On the Usage of Virtual Paths, Virtual Channels, and Buffers in ATM Traffic Management

Åke Arvidsson*

Abstract

Traffic in ATM networks can be described by numerous parameters. On a per session basis, one may use peak rate, average rate, sustainable rate, average burst duration, average silence duration, and others. In a longer time scale, parameters like the average and peakedness of connection request interarrival times, the average and variance of session holding times, and so on are proposed.

Theoretically, users should provide the former parameters and network operators the latter. In reality, however, few users can be expected to provide all sorts of statistical information about their traffic in advance, and operators do not have enough experience to prepare traffic forecasts for new services and applications. Moreover, even if the information could be provided, the lack of simple yet valid traffic models for ATM networks means that it is far from clear how such information should be used.

Realising that ATM networks, which are already being built, at least for the next few years will have to operate under these uncertain conditions, we propose a robust and forgiving network design and traffic management strategy. The idea is to use only little information about offered traffics and then dynamically control resource allocations, so as to provide acceptable quality of service combined with high utilisation.

The network design is based on the idea of keeping congestion at the edges of the network, so that the operating areas of fast congestion control mechanisms are minimised. Traffic variations are characterised by a model of six layers, each of which is assigned a corresponding layer in a traffic management model. We define the functionalities of each level in the latter modell and review possible implementations. In particular, we report on an implementation of two of these layers and discuss in some detail how a third one could be added.

1 Introduction

Public switched telephone networks (PSTNs), including integrated services digital networks (ISDNs), are based on the synchronous transfer mode (STM). From the point of view of network dimensioning and traffic management, STM networks are fully characterised by the number of circuits between all origin/destination pairs (OD pairs), a traffic by its origin, destination and amount in Erlangs, and the quality of service (QoS) by the probability of rejecting a connection request due to lack of idle circuits. Solutions to classical network design problems often rely on satisfying QoS under the so called busy hours. The associated traffic management problem with respect to QoS refers to more or less complicated overflow facilities to improve QoS, or to alleviate temporary situations of congestion by rerouting overflowing requests.

However, things are changing with the introduction of the broadband ISDN (BISDN), and the asynchronous transfer mode (ATM). In ATM networks there are no explicit circuits, but only raw transmission capacities. Traffic metrics refer to various statistical properties of connection requests, holding times, and information generation within sessions such as peak cell rate, average cell rate, *etc.* Finally, QoS is defined not only as a probability of rejecting

*Postal address: Department of Telecommunication and Mathematics, University of Karlskrona/Ronneby, S-371 79 Karlskrona, Sweden. Telephone: +46 455 780 53. Facsimile: +46 455 780 57. Email: akear@itm.hk-r.se.

a connection request, but also in as the service offered to established connections in terms of information cell loss and delay. Consequently, network dimensioning and traffic management have become more complex issues and old methods are no longer sufficient for solving them.

As for new methods, two conceptually different approaches can be taken: The first one is to develop detailed models of traffics and networks, which at least theoretically permit analysing the network or parts of the network with a high degree of precision. The second approach is to use simpler models and worst case approximations, from which only more simple network analyses can be made.

The difficulties with the first approach are that detailed models do not exist and that they, if they existed, would require equally detailed information on traffic characteristics. A further point is that even if models existed and if detailed and accurate information somehow was obtainable, it is far from trivial to communicate large amounts of information with and within the network in real time, and to process it for real time traffic management purposes.

We believe that the most convenient way to proceed, not to say the only possible one, is to accept that traffic models and parameter information are limited and uncertain. Consequently, networks must be dimensioned and traffic managed in the presence of uncertain data. A key issue is then real time control of network resources and traffic flows to maintain acceptable network performance. It is the aim of this paper to outline such a traffic management strategy. The strategy is based on simple traffic descriptors and simple real time traffic management controls.

2 Networks

2.1 Nodes and Links

Our physical network structure, figure 1, is based on two types of nodes: ATM switching systems (ASSs) and ATM cross connects (ACCs). The subscribers are connected to the ASSs, which have full switching capabilities for establishing and clearing connections, charging, and cell switching on the whole address field, *i.e.* both the virtual path identifier (VPI) and the virtual channel identifier (VCI). ACCs are cheaper and much simpler systems, with no processing and cell switching only on the VPI. ASSs and ACCs may be located at the same physical place, but we still consider them as two distinct nodes. To give an idea of the number of nodes in a real network, a nation wide network in Sweden would have some 100–300 ASSs and 10–30 ACCs.

The ACCs are interconnected by a sparse network of high capacity transmission links. We refer to this part of the network as the *transit network*. The physical transit network is enhanced by establishing a mesh network of direct logical links between the ACCs. A logical link is a an abstract entity consisting of capacity reservations on a series of incident physical links and cross connected through intermediate ACCs, so the two end point ACCs are interconnected by a direct (logical) path. ASSs are not interconnected, but can only access the transit network through links to the ACCs. For reasons of reliability, each ASS is connected to two different ACCs through two distinct links. This part of the network is called the *access network*. Finally the links between the users and the ASSs, which only rarely are doubled, form the *subscriber network*.

We also mention that the network we refer to as a “physical network” may in fact be a logical transport network based on the synchronous digital hierarchy (SDH) or the slotted envelope network (SONET). This allows for protection switching and network reconfiguration on two levels, as both SDH/SONET VPs and ATM VPs can be reconfigured. We propose to use the former reconfiguration for what essentially are fault situations, and the latter ones for what can be classified as traffic management purposes. Faults and protection switching are not considered further herein.

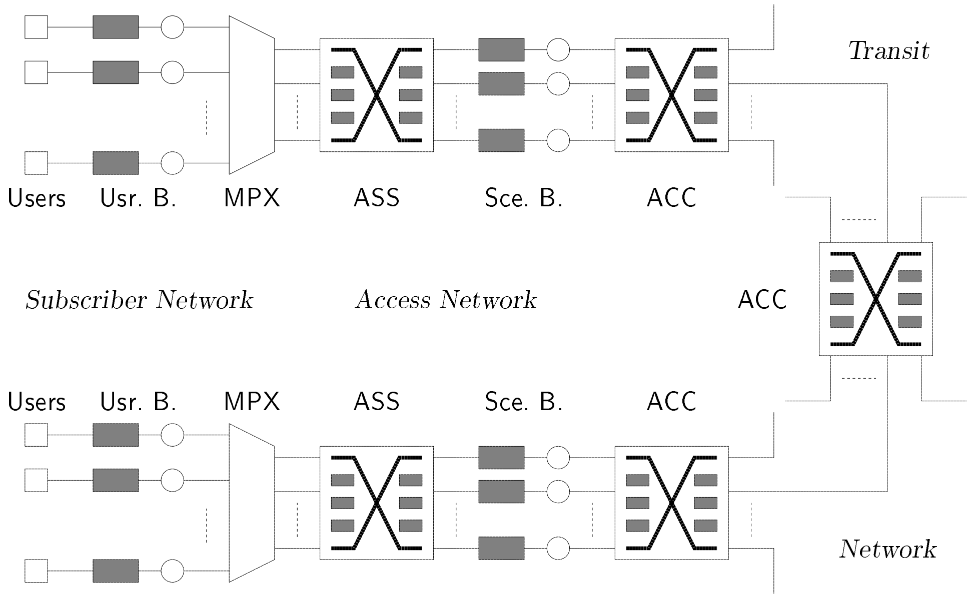

Figure 1: *Definition of physical network structure.*

2.2 Buffers

Users might have to buffer data in order to produce a traffic which is in accordance with a negotiated contract. We call these buffers user buffers, "Usr. B." in the figure, and they fall beyond the scope of the present work.

Within networks, buffering is required to to smooth out (short) periods of overload to avoid cell losses. The requirements of these buffers often vary from one type of service to another: For delay sensitive traffic such as voice, buffers must be kept short in order not cause unwanted delays. On the other hand, for loss sensitive traffic such as data, buffers must be made long to avoid unwanted losses. Obviously these demands are contradictory, and there are two possible ways out: Either to use separate buffers for different services, or to go for a compromise such as some sort of medium long buffer.

We use separate buffers on the end points of the networks and common buffers within the network. Separate buffers are more complicated to handle and longer buffers might have to rely on low speed technology, hence such buffers are useful only when transmission speeds are not too high. This is typically the case at the edges of the network, *e.g.* in the access network, where traffic volumes and the levels of multiplexing are not too high. Such buffers are referred to as service buffers and shown as "Sce. B." in the figure. Common buffers are simpler to handle and small lengths promote the use of high speed technology. These will typically be used inside the network, *i.e.* in the ASSs and ACCs. These buffers are referred to as switching buffers and shown inside the ASSs and ACCs.

3 Traffics

3.1 Service Classes

Each type of service in a BISDN, *e.g.* voice, interactive data, video on demand, typically has its own characteristics in terms of required QoS and important traffic descriptors. We define a service class (SC) as a collection of one or more services which have the same or equivalent QoS requirements and traffic characteristics. For each SC we assign a devoted virtual service network (VSN). VSNs are logical entities on the physical network made from logically reserved transmission capacity and switching resources. VSNs can therefore be seen as a logically separate, variable networks devoted to particular SCs.

By logical subdivision of a network into VSNs are the following advantages obtained: Similar requirements within SCs allow for straight forward statistical multiplexing from a QoS point of view; Uniform traffic behaviour simplifies the associated modelling and analysis, both for network dimensioning and traffic management; Upper bounds on the resources of VSNs prevents the consequences of incorrect dimensioning for one SC to spill over and influence the performance of other SCs, and restricts the possible impact of misbehaving users to within SCs.

We also mention that VSNs may be value added networks (VANs). A typical example of such a network isa transport network for TCP/IP-information, where intermediate nodes have been equipped to support retransmissions *etc.* These nodes are then seen as a part of the VSN/VAN, but not of the official network. Such a network could support the connectionless feature of a TCP/IP service in its own way, and arrange its own bandwidth needs, *e.g.* by exploiting surplus bandwidths which might be available at low cost or by dialing underused lines of other SCs. To do this, it needs to provide its own information storage and traffic transformation.

3.2 Traffic Flows

A connection from one user to another is normally set up between the users and their ASSs through the subscriber network, between the ASSes and their respective ACCs through the access network, and finally between the two ACCs through the transit network.

In the subscriber network, no statistical multiplexing takes place. Each connection is carried on a distinct virtual channel (VC). The VC is identified by the VCI. In the access network, ASSes statistically multiplex all VCs with the same destination ASS. Such a bundle of VCs with a common destination ASS forms a virtual path (VP), here referred to as an access network VP (ANVP). An ANVP is identified by its VPI. Finally, in the transit network, ACCs statistically multiplex all ANVPs with the same destination ACC. The resulting bundle of ANVPs with a common destination ACC is referred to as a transit network VP (TNVP). TNVPs are also identified by their VPIs.

It is noted that there is no statistical multiplexing between SCs or between different pairs of ACCs. We believe, however, that this is no great loss:

- For a Swedish national network with the above numbers, approximately 10,000,000 users would share say 200 ASSs and 20 ACCs. With evenly distributed load, this means that each ASS multiplexing point potentially include 50,000 users, and each ACC multiplexing point 500,000 users. Even if individual users only generate as little as 0.01 Erlang each, the number of simultaneous sessions at the two multiplexing points amounts to 500 and 5,000 respectively. It is thus clear that there is only limited marginal statistical gain in increasing the number of simultaneous sessions at each multiplexing point.

- Uniform traffic characteristics within VSNs opens up possibilities for using more accurate multiplexing models and therefore also for safe deployment of more aggressive statistical multiplexing. This means that we might in fact obtain higher utilisation by multiplexing a smaller number of statistically identical connections than we would by multiplexing a larger number of statistically disparate connections.

- The most striking argument is, however, that for different SCs there is little gain to be had. This became obvious already with the studies of the SENET concept, where the so called strategy of movable boundary corresponds to statistical multiplexing of different SCs. In [11] it was concluded that statistical multiplexing cannot be be used to increase utilisation since the variability of the two processes take place in different time scales. In [8] a similar conclusion is reached from studies of ATM multiplexing.

3.3 Control by User Parameter Control and Tagging

User parameter control (UPC) or policing refers to supervising the traffic process from a single user to ensure that no traffic exceed the negotiated characteristics. For our network concept, each SC will typically have its own associated set of leaky bucket parameters. These are selected to meet the traffic behaviour and QoS demands of that particular SC. Another important factor in tuning the leaky bucket is, however, to make the resulting output process suitable for statistical multiplexing, and hence indirectly also to make the process mathematically tractable.

The cell header contains a simple prioritisation mechanism by means of the cell loss priority (CLP) bit. There has been a number of proposals made for how this bit should be used, many of which suggest it should be used to distinguish between loss sensitive and non loss sensitive traffics. We think this idea has two drawbacks: Firstly, a mechanism is needed to tag cells which do not conform to a traffic contract. Such cells might be a regular part of a service, where exceeding cells are treated on a best effort basis with no guarantees, but could of course also be issued by malicious users. Secondly, with the wide range of services expected in a BISDN, a simple classification of cells into "important" and "not important" leaves too little room for QoS discrimination between SCs.

In our proposal, the CLP bit can be used freely by the UPC within SCs since VSNs means that SCs can be identified by VPIs. The usage may thus very between SCs, but with the common interpretation that tagged cells are less valuable to the service in question, and should therefore be discarded first in case of congestion.

3.4 Control by Call Acceptance Control and Routing

Call acceptance control (CAC) refers to the decision of whether or not to accept a connection request. Requests are accepted if the connection is not expected to degrade the QoS for other, already accepted connections. This means that a certain amount of bandwidth, the so called equivalent bandwidth (EB), must be available on all links from origin to destination in order to accept the request.

An upper bound for the EB is defined by the peak rate p. Setting EB equal to p will minimise queuing delays and the associated service buffer overflows, but may result in poor utilisation. Similarly, a lower bound is given by the average rate a. Reserving just above a means that queuing delays will be considerable and service buffer overflows frequent. Thus, the EB is somewhere between p and a. The exact value is a complicated function of the burstiness behaviour of the session in question and of all other sessions with which it is statistically multiplexed.

VSNs restrict statistical multiplexing within sessions to distinct SCs. This means that we can precompute tables which for each SC relate the number of established connections of a that SC to the EB of one more connections of the same SC. Note that SC separation is a prerequisite for such tables, as they would be multi dimensional, and therefore practically unmanageable, if different SCs were statistically multiplexed. Reversing the tables, we can relate a bandwidth to a number of connections, or a number of equivalent circuits (ECs). The number of ECs is typically a concave function of the bandwidth, *i.e.* there is a statistical gain which results in that the marginal EB decreases the more connections that are multiplexed. For the first connection we may have marginal EB close to p, while as the number of connections approaches infinity, the marginal EB will gradually approach a.

CAC is in this way carried out simply by means of tables of ECs, which directly gives the maximum number of connections that can be accepted on a VP of given capacity. Note that only three EC table look ups must be made for a CAC-decision: ASS to ACC, ACC to ACC, and ACC to ASS. The request is accepted if ECs are available, otherwise the request rejected, queued, or rerouted. The latter option is similar to alternative routing in circuit switched networks: When no ECs are available on the AB VP, we may try ACB if ECs are available on the AC and CB VPs. Note that in node C, the connection must be looped through an ASS, to get the VC out of the AC VP and into the CB VP. It is noted that the increased number of switching points will increase both the number of table look ups and

Layer	Variations	Time scale
Network	Traffic growth	Days
Traffic	Users preferences	Hours
Session	Call placements	Minutes
Activity	Applications employed	Seconds
Burst	Information generation	Milliseconds
Cell	Cell delivery	Microseconds

Figure 2: *Layered model of traffic variations.*

the information delay for connections set up this way.

4 A Layered Model of Variations

Traffic variations, or variations in demand for network resources, can occur in virtually any time span: From one year to the next, and from one millisecond to the next. It is the task of network management to contol these. To solve this task, layered approaches have been proposed by [6] and others. We propose a classification model consisting of six layers, figure 2. Starting with the fastest variations these are the cell layer, the burst layer, the activity layer, the session layer, the traffic layer, and the network layer.

We will now elaborate on these layers from the point of view of the network and the network operator. For each layer we give its scope, time span and typical factors causing the variations.

4.1 The Network Layer

The network layer captures traffic growth, *i.e.* the long term increase in demand for network resources. The nature of these variations is typically a steady, but somewhat variable increase over a time scale on the order of weeks or more.

It is well known that the usage of existing services such as telephony and facsimile tends to increase over time. We also witness the introduction and increasing popularity of new services such as high speed data transfer as users upgrade and extend their equipment. This causes a double growth in traffic demand as new generations of hardware and software tend to handle more and more information to support new features and fancier interfaces. Other factors behind traffic growth include new business concepts, *e.g.* decentralised offices, and new working habits, *e.g.* telecommuting.

4.2 The Traffic Layer

The traffic layer captures slower variations in the behaviour of users as a group. It refers to a time scale of hours or days. Traffic layer variations basically tend to repeat themselves from one period to another, but with somewhat different demands each time.

Typical factors behind these variations are the hour of the day, and the day of the week. For example, in business telephony, peaks in traffic demand regularly occur around 9.30 *a.m.* and 2.30 *p.m.* For new services, we can foresee regular peaks in demands for interactive data traffic in business areas during working hours, and for entertainment video in residential areas during evenings and weekends. Special events might also temporarily cause demands for particular services at particular hours and from particular areas.

4.3 The Session Layer

The session layer captures the momentary variations in the behaviour of groups of users. The time scale is on the order of minutes. It has a considerable degree of randomness in it, hence predicting demands from one minute to the next is usually not very meaningful.

This level corresponds to the time scale of arrivals for call requests in traditional telephone networks, or session requests in computer networks. Observing the number of call arrivals over a few minutes, one will notice variations from one minute to the next. These variations are often referred to as random and modelled by Poisson processes. New services for non-human users might experience quite different connection request arrival patterns for which Poisson models are not applicable. In this case, classical models of bursty traffic such as the interrupted Poisson process (IPP) might be used.

4.4 The Activity Layer

The activity layer captures the variations in user activity within an ongoing session. Such variations are partly predictable, but will also contain a considerable degree of randomness. A typical time scale for the activity layer is on the order of seconds.

Data users tend to work with different applications from time to time, and also alternate between different usages of a particular application. For example, one might switch between transferring and inspecting documents, or between entering data and proof reading. Such variations can be anything from fairly regular to completely random. As for video traffic, coding schemes which deploy image compression might need more bandwidth at scene changes than within scenes. Other schemes prescribe repeated updates of the complete image, hence periodic changes of activities take place. Clearly the former will produce random variations within the activity layer, while the latter will exhibit a regular behaviour. Lastly, if speech compression is used, voice traffic exhibits activity variations as various parties are active in a conversation. Typical models include markov modulated processes such as the Markov modulated Poisson processes (MMPPs), Markov modulated Bernoulli processes (MMBPs), Markovian arrival processes (MAPs), time series models such as the autoregressive modell (AR), the moving average modell (MA), and the autoregressive moving average model (ARMA), and lately also fractal models. As for the latter, fractal models capture by definition variations in more than one layer.

4.5 The Burst Layer

The burst layer corresponds to variations due to information generation during active periods within sessions. This includes processes such as packetising and framing, which may be predictable as well as random. In any case, a typical time scale is on the order of milliseconds.

For data users, burst layer variations refer to packet assembly and high level flow control by the information transfer protocol. This suggests variations composed of both predictable variations and those which are purely random. In video traffic, burst layer variations may be related to the rate at which new frames are being delivered from the source. This process will typically be completely deterministic if no compression is used, but can be made more random by applying compression of the video signal. The list of suitable models is the same as for the activity layer, but we note that the average rate at which information is generated on the burst layer varies in the time scale of the activity layer.

4.6 The Cell Layer

The cell layer contains the variations in the rate at which bursts from sessions in progress present new cells to the network. This process includes both the formation of cells and passing them to the network. A reasonable time scale is therefore on the order of microseconds.

Cell delivery is typically controlled by the ATM adaptation layer of users' applications. Some form of protocol data units are received from the upper layers, disassembled into cell size segments and packet into cells with an appended adaptation layer header. The rate at which cells are forwarded to the network thus depends on the arrivals of protocol data units, which are captured by the burst layer, the cell formation process, and the policing and shaping functions present between the user and the network. Again, the same models as for

Layer	Means	Time scale
Network	Expanding facilities	Days
Traffic	VP assignments	Hours
Session	VC assignments	Minutes
Activity	Flow operations	Seconds
Burst	Large buffers	Milliseconds
Cell	Small buffers	Microseconds

Figure 3: *Layered model of traffic management.*

the activity layer can be used and we note a similar relationship of non-stationarity: The mean on the cell layer varies in the time scale of the burst layer.

5 A Layered Model of Management

Clearly there is no single means by which all the above variations can be handled. For example, using buffers to compensate for traffic variations on the order of hours, would require buffers that can accumulate excess day time traffic and then release it during the night. Such buffers are not only infeasible because of their size, but would also cause unacceptable delays. Similarly, virtual paths cannot be rearranged in order to compensate for burst variations on the order of milliseconds. Rearranging virtual paths affects numerous sessions in progress throughout the network, and is therefore too time consuming and too costly to be applicable to the fast and frequent variations on the burst scale.

As mentioned above, we propose to use a layered approach to traffic management. Each layer in the management model, figure 3, corresponds to a similar layer in the variation model, and fluctuatins on a certain variation level are handled by controlling entities on the corresponding management level.

The task of each management level is to reconfigure adequate network resources and traffic flows in accordance with demand variations on the level specific time scale. In doing so, it is presumed that slower variations are handled by the upper layers and faster variations by the lower layers. Each level has specific means assigned to it, such as virtual paths or service buffers, and level specific performance criteria, such as link utilisation or cell loss, to meet.

Below we give prerequisites and offerings of each of these and discuss the means by which these are accomplished.

5.1 The Network Layer

Offering The network layer offers management of extremely long term variations at the edge of the network, such as monthly or yearly growth of users' communication requirements.

Upper layer requirement Not applicable.

Lower layer requirement The traffic layer is supposed to pass on reports on demands on the traffic layer time scale.

Means Growing traffic demands can only be met by adding sufficient resources to the network. This can be done by putting new switching systems and fibres in place, deploying fibres placed into the ground before, or temporarily by leasing capacity or using resources that are intended for protection in case of network failures.

Actions are typically taken on forecasts prepared from traffic layer reports, market analyses, tariff policies *etc.*

5.2 The Traffic Layer

Offering Management of very long term variations at the edge of the network, such as daily or weekly peaks in users' requests for connections. Reports on demand variations are issued to the network layer.

Upper layer requirement The network layer is supposed to ensure that the available resources are sufficient for the average demand over the time scale of the traffic layer.

Lower layer requirement The session layer is supposed to pass on reports on demands on the session layer time scale.

Means The time scale of the traffic layer is such that only major shifts in transmission capacity are of interest.

The assignment and routing of virtual paths are the domain of the traffic layer. These can be established in accordance with forecasts of traffic layer demands, possibly with different forecasts applying to different periods such as morning, afternoon, evening and weekend. Another alternative is to provide on-line dynamic reallocation in accordance with reports from surrounding layers.

It is noted that updating the virtual path network is complicated and costly. Hence it must be restricted to occasions where the effort is duly paid of in terms of improvements in quality of service (revenues).

5.3 The Session Layer

Offering Management of medium to long term variations at the edge of the network, such as irregular connection requests on the session time scale. Reports on demand variations are issued to the traffic layer.

Upper layer requirement The traffic layer is supposed to ensure that the available resources are sufficient for the average demand over the time scale of the traffic layer.

Lower layer requirement The activity layer is supposed to pass on reports on demands on the activity layer time scale.

Means The time scale of the session layer excludes buffers and restricts operations to capacity reallocations.

The session layer controls the assignment and routing of virtual channels. It executes the call acceptance control in order to keep the traffic flow of the underlying layers within given bounds. A considerable improvement in grade of service and robustness can be obtained by allowing advanced routing mechanisms, such as those developed for the PSTN, to search for free capacity when establishing new connections.

The CAC admits calls up to the limits of the underlying layers. If a direct path between an origin and a destination cannot provide this capacity, alternative routes can be tried in much the same way as alternative routing schemes in present telephone networks, see also section 3.4. Note, however, that contrary to some telephone network routing schemes, the session layer is not responsible for very long term variations on the order of hours or days.

5.4 The Activity Layer

Offering Management of short to medium term variations at the edge of the network, such as repeated, correlated sequences of bursts. Reports on demand variations are issued to the session layer.

Upper layer requirement The session layer is supposed to ensure that the available resources are sufficient for the average demand over the time scale of the traffic layer.

Lower layer requirement The burst layer is supposed to pass on reports on demands on the activity layer time scale.

Means This level handles traffic variations related to the activity layer. It supplements the session layer in that it may dynamically alternate the number of active ECs to efficiently handle traffic correlation within sessions.

Activity layer peaks, which last on the order of seconds, cannot be absorbed by buffers inside the network because of buffer sizes and associated delays. Two alternatives are then available: The first is to find an acceptable steady state trade-off between periods of low loss and low utilisation and periods of high utilisation and high loss, and the second is to add control functions to dynamically adjust the assigned capacity so that high utilisation is combined with low loss.

In the latter approach, both the input and the output of information can be controlled. Input control refers to issuing backward congestion notifications to users which respond by stopping or slowing down the rate at which information is delivered, *e.g.* by applying increased compression. Output control refers to dynamic allocation of transmission capacity by issuing connection requests to the session layer.

Input control is suitable for faster actions of limited duration, while output control is more apt for slower actions of longer duration. By combining the two, one can combine fast actions with long duration.

If a dual leaky bucket is used as UPC-device, can one bucket be set to control traffics on the time scale of the activity layer.

5.5 The Burst Layer

Offering Management of very short term variations at the edge of the network, such as bursts of cells. Reports on demand variations are issued to the activity layer.

Upper layer requirement The activity layer is supposed to ensure that the available resources are sufficient for the average demand over the time scale of the activity layer.

Lower layer requirement The cell layer is supposed to pass on reports on demands on the cell layer time scale.

Means The burst layer activities are restricted in time and, due to propagation time also in space, hence no complicated actions can be taken.

The burst layer relies on the burst buffer, shown as service buffer in figure 1, which levels out variations within the time scale of the layer. The performance can be improved by dynamic reallocation of buffer space within the physical unit, either between a burst layer entity and a common pool, or between two different burst layer entities. VSNs permit the size of the service buffers to be different for different SCs, and we may also apply different discarding strategies in case of unresolved congestion.

If a dual leaky bucket is used as UPC-device, can one bucket be set to control traffics on the time scale of the burst layer.

5.6 The Cell Layer

Offering Management of extremely short term variations within the network, such as sessions in progress causing coincident arrivals in switching networks. Reports on demand variations are issued to the burst layer.

Upper layer requirement The burst layer is supposed to ensure that the available resources are sufficient for the average demand over the time scale of the burst layer.

Lower layer requirement Not applicable.

Means Being concerned with variations on the order of microseconds, the cell layer activities are even more restricted in time and space than the burst layer.

The basic element is the cell buffer, in figure 1 depicted as switching buffers. Cell buffers absorb cell layer activity peaks and empty them during dips. The rate at which the buffer is emptied is fixed and amounts to the capacity of the particular link or switching element.

To further enhance the quality of service, one might also employ congestion notifications or advance reservations between two entities on this level.

Cell buffers must be kept short for a number of reasons: The speed at which cells are shuffled through the buffer necessitates high speed technology hence long buffers become expensive; the delay per buffer must be small since a large number of buffers, typically one or more per switching network, are passed; and the delay must comply with the most stringent requirements since the high speed necessitates a one-buffer-fits-all approach, and there is no time for complicated priority arrangements that discriminate between cells with various degrees of urgentness. Figure 1 shows shapers after each service buffer, which are intended to ensure that traffic delivered to the transit network, which can be seen as a network of cell buffers, is sufficiently well-behaved.

6 Implementing a Strategy for Traffic Management

6.1 General

Although ideally layered models such as the one in figure 3 should permit independent implementations on each level, this is not simple and often leads to inefficient solutions with ample overhead due to multiplication of functionalities. To implement a cost effective traffic management strategy, we propose to consider a number n of adjacent levels simultaneously. The case $n = 2$ is particularly simple, since we only consider one layer boundary at a time.

6.2 Traffic and Session Layers

We will now briefly review a strategy for the traffic and session layers presented in [1, 2] and further refined in [3]. The strategy, figure 4, forms a continuous loop with the following steps: Reports on the current demands are passed from the session layer to the traffic layer; The traffic layer computes the optimal design of the VP network; A cost comparison is made where the the currently existing network is compared against the newly computed one; The new design is implemented if this appears to be more profitable than keeping the old one. As the figure shows, traffic estimation and actual implementation are carried out in the ASSs and ACCs while VP network design and the decision of whether to implement a new design or not are made at a central network management centre.

Demand estimation is a process of measuring current traffics (and psossibly predict future traffics). In our implementation this is done by observing counting arrivals of connection requests during a period of t_{M} time units. This method was chosen as a preferred alternative to measuring carried traffic, since the latter is not an unbiased estimator of the demands. The counting period t_{M} is chosen so that the expected squared error is minimised: the larger t_{M} the smaller is the expected variance, but the lower t_{M} the more accurate is the expected mean (remember that the connection request arrival process is non-stationary). Lastly, we add that a further improvement has recently been introduced by dividing the counting period in two halves of $t_{\mathrm{M}}/2$ each, and then computing a total estimate by assigning the two consecutive estimates different weights. Other methods include Kalman filtering, neural networks and possibly any method for time series prediction.

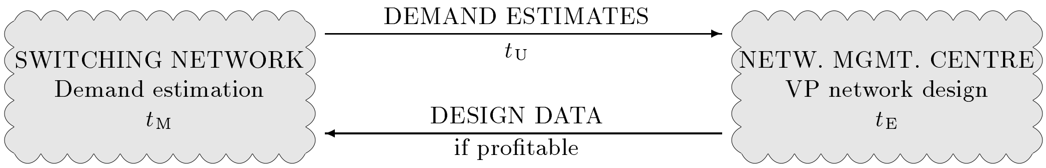

Figure 4: *Management strategy for traffic and session layers.*

The task of the VP network design algorithm is to fit appropriately dimensioned VSNs into the limited resources of the physical network so that some kind of profit function is maximised. The profit function can be defined in various ways, partly depending on whether all demands can in fact be accomodated or not. In former case is maximum profit equal to *e.g.* minimum path lengths or maximum amount of unassigned capacity, in the latter *e.g.* maximum carried traffic or maximum revenues. The latter criteria are obvious, but to motivate the two former ones, we mention that minimum path lengths give minimum delay and loss for accepted connections, and maximum amount of unassigned capacity gives the highest potential to handle unexpected traffic peaks or equipment faults.

A large number of algorithms and associated optimisation criteria have been proposed in the literature. We currently use a simple greedy algorithm, though a possibly better alternative called simulated allocation, which has some similarities to simulated annealing, has been proposed in [10]. The parameter t_{E} is the time required to compute and possibly implement a VP network design. It is obvious that the smaller t_{E}, the faster the response to the new conditions, and hence the better the long term performance. This can also be expressed in terms of network design: There is a cost in performance for every time unit spent on design, hence the cost-optimal algorithm is a trade-off between design optimality and execution time optimality.

The decision of whether to implement a new VP network design or to keep the old one is also a matter of cost optimisation: Since traffics vary more or less constantly, new designs will most likely result in shorter paths, more free capacity, more carried traffic or higher revenues. On the other hand, implementing a new design involves processing in the switches as VCs and VPs are redirected, and this does not come for free. Hence, a new design should be implemented only if the improvement in terms of path lengths, spare capacity, carried traffic or network revenues exceeds processing costs. More precisely, the costs for a VP network update must be paid back no latter than when the next update is made. Since the benefits of an update may be uncertain, *e.g.* future carried traffic or future revenues, the final decision must be made in a game theoretic way where the expected gain is maximised. Our works consider only these uncertain cases, but we believe that the case of certain gains, *e.g.* shortest paths and most free capacity, reduce to special cases of our results.

Finally, the cycle repetition interval t_{U} must be selected. Again the choice is a matter of cost optimisation: The more frequent the updates, *i.e.* the smaller t_{U}, the earlier are significant changes detected and redesigned for, but the fewer the updates, *i.e.* the larger t_{U}, the less resources are spent on management overhead. In conclusion, the cycle should be repeated at a rate where the sum of costs for unappropriate network designs and costs for management overhead is minimised. A model for the optimal choice of t_{U} in this respect has been developed in our works.

An example of the results obtained when applying our complete strategy is given in figure 5. The figure shows relative network profit versus the degree of determinism in demand variations on the traffic level. Traffic determinism can be interpreted as the standard deviation of the offered traffic relative to a predicted value, *i.e.* 1% uncertainty means that offered traffics can be predicted at any time with an average error of 1%. Relative profit is defined as carried traffic minus management costs over offered traffic, *i.e.* 100% means that all offered traffic is carried at no management costs. The different curves refer to the case of no actions on the traffic level (fixed VP network), deterministic actions (precomputed VP networks implemented at fixed instants) and dynamic actions (our strategy) for various values of t_{E}.

The figure shows that for almost perfect determinism in traffic level demand variations, fixed updating is the best alternative. With some degree of of uncertainty, however, is dynamic updating the best. The smaller the design and implementation time t_{E}, the less uncertainty is required for dynamic updating to achieve the best result. We also note that for large uncertainty, fixed updating is outperformed even by the alternative of no updating at all.

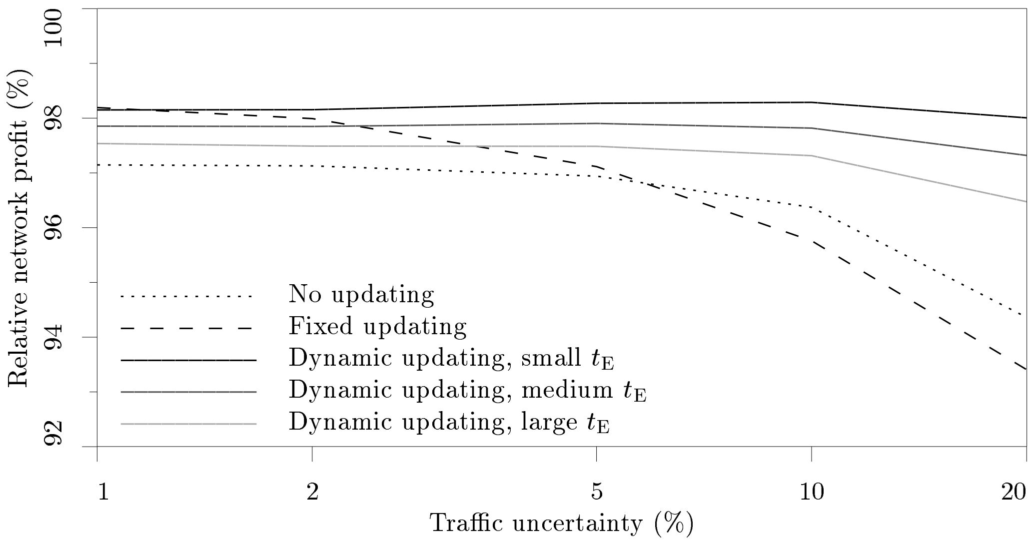

Figure 5: *Performance of traffic and session traffic management.*

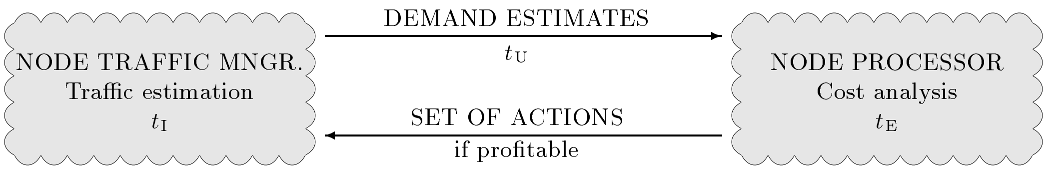

Figure 6: *Management strategy for session and activity layers.*

6.3 Session and Activity Layers

We will now briefly outline a proposal a sort of similar strategy for the session and activity layers. The strategy is depicted in figure 6 and forms a continuous loop consisting of the following steps: The activity layer reports current demands to the session layer: The session layer decides whether to increase the number of ECs allocated to the origin-destination pair in question, keep the current number, or decrease it. The decision is based on cost optimisation where management costs are traded off against traffic revenues, including QoS related costs and benefits. Whatever the action, these may be combined with decisions to activate or passivate users' abilities to slow down input rates.

Again we need to find methods for demand estimation, find optimal strategies for collecting demand estimation data, develop decision criteria for various actions which take both costs and benefits into account, and find the optimal updating repetition time or triggering mechanism that balances costs for degraded QoS against costs for management actions.

This work has just started. A number of methods for short term prediction of offered traffic, *i.e.* burst prediction, has recently been published by many authros. These are often based on Kalman filtering, *e.g.* [9], neural networks, *e.g.* [4, 5], and general time series methods, *e.g.* [7].

7 Conclusions

Motivated by the difficulties encountered in characterising, modelling, measuring and predicting broadband traffics, we see a need to construct adaptable ATM networks and that rely on dynamic controls. In this paper, we have outlined a first proposal for such a system. The network structure and management strategy are designed to provide tolerance and forgiveness against imperfect traffic information. The strategy is described by means of a layered model, where each layer is responsible for handling traffic variations in a certain time scale and each level handles its own portion of the network resources.

Both the network structure and the management model are designed to obtain an almost transparent core network which is fast and simple. Hence only high speed, low level controls are applied here. The core network supports an arbitrary number of logically distinct service separated networks. Each of these can have its own tailored equipment and strategy, such as different buffers and control methods.

Finally we have reported on a successful implementation of two of the layers in the traffic management model, and outlined how the next layer can be included.

References

[1] Arvidsson, Å.: *Management of Reconfigurable Virtual Path Networks*, pp. 931–940 in Proc. 14th Int. Teletraffic Cong., Antibes Juan-Les-Pins, France (1994).

[2] Arvidsson, Å.: *Real Time Management of Virtual Paths*, pp. 1399–1403 in Proc. Globecom '94., San Francisco, California, U.S.A. (1994).

[3] Arvidsson, Å.: *Real Time Traffic Management by Dynamic Bandwidth Allocation*, in Proc. ITC Specialist Seminar on Teletraffic Modelling and Measurement, November 7–9, Leidschendam, The Netherlands (1995, to appear).

[4] Chong, S., Li, S.-Q. & Ghosh, J.: *Dynamic Bandwidth Allocation for Efficient Transport of Real-Time VBR Video over ATM*, pp. 81–90 in Proc. Infocom '94., Toronto, Ontario, Canada (1994).

[5] Fan, Z. & Mars, P.: *ATM Traffic Prediction Using FIR Neural Networks*, pp. 34.1–34.10 in Proc. 3rd IFIP TC6/WG6.3, WG6.4 Workshop on Perf. Modelling and Eval. of ATM Netw., Ilkley, U.K. (1995).

[6] Filipiak, J.: *M-Architecture: A Structural Model of Traffic Management and Control in Broadband ISDNs*, IEEE Commun. Mag., vol. 25, no. 5, pp. 25–31 (1989).

[7] Hu X., Lambert, J. & Pitsillides, A.: *Dynamic Forecast of Boundaries of Queue in ATM Networks*, pp. 387–392 in Proc. Australian Telecommun. Netw. and Appl. Conf., Melbourne, Victoria, Australia (1994).

[8] Kim, Y. & Un, C.: *Performance Analysis of Statistical Multiplexing for Heterogeneous Bursty Traffic in an ATM Network*, IEEE Trans. on Commun., vol. 42, no. 2/3/4, pp. 745–753 (1994).

[9] Kolarov, A.: *Application of Kalman Filters in High-Speed Networks*, pp. 624–628 in Proc. Globecom '94., San Francisco, California, U.S.A. (1994).

[10] Pióro, M. & Gajowniczek, P.: *Stochastic Allocation of Virtual Paths to ATM Links*, paper no. 2.1 in Proc. 2nd IFIP TC6/WG6.4 Workshop on Perf. Modelling and Eval. of ATM Netw., Bradford, U.K. (1994).

[11] Weinstein, C., Malpass, M. & Fischer, M: *Data Traffic Performance of an Integrated Circuit- and Packet-Switched Multiplex Structure*, IEEE Trans. on Commun., vol. 28, no. 6, pp. 873–878 (1980).