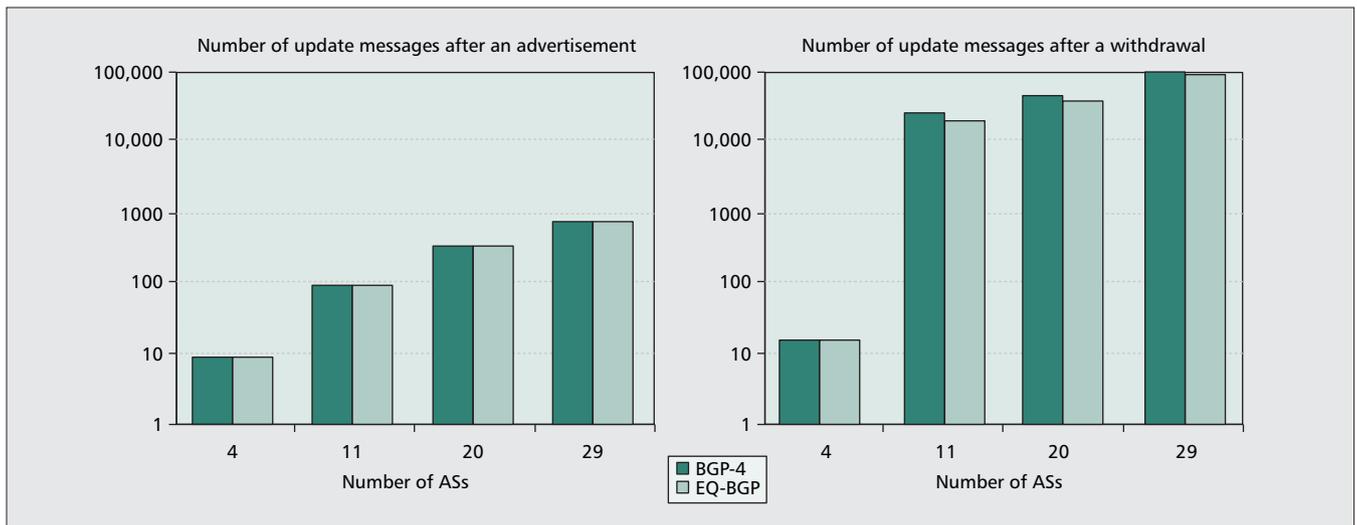
**■ Figure 6.** *Scalability of EQ-BGP vs. BGP-4.*

the ASs support the IP telephony CoS, supplying different values of IPTD. The IPTD values were randomly chosen from 1 to 10 ms, whereas the nominal values of IPDV, IPLR, and the corresponding parameters on interdomain links were the same for all domains.

Our experiments were performed using the ns2 simulator, in which the EQ-BGP protocol has been implemented. All experiments were performed assuming that the advertisement or withdrawal of a prefix occurs when the network is in a stable state (i.e., after it has already converged). Each simulation run was stopped when the last update message originated by the considered stressing event was processed. The results presented here were collected from 10 simulation runs, in which a randomly chosen AS advertises or withdraws a route. The reported values of convergence time include the 95 percent confidence interval. The next subsections present the results obtained in terms of both the network convergence time and the number of messages exchanged during a convergence.

### NETWORK CONVERGENCE TIME EVALUATION

Figure 5 shows the results for the NCT of the ring, full mesh, and Internet-like network topologies after the advertisement (or withdrawal) of a route. For the advertisement case, the full mesh topology exhibits the same NCT for EQ-BGP and BGP-4 regardless of the number of ASs. This can be explained by considering that all the ASs prefer the paths through direct links, given that the QoS level is usually better than that offered through alternative paths. Accordingly, the routing process ends at the same time for all the cases assessed. On the other hand, for the ring and Internet-like topologies, the NCT increases with both the number of ASs and the number of interdomain links. Moreover, the EQ-BGP protocol needs more time to converge. This is caused by the introduction of an additional degree of freedom compared to BGP-4, stemming from the possibility of assigning arbitrary QoS parameters in an AS instead of a single metric, as in the case of the AS path length with BGP-4. This new degree of freedom increases

the chances that the advertised path does better than the one currently used by several routers. These routers now switch their best path selection and advertise the change to their neighbors, which might in turn switch their best path selection as well, causing slower convergence.

The opposite effect can be observed in the case of a route withdrawal. EQ-BGP generally converges slightly faster than BGP-4. This is because alternative paths have assigned more information about their capabilities, and hence less suitable paths are removed faster.

Within the limits of the preliminary evaluations performed so far, the EQ-BGP protocol has proven to be stable and to exhibit a convergence time comparable to that of BGP-4.

### NUMBER OF MESSAGES EXCHANGED DURING CONVERGENCE

EQ-BGP is designed for multidomain networks, so assessing its scalability is an important part of performance evaluation. To achieve this, we compare the number of update messages processed by EQ-BGP and BGP-4 during a convergence in the worst-case scenario (i.e., the full mesh topology). Figure 6 confirms that EQ-BGP and BGP-4 require a similar number of messages to converge.

## DISCUSSION AND FUTURE WORK

This article introduces a complete system to solve the problem of finding and providing end-to-end QoS paths between users connected through heterogeneous access networks. The design, implementation, and test of such a system is the main target of the EuQoS research project. At present, a first prototype of the EuQoS system has been designed and developed, and is starting to be deployed on a real testbed throughout Europe. This prototype includes different access technologies such as WiFi, LAN, xDSL, and UMTS, for which specific solutions have been implemented, and incorporates key functions such as routing, signaling, and resource reservations between end users.

In terms of QoSR, most of the research efforts at this stage of the project have focused

on interdomain issues. We have designed and preliminarily tested an extended version of BGP-4, EQ-BGP, which is proved to be as scalable as BGP-4. Several issues, however, are still open. The project team is currently working on successive versions of the current prototype tackling such issues. Furthermore, we are refining the architecture design with the purpose of devising new and more advanced QoS solutions. We are also actively working on the design of a sound business model for deployment of the system.

## ACKNOWLEDGMENTS

## REFERENCES

[1] P. Van Mieghem *et al.*, "Quality of Service Routing," Ch. 2, *Quality of Future Internet Services: COST 263 Final Report*, Springer-Verlag, Oct. 2003.
[2] X. Masip-Bruin *et al.*, "Research Challenges in QoS Routing," *Comp. Commun. J.*, Elsevier, vol. 29, no. 5, Mar. 2006.
[3] "End-to-end Quality of Service Support over Heterogeneous Networks (EuQoS)," IST FP6-IP 004503, http://www.euqos.org/
[4] J. Rosenberg *et al.*, "SIP: Session Initiation Protocol," IETF RFC 3261, June 2002.
[5] Next Steps in Signaling (NSIS) IETF WG: http://www.ietf.org/html. charters/nsis-charter.html
[6] H. Schulzrinne and R. Hancock, "GIST: General Internet Signaling Transport," IETF draft-ietf-nsis-ntlp-09, Mar. 2006.
[7] J. Manner *et al.*, "NSLP for Quality-of-Service Signaling," IETF draft-ietf-nsis-qos-nslp-08, Apr. 2006.
[8] L. Cordeiro *et al.*, "Hybrid on-Path Off-Path Approach for End-to End Signaling Across NSIS DOMAINS (HyPath)," IETF draft-cordeiro-nsis-hypath-00, Feb. 2006.
[9] D. Durham *et al.*, "The COPS (Common Open Policy Service) Protocol," RFC 2748, Jan. 2002.
[10] G. Cristallo and C. Jaquenet, "An Approach to Interdomain Traffic Engineering," *XVIII World Telecommun. Cong.*, Paris, France, Sept. 2002.
[11] A. Beben, "EQ-BGP: an Efficient Inter-domain QoS Routing Protocol," *Proc. 20th Int'l. Conf. Advanced Info. Networking and Apps.*, Vienna, Austria, Apr. 2006.
[12] B. Premore, "Multi-AS Topologies from BGP Routing Tables," http://www.ssfnet.org/Exchange/gallery/asgraph/index.html, 2005.

## BIOGRAPHIES

XAVI MASIP-BRUIN received M.S. and Ph.D. degrees from the Technical University of Catalonia (UPC), Spain, both in telecommunications engineering, in 1997 and 2003, respectively. He is with the Advanced Networks Architectures Laboratory at UPC, where he is an associate professor of computer science. His current research interests lie in broadband communications, QoS management and provision, and traffic engineering.

MARCELO YANNUZZI (yannuzzi@ac.upc.edu) received a degree in electrical engineering from the University of the Republic, Uruguay, in 2001, and a DEA from UPC in 2005. He is with the Advanced Networks Architectures Laboratory at UPC, where he holds a research and teaching position. His research interests are in the areas of interdomain routing, traffic engineering, and network stability in both IP and optical networks.

RENÉ SERRAL-GRACIÀ is a Ph.D. student and assistant professor in the Department of Computer Architecture at UPC. There he received an engineering degree in information technologies (2003). He is also a researcher at the Advanced Broadband Communications Center (CCABA) of UPC. His main research topics are QoS management and provision, traffic engineering, and IP traffic analysis and characterization.

JORDI DOMINGO-PASCUAL is a full professor of computer science and communications at UPC. He is co-founder and researcher of the Advanced Broadband Communications Center of the university (CCABA). His research topics are broadband communications and applications, IP/ATM integration, QoS management and provision, traffic engineering, IP traffic analysis and characterization, and QoS measurements.

JOSÉ ENRIQUEZ GABEIRAS received a Master's degree in engineering from the Technical University of Madrid, Spain, in 1991. He is head of the Advanced Networks Planning Department of Telefónica R&D. This group is responsible for studies on planning, dimensioning, and management of broadband networks in Telefónica's operations. He is presently the leader of project EuQoS (end-to-end QoS over heterogeneous networks) under the FP6 IST program of the European Commission.

MARIA ÁNGELES CALLEJO RODRIGUEZ received a Master's degree in engineering from the Technical University of Madrid in 2004. Since 2004, she has been working for Telefónica R&D in the Advanced Networks Planning Department on the following topics: traffic engineering and traffic control alternatives for IP networks and ambient networks environments, multipath routing with dynamic variance, and provision of QoS over heterogeneous networks.

MICHEL DIAZ [SM] is research director at the French CNRS. He was director of French research on architecture, networks, systems, and parallelism from 1994 to 1999, and is now coordinating French research groups on communication networks. He is expert for many European and French programs and projects. He served as a Chairman or member of many Program Committees. He is the editor of 11 books, and has written one book and more than 200 technical publications. He received the Silver Core of the IFIP, is member of the New York Academy of Sciences, and is listed in *Who's Who in Science and Engineering*.

FLORIN RACARU is preparing his Ph.D. at the Laboratory for Architecture and Analysis of Systems of the National Center for Scientific Research (LAAS-CNRS), France. At present, he is working on the Tools and Software for Communication team. His research interests include QoS over the Internet, network architectures, and protocols.

GIOVANNI STEA is an assistant professor with the Department of Information Engineering of the University of Pisa, Italy. He has been and is involved in national and European research projects. He has served as a member of the technical and organization committee for several international conferences, including European Wireless, SIGCOMM, and VALUETOOLS. His current research focuses are QoS and traffic engineering.

ENZO MINGOZZI has been an associate professor on the Faculty of Engineering of the University of Pisa since January 2005. His current research interests focus on the design and analysis of MAC protocols for wireless networks, and QoS provisioning and service integration in computer networks.

ANDRZEJ BEBEN received M.Sc. and Ph.D. degrees in telecommunications from Warsaw University of Technology (WUT), Poland, in 1998 and 2001, respectively. Since 2001 he has been an assistant professor with the Institute of Telecommunications at WUT, where he is a member of the Telecommunication Network Technologies research group. His research areas include ATM, IP networks (fixed and wireless), traffic engineering, simulation techniques, measurement methods, and testbeds.

WOJCIECH BURAKOWSKI received M.Sc., Ph.D. and D.Sc. degrees in telecommunications from WU) in 1975, 1982, and 1992, respectively. Since 1997 he has been a professor at WUT (tenured since 2003). Since 2002 he has been the research director of the Institute of Telecommunications at WUT. Moreover, he leads the Telecommunication Network Technologies (TNT) research group that has been involved in several EU research projects. His research areas include ATM, QoS IP, heterogeneous networks (fixed and wireless), traffic engineering, simulation techniques, measurement methods, and testbeds.

EDMUNDO MONTEIRO is an associate professor at the University of Coimbra, Portugal, from which he graduated in 1984 and got a Ph.D. in electrical engineering — computer specialty in 1995. His research interests are computer communications, QoS, mobility, routing, resilience, and security.

LUIS CORDEIRO has a computer engineering degree from the University of Coimbra where he is currently a Ph.D. student working on IP network signaling. He is involved in the EuQoS European Integrated Project and the Next Steps In Signaling (NSIS) IETF Working Group where he is promoting discussion of the HyPath draft as first author.

*In terms of QoSR, most of the research efforts at this stage of the project have focused on interdomain issues. We have designed and preliminarily tested an extended version of BGP-4, EQ-BGP, which is proved to be as scalable as BGP-4.*