

Enhancing IP Service Provision over Heterogeneous Wireless Networks: A Path toward 4G

Luca Becchetti, Francesco Delli Priscoli, and Tiziano Inzerilli, University of Rome "La Sapienza," Italy
Petri Mähönen, University of Oulu, Finland
Luis Muñoz, University of Cantabria, Spain

ABSTRACT

Second-generation mobile radio systems have been deployed successfully worldwide and have evolved to higher data rates and services. Third-generation mobile radio systems are currently starting to be developed in different regions of the world. Today, the open question is how the third-generation systems will evolve. It is very likely that fourth-generation systems will not be a single standardized air interface, but a set of different technologies and standards. In particular, wireless LAN/wireless PAN type systems are designed for high/medium-data-rate access, low range, and, in general, low mobility. They are applicable to corporate networks and public access as a complement to cellular mobile radio systems for hot spot applications such as airports, hotels, and campuses. In this specific WLAN/WPAN framework and to guarantee an agreed QoS provision over such infrastructures, we propose a solution based on the wireless adaptation layer approach. In particular, aspects related to wireless link impairments and traffic requirements are approached by the implementation of configurable, modular software that is adapted to the specific conditions and needs of the particular wireless infrastructures.

INTRODUCTION

In recent years, Internet technology has emerged as the major driving force behind new developments in the area of telecommunications networks. The volume of data packet traffic has increased at extreme rates to the point that voice traffic, although a dominant source of revenue in the past, is likely to be supplanted by multimedia and other services supported by the World Wide Web.

Meanwhile, wireless networks face a similar trend of exponential traffic increase and growing importance to users. In some countries, such as Finland, Italy, and Spain, the number of mobile subscriptions has already exceeded the number of fixed lines.

The combination of both developments, the growth of the Internet and the success of wireless networks, suggests that the coming trend will be an increasing demand for wireless access to Internet applications.

The Wireless Internet Network (WINE) [1, 2] project was born to study this scenario, that is, to find transparent compensation methods for the poor performance of Internet wireless links in an indoor environment (either home or business) where mobility of users has to be considered. In such an environment, users with mobile Internet-based equipment (e.g., portable computers, mobile devices) require acceptable performance when communicating over the Internet, to perform their everyday work. It is highly desirable to design an Internet-oriented wireless communication system that behaves in the same way as a wire-dominated system to a great extent. Background work in the field proposes adaptive link layer protocols in terms of error control for Internet wireless links [3–6]. WINE augments existing approaches by introducing the wireless application layer (WAL), which is a generic framework combining adaptive local error control and traffic control.

This article is mainly devoted to introducing a likely WAL fourth-generation (4G) scenario focusing on the key aspect of providing quality of service (QoS) in underlying wireless networks carrying IP traffic. From a 4G perspective, for each connection a specific *QoS contract* will be established, defining the QoS requirements (in

The authors note that their names are appearing in alphabetical order.

terms of both maximum tolerated loss probability and maximum tolerated transmission delay) that have to be fulfilled by so-called *compliant traffic*. The statistical definition of compliant traffic is part of the QoS contract. In this respect, this article focuses on two key aspects:

- Link reliability, to guarantee that the maximum tolerated IP datagram loss probability established in QoS contracts is fulfilled
- Traffic control, to ensure the admission of compliant traffic in the underlying wireless networks, to guarantee that the maximum tolerated delays established in the QoS contracts are respected, and to maximize the exploitation of the scarce bandwidth

The article is organized as follows. In the next section, the WAL and its synergies with 4G are described. Following that, the modules related to link reliability and traffic control are described. Finally, some conclusions are drawn.

THE WAL: A STEP TOWARD 4G

Almost every decade a new mobile system is developed and becomes a commercial service. Thus, International Telecommunication Union — Radiocommunication Standardization Section (ITU-R) discussions on the 3G International Mobile Telecommunications 2000 (IMT-2000) system began more than 10 years before commercial services. Hence, the next system to provide present and completely new services may be expected by 2010. It is quite probable that 3G systems will be in wide use only around 2005, although deployments of networks have already begun in 2001.

The present-day 2G and emerging 3G mobile cellular networks are only some of the technologies moving toward mobile IP infrastructures. We consider 3G alone not to be enough for a ubiquitous multimedia-capable IP infrastructure. Moreover, we believe that 4G is not likely to be a single standardized air interface and networking infrastructure like 3G. Instead, 4G will be a system that will include several different networking technologies, and this heterogeneous architecture will interoperate through IP and the WAL in order to provide the best possible networking services wherever the user is located. Strictly speaking, it is too early to define 4G, since we have not even tested the boundaries of 2.5G, 3G, and wireless LANs (WLANs) yet. Nevertheless, it seems that some outlines of 4G are already emerging. Although we do not know what 4G is exactly, and there are competing definitions in the community, we try to list some possible trends. The future 4G networks are heterogeneous networks, which include a large number of different access networks. The terminals and base stations are using the software radio approach, although they are not yet full software radio. Moreover, it looks as if the one common factor is that 4G networking requires us to provide all-IP architecture and connectivity to anywhere at any time. There will also be, no doubt, new radio access methods and software radio breakthroughs in the future, so new transmission methods will also be incorporated into the heterogeneous 4G scenario. One issue we are interested in is how one can use IP to

“glue” together different radio networks to provide pervasive access to the Internet. Thus, the WAL is a step toward ubiquitous IP access over heterogeneous networks. It can be seen as one of the building blocks of the forthcoming 4G systems.

QOS AND TRANSPARENCY

One of the main challenges of the heterogeneous systems beyond 3G is efficient provision of an end-to-end QoS tailored to the specific application requirements. In general, a specific IP stream flows through more than one (wired or wireless) network. For each (wired or wireless) crossed network, a QoS contract is agreed, establishing the statistical characteristics of the traffic that has to be admitted in the considered network (e.g., in terms of average bit rate), as well as the QoS requirements characterizing such traffic (e.g., in terms of maximum delay tolerated by the IP packets and maximum loss probability of the IP packets) within the considered network. The fulfillment of the various QoS contracts in all the networks utilized by the connection entails satisfactory end-to-end QoS for the connection in question.

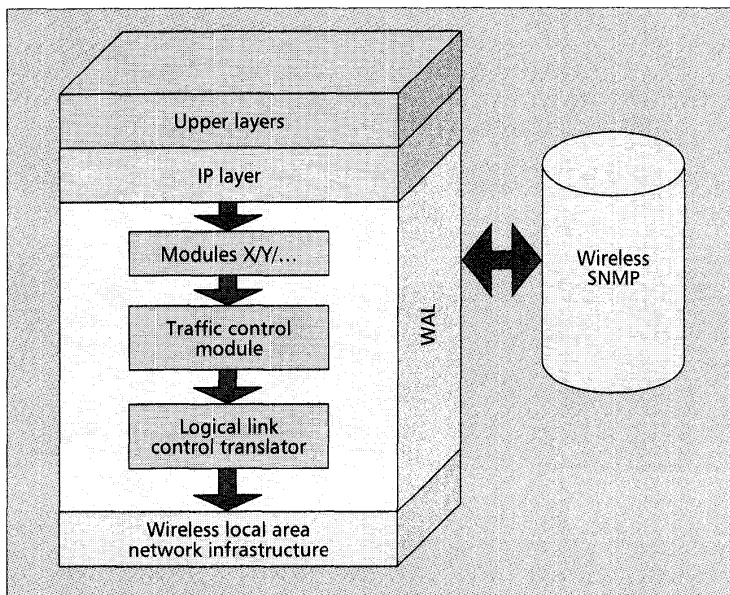
QoS contracts are already difficult to handle in wired networks, and even more challenging in wireless networks, in which this goal has to be achieved in conjunction with efficient exploitation of the scarce bandwidth. As a result, in wireless networks, traffic control strategies can be key factors for fulfilling QoS contracts and, at the same time, efficiently exploiting the available bandwidth.

In addition, two main problems arise when considering heterogeneous access for Internet applications. The first is that IP provides only best effort packet delivery service and may consequently be inadequate to enable the fulfillment of QoS contracts; in addition, it might make inefficient use of the available bandwidth. The second problem derives from the fact that the various wireless infrastructures, in general, adopt different mechanisms, not without remarkable deficiencies, in order to fulfill a QoS contract. As an example, Bluetooth offers a connection-oriented service and the possibility of negotiating QoS at connection setup on a per-flow basis. On the other hand, typical IEEE 802.11 hardware drivers lack support for a specified QoS differentiation at the medium access control (MAC) layer.

Another key factor related to QoS is the TCP/IP datagram loss probability in the wireless link. Therefore, link layer mechanisms that allow TCP/IP transmission over heterogeneous infrastructures, guaranteeing an agreed datagram loss probability, must be contemplated. These mechanisms should be adaptive so that any change in channel conditions will generate readjustment of the transfer parameters. The WAL supports automatic mechanisms for renegotiation of transfer parameters on channel state changes.

One last aspect to consider is that standardization is well established for the IP layer, and the wireless local area infrastructure is already widely deployed and fixed. Therefore, there is little flexibility for change in either the IP or wireless network infrastructure layer. Because of the wide acceptance and large number of installed devices, any major change is almost

QoS contracts are already difficult to handle in wired networks, and even more challenging in wireless networks, in which this goal has to be achieved in conjunction with efficient exploitation of the scarce bandwidth.



■ **Figure 1.** The WAL and its interaction with both the IP layer and wireless infrastructures.

impossible, or would take a very long time to implement on a large scale.

In order to overcome these problems, the WAL is to be added between the IP and wireless network infrastructure layers, maintaining transparency with respect to both the IP layer and the wireless network infrastructure layers; that is, its insertion between the IP and wireless infrastructure network layers does not modify either the usual IP or wireless network infrastructure protocols.

Maintaining this transparency with respect to the wireless infrastructures implies that the same WAL architecture must be able to work in conjunction with all “similar” wireless network infrastructures. Thus, in the WLAN case, the WAL has to provide the IP layer with a single interface independent of the particular WLAN standard. Thus, the same WAL architecture can be used in different access points (APs) or mobile terminals (MTs) in conjunction with different wireless network infrastructures. Figure 1 shows the WAL architecture and its interaction with both the wireless network infrastructures, through the logical link control translator (LLCT) module, and the IP protocol stack.

In the figure above, the set of modules X/Y/... corresponds to those related with transport/link layer control mechanisms. The traffic control module is responsible for guaranteeing the agreed traffic parameters.

WAL OPERATION OVERVIEW

To guarantee that the agreed QoS requirements are satisfied, the WAL selects a set of modules for packet processing before their transmission and on their reception. The set of modules and their parameters are adjusted to channel state conditions through link monitoring.

Each time an IP datagram is intercepted by the WAL (Fig. 2), it is classified by the WAL

coordinator module — the WAL coordinator is the “intelligence” of the WAL — which, in turn, calls modules sequentially. That is, the WAL coordinator accesses a lookup table to discover the class assigned to the packet’s flow, which indicates the appropriate module to be applied and its arrangement. Then it calls the first WAL module. The first module performs its processing task and returns control to the WAL coordinator. Next, the WAL coordinator calls the second WAL module, and so on. Finally, the WAL coordinator prepares the WAL header, which in turn may contain the modules’ sub-headers, and calls the LLCT module for further data processing. On reception, the QoS module drives the IP datagrams on the reverse chain of modules.

ASSOCIATION-ORIENTED SERVICE PROVISION

A novel feature of the WAL is its separation of traffic packets into classes and associations. Service provision in the WAL is, in fact, based on these two concepts.

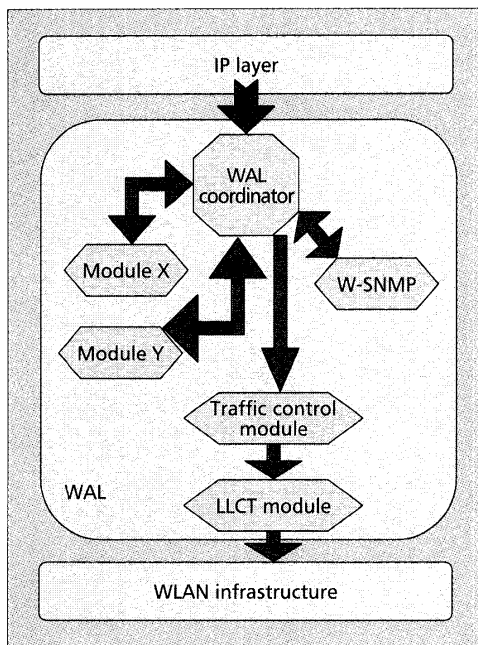
A WAL traffic class defines the type of traffic (audio, video, FTP/HTTP, etc.), and a list of functional modules (in the required order) which must be applied to that type of traffic. The connections belonging to the same WAL traffic class are characterized by the same QoS contract. The module list for every class is completely defined so that every WAL terminal uses the same module arrangement within the same class. This approach allows the WAL packet classification to be mapped onto existing Internet QoS models; specifically, differentiated services (DiffServ) are considered here, and the WAL packet class is determined by the DS (DiffServ) field and/or protocol type field of IP datagram headers.

An association identifies a particular WAL class in conjunction with a specified destination (i.e., $WAL_Association = WAL_Class + MT_Id$). As the condition of each wireless channel between the AP and an MT can vary, the parameters of the modules defined for the corresponding class have to be adjusted to adapt to these changes.

THE WAL MODULES

As described in previous sections, there are two main responsibilities of the WAL: first, to compensate for the impairments of the wireless medium and enhance throughput by incorporating transport/link layer control mechanisms; and second, to control the traffic characteristics according to a set of specified parameters in order to cope with congestion states and achieve a fair allocation of link resources. In this section the main modules dealing with these aspects are described.

The WAL approach is flexible enough to contemplate the use of an undetermined set of modules for a WAL class which can be particularized to a specific wireless platform so that it can always bring enhancements and avoid unnecessary functional duplications. Examples of WAL modules can be forward error correction (FEC), automatic repeat request (ARQ), segmentation and reassembly (SAR), traffic control, and header compressor.



■ Figure 2. The detailed WAL architecture.

In this section a basic choice of modules that cope with link reliability, throughput enhancement, and traffic control aspects is described.

ENHANCING THROUGHPUT AND LINK RELIABILITY

Most current network applications that require reliable transmission use TCP. TCP achieves reliability by requiring that the TCP sender retransmit lost packets. For this purpose, the TCP receiver sends back an acknowledgment (ACK) for every data packet it receives correctly. ACKs are cumulative; that is, an ACK carrying sequence number m acknowledges all data packets up to and including the data packet with sequence number $m - 1$. A new packet received by the receiver is said to be out of order if it is not the next in-sequence packet expected by the receiver. On receiving this out-of-order packet, the receiver sends a duplicate acknowledgment (dupack), acknowledging all the in-sequence bytes received so far. The TCP sender determines that a packet is lost using one of the following two mechanisms:

- Retransmission timeout, in which the packet is presumed lost if a timer, set when a packet is transmitted, expires before acknowledgment of the packet is received.
- Fast retransmit: If three dupacks containing the same sequence number are received, the packet which includes that sequence number is presumed lost, and retransmitted.

The TCP sender maintains a congestion window that determines the maximum amount of unacknowledged data sent by the sender. When a packet loss is detected, the TCP congestion control mechanism drastically reduces the congestion window size, effectively reducing the amount of data sent by the sender in one round-trip time (RTT). Thus, due to the TCP assump-

tion that the primary cause of packet loss is congestion, TCP can perform poorly over wireless links. Since the main reason for packet losses over wireless links can be due to transmission errors, the TCP sender, on detecting such a packet loss, reduces its congestion window, unnecessarily degrading throughput. Therefore, it is desirable to achieve the goal of improving TCP performance in the wireless hop without modifying the existing implementations elsewhere in the Internet. The Snoop protocol [7] is a good candidate to cope with these two requirements. It should be noted that the WAL is not a Snoop-specific solution; in fact, any other solution can (and will) be included in the module library, and the best solution for the specific problem selected from among these. We use Snoop here as an example.

For data transfer to a mobile host, the protocol improves TCP performance by deploying an agent at the base station called a *snoop agent*. This agent allows fast local recovery of lost segments by using the information conveyed in TCP ACKs sent by the mobile host. The only modifications needed at the base station are caching of unacknowledged TCP segments and performing local retransmissions based on policies that depend on TCP ACKs (from the mobile host) and timeouts. The module shields the sender from the vagaries of the wireless link by using dupacks to identify packet loss and performing local retransmission as soon as this loss is detected.

For data transfer from a mobile host, the snoop agent detects corrupted segments in the wireless link at the base station by snooping on packets arriving from the mobile host and identifying gaps in the transmission sequence. Then, when a packet is dropped on the wireless link, future cumulative ACKs corresponding to the lost packet are marked (at the base station) with explicit loss notification (ELN) information, to identify that a non-congestion-related loss has occurred. Upon receiving this information with dupacks, the sender uses it to identify that the loss was not caused by congestion and may perform retransmissions without invoking the associated congestion control procedures. ELN support implies a slight modification of the TCP code in the mobile host.

Wireless infrastructures, in particular the WLAN and wireless personal area network (WPAN), implement various error control mechanisms. As an example, IEEE 802.11b implements only error detection techniques, while HIPERLAN-2 incorporates powerful error detection and correction schemes. To cope with these differences in implementations and guarantee an agreed QoS in the wireless link, whatever the infrastructure is, an FEC module becomes essential in order to support not only real-time applications (carried on UDP) but also non-real-time streams (carried on TCP).

Many linear codes can be found in the literature implementing an FEC strategy. The majority of them could be good candidates to be included in the corresponding WAL module. The use of one or another, as well as the parameters of the selected code, can be negotiated according to channel conditions by using appropriate WAL signaling exchange.

It is desirable to achieve the goal of improving TCP performance in the wireless hop without modifying the existing implementations elsewhere in the Internet. The Snoop protocol is a good candidate to cope with these two requirements.

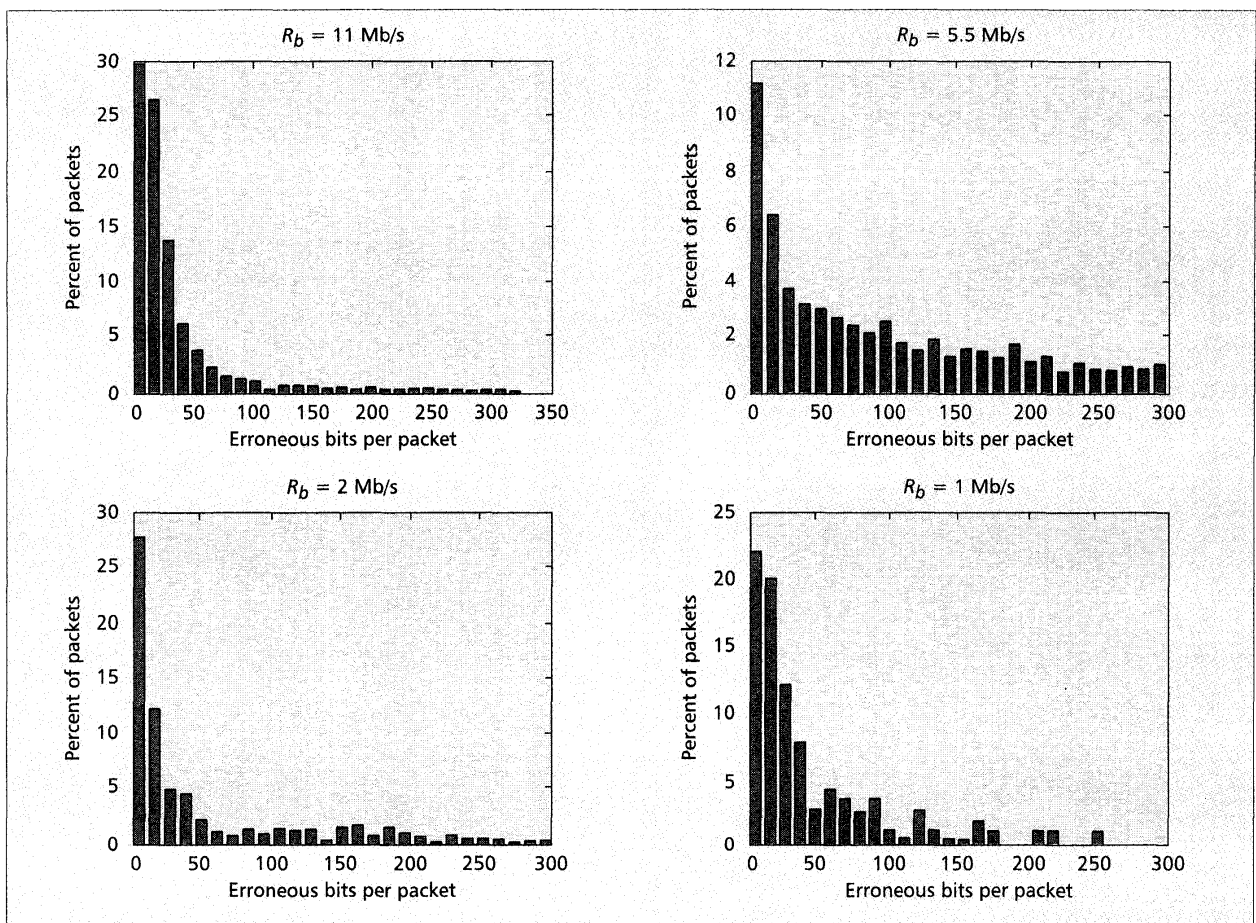


Figure 3. Error distribution per packet measured in IEEE 802.11b.

As an example, Fig. 3 shows the error distribution per packet for a maximum transmission unit (MTU) of 1500 bytes. These results are obtained from extensive measurements done in the WINE test networks using IEEE 802.11b operating at different bit rates. There are two main reasons for the differences in the error distribution functions for the corresponding bit rates. First, for each corresponding bit rate, there is a different relationship between bit duration (packet duration) and average fade duration of the radio channel. Second, the infrastructure uses different modulation schemes at different bit rates, which have different sensitivities.

As can also be seen in the figure, for $R_b = 11$ Mb/s, approximately 30 percent of the packets exhibit a number of erroneous bits that never exceeds 10 b/packet. Once an error occurs in the packet, the average error burst length associated with that error is approximately 1.5 bits. The use of a Reed-Solomon code, with capability to correct up to eight erroneous symbols, enables the correction of more than 25 percent of the erroneous packets with the consequent TCP increase of throughput or subjective quality. The reason to not use a larger error correction capability is to avoid increasing the computational overhead of the decoding algorithm, which would limit the final performances. In this sense, it is important to remark that the overhead, mainly introduced by

the decoding algorithm, suggests the use of FEC only in those situations in which packet error rate (PER) is high enough. Thus, the throughput gain due to packet error recovery is more significant than the degradation derived from the computational overhead introduced by the decoder.

TRAFFIC CONTROL

A traffic control module aims to regulate the admission of traffic compliant with the agreed QoS contract in the underlying wireless network. In addition, the traffic control module aims to maximize exploitation of the scarce bandwidth and, at the same time, to guarantee that the IP datagrams admitted in the underlying wireless networks are carried from the AP to the MT and vice versa within the maximum tolerated delay agreed in the QoS contract.

The structure of the traffic control module running in an AP is shown in Fig. 4. The figure shows that the traffic control module consists of the following building blocks:

- A *classifier*, in charge of sorting the IP datagrams arriving at the traffic control module according to the associations they belong to; this sorting is carried out by reading the headers of the IP datagrams.
- A set of *MT schedulers*. Each MT scheduler handles queues storing datagrams destined to the same MT. The role of an MT sched-

uler is that whenever it is selected from among the other MT schedulers to pass an IP datagram to the LLCT, it chooses the corresponding IP datagram among those stored in its queues. MT schedulers perform traffic regulation allocating a minimum bandwidth for each association, smoothing peaks of traffic, and controlling packet delay. Traffic overflowing out of regulators can be either dropped or delayed.

- A *main scheduler algorithm* whose role is, whenever an IP datagram has to be passed to the LLCT, to select the MT scheduler that must transmit the datagram. The main scheduler makes the choice of the MT scheduler following a general principle of fairness in allocating resources among the flows destined to various MTs and contending for the wireless medium. In addition, it gives priority to packets destined to MTs with the best channel quality, temporarily excluding disturbed channels from transmission.

It is important to note that the traffic control implemented in the MT will typically lack the main scheduler algorithm and will include a single scheduler with the same architecture as the AP's MT scheduler. This is because MTs in the same cell are not supposed to exchange packets directly, but will preferably use the serving AP as an intermediate hop for communication.

In the following, the basic principles of operation of the traffic control module are described, considering the main scheduler and the MT schedulers separately.

The main scheduler selects MT schedulers with the aim of achieving fairness among the various MTs having connections in progress with the AP, while considering channel state conditions in order to increase the overall wireless medium throughput. More specifically, a channel from the AP to an MT is assigned a portion of the total bandwidth available in the wireless medium. This is, in turn, used to calculate a target quantity of bytes that each channel is allowed to use over a certain interval of time.

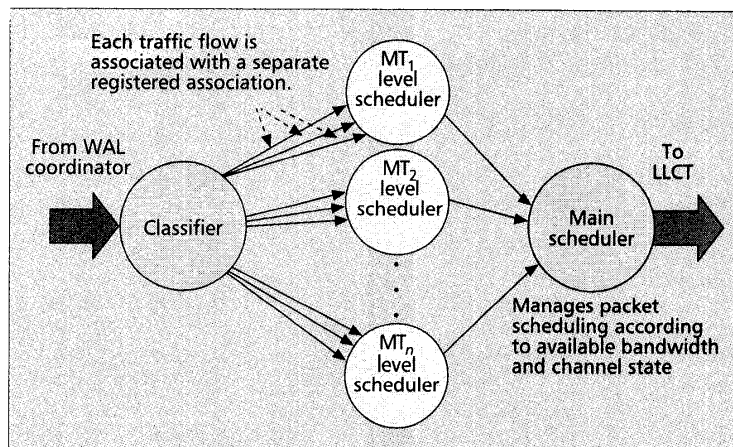
The effective utilization of a channel can differ from its target value for many reasons. When performing its task, the main scheduler has to take into account the following issues:

- IP datagrams are variable-length.
- Some selected MT schedulers can have their buffers empty due to statistical oscillations of traffic.
- At certain times, some channels toward MTs can be disturbed.

In this last respect, the main scheduler algorithm does not select the MT schedulers corresponding to the MTs whose channels are disturbed. The information concerning the quality of the various links is taken from the underlying wireless networks by means of a wireless Simple Network Management Protocol (SNMP) agent [8].

It is the task of the main scheduler to try to minimize the differences between target values and effective values, dynamically selecting MT schedulers among those currently associated with "good channels."

Each MT scheduler has the internal structure shown in Fig. 5. The figure shows that an



■ Figure 4. The traffic control module architecture.

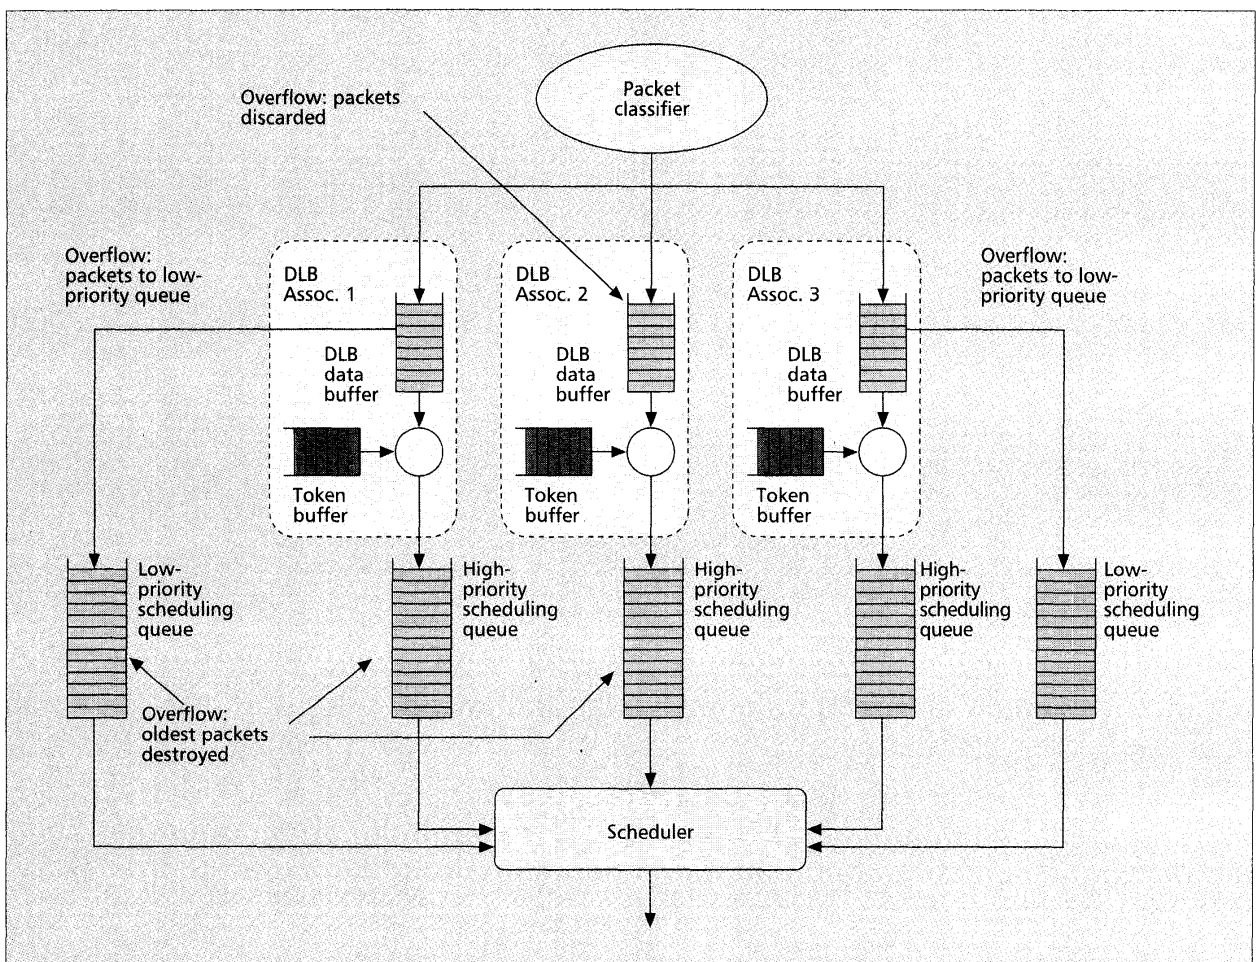
MT scheduler consists of a set of dual leaky buckets (DLBs), each regulating the traffic corresponding to an association. After being classified, IP datagrams arrive at the DLB data buffer of the corresponding association. The DLB data buffer is a FIFO buffer in which IP datagrams can wait before they are let out (*admitted*) according to standard DLB operation [9]. It is important to note that IP datagrams can be stored in the DLB data buffer only if this buffer is not full; otherwise, such IP datagrams overflow (Fig. 5).

DLBs perform policing, traffic shaping, and congestion control functions. In particular, the DLB control parameters are set consistently with the QoS contracts so that the traffic admitted by the DLBs is only the compliant traffic. The compliant traffic admitted by the DLBs is stored in the *scheduler queues*, whereas the traffic that overflows from the DLB data buffer (i.e., the noncompliant traffic) is not necessarily discarded. As a matter of fact, one of the goals of the traffic control module is to pass as much non-compliant traffic as possible to the LLCT, provided the QoS requirements of the compliant traffic are respected.

Whenever (following the decisions of the main scheduler) an MT scheduler has to pass an IP datagram toward the LLCT, the *MT scheduling algorithm* is run. This algorithm decides which IP datagram among those presently stored in the various buffers of the MT scheduler has to be passed to the LLCT. This algorithm acts by distinguishing high-priority from low-priority queues. An IP datagram is selected from a low-priority queue only if all the high-priority queues are empty. Moreover, in both sets of high-priority and low-priority queues, IP datagrams are selected according to an earliest deadline first (EDF) discipline [10].

CONCLUSIONS

One of the defining factors of fourth-generation systems will be the cooperation of different access technologies such as cellular, cordless, and WLAN type systems. Different access networks will be combined in a common platform to complement each other optimally, both from



■ Figure 5. The architecture of the MT scheduler.

the user's perspective and for different service requirements as well as radio environments. In this respect, in 4G systems each connection is expected to establish a particular QoS contract with the network stating the characteristics of compliant traffic, as well as its QoS requirements; the latter are expressed in terms of maximum tolerated packet loss probability and maximum tolerated transmission delay.

Therefore, platform independence and QoS provision are key factors in 4G systems. The WAL approach addresses both questions specifically and proposes configurable, modular software as a solution for fast deployment of a 4G infrastructure.

The WAL consists of a set of modules whose basic aim is to assist the underlying wireless network in *efficiently* fulfilling QoS contracts. Thus, the WAL modules aim to keep both the loss probability and the delays experienced by the IP traffic below maximum tolerated values (established in the QoS contract). At the same time, the WAL modules cooperate to maximize the exploitation of the available bandwidth.

The WAL structure and organization are flexible enough to enable both the addition of new modules and the upgrade of the algorithms on which the various modules are based. It can

adapt itself to a variety of wireless platforms, offering a wide range of classes which can be customized for specific underlying networks.

ACKNOWLEDGMENTS

This work has been performed in the framework of the project WINE IST-1999-10028, which is funded by the European Community. The authors would like to acknowledge the contributions of their colleagues from VTT, Philips, Cefriel, University of Rome "La Sapienza," AQL, Acorde, University of Cantabria, Intracom, University of Athens, and The Queen's University of Belfast. Petri Mähönen and Luis Muñoz would also like to acknowledge the partial financial support of the Academy of Finland (Grant 50624) and the Spanish CICYT (project 1FD97-0960-C05-03), respectively.

REFERENCES

- [1] Z. D. Shelby et al., "Wireless IPv6 Networking — WINE," *IST Mobile Commun. Summit 2000*, Galway, Ireland, Oct. 2000, pp. 575–80.
- [2] WINE Project: <http://www.vtt.fi/ele/projects/wine>
- [3] D. Eckhart and P. Steenkiste, "Improving Wireless LAN Performance Via Adaptive Local Error Control," *Proc. ICNP '98*, 1998, pp. 327–38.
- [4] R. Ludwig, "A Case for Flow-Adaptive Wireless Links," Tech. rep. UCB/CSD-99-1053, UC Berkeley, May 1999.

- [5] G. C. Polyzos and G. Xylomenos, "Enhancing Wireless Internet Links for Multimedia Services," *Proc. 5th Int'l. Wksp. Mobile Multimedia Commun.* '98, Berlin, Germany, Oct. 1998, pp. 379–84.
- [6] R. Ludwig, A. Konrad, and D. Joseph, "Optimizing the End-to-End Performance of Reliable Flows over Wireless Links," *Proc. ACM/IEEE MobiCom*, 1999, p. 113–19.
- [7] H. Balakrishnan, S. Seshan, and R. Katz, "Improving Reliable Transport and Handoff Performance in Cellular Wireless Networks," *Wireless Networks*, 1995, pp. 469–81.
- [8] P. Bhagwat et al., "Using Channel State Dependent Packet Scheduling to Improve TCP Throughput over Wireless LANs," *ACM/Baltzer Wireless Networks J.*, vol. 3, no. 1, Jan. 1997, pp. 91–102.
- [9] L. Georgiadis et al., "Efficient Network QoS Provisioning Based on Per Node Traffic Shaping," *IEEE/ACM Trans. Net.*, vol. 4, no. 4, Aug. 1996, pp. 482–501.
- [10] F. Delli Priscoli, "Design and Implementation of a Simple and Efficient Medium Access Control for High Speed Wireless Local Area Networks," *JSAC*, vol. 17, no. 11, Nov. 1999, pp. 2052–64.

BIOGRAPHIES

LUCA BECCHETTI (becchett@dis.uniroma1.it) graduated in electronic engineering at the University of Rome "La Sapienza" in 1995. He received a Ph.D. in computer science from the University of Rome "La Sapienza" in 1998. He is currently assistant professor at the Computer Science Department of the University of Rome "La Sapienza." He studies theoretical aspects of efficient network management and scheduling. In this framework, he is involved in the European Union IST project ALCOM-FT. He is also involved in two European Union IST projects dealing with resource management for broadband satellite access networks and wireless systems.

FRANCESCO DELLI PRISCOLI (dellipriscoli@dis.uniroma1.it) graduated in electronic engineering from the University of Rome "La Sapienza" in 1986. He received a Ph.D. in system engineering from the University of Rome "La Sapienza" in 1991. From 1986 to 1991 he worked in the Studies and Experimentation Department of Telespazio (Rome). Since

1991 he has been working at the University of Rome "La Sapienza" where, at present, he is associate professor and gives the course "Automatic Controls." In the framework of his activity, he has researched in nonlinear control theory, in the area of access techniques and mobility procedures for third- and fourth-generation mobile systems.

TIZIANO INZERILLI (inzerilli@tiscalinet.it) graduated in electronic engineering at the University of Rome "La Sapienza" in 2000, defending a dissertation on the design of a micromobility protocol for the Internet (SIMPLE). He is currently working toward a Ph.D. in computer engineering and carrying out research on mobility and traffic control in the Internet while participating in the IST projects WINE and Broadband Access for High-Speed Multimedia via Satellite (BRAHMS).

PETRI MÄHÖNEN, Ph.D. (Petri.Mahonen@oulu.fi) [M'96, SM '01] was a research professor at the VTT, Technical Research Center of Finland. In March 2001 he moved to the University of Oulu, Telecommunications Laboratory. Before this he studied and worked at the Academy of Finland, University of Helsinki, and Oxford University. At VTT he was a founding scientific director of the Wireless Internet Lab. He has been principal investigator in different national and international projects, including European Union research projects such as WINE and WIND-FLEX. His current research focuses on wireless Internet, low-power medium access control, ad hoc networks, broadband wireless access, and personal area networking.

LUIS MUÑOZ (luis@tmat.unican.es) is associate professor at the University of Cantabria. He received a degree in telecommunications engineering and a Ph.D. degree from the Technical University of Cataluña (UPC), Spain, in 1990 and 1995, respectively. He joined the Department of Communications Engineering in 1990 where he has been involved in different national and international projects, the latter within the ACTS and IST programs, related to voice and data transmission over wired and wireless infrastructures. His current research interests lie in the design and implementation of link layer control mechanisms for wireless Internet infrastructures, channel coding, and medium access protocols.

Platform-independence and QoS provision are key factors in 4G systems. The WAL approach addresses both questions specifically and proposes configurable, modular software as a solution for fast deployment of a 4G infrastructure.