♦ Supporting QoS in Broadband Wireless and Wired Access

Bastien Peelen, Miroslav Zivkovic, Dennis Bijwaard, and Harold Teunissen

Wireless local area network (WLAN), cable, and digital subscriber line (xDSL) are among the most popular broadband access technologies in use today. In all such technologies, the transport capacity provided at the level of the physical medium is non-deterministically shared by different traffic streams generated by a multitude of applications. These traffic streams may have different or even incompatible characteristics (e.g., one may contain bursty best-effort traffic generated by file transfer, another streaming quality of service [QoS] traffic generated by video-on-demand), and they may interfere with one another. To accommodate admitted QoS traffic in a fluctuating available bandwidth and to protect it from high-load statistical traffic patterns, traffic must be regulated. This paper describes a bandwidth-distribution mechanism for broadband access technologies that uses real-time characteristics of both the active-medium-sensing and the feed-forward control mechanisms. To validate this mechanism, two prototypes are developed, one based on wireless, the other on wired shared media. These prototypes employ legacy network elements without intrinsic QoS capabilities. Finally, we present the results of tests run on these prototypes and draw conclusions from our work. © 2003 Lucent Technologies Inc.

Introduction

Network quality of service (QoS) can be thought of as the process of classifying packets for the purpose of treating some packet flows differently than other packet flows. This classification process involves all the network layers in every network element in the communication path, as illustrated in **Figure 1**. End-to-end QoS is determined by the weakest link among all network elements between the sender and the receiver.

Some QoS mechanisms conceived in the past addressed end-to-end QoS, but only in recent years has the QoS community come to understand that the problem of providing end-to-end QoS must be solved by dividing the problem along network domain boundaries [8], as illustrated in Figure 1. Within these network domains, there are various ways to provide QoS. One approach—known as over-engineering or over-provisioning the network—is simply to provide enough bandwidth. Another approach is a simple form of prioritization, without admission control. To avoid overload and congestion, both approaches require that traffic remain within the bounds of statistical traffic patterns for each priority class. However, when there is sharing of physical and access media

Bell Labs Technical Journal 8(2), 65–81 (2003) © 2003 Lucent Technologies Inc. Published by Wiley Periodicals, Inc. Published online in Wiley InterScience (www.interscience.wiley.com). • DOI: 10.1002/bltj.10064

and fluctuating available bandwidth, both these approaches are unsatisfactory, as will be shown in the course of this paper.

Our approach addresses situations in which:

- Over-provisioning is not an option,
- It is impossible or undesirable to use the layer 2 prioritization mechanisms offered by shared-medium access technologies,
- The level of guarantee offered to QoS traffic by distributed prioritization mechanisms is not high enough, and
- The stochastic nature of the traffic is causing overload in both the access network and the queues of the priority classes of the layer 2 prioritization mechanisms.

The remainder of this paper will discuss our approach to constructing two prototypes that make possible the admission of QoS traffic sessions to broadband access networks based on shared media. The requirements for this approach are based on the idea that the solutions it provides should:

- Allow QoS-enabled services to reserve bandwidth in the access network,
- Protect higher-priority (i.e., QoS) traffic from lower-priority (i.e., best-effort) traffic,
- Provide safeguards against fluctuating medium bandwidth for higher-priority (i.e., QoS) traffic,
- Support legacy hardware,
- Allow the use of complementary prioritization QoS mechanisms, and
- Incorporate standards into the solution, where appropriate.

Panel 1. Abbreviations, Acronyms, and Terms

ADSL—asymmetric digital subscriber line
AP—access point
API—application programming interface
ATU-R—ADSL termination units for the remote site
BER—bit error rate
BRAS—broadband remote access server
CoS—class of service
DiffServ—differentiated services
DSL—digital subscriber line
DSLAM—digital subscriber line access multiplexer
FIR—finite impulse response
FTP—file transfer protocol
GUI—graphical user interface
IEEE—Institute of Electrical and Electronics Engineers
IETF—Internet Engineering Task Force
IP—Internet protocol
ISDN—integrated services digital network
LAN—local area network
MAC—medium access control
MPLS—multiprotocol label switching
OS—operating system
PC—personal computer
PPP—point-to-point protocol
QoS—quality of service
RADSL—rate-adaptive DSL
RSVP—resource reservation protocol
RSVP-TE—resource reservation protocol traffic engineering
SBM—subnet bandwidth manager
SNR—signal-to-noise ratio
SOHO—small office home office
TCP—transmission control protocol
UDP—user datagram protocol
WAN—wide area network
WLAN—wireless LAN
xDSL—any of various digital subscriber line technologies

Overview of Different QoS Models

Network QoS mechanisms adhere to one of two models: the reservation model or the prioritization model. The following two sections offer brief descriptions of these models.

Reservation Model QoS

One way to divide scarce resources is to allow the parties that use them to reserve them. The two main network resources that are available for QoS reservation are bandwidth and low-latency data paths. The Integrated Services Working Group [2] of the Internet Engineering Task Force (IETF) has focused on mechanisms for reserving these resources. Some protocols that adhere to the reservation model are: resource reservation protocol (RSVP) [3], multiprotocol label switching (MPLS) [10], and subnet bandwidth manager (SBM) [14].

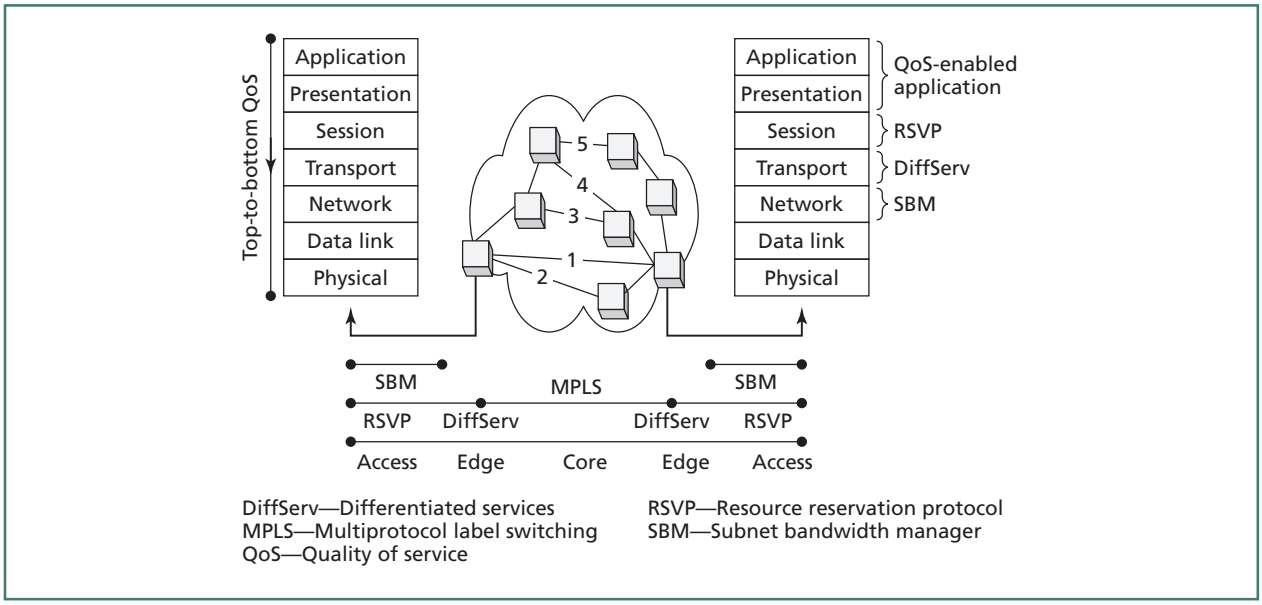

Figure 1.
End-to-end QoS across access, edge, and core domains.

Prioritization Model QoS

Data prioritization is another way of providing QoS [1]. It is complementary to bandwidth reservation in most network contexts. This type of QoS—sometimes referred to as class of service (CoS) QoS—provides QoS by treating higher-priority packets better than lower-priority ones. Because it handles aggregated rather than separate flows, this type of QoS enables looser prioritization of traffic than reservation-based QoS. Some protocols that adhere to the prioritization model are: differentiated services (DiffServ) [1], Institute of Electrical and Electronic Engineers (IEEE) 802.1D annex H [5], and IEEE 802.11e [7].

There are two different types of prioritization: centrally controlled and stochastically distributed. With centrally controlled prioritization, higher-priority traffic can be allowed to use more timeslots and more bandwidth per end node than lower-priority traffic. With stochastically distributed prioritization, higher-priority traffic has a greater chance of being sent than lower-priority traffic; however, because stochastically distributed prioritization is stochastic, there is no guarantee that higher-priority traffic will always get through first, especially when there is a lot of lower-priority traffic. In both types of prioritization, best-effort traffic usually gets the lowest priority, and there is generally no limit to the amount of best-effort traffic that applications can try to send (i.e., there is no shaping). Furthermore, the relative weights given to higher-priority traffic, optimization of overall throughput, and fairness are generally the same in all end nodes. There are situations (e.g., one endpoint uses high-priority traffic, and 100 others use best-effort traffic) in which such weighting may not be ideal. For prioritization to work under such conditions, each node would have to know how much traffic other nodes are sending.

Synergy between the Reservation and Prioritization Models

Reservation and prioritization QoS mechanisms have primarily been deployed in disparate network domains, as illustrated in Figure 1. For example, traffic from access networks is aggregated at the edges, which makes the edge more suitable for prioritization mechanisms (e.g., DiffServ). (Typically, higher priority classes are assigned to packets traveling on low-latency and low-jitter data paths.) The aggregated traffic streams then travel through the core network using reservation mechanisms (e.g., MPLS in conjunction with RSVP traffic engineering [RSVP-TE]). Neither the statistical aggregation principle employed in the edge, nor the

semi-static reservations of aggregate tunnels employed in the core are of use in the access domain. In fact, the lower order of granularity of the traffic load in the access domain means that any prioritization mechanisms in this domain must rely heavily on reservation mechanisms; otherwise, the scarce resources in each priority class will be depleted in high-load situations by the overloading of their queues. The rest of this paper discusses the use of reservation model QoS mechanisms in broadband shared-media access networks.

QoS in Broadband Shared-Media Access Networks

Before we describe our approach, we will discuss the most significant characteristics of the shared media we use as broadband access networks in our prototypes.

Popular broadband access technologies (e.g., wireless local area network [WLAN], cable, and digital subscriber line [xDSL]) have in common the nondeterministic nature of their traffic. They also have in common a shared medium in which bandwidth is shared by a multitude of applications that generate traffic streams with different and sometimes incompatible characteristics. For example, one stream may contain bursty best-effort traffic generated by file transfer, another streaming QoS traffic generated by video-on-demand. Finally, in each of these technologies, changing characteristics of the physical medium and high traffic load cause fluctuations in the available bandwidth per application over time.

Reservation and prioritization model QoS protocols have been designed to deliver data with certain QoS characteristics in a high-capacity, packet-switched, wireline environment. Employing these protocols on WLAN, cable, and xDSL networks is difficult, because of the physical characteristics of these networks and their use of a shared medium. (Although it is common knowledge that cable and WLAN both use a shared physical medium, this is not as obvious in the case of broadband xDSL, particularly because different xDSL endpoints connected to one digital subscriber line (DSL) access multiplexer (DSLAM) do not share wires. However, such endpoints do share available bandwidth in the physical medium, because of crosstalk and electrical interference.) Nevertheless, because these networks use shared media, appropriate QoS provisions are essential; otherwise, applications that generate high loads of best-effort traffic will consume too much bandwidth and will degrade the service levels of QoS-sensitive applications.

Characteristics of Shared Media

In shared media, numerous terminals and applications (the nodes in **Figure 2**) compete for shared bandwidth. In circuit-switched shared media, bandwidth distribution is centrally controlled, but in packet-based shared media most elements of the network infrastructure do not contain the functionality required to collaborate in the end-to-end QoS that would be necessary to control bandwidth distribution.

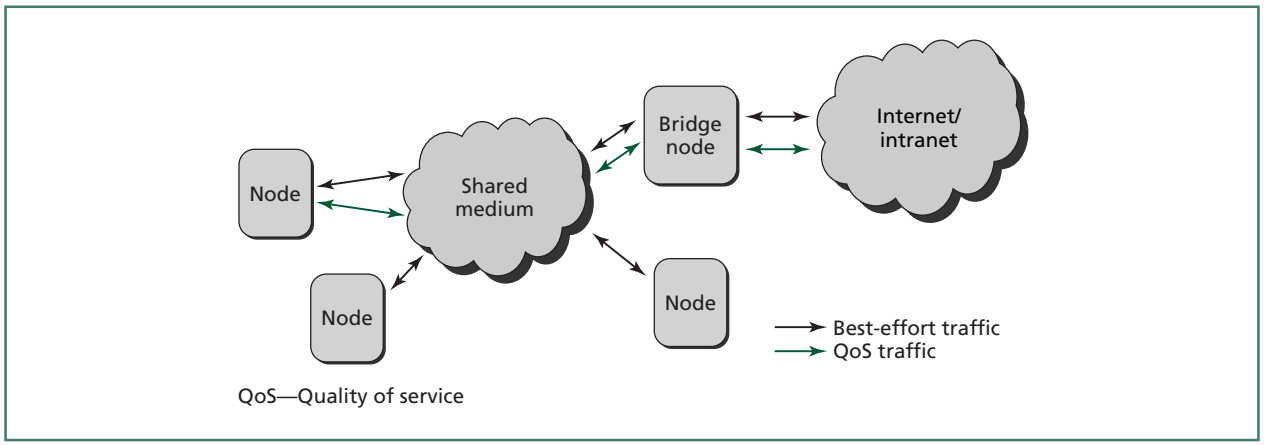

Figure 2.
Traffic of multiple nodes in a shared medium access network.

Unlike the case in packet-switched networks, in which data in excess of a certain threshold can simply be dropped at a switch or router, all data sent on shared media use part of the available bandwidth. Because there are no commonly available bandwidth-distribution mechanisms for shared media, it is currently impossible to offer a high guarantee for QoS traffic or to divide the available bandwidth between all nodes accessing the media fairly.

In shared-media access networks, a bridge node connects the medium to an external network, such as the Internet or an intranet. In most situations, this bridge node uses the greater part of the available bandwidth, because most network traffic flows between the access network and the edge network (i.e., the wide area network [WAN]). In a QoS-enabled shared medium, this node also functions as a bridge between the QoS mechanisms of its shared medium and the QoS mechanisms of the edge network to which it provides connectivity.

Fluctuating Available Bandwidth

Another important characteristic of shared media is fluctuating available bandwidth, which is illustrated in **Figure 3**. Every access network has a theoretically defined amount of available bandwidth (1). Unfortunately, in practice, the bandwidth that is actually available is less than this amount. The available bandwidth (3) is the difference between the theoretical bandwidth (1) and the unavailable bandwidth (2). The available bandwidth fluctuates because of:

- The characteristics of the physical medium
  - Electromagnetic interference (e.g., crosstalk, signal-to-noise ratio [SNR], and shared frequency bands)
  - Atmospheric influences (e.g., rain, atmospheric humidity, and lightning)
- The characteristics of the protocol mechanism of shared media
  - Collisions, back-offs, overhead, and slow starts
  - Dynamic channel selection and rate adaptation

Because the unavailable bandwidth is inherently random, the available bandwidth is also random (i.e., it fluctuates unpredictably). This can be disastrous for QoS-sensitive applications. The bandwidth used by QoS traffic (5) is an adaptive and configurable threshold in our QoS-enabled admission-control mechanisms. To maximize the probability that QoS bandwidth will be available, the maximum amount of QoS traffic to be admitted should be configured so as to leave a safety margin, as a protection against the minima of the fluctuating available bandwidth. The bandwidth that can be allocated to best-effort traffic (4) is the difference between the fluctuating available bandwidth (3) and the bandwidth used by QoS traffic (5).

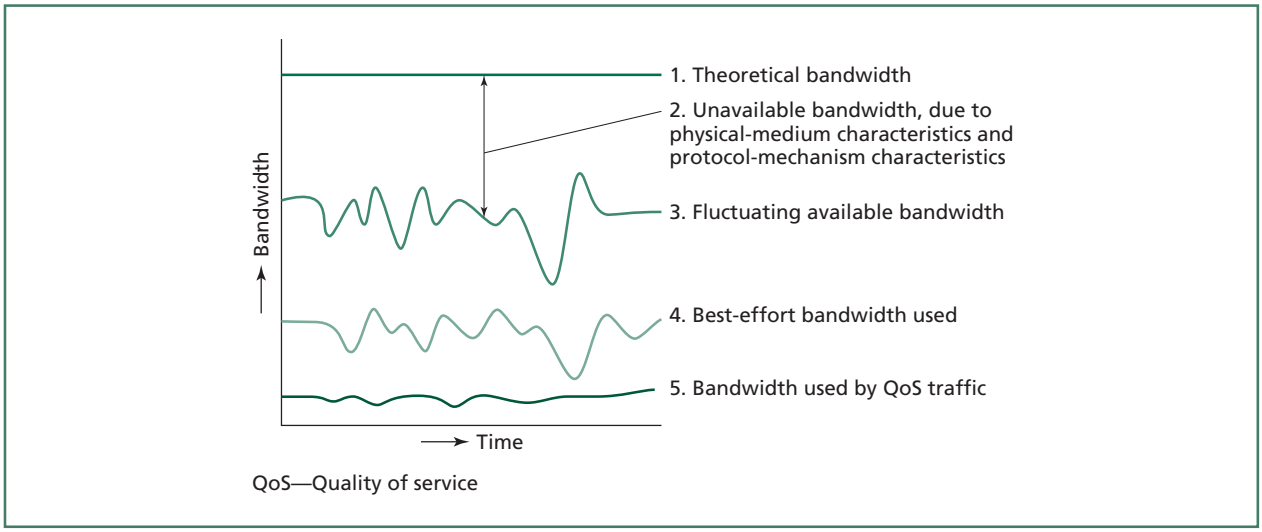

Figure 3.
Fluctuating available bandwidth at medium level.

The nature and the range of bandwidth fluctuation are determined primarily by the access network technology used. The following sections discuss the most significant factors contributing to bandwidth fluctuation in the WLAN and xDSL access technologies.

Bandwidth fluctuation in WLAN access technology. A wireless medium has a much higher bit error rate (BER) than a wired medium of comparable bandwidth. The higher BER is caused by the characteristics of the wireless medium (e.g., path loss, shadowing, multipath fading, interference, and hidden terminals).

The fluctuation of the available bandwidth for a terminal in a WLAN is related to the SNR measured at the terminal. When the SNR drops, the BER will probably increase, and the number of packet retransmissions will also increase, effectively reducing the available bandwidth for all connected terminals. Because the SNR is random, the fluctuations of the available bandwidth will also be random, as illustrated in Figure 3 (3).

A back-off mechanism is used within a WLAN to avoid collisions as much as possible, because each collision that occurs wastes bandwidth (2). There are ways to provide central control over the timeslots in which remote nodes are allowed to send data, but for various reasons this functionality is not widely implemented in WLAN cards.

WLAN networks, with the exception of ad-hoc networks, have a relay node known as an access point (AP). All traffic sent by the nodes in a WLAN network, including inter-node traffic, must first be sent to the AP. The AP then forwards the traffic to its destination. Because the relay in a WLAN is done at the link layer, it cannot be controlled without changing the network devices, which would be technology-intrusive. Also, because of the nature of the relay mechanism, inter-node packets occupy the medium twice, which must be taken into account in bandwidth calculations.

IEEE 802.11e proposes a prioritization model QoS mechanism for a WLAN [7]. However, this technology is not yet widely available on the market. Assigning different priority classes to different kinds of traffic would solve part of the QoS problem, but QoS still cannot be guaranteed when there is excessive traffic in one of the priority classes.

Bandwidth fluctuation in xDSL access technology. DSL technology was introduced to make it possible to use legacy telephone cables for higher-capacity data transfer than was previously possible with other modulation standards (e.g., integrated services digital network [ISDN]). Still, the electromagnetic characteristics of this medium cause high and highly varying error rates. To enhance the performance and reliability of DSL, rate adaptation has been incorporated into it. Rate-adaptive DSL (RADSL) is a version of DSL that is able to determine—in real time—the available line rate (i.e., the available bandwidth) of the wires between a DSL modem and a DSLAM. This rate-adaptive mechanism can be enabled either when the connection is set up, or at any time during the lifetime of the connection.

Available bandwidth in xDSL access networks depends on several factors:

- The distance between the subscriber and the central office;
- The electromagnetic interference, both external (i.e., interference) and internal (i.e., crosstalk); and
- The physical characteristics of the copper line wire:
  - The diameter of the signal-carrier wires,
  - The quality of the material, and
  - The shielding.

Recent implementations of DSL incorporate the rate-adaptive mechanism. In fact, in the majority of the xDSL installed base, rate-adaptation functionality is already present in the equipment. Although rate adaptation improves the reliability of xDSL in terms of error rate and packet loss, it does not solve the QoS problem.

An Admission Control Mechanism for QoS Traffic

Our approach to supporting QoS in broadband wireless and wired access has been to create an access control mechanism that provides an admission control function for QoS traffic. However, the observations we have made concerning the characteristics of shared media lead to the conclusion that there are situations in which it is not possible to guarantee QoS effectively without an access control mechanism for lower-priority (i.e., best-effort) traffic too. In other

words, the total traffic output of each node, which consists of both admitted QoS traffic and best-effort traffic, must be controlled by an access control mechanism at the source node. The input to this access control mechanism consists of static and dynamic medium characteristics, medium access control (MAC) protocol characteristics, and real-time information about required bandwidth. All parameters can be either measured or provided manually. Furthermore, the system can be configured by parameters that allow several operating preferences, such as the balance between minimal bandwidth waste and a high safety margin for reserved QoS bandwidth, or the degree of fairness in bandwidth distribution among nodes and traffic classes.

Using an access control mechanism, we have developed two prototypes that are capable of supporting QoS and best-effort applications concurrently for multiple terminals and applications. The requirements for our prototypes were that they:

- Be non-intrusive to the network hardware, the drivers, and the operating system,
- Not be biased in favor of any network technology vendor,
- Not be biased in favor of any operating system (OS), and
- Be able to support legacy applications without modifications to the applications themselves.

By integrating and combining existing mechanisms and network devices, the prototypes provide solutions to the QoS problems we have discussed, while meeting these requirements.

Operation of the Admission Control Mechanism

The QoS-enabled admission control mechanism that we have developed is illustrated in **Figure 4** and explained below. Figure 4a depicts the available bandwidth for QoS reservations. This available bandwidth is configurable and is mainly determined by the dynamics of the total available shared medium bandwidth. Figure 4b shows a high-quality video QoS service being added successfully. In Figures 4c and 4d, additional services are added. In Figure 4e, a node attempts to add another high-quality video service, but is denied permission by the admission control mechanism, because of insufficient bandwidth. In figure 4f, a low-quality video QoS service is added successfully, while the admission of the high-quality video QoS service is refused.

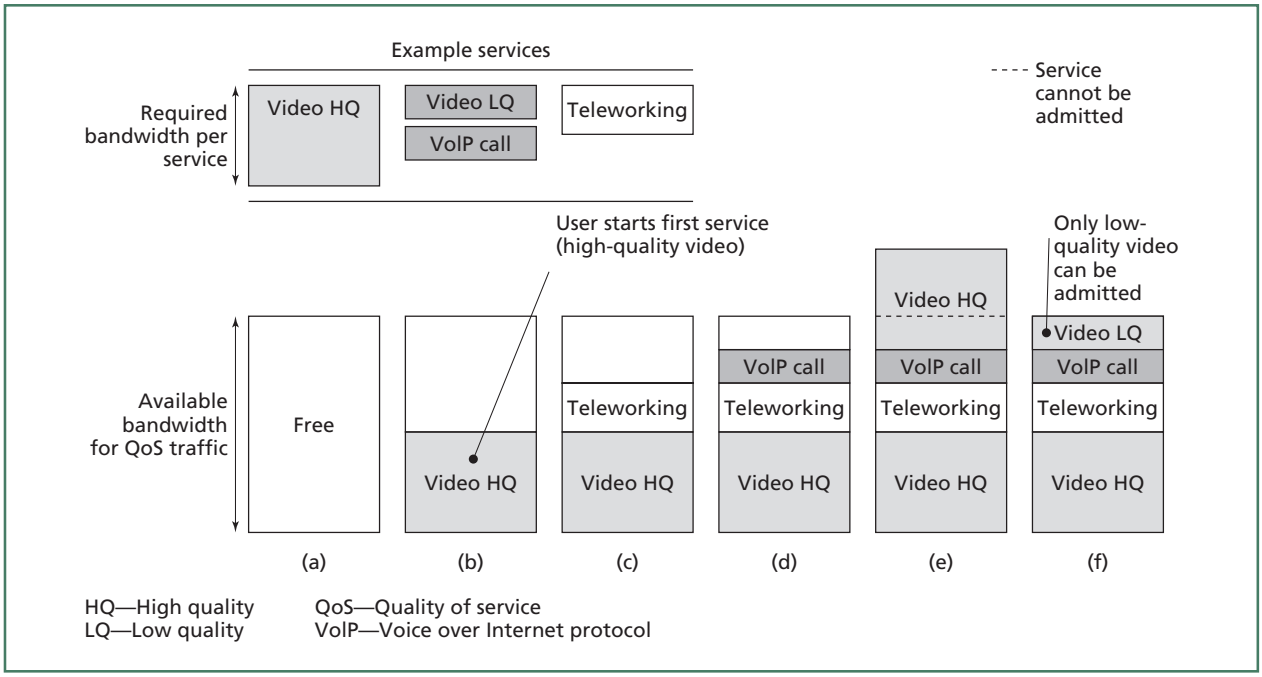

Figure 4.
QoS-enabled admission control.

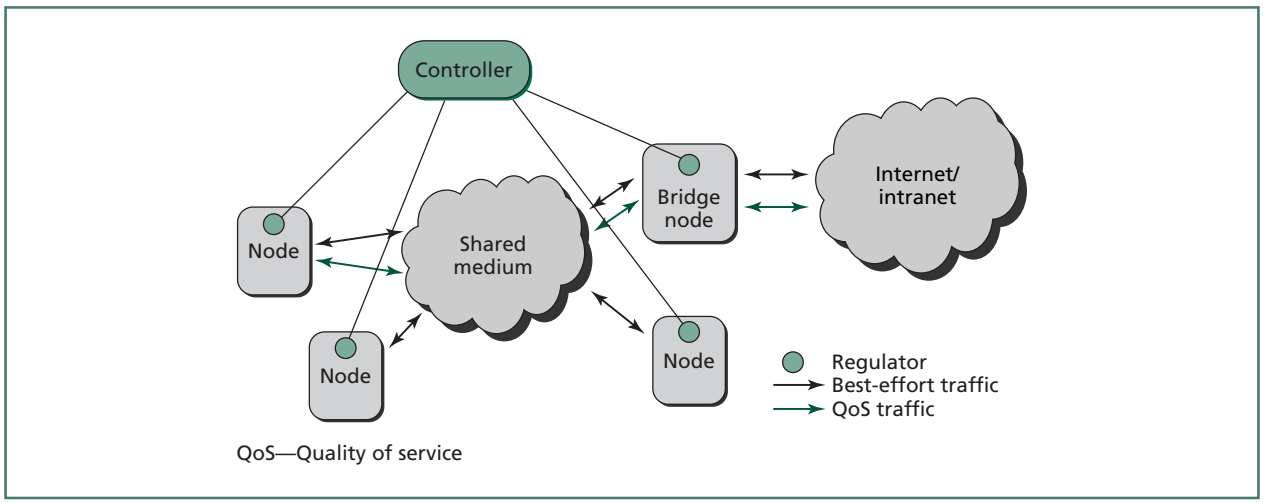

Figure 5.
A control system for a shared medium access network.

Best-effort traffic is controlled by distributing the remaining shared-medium bandwidth among the best-effort traffic sources, according to their bandwidth needs. These needs are either measured in real-time or taken into account implicitly, depending on the specific prototype implementation, as will be discussed in the following sections.

Architecture of the Admission Control Mechanism

The architecture of our admission control mechanism is illustrated in **Figures 5** and **6**. Figure 5 shows where the controller and the regulators are positioned in the network; both are technology-neutral. Figure 6 shows some details of the regulator, which shapes the best-effort traffic of a node according to the traffic-shaping specifications sent by the controller for this node. It is important to note that our solution does not affect regular QoS packet-forwarding. Our solution (with its controller and regulators) can be characterized as a feed-forward control system in which the amount of best-effort traffic occupying the network is controlled. It is based on the following control engineering [4] concepts:

- *Responsiveness.* The regulator values should respond quickly to changing circumstances. It should be possible, as in our tests, to configure the responsiveness of the control system in such a manner that the regulators reach a percentage (e.g., 90%) of their target values within 1 second. Doing so validates the subjective experience of the users that the system is indeed responsive. Furthermore, the minimum regulator value should not be zero; rather, it should be set to a small value that allows terminals to send and receive best-effort traffic with small bandwidth requirements immediately. If necessary, the regulator should be able to

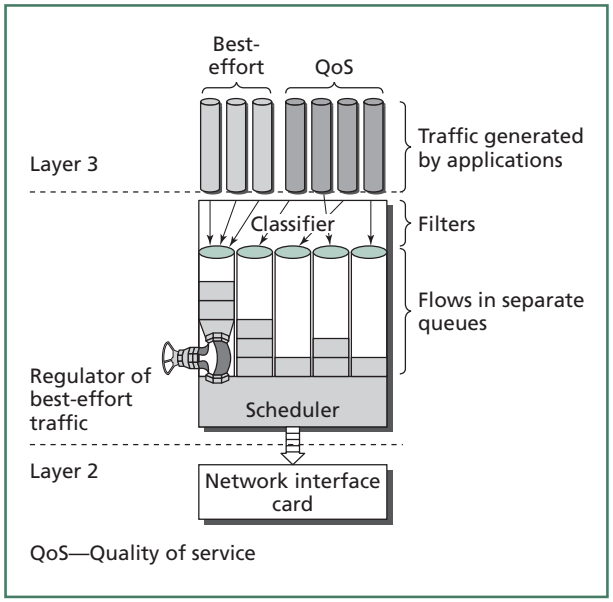

Figure 6.
A regulator of best-effort traffic.

increase the value, allowing more best-effort traffic to be sent. These conditions make possible increased responsiveness.

- *Steadiness.* It should be possible to maintain control of the control system. In order to do this, some configurable and adaptive mechanism could be applied (e.g., filters) to compensate for forces that could cause uncontrollable oscillations. This mechanism should be designed in such a way that the impact of regulator dynamics and the responsiveness requirements of the control system are brought into harmony.
- *Fairness.* All users should be treated equally (unless otherwise configured); no starvation should be possible. All users who have the same transmission conditions (e.g., an equal amount of data to send and the same SNR) should have the same regulator value. Furthermore, the architecture of the control system should be able to adopt other definitions of fairness.
- *Efficiency.* The algorithm should cause a minimum of computational, timing, and transmission overhead. The messages exchanged between the controller and the regulators should be as small as possible. Bandwidth that is not used for QoS traffic should be used as much as possible for best-effort traffic, subject to the restrictions of the configurable balance between minimal bandwidth waste and a high safety margin for reserved QoS traffic.

The input data for the control system algorithm are the fluctuating available bandwidth (measured in real time), the minimum regulator value, and the number of source nodes in the shared medium. The amount of bandwidth reserved for QoS traffic is subtracted from the fluctuating available bandwidth. The result is the bandwidth to be distributed for the best-effort traffic of the nodes.

The steps of the control algorithm we use can be described as follows:

1. Each node regularly sends information about best-effort traffic characteristics, network quality that has been measured during a configurable interval, and active QoS reservations.
2. The controller calculates how much best-effort bandwidth can be allocated for each separate node for the next interval, based on the inputs of step 1, and in accordance with the configured balance of the control-engineering concepts discussed above.
3. This bandwidth is divided by the access control mechanism in accordance with both the users' needs and their experienced network quality, as measured in step 1. The new regulator values are calculated.
4. Regulator values are sent from the controller to the users' regulators.
5. Steps 1 through 4 are repeated endlessly. The frequency of this control loop is determined by the responsiveness requirements of the deployment situations. In our prototypes, the interval was configured to be about 1 second.

WLAN Prototype

In our WLAN prototype, the best-effort data that terminal nodes send is regulated at the source. Each node is equipped with a regulator that controls the amount of best-effort traffic it is allowed to send. The regulators are positioned between layers 2 and 3 of the protocol stack. The throughput control of a regulator is described by token-bucket parameters. The Microsoft Windows* traffic control application programming interface (API) and a similar package (based on kernel sockets) for the Linux* operating system are used for the throughput control implementations of the regulators in our prototypes. Bandwidth for QoS traffic is reserved and handled by RSVP. The regulators are centrally operated by a controller that has real-time knowledge about the amount of QoS traffic and the fluctuating available network bandwidth. The controller calculates the bandwidth available for best-effort traffic, distributes it among the active nodes, and controls the regulators accordingly. In this process, an active node can acquire all the available bandwidth for best-effort bursts, as long as other nodes are inactive. If multiple nodes are active, the control mechanism will distribute the available bandwidth fairly. The controller takes into account the fact that data sent from one wireless node connected to an AP to another wireless node connected to the same AP travels the wireless medium twice, and so consumes twice the amount of bandwidth. We will call such traffic

internal traffic, and we will call other traffic to and from the AP external traffic.

In order to validate our WLAN approach, we performed a number of tests; we describe a few of them in the following sections.

WLAN Test Setup

All tests were performed using the basic configuration shown in **Figure 7**. The number of terminals varied in testing; the minimum number was 2. All terminals had a regulator, and the terminal physically connected to the AP hosted the controller. We implemented the controller software in the Java programming language and the regulators in C++. If not otherwise specified, the distance between each laptop and the AP was about 5 meters.

WLAN Test Results

All the tests depend on knowing how much of the theoretical maximum bandwidth of 1.1 Mb/s is available for use. To determine this, we had the terminals send as much data as possible, and we measured the throughput. The available bandwidth varied around 800 KB/s, depending on the number of terminals we used in the experiment. With two terminals, we measured throughput of around 830 KB/s, which is 60% of the theoretical maximum. We then used this figure as the maximum available bandwidth throughout the tests. Therefore, in all the tests, the maximum available bandwidth, which must be split between QoS and best-effort traffic, is equal to 60% of the theoretical maximum bandwidth. The exact relationship between the number of terminals and the available bandwidth can be found in [11].

WLAN test 1. This test shows that the responsiveness, steadiness, and fairness requirements are met. When two terminals try to send best-effort traffic with the same characteristics, and the SNR measured at the terminals is approximately equal, then the available best-effort bandwidth is divided equally between them.

Figure 8 shows a screenshot of the graphical user interface (GUI) of the system. The graph shows the regulator values for internal best-effort traffic. The horizontal grid lines are one tenth of the maximum, which is 900 KB/s. The vertical grid lines are 5 seconds apart. The dark green graph line shows the regulator values for terminal 1, while the light green line shows the regulator values for terminal 2. The time line in the graph is from right to left, so the part of the graph numbered 1 represents the oldest regulator values.

The following events can be distinguished in Figure 8:

1. The control system is disabled, which means the prototype runs in a standard best-effort traffic

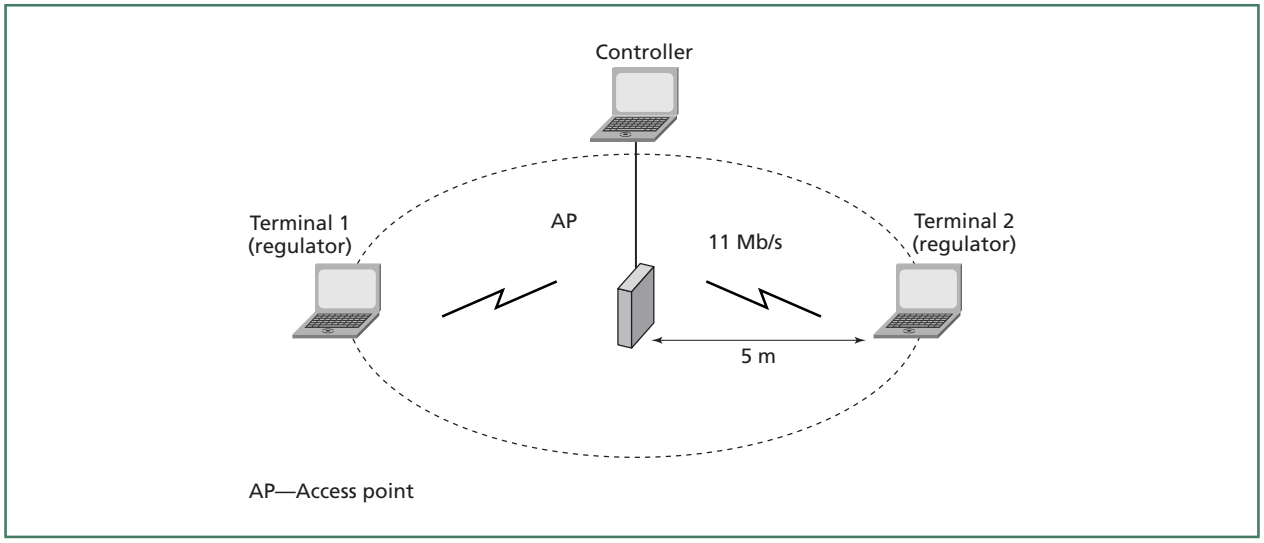

Figure 7.
Wireless LAN prototype.

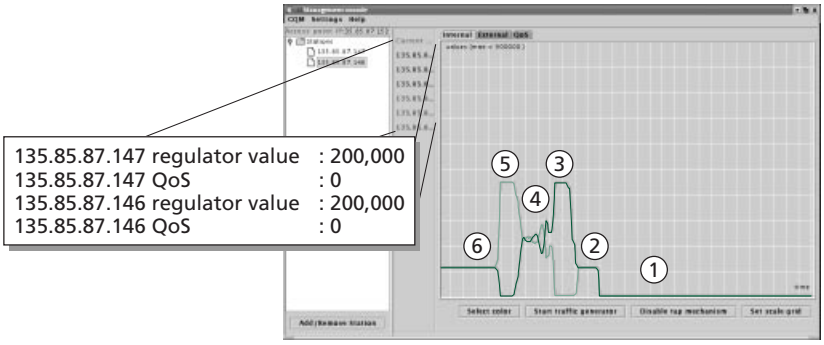

Figure 8.
Wireless LAN test 1 results.

configuration. In this situation, video streams will be disrupted by high-load best-effort traffic.

2. The control system is enabled. No traffic is transmitted by the terminals. This causes the control system to distribute the available best-effort bandwidth to the regulators. Both regulators are configured to a quarter of the maximum available bandwidth for internal traffic and a quarter of the maximum available bandwidth for external traffic. This accounts for the total amount of available best-effort bandwidth.
3. Terminal 1 starts transmitting internal user datagram protocol (UDP) traffic. The controller configures the internal traffic regulator for terminal 1 to the maximum amount of available best-effort bandwidth, increasing the traffic by a factor of 4. The internal traffic regulator for terminal 2 is configured to zero. The external traffic regulators for both terminals are also configured to zero. The responsiveness of the system is around 1 second, as can be seen in the figure.
4. Terminal 2 starts transmitting internal UDP traffic. The available best-effort bandwidth is divided equally between the internal traffic regulators for the two terminals. The regulator values become steady in a few seconds. In contrast to step 1, video streams are no longer disrupted by high-load best-effort traffic.
5. Terminal 1 stops transmitting. Terminal 2 gets all the available best-effort bandwidth for its internal UDP traffic.
6. Terminal 2 stops transmitting as well. Because neither terminal is transmitting, all regulators have the same value. The system is in the same state it was in step 2.

WLAN test 2. This test, the results of which are depicted in **Figure 9**, shows what happens when two terminals transmit different amounts of traffic. As in the previous test, the regulator values for terminal 1 are indicated by the dark green line, the regulator values for terminal 2 by the light green line. The following events can be distinguished in Figure 9:

1. The control system is disabled. (See step 1 of test 1.)
2. The control system is enabled. No traffic is transmitted. All regulators have the same value.
3. Terminal 1 starts to transmit external UDP traffic at a rate of 256 KB/s.
4. Terminal 2 starts to transmit external UDP traffic at a rate of 128 KB/s. Terminal 2 gets less bandwidth than terminal 1. However, the regulators are not configured to the precise bandwidth requirements of the terminals, which total 384 KB/s. Since more best-effort bandwidth is available (i.e., 900 KB/s), the controller distributes the remaining 516 KB/s best-effort bandwidth equally to the regulators, as can be seen in Figure 9.

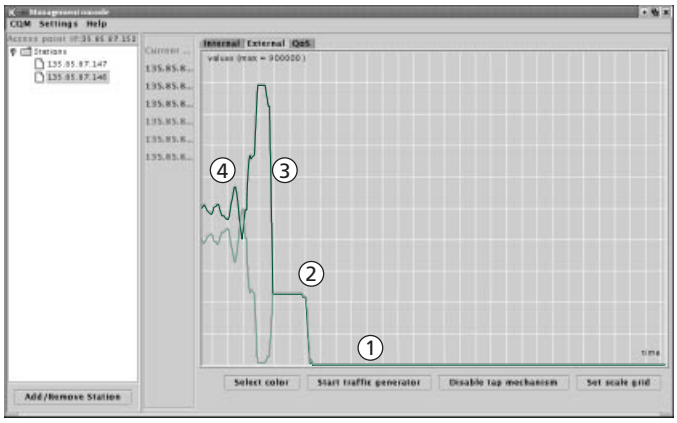

Figure 9.
Wireless LAN test 2 results.

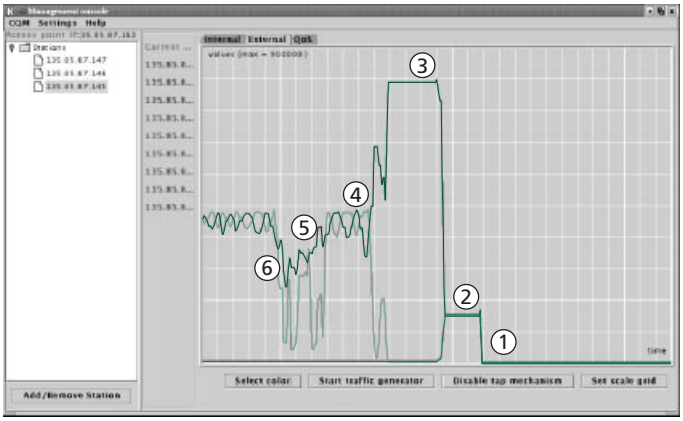

Figure 10.
Wireless LAN test 3 results.

WLAN test 3. The third test, the results of which are depicted in **Figure 10**, shows the impact of the distance from a terminal to the AP on the available bandwidth for that terminal. As in the previous test, the regulator values for terminal 1 are indicated by the dark green line, the regulator values for terminal 2 by the light green line. When a terminal moves away from the AP, its SNR decreases. This means that it is more difficult to separate the received signal from the background noise. Also, as the distance increases, the BER increases, more retransmissions are necessary at the MAC level [6], and the throughput at the Internet protocol (IP) level drops. The results of extensive testing of the relationship between the SNR and the regulator value can be found in [11]. The following events can be distinguished in Figure 10:

1. The control system is disabled, as in step 1 of tests 1 and 2. Both terminals are 5 meters from the AP.
2. The control system is enabled. No traffic is transmitted. All regulators have the same value.

3. The first terminal starts to transmit external UDP traffic.
4. The second terminal starts to transmit external UDP traffic.
5. The second terminal is placed 20 meters from the AP, while the first terminal is moved slowly from its initial position 5 meters from the AP to a position 10 meters from the AP.
6. Both terminals are moved back to their initial positions.

In the WLAN tests, the minimum amount of best-effort traffic that could be transmitted by a node (i.e., the minimum regulator value) was set to 4,000 bytes per second, so a terminal could send up to 4,000 bytes per second of best-effort traffic without the value of the regulator having to be increased. This further improved the responsiveness of the system.

xDSL Prototype

Figure 11 shows the architecture of our xDSL prototype. The network domain labeled "Access network" is the domain that we refer to in this paper as the broadband shared-medium access network.

The access network provides physical and logical connections from multiple nodes to a broadband remote access server (BRAS). These nodes can be single end-user personal computers (PCs) or small-office home-office (SOHO) local area networks (LANs). Node connections are implemented using DSL modems and DSLAMs. The BRAS functions both as an edge router that connects the access network to the core network and as a QoS-enabled admission controller. QoS-enabled admission control is implemented by extending the standard point-to-point protocol (PPP)-based authentication and admission control mechanism of the BRAS to include an additional verification procedure. The available medium bandwidth and already-existing QoS sessions are taken into account during the execution of this verification procedure, to ensure that sufficient bandwidth is available in all connections of the access network. Because the BRAS is an edge router that functions like the bridge node in Figure 2, it regulates best-effort traffic from the core to the access network by means of integrated-shaping and queuing functionality.

In order to validate our xDSL approach, we performed a number of tests. The tests that we describe here are concerned with download traffic (i.e., traffic from the core to the access network). These tests complement the WLAN tests of our general QoS-enabled admission control mechanism.

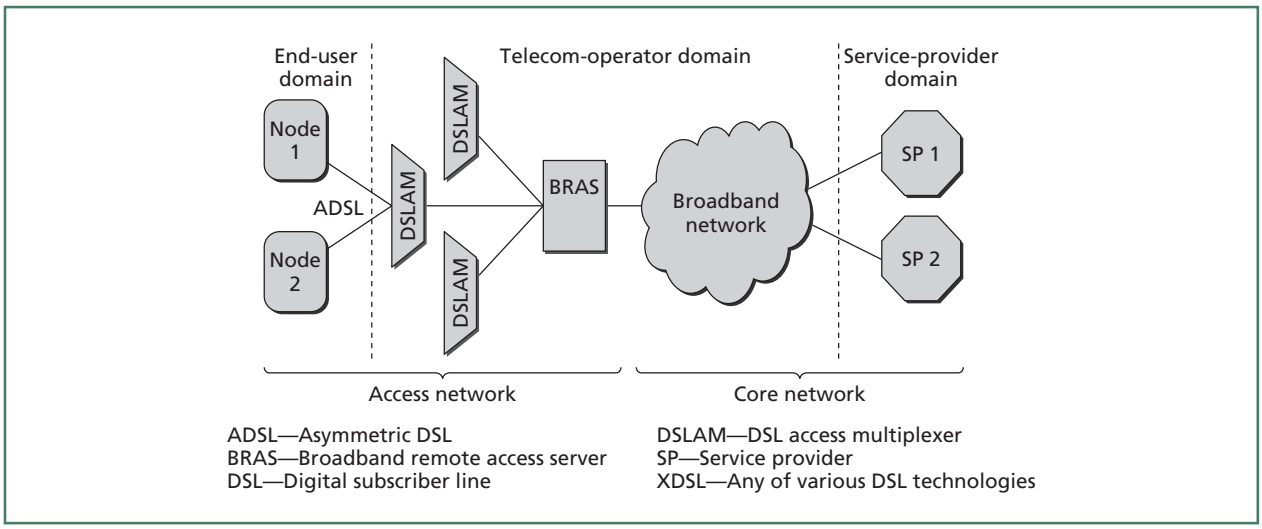

Figure 11.
xDSL prototype.

xDSL Test Setup

The test setup is based on asymmetric digital subscriber line (ADSL), the asymmetrical variant of xDSL, which uses the greater part of the physically available bandwidth for download traffic and the remaining (smaller) part for upload traffic.

Three types of ADSL termination units for the remote site (ATU-Rs) were used to reduce the risk of device-specific performance characteristics; Lucent Technologies' CellPipe™, Efficient Networks' SpeedStream* 5621, and Alcatel's Speed Touch Home.* We used the Lucent Technologies' AnyMedia™ access system as the DSLAM, and the Redback SMS 1800 and the Lucent Technologies' Springtide™ 5000 as BRAS. For the subscriber profile database we used a RADIUS server on a SUN SparcStation.* We implemented admission control extensions in the NavisRADIUS™ package from Lucent Technologies. To obtain quantitative test results, we used a Netcom Smartbits* network test unit and a Linux PC equipped with a tcpdump setup. For information regarding tcpdump, see [12] and [13].

xDSL Test Results

Our tests demonstrated that a QoS service that requests more bandwidth than is available in the xDSL communication path is correctly denied. The denial occurs regardless of whether the node already has any active QoS sessions engaged. Where QoS traffic is concerned, the test results of our xDSL prototype are hard to quantify. We have seen that the basic characteristics of QoS traffic (e.g., throughput, latency, and packet loss) are not adversely affected by the best-effort traffic load. In fact, the throughput of the QoS sessions remains exactly as specified during admission of the sessions, and the packet loss is 0%, regardless of the best-effort traffic load being applied concurrently. Thus, the QoS sessions are not affected by existing or new best-effort traffic.

Figure 12 shows best-effort traffic latency while different QoS sessions are engaged. These results were obtained using tcpdump. The QoS sessions occupied less than 70% of the available medium bandwidth at any time during the tests. This left enough bandwidth for best-effort traffic to allow us to measure the

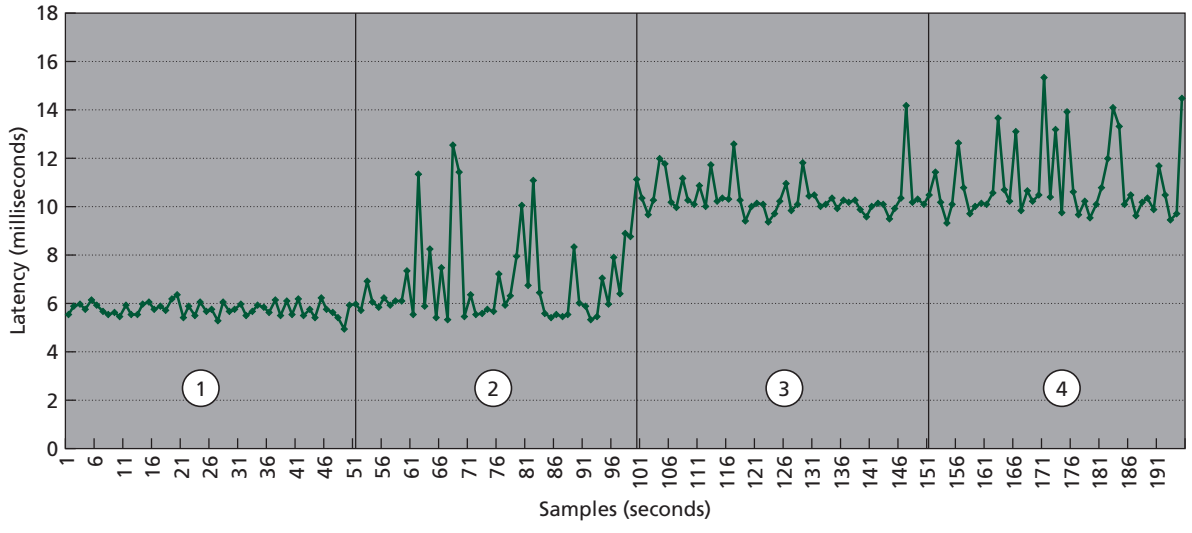

DSL—Digital subscriber line
xDSL—Any of various DSL technologies

Figure 12.
Latency of best-effort traffic in the xDSL communication path.

latency effects on best-effort traffic. The following events can be distinguished in Figure 12:

1. Only the best-effort test traffic from tcpdump is present in the ADSL connection path between the ATU-R and the BRAS. Measured latency is around 6 ms.
2. A QoS-enabled file transfer protocol (FTP) session using 500 KB/s is started; this causes the latency of the best-effort test traffic to become higher and to behave erratically. The reason for this is that FTP uses transmission control protocol (TCP) as its underlying transport protocol. The erratic effects are caused by the TCP protocol mechanisms interacting with the shaping and queuing in the communication path between the sender and the receiver. For a discussion regarding these effects, see [9].
3. The QoS-enabled FTP session is stopped and a QoS-enabled UDP stream transporting video data is started. This QoS traffic consumes 1.5 Mb/s and causes the average latency of the best-effort traffic to become higher. The latency becomes less erratic because of the constant UDP stream from the video server and the lack of TCP traffic-management mechanisms in the UDP protocol.
4. The QoS-enabled FTP session of step 2 is started again, while the video session remains, producing a total of 2.0 Mb/s of QoS traffic. The same erratic latency reappears, and the average and the bottom line are raised by the QoS traffic of the video session.

These tests show that the BRAS is capable of policing and shaping the network traffic for individual QoS sessions, effectively isolating the sessions from each other. The traffic flows of other QoS sessions for the same node are not affected in terms of packet loss, delay, or jitter. The effects on best-effort traffic characteristics (i.e., packet loss, delay, and jitter) also remain unchanged. This indicates a correct implementation and coupling of the QoS-enabled admission control with the traffic-regulator functionality of the BRAS.

Conclusions

This paper has addressed the lack of QoS in today's broadband shared-media access networks. Our approach supports QoS for multiple sessions per node by incorporating a QoS mechanism that uses admission control based on reservations. Our prototype tests have shown two successful implementations of this mechanism, one in an xDSL, the other in a WLAN environment. In our prototypes, QoS sessions are denied if sufficient bandwidth is not available. Furthermore, services with different traffic characteristics can be used simultaneously, without quality degradation. We have also seen that the shaping functionality of the regulators in the BRAS and the WLAN source nodes protect active QoS sessions from otherwise uncontrolled best-effort traffic. Combined with edge- and core-network QoS capabilities, the reservation model QoS solution presented in this paper offers a straightforward mechanism that can be applied to any transport protocol, as long as the start-up and termination of all QoS sessions in the access domain can be determined by the controller. Furthermore, this approach is complementary to prioritization model QoS mechanisms in the same broadband shared-medium access network.

Acknowledgments

The authors would like to thank the following people for their contributions: Ko Lagerberg, Jan-Peter Sanderman, Peter Leijdekkers, Jacco Brok, Jeroen van Bemmel, Sjoerd Talsma, Dirk Jaap Plas, Ronald de Man, Ronald van Haalen, Richa Malhotra, and Jeroen Schot.

References

[1] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An Architecture for Differentiated Services," IETF RFC 2475, Dec. 1998, <http://www.ietf.org/rfc/rfc2475.txt?number=2475>.

[2] R. Braden, D. Clark, and S. Shenker, "Integrated Services in the Internet Architecture: an Overview," IETF RFC 1633, June 1994,

<http://www.ietf.org/rfc/rfc1633.txt?number=1633>.

[3] R. Braden, L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource reSerVation Protocol (RSVP)—Version 1 Functional Specification," IETF RFC 2205, Sept. 1997, <http://www.ietf.org/rfc/rfc2205.txt?number=2205>.

[4] K. Dutton, S. Thompson, and B. Barraclough, The Art of Control Engineering, Addison-Wesley, Boston, 1997.

[5] Institute of Electrical and Electronics Engineers, "MAC Bridges," 802.1D Annex H, <http://www.ieee802.org/1/pages/802.1D.html>.

[6] Institute of Electrical and Electronics Engineers, "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications," ISO/IEC 8802-11:1999(E), ANSI/IEEE Std 802.11, 1999.

[7] Institute of Electrical and Electronics Engineers, P802.11 Task Group E, "Draft Supplement to Standard for Telecommunications and Information Exchange Between Systems—LAN/MAN Specific Requirements—Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)," IEEE 802.11e/D2.0, Nov. 2001.

[8] A. C. Milonas, "Enterprise Networking for the New Millenium," Bell Labs Tech. J., 6:1 (2001), 73–94.

[9] R. Morris, "Scalable TCP Congestion Control," IEEE INFOCOM 2000 Conference, (Tel Aviv, Israel, 2000).

[10] E. Rosen, E. A. Viswanathan, and R. Callon, "Multiprotocol Label Switching Architecture," IETF RFC 3031, Jan. 2001, <http://www.ietf.org/rfc/rfc3031.txt?number=3031>.

[11] J. P. Sanderman, "A Traffic Control Algorithm for Wireless LANs," Thesis, University of Twente, 2001.

[12] tcpdump, <http://www.tcpdump.org>.

[13] tcpdump, Network Research Group (NRG), Information and Computing Sciences Division (ICSD), Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA, <http://ee.lbl.gov>.

[14] R. Yavatkar, D. Hoffman, Y. Bernet, F. Baker, and M. Speer, "SBM (Subnet Bandwidth Manager): A Protocol for RSVP-based Admission Control over IEEE 802-style Networks," IEEE RFC 2814, May 2000, <http://www.ietf.org/rfc/rfc2814.txt>.

(Manuscript approved April 2003)

BASTIEN PEELEN is a member of technical staff in the

Bell Labs Advanced Technologies EMEA department, at Lucent Technologies in The Netherlands. He holds an M.Sc. degree in electrical engineering from the University of Twente, The Netherlands, with a specialization in network theory and signal processing. At Bell Labs he specializes in network infrastructure, QoS, and service platforms for wireless and wireline systems. Mr. Peelen has received three Lucent and Bell Labs recognition awards and has two patents pending for his work in this field.

MIROSLAV ZIVKOVIC is a member of technical staff in

the Bell Labs Advanced Technologies EMEA department, at Lucent Technologies in The Netherlands. He holds a Dipl. Ing. degree in electronics and telecommunications from the Faculty of Electrical Engineering in Belgrade, Serbia and Montenegro. His current research interests are in the area of mobile and wireless system security. He received a central Bell Labs Teamwork Award for his work as a member of the 3G Services Platforms Team.

DENNIS BIJWAARD is a member of technical staff in the

Bell Labs Advanced Technologies EMEA department, at Lucent Technologies in The Netherlands. He received an M.Sc. degree from the University of Twente, The Netherlands, working on language parsing and function approximation using artificial intelligence techniques. He now works on next-generation networks and the visualization of software and protocol dynamics. He received a central Bell Labs Teamwork Award for his work as a member of the 3G Services Platforms Team. His main interests are software reverse engineering, distributed computing, QoS, and artificial intelligence.

HAROLD TEUNISSEN is a member of technical

staff in the Bell Labs Advanced Technologies EMEA department at Lucent Technologies in The Netherlands. He holds a B.Sc. degree in electrical engineering from the Polytechnical College, Heerlen, and an M.Sc. degree in

computer science from the University of Twente. His current research interests are in the areas of mobile and wireless systems beyond 3G. Currently, he is responsible for research on service roaming and mobility management for integrated wireless systems. Mr. Teunissen has received three Lucent and Bell Labs recognition awards and has one patent pending for his work in this field. ◆