A Management System for Load Balancing through Adaptive Routing in Multi-Service ATM Networks

Panos Georgatsos
ALPHA Systems S.A.,
3, Xanthou str., 1777 78 Tavros, Athens, Greece.
Tel: +30 1 48 26 014-16, Fax: +30 1 48 26 017
E-mail: panosg@alpha.com.gr

David Griffin
Institute of Computer Science, Foundation of Research and Technology-Hellas,
PO Box 1385, Heraklion 717-10, Crete, Greece.
Tel: +30 81 391722,, Fax: +30 81 391601
E-mail: david@ics.forth.gr

Abstract

This paper proposes a hierarchical approach to the routing issue in ATM-based broadband multi-service networks. We show how distributed Route Selection Algorithms embedded within network nodes may be managed by a TMN system, by virtue of a management service called Load Balancing, to increase their adaptivity to network wide conditions and therefore their effectiveness. The paper elaborates on the issues involved, presents specific algorithms and a TMN-compliant management architecture. The proposed management service influences routing decisions by conveying network-wide information and contributes to network load balancing by directing the routing algorithms to route traffic to the least congested network areas. Results regarding the effect of the proposed management system to network operation are also presented.

1. Introduction

Routing is a critical network design issue with the overall aim of maximizing network throughput in terms of service call admissions, while guaranteeing the performance of the network services within specified levels. Routing in ATM-based networks is based on Virtual Path Connections (VPCs). A route is defined as a sequence of VPCs, where each VPC is defined as a sequence of links being allocated a specific portion of link capacity. Multiple routes for a given source-destination (s-d) pair and for a particular service may be available (alternative routing). Alternative routing is desirable since it reduces the likelihood of blocking, it reduces network vulnerability and enhances routing adaptivity to topological and traffic changes. Experimental studies have verified network performance improvement with alternative routing [3],[14]. The definition of a routing policy involves:

- the definition of a suitable network of VPCs;
- the definition of a Routing Plan comprising a set of admissible routes for each (s-d) pair and CoS (connection type, network bearer service type, making up the network services), based on the defined network of VPCs;
- the definition of a policy for route selection in case of alternative routing.

The design of an efficient routing policy is of enormous complexity, since it depends on a number of variable and sometimes uncertain parameters. The complexity is even greater, considering the diversity of bandwidth and performance requirements of the services that the network must support. The routing policy should therefore be adaptive to cater for traffic and topological changes.

The ITU-T introduced the Telecommunications Management Network (TMN) [20], as a means of provisioning the required network management intelligence and have distinguished between the management and control planes in the operation of communication networks [18], [19]. Following this distinction, the routing functionality introduced previously is spread over the control and management planes as follows [7]: the definition of the VPC network and the

Routing Plan are of concern to the management plane. On the other hand, route selection functionality invoked at call request epochs, is control plane functionality performed within the network itself.

All possible routes for a given (s-d) pair and particular CoS are downloaded by the TMN to network switches where the actual routing decisions are taken at call set-up time. Route selection is done by means of a *Route Selection Algorithm* (RSA). Without loss of generality, it is assumed that RSAs operate on the basis of parameters (*route selection parameters*) associated with the available routes. Following the ideas on the taxonomy of routing algorithms [11], [12], several types of RSAs can be distinguished according to the selection method they employ, the information they utilize and the degree of adaptivity they offer.

According to the selection method employed, a RSA can be: deterministic, whereby route selection is made according to a predefined order, random, whereby route selection is made based on probabilistic criteria, and locally adaptive, whereby route selection is made based on a policy taking into account the current load on the VPCs, as seen locally (at the switch level).

The information that a RSA utilizes may be: local information, global network information, or no information at all. According to its degree of adaptivity (rate at which used information is renewed), RSAs can be: static (no adaptive at all) or dynamic. Another parameter associated with adaptivity is how adaptivity is provided. It can be provided through inter-node exchange, periodically or at exception, or only locally at connection acceptance/release times, or from TMN periodically or at exception.

Examples of random, dynamic (at the order of connection acceptance/release) RSAs without requiring any information are the Dynamic Alternate Routing (DAR), Linear Reward Penalty and Linear Reward Inaction algorithms, proposed for telephone traffic routing [6], [8]. An example of dynamic deterministic algorithm is the DNHR algorithm used by AT&T long distance telephone network [15].

The adaptivity of RSAs should not be confused with the quasi-adaptive nature of the Routing Plan. The Routing Plan has been constructed on the basis of predicted network usage; and it is redefined whenever, significant changes in network predictions are verified. The adaptivity of RSAs refers within the time-frame of network usage predictions, where the Routing Plan is stable. Such adaptivity is desirable since it compensates for inaccuracies in traffic predictions and/or network usage fluctuations around the predicted values.

The above analysis indicates that there is scope for RSA management and proposes that the issue of routing management encompasses two levels: a higher level for the management of the Routing Plan and a lower level for the management of RSAs. This view was first suggested in [7] where a hierarchical management system for VPC and routing management was proposed. Within this framework, the paper focuses on policies for managing RSAs, assuming a given VPC network and a specified Routing Plan.

There is a significant research in the area of network routing and the problem of RSA management has been tackled in the overall context of routing algorithms (e.g. [2], [3], [5], [6], [10]-[14]). However, the majority of these studies do not take into account the different bandwidth and performance requirements of the multi-class network environment. Moreover, these studies do not address the issue of RSA management in the overall context of network management and they do not offer a clear distinction between the management and control plane functionality. There is an emerging trend [4], [7], [16], [17], to move towards the automation of the monitoring, decision making and configuration management loop by enhancing the intelligence of the management functions.

Recognizing the need for enhanced network management systems and adopting the framework of routing management presented previously (cf. [7]), the paper proposes a management architecture and specific management algorithms for managing the RSAs run in the network switches.

A TMN approach is adopted. Given the VPC network and the Routing Plan, the paper defines an appropriate management service, the Load Balancing (LB) management service, for RSA management. The proposed management service, taking a network-wide view, makes the RSAs network-state adaptive, by conveying global network information and contributes to network load balancing, by regulating load distribution. The paper proposes specific algorithms for route selection management taking into account the wide range of traffic types coexisting in broadband multi-service networks. The proposed algorithms are built around the concept of ‘route potentiality’ (for setting-up new connections). The paper also presents a management architecture fulfilling the objectives of the LB management service.

The paper is organized as follows: Section 2 puts the LB management service into the perspectives of network management. Section 3 describes its functional principles, proposing specific algorithms. Section 4 presents a TMN-compliant functional architecture. Section 5 presents comprehensive performance results and section 6 presents the conclusions and highlights aspects of future work.

2. The Load Balancing management service

Within a multi-class ATM network environment, the objectives of the LB management service are:

- to manage the RSAs run in the network switches according to network-wide traffic conditions;
- to monitor the network with the purpose to provide warnings of deterioration in route availability and load deviations at link level, indicating inefficient use of transmission facilities.

The LB management service operates within a defined set of VPCs and routes per CoS. The scope for RSA management has been introduced in the previous section. It aims at making efficient use of the network resources defined for routing (VPCs) and it is achieved by tuning the route selection parameters. Taking into account the multi-class network environment and the fact that the routes of all network CoSs share the same VPC infrastructure, route selection management should not aim at making efficient use of the VPC resources only at the level of a single CoS, but at the overall CoS level.

The LB management service is regarded as a component of an overall routing management service. In [7] the above LB management service has been proposed as the lower level of a two-level hierarchy for routing management (see figure 1). The higher level of the routing management hierarchy, comprising the Routing Plan component, operates at epochs where network usage predictions change, producing new sets of routes per CoS, based on the current set of VPCs. The lower level, comprising the LB component, operates within the time-frame of network usage predictions, and based on the actual network usage influences route selections within the network of routes established in the higher level. The lower management level is introduced to compensate for inaccuracies in network usage predictions and sudden fluctuations of the load around the predictions.

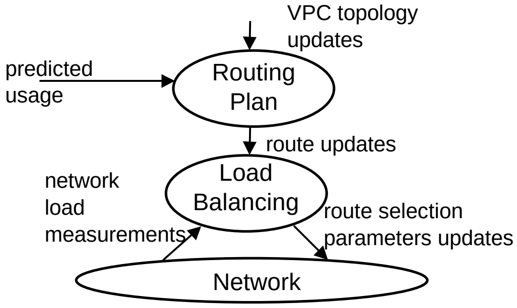

Figure 1: A two-level hierarchical system for routing management.

The LB management service contributes to the efficient operation of networks from several aspects which in turn further justify its existence. Through its actions LB makes routing decisions network-state adaptive. Network-state adaptive routing has been recognized as a useful merit of routing algorithms as it is proved by the huge quantity of literature in the subject; indeed, network performance improvement has been verified under adaptive routing [3], [8], [11], [12]. Moreover, through RSA management, distribution of network load may be regulated; therefore enabling network load balancing. Balanced networks have been widely accepted as a valid objective of network design and routing policies [5], [9]. Furthermore, by aiming at influencing the routing decisions so that the most advantageous route is selected, the signaling overhead and hence the connection set-up time is reduced. Apart from its active role in routing management, LB contributes to preventive management as well. By taking a future perspective, it notifies the management functions responsible for route definition, of undesirable trends in network availability (to accept new connections).

The LB management service belongs to the performance management functional area. Figure 2 shows the relationship of the LB management service with other management services/components and the TMN users. The boundaries of the management responsibility of the LB management service are shown in Figure 3 which depicts the interactions between the management and control planes from the point of view of LB management service.

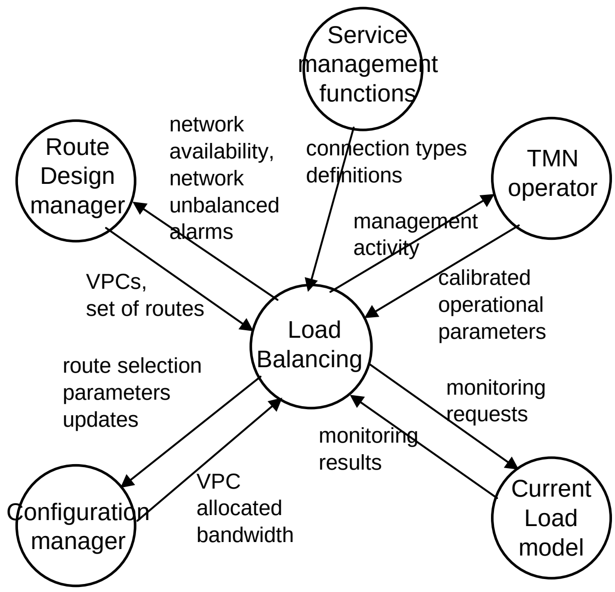

Figure 2: Enterprise view of the LB management service.

It should be stressed that the introduction of a management service like the LB management service, does not make the RSAs obsolete nor does it imply that the management plane is involved in routing decisions at call set-up times. On the contrary, it enhances them by conveying network wide information. By placing it in the management plane, the network switches are relieved from the burden of processing their routing information as well as the inter-node exchange of routing information is avoided. Therefore, the required routing intelligence in the network switches is reduced, resulting in faster routing decisions; an essential target of future broadband networks. In this sense, the LB management service implies a semi-dynamic type of routing policy combining the merits of centralized and decentralized (based on local information) routing policies. Semi-dynamic routing policies have been introduced since the traditional data networks and improvement in network performance has been confirmed under such schemes [11], [14].

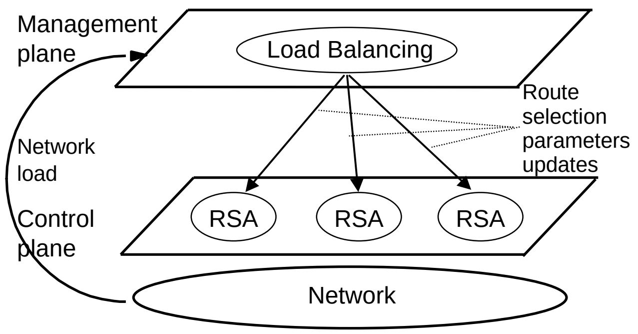

Figure 3: Management and Control plane interactions.

3. The Load Balancing algorithm

This section describes the main aspects and proposes specific algorithms for the functionality of the LB management service. As implied by its objectives, the LB functionality encompasses: management of RSAs and network load surveillance for reporting network unavailability and link load deviations.

3.1 The route selection management algorithm

The essence of the proposed algorithm is to assign to routes a figure of merit and subsequently influence the RSAs to route traffic towards routes of higher figures of merit. This view is in accordance with the traditional view of routing according to which routing schemes are variants of shortest path algorithms [12].

In the context of connection-oriented networks, like ATM-based broadband networks, it is natural to consider that route merit refers to the *potentiality* of the route to accommodate new connections. In traditional data networks route merit relates to the network delay implied by the route. The route figure of merit should be a function of the spare capacity along the route and it should take into account the fact that different CoSs may share the same (part of) route(s).

Adopting the above approach, *route potentiality* is calculated for all possible routes that exist between a given network switch and a particular network destination, for each CoS.

The routing information available at a network switch associates a particular network destination and CoS with a VPC starting from the switch. Therefore, route selection in fact refers to the selection of a particular VPC. Considering a network switch and a specific VPC starting from this switch, for a given network destination and CoS, more than one route may have been defined on this VPC; all these routes use this VPC as an exit from this node. The potentialities of all these routes can therefore be accumulated, giving rise to a figure of merit of selecting this VPC as the next step in the route. The figure of merit of VPC selection reflects the potentiality of the network to accommodate new connections in the route(s) starting from this VPC.

The VPCs at each switch are therefore graded with a figure of merit, *VPC selection potential*, according to the potentiality of the routes their selection indicates. The algorithm then recommends VPCs for routing according to their figure of merit (selection potential). This is done by setting appropriately the route selection parameters so that VPCs with the higher figure of merit have advantage over the ones with lower figure of merit. In case of deterministic RSAs, VPCs are prioritized in the order of their merit; and RSAs make the selection according to this order. In case of random RSAs, VPCs with higher figure of merit are assigned higher frequency. In case of locally adaptive RSAs, VPCs are classified into equivalent groups according to the significance in the difference of their figure of merit; VPC selection is done in the order of the equivalent sets and by applying local criteria for the VPCs within a set. The latter routing policy enhances the concept of δ-routing [11] proposed for data networks.

It could be argued that with the proposed algorithm for route selection management, network load is spread as evenly as possible, therefore network availability for new connections is as even as possible. Hence load balancing is obtained. Moreover, adaptivity to network conditions is achieved, since network load is taken into account in the calculation of route potentialities.

It should be noted that the proposed algorithm is not simply a widest path routing algorithm trying to route traffic over routes with the highest potentiality. It is a *highest potentiality(HP) path* routing algorithm, trying to achieve routing over the network part(s) that have the highest potentiality to accommodate new connections. Figure 4 outlines this point. Under widest path routing, VPC A corresponding to route R1 should be selected for going from node 1 to node 6. Under highest potentiality path routing, VPC B, corresponding to route R2 or R3, is selected. Therefore, under highest potentiality path routing, full advantage of route alternatibility, not only locally (at the vicinity of a switch) but also remotely, is taken. In this sense highest potentiality path routing outperforms widest path routing.

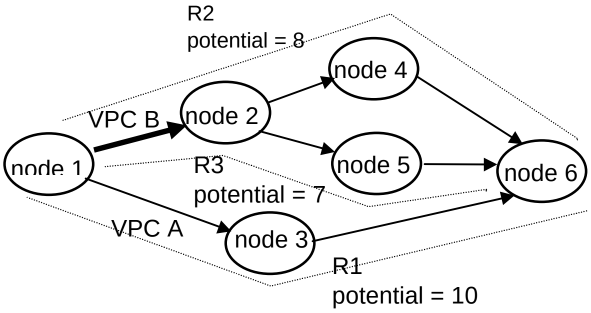

Figure 4: Widest path vs. HP routing.

In the following, we formulate the notions of *route potentiality* and *VPC selection potential* introduced previously and propose specific formulas for their calculation. *Route potentiality* is defined in terms of the notion *VPC acceptance potential* which is introduced next.

For a given VPC, say v, *VPC acceptance potential*, denoted by $VPt_a(v; c)$ for CoS c, is defined as the number of VCCs of this CoS that can be potentially accommodated in the VPC, taking into account the current load on the VPC. The fact that a number of CoSs may share the same VPC in their routes needs also be taken into account. The following heuristic is proposed for their calculation.

Considering a VPC, the $VPt_a(.)$s are calculated as solutions of the following linear system with respect to k(.)s:

$$\sum_{i=1}^{C} \delta(i)B(i)k(i) = S \qquad (1)$$

$$\frac{B(i)\,k(i)}{B(j)\,k(j)} = \frac{g(i)}{g(j)}, \; i \neq j,\; i,j = 1...C,\; \delta(i)=\delta(j)=1 \qquad (2)$$

where

δ(i) is a Boolean taking the value 1 if CoS i uses the VPC in its routes, and the value 0 otherwise.

B(i) is an estimate of the VPC bandwidth that CoS i will consume when it is accepted on the VPC. It can be the mean or the peak bandwidth requirement of CoS i, or it can be its effective bandwidth as calculated by the CAC network algorithm.

S is an estimate of the VPC spare bandwidth. It is recommended to be in the form of a moving average and not in the form of instantaneous value for reducing sensitivity to traffic fluctuations and increasing estimate accuracy.

C is the number of the different CoSs supported by the network.

g(i) are weights differentiating the access of each CoS on the VPC. They reflect the frequency with which CoS i is using the VPC.

The weights g(.)s are closely related to the Routing Plan. They can be calculated either dynamically, from network measures, or statically, directly from the Routing Plan. In the latter case, the following is proposed:

$$g(i) = \sum_{j=1}^{O(i)} 1/a(j) \qquad (3)$$

where

O(i) is the number of selection occasions on which CoS i may select the VPC. Note that for a given CoS, a particular VPC may belong to routes to one or more destinations. Therefore, O(i) equals the number of possible destinations that CoS i may reach through the given VPC.

a(j) is the alternatibility factor i.e. the number of routing alternatives for routing occasion j.

Equation (1) is intuitively evident, taking into account that a VPC can accommodate connections of different classes. Equation (2) says that the ratio of the bandwidth to be consumed by two different CoSs is taken to be proportional to the visit ratio of these CoSs to the VPC; note that the product B(i)k(i) is the amount of VPC's bandwidth to be given away to CoS i connections.

The system of the equations (1), (2) yields the following solution

$$VPt_a(v;i) = \frac{g(i)}{B(i)} \frac{S}{\sum_{j=1}^{C} g(j)}, \; i=1..C \qquad (4)$$

As it can be seen from (4), the *VPC acceptance potential* for a CoS depends on the bandwidth characteristics of the CoSs, on the spare bandwidth of the VPC and on the alternatibility with which the VPC is used for routing. Note that because the alternatibility with which a CoS uses a VPC has been taken into account, a CoS with alternative routes is discouraged to occupy a VPC at the expense of the CoSs that use that VPC as a unique option.

Having defined the notion of *VPC acceptance potential*, the notion of *Route potentiality* is defined next.

For a given route, say r, *Route potentiality*, denoted by RPt(r; c) for CoS c, is defined as the number of VCCs of this CoS that can be potentially established in the route, taking into account the current load of the VPCs along the route. The fact that a number of CoSs may share parts of the route needs also be taken into account. *Route Potentiality* is defined in terms of *VPC acceptance potential* as follows:

$$RPt(r; c) = \min \{ VPt_a(v;c) \ \forall v \in V(r) \} \qquad (5)$$

where V(r) denotes the set of VPCs that make up route r.

Finally, the notion of *VPC selection potential* is defined. Considering a VPC, say v, starting at a specific network node, say n, *VPC selection potential*, denoted by $VPt_s(v; c,d)$ for CoS c and destination node d, is defined as the number of VCCs that can be potentially established in all possible routes to the particular destination starting with this VPC. It can be defined as follows:

$$VPt_s(v; c,d) = \min \{ VPt_a(v; c), \sum_{r \in R(n,d;c)} RPt(r;c) \} \qquad (6)$$

where R(n,d;c) denotes the set of all routes from node n to destination node d defined for CoS c.

3.2 Network load surveillance

Based on the analysis presented in the previous section, network availability for new connections can be estimated by extending the notion of potentiality at the node level.

For a given network node, say n, *Node Potentiality*, denoted by NPt(n; c,d) for CoS c and destination node d, is defined as the number of VCCs that can be potentially established from node n to node d over all possible routes starting at node n, taking into account the load in the network i.e.

$$NPt(n; c,d) = \sum_{v \in V(n)} VPt_s(v;c,d) \qquad (7)$$

where V(n) denotes the set of VPCs starting from node n.

Considering access nodes, the above formula provides a measure for network availability for specific source-destination pairs and CoSs. Load deviations at the link level can be measured by calculating the difference of link utilization from the network-wide average value.

Warnings are emitted as threshold crossings events on the previous measures. The threshold values as well as the parameters of the measurements (e.g. moving average method, observation period) are provided as input to the LB management service by the Routing Plan management components.

3.3 The Load Balancing algorithm

The bullets below summarize the previous algorithms, offering a complete view of the LB algorithm.

- For each VPC calculate its acceptance potential for each CoS (cf. (4)).
- For each CoS and node from which there are more than one VPCs to a given destination.
  - Find the HP paths and calculate the VPC selection potential for each destination (cf. (6)). A centralized algorithm for finding HP paths can be found in [22].
  - Grade the VPCs and determine the new values of the associated route selection parameters and if required send appropriate management actions.
- Calculate node potentiality measure (cf. (7)) and link load deviation and emit appropriate alarms if necessary.

4. Functional architecture

Adopting a TMN approach [20], this section presents the functional architecture of the LB management service. Based on the methodology of the ITU-T Recommendation M.3020 [21], the LB management service is decomposed into management service components (MSCs) which in turn are decomposed into management functional components (MFCs) which are then mapped to the layers of the TMN hierarchy. The following decomposition into MCSs and MFCs is proposed:

- a *network load balance* MSC consisting of two MFCs: the RSA Management MFC and the Link Load Deviation Surveillance MFC;
- a *configuration manager* MSC, consisting of a single MFC which includes the network model;
- a *current load model* MSC, consisting of a single MFC which is responsible for providing the required statistics;
- a *connection type model* MSC, consisting of a single MFC which is the repository of the connection classes that the network supports.

The first MSC is specific to the LB management service. The latter three MSCs are generic and can be considered as individual MSs in their own right. Figure 5 shows the allocation of MFCs to OSFs and their allocation to the TMN architectural layers.

The functionality of the LB management service has been placed at the network management layer following the directives implied by the decomposition of the logical TMN architecture; for, it concerns with management of network entities, relating information from a number of network elements. Alternatively, it could be placed at the network element management layer, requiring the existence of a distributed algorithm for finding HP routes. In this case, the interactions between the network and network element management layers for conveying network measures would be substituted by the interactions between the network element management layers for the exchange of the information required by the distributed algorithm. Additionally in this case, the interactions with the other components of the network management layer (e.g. the ones responsible for managing the set of VPCs and the routes) as well as the cost of meta-management (e.g. s/w maintenance) would be increased. The magnitude of these tradeoffs depends on the TMN transmission infrastructure and the physical location of the TMN hosts. The decision as to which architectural option to choose is therefore left open to the TMN system designers.

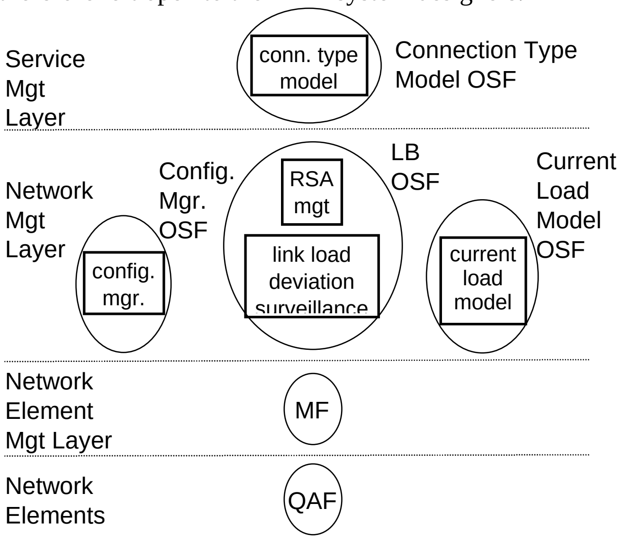

Figure 5: LB TMN functional architecture.

5. Results

The performance of the proposed LB management system at network level has been assessed. The performance tests were carried out in the TMN testbed developed by the RACE II R2059 project ICM (Integrated Communications Management), using an ATM network simulator. The testcases considered covered a number of different network configurations regarding logical topologies and network offered load patterns.

The results indicated network performance improvement under the proposed LB management system. Improvement is achieved in terms of connection rejection ratios and load deviation at link level. In a multi-class network environment the proposed LB management system balances connection blocking probabilities per CoS, so that to improve overall network performance in terms of *offered bandwidth loss probability*. Marginal connection blocking probabilities are also improved for some CoSs, especially for those of higher bandwidth requirements. The notion of offered bandwidth loss probability is defined as the aggregate connection blocking probability, where the marginal blocking probability of each CoS is weighted with a weight proportional to its bandwidth requirements. Considering that whenever a connection is rejected, the network looses an amount of revenue proportional to the bandwidth requirements of the connection, this measure is indicative of the losses in network revenue, considering that each connection is charged only on the basis of its bandwidth requirements. This result shows how routing management activities may be used to achieve the business objectives of the network operators.

Furthermore it was shown that the gains in network performance under the LB management system depends on the following factors: route alternatibility between the route networks of each (s-d) pair and CoS, route sharing and the activation epochs of the LB functionality.

With respect to route alternatibility it was verified that the higher the alternatibility in the routes, the higher the performance gains. However, performance deteriorates in cases where there is asymmetry in the alternatibility amongst the defined route networks per CoS and (s-d) pair, especially when combined with traffic generation asymmetry in favor of the (s-d) pairs with the higher alternatibility in their route networks. With respect to route sharing, it was shown that network performance under the LB management system improves more in cases where there is partial sharing of the defined routes amongst the CoSs. With respect to the LB activation epochs, it was verified -as expected- that the triggering mechanism should be asynchronous based on actual network conditions. Moreover the results indicated that the intuitive argument that the more frequent LB is activated the better the performance will be, it is not generally true. The LB system should base its functionality on some sort of predictions regarding anticipating network usage. These predictions should refer within the time-frame of the predictions based on which the VPC and route networks were built. From the insight gained from the results, the information that should be utilized in making such predictions should concern actual VPC usage and actual number of active connections per (s-d) pair and CoS. The latter information, combined with the maximum number of active connections that the network should provide per (s-d) pair and CoS (derived by the Routing Plan management functions), should enable the LB functionality to make medium-term predictions regarding the anticipating offered load per (s-d) pair and CoS. Details on the results can be found in [23].

6. Conclusions and future work

In this paper we dealt with the issue of load balancing in multi-service ATM multi-service networks. Adopting a hierarchical framework for routing management, the issue of load balancing was tackled through adaptive routing, namely through management of the route selection algorithms (RSAs) which run in the network switches and actually take the routing decisions. The paper introduced the LB management service for RSA management. The proposed management service, taking a network-wide view, makes the RSAs network-state adaptive by conveying global network information and contributes to network load balancing by regulating load distribution.

The paper proposed specific algorithms for route selection management taking into account the wide range of traffic types coexisting in broadband multi-service networks. The proposed algorithms are built around the concept of 'route potentiality' (for setting-up new connections). The routes with the highest potentiality are established and subsequently are recommended for routing to the RSAs. The paper also elaborated on a management architecture fulfilling the objectives of the LB management service. By placing the LB functionality in the management plane, the switches are relieved from the burden of managing their RSAs. Note that this is essential in IBCN since the switches should incorporate very fast decision algorithm without causing any routing overhead in the network. The proposed architecture and algorithms has been prototyped, demonstrated and tested in real and simulated network environment. Results regarding the performance of the proposed system have also been presented. The results showed network performance improvement under LB, in terms of connection rejection ratio and load deviation at link level.

Future work is concerned with further testing of the proposed system both at architectural and algorithmic levels. Other aspects of future work include research on some functional issues such as: determining appropriate activation periods or conditions and refinement of VPC metrics for deriving route potentiality. The results already taken so far encourage the undertaking of such tasks.

Acknowledgments

The work described in this paper has been carried out by the authors in the course of the RACE II Integrated Communications Management (ICM) project (R2059). The RACE II programme is partially funded by the Commission of the European Union. The authors wish to acknowledge Peter Baxendale (of University of Durham), Andy Carr (of Cray Communications Ltd), Kostas Kassapakis (of ALPHA Systems S.A.) and Bruno Rossi (of ASCOM/Monetel) for designing and implementing the components of the proposed management system, and Nikos Balatsoukas (of University of Athens) for carrying out the performance tests.

References

1. A.Aho, J.Hopcroft, J.Ulman "*Data structures and algorithms",* Addisson-Wesley publ. com., 1983.
2. G.Ash, B.D.Huang "*An analytical Model for Adaptive Routing Networks*", IEEE Trans. on Comm., Vol.41, No.11, Nov.1993.
3. R.R.Boorstyn, A.Livne "*A technique for adaptive routing in networks*", IEEE Trans. on Comm., Vol.29, 1981.
4. K.Geiths, P.Francois, D.Griffin, C.Kaas-Petersen, A.Mann "*Service and traffic management for IBCN*", IBM Systems Journal, Vol.31, No.4, 1992.
5. E.Gelenbe, X.Mang "*Adaptive Routing for Equitable Load Balancing*", ITC 14/ J.Labetoule and J.W.Roberts (Eds), 1994 Elsevier Science B.V.
6. A.Girard, S.Hurtubise "*Dynamic Routing and Call Repacking in Circuit-Switched Networks"*", IEEE Trans. on Comm., Vol.31, ,No.12, Dec.1983.
7. D.Griffin, P.Georgatsos "*A TMN system for VPC and routing management in ATM networks*", Integrated Network Management IV, Proc. of 4rth intern. symposium on integrated network management, 1995, ed. Adarshpal S. Sethi, Yves Raynaud and Fabienne Faure-Vincent, Chapman & Hall, UK, 1995.
8. N.Eshragh, P.Mars "*Study of dynamic routing strategies in circuit-switched networks*", 3rd UK comp. and telecom. perf. engin. workshop, Edinburgh, Sept.1987.
9. A.Kershenbaum, P.Kermani, G.Grover "*MENTOR: An Algorithm for Mesh Network Topological Optimization and Routing*", IEEE Trans. on Comm., Vol.19, No.4, April 1991.
10. K.R.Krishnan "*Markov Decision Algorithms for Dynamic Routing*", IEEE Communications Magazine, Oct.1990.
11. H.Rudin "*On routing and 'delta routing': A taxonomy and performance comparison of techniques for packet-switched networks*", IEEE Trans. on Comm., Vol.24, No.1, 1976.
12. M.Schwartz, T.E.Stern "*Routing techniques used in computer communication networks*", IEEE Trans. on Comm., Vol.28, 1980.
13. E.Sykas, K.Vlakos, E.Protonotarios "*Simulative Analysis of Optimal Resource Allocation and Routing in IBCNs*", IEEE Journal Select. Areas Comm., Vol.9, No.3, April 1991.
14. T.P.Yum, "*The Design and Analysis of a Semidynamic Deterministic Routing Rule*", IEEE Trans. on Comm., Vol.29, No.4, April 1981.
15. Y.Watanabe "*Dynamic Routing Schemes for International Networks*", IEEE Comm. Magazine, Oct.1990.
16. M.Wernic, O.Aboul-Magd, H.Gilbert "*Traffic management for B-ISDN Services*", IEEE Network, Sept.1992.
17. G.Woodruff, R.Kositpaiboon "*Multimedia traffic management principles for Guaranteed ATM Network Performance*", IEEE J. Select. Areas of Comm., Vol.8, No.3, July 1992.
18. ITU-T Recommendation I.320 "*ISDN protocol reference model*"
19. ITU-T Recommendation I.321 "B-*ISDN protocol reference model and its application*"
20. ITU-T Recommendation M.3010 "*Principles of a telecommunications management network*"
21. ITU-T Recommendation M.3020 *"TMN interface specification methodology"*
22. P.Georgatsos, D.Griffin "Load Balancing in Broadband Multi-Service Networks: A management Perspective", Procs of 3rd int. workshop on Performance Modelling and Evaluation of ATM networks, Bradford, UK, July 1995.

23. ICM Deliverable 21 *"Integration, Performance Verification and Evaluation of Results"*, Feb.96.