A TMN system for VPC and routing management in ATM networks

D. P. Griffin
Institute of Computer Science,
Foundation for Research and Technology-Hellas,
PO Box 1385, 711-10 Heraklion, Crete, Greece.
Tel: +30 81 391722, Fax: +30 81 391601
email: david@ics.forth.gr

P. Georgatsos
ALPHA Systems S.A.,
3 Xanthou Str., 177-78 Tavros, Athens, Greece.
Tel: +30 1 482 6014, 15, 16, Fax: +30 1 482 6017
email: panos@alpha.ath.forthnet.gr

Abstract

In this paper we present a VPC and Routing Management Service for multi-class ATM networks. Considering the requirements, we decompose the Management Service into a number of distinct but cooperating functional components which we map to the TMN architecture. We describe the architectural components and analyse their operational dependencies and information exchange in the context of the overall system operation.

The proposed system offers the generic functions of performance monitoring, load monitoring and configuration management in ATM networks. In addition, it provides specific functions for routing and bandwidth management in a hierarchical structure.

Keywords

ATM, TMN, performance management, routing, VPC, multi-class environment.

1 INTRODUCTION

The efficient operation of a network depends on a number of design parameters, one of them being routing. The overall objective of a routing policy is to increase the network throughput, while guaranteeing the performance of the network within specified levels. The design of an efficient routing policy is of enormous complexity, since it depends on a number of variable and sometimes uncertain parameters. This complexity is even greater, taking into account the diversity of bandwidth and performance requirements that the network must support. The routing policy should be adaptive to cater for traffic and topological changes.

Routing in Asynchronous Transfer Mode (ATM) (ITU I.150) is based on Virtual Path Connections (VPCs). A route is defined as a concatenation of VPCs, where each VPC is defined as a sequence of links being allocated a specific portion of the link capacity. It has been widely accepted that VPCs offer valuable features that enable the construction of economical and efficient ATM networks, the most important being management flexibility. Because VPCs are defined by configurable parameters, these parameters and subsequently the routes based on them can be configured on-line by a management system according to network conditions.

Since user behaviour changes dynamically there is a danger that the network may become inefficient when the bandwidth allocated to VPCs or the existing routes are not in accordance with the quantity of traffic that is required to be routed over them. To combat this, the VPC topology, the routes, and the bandwidth allocated to VPCs must be dynamically re-configured. A VPC and Routing management system is required to take advantage of the features of VPCs while ensuring that the performance of the network is as high as possible during conditions of changing traffic.

The ITU-T have distinguished between the management and control planes in the operation of communications networks (ITU I.320, I.321) and introduced the Telecommunications Management Network (TMN) (ITU M.3010) as a means of provisioning management systems with standard interoperable components according to the ISO systems management standards. The TMN should compliment and enhance the control plane functions by configuring operational parameters. The TMN should not replace the control plane and in general it has less stringent requirements on real-time response.

Although there is a significant research interest in the area of performance management on ATM particularly in routing (Sykas 1991, Gelenbe 1994), bandwidth assignment (Hui 1988, Saito 1991) and VPC management (Ohta 1992, Sato 1991), the problem of VPC and routing management remains largely open. The majority of management systems deployed today are concerned with network configuration and network monitoring and the management intelligence is provided by the human users of the management systems. There is a trend (Woodruff 1990, Wernic 1992, Geihs, 1992) to increase the intelligence of the management functions to encapsulate human management intelligence in decision making TMN components to move towards the automation of the monitoring, decision making and configuration management loop.

Within the framework of performance management this paper investigates the requirements of VPC and routing management functions for ATM based B-ISDN networks and proposes a TMN system for implementation. The ITU-T terminology (ITU M.3020) for describing Management Services is adopted. In particular the paper proposes a Management Service for VPC and routing management and decomposes it into a number of components. The design is mapped to the TMN architecture for implementation using TMN and OSI systems management principles.

Section 2 defines the VPC and Routing Management Service and section 3 discusses the environmental assumptions and constraints. Section 4 presents the decomposition into management components and outlines the rationale behind it. The mapping to the TMN architecture is also presented in this section. Section 5 details the management components and section 6 describes their interactions and their relationships. Finally section 7 presents the conclusions and identifies future work.

2 THE MANAGEMENT SERVICE

Within a multi-class ATM network environment the objective of the VPC and Routing Management Service is to guarantee network availability whilst guaranteeing that the network meets the performance requirements of the different service classes. This Management Service is beneficial to the network operator since it ensures that the network resources are used as efficiently as possible.

The VPC and Routing Management Service has both static and dynamic aspects. The static aspect is related to the design of a VPC network and a routing plan (the set of routes and selection criteria for each source-destination pair and service class) to meet predicted demand. In fact the static aspect is of quasi-static form in the sense that is invoked whenever the traffic predictions change significantly. The dynamic aspect manages the VPC network and the routing plan to cater for unpredictable user behaviour within the epoch of the traffic predictions.

This Management Service belongs to the performance and configuration management functional areas and specifically covers traffic management while its static aspects are related to the network planning functions. Figure 1 shows the relationship of VPC and Routing Management with the network, human managers (TMN users), other management functions, network customers and other network operators.

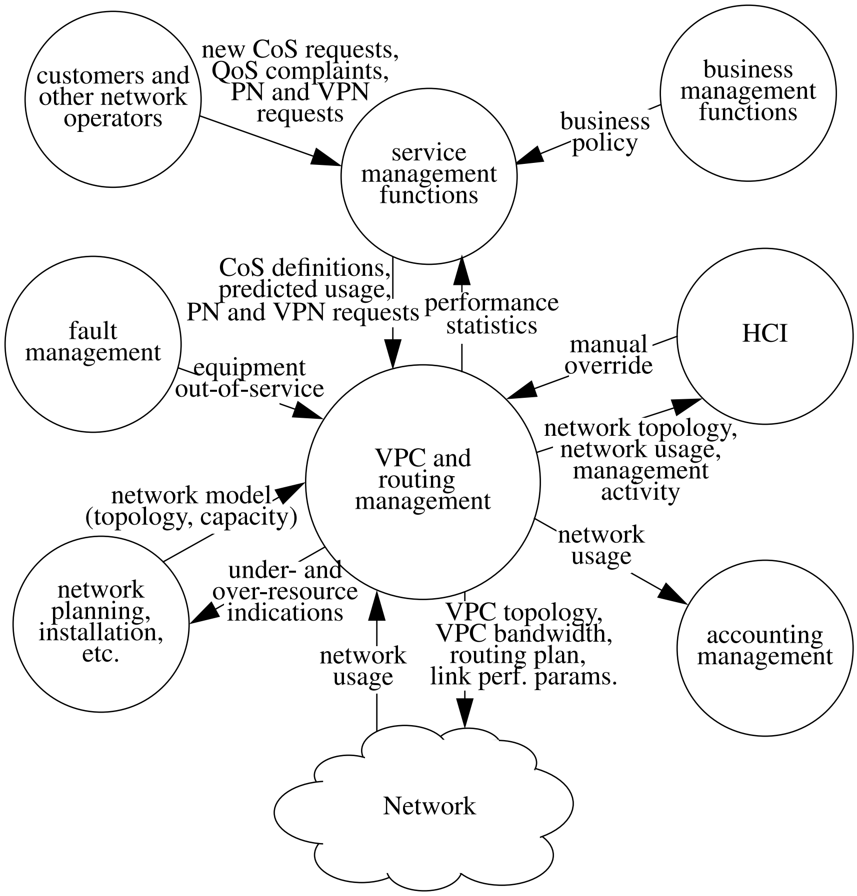

Figure 1 Enterprise view of VPC and Routing Management.

The methodology of ITU-T Recommendation M.3020 is adopted. According to this Recommendation Management Services are composed of Management Service components (MSCs) which in turn are constructed from management functional components (MFCs). MFCs are themselves constructed from management function sets (MFSs) which are groups of management functions which logically belong together. In this paper we will decompose the Management Service to identify the constituent MSCs and MFCs and show how these can be mapped to the logical TMN architecture.

3 THE ENVIRONMENT

This section describes the network environment from the perspectives of the VPC and Routing Management Service.

The managed environment is assumed to be a public ATM network offering switched, on-demand services ranging from simple telephony and file transfers to multi-media conferences.

3.1 Assumptions on the network services

Service calls are decomposed into a number of unidirectional connections. A large number of connection types are supported, for example, telephony services may be supported by a number of connection types which offer a range of qualities - different delays or call blocking probabilities. The term class of service (CoS) is used to denote a particular connection type. The CoSs are the bearer services provided by the network. The CoS definition characterises the connection type in terms of bandwidth and performance requirements.

Our work assumes that the bandwidth requirements can be characterised by mean and peak values. Alternative bandwidth parameters may be used according to the specific connection admission control (CAC) algorithms employed in the switches.

We assume the following performance parameters: cell loss probability; delay; delay jitter; and connection blocking probability (or availability). In fact these are the performance parameters the Management Service is able to influence and are of direct interest but other performance parameters may be included in the CoS definition, connection release delay for example, but these may not be influenced by this Management Service and are not considered further here.

An issue to be clarified is the relationship between the classes of the bearer services provided by the network and the four AAL classes recommended by the ITU-T (ITU I.362). The AAL provides a limited range of services, e.g. connection-oriented vs. connectionless, error recovery, re-transmissions with the assumption of a given performance of the underlying bearer service.

Our view is that there needs to be a range of bearer services of different qualities and costs to support the AAL services. This will allow decisions to be made on whether to use a comprehensive AAL with a cheap, low performance bearer service or a lightweight AAL with a higher performance bearer service (e.g. smaller cell loss ratio). This view is in accordance with the views of the ATM Forum (ATM Forum 1993) where they explicitly recommend the augmentation of the AAL service classes with a range of quality of service classes. The AAL exists in the user terminals whilst the underlying bearer service is provisioned by the network operators.

Our work concentrates on the management of the bearer services from the viewpoint of the network operator. Although AAL issues are considered from the perspective of the requirements they impose on the underlying bearer services, the end-to-end management issues of layer 4 and above are not the focus of our work.

Another important point is the role of connection oriented services with a predefined bandwidth and performance compared to that of best effort (no performance requirements) or available bandwidth (ABR) services. We recognise the requirement for all types of services but our work concentrates on the management needs of the services with predefined bandwidth and performance. Best effort and ABR services are controlled via the signalling protocols. If they are to coexist with services of defined quality on the same network, there is a necessity for the bandwidth and routing management functions to dynamically manage the partitioning of the network resources. However this is an issue for future work.

3.2 Assumptions on user behaviour

User behaviour is not constant and changes dynamically. There are two sources of variation: the type of user and the population of the users. There are potentially many different types of users characterised by the types of service they use and also by their usage patterns. The behaviour of individual users changes over time with respect to the services they use and the way that they use their services.

We assume, by virtue of the law of large numbers, that estimates of aggregate user behaviour can be made and trends can be identified in the short term (e.g. business vs. domestic traffic throughout the working day) and the medium to long term (e.g. seasonal variations, new service introduction, competition).

3.3 Assumptions on network operation

ATM networks are connection-oriented networks. Each node basically provides switching and call control (CC) functionality which includes route selection and connection admission control (CAC). Switching is done at two levels: VP cross-connects switch cells within a VPC based on the VPI; VC switches switch cells of a particular VCC between VPCs based on their VCI. VPCs are created on a semi-permanent basis by management actions whereas VCCs are created dynamically by the control plane of the network via UNI and NNI signalling. Route selection refers to the selection of a particular route upon receipt of a connection establishment request by means of a Route Selection Algorithm (RSA). CAC is required in order for the node to determine whether the connection can be accommodated on the selected route. This is done by means of a CAC algorithm that controls the VPC loading in the admissible region i.e. in the region where buffer overflow is within the bounds of a pre-defined probability (cell loss target of the CAC).

The latter two functions, route selection and CAC, are part of the control plane. However their behaviour is specified according to operational parameters which are defined and managed by the TMN.

In order to accomplish routing all possible routes towards a given destination and a particular CoS are stored local to the switches in a route selection table. For a number of reasons (increased availability, reduced vulnerability to failures, adaptivity) more than one route may exist to a destination for a CoS. The RSA in each switch searches its route selection table for entries satisfying the destination and CoS. We assume that the RSA is based on route selection parameters associated with the available routes. These parameters reflect the preference of selecting one route over another. The RSA should be as fast as possible and cause the minimum overhead to the network.

4 DECOMPOSITION

4.1 The rationale

Connection rejection is affected by two factors: the number of alternative routes and the available capacity on the VPCs. These two factors cannot be treated in isolation and the VPC and Routing management system must therefore ensure that there are sufficient numbers of routes and bandwidth on the VPCs forming the routes to guarantee network performance and availability.

As mentioned previously the Management Service should provide adaptivity to changing traffic conditions. There are two levels at which the traffic can change: cell level variations within the scope of a single connection; and connection level variations as users establish and release calls. The former is considered to be dealt with by the CAC and UPC functions of the control plane. Connections can never exceed the bandwidth parameters defined for a CoS due to the role of the UPC functions. If connections do not consume the full bandwidth the shortfall cannot be used by other connections because of the concept of pre-defined bandwidth reservation at connection set-up time which is paid for by the users. For this reason cell level variations are of no concern to this Management Service and the management of connection level variations is the main focus.

The following views of the network are useful for offering different levels of abstraction to assist the task of formulating the problem faced by the VPC and Routing Management Service.

- *The physical network* consisting of the network nodes and the transmission links.
- *The VPC network* consisting of the VC switches interconnected by VPCs.
- *The ClassRoute networks*. For each CoS, the ClassRoute network is the sub-network of the VPC network which consists only of the VPCs that belong to routes of that CoS.
- *The SDClassRoute networks*. For each CoS and a given source-destination (s-d) pair, the SDClassRoute network is the sub-network of the ClassRoute network consisting only of the VPCs that belong to the routes of the given (s-d) pair.

Having introduced the above network views the goal of the VPC and Routing Management Service can be formulated as follows:

- Given the physical network and the traffic predictions per s-d and CoS, define VPC and SDClassRoute networks so that the traffic demands are met and the performance levels specified per CoS are guaranteed.

The solution requires answers to the following questions:

- How is the VPC network constructed and how frequently will it change?
- How are the ClassRoute networks constructed and how frequently will they change?
- According to what criteria will routing be achieved in the ClassRoute networks? i.e. Given the VPC and ClassRoute networks how are the route selection parameters assigned and how frequently will they change?

The definition of the VPC and ClassRoute networks is an iterative procedure which cannot separate the two tasks involved. Routes are defined in terms of VPCs and the VPCs have been defined in order to support routing.

The VPC and the ClassRoute networks are constructed using, as input, estimates for the network traffic per s-d pair and CoS. The construction of these two networks is related to the network planning activity, whereby the topology of the physical network is defined based on longer term network traffic predictions. The design of the VPC and Routing management system should therefore cater for changes in the predictions and inaccuracies in the predictions.

Whenever the traffic predictions change, the VPC and ClassRoute networks need to be reconstructed. The level of reconstruction obviously depends on the significance of the changes. As a result, new values for VPC bandwidth may be given, or the topology of the VPC network may change (by creating and deleting VPCs) or the topology of the ClassRoute networks may change (by creating and deleting routes). Each of these reconfigurations deals with a different level of abstraction according to the network views described above. Moreover they may be performed within different time scales and they require different levels of complexity and hence computational effort. We envisage that an efficient way to deal with such reconfigurations is through a hierarchical system.

The essence of the hierarchy we propose is as follows. First the VPC bandwidth is reconfigured within the existing SDClassRoute networks. If it is not possible to accommodate the traffic predictions within the SDClassRoute networks, the SDClassRoute networks are reconfigured within the existing VPC network. If it is found that the VPC topology is insufficient for the predicted traffic then finally the VPC network is reconfigured. Ultimately it may be discovered that the physical network is unable to cope with predicted traffic and the network planning functions are informed to request that additional physical resources are deployed.

This indicates the need for having three management components: Bandwidth Allocation (for VPC bandwidth updates given SDClassRoute networks), Route Planning (for route updates given the VPC network) and VPC Topology (for VPC topology updates).

The above assumes that the traffic predictions are accurate, but as mentioned previously, this cannot be taken for granted. For this reason we introduce a lower level into the hierarchy which tries to make the initial estimates more accurate by taking into account the actual usage of the network. The lower level functionality operates within the SDClassRoute networks and redefines the VPC bandwidth and route selection parameters taking into account the actual network load. Redefinition of SDClassRoute networks and VPC topology is not done at this level since it must be as lightweight as possible. However this level will provide triggers to the higher level when it is proved that the first level estimates under or over estimate the actual situation and this cannot be resolved at this level. Even if the predictions are accurate there is still a case for lightweight lower level functions to cater for traffic fluctuations within the timeframe of the predictions.

This indicates the need for two components in the lower level: Bandwidth Distribution (for updating VPC bandwidth) and Load Balancing (for updating route selection parameters).

The proposed hierarchical system exhibits a fair management behaviour whereby initial management decisions taken with a future perspective are continuously refined in the light of current development. Apart from its fairness, such a behaviour provides a desirable level of adaptivity to network conditions.

4.2 MSCs and MFCs

The previous section indicates the following decomposition of the VPC and Routing Management Service into MSCs:

- *management of VPC topology* which is placed in a VPC Topology MFC
- *management of VPC bandwidth* which is further decomposed into:
  - a VPC Bandwidth Allocation MFC
  - a VPC Bandwidth Distribution MFC
- *management of the routing plan* which is placed in a Route Planning MFC
- *network load balancing* which is placed in a Load Balancing MFC
- *performance verification* which is placed in a Performance Verification MFC
- *traffic predictions* which are placed in a Predicted Usage Model MFC

Additionally, the following support MFCs are required:

- a Configuration Management MFC which includes the network model
- a Current Load Model MFC for providing the required network statistics
- a CAC Manager MFC for the TMN to model the CAC behaviour for dimensioning purposes
- a CoS Model MFC

4.3 Mapping to the TMN architecture

The functional architecture is based on the principles of ITU-T recommendation M.3010. Figure 3 shows the allocation of MFCs to OSFs and also places the OSFs into the architectural layers.

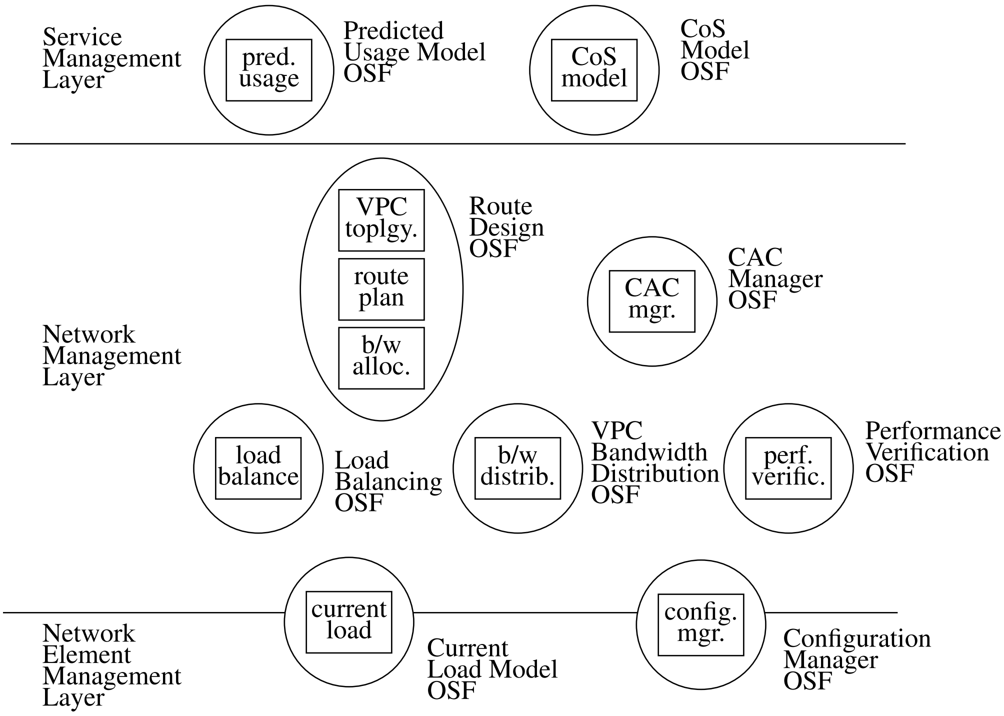

Figure 2 Mapping of MFCs to OSFs and OSFs to the TMN hierarchical layers.

By adopting a hierarchical TMN architecture we take advantage of a centralised management approach in the sense of reducing the placement of intelligence in the managed elements and so burdening their design and eventually their cost. But at the same time we use a hierarchical system to push management intelligence and frequently used management functions as close as possible to the network elements to avoid the management communications overhead inherent in centralised systems.

5 DESCRIPTION OF THE ARCHITECTURAL COMPONENTS

The functionality of the identified OSFs is briefly discussed in this section. The description is at a high level as the paper focuses on architectural rather than design issues. The OSF problems resemble the well known problems of network design, capacity sharing, bandwidth management and routing. However, these problems need to be consolidated and put into the perspective of the proposed architecture.

5.1 Route Design OSF

This OSF has both static and dynamic aspects. The static aspect is related to the network planning activity and is used to initially configure the network in terms of VPCs and routes. This part is performed at network start-up time. The dynamic aspects of its operation cater primarily for changes in the predicted network traffic and for prediction inaccuracies that could not be resolved by the lower level OSFs. As a result, the VPC and ClassRoute networks are reconfigured. The dynamic part consists of the functionality of the VPC Topology, Route Planning and Bandwidth Allocation MFCs.

The Bandwidth Allocation MFC is the first function to be invoked whenever the predicted traffic changes significantly. Based on the predicted usage, the s-d predictions are mapped to VPCs within the existing SDClassRoute networks, and the minimum bandwidth required by each VPC in order to meet the predicted demand is identified. If it is impossible to allocate sufficient bandwidth for the predicted traffic within the constraints of the current SDClassRoute networks and the link capacities, the Route Planning MFC is notified.

The Route Planning MFC attempts to redesign the SDClassRoute networks on the existing VPC network, to remove bottlenecks for example. It tries to increase the number of alternative routes, using the current VPC topology. This process also identifies the new bandwidth requirements on the VPCs. In order to enhance alternative routing and to compensate for inaccuracies in the routing estimates, Route Planning may assign a set of 'back-up' routes to each CoS in addition to the primary set of routes. For a given CoS, the set of 'back-up' routes consists of the routes allocated to the higher quality CoSs. If the Route Planning MFC cannot design a new set of SDClassRoute networks to accommodate the predicted traffic due to limitations in the existing VPC network topology, the VPC Topology MFC is invoked.

The VPC Topology MFC redesigns the VPC network to meet the new requirements. New VPCs may be created to coexist with the current ones and new SDClassRoute networks will be defined so that the new VPC topology may be introduced gradually for new connections. The bandwidth requirements for the VPCs in the final VPC topology are identified and passed down to the lower MFCs. If it is not feasible to design a VPC network to satisfy the traffic demand because of limitations in the underlying physical network, e.g. not enough links, the network planning function is notified.

The Route Design OSF should cater for designing SDClassRoute networks according to the CoS requirements. CoS cell losses targets can be met by adjusting the CAC cell loss targets appropriately so as to ensure that accumulated cell losses over the links of the SDClassRoute network do not exceed those defined for that CoS. Guarantees for delay and jitter can be provided by identifying the maximum number of buffers and switches and ensuring that the SDClassRoute networks do not exceed these values. Finally CoS availability is guaranteed by being an overall optimization constraint that the iterative procedure for defining VPC and SDClassRoute networks should meet.

5.2 VPC Bandwidth Distribution OSF

Taking the current load into account, the VPC Bandwidth Distribution OSF implements the allocation of bandwidth to VPCs as requested by the Route Design OSF. The current load must be considered to avoid situations where the predicted required bandwidth is lower than the current load and hence the new bandwidth allocation would violate the assumptions made by the CAC algorithms in the network and possibly cause excessive cell losses.

In addition to implementing the policies of the Route Design OSF, the VPC Bandwidth Distribution OSF attempts to compensate for inaccuracies in the Predicted Usage Model by distributing any unallocated link bandwidth (viewed as a common pool) among the VPCs.

Unused bandwidth (allocated bandwidth minus current load) in each VPC is the criterion for redistribution to avoid situations where some VPCs are heavily utilised (and consequently there is little bandwidth available for new connections) whilst other VPCs on the same links are lightly utilised. Unused bandwidth is distributed as evenly as possible within certain constraints. For example VPCs can be assigned a class or priority attribute to indicate which VPCs should gain unused bandwidth at the expense of lower priority VPCs. VPCs used for CoSs with low blocking probabilities will be assigned higher priorities.

By varying the averaging interval for calculating the required measures, the sensitivity of the

VPC Bandwidth Distribution OSF can be controlled.

5.3 Load Balancing OSF

The Load Balancing OSF operates within the SDClassRoute and VPC networks defined by the Route Design OSF. The Route Design OSF reserves network resources (VPCs) and indicates their use (by defining routes); The Load Balancing OSF tries to make the best possible use (most efficient utilisation) of the reserved resources.

To achieve this, the Load Balancing OSF takes a network-wide view and tries to influence the routing decisions so that arriving connections use the routes with the highest availability. The view taken is that the routes at a node are prioritised according to their potential as being good routes. Since in our case we deal primarily with a connection-oriented network, potential refers to the availability (spare capacity) of the route to accommodate connections. This way the network load is spread as even as possible and the network availability for new connections is as even as possible (hence the name Load Balancing). Note that the above view is in accordance with the traditional view of routing according to which routing schemes are variants of shortest path algorithms

The multi-class environment that the network operates in, should be taken into account. The Load Balancing OSF should not only aim at optimising routing in the ClassRoute networks but also in the VPC network. This further justifies the need of a central component, like the Load Balancing OSF, which utilising network-wide information about every class tries to harmonise routing within each class and between classes.

5.4 Performance Verification OSF

The Performance Verification OSF is concerned with ensuring that the network meets the performance targets for the different CoSs. This is done at two levels: by monitoring the network and by accepting customer QoS complaints via the service layer's customer interface.

The connection rejection ratios per CoS and per source destination pairs are retrieved from the Current Load Model and compared to the rejection targets as specified per CoS. Customer complaints are analysed and if they are justified the Route Design OSF will be triggered. If CoSs are found to experience connection rejection ratios in excess of the target, an indication is sent to the Route Design OSF to cause the number of routes, or the bandwidth required by the routes, to be updated. The Performance Verification OSF quantifies the performance of the Route Design, the Load Balancing and the VPC Bandwidth Distribution OSFs, being the indisputable measure of their efficiency.

5.5 Predicted Usage Model OSF

This models the predicted usage of the network in terms of the numbers of connections of each CoS required between s-d pairs. The model details how the number of connections changes: hour by hour over the day; day by day over the week; and week by week over the year

Initially this is configured by the service level of the TMN but it is modified by the actual usage of the network via the Current Load Model. This is so that the predicted model becomes more accurate as experience of the usage of the network is gained.

Whenever the predicted load model indicates that the traffic will change significantly the Route Design OSF will be provided with a prediction of traffic for the next time interval. The exact definition of a significant change is a design variable to be experimented with according to the performance of the system as a whole.

5.6 Configuration Management OSF

The configuration manager is responsible for maintaining a consistent model of the physical and logical configuration of the network. It will receive configuration actions from the other OSFs and be responsible for implementing the changes in the network. This task may involve coordination of configuration actions over a number of network elements, for example when a VPC is created. The configuration manager can provide event reports to the other OSFs whenever a configuration action has succeeded.

5.7 Current Load Model OSF

The Current Load Model monitors the network usage and calculates usage statistics according to the requirements of the other OSFs. The Current Load Model is capable of calculating peak, mean, EWMA, etc. statistics according to the specifications of the other components. It will identify the minimum number of network probes and measurements to meet the varied demands of its users.

5.8 CAC Manager OSF

The CAC Manager reproduces the CAC algorithm in the network. When supplied with a traffic mix in the form of a list of the number of connections of each CoS the CAC Manager returns the effective bandwidth of that traffic mix. The calculation has exactly the same result as the equivalent CAC algorithm in the network.

5.9 CoS Model OSF

This models the bandwidth and performance targets for each CoS (see section 3.1).

6 INTERACTIONS BETWEEN THE ARCHITECTURAL COMPONENTS

6.1 Manager-Agent relationship

Figure 3 shows the manager-agent relationships between the derived components.

The VPC Bandwidth Distribution OSF and the Load Balancing OSF are agents of the Route Design OSF. However, their operation is not totally independent, since the effect (in the network) of one of them is taken into account by the other. The VPC Bandwidth Distribution OSF looks at the current load of the VPCs, which is determined by the routing decisions, and the Load Balancing OSF looks at the availability of the VPCs which is determined by the VPC Bandwidth Distribution OSF. This indicates that some coordination needs to exist among them, to avoid possible contradictions.

The Route Design OSF and the VPC Bandwidth Distribution OSF manage the VPC network whereas the Load Balancing OSF determines how to optimise its use. It can be argued that the Load Balancing OSF complements the VPC Bandwidth Distribution OSF, in the sense that it takes advantage of the VPC bandwidth increase.

When the Load Balancing OSF is activated it assumes a stable VPC network. This implies that during the operation of the Load Balancing OSF, the VPC Bandwidth Distribution OSF and the

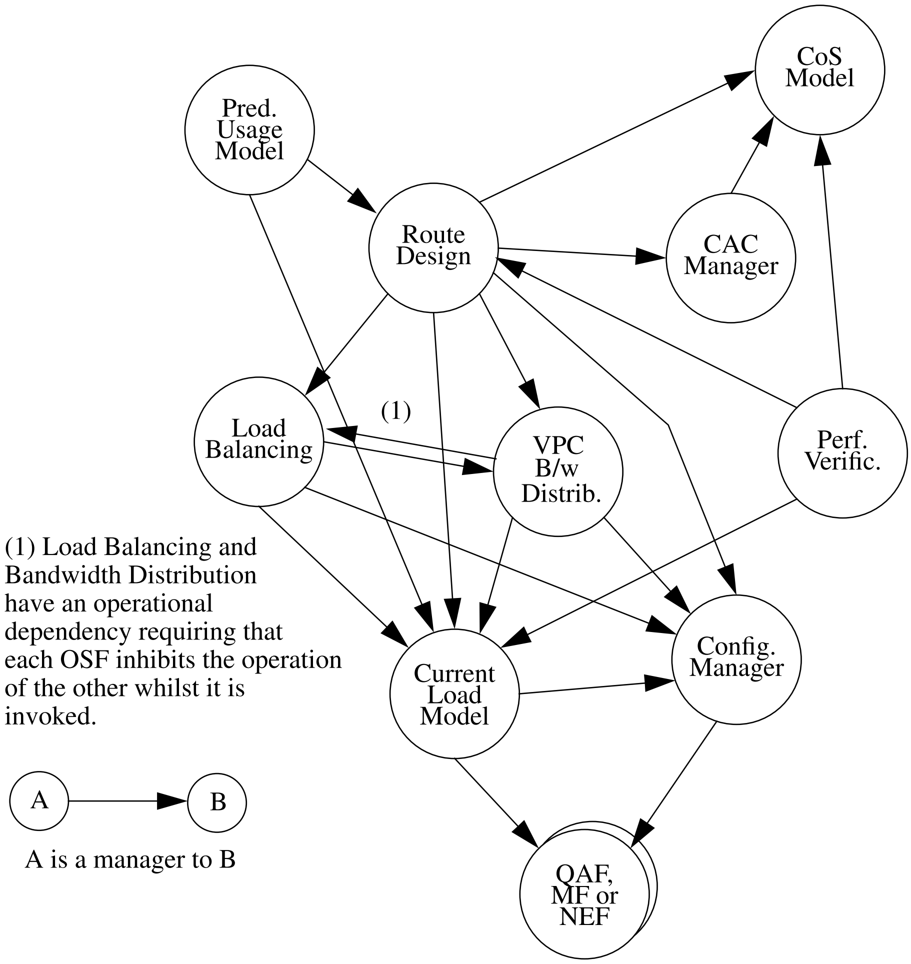

Figure 3 Manager-agent relationships between the OSFs.

Route Design OSF should be prohibited from taking actions. And conversely when the VPC Bandwidth Distribution OSF or the Route Design OSF are about to change VPC bandwidth or topology, the Load Balancing OSF should not be activated until the change has been made.

7 CONCLUSIONS AND FUTURE WORK

In this paper we dealt with a VPC and Routing Management Service for multi-class ATM networks. The proposed system offers the generic functions of performance monitoring, load monitoring and configuration management on ATM networks. In addition, it provides specific functions for routing and bandwidth management in a hierarchical structure. The components in the hierarchy differ in terms of the level of abstraction, complexity and timescale. The management functions to be invoked most frequently are close to the NEs and are as lightweight as possible to reduce management overhead. The more comprehensive functions are placed in the higher levels of the hierarchy and are only invoked when the lower levels are unable to resolve issues within the scope of their functionality and operational parameters. Such a hierarchy provides for continuous refinement of the management decisions and avoids the problems of a fully centralised approach.

The VPC and routing management system provides the following benefits to the network operator:

It allows the network to be used as efficiently as possible within the constraints of the physical resources. It will indicate when the network resources are insufficient for the traffic and hence additional resources need to be deployed. Alternatively it will show when resources are under-used and may be taken out of service or redeployed to avoid congestion elsewhere.

It implements the requirements of the service management layer to provide for users according to the business policy of the network operator. A range of service qualities and types (CBR and VBR) can be implemented for which the service management layer may charge different prices. It designs logical overlay VPC and routing networks so that the different service types can exist on the same physical network.

It distributes load as evenly as possible throughout the network to maximise the network availability and minimise disruptions in the case of failures. It can make dynamic configurations to adapt the network configuration to fluctuating traffic and make changes before they actually happen based on a Predicted Usage Model.

By building intelligence into the TMN the requirements on the NEs are simplified. The TMN functions replace the alternative of elaborate algorithms in the switches that must interact via signalling procedures to allow global network conditions to influence local algorithms. In a multi-class environment the inter-node exchange of routing information is prohibitive simply by the large number of CoSs. Therefore it increases the capacity for revenue earning traffic. By placing these functions in the TMN no additional requirements are placed on the NEs apart from the most basic of management interfaces.

The design is flexible enough to incorporate different algorithms or different levels of functionality to adapt to the specific CAC and RSAs in the network elements. Static algorithms in the elements can be transformed to quasi-static algorithms by TMN actions.

The proposed system can be used for implementing private and virtual private network services since it manages bandwidth reservation and routing within specified performance targets. Provision has been made (see Section 5.4) to provide an abstract interface to the service management functions responsible for the private services to implement their requests.

The architectural framework can be used as a testbed for testing and validating bandwidth management, routing management and load balancing algorithms.

At the time of writing algorithms for the architectural components described in this paper have been developed and the detailed design of prototypes has been completed. This work being undertaken by the RACE II ICM project. A significant portion of the system has already been implemented and demonstrated. Future work includes testing and validation of the components, the system and the architectural concepts on a real ATM testbed provided by another RACE II project (EXPLOIT) as well as in a simulated environment for scalability and extended testing purposes. The information modelling of the interfaces is based on the existing and emerging standards and where necessary, object definitions were expanded and new managed objects were defined. These extensions will be fed back into the standardisation activities.

8 ACKNOWLEDGEMENTS

This paper describes work undertaken in the context of the RACE II Integrated Communications Management project (R2059). The RACE programme is partially funded by the Commission of the European Union.

9 REFERENCES

E.Sykas, K.Vlakos, E.Protonotarios, "Simulative Analysis of Optimal Resource Allocation and Routing in IBCNs", IEE J. Select. Areas Comm., Vol.9, No.3, April 1991.

J.Y.Hui, "Resource Allocation for Broadband Networks", IEEE J. Select. Areas Commun., Vol.6, No.9, Dec.1988.

S.Ohta, K.Sato, "Dynamic Bandwidth Control of the Virtual Path in an Asynchronous Transfer Mode Network", IEEE Trans. Commun., Vol.40, No.7, July 1992.

G.Woodruff, R.Kositpaiboon, "Multimedia Traffic Management Principles for Guaranteed ATM Network Performance", IEEE J. Select. Areas Commun., Vol.8, No.3, April 1990.

Y.Sato, K.Sato, "Virtual Path and Link Capacity Design for ATM Networks", IEEE J. Select. Areas Commun., Vol.9, No.1, Jan.1991.

M.Wernic, O.Aboul-Magd, H.Gilbert, "Traffic Management for B-ISDN Services", IEEE Network, Sept.1992.

H.Saito, K.Shiomoto, "Dynamic Call Admission Control in ATM Networks", IEEE J. Select. Areas Commun., Vol.9, No.7, Sept.1991.

E.Gelenbe, X.Mang, "Adaptive Routing for Equitable Load Balancing", ITC 14 / J. Labetoule and J.W.Roberts (Eds), 1994 Elsevier Science B.V.

ATM Forum, "ATM User-Network Interface Specification", Version 3.0, Sept. 1993.

K. Geihs, P. Francois, D. Griffin, C. Kaas-Petersen, A. Mann, "Service and traffic management for IBCN", IBM Systems Journal, Vol. 31, No. 4, 1992

ITU-T Recommendation I.320 - ISDN protocol reference model

ITU-T Recommendation I.321 - B-ISDN protocol reference model and its application

ITU-T Recommendation I.150 - B-ISDN asynchronous transfer mode functional characteristics

ITU-T Recommendation I.362 - B-ISDN ATM Adaptation Layer (AAL) functional description

ITU-T Recommendation M.3010 - Principles for a telecommunications management network

ITU-T Recommendation M.3020 - TMN interface specification methodology

David Griffin received the B.Sc degree. in Electronic, Computer and Systems Engineering from Loughborough University, UK in 1988. He joined GEC Plessey Telecommunications Ltd., UK as a Systems Design Engineer, where he worked on the CEU RACE I NEMESYS project on Traffic and Quality of Service Management for broadband networks. He was the chairperson of the project technical committee and worked on TMN architectures, ATM traffic experiments and system validation. In 1993 Mr. Griffin joined ICS-FORTH in Crete, Greece and is currently employed as a Research Associate on the CEU RACE II ICM project. He is the leader of the project group on TMN architectures, performance management case studies and TMN system design for FDDI, ATM and optical networks.

Panos Georgatsos received the B.S. degree in Mathematics from the National University of Athens, Greece, in 1985, and the Ph.D. degree in Computer Science, with specialisation in network routing and performance analysis, from Bradford University, UK, in 1989. Dr. Georgatsos is working for ALPHA Systems SA, Athens, Greece, as a network performance consultant. His research interests are in the areas of network and service management, analytical modelling, simulation and performance evaluation. He has been participating in a number of telecommunications projects within the framework of the CEU funded RACE programme.