

# Simulation of a Campus Backbone Network, a case-study

J. Potemans<sup>1</sup>, J. Theunis, B. Rodiers, B. Van den Broeck, P. Leys, E. Van Lil, A. Van de Capelle

Katholieke Universiteit Leuven (K.U.Leuven)  
Department of Electrical Engineering  
ESAT-TELEMIC division  
Kasteelpark Arenberg 10  
B-3001 Heverlee (Leuven), Belgium  
E-mail: [jan.potemans@esat.kuleuven.ac.be](mailto:jan.potemans@esat.kuleuven.ac.be)

## Abstract

This paper describes the work performed in the framework of a Master's thesis. The aim of the thesis was to make a simulation model in *OPNET Modeler* of the backbone at the K.U.Leuven, Belgium. The backbone consists of five *Cisco Catalyst 5500* switches in a ring topology interconnected by full-duplex Gigabit Ethernet links. Accurate real-life measurements are used to add background traffic to this topology. The use of VLANs to separate traffic generated by students and personnel adds extra difficulties (the traffic is routed at different locations). This paper gives an overview of the results and the problems encountered.

## 1. Introduction

Computer networks have become extremely important in our present-day society. A majority of companies depend on the proper functioning of their networks for communications, administration, automation, e-business solutions etc. Outage of the network can cost a company millions of euros.

To avoid these problems, system managers need professional tools to help them with the design and maintenance of computer networks. A simulation tool offers a way to predict the impact on the network of a hardware upgrade, a change in topology, an increase in traffic load or the use of a new application. *OPNET Technologies* offers a number of tools specifically designed for these purposes.

This paper describes how the *Modeler* tool can be used to make a simulation model for the campus backbone network of the K.U.Leuven, Belgium. It is a summary of the work performed in the framework of a Master's thesis [1].

It is clear that this thesis was not proposed in the perspective described above. Our university won't lose that much money when the network is down. The aim of the thesis is rather a study of the possibilities, problems and limitations of today's simulation tools. This study fits into the doctoral research on hybrid simulation techniques performed in our group. *ESAT-Telemic* is a partner in *OPNET's University Program* and uses the *Modeler* simulation tool both for research [2,3,4] and education [5,6,7].

This paper is organized as follows: the topology of both the backbone network and the Internet access network of the university is discussed in section 2.

Section 3 gives an overview of the traffic measurements on these networks together with the problems encountered. The *OPNET* model of these networks is dealt with in section 4. The integration of the measured data into the *OPNET* topology is discussed in detail. Section 5 describes how this model can be fine-tuned. A VoIP case study is worked out in section 6. Finally, section 7 sums up the conclusions of this paper.

## 2. Topology

The backbone network consists of five *Cisco Catalyst 5500* backbone switches, interconnected by a full-duplex Gigabit Ethernet ring. Figure 1 gives an overview of this ring topology with the five backbone switches: *laswitch-ludit*, *laswitch-cw*, *laswitch-stra*, *laswitch-hh* and *laswitch-esat* (according to their location). All these backbone switches have routing capabilities.

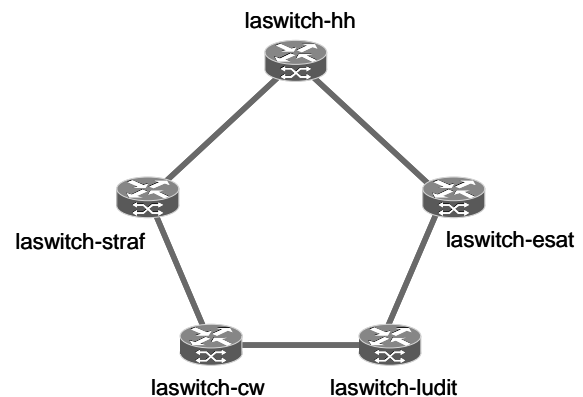


Figure 1: Backbone topology

Several peripheral switches are attached to these backbone switches to offer connectivity to local sites, which can be administrative services, research groups, auditoria, classrooms, student dorms etc. Figure 2 shows how the *Local Area Networks* (LANs) are linked to the backbone.

*Virtual LANs* (VLANs) [8] are used to define logical units to which people of the same administrative service, research group, etc. belong. This makes it possible for users of the same unit to connect to the network from geographically distributed locations (which is important because the university buildings are spread over town).

<sup>1</sup> Research Assistant of the Fund of Scientific Research – Flanders (FWO-Vlaanderen)




Figure 2: Local sites connected to the backbone via peripheral switches

A clear distinction is made between units of university personnel and units of students. Student VLANs (auditoria, classrooms, university student dorms) are trunked over the backbone and centrally routed by the *Kotnet* routers. Personnel VLANs can also be routed by the router modules in the backbone switches.

Most of the peripheral switches are connected to two backbone switches (or to each other) to introduce redundancy in the system. The *Spanning Tree Protocol* (STP) is used to obtain a loop free topology: all redundant links are put in standby. The STP protocol is executed per VLAN.

The distinction between students and personnel also becomes clear in the access network to the Internet. Figure 3 depicts this access network for the *KULeuvenNet*, the network for university personnel. Connection to the Internet is made through two routers: *Cisco-Access* and *Cisco-Kulnet*. Both are *Cisco 7206* routers. A firewall in between protects the network against unauthorized access. A web caching system is provided to relieve the connection to the Internet.

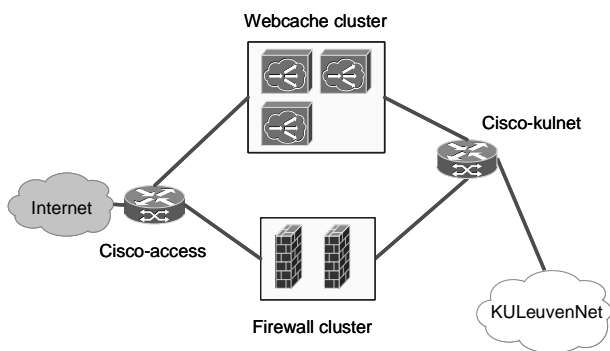


Figure 3: Internet access for KULeuvenNet

Figure 4 shows the access network to the Internet for *Kotnet*, the student network of the university. Instead of the *Kulnet* router, three *Kotnet* routers are used. *Kotnet-1* is used for routing traffic to and from two cable operators (offering access in private dorms). *Kotnet-2* serves all the student dorms of the university. These are *Cisco 7507* routers. Next to routing they also need to control access and collect network usage statistics for every single student, which explains the choice for very powerful routers. *Kotnet-3*, a *Cisco 4700* router, is used for dialin, auditoria and PC classrooms. A *Network Address Translation* (NAT) cluster is used because students are assigned private IP-addresses (by a DHCP server).

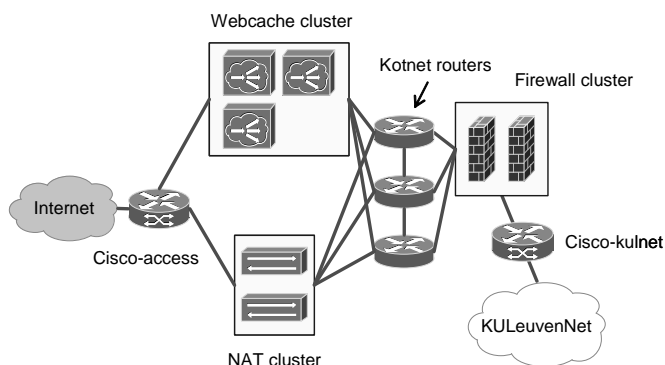


Figure 4: Internet access network for KotNet

### 3. Measurements

The performance of a network heavily depends on the amount of traffic occupying the links and the queues (in routers and switches) throughout the network. To get a clear understanding of the traffic load on the backbone network, accurate measurements are needed. The *Simple Network Management Protocol* (SNMP) [9,10] can be used to get information stored in switches and routers. The *Management Information Base* (MIB) [11] specifies which data is stored in the network equipment. In this project the counters for the incoming and outgoing traffic per interface – both in packets and bytes – are queried using the free *snmpget* tool [12].

The scripting language *Perl* is used to automate the measurements. A first set of scripts queries every two minutes all the interfaces of every router and switch of the backbone and the Internet access network. Four counter values per interface are saved in a text file: the number of incoming packets, the number of incoming bytes, the number of outgoing packets and the number of outgoing bytes. A second set of scripts processes the saved data. For each interval of two minutes, the average number of incoming and outgoing packets and bits per second is calculated.

The major problem is that one could only measure the amount of traffic that entered or exited via a particular interface (the software installed on the backbone switches did not allow to measure inter-VLAN traffic). In most cases it is not clear how the incoming traffic is distributed over the outgoing interfaces.

Figure 5 shows a backbone switch with five interfaces. It is possible to measure the data streams 1 to 10. Assuming that most of the traffic comes from or goes to the backbone, this picture can be simplified as in Figure 6 with only four data streams. The traffic between two backbone switches is simulated using point-to-point traffic between two stations connected to these switches.

There is however a difference between both pictures. In Figure 5 streams 1 and 4 enter the switch, while streams 2 and 3 exit the switch. In Figure 6 all four streams enter (and exit) the switch. If no traffic leaves the backbone ( $1 = 3$  and  $2 = 4$ ) the switch in Figure 6 is loaded twice as much. However if all traffic leaves the backbone, the switch will be loaded correctly.

In reality the load will be somewhere between these two extremes. Therefore fine-tuning of the model will be necessary (see section 5).




Figure 5: Measured traffic streams in a backbone switch

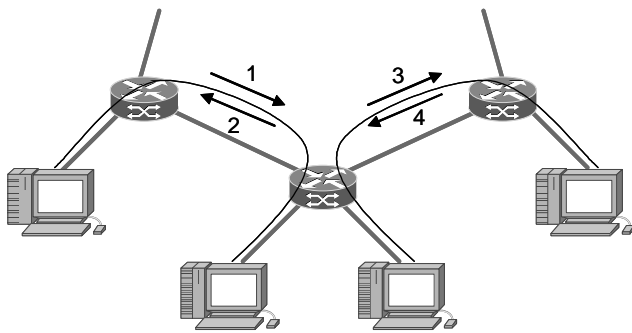


Figure 6: Corresponding data streams in OPNET

There is one case in which the routing is known. *Kotnet* traffic (caused by students) is trunked over the backbone to one of the *Kotnet* routers (and vice-versa). The easiest way to measure the amount of traffic to and from each student site (a student house, an auditorium, ...) is by querying these *Kotnet*-routers. Measuring on the peripheral switches, to which the sites are directly connected, is much more difficult because internal traffic of the site is also visible (and accounted for) on the respective interface of the switch.

#### 4. OPNET model

Figure 7 shows the *OPNET* topology of the backbone in case the *Kotnet* traffic is subtracted from the other traffic and dealt with separately. By giving this traffic a VLAN identifier it is trunked over the backbone and routed in the *Kotnet* routers.

Figure 8 depicts the Internet access network topology, in which the icon “*KULeuvenNet*” represents the topology of Figure 7. Every node comes directly from the component library or is generated using the *device creator* (to have the appropriate number and type of interfaces). For the routers and switches the packet-forwarding rate, the memory size and the capacity of the backplane can be adjusted.

Only for the webcaching infrastructure no appropriate model could be found in the library. Instead of using the available cache-server, a workstation is used in combination with a basic switching device (without specialized features like VLANs). This way all the *kotnet* routers, the *kulnet* router and the *access* router could use the webcaching device. Measurements show that a byte-hit ratio of 26% is attained.

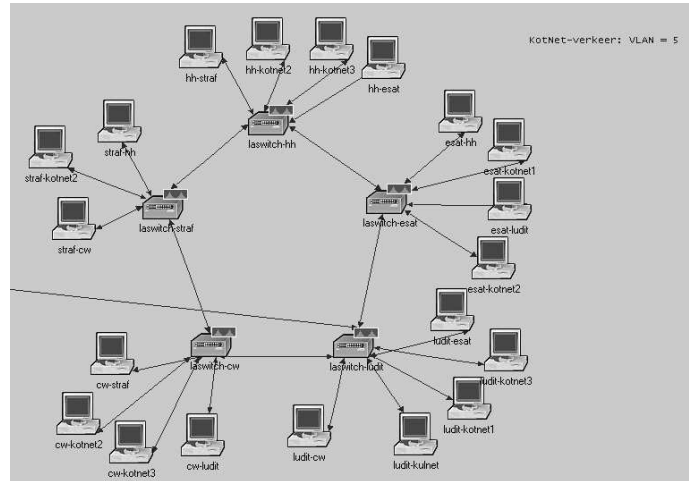


Figure 7: OPNET topology for the backbone network

To add traffic to these topologies, pairs of workstations are connected to the backbone switches: e.g. workstations *esat-hh* and *hh-esat*, connected to *laswitch-esat* and *laswitch-hh* respectively are used to add traffic on the backbone link between these two switches.

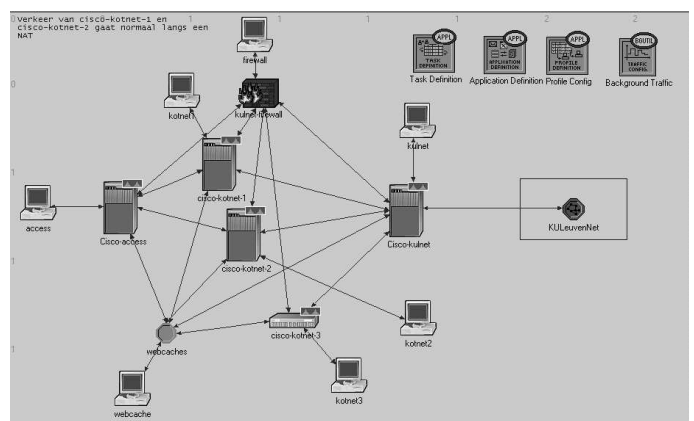


Figure 8: OPNET topology for Internet access

Analogously workstation *esat-kotnet1* exchanges traffic with a workstation attached to the *kotnet-1* router. The webcache workstation is used to represent the traffic coming from and going to the webcache device.

The measured traffic is integrated in the topology as background traffic to enable *hybrid simulations*. This type of simulations is far more efficient than discrete event simulations because only explicit traffic needs to be simulated. The queuing behavior of this traffic is based on the background traffic load on the queues throughout the topology. Background traffic can be inserted using the *conversation pair browser* (see Figure 9).

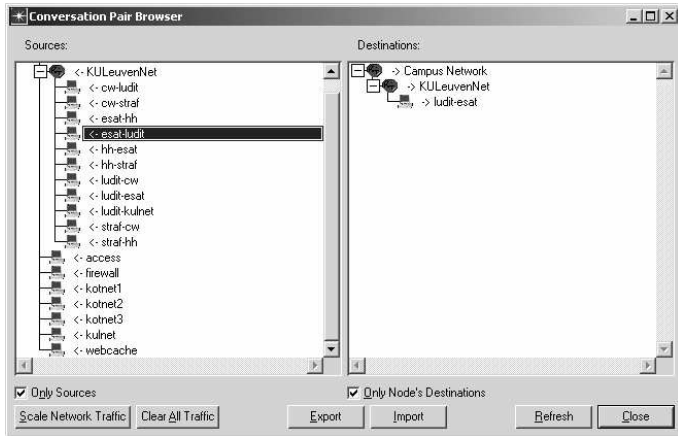


Figure 9: Conversation pair browser

For each couple of workstations one can configure the appropriate background traffic profile. These profiles were imported from a text file, composed by a *Perl* script that transforms the measured data into the correct format. This format was obtained by exporting an empty conversation pair matrix to a text file. Figure 10 shows the imported traffic on the backbone link from *ludit* to *esat*.

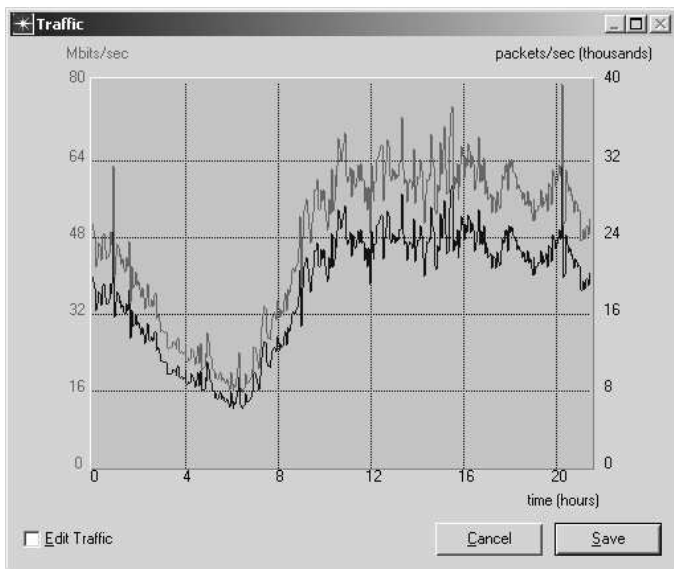


Figure 10: Imported traffic from ludit to esat

The measured data could be easily inserted in the model, except for the *Kotnet* traffic. The specific trunking behavior (caused by the use of a VLAN) was not supported in background.

It was impossible to treat the *Kotnet* traffic separately. The topology of Figure 7 is simplified to the topology of Figure 11 in which the *Kotnet* traffic is added to the remaining workstation pairs. The trunking behavior is lost which makes the model far less accurate (even after fine-tuning).




Figure 11: Backbone topology without separate KotNet traffic

Another problem encountered concerns the *Spanning Tree* protocol. *OPNET* doesn't support the *Per VLAN Spanning Tree* (PVST) protocol yet. This means that always one link of the Gigabit Ethernet ring remains unloaded whereas in reality, there is traffic on every link. To prevent a heavily used link from being put in standby, we marked the least used link as *failed* (see Figure 11).

Analysis of the measured data also showed that the backbone traffic was *self-similar* [13]. This means that the traffic is much more bursty than in traditional models like Poisson. The background traffic configuration did not support self-similar packet arrival. The *Raw Packet Generator* (RPG) [3] model, which can be used to generate self-similar traffic, only works for explicit traffic. Due to the division of the traffic in small intervals of 2 minutes, a Poisson approximation (exponential interarrival time variability) seems reasonable.

The packet size variability should also be specified in the background traffic configuration utility. A packet size probability distribution function could not be extracted out of the SNMP measurements. By dividing the number of bytes/sec by the number of packets/sec, one can only calculate the average packet size in a time interval. When averaging out over the complete time range, a single (needed by *OPNET*) mean packet size is obtained which is entered in the configuration box together with a constant distribution as approximation.

## 5. Fine-tuning of the model

Due to the unknown routing behavior of the traffic (explained in section 3), the load of the routers and the switches in the topology is no longer correct. This problem was compensated for by adjusting the packet-forwarding rate of these devices. Note that this scaling is only an approximation.

During one day of measurements, the *round trip times* (RTT) of some *ping* (ICMP) packets were determined every hour. These packets originated from a workstation connected to *laswitch-esat* and had one of the other switches or routers as destination.

The fine-tuning of the model consisted of simulating the same ping application in *OPNET* and adjusting the packet-forwarding rates of the switches and routers until the round trip times matched reality. This procedure was performed for each node separately. Figure 12 shows how ping packets are sent from *esat* to a ping-server at *cw*.

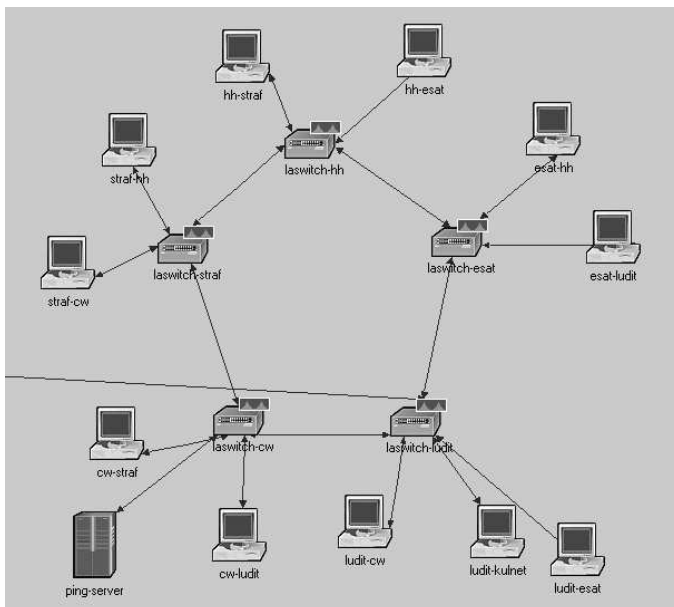


Figure 12: A ping-server connected at cw

*OPNET* has no built-in ICMP application. However, a custom application was built very quickly.

## 6. Case-study: Voice over IP

To illustrate the benefits of the designed simulation model, a more detailed case-study was performed. The aim is to determine whether the backbone infrastructure is powerful enough to integrate the phone traffic with the current data traffic [14, 15].

Figure 13 shows the modified topology for the integration of voice traffic. A LAN representing VoIP users, is connected to each backbone switch. A gateway to the PSTN is represented by a switch and connected to the backbone switch at *ludit*. The *voice\_extern* LAN represents people calling to the university while the *voice-extern* server is used as destination for outgoing calls.

The predefined IP Telephony application is used to configure the calling profiles. The ITU G.729 standard with *Voice Activity Detection* is used as codec. The distribution of the call duration together with the average number of calls per second are derived from call records during a peak hour. These numbers are used to adjust the application and profile settings. It is important to mention that the number of VoIP users is taken arbitrarily but that the profiles are adjusted to generate the total number of calls measured.




Figure 13: Voice users added to the topology

The quality of a VoIP call depends on the end-to-end delay, the jitter and the packet loss. Simulations show that the average (after an initialisation period) network delay for the VoIP packets is only 0.474 milliseconds which can be neglected. Also the jitter is very limited and there is even no packet loss. One can conclude that the integration of voice traffic will cause no problems for the backbone network. If problems are encountered they will most probably originate from the access network.

## 7. Conclusion

This paper showed how the *OPNET Modeler* tool can be used for the simulation of a backbone network. The Campus backbone network of the K.U.Leuven, Belgium, served as a case-study for this paper.

It turned out that not all information necessary could be measured by the software on the switches and routers. This lack of information (about routing behavior) strongly compromises the accuracy of the simulation model.

Furthermore, the *Modeler* software did not allow to include all features present in the backbone network. It was impossible to allocate VLANs and self-similar behavior to background traffic. The *Per VLAN Spanning Tree* protocol was not supported.

However, if all these problems were to be overcome, the *OPNET Modeler* simulation environment would offer a very useful tool for the design and management of our backbone network.

## References

- [1] B. Rodiers, "Analysis & Simulation of the K.U.Leuven-network" (in Dutch), thesis to obtain the degree of Master in Electronics Engineering, K.U.Leuven, Belgium, June 2002.
- [2] B. Van den Broeck, P. Leys, J. Potemans, J. Theunis, E. Van Lil and A. Van de Capelle, "Validation of router models in OPNET", OPNETWORK 2002, Washington D.C., USA, August 2002.
- [3] P. Leys, J. Potemans, B. Van den Broeck, J. Theunis, E. Van Lil and A. Van de Capelle, "Use of the Raw Packet Generator in OPNET", OPNETWORK 2002, Washington D.C., USA, August 2002.
- [4] M. Teughels, E. Van Lil, and A. Van de Capelle, "Backbone Network Simulation: a Self-Similar Perpetuum Mobile", OPNETWORK 1999, Washington D.C., USA, August 1999.
- [5] J. Theunis, B. Van den Broeck, P. Leys, J. Potemans, E. Van Lil and A. Van de Capelle, "OPNET in Advanced Networking Education", OPNETWORK 2002, Washington D.C., USA, August 2002.
- [6] J. Potemans, J. Theunis, M. Teughels, E. Van Lil and A. Van de Capelle, "Student Network Design Projects using OPNET", OPNETWORK 2001, Washington D.C., USA, August 2001.
- [7] J. Theunis, J. Potemans, M. Teughels, A. Van de Capelle and E. Van Lil, "Project Driven Graduate Network Education", Proc. to the International Conference on Networking ICN'01, Colmar, France, pp.790-802, 2001.
- [8] S. McQuerry, "Interconnecting Cisco Network Devices" (Chapter 6: Extending Switched Networks with Virtual LANs), Cisco Press, 2000.
- [9] J.D. Case, M. Fedor, M.L. Schoffstall, C. Davin, "Simple Network Management Protocol (SNMP)", RFC 1157, May 1990.
- [10] J. Case, K. McCloghrie, M. Rose, S. Waldbusser, "Introduction to version 2 of the Internet-standard Network Management Framework", RFC 1441, April 1993.
- [11] K. McCloghrie, D. Perkins, J. Schoenwaelder, "Structure of Management Information Version 2 (SMIv2)", RFC 2578, April 1990.
- [12] NET-SNMP home page: <http://net-snmp.sourceforge.net>
- [13] W. Leland, M. Taqqu, W. Willinger, D. Wilson, "On the Self-Similar Nature of Ethernet Traffic (Extended Version)", IEEE/ACM Transactions on Networking, vol.2, no.1, pp. 1-15, February 1994.
- [14] J. Bankston et al., "Cisco AVVID and IP Telephony Design and Implementation", Sygness Publishing, 2001.
- [15] R. Caputo, "Cisco Packetized Voice & Data Integration", McGraw-Hill Professional Publishing, 1999.