

An On-line Managed Bandwidth Service for ATM Networks

C. Bouras C. Chantzi V. Kapoulas I. Samprakou A. Sevasti

*Computer Technology Institute-CTI, Riga Feraiou 61, 26221 Patras, Greece, and
Computer Engineering and Informatics Dept., University of Patras, 26500 Rion, Patras, Greece*

e-mail : bouras@cti.gr

ABSTRACT

In our days, the number of demanding, in terms of network bandwidth, applications increases dramatically as the number of users that obtain access to networking facilities augments. As a result, efficient management of the available network resources becomes a critical issue both from a functional point of view, so that users can be provided with the bandwidth they need, and an economical point of view, so that carriers can satisfactorily and efficiently serve as many customers as possible. The work presented here attempts to face from a practical point of view the problem of bandwidth management in ATM networks, especially in cases where the available resources have to be shared among many users with conflicting needs. We have designed and implemented a Managed Bandwidth Service (MBS) that allows users to make multiple bandwidth reservations throughout the topology of the network infrastructure to which they have access. This MBS also provides an administrative module for configuring and monitoring the service's operation. In this paper we present the MBS developed and all the implementation issues related to it, as well as our on-going and future work on the service.

I. INTRODUCTION

In the last few years, great attention has been given to the design and implementation of mechanisms for network bandwidth management. The number of applications with remarkable demands in terms of network resources and quality of service is constantly rising, resulting in underlying networks exhibiting denial of service and long delays, and facing several bandwidth availability problems. Such applications require i.e. support for the transmission of voice, video and all kinds of multimedia data. In cases of lack of resources, the service provided does not meet the desired, promised or expected quality levels. Therefore, it is becoming increasingly vital to ensure high-quality network services for bandwidth critical applications.

At the same time, lack or ineffective allocation of resources results in limitations to the number of users that can be simultaneously served, regardless of whether they are running bandwidth critical applications over the underlying network or not. These limitations can prove to be extremely hazardous for the networking facilities provider, especially when users are charged for their

access to the network. In this case, efficient management of bandwidth can significantly contribute to the improvement of network use from an economical point of view. In this paper, we present the design and implementation of an on-line Managed Bandwidth Service (MBS) for ATM networks. This work took place within the pilot project 'Managed Bandwidth Services of GUNet: Greek Universities Network' ([20]) and the project 'Management of Additional Services-MBS' of the Greek Research and Technology Network-GRNET ([19]). It was also partially funded by the project 'Management of network services with Quality assurance' of the Support for Research Activities Programme (PENED'99-EPET II).

The main goal of this MBS is to provide users with a mechanism to reserve network resources for applications such as 'real time' transmission of data/image/voice. The service allows the definition of virtual connections with reserved bandwidth between two network nodes taking into consideration the requested lifetime of the established connections, the requested traffic characteristics profile and a complete set of network parameters.

The requests for virtual connections can be made during a user session, throughout which the user can ask for more than one virtual connection. Our MBS is an end-to-end service, and point-to-point or point-to-multipoint connections are established as soon as the users' corresponding requests are declared as accepted. The user can make his requests via a user-friendly Web interface. In addition, the service provides the network manager with a Web interface for the management of users' requests and network topology. The connections are implemented either as permanent virtual paths or as permanent virtual circuits, depending on the users' requirements and the availability of the system.

The rest of the paper is organised as follows: In section 2 a brief overview of the initiatives on the issue of network bandwidth management is provided. This overview contains references to research work, international projects and commercial products that deal with bandwidth management. Section 3 describes our MBS from an operational point of view, while section 4 deals with the implementation issues of the service. Finally, we describe our conclusions with respect to the service as well as our on-going and future work on it, in section 5.

II. RELATED WORK

Bandwidth allocation has been a popular research subject during the last years. More specifically, in the case of ATM networks, a variety of related algorithms that aim to the best way of implementing resource allocation and flow control have been proposed. These algorithms can be classified in the following five levels: cell level (priorities, cells flow control), flow level (burst control, etc.), connection level (Connection Admission Control-CAC that ensures the existing bandwidth is enough for guaranteed quality), Virtual Path level (bandwidth allocation), network level (routing of Virtual Paths). An integrated solution to the bandwidth management with guaranteed quality problem, needs to tackle at least with the partial problems of connection (CAC), Virtual Path and network (routing) levels.

As far as the connection level is concerned, the authors of [1] are presenting a multi-class CAC policy for high-speed ATM switches. According to this policy, traffic is described with respect to the usage parameter control (UPC) parameters. In [2], the authors are proposing three algorithms for implementing efficient bandwidth allocation and call admission control for VBR service using UPC parameters. Concerning the Virtual Path and network levels, the authors of [3], are dealing with Virtual Path management in ATM networks and proposing a Virtual Path Bandwidth Resizing Algorithm that conducts utilisation driven virtual path bandwidth adjustment. Finally, the authors of [4] are presenting the subject of Bandwidth Allocation for Virtual Paths (BAVP) that aims at the optimal bandwidth allocation (capacity, service rate) in Virtual Paths, based on observations of the overall network status.

In order to solve the bandwidth availability problems of today's networks, many companies offer several products/solutions for managing network bandwidth. These products are known as traffic tuners, traffic shapers or bandwidth managers.

WiseWan, by NetReality ([5]) falls into the category of hardware-based bandwidth managers. Software-based traffic shapers demonstrate slower performance and longer set-up time, but easier upgrading. Check Point's FloodGate-1 ([6]), Elron Software Inc.'s CommandView Bandwidth Optimizer ([7]) for monitoring, analysing and controlling the bandwidth utilisation of Internet connections and NetGuard's Guidepost ([8]), which provides users with bandwidth allocation priority according to task importance, are such traffic shapers. Ukiah Software Inc.'s Netroad Trafficware ([9]) is another software-based bandwidth manager that can apply bandwidth policies to users or groups already defined by a standard directory service, while Sun Microsystems Inc.'s Sun Bandwidth Allocator 1.0 ([10]) is a software product that controls traffic sent over a network link and QoS. Solaris Bandwidth Manager 1.6 ([11]) provides network administrators with full control of network resources, enabling them to provide differentiated classes of service.

The third category of bandwidth managers that combine hardware with software and perform shaping functions

using software embedded in dedicated hardware devices, contains products such as Allot Communications Inc.'s Netenforcer ([12]), which is a hybrid product that allows bandwidth sharing between many ISPs and domains. Packeteer Inc.'s Packetshaper ([13]) is another hybrid bandwidth manager that offers bandwidth management, classification of network traffic and bandwidth allocation based on predefined policies.

Similar, to the MBS presented in this paper, services have been developed in several national research networks. One of the most known similar services has been developed by the TEN-155, a high-speed network for the pan-European academic and research community. The TEN-155 Managed Bandwidth Service ([14]), is an end-to-end service which enables connectivity among research projects in countries connected to TEN-155. UKERNA has also developed a MBS, called the SuperJANET Managed Bandwidth Service ([15]), in order to meet the increasing demands for guaranteed QoS over the SuperJANET III backbone and the requirements of the research community. Another national network where an MBS is being developed is the SWITCHlan network, which is also based on an ATM backbone. The SWITCH Managed Bandwidth Service ([16]), allows the creation of VPNs with committed bandwidth between interested parties on top of the existing network infrastructure. Finally, the VCCS (Virginia Community College System) Intranet ([17]), has adopted Policy-Based Management that allows for central control of bandwidth management and the enforcement of QoS policies. A similar service implemented on an IP network, is the 'Reserved Bandwidth' service of the vBNS backbone network ([18]).

III. DESCRIPTION OF THE MBS SERVICE

The MBS presented here has as its main goal to provide facilities for managing the bandwidth of GRNET, the Greek Research and Technology Network that provides Internet services to the Greek Academic and Research community. GRNET interconnects Universities and Research Centres in Greece, as well as other R&D Departments of industrial organisations, through an advanced, high speed network. The GRNET backbone includes, at this stage, network nodes in 6 major Greek cities and participates actively in the Pan-European TEN-155 high speed computer network. It is based on an ATM backbone at the edges of which, IP routers are connected.

In the MBS, bandwidth reservation is implemented by the establishment of virtual connections between two network nodes. The service has the ability to provide different connection parameters for each one of the requested virtual connections and per different directions between the nodes. The connection parameters refer to the various types of service that can be provided in an ATM network (CBR, VBR, ABR and UBR). It is not initially feasible for the application to support all different types of service (e.g. ABR). That is either because some of them are not supported by all the ATM

switches of GRNET or because the required information as to the CAC procedure of ATM switches is not available. At the moment, MBS provides CBR connections while in the future it will be upgraded in order to provide more service types.

The service is based on users' requests, which give information about each virtual connection requested. These include the origin and destination of a virtual connection, the duration of bandwidth reservation, the usage type of the virtual connection (what type of traffic is it used for, i.e. MPLS, RTS, DiffServ, Voice over ATM, IP, Video on Demand etc.), the requested bandwidth, the service type and the connection type: Permanent Virtual Path Connection (PVPC) or Permanent Virtual Channel Connection (PVCC). For simplicity purposes, these two types will be referred to as VP and VC from now on.

MBS provides users with a number of predefined contracts for the requested connection. These contracts describe the requested bandwidth to be reserved, the connection service type and one or more parameters for the connections to be established, so that the connections will satisfy both the user's requests and the network performance's improvement objectives. Depending on the service type of the connection to be established (CBR, VBR, ABR, UBR), the contract provided by the application and selected by the user is used for specifying the values of the following parameters:

Service type	Parameters specified by a corresponding contract
CBR	<i>Peek Cell Rate</i> <i>Cell Delay Variation Time</i>
VBR	<i>Peek Cell Rate</i> <i>Sustainable Cell Rate</i> <i>Maximum Burst Size</i> <i>Cell Delay Variation Time</i>
ABR	<i>Peek Cell Rate</i> <i>Minimum Cell Rate</i> <i>Minimum Cell Rate</i>

Table 1. Parameters determining the predefined contracts in each case of connection

Once the user provides all necessary information and selects the predefined contract for the requested connection, all requests are being examined by the system and if the network load can bear the specified connection, the system makes the reservation. If the application for a connection is rejected, the user is able to reapply for a new request (as far as the requested bandwidth and requested time of the connection are concerned) within the same session. In order to help the user apply for an acceptable request, the application provides him with a network topology map that depicts network load and is being dynamically adjusted each time a new reservation is made.

In case a request is accepted, then the internal updating procedure of the appropriate modules begins. This procedure includes the notification of the MBS manager, programming of the ATM switches, and notification of

the involved parties. During the current phase of the service, the programming of the ATM switches is "manually" performed by the manager, according to the specifications of each accepted virtual connection that can be accessed by the manager interface. The activation of new connections includes the insertion of information concerning VPI/VCI, input and output ports at the switches participating in a connection, through the MBS manager interface.

MBS service at the moment offers only CBR contracts in 512Kbps, 1Mbps and 2Mbps. The service is provided for a specific amount of the GRNET capacity together with a number of virtual paths and circuits. The Web-based interfaces provided to the users and the manager of MBS, can be used for applying for connections, responding to requests, presenting the network status, insertion of VPI/VCI information to the ATM switches, etc. Applications for virtual connections must be submitted by one of the involved parties, in case more than one connections are requested (point-to-multipoint service) and must be made at least one day and at most one month before the date for which the virtual connection is requested to start. The shortest allowed duration for the duration of a requested virtual connection is one day, and the longest is two months.

IV. IMPLEMENTATION ISSUES

A. Service's Database

The database of the service contains information related to the topology and the load of the GRNET topology along with details concerning the individual nodes and links. The contents of the database are constantly being updated in order to reflect the current network state and are being used by the decision-making module of the system. The database also contains the tables where the data of each session are stored: the virtual connections that have been reserved within a session and a detailed description of the implementation of each connection (origin-destination-path, VPIs-VCI, etc.). In addition, the database contains the predefined contracts by the MBS, stores data concerning the users and the managers of the system and maintains statistics about the reservations. The database of the MBS service is used as the intermediate interface between the other parts of the service.

B. User Interface

As it has already been mentioned, the MBS service provides users with a Web-based interface through which, they can make their requests for virtual connections and receive the system's responses. The user forms can be used in order for the user to insert the information concerning his requests for virtual connections and get an acknowledgement relatively to the results of his requests. In Figure 1, the form that the user has to use in order to make an application for a virtual connection is presented.

Request for MBS connection

Session data

The first part of the following form is filled in with data in respect of an MBS session and is being filled in only once at the beginning of a session. (One session may consist of more than one requests for MBS connections, therefore you need not fill in the session's data every time you make a new request)

User name:
Sampraku Ioanna

User e-mail:
sampraku@ceid.upatras.gr

User access node to GRnet:
Patras node

Beginning of the session:
Time: 9:00
Day: 25 Month: April Year: 2001

End of the session:
Time: 12:00
Day: 30 Month: May Year: 2001

MBS connection data

Fill in this part of the form with the desirable features of the new MBS connection, which you submit within the present session.

Origin of MBS connection (select the node of GRnet backbone from which your MBS connection starts):
Patras node

E-mail of the administrator of MBS service in the origin:
sevastia@cti.gr

Destination of MBS connection (select the node of GRnet backbone at which your MBS connection ends):
Athens node

E-mail of the administrator of MBS service in the destination:
hantzi@ceid.upatras.gr

Usage type (select the usage type of the requested MBS connection from the following list):
MPLS

Predefined contracts (the requested bandwidth and service are provided by the following contracts):
Predefined contract for CBR, 2.000 Mbps with PCR=37732, CDTV=-, SCR=-, MBS=-

Connection type (select the desirable type of connection from the following list):
Virtual Path

Submit Reset

End of session

Figure 1. The user's application form for requesting a virtual connection

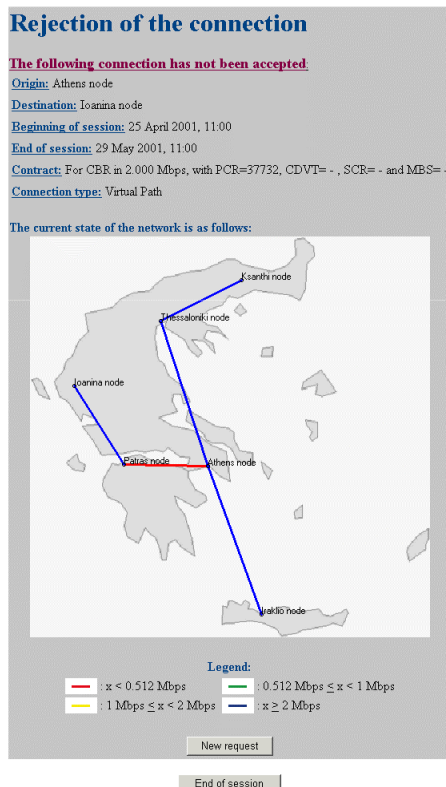


Figure 2. The user's feedback form concerning a previously made application

In case a request is rejected (see Figure 2), the user has the ability to formulate it again. Before making his new

request, he must take into account the current state of the network as it is being shown on the map of the network topology, where the links of the network are painted as follows:

- If the last request concerned a VP, the link is painted so that it reflects the available bandwidth of the link, that is, the total bandwidth of the link minus the sum of the bandwidths of the VPs that already exist in the link.
- If the last request concerned a VC, the link is painted so that it reflects the largest bandwidth value among:
 - the total bandwidth of the link minus the sum of the bandwidths of the VPs that already exist in the link
 - the highest available bandwidth within a VP (for all the VPs that already exist in the link)

C. Decision-making module

This part of the application is responsible for deciding whether a user's request for the reservation of a virtual connection will be accepted or not. The decision-making module is based on a simple yet efficient algorithm that combines the principles of least-cost-path routing and FCFS in order to make that decision. The pseudocode in Figure 3 describes the major steps followed throughout the algorithm designed.

The process described in this algorithm is being executed by the MBS whenever the user makes a request for the reservation of a virtual connection via the user's application form. The third step of the algorithm depicted in Figure 3 involves the computation of weights or values for all the links of the network topology. These values are being used as input data to the fourth step of the algorithm: the execution of the Dijkstra routing algorithm in order to find the 'best' path between the origin-node and the destination-node of the requested virtual connection.

More specifically, whenever the Dijkstra algorithm is executed (in the context of the decision-making module's algorithm), the algorithm itself assigns a numerical value (link_value) to every link of the network topology. This value represents the current traffic load of the link. The computation of each link's value is based on the currently available bandwidth of the link (BW_avail) and the minimum bandwidth of the link (Min_BW_reserved) a user may reserve. At the moment, the latest value is equal to:

$$\text{Min_BW_reserved} = 512 \text{ Kbps} \quad (1)$$

The value of BW_avail is computed by removing the total bandwidth reserved by virtual paths from the total bandwidth of the link. To be more specific, the value of BW_avail for link i is calculated as follows:

$$\text{BW_avail}_i = (\text{capacity}_i - \sum_j \text{total_BW}_j) + \sum_j \text{avail_BW}_j \quad (2)$$

where j indicates each VP in the link_i.

- 1 *Extraction of network traffic data from the system database*
- 2 *Temporary removal of all network links, that cannot accommodate the requested connection bandwidth, from the network topology*
- 3 *Calculation of values that represent current workload for the remaining network topology*
- 4 *Execution of the Dijkstra algorithm for the current network topology, starting from the origin node of the requested connection:*

IF {the emerging routing table contains the destination node of the requested connection}
THEN

- 4a *the request for connection is accepted*
- 4b *the database is locked for the processing of other applications*
- 4c *the path from the origin node towards the destination node of the requested connection is created according to the existent routing table*
- 4d *get the next connection path link (starting from origin node, the connection path is being traversed and appropriate resources are allocated)*
- 4e **IF** {the current link is terminating at the destination node}

THEN

all necessary resources have been allocated, go to 4f

ELSE

IF {the connection requested was a Virtual Path}

THEN

a new VP link is created between the nodes of the current link

ELSE

IF {there exists a VP that can accommodate the requested VC on the current link }

THEN

the new VC link is created in this VP and stored in the system database

ELSE

A new VP link is created between the nodes of the current network link and stored in the system database

A new VC link is created inside the new VP link and stored in the system database

- 4f *the database is unlocked for the processing of other applications*

ELSE *the connection is rejected*

- 5 *The algorithm's decision data is presented to the user*

Figure 3. The service's decision-making algorithm

ATTENTION:

For the computation of the $\sum_j \text{avail_BW}_j$ we take into account only the avail_BW_j , where $\text{avail_BW}_j \geq$ (requested bandwidth by the current request for virtual connection)

Thus, the link_value_i of each link i of the topology is computed as follows:

$$\text{link_value}_i = \text{Min_BW_reserved} / \text{BW_avail}_i \quad (3)$$

D. Manager Interface

MBS provides the service manager with an interface that allows him to perform the following operations:

- Activate/deactivate virtual connections
- Update/configure the topology of the network
- Configure the user's form

The manager is responsible for the 'manual' reservation of the virtual connections that have been requested and accepted by the service. The reservation of the virtual connections is based on the features of the connection and the contract selected by the user, for each accepted connection. At any moment, the manager has the ability to access the list of the pending reservations and select all or a subset of them in order to activate them, that is to appropriately configure the ATM switches. As the manager implements the reservations at the switches, he has the ability to store detailed information about the implementation of virtual connections, like the VPI/VCI, at the switches that are being used by a virtual connection. This information can be inserted by using the form presented in Figure 4.

Activation of virtual connections

PENDING SESSIONS:

SESSION 13:

Session beginning:
Date 2001-02-18
Time 03:00

Session ending:
Date 2001-02-20
Time 04:00

The session above consists of the following connections:

Connection 13 - Connection data:

Origin	Destination	Bandwidth	Service	Contract	Connection Type
Patras node	Athens node	2.000 Mbps	CER	3rd contract for CER	Virtual Path

Athens node

Input port	Output port	VPI in	VCI in	VPI out	VCI out
1	2	220		230	

Patras node

Input port	Output port	VPI in	VCI in	VPI out	VCI out
2	8	250		240	

Entry of connection data

Submit

Back

Figure 4. The form through which the MBS manager can store detailed information about the implementation of virtual connections

Of course, the service manager can deactivate all expired connections or any connections he wishes to, through specific forms of the manager interface. As it has already been mentioned, the MBS manager can also update/configure the topology of the network, by using certain forms of the MBS manager interface. Additionally, he may modify/delete a node and add/modify/delete a network link. Finally, the manager can configure the user's form, by adding/deleting a new

value to the set of bandwidths provided by the service, a predefined contract that is provided to users, a bandwidth usage type and a service type provided.

V. FUTURE WORK – CONCLUSIONS

Our future work on this, will consider the implementation and experimental use of more sophisticated algorithms by the decision-making module of the system. At the moment, the service is deciding upon the arrival of a request, based on the current availability of the network and trying to route an accepted connection through the most appropriate, least loaded path. In this way, the MBS presented here is distributing the bandwidth reservations throughout the network backbone, thus reducing the possibility for denial of a request in overloaded areas of the network. However, the relevant algorithm is entirely based on the FCFS principle and does not allow for any kind of dynamic reallocation in time or of bandwidth that could be adopted in order for the number of accepted requests to increase.

Apart from that and based on the fact that the tendency for WAN backbones is to adopt mechanisms for bandwidth management on the IP level, we will also investigate the mechanisms to transfer the bandwidth management procedure from ATM oriented methodology to the IP level, by suggesting relevant methodologies in VPN and MPLS technologies. The potentials of a new architecture that places the IP level directly over a WDM infrastructure, with respect to the management of the available bandwidth, will also be investigated as part of our future work.

The MBS presented in this paper can be accessed in <http://gaia.cti.gr/mbs/english/service.php>, strictly for demonstration purposes. We would also like to thank all our colleagues in GUNet and GRNET for their cooperation, assistance, comments and evaluation efforts during the deployment of MBS.

REFERENCES

- [1] G. Ramamurthy and Q. Ren, "Multi-Class Connection Admission Control Policy for High Speed ATM Switches," in Proceedings of the Conference on Computer Communications (IEEE Infocom), Kobe, Japan, April 1997
- [2] D. Wu and H. J. Chao, "Efficient Bandwidth Allocation and Call Admission Control for VBR Service Using UPC Parameters," in Proceedings of the Conference on Computer Communications (IEEE Infocom), New York, USA, March. 1999
- [3] Y.-D. Lin, W.-J. Su, and C.-C. Lo, "Virtual Path Management in ATM Networks", in Proceedings of International Conference on Communications, Dallas, USA, June 1996
- [4] A. Pitsillides, G. Stylianou, C. Pattichis, A. Sekercioglu, and A. Vasilakos, "Bandwidth Allocation for Virtual Paths (BAVP): Investigation of Performance of Classical Constrained and Genetic Algorithm Based Optimisation Techniques", in Proceedings of Infocom 2000, Tel-Aviv, Israel, March 2000
- [5] WiseWan Corporate Brochure, "Multi-Service WAN Platform for Today's Global Networks – Effective Bandwidth Management for Wide Area Networks", <http://www.nreality.com/products/literature.htm>
- [6] FloodGate-1 datasheet, "Secure Virtual Network Architecture", CheckPoint Inc., <http://www.checkpoint.com/products/floodgate-1/>
- [7] CommandView Bandwidth Optimizer, <http://www.tlic.com/connectivity/bandwidthoptimizer.cfm>, Elron Software Inc.
- [8] Guidepost, <http://www.netguard.com>, NetGuard Inc.
- [9] Netroad Trafficware, Ukiah Software Inc., "NDS - Enabled Applications – Enabling Policy-Based Traffic Management", Linda Boyer, August '98, <http://www.nwconnection.com/aug.98/nds/>
- [10] Sun Bandwidth Allocator 1.0 Brochure, Sun Microsystems Inc., "Bandwidth Management for IP Networks, Sun Bandwidth Allocator", White Paper, April 1998, <http://www.sun.com/software/bandwidth/allocator/>
- [11] Solaris Bandwidth Manager 1.6 Brochure, Sun Microsystems Inc., "Solaris Bandwidth Manager 1.5 – Bandwidth Management for IP Networks", White Paper, 1999, <http://www.sun.com/software/solaris/ds/ds-bandwidth/>
- [12] NetEnforcer Datasheet, Allot Communications Inc., June 2000, "Policy-Based Network Architecture" http://www.allot.com/html/products_netenforcer.shtml
- [13] PacketShaper Brochure, Datasheet "Strategic Bandwidth Management, Designed for the ISP", Packeteer, <http://www.packeteer.com/products/packetshaper/index.cfm>
- [14] DANTE, "External Operational Procedures for the TEN-155 Managed Bandwidth Service", 2 August 1999, QUA-99-063 OPS-99-071, Version 1
- [15] "SuperJANET Management Bandwidth Service Terms and Conditions", ND/STG/MBS/DOC/001, Version 2, March 2000
- [16] "SWITCH Managed Bandwidth Service", <http://www.switch.ch/lan/mbs/>
- [17] Roney E. Boyd Jr., "VCCS Network management White Paper", August 12, 1999, <http://www.so.cc.va.us/its/projects/VCCS-NW-MANGT-WP1.htm>
- [18] Laura Cunningham, "QoS Development in the vBNS", in the Communications Magazine, Vol. 36, No. 5, May 1998
- [19] GRNET: The Greek Research and Technology Network, http://grnet.gr/index_en.html
- [20] GUNet: Greek Universities Network, <http://www.gunet.gr/>