


A FRAMEWORK FOR DYNAMIC SLA-BASED QoS CONTROL FOR UMTS

RAJIV CHAKRAVORTY, IAN PRATT, AND JON CROWCROFT
UNIVERSITY OF CAMBRIDGE COMPUTER LABORATORY
MAURIZIO D'ARIENZO, UNIVERSITY OF NAPLES ("FEDERICO II")



With the evolution of QoS-capable 3G wireless networks, the wireless community has been increasingly looking for a framework that can provide an effective, network independent, end-to-end Quality of Service QoS control.

ABSTRACT

With the evolution of QoS-capable 3G wireless networks, the wireless community has been increasingly looking for a framework that can provide effective network-independent end-to-end QoS control. In this article we first construct such a framework and then describe how dynamic SLA-based control can be used to achieve end-to-end QoS in a wired and wireless (UMTS) environment. The proposed framework, which is an extension to the IST CADENUS project, offers effective wired-wireless QoS translation, efficient QoS control and management, and dynamic SLA policy-based QoS provisioning.

INTRODUCTION

The Universal Mobile Telecommunications System (UMTS) standard for third-generation (3G) wireless networks is gradually evolving, and networks based on this standard are beginning to be deployed throughout Europe. With the Internet's ever expanding dimensions, constituting heterogeneous wired and wireless networks, focus is shifting toward a model that can achieve end-to-end quality of service (QoS) control rather than QoS control models for specific networks. Take, for instance, European Telecommunications Standard Institute (ETSI) UMTS Release 99 [1], where work based on similar lines to that of the Internet Engineering Task Force (IETF), with a penchant for an "all-IP" architecture, targets achieving end-to-end wired-wireless QoS control rather than plain wireless QoS.

However, despite the efforts made in this direction, the goal of realizing an effective end-to-end QoS control scheme in a wired-wireless environment is still unfulfilled. There are some important challenges that have to be overcome:

Network-level QoS translation (mapping): Service differentiation based on a set of traffic classes, for both wireless as well as IP core, needs an effective and reliable QoS translation (or QoS mapping) mechanism.

Dynamic QoS management: This includes dynamic QoS monitoring and control: proper

monitoring of the network as per user agreements, and control that can dynamically arbitrate or modify the QoS available from the network.

Control management infrastructure: Networks will have to ensure QoS control using their own sets of control mechanisms. These mechanisms will enable establishment, maintenance, and termination of underlying network QoS. To achieve interoperability between these control mechanisms in the quest for end-to-end QoS control would necessitate a framework that takes care of both wired as well as wireless QoS.

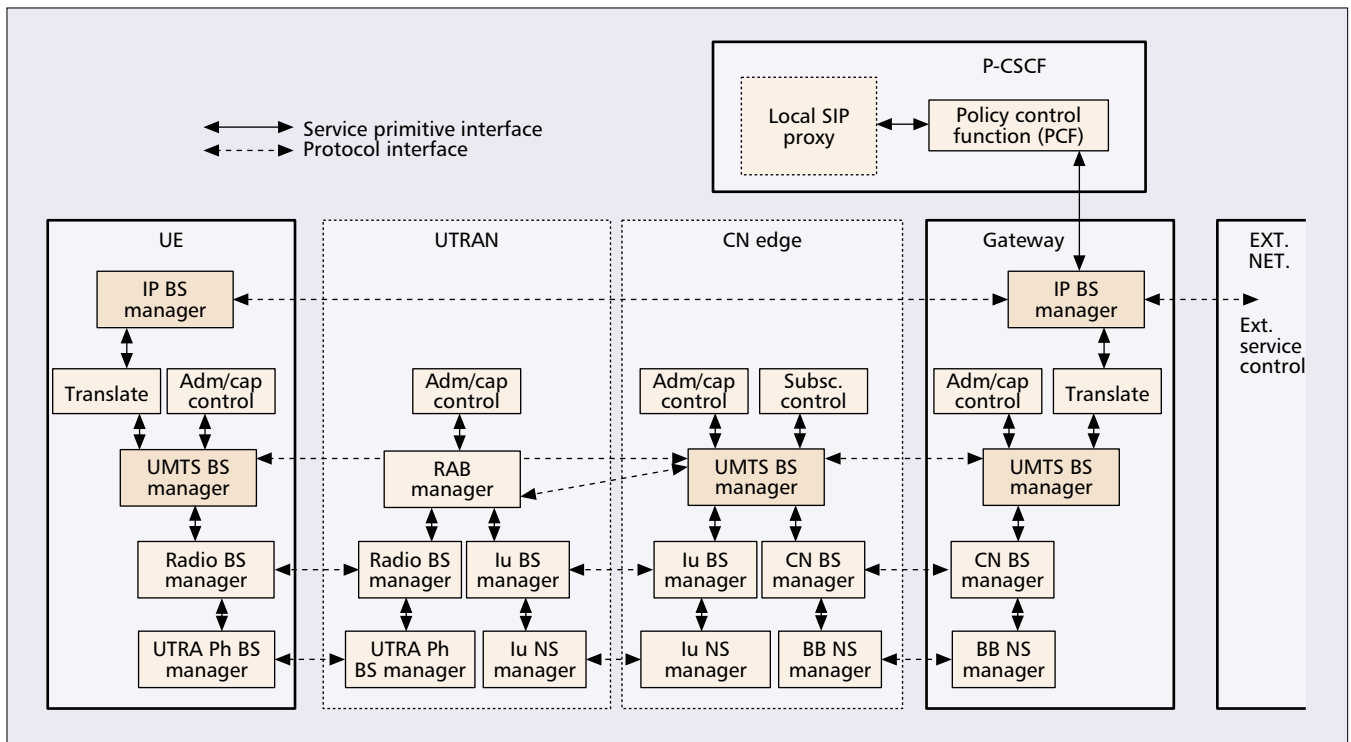
In this article we address the aforementioned issues with a simple extension to the CADENUS framework [2] to achieve simple yet effective end-to-end QoS control in a wired-wireless environment. The CADENUS framework currently only defines an integrated architecture for creation, configuration, and provisioning of end-user services for wired networks that offer some kind of service differentiation (e.g., premium IP networks) [3].

A first step in this direction would be to interwork QoS of both wired as well as wireless (UMTS) networks. ETSI specifies different ways of interworking UMTS QoS with Internet QoS. The Third Generation Partnership Project (3GPP) [4] currently specifies the use of:

- Signaling; for example, Resource Reservation Protocol (RSVP) and Label Distribution Protocol (LDP) along the flow path
- Packet marking or labeling, such as differentiated services (DiffServ) and multiprotocol label switching (MPLS)
- Interworking policy control with network resource mediators
- Using service level agreements (SLAs) enforced by the network border routers

In this article we discuss how we can achieve end-to-end QoS control using a combined mix of dynamic SLA-based and policy control schemes.

The article is structured as follows. We first review QoS related issues pertaining to UMTS. We then delineate the CADENUS framework and describe the proposed CADENUS extension for UMTS. We discuss our initial evaluation through simulations, and our ongoing research over a real experimental testbed.



■ Figure 1. UMTS QoS management (control plane).

UMTS QUALITY OF SERVICE ISSUES

UMTS ARCHITECTURE

UMTS defines a system architecture that consists of a number of logical network elements: the UMTS terrestrial radio access network (UTRAN) that handles all radio-related functionality, the circuit-switched (CS) and packet-switched (PS) domains with two support nodes: the serving GPRS support node (SGSN) and gateway GPRS support node (GGSN), and the core network (CN) adopted from the General Packet Radio Service (GPRS) architecture [4].

The core network in the packet-switched domain connects UTRANs with the external networks through the GPRS support nodes. User equipment (UE) that interfaces the user and radio interfaces is also defined, which consists of mobile equipment (ME) and the UMTS subscriber identity module (USIM). Thus, packet-switched UMTS can support end-to-end IP services with guaranteed QoS as defined by the 3GPP [1, 4]. These end-to-end services involve the UE, UMTS access, and CN.

UMTS QoS MANAGEMENT

UMTS achieves QoS management using a layered architecture, with bearer services (BS) established between UMTS modules at different layers. Each BS deals with control signaling, user-plane transport, and QoS management. End-to-end QoS control is possible by interworking TE/MT BS, UMTS BS, and external BS. To provide service differentiation, a UMTS network supports different BS that correspond to similar differentiation to that applied in the IP CN. For QoS control with an external IP network, 3GPP [1] dictates use of an IP BS manag-

er within the UMTS GGSN node (Fig. 1). The IP BS manager provides QoS control for an IP core using DiffServ edge functionality (or an RSVP function).


ETSI 3GPP Release 99 [1] specifies that the GGSN *should* implement DiffServ edge functionality, along with an IP policy enforcement point (PEP) and optionally RSVP/integrated services (IntServ) function at its IP BS manager. On the other hand, UE may support the same in its IP BS manager, depending on what capability it wants its IP BS manager to have. While the UE may support all or none of the functionality, it is the responsibility of the UMTS network to ensure end-to-end QoS [4].

From the 3GPP standardization point of view, the UE determines its QoS requirements using an application layer scheme (e.g., Session Description Protocol, SDP, used with Session Initiation Protocol, SIP), which then maps QoS requirements to the Packet Data Protocol (PDP) context parameters or IP layer parameters (e.g., RSVP messages).

When a GGSN learns about QoS requirements (through PDP context or RSVP messages), it may use this IP level information to configure the DiffServ classifier and provide interworking between the PDP context and backbone IP network. This information is made available using an *authorization token* located in the PDP context messages, or from the RSVP messages directly. An authorization token consists of the policy and QoS interworking functions such as packet classifier, QoS information (SDP-based), packet handling action, and event generation information (for usage recording purposes).

If a UE supports RSVP, the GGSN can use RSVP instead of the PDP context to control QoS through the backbone IP network. In this

When strict guarantees are necessary, for instance in real-time voice and video applications, using the conversational traffic class can be a good choice. The streaming class is slightly relaxed in terms of delays, to which other streaming applications could be mapped.



■ Figure 2. The CADENUS architecture.

case, RSVP sessions from a UE may be used to configure the DiffServ classifier functionality directly. However, if the GGSN does not support RSVP, the messages may transparently pass through the GGSN. For UE supporting only UMTS QoS mechanisms, application QoS requirements can be signaled to the IP BS manager at the GGSN, by mapping the UMTS QoS parameters to the PDP context activation messages. The admission/capability function takes care of resource allocation based on availability, some other policy, or service invocation-based administrative issue. The subscription function authorizes usage of the service for a particular user, while the translation function at the edge converts service primitives of the UMTS BS to primitives of the external network.

On the other hand, the UMTS data plane is responsible for traffic classification, mapping, conditioning, and resource management. The classification function accommodates a specific flow of packets based on packet headers or some other means of classification, while the mapping function maps service classes in order to derive the intended QoS from the core networks. The policing function checks if the traffic profile is consistent with what was negotiated, and if it is not, the shaping function shapes (marks, drops, or delays) the packets for traffic compliance. The resource manager schedules packets, and performs bandwidth and queue management functions. However, unlike the downlink, the traffic conditioning for the uplink is performed directly by the mobile terminal.

UMTS QoS TRAFFIC CLASSES

UMTS defines four QoS traffic classes: *conversational*, *streaming*, *interactive*, and *background*. The conversational class provides strict delay guarantees, while the background class offers no qualitative or quantitative guarantees. It is a *best effort* class. When strict guarantees are neces-

sary, as in real-time voice and video applications, using the conversational traffic class can be a good choice. The streaming class is slightly relaxed in terms of delays, to which other streaming applications could be mapped.

While interactive traffic follows a request-response pattern and can only justly provide qualitative guarantees, the background class is similar to best effort traffic consisting of bulk (e.g., ftp) and asynchronous traffic flows (e.g., email) that fall into this category. Besides QoS classes, attributes have also been defined in support of the UMTS QoS classes.


THE CADENUS FRAMEWORK

The CADENUS framework investigates service and resource management aspects of networks that an Internet service provider (ISP) or an operator may consider for possible service differentiation [3]. The framework incorporates two key ideas: a generalized mediation concept, and a contract negotiation and translation feature for all the components.

CADENUS COMPONENTS

The CADENUS framework proposes a CADENUS Mediation Component Architecture (CMCA) [3] that partitions system functionalities into three major components: an access mediator, a service mediator, and a resource mediator (Fig. 2).

Access mediator (AM): The AM is the device into which users input their requests to the system. It adds value to the user by presenting a wider selection of services, ensuring lowest cost and offering a harmonized interface to the user. It serves to assist and ease the service selection process. It is also responsible in selecting the appropriate service mediator(s), according to the request made by the users. The other main functions include automation of r-SLAs (retail



■ Figure 3. CADENUS extension for UMTS.

SLAs), *authorization-authentication-accounting* (aaa), static or dynamic negotiation of r-SLAs, and interaction with directory services for service specific information.

Service mediator (SM): The SM is the place where services are created and from which the impact of service reconfiguration is communicated to network resource management. It is responsible for finding the service by requesting information from the appropriate resource mediator(s). In some cases, it may also build services from individual elements — the service itself. It maintains no direct contact with end users for SLAs, or for authentication or accounting. Instead, it is a unique entry point for presentation of the service and subscriptions, the contract profile, and access to the service chosen.

Resource mediator (RM): The RM selects appropriate network capabilities from multiple network providers and network technologies, given several available options. The RM is a key functional entity in CADENUS architecture that has a complete end-to-end view of the QoS setup process. From an application viewpoint, the RM only provides an abstract view of the underlying network and the differentiation it can offer.

The combined role of these components is to efficiently manage user access to the service, present the portfolio of available services, and appropriately configure and manage the QoS-aware network elements available in the underlying network infrastructure.

CADENUS SERVICE DEFINITION

The CADENUS architecture offers QoS services *on demand* to end users. The AM component allows users to select a service with a specific QoS. This process invokes the SLA, and consequently one or more service level specifications (SLSs) are created in the transaction. SLSs help


fulfil the technical requirements of the SLA, and their scope includes flow identification, traffic conformance, performance guarantees, and reliability information.

In the CADENUS framework, service contract definitions are of two types: retail SLAs (r-SLA) as a unitary service between customer and service provider, and wholesale SLAs (w-SLA) for contracts between service providers. An SLA is contracted (signed) by a customer or, more exactly, when a customer subscribes or modifies his/her service contract with his/her service provider. This modification step is offline; however, from the perspective of wireless networks like UMTS, this presents a difficulty since QoS available in wireless networks can vary considerably during service usage. Hence, it should be possible to be able to perform QoS (re)negotiation on the fly, necessitating use of dynamic SLAs within the CADENUS framework. The benefit of using dynamic SLAs as opposed to simple static ones is that decisions about user QoS can be made in real time.

CADENUS EXTENSION FOR UMTS

Our proposed extension, the CADENUS-UMTS Extension (CUE), consists of two new independent functional components: the CUE-SM and CUE-RM (Fig. 3). The extension lacks a third component from the original CADENUS framework, the AM. This is because the AM is used where users have to explicitly map their service requests to the system (as in premium IP networks), whereas in UMTS such requests typically take the form of PDP context activation/change request messages in real time. This is also one reason why dynamic SLAs are more appropriate in the context of a UMTS framework. The SLA “dynamicness” in UMTS is such

Once the authorization token is available to GGSN, it can trigger an event which is forwarded to the CUE-RM that decides, based on some admission control scheme, whether to accept or reject the user application QoS requirements.



■ **Figure 4.** A scenario for control management in CUE.

that SLAs can change *on the fly* during a live session (e.g., during on-the-fly codec changes in a voice over IP session). SLAs can be negotiated per session (during session initialization using PDP) or in the midst of a session typically through (re)negotiations (PDP context modify requests). However, the concept of static SLAs for UMTS users is still applicable. This can correspond to the contract negotiated by the customer at the start. Such contracts may provide SLAs that are *time-varying* but essentially static, yet application requirements may change the SLA depending on user demand. Once a contract is negotiated by a user, SLAs (in particular r-SLAs) are made available in the form of a residual user QoS profile located in the user's home UMTS network; mostly profiles will be located in the home location register or location register (HLR/LR) of the UMTS network. Once a user profile is updated, it is converted into one or more SLSs by the CUE-SM, and communicated to the CUE-RMs. The CUE-RM will use the SLS to statically configure the UMTS GGSNs accordingly. Notice that this step is indicative of a provisioning based resource allocation model rather than a simple outsourcing one.

Also, there exists little functional distinction between the CUE framework and IETF's policy-based framework [5]. In terms of functionality, an SM is a policy repository and policy translation system, while an RM can be compared to a policy decision point (PDP). Contrast this with a 3GPP local policy-based scheme [1], where a proxy-call center control function (P-CSCF) with a built-in policy control function (PCF) has functionality similar to that of a PDP. The PEP in this case is the GGSN. A PCF (here the PDP) can be located anywhere within the provider's network. Upon reception of a request from a user, the GGSN (PEP) can retrieve the authorization token from a PDP context activation

message and trigger the request to the PDP (P-CSCF) that performs QoS authorization (bandwidth, delay, etc.) for the session. The interface used between GGSN and PCF is known as the Go interface [4], which allows service-based local policy and QoS interworking based on Common Open Policy Service (COPS) compliance. The initial authorization operation specified by the 3GPP is a pull operation (outsourcing based), but subsequent operations (modify) may be push or pull.

In CUE, however, service-based policy and QoS interworking is based purely on a policy outsourcing (pull) model [6]. Also, a PEP as in CUE need not necessarily be located in the GGSN as specified by 3GPP; however, the basic functionality applicable to the data and control plane specific to the 3GPP standards would still remain. For instance, policy decisions may be stored by the COPS client in the CUE-RM to achieve a degree of separation. Furthermore, events in GGSN that demand change/modification (in a PDP change request message) will also trigger events to the CUE-RM, which will then decide on the admission of this new session and respond appropriately. A pure outsourcing based policy control model demands a new COPS client type that makes it possible to have dynamic policy-based resource allocation.

CONTROL MANAGEMENT IN CUE

As per the 3GPP specifications, it is the responsibility of the UMTS network to ensure QoS from the core IP network even if UE is not capable of it. A UE may typically convey its application QoS requirements to the GGSN IP BS manager using some appropriate signaling mechanism. This is possible using RSVP control messages, or PDP context activation/modify request messages consisting of the authorization token.

Once the authorization token is available to GGSN, it can trigger an event that is forwarded to the CUE-RM, which decides, based on some admission control scheme, whether to accept or reject the user application QoS requirements (Fig. 4). If the QoS requested by the application is below a certain threshold originally contracted by the user, static SLAs can take over. In this case, the QoS profile will be retrieved by the CUE-SM from the user profile database and communicated to the CUE-RM. The final decision about the admission is conveyed to the GGSN by the CUE-RM, which then responds to the user in a PDP response message. This whole (re)negotiation process is based on the dynamic SLA-based policy outsourcing model, and is similar to the one used in [6].

From the 3GPP viewpoint, the control and data plane functionality is an integral part of the network elements. While the control plane functionality has to be tied to the GGSN, it may be a good idea to have part of the data plane functionality collocated in an edge router. Specifically, since data traffic handling and scheduling can have additional computational overhead, this job can become part of a collocated edge router (Fig. 3). Furthermore, the edge router can also take decisions on the traffic flows based on some sort of feedback available from the CUE-RM or, alternatively, from the GGSN.

CUE QoS TRANSLATION

An important goal in end-to-end QoS control over wired and wireless networks is to ensure effective translation of QoS traffic classes. Hence, mapping rules have to be defined so that functions can interoperate efficiently. The DiffServ framework offers assured forwarding (AF) and expedited forwarding (EF) as *per-hop behaviors* (PHBs) for a DiffServ (DS) compliant node. The EF PHB group is used when low loss, low latency, low jitter, and assured bandwidth is required for end-to-end service. This group can offer deterministic service guarantees and a service likened to that of a virtual leased line. A single codepoint is defined for the EF class. The AF (read as *AFij*) class allows a DS domain to provide different levels of guarantees for forwarding IP packets. Currently, four classes ($N = 4$, $1 \leq i \leq N$) with three levels of drop precedence in each class ($M = 3$, $1 \leq j \leq M$) are defined for general use. An example usage in AF is that each class represents a higher level of service (e.g., platinum = 1, gold = 2, silver = 3, bronze = 4) with low, medium, and high ($j = 1, 2, 3$, respectively) drop precedence levels. Thus, AF11 represents the best, AF43 the worst service level.

For interworking purposes, multiple QoS mapping levels can be defined. Here, two approaches can be envisaged.

One-to-one QoS mapping: This is a strict one-to-one QoS mapping, which maps each UMTS QoS class to a corresponding DiffServ QoS class. However, one-to-one mapping might not always be possible since networks may support different sets of QoS classes. A sample one-to-one mapping for DiffServ-UMTS QoS interworking is given in [7]. In this case, the traffic handling priority attributes of the UMTS

UMTS QoS class	Mapped DiffServ class	Mapping criteria
Conversational	EF only	Strict mapping for deterministic QoS guarantees, low latency, and jitter
Streaming	AF11, AF12, AF21, AF22, etc.	Low jitter
Interactive	AF13, AF23, AF22, AF32, etc.	Relatively low latency and loss probability
Background	AF42, AF43, etc.	Best effort

■ **Table 1.** Sample mapping of UMTS QoS traffic classes with DiffServ IP core.

interactive class map to the drop precedence of the AF class as they both share the same number of priority levels internally.

Many-to-one mapping: This can map a number of DiffServ QoS traffic classes into a single UMTS QoS traffic class. A DiffServ core can define many QoS traffic classes (using AF PHB) when compared to only limited QoS classes supported by UMTS; then a close set of DiffServ QoS traffic classes having almost similar QoS requirements can be merged into a single UMTS QoS class.

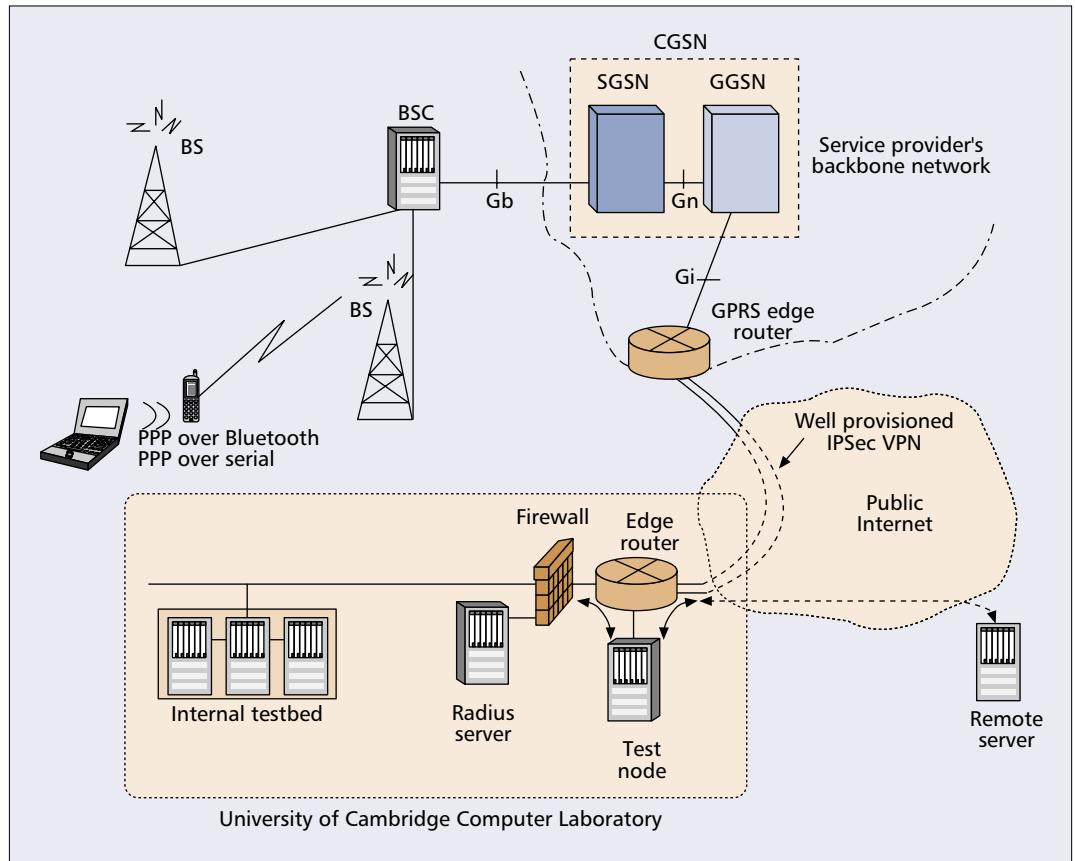
A sample interworking of QoS traffic classes between DiffServ IP core and UMTS QoS classes is shown in Table 1. Note that a number of such QoS mapping levels can be defined depending on how many QoS classes are supported by the UMTS and DiffServ IP core. Some other studies consider strict one-to-one QoS mapping (e.g., [7]); however, ours is a more *flexible* approach where strict one-to-one mapping is enforced only for certain traffic class to support strict QoS translation where necessary (e.g., between EF-conversational class for deterministic QoS guarantees), while one-to-many mapping may be used where QoS requirements are somewhat more relaxed. Such a scheme is able to offer adequate flexibility and control of QoS translation as desired by network providers, while not trading any of the QoS requirements for the supported traffic classes.

A COMPARATIVE SIMULATION STUDY

While the use of QoS profiles has been specified by ETSI for GPRS/UMTS traffic classes, how QoS should be managed by means of traffic scheduling, shaping, and connection admission control remains an open implementation issue.

In order to understand these issues, we performed a simulation study for GPRS consisting of different scheduling schemes: first in first out (FIFO), priority FIFO, static priority scheduling (SPS), shortest job first (SJF) and its two variants, earliest deadline first (EDF), weighted round-robin (WRR) and the token bank leaky bucket (TBLB) algorithm. A separate report on this comparative simulation study is available [8]. By simulating traffic related to an advanced travelers information system (ATIS) at the IP level, we compared scheduling algorithms for use in GPRS across different performance metrics such as the average packet data channel (PDCH, in GPRS) occupation, average waiting time in a scheduler, and also packet drop probability. We used four

For the case of the average time in the scheduler, we found a variant of the shortest job first to give good results. We plan to extend this simulation study using a parameterized model based on the real-world, long-term, traffic traces that we are collecting over a commercially deployed GPRS network.



■ **Figure 5.** The testbed for CUE evaluation.

different traffic classes based on different priority levels, and an ATIS traffic scenario based on the urban (80 percent of mostly light messaging traffic) and suburban (80 percent of session-oriented) traffic patterns. We used a dynamic channel allocation scheme based on the *capacity on demand* principle between GSM and GPRS users.

Results from the comparative simulation study show that in terms of packet drop probability, earliest deadline first (EDF) performs reasonably well for the majority of the priority traffic classes and types (urban and suburban), while FIFO and shortest job first (SJF) are suitable for the best effort class. In terms of average PDCH occupation, most schemes give acceptable performance, the only exception being TBLB, where performance can decline with increase in user arrival rates. For the case of the average time in the scheduler, we found a variant of the shortest job first to give good results. We plan to extend this simulation study using a parameterized model based on real-world long-term traffic traces we are collecting over a commercially deployed GPRS network.

ONGOING CUE TESTBED EVALUATION

Besides simulation study, our current research involves real evaluation using an experimental GPRS testbed. In the GPRS network testbed from a commercial operator we have already used for previous work [9], both the SGSN and GGSN nodes are collocated in a combined GPRS support node (CGSN). We have access to

the GPRS GGSN node (Fig. 5) through a separate test node located in our laboratory (but connected to the GPRS CGSN using an IPsec virtual private network, VPN), with the routing configured such that all packets flowing to and from our GPRS mobile hosts are passed to the test node for processing. Such a test node can be made to perform control and data plane functions expected of a GGSN (traffic classification, mapping, and scheduling). Ongoing research integrates the CADENUS components (CUE-SM and CUE-RM) atop this infrastructure.

RELATED DISCUSSION

Related research has also addressed QoS mapping; for instance, S. Maniatis *et al.* [7] introduce the *RCL architecture* that presents such a mapping for traffic classes between UMTS and a core IP network, and largely trades end-to-end QoS based on static SLAs.

Also, most proposals that consider resource negotiation in DiffServ networks are essentially *static*, meaning QoS can be negotiated only for the long term (e.g., days); hence the term static SLAs. IETF's RAP group defines the COPS protocol and its extension, COPS provisioning model (COPS-PR) that can be used to statically provision resources with a DiffServ network. A static resource provisioning model is simple; however, it can lead to underutilization of network resources and may compound the adaptation process at times of high variability in the traffic demand.

To further evolve schemes that allow dynamic provisioning, S. Salsano *et al.* [6] propose a new COPS client type, COPS-DRA, that allows dynamic resource allocation and policy control for DiffServ-based networks using COPS for signaling. The new client supports combinations of both resource allocation models: outsourcing and provisioning. It uses signaling for resource admission control and QoS-aware call setup for SIP-based applications. It enables dynamic provisioning by exchanging resource allocation requests from the edge routers to a logically centralized bandwidth broker (BB), which it finally maps in a policy-based PDP-PEP relationship.

L.-N. Hamer *et al.* [10] consider COPS-PR for use in the UMTS framework. Their proposal considers a policy outsourcing scheme, but using COPS-PR, where a PEP informs the PDP (during initialization) what types of event it should trigger during policy outsourcing. Thus, when an event occurs the PEP triggers a COPS REQ to the PDP (here a UMTS PCF), which then retrieves session information and makes an informed decision that is finally pushed to the PEP using standard COPS-PR semantics. They show that suitably sharing the events between a PEP and PDP using COPS-PR can help redefine the degree of “dynamicness” of a resource provisioning model.

CONCLUSIONS

In this article we provide an architecture for end-to-end QoS control in a wired-wireless environment with effective QoS translation, proper control management, and dynamic SLA-based resource provisioning. We achieve this in our CUE framework, which is an extension of the CADENUS architecture. To derive all the benefits of the CADENUS framework, the CUE architecture adds two new components, CUE-SM and CUE-RM, that can be used to provision end-to-end QoS in a wired-wireless network. The framework makes use of dynamic QoS arbitration, by using PDP context activation/modify messages, which can be changed in real-time session. Ongoing research involves a thorough study of wired-wireless QoS interworking issues through simulations, and a practical performance evaluation of the framework over our testbed.

ACKNOWLEDGMENTS

The authors wish to thank Vodafone Group R&D, Sun Microsystems Inc., and Benchmark Capital for supporting this work. The research is also partially funded by IST project CADENUS IST-1999-11017 “Creation and Deployment of End-User Services in Premium IP Networks” and the ongoing COST-263 action on “Quality of Future Internet Service.”

REFERENCES

- [1] ETSI 3GPP, “QoS Concept and Architecture (Release 1999),” 3GPP TS 23.107 v. 3.9.0, Sept. 2002.
- [2] IST’s CADENUS Project, <http://www.cadenus.org>
- [3] G. Cortese *et al.*, “CADENUS: Creation and Deployment of End-User Services in Premium IP Networks,” *IEEE Commun. Mag.*, vol.41, no. 1, Jan. 2003, pp. 54–69.
- [4] ETSI 3GPP, “End-to-End QoS Concept and Architecture (Release 5),” 3GPP TS 23.207 v. 5.5.0, Sept. 2002.
- [5] K. Chan *et al.*, “COPS Usage for Policy Provisioning,” IETF RFC 3084, Mar. 2001.
- [6] S. Salsano *et al.*, “QoS Control by Means of COPS to Support SIP-Based Applications,” *IEEE Network*, Mar.-Apr. 2002, pp. 27–33.
- [7] S. Maniatis, E. G. Nikolouzou, and I. S. Venieris, “QoS Issues in the Converged 3G Wireless and Wired Network,” *IEEE Commun. Mag.*, Aug. 2002, pp. 44–53.
- [8] M. D’Arienzo *et al.*, “A Comparative Simulation Study for Multiple Traffic Scheduling Algorithms over GPRS,” Tech. rep., Comp. Sci. Dept., Unive. of Naples (“Federico II”), Italy, <http://www.grid.unina.it/>
- [9] R. Chakravorty and I. Pratt, “Performance Issues with General Packet Radio Service,” *J. Commun. and Networks*, Special Issue on Evolving from 3G Deployment to 4G Definition, vol. 4, no. 2, Dec. 2002, pp. 266–81.
- [10] L.-N. Hamer *et al.*, “COPS-PR for Outsourcing in UMTS: UMTS Go PIB,” IETF draft, draft-hamer-rap-cops-umts-go-00.txt

BIOGRAPHIES

RAJIV CHAKRAVORTY (Rajiv.Chakravorty@cl.cam.ac.uk) is a second-year Ph.D. student at the University of Cambridge Computer Laboratory, where he holds a SUN Microsystems student fellowship. Prior to joining his Ph.D., he was working with Philips Research — ASA Labs, Eindhoven, The Netherlands. He earned his M.Tech. degree from the Indian Institute of Technology (IIT) Delhi, where he was a recipient of the DAAD Scholarship Award from Germany for which he carried out his research at the Chair of Communication Networks (ComNets), RWTH-Aachen, Germany. His current interests include mobile and ubiquitous systems, wireless and wired-wireless integration, and QoS.

MAURIZIO D’ARIENZO is a research fellow at the Computer Engineering and Systems Department, University of Naples (“Federico II”), Italy. After a Laurea degree in electronics engineering, he earned a Ph.D. degree in computer science, both from the same university. His main research interests are focused on communications and networking, in particular with respect to distributed systems, network management, and QoS in IP networks.

IAN PRATT is a senior lecturer at the University of Cambridge Computer Laboratory. He holds a Ph.D. in computer science and was elected a Fellow of King’s College, Cambridge in 1996. As a member of the Laboratory’s Systems Research Group for over 10 years, he has worked on a number of influential projects, including the Desk Area Network workstation, the Nemesis operating system, the Cambridge Open Mobile Systems, and the XenoServer infrastructure for public distributed computing. His research interests cover a broad range of systems topics, including computer architecture, operating system design, mobile systems, and networking.

JON CROWCROFT [SM] is the Marconi Professor of Networked Systems in the Computer Laboratory of the University of Cambridge. Prior to that he was professor of networked systems at UCL in the Computer Science Department. He is a Fellow of the ACM, the British Computer Society, the IEE, and the Royal Academy of Engineering. He was a member of the IAB and was general chair for ACM SIGCOMM ’95–’99. He is on the editorial team for COMNET, and on the program committee for ACM SIGCOMM and IEEE INFOCOM. He has published five books; the latest is *Linux TCP/IP Implementation* (Wiley, 2001).

Ongoing research involves a thorough study of wired-wireless QoS interworking issues through simulations, and a practical performance evaluation of the framework over our testbed.