# Dynamic Bandwidth Management in ATM Networks

Matthieu Verdier, David Griffin
*Department of Electronic and Electrical Engineering,*
*University College London, Torrington Place, London WC1E 7JE, UK,*
*M.Verdier@ee.ucl.ac.uk, D.Griffin@ee.ucl.ac.uk, Tel.: +44 171 319 3557*

Panos Georgatsos
*Algosystems S.A., 4, Sardeon Str., 171 21 N. Smyrni, Athens, Greece,*
*pgeorgat@algo.com.gr, Tel.: +30 1 93 10 281*

**Abstract**

This paper describes a system for dynamic bandwidth management of Virtual Path Connections in ATM networks. A two level hierarchical approach is taken consisting of network planning functions which take place on a longer time scale than the bandwidth distribution functions which attempt to track the usage of the VPCs in terms of the switched VCCs admitted onto them. The paper concentrates on approaches for the latter of these sets of functions and places them in the context of a larger system for network survivability and availability as studied by the ACTS REFORM project.

## 1. Introduction

The aim of the REFORM project (EU ACTS project AC208) [7] [8] [9] is to specify, implement and test a *reliable* system that offers ATM, multi-class, switched services. Generally speaking, network *reliability* entails: network *survivability* and network *availability.* Network *survivability* refers to the necessary functions to guarantee a continuous service for established connections in cases of failures occurring within the network. Network *availability* refers to the optimal configuration and operation of the network at all times, to accept successfully the highest potential amount of new service requests.

Within the REFORM system, network *survivability* is implemented by means of an ATM layer protection switching mechanism. This mechanism targets at the reconfiguration of the VP layer infrastructure by switching the failed VPCs to standby (predetermined) alternative VPCs. The full methodology, including restoration resource control protocols and network reconfiguration algorithms that go along with this mechanism are documented in [8] but this aspect of the REFORM system is not the subject of this paper.

Network *availability* is concerned with the cost-effective planning and maintenance of network resources so that to maximise user connection admissions. The planning aspect is not covered by this paper but the project has described suitable VP layer design algorithms needed to originally configure the ATM layer (given a physical network) so as to meet the traffic predictions [8]. These algorithms include the configuration of the necessary protection resources required by the protection switching mechanism.

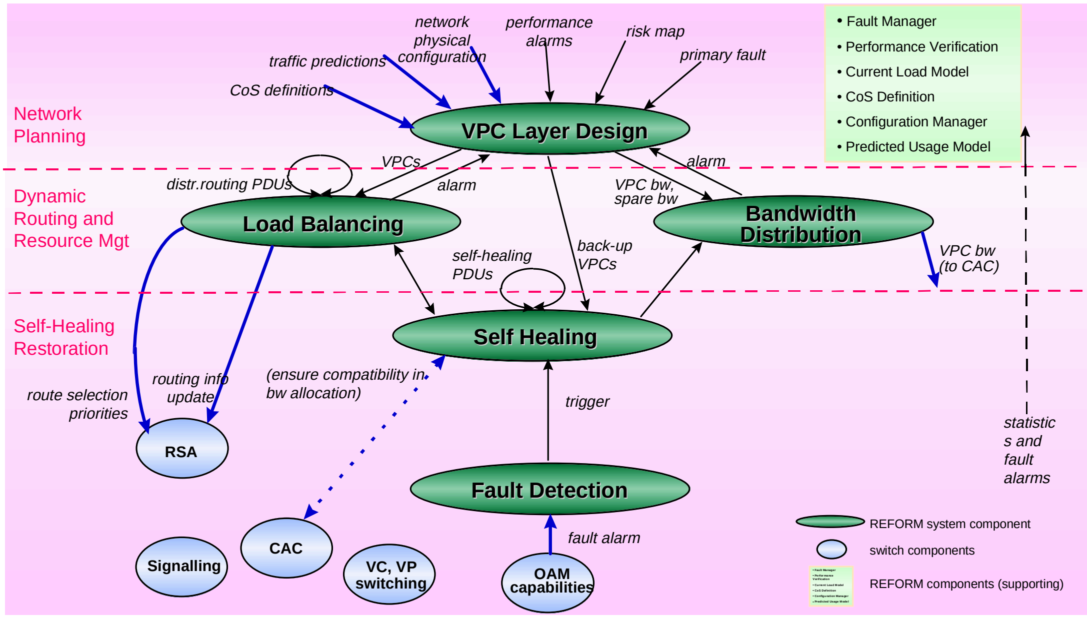


**Figure 1 REFORM functional model**

For the purposes of this paper we can consider that network operation is generally decomposed into two distinct operational phases. During the *initialisation* phase, the network is prepared for service provisioning at a certain service level. During the *normal* phase, the network delivers services, and sees to the active management of its resources so as to guarantee its service levels under deviations of offered traffic at its edges. This paper is concerned with the dynamic management of bandwidth during the normal phase according to network designs created during the initialisation phase.

The initialisation phase results in the definition of a suitable network of working VPCs (for carrying user traffic) and admissible routes based on them per source-destination and Class of Service[1] (CoS) so that to preserve the performance characteristics of each CoS. Furthermore, for the network to cope gracefully (without affecting the integrity of existing services and its availability to future services) with fault conditions, protection VPCs need also to be planned and the appropriate restoration bandwidth need to be allocated. The initialisation activities are undertaken within a single functional component, called *VPC Layer Design (VPC_LD).*

Since user behaviour changes dynamically there is a chance that the network may become unbalanced when the bandwidth allocated to VPCs on the existing admissible routes are not in accordance with the quantity of the traffic that it is actually offered to be routed over them.

There are basically two levels at which adaptivity to traffic variations should be provided, one at a level of (structural) traffic prediction changes and one at the level of actual traffic fluctuations around the predictions. Therefore, it is reasonable to consider that VPC and routing management is achieved through a two level hierarchy (see Figure 2).

The higher level of the hierarchy is undertaken by the VPC_LD component which reconfigures the VPC and routes per class of service whenever the traffic predictions change significantly. The level of reconstruction obviously depends on the significance of the changes.

[1] Each service type belongs to a Class of Service which defines the bandwidth and performance characteristics of the traffic.

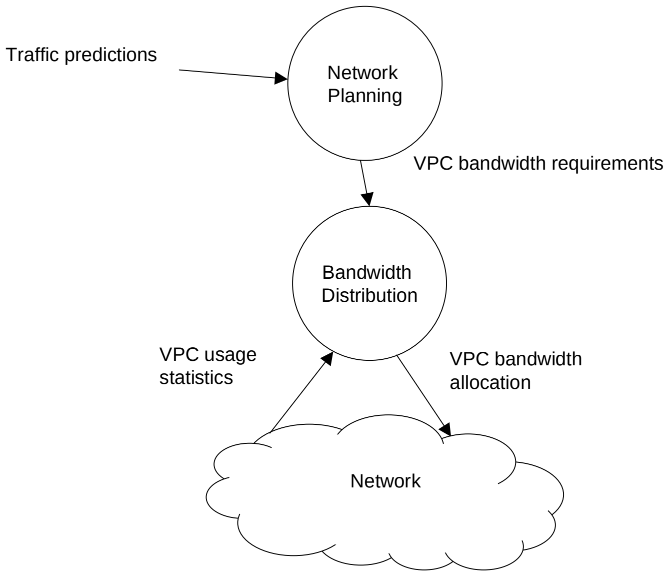


**Figure 2 Hierarchical approach to bandwidth management**

The lower level in the hierarchy is introduced to compensate for inaccuracies in the traffic predictions and short term fluctuations in actual load around the predictions. The lower level functionality operates on the set of existing routes and redefines the working VPC bandwidth and route selection parameters.

The above discussion indicates that two further functional components need to play part in the resource and routing management hierarchy: *Bandwidth Distribution* (for updating working VPC bandwidth) and *Load Balancing* (for updating route selection parameters). This paper discusses the objectives, design constraints and proposes an algorithm for the first of these two components: Bandwidth Distribution.

## 2. Bandwidth Distribution

The Bandwidth Distribution (BD) component is responsible for the management of the bandwidth allocated to working VPCs according to actual traffic conditions. That is, it adjusts the bandwidth allocated to the working VPCs to their actual usage to avoid situations where in the same links some VPCs tend to become over-utilised while other VPCs remain under-utilised. The dynamic management of the working VPC allocated bandwidth is achieved by distributing portions (viewed as a common pool) of the link working bandwidth (link capacity minus restoration bandwidth) among the working VPCs. Specifically, the management of the VPC allocated bandwidth is done within specific (upper and lower) bounds on the VPC bandwidth as originally estimated by VPC_LD (VPC *required bandwidth*).

The activities of BD are required to compensate for inaccuracies in traffic predictions and in the VPC bandwidth as estimated by VPC_LD as well as to withstand (short to medium) actual traffic variations. This is so, since it cannot be taken for granted that the traffic predictions will be accurate and furthermore even if they are accurate, they are accurate within statistical range. In this respect i.e. considering the random nature of the arriving traffic, the bandwidth that needs to be allocated to VPC, so that certain objectives (regarding connection admission) to be met, is a stochastic variable, depending on the connection arriving pattern. The VPC_LD component estimates originally the bandwidth that needs to be allocated to the VPCs (*required bandwidth*) so that satisfy traffic predictions. This is viewed as the mean value of the (stochastic in nature) bandwidth that needs to be allocated to the

VPCs. It is the task then of the BD component, to manage the allocated bandwidth of VPCs, around the (mean) *required bandwidth*, according to actual traffic conditions.

By monitoring the usage on working VPCs, the BD component also emits warnings to VPC_LD indicating insufficient usage of the planned network resources. The warning are issued in cases where some VPCs remain under-utilised (with respect to their required bandwidth as specified by VPC_LD) for a significant period of time. This implies that these resources cannot be utilised in the routes by the Load Balancing component and therefore such cases are interpreted as indicating overestimation of network resources.

The proposed algorithm for bandwidth redistribution assumes that there is a common pool of bandwidth per link to be redistributed to the VPCs when necessary. The algorithm assumes that this common pool of bandwidth per link is the links' unallocated bandwidth. Note, that by its definition, this pool of bandwidth is not totally allocated to the VPCs at any instant; but it is there to be allocated to the VPCs that go highly utilised only when such conditions occur. Each VPC grabs or returns portions of its allocated bandwidth to the common pool of bandwidth according to its congestion level. We assume here that the modification of the bandwidth of a VPC does not impact on the traffic parameters (QoS) of the VCs using this VPC or other VPCs sharing the same links or nodes.

## 2.1 Possible approaches

In distributing resources we have three main approaches:

### 2.1.1 Generous

The generous approach in the BD context would consist in allocating as much bandwidth possible to every VPC in order that for example all VPCs have an equal share of the spare bandwidth. In fact this scheme would be very reasonable except for one thing: It would make creation of new VPC, and thus reconfiguration impossible or very difficult.

### 2.1.2 Greedy

In this strategy we allocate "just enough" bandwidth to each VPC, so to maintain spare link bandwidth for new VPCs.

If you say that "just enough" is the used VPC bandwidth + 20% then a small VPC might not be able to accept any new VC. And a big VPC will have a large spare bandwidth which will imply a high acceptance potential and might lead to an unbalanced network.

If you say that "just enough " is the used VPC bandwidth + Constant, then you have to be careful that this constant is large enough to allow creation of new VC and not too large to avoid wasting resources.

### 2.1.3 Fair

Fairness here means taking into account information and the dynamics of the other components of the system, here mainly from VPC_LD, Reconfiguration and LB.

1. VPC_LD provides BD with the required bandwidth for each VPC: Vreq.. Either:
    - Hypothesis 1: Vreq is only an initial value.
    - Hypothesis 2: The traffic predictions shows that Vreq (or an interval around Vreq) is the probable bandwidth needed at some point by this VPC. In that case even if the used bandwidth is low, BD shouldn't allocate the VPC less bandwidth than Vreq
2. Reconfiguration needs spare link bandwidth.

3. LB is most concerned with the overall distribution of spare bandwidth within the VPCs, since it determines how the VC will be routed. Either
    - it is important that LB should try to fulfil the prediction made by VPC_LD. In that case even if a VPC is empty you should allocate him Vreq so that LB will have the incentive to give him traffic.
    - VPC_LD only provides a topology. In that case BD should shape the spare bandwidth resources so that the spare bandwidth is evenly distributed between the different VPCs.

## 2.2 Chosen solution

The approach consists in defining for each VPC an Admissible Zone AZ in which the VPC bandwidth is allocated. This Admissible Zone is defined around the Vreq value given by the VPC_LD component. BD is not allowed to allocate bandwidth outside these boundaries. We also define the operating point of the VPC as the last measured Vused send to BD by NRM, this is the freshest actually measured used bandwidth that BD knows. Around the operating point of the VPC we define a working zone WZ which defines an upper and lower threshold. Above this working zone we define a buffer zone BZ.

The general rule of operation goes as follows: Each time the bandwidth used crosses one of the working zone boundaries an alarm is raised. BD then centres the working zone around the new operating point, and allocates the VPC a new bandwidth equal to the upper boundary of the working zone augmented by the buffer zone. Special cases occurs when this bandwidth reaches the upper limit of the admissible zone or when this bandwidth is not available because of link capacity restrictions along the VPC path or its protection path. Other special cases occurs when reaching the lower bound of the admissible zone.

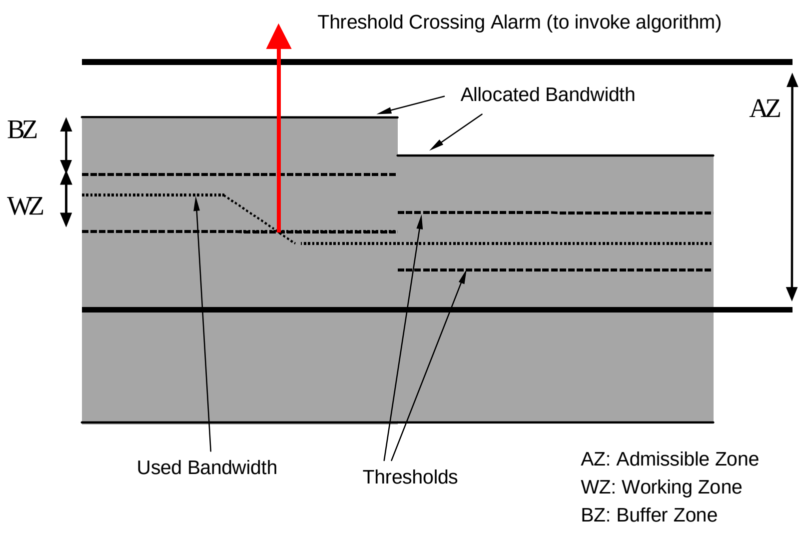


**Figure 3 Tracking used bandwidth on VPCs**

Advantages seen to the moving working zone and buffer zone approach:

- A VPC wants more bandwidth but we can't allocate any to him: The WZ being centred around Vused, BD will not repeatedly receive alarms since the upper boundary of the WZ will be above Valloc. What should be done is to store the fact

that this VPC needs more bandwidth as soon as possible, in a list of pending issues, until the full WZ and BZ can be allocated.
    - If the VPC is empty or has very low traffic: again BD will not receive repeated alarms since the lower boundary of the WZ will be negative (or zero).
- We have a fairly good idea of how much unutilised bandwidth there is on each link: spare link bandwidth + protec bandwidth + sum(BZ) + sum(1/2WZ)
    - The unutilised bandwidth in each VPC is approximately constant: BZ+1/2WZ and bounded: max=BZ+WZ, min=BZ.

### 2.2.1 The Admissible Zone

The Admissible Zone of one VPC defines the upper and lower limit of the VPC allocated bandwidth. Its upper and lower value should be determined by a "VPC class". The width of AZ determines how close the bandwidth allocated to the VPC should be to the Vreq value determined by VPC_LD: a narrow AZ means little flexibility in altering the VPC's bandwidth while a wide AZ means that Vreq shouldn't be taken too seriously by BD. The exact use and range of the "VPC class" is for further study during the detailed design of BD and VPC_LD.

### 2.2.2 The Working Zone width

The width of the Working Zone determines the frequency with which threshold crossing alarms will be generated. The Vused value fluctuates as connections are made and dropped on the VPC in question. The magnitude and frequency of the fluctuations depends on the bandwidth of the VCs and on the connection arrival and drop patterns (which is influenced by: the arrival and drop rates to the network in terms of CoS and source-destination pairs; the admissible routes defined for VCs of a particular CoS; and the behaviour of load balancing in selecting routes for VCs among the set of admissible ones).

The width of WZ is therefore dependent on two factors: the nature of the fluctuations in the Vused signal; and the required sensitivity of BD to used bandwidth changes. Regarding the Vused fluctuations, exact traffic models are not known at the time of writing. Further work on traffic simulation is required under varying conditions of sets of admissible routes and LB behaviour.

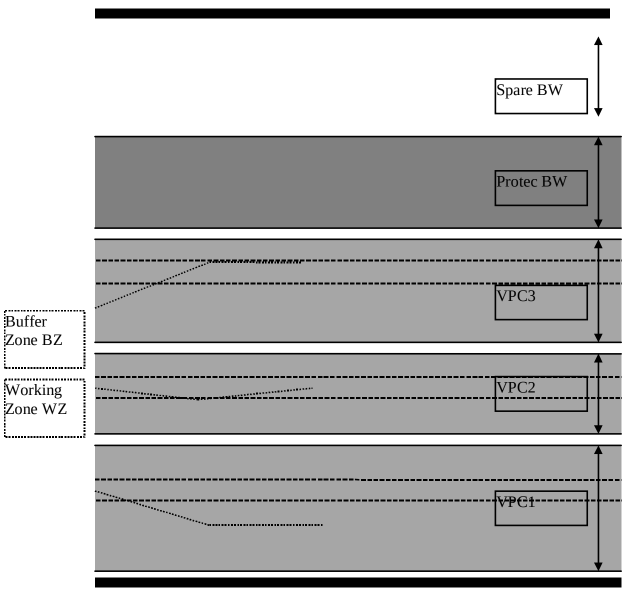


**Figure 4 Distributing link capacity to VPCs**

Given a specific model of Vused variations over time (to be established experimentally) we can consider the second factor: the required sensitivity of BD to variations in Vused. If WZ is made small then a large number of threshold crossings will take place, requiring BD to be invoked frequently. On the one hand this is positive in the sense that it will cause the bandwidth allocated to VPCs to closely follow the Vused value thereby ensuring that link capacity isn't needlessly allocated to VPCs. On the other hand this is negative in the sense that it will increase the quantity of management traffic (alarms) and the processing overhead of BD. If WZ is made large the number of alarms and hence the processing load on BD will be reduced but at the cost of a less accurate reflection of Valloc to Vused. There is clearly a trade off between these two approaches and this will be experimented with during the testing phase of BD both via simulations and testbed runs.

### 2.2.3 The Buffer Zone width

The width of the Buffer Zone is related to the widest bandwidth of the VC class supported by the VPC and will have to be given by a VPC class parameter. In fact the Buffer Zone has to be at least the widest width of the VC supported by the VPC to ensure that the VPC which is has a Vused of slightly below the upper WZ threshold can accept (at least) one more VC. This VC will cause the upper threshold to be crossed, which, in turn, will invoke the BD algorithm to increase the VPC bandwidth (if there is available link capacity). If the interval between the time the VC is accepted and the time when a successful increase in VPC bandwidth has taken place is longer than the inter-connection arrival time then VCs may be blocked if the BZ isn't wide enough to accommodate the expected VC arrivals in that period. In this case BZ should be wider than a single VC.

The exact value will have to be determined experimentally when traffic arrival patterns are known.

If BZ is too small it could lead to unacceptable Cell Blocking Rates (although these may be mitigated by alternate routes using different VPCs).

If BZ is too large it would lead to an inefficient use of the link bandwidth.

## 3. Conclusions

This paper presents the dynamical bandwidth distribution algorithm implemented by the REFORM project. The REFORM ATM network management and control architecture is described, then the Bandwidth Distribution algorithm is detailed. The algorithm is based in the utilisation of moving windows defining thresholds on the used bandwidth of a VPC which when crossed triggers the reallocation of bandwidth process, and the redefinition of the thresholds. The allocation of bandwidth to a VPC is constrained by the available resources. Simulation of the dynamical bandwidth distribution algorithm will be done as part of the REFORM project. Results of the simulation will be presented if available for the presentation.

## References


[1] Friesen, V., Harms, J., Wong, J., "Resource Management with Virtual Paths in ATM Networks," IEEE Network, September 1996.

[2] Georgatsos, P. Griffin, D., "A Management System for Load Balancing through Adaptive Routing in Multi-Service ATM Networks," in IEEE INFOCOM'96, Proceedings Vol. 2, IEEE Computer Society Press, Los Alimitios, CA, USA, 1996.

[3] Georgatsos, P., Griffin, D., "A General Framework for Routing Management in Multi-Service ATM Networks," in "Integrated Network Management V: Integrated management in a virtual world," proceedings of the fifth IFIP/IEEE international symposium on Integrated Network Management, Lazar, Saracco, editors, Chapman & Hall, London, UK, 1997.

[4] Griffin, D., et al, "ATM Virtual Path Connection and Routing Management", in Integrated Communications Management of Broadband Networks, ed. D. Griffin, pp. 73-145, Crete University Press, 1997.
[5] Griffin, D., Georgatsos, P., "A TMN System for VPC and Routing Management in ATM Networks," in "Integrated Network Management IV," proceedings of the 4th international symposium on Integrated Network Management, Sethi, Raynaud and Faure-Vincent, editors, Chapman & Hall, London, UK, 1995.
[6] Medhi, D., "Models for Network Design, Servicing and Monitoring of ATM Networks based on the Virtual Path Concept," in Computer Networks and ISDN Systems 29, pp 373-386, 1997.
[7] REFORM Deliverable D1, “REFORM System Requirements and Analysis,” A208/BELL/WP1/DS/P/002/b1, January 1997.
[8] REFORM Deliverable D7, “Description of self-healing mechanisms, load balancing and dynamic routing capabilities of the REFORM system,” September 1997.
[9] Sartzetakis, S., Georgatsos, P., Konstantoulakis, G., Pavlou, G., Griffin, D., "Unified Fault, Resource Management and Control in ATM-based IBCNs," in "Integrated Network Management V: Integrated management in a virtual world," proceedings of the fifth IFIP/IEEE international symposium on Integrated Network Management, Lazar, Saracco, editors, Chapman & Hall, London, UK, 1997.
[10] Shioda, Saito, Yokoi, IEICE Trans. Communications, vol. E-80-B, no. 2, February 1997.
[11] T'Joens, Y, Georgatsos, P., Georgiadis, L., "Network reliability in ATM based IBCN, a functional description of the REFORM system", Proceedings of DRCN'98, Belgium, 1998.