

Management of Networks that provide QoS Guarantees

Ron Sprenkels, Bram van der Waaij, Bert-Jan van Beijnum, Aiko Pras

Centre for Telematics and Information Technology (CTIT), University of Twente

PO Box 217, 7500 AE Enschede, The Netherlands

sprenkel@cs.utwente.nl - tel.: +31-53-4894663 - fax: +31-53-4894524

Abstract

This paper presents the results of a case study to the feasibility of introducing ATM SVCs into the Dutch SURFnet research ATM network. The key issue that is examined are the implications of the Quality of Service support of ATM. QoS guarantees for a connection require a portion of the finite ATM network resource. Once all network resource is allocated to connections no new connections will be accepted, and users will start experiencing denial of service. The key research question here is if and how this denial of service probability can be kept to a minimum.

Keywords: case-study, ATM, SVCs, SURFnet, CAC, Quality of Service, denial of service.

1 Introduction

SURFnet bv is the organisation responsible for running the research network between the Dutch universities. To enable the introduction of advanced multi-media applications like video conferencing and tele-education, SURFnet investigates the introduction of new network technologies. At this moment, one of the most promising technologies to support multi-media applications, is Asynchronous Transfer Mode (ATM). ATM allows users to set up Switched Virtual Circuits (SVCs) that have a guaranteed Quality of Service (QoS) and is therefore perfectly suited to support multi-media applications.

To deliver the guaranteed QoS for a connection, an ATM network allocates network resources (like link bandwidth and cell buffer space) for that connection. Current ATM networks allocate the scarce network resource to incoming connection requests on a first-come first-served basis. This can result in an unwanted distribution of network resource over connections, as illustrated by the following example. On a university ATM network, a student starts a video conference to discuss last night's soccer match with a friend. Next a teacher tries to start a video conference for a remote guest lecture; this fails due to lack of resource. The teacher experiences the failure to establish an ATM connection as denial of service. Denial of service due to an unwanted distribution of network resource is therefore an important problem to be addressed when deploying network technologies like ATM that provide QoS guarantees.

This paper presents the results of a study we have performed in 1997 for SURFnet to investigate how problems like service denial can be managed, and to determine whether the time is ripe to introduce ATM SVCs. Because the SURFnet ATM network consists of equipment from multiple vendors, the emphasis of the study had to be on standardized solutions.

The approach taken in the study is to address the problem from three different perspectives (Figure 1):

- Is it possible to implement different admission policies for different types of users? Such policies must be enforced in the connection setup phase to prevent that certain (categories of) users establish connections at less desirable moments, for example at working hours or at times the network is highly loaded. Section 2 discusses this question.
- Is it possible to interfere with existing connections, for example to reduce the resource allocation for a connection or to abort a connection? Section 3 discusses this question.
- Is it possible to take measures “after the fact” that make denial of service less likely? A common strategy to achieve this is to call users to account. Section 4 therefore discusses current options to do ATM-SVC accounting.

The conclusions of this feasibility study are presented in Section 5.

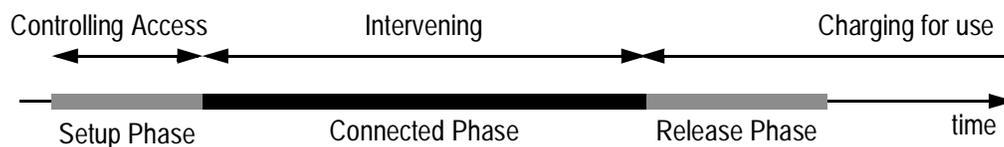


Figure 1: The Life-cycle of an ATM SVC connection

2 Controlling Access in the Setup Phase

Denial of service occurs when the network does not have sufficient resource left to support the new connection. The first opportunity to minimize denial of service due to lack of resource is therefore to control the user’s access to network resource in the connection setup phase. This controlled access should result in a good distribution of the scarce network resource over connections.

When investigating capabilities of some of today’s ATM switches we found no switches that are able to perform such controlled access to network resource. We found switches to have only a CAC function. The CAC function accepts a new connection if the required resource for it is currently available, irrespective of the amount of requested resource or of the requesting user. So with only CAC, any user can claim all remaining network resource at any time. If that happens, all subsequent connection requests will fail. This shows that CAC alone is not suffi-

cient to prevent unwanted distributions of network resource over connections from occurring. We therefore investigated a solution to this problem.

First possible criteria for a connection admission policy will be discussed in section 2.1. Next an approach to verify for each request if it complies with the admission policy is discussed in section 2.2.

2.1 Criteria for a connection admission policy

We have identified at least four criteria on which to base connection admission decisions. Each of these criteria is illustrated by an example.

- **Individual Users**

The maximum amount of network resource a user can claim for a connection can be specified. This can be done per individual user or per group of users. This allows users to get a different maximum QoS from the network. Example: users in the “economy” class can use at most 1 megabits per second (Mbps) per connection, users in the “premium” class can use up to 3 Mbps.

- **Remaining Credits**

A user or group of users can be allowed to use a limited amount of resource per measurement period. Example: A user is allowed 10 hours of 1 Mbps connections per week.

- **Network Status**

Policies can depend on the current overall network status. Example: “economy” users can not set up new connections when the network is over 60% loaded.

- **Current Time**

Policies can depend on the current time of day. Example: “economy” users can use up to 1 Mbps connections during business hours, and up to 3 Mbps connections in other hours.

Enforcement of an admission policy based on these criteria results in users experiencing different QoS from the network. This should be reflected in the Service Level Agreement between users and provider of the ATM network service.

2.2 An Approach to Enforcing a Connection Admission Policy

This section proposes an approach to verifying the connection admission policy for each connection request. An admission policy server is introduced into the network for this purpose.

A connection request arrives over a User Network Interface (UNI) at a switch in the ATM network. This switch performs its normal CAC function to see if sufficient resource is available to accept the connection. It must also be verified if the request meets all the criteria prescribed by the connection admission policy. The signalling system of the switch should therefore forward the request to the central admission policy server in the network. This server verifies if the request meets all criteria, and either accepts or rejects the request. The server sends the outcome of the verification back to the switch. If the switch has sufficient resource and the admission policy server agrees then the connection is set up.

The admission policy server needs to have access to all the information needed to verify the criteria identified in section 2.1. This is shown in Figure 2 below.

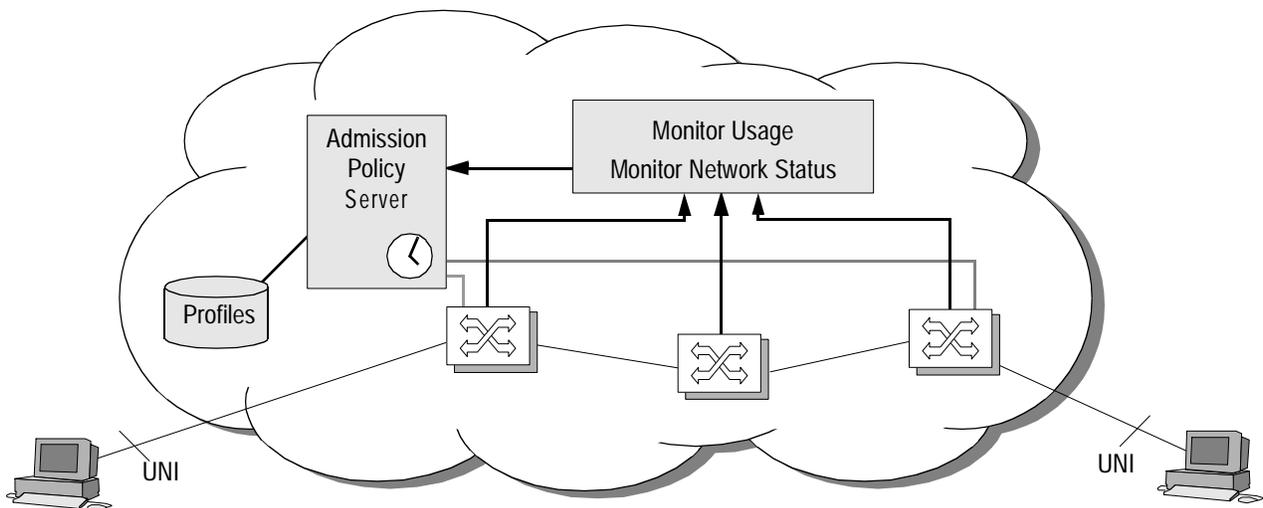


Figure 2: ATM network with Admission Policy function

The admission policy server has access to a database of user profiles. Such a profile contains for each user the information needed to check the admission policy criteria. Examples are the maximum allowed bandwidth per connection, remaining usage credits and allowed times of day for connections.

The network will have a monitoring process to gather per-user information on usage of the network resource. This information is available to the admission policy server for admission criteria that depend on how much use a user has made of the network in the past. Such criteria in fact make use of accounting information (discussed in section 4) for connection admission decisions.

Another process monitors the status of the network, e.g. the load percentage of the network. The total amount of bandwidth reservations on each link or the average fill level of the switch cell buffers could be used to calculate a value for the network load percentage.

Finally the admission policy server will have access to the current time, to allow for policies that depend on the time of day, day of the week etc.

We did this work on connection admission policies in the summer of 1997 [8]. Currently the RSVP Admission Policy group of the IETF is doing similar work on this subject (admission control framework [11], COPS [3] and OOPS [7]).

3 Intervening During the Connected Phase

As part of our work for SURFnet, we investigated the possibility to abort existing connections. A human network manager might want to abort an unwanted existing connection or reduce the amount of resource allocated for it, to make network resource available in favour of a different connection. This may be necessary in cases where controlled access to network resource (section 2) is not possible, and accounting (to be discussed in section 4) is either not possible or does not have the desired effect. Aborting connections may have far-reaching consequences; the provider should therefore discuss the possible use of this option with its customers as part of the Service Level Agreements (SLAs).

3.1 IETF MIB Support for Intervening

Based on the MIBs defined by the IETF, this section explains how managers should be able to abort existing SVCs. The MIBs that can be used for this purpose are the 'AToM-MIB' [1] and the 'Supplemental MIB' [2]. The AToM-MIB defines general ATM management information and has currently the status of proposed standard; the supplemental MIB defines information specific for SVCs and has currently the status of working draft.

Whenever a connection is set up by the signalling system, a row is created in the *Cross-connect table* of each switch on the path between the source and destination. The row contains, amongst others, the Interface number (IF), Virtual Path Identifier (VPI) and Virtual Channel Identifier (VCI) of the incoming as well as outgoing link, but also an index value and a row status (Figure 3). The ATM addresses that belong to the connection can be found by using the row's {IF, VPI, VCI} triples as pointers to the *Address table*. The resource reserved for this connection can be found by using the row's index parameter as pointer to the associated entries in the *Traffic parameter table* (this

mapping uses another table, which is not shown in the figure). The connection can be aborted by setting the row status parameter to *destroy*, which removes the row from the cross connect table.

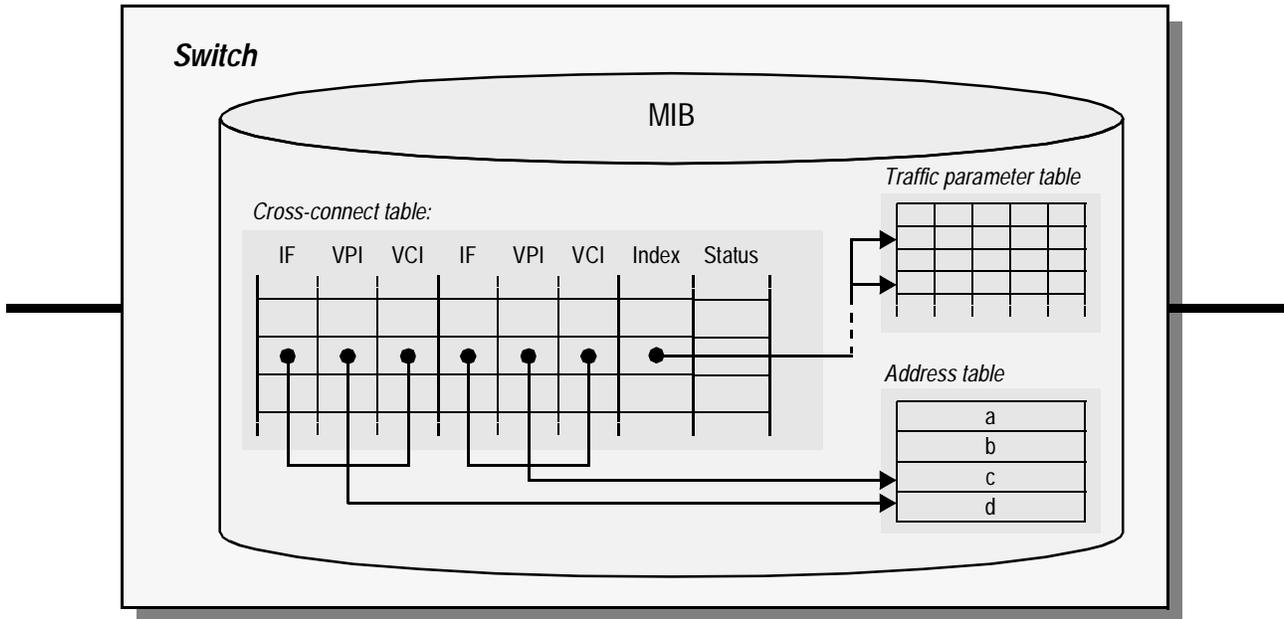


Figure 3: Aborting connections, based on IETF MIBs

After the manager removes a row in one of the switches along the path, the signalling system should inform the other switches along the path that the connection has been aborted and, as a result, all resources that have been reserved for the connection should be released.

As part of our research for SURFnet, we wanted to test this theory on commercially available switches. Since these switches do not (yet) implement the supplemental MIB, we had to use proprietary MIBs for these tests. Unfortunately the implementation of these MIBs contained a number of bugs that prevented us from performing the tests.

4 Charging for Network Use

In general a good approach to discourage the use of a scarce resource is to make users pay for it. Paying for the use of network resource will encourage them to set up only those connections they really need, and release them immediately after use. The mechanism that is needed to let users pay for their connections is called accounting, and is generally divided into three phases [5].

- The first phase is to measure, at specific locations in the network, the use of network resource for each connection; this is called usage metering (section 4.1).
- The second phase is to collect all per-provider usage metering information for one connection and calculate a price for it; this is called charging (section 4.2).

- The third phase is to collect all the charging information for one user and to put it on a bill. This bill is sent to whoever has to pay for the connections set up by this user; this is called billing (section 4.3)

4.1 Usage metering

The usage metering process is responsible for the collection and storage of usage information per user per connection. This information is stored in a usage record. Usage metering shall in most cases be performed on the borders of the provider's network. Because the service is defined on those borders, this ensures that the obtained metering information will be in complete agreement with the delivered service. A neighbour at the other side of the network border can either be a user or another provider.

In the connection setup phase the setup request passes a number of switches between the end points of the connection. Each of those switches that does usage metering will create a separate usage record for the connection.

Initially each usage record contains:

- the precise location where the usage information was metered;
- the time the connection was created;
- information describing the connection itself, like the end points involved in the connection and the traffic contract.

During the lifetime of the connection, usage information (e.g. amount of transported data) will be collected and also added to the record. When the connection finally is released the last usage information and the time of connection release are added to the record. Now the usage metering process for this connection is finished and the usage record is ready for further processing (see section 4.2).

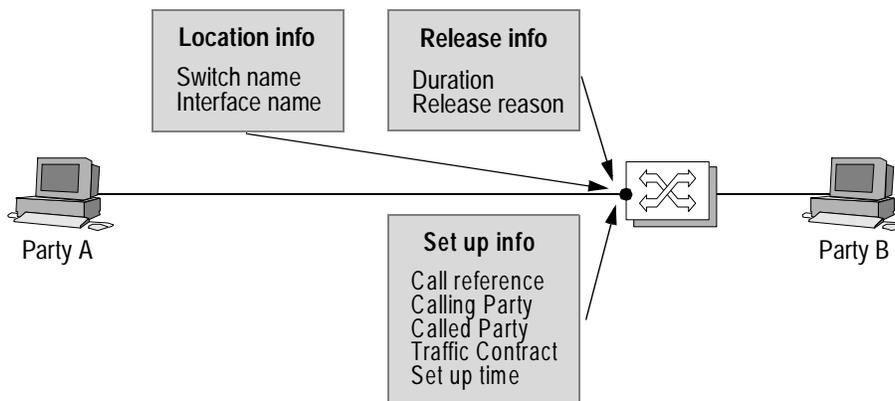


Figure 4: Usage Metering for a Connection

Because SURFnet operates a multi vendor ATM network we had to look at a standard solution for usage metering. The IETF is one of the organisations that have drafted a proposal for the necessary parameters for a usage record [9]. Below we give an overview of the most important parameters in this proposal (see Figure 4):

- The metering location

To be able to determine the location on the network border where the usage record is created, the name of ATM switch must be preserved. The name of the interface is necessary to pin-point the exact location within the switch.

- Connection setup information

The connection can be uniquely identified by the calling party, the called party and the call reference value which is coming from Q.2931 [10]. The traffic contract describes the connection itself. The set up time specifies the moment the connection was created, which is useful in the charging phase to make a distinction between different times of the day.

- Connection release information

When the connection is released, the duration - or life time - of the connection can be determined. The release cause is also useful, because a user would probably want to be charged less for a SVC terminated due to a network error, than for one released by himself.

At the moment there are no implementations of this IETF proposal, but a number of vendors has propriety solutions for usage metering [15][16][17][18].

4.2 Charging

The purpose of charging is to calculate a price for the use of network resource, and to generate charging records for that. We have split this process into three sections:

- The installation costs

A charging record for the installation costs is created once, when the user first subscribes to the ATM network service. The user is charged for the fact that from now on he has the possibility to use the network.

- The subscription costs

A charging record for these costs is issued periodically, and is used to charge the user for having access to the network during this subscription period.

- The per-call costs

The per-call costs can be calculated in two different ways:

- A flat rate charging scheme can be used. In a charging period each user pays the same fixed price each billing period, corresponding to the average user's use of the network. On balance this scheme reduces the per-call costs to zero; the height of the bill will be independent of how much a user has actually used the network. As such in a flat rate charging scheme there won't be a charging record per call, and the call related costs can be considered a component of the subscription costs. Note that in this case no metering records per call are needed either.
- The actual use of network resource of each individual call can be charged. A call can consist of multiple connections, so the provider first has to collect all usage metering records for a call. Based on these usage metering records the price of the call is calculated. This can e.g. be a fixed price per call (reflecting the average costs per call), or it can be determined by looking at (a combination of) the duration of the call, the transmitted and received volume of data, the amount of resource that was reserved for the call, time of day the call occurred, etc.

At the moment there are two ACTS [12] projects working on this topic: Contract Negotiation and Charging in ATM Networks (CANCAN, [13]) and Charging and Accounting Schemes in Multi-Service ATM Networks (CA\$hMAN, [14]).

4.3 Billing

The billing process is responsible for creating a bill for each user. In order to do this, a list must be made of all charging records belonging to one user. From this list a bill must be constructed for each user, specifying the installation costs, subscription costs and the costs for all the calls the user has made.

When all the connections of all calls of a user stay within the boundaries of the user's own provider, charging records will only be generated within the domain of the provider. In this case the provider has all the charging information needed to send his subscriber the bill for his connections.

When a connection crosses one or more provider boundaries, the billing process becomes more complex. For a single call, costs (and thus charging records) are generated within different providers, and there are now different ways in which the user can finally receive the bill for all of those costs. We have identified four options to deal with this problem, but other solutions are also possible [4]:

- Settlement with the predecessor

A provider can send the bill for its part in the call to the provider the call is directly coming from. The originating provider then finally charges the user for the complete call.

- Settlement with the originating provider

Another solution is to let each provider send his bill directly to the originating provider. The originating provider collects all bills and charges the user for the complete call.

- Settlement with the originating user

A third solution is to let each provider, including the originating provider, send their bill directly to the user.

- Central billing system

And finally it is also possible to establish a central billing system that collects all the bills from the providers and combines them into one bill the user.

5 Conclusions

This paper discussed the state of the art in ATM SVC management. It focused on the problem of service denial, which occurs in case the requested network resource exceeds the available resource. As shown in Section 2, existing Connection Admission Control (CAC) mechanisms are able to detect such situations, but are unable to resolve such situations in ways that critical users get precedence over others. Preventing 'low priority' users from (deliberately) taking all available resource is therefore not possible. Given this impossibility, the question arises whether it is possible to abort such connections. As shown in Section 3, aborting connections is possible, but cumbersome. Finally Section 4 discussed a possible way to reduce the amount of requested network resource by calling users to account. Accounting can be divided into three different aspects: usage metering, charging and billing. For ATM usage metering a number of MIB standards are being developed by the IETF. Although these standards have not yet been implemented in existing switches, vendor specific MIBs with similar functionality have already been implemented. Usage metering is therefore possible, although not in a vendor independent way. Standards for ATM charging and billing do not yet exist.

Incomplete SVC management, in particular the impossibility to prevent users from (deliberately) capturing all network resource, was one of the important reasons for SURFnet to postpone the introduction of ATM SVCs. Although SURFnet continues to use PVCs, it currently investigates alternatives to SVCs, such as IPv6 in combination with RSVP. Depending on the outcome of this investigation, it is not unlikely that SVCs will never be introduced.

IPv6 and RSVP may face a similar fate as SVCs, in the sense that they too guarantee a certain QoS by reserving network resource. The same denial of service problem therefore exist with IPv6 and RSVP. In fact every network that provides QoS guarantees faces this problem. Apparently the IETF has learnt from the problems with managing SVCs, since one of the IETF's working groups has recently started to address policy issues (rapwg) and work is in progress on defining a framework for policy control [11] and policy service protocols (COPS [3] and OOPS [7]). Bringing forward standards and implementations for policy management will be a key issue for the success of these new Internet protocols!

Acknowledgements

The research reported on in this paper has been performed for SURFnet bv who also funded it. In particular we would like to thank Victor Reijns, Niels den Otter and Ronald van der Pol for bringing their expertise into this research project and for the comments and suggestions they gave on the material presented in this report.

6 References

- [1] M. Ahmed, K. Tesink, *Definitions of Managed Objects for ATM Management Version 8.0 using SMIPv2*, RFC 1695
- [2] F. Ly, M. Noto, A. Smith, E.M. Spiegel, K. Tesink, *Definitions of Supplemental Managed Objects for ATM Management*, IETF Internet-Draft, <http://www.ietf.org/html.charters/atommib-charter.html>
- [3] J. Boyle, R. Cohen et al., *The COPS (Common Open Policy Service) Protocol*, IETF Internet-Draft, <http://www.ietf.org/html.charters/rap-charter.html>
- [4] A. van Dijk, *ATM Accounting Management*, University of Twente, M.Sc. Thesis, <ftp://ftp.cs.utwente.nl/pub/src/snmp/ut-thesis/dijk.ps>
- [5] A. van Dijk, *ATM Accounting Management website* (<http://wwwsnmp.cs.utwente.nl/~nm/projects/ut-atm/accounting/>).
- [6] W. Dijkstra, *Multi-Layered Connection Admission Control for ATM Networks*, University of Twente / British Telecom Labs, August 1997 (<http://www.snmp.cs.utwente.nl/~nm/assign/>).
- [7] S. Herzog, D. Pendarakis et al., *Open Outsourcing Policy Service (OOPS) for RSVP*, IETF Internet-Draft, <http://www.ietf.org/html.charters/rsvp-charter.html>

- [8] R. Sprenkels, B. van der Waaij, B.J. van Beijnum, A. Pras, *Results of the SURFnet4 management project 1997*, CTIT Technical Report Series No. 98-08, Centre for Telematics and Information Technology, University of Twente, the Netherlands, 1998.
- [9] K. McCloghrie, J. Heinanen, W. Greene, A. Prasad, “*Accounting Information for ATM Networks*”, IETF Internet-Draft, <http://www.ietf.org/html.charters/atommib-charter.html>
- [10] ITU-T, Recommendation Q2931: “*B-ISDN Application Protocols for Access Signalling*”, February 1995.
- [11] R. Yavatkar, D. Pendarakis et al., *A Framework for Policy-based Admission Control*, IETF Internet-Draft, <http://www.ietf.org/html.charters/rap-charter.html>
- [12] ACTS: *Advanced Communication Technologies & Services*, European Union, <http://www.infowin.org/ACTS/>
- [13] CANCAN (AC014): *Contract Negotiation and Charging in ATM Networks*, ACTS Project, <http://www.teltec.dcu.ie/cancan/>
- [14] CA\$hMAN (AC039): *Charging and Accounting Schemes in Multi-Service ATM Networks*, ACTS Project, <http://www.isoft.intranet.gr/cashman/>
- [15] Fore, <Manu0147-02>, *ForeView Network Management User’s Manual*, Chapter 14, April 1997
- [16] Cisco, *LightStream 1010 ATM Switch Software Configuration Guide 1997*
- [17] NewBridge, *CrossKeys, KeyBill management software*, <http://www.crosskeys.com/cross/services/keybill/summary.html>
- [18] General DataCom (GDC), *APEX Accounting Management System*, <http://www.gdc.com/inotes/pdf/ams.pdf>