An Integrated Approach to Switched VC ATM Restoration in the REFORM system

Yves T'Joens (Alcatel Bell, Belgium), Panos Georgatsos (AlgoSystems, Greece), David Griffin (University College London, UK), Per-Thomas Huth (Telenor, Norway) , Leonidas Georgiades (University of Thessaloniki, Greece), George Pavlou (Univesity of Surey, UK), Dimitris Manikis, Eleni Mykoniati (National Technical University of Athens, Greece)

Alcatel Corporate Research Center, Francis Wellesplein 1, B-2018 Antwerp, Belgium, tel:+3232407890, fax :+3232409932, e-mail : yves.tjoens@alcatel.be

Abstract: This paper presents the results of the ACTS-REFORM project. This project has built a system for the provisioning of reliable ATM VC switched services with QoS guarantees on a given and fixed physical network topology. We discuss the overall system specification where advanced network planning and adaptive network operation have been integrated with a high speed protection switching scheme in order to optimize network operation in a cost effective manner. We further highlight some of the most significant results obtained by testing the system on diverse testbeds and simulation environments.

1. Introduction

Today, an increasing amount of operators start commercial deployment of ATM switched services, fuelled by the clear need for service integration on their networks and the introduction of ATM based access technologies such as xDSL and PON. As inherent to commercial service offerings, core network reliability is a key requirement to network operation. The term reliability, as used throughout this paper, encompasses both survivability and availability aspects. Today's operators are less and less interested as to the specifics of a single technology for restoration of traffic, but turn to manufacturers and integrators for total network solutions for resilient commercial deployment of ATM services. In this context the REFORM project has specified, built and demonstrated its total integrated concept for ATM layer resilience.

This paper presents the specifications and experimental results of the REFORM system, where advanced network planning and adaptive network operation are integrated with a high speed protection switching scheme in order to optimize in a cost effective way the provisioning of reliable VC switched services, given an operator investment in physical network topology.

Chapter 2 introduces the requirements and constraints to which the system has been built and presents the functional architecture of the REFORM system. Chapters 3 and 4 discuss the planning and VC routing architecture respectively, while chapter 5 explores the bandwidth distribution in the planned VP technology. Chapters 6 and 7 discuss respectively the restoration and full dynamic behaviour of the REFORM system. Chapter 8 concludes the paper, by summarising the main results.

2. Resilient ATM networks, the REFORM approach

Considering that the main threats to network integrity are topology and traffic variations, the main functional requirement underlying the REFORM system, is to offer a complete and integral treatment to the problem of network resilience. As such, the REFORM system was designed with the objectives to cover both control and management planes of network operation, during all phases of the network's day to day operation (including network failures). A detailed account of the REFORM system functional and non-functional requirements is presented in [D1].

To the above end, the REFORM system functional model provides an implementation-independent view of the required functionality and their interactions for ensuring network resilience (reliability) in a cost-effective manner. The term resilience encompasses two aspects: availability and survivability. Availability is mainly concerned with minimising the call/connection blocking probabilities as perceived by the user, that is of optimising network throughput in terms of successful call/connection attempts, under QoS constraints. Given a physical network topology, this is realised by optimally planning the ATM VP layer (based on medium to long term anticipated network usage predictions), and by intelligent adaptive routing and dynamic VPC bandwidth management functions to cater for small to medium term traffic fluctuations (around predicted values) in actual network usage. Survivability on the other hand, refers to the ability of the network to recover existing connections from outages and builds upon distributed protection switching mechanisms at the Virtual Path (VP) layer.

Figure 1 depicts the REFORM system functional model. The REFORM network architecture operates in a hierarchy, at different levels of abstraction and time-scale, offering different degrees of adaptivity (from proactive to reactive response) to changing traffic and topology conditions. More details on the behaviour and interactions of the functional components and their algorithms can be found in [D4], [D14].

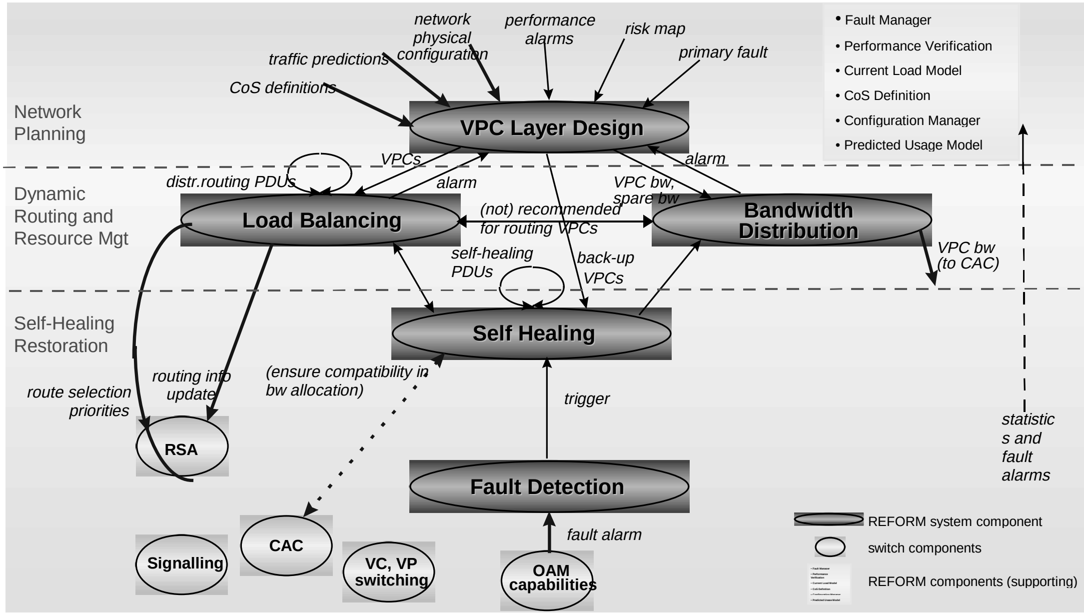

Figure 1 : The REFORM functional model

The next sections introduce the system's key functional components and present experimental results.

3. Network Planning

Network initialisation starts with planning the resources necessary for service provisioning.

The planning process consists of the following components

Information Input. The input the operator has is the physical topology, the Class of Service (CoS) model and the traffic predictions associated with the class. Specifically, each connection belongs to a CoS and with each CoS certain performance objectives (blocking probability, bandwidth, etc) are associated. The statistical predictions of each CoS refer to the arrival rate at each origin-destination pair using the specified CoS, bandwidth and its duration.

Route Design**.** This component determines the candidate routes for each origin-destination pair so that the performance objectives of each CoS are satisfied. The route determination takes into account the CoS statistics.

Working VPC Design**.** This component builds VPCs that are able to accommodate the traffic induced on the routes determined by the Route Design component. The Working VPC design is based on the effort to satisfy a number of conflicting objectives, such as call set-up delay, cell processing load, restoration processing load etc.

Protection VPC Design**.** These component designs VPCs to which traffic from failed working VPCs is redirected in case of network component failures. The restoration objective of the design is to provide a set of Protection VPCs that are able to provide protection and complete restoration in case of

any single link failure. The restoration objective allows for bandwidth sharing among protection VPC, and this feature is used to minimise the restoration bandwidth needed by the network.
The output of the network planning process is the VP topology, both for the active traffic paths and the backup resource paths, as well as a set of what is denoted ‘admissible’ routes, indicating the set of optimal routes between ingress and egress switches of the domain. Traffic between each ingress and egress point is following one of these routes. The specific route to be followed is based on the actual system load and is done in an adaptive fashion, as will be explained in Section 4.
The network planning process has been implemented, validated and tested on a number of random test networks, the project testbed, as well as on a real European network topology (COST 239) consisting of 19 nodes and 37 links. The execution time of the algorithms was in the range of minutes, which is quite acceptable for off-line algorithms. It was found that the bandwidth sharing property of protection VPCs can be exploited to provide bandwidth-efficient protection of working VPCs. For example, in the network topology of the COST239 network, the required total protection VPC bandwidth was about 65% of the total bandwidth of the protected working VPCs. Note also that the protection VPC bandwidth is needed only in the case of link failures. Under normal conditions the protection VPC bandwidth can be used to carry non-protected traffic, such as best-effort.

4. Adaptive Routing

In the REFORM system, a hybrid routing scheme is adopted. The centralised part (see network planning) defines the admissible (explicit) routes per source-destination and CoS (distinct class of service) so that the QoS requirements of the CoSs and certain network-wide cost-effective criteria are met.
The distributed part, executed in the ingress VC switches selects appropriate routes[1] for new connections according to their routing table at call set-up times. The objective of introducing dynamic routing is to manage the route selection process used by the VC so as to balance the load over the VP layer network as evenly as possible. The management of the routing tables makes the routing decisions network state dependent, providing the desired levels of adaptability to avoid congested and/or damaged network areas. The network topology information required for routing purposes is disseminated by means of a distributed Topology Update Protocol (TUP). The TUP used within the REFORM system [D7], [D14], has been derived from the PNNI version 1 specification, however adopted to carry the Class of Service information (denoting the distinct QoS-based service classes offered by the network), and limited to single level operation.
Due to the statistical nature of the input (traffic predictions) and the errors inherent to the involved planning optimisation tasks, variations in traffic over certain source destination pairs may be experienced during network operation. The actual load on some explicit routes and therefore VPCs may experience higher or lower loads as compared with the initially planned load. It is then the role of the adaptive routing function to dynamically adjust to actual traffic conditions, by appropriately influencing the routing decisions, so that to continuously optimise network throughput in terms of connections accepted. Furthermore, if the ingress VC switches should see that for certain destination/CoS combinations the network has no longer resources to accommodate new service request, it issues appropriate alarms to the network planning layer, since this indicates potential increases in call/connection blocking.
The adaptive routing algorithm in the REFORM system is a QoS-based multi-class, multi-path, link-state-dependent, distributed routing scheme, based on the *highest potential area (HPA)* routing concept [D7], [D14]. Assuming that the ability to accept new connections (per CoS) over a specific route could be quantified, the *highest potential area* routing concept suggests that a coming connection will not be routed over the route with the highest ability to accept new connections (w*idest path* routing), but over a route that crosses these network areas with the highest ability to accept new connections.
The performance of the developed algorithm has been tested by means of simulation. The simulator used in the Reform project is especially developed for studying routing methods in ATM networks,

[1] The choice is limited to the set of admissible routes, as calculated during network planning. Source node routing is applied, implying that the ingress node selects the full path, and feeds this information to the other nodes during call establishment through a path descriptor in the SETUP message.

and allows for comparison to shortest path routing, shortest path routing with alternatives and least loaded path routing.

One of the parameters to be tuned within the load balancing functionality is the change in available capacity per VP that is considered significant and triggers an update of topology information throughout the network. The decrease in the update frequency has been shown to be most significant when we increase the significant change level from 0 to 20%, therefore, 20% seems to be an optimal level for this parameter. A simulated 4-node network and 8-node network shows that the HPA algorithm behaves well for low and medium traffic load, but not significantly better than the traditional methods. When the network gets heavy loaded the throughput under the LB-algorithm decreases. Tests in bigger networks, and active interaction with the Bandwidth Distribution algorithm are however expected to show the full strength of the HPA algorithm, studies on this however are still ongoing.

5. Bandwidth Distribution

It is the task of the dynamic routing functionality to balance the VC connection load over the network. On the other hand, the bandwidth distribution function (BD) dynamically adjusts the bandwidth allocated to working VPCs (and their protection resources) according to the actual load observed in the network. BD distributes unused bandwidth on a physical link to the higher utilised VPCs at the expense of the lower utilised ones. To achieve this, the BD algorithm is invoked on certain threshold crossings corresponding to high or low VPC utilisation. It is important to state that BD is not entirely free to allocate as much or as little bandwidth at it thinks appropriate according to its short term view of network load. BD always operates within constraints imposed by the network planning layer. Specifically these are upper and lower bounds which are defined by network planning according to its view of longer term traffic predictions and the quality levels, in terms of call-blocking, associated with each traffic type.

Simulation results have shown significant improvements on network utilisation when applying a dynamic management scheme to VPC bandwidth allocation via the REFORM BD algorithm. The graph below plots the traffic transported (offered traffic minus blocked calls) against the offered load in a number of cases. The results refer to a single link interconnecting two ATM nodes.

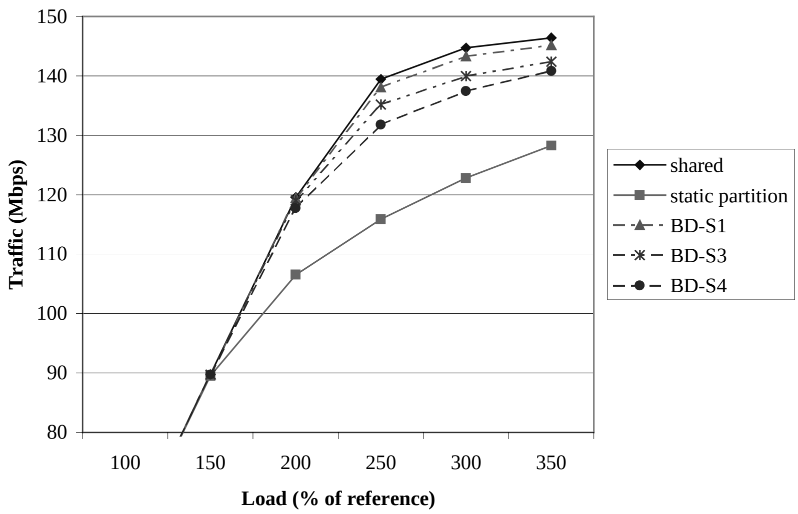

The first case, "shared", is where all CoSs (telephone, video-on-demand, and video-conference in our simulation studies) are transported on a single VPC which occupies the entire link capacity. In this case we see the maximum load transported and it may seem, at first glance, to be the most efficient scheme. However, there is no differentiation between the CoSs - calls are accepted or rejected on a first-come first-served basis and it is very much a best-effort service in terms of call-blocking or network availability.

The second case, "static partition", is where separate VPCs are defined per CoS but bandwidth is allocated statically. In this scenario bandwidth can be reserved per CoS and the network may now offer different levels of call-blocking to different classes. In this static partitioning case we can achieve

the required QoS differentiation but this is at the expense of network utilisation and a higher proportion of blocked calls - as can be seen on the graph.
By employing the BD algorithm and a dynamic bandwidth allocation scheme we can achieve a utilisation close to the "shared" case while retaining the QoS levels required per CoS. The graph shows 3 cases labelled "BD-S1" to "BD-S3". In each case different values are used for the sensitivity of BD's response to traffic fluctuations. "BD-S1" requires that BD is invoked more frequently, thereby imposing a greater management overhead, while "BD-S4" is less efficient in terms of network utilisation but it requires less management traffic and processing capacity.
The simulation results are being validated through comparisons with the results obtained from experiments running in real ATM switches with the full REFORM prototype system. Further work includes investigation of the dynamic relationship between BD and load balancing and their combined effect on network operation.

6. VP Protection Scheme

In order to recover from single link failures in the network, the system adopts a static shared VPC protection scheme. Static in the sense that the protection VPCs that are associated with working VPCs are pre-planned and installed within the network by the network planning layer. Shared since the restoration resources that are available on a physical link are not dedicated to a protection VPC, but can potentially be captured by different protection VPCs recovering from non-related physical link failures.
The project has during its experimentation improved on the VP protection scheme, thereby significantly reducing the overall protection switching time.
In the first implementation of the restoration mechanism, the restoration resources are captured sequentially in each node along the protection path. The overall restoration time comprises the restoration message propagation delay and the bandwidth allocation time in every node along the protection path. Thus, the restoration time is proportional to the restoration path length. The project has revised the VP restoration scheme during its experimentation. The improved VP restoration scheme eliminates the impact of the restoration path length to the overall restoration time by capturing the restoration resources in parallel on every node, thereby significantly reducing the overall protection switching time.
The restoration protocol is designed to overcome unexpected situations such as loss or corruption of restoration protocol data units in a fast and efficient manner. Moreover, n the case of restoration failure a fast crank-back mechanism is activated in order to release the captured resources and make them available to other possible concurrent restoration sessions.
The above implementations of the protection scheme were extensively tested and demonstrated in diverse testbeds and in an emulated network environment. The outcome of the experiments is that the aforementioned protection schemes can achieve restoration times at the order of hundreds of milliseconds, while the revised implementation performs about 40% better than the first one.

7. VPC Layer Reconfiguration

During network operation, the network will be confronted with variations in traffic, and possibly variations in physical topology. The variations in traffic load may be explicitly provided to the network reconfiguration function by the operator, based on business forecasts, but may also be triggered from within the network, more specifically by the alarms described in the previous sections. Variations in topology occur when introducing or removing (single link failure) physical links. The reconfiguration process resulting from the latter is more generally known as network normalization. Note here that introduction/removal of switching equipment has been outside the scope of the system specification.

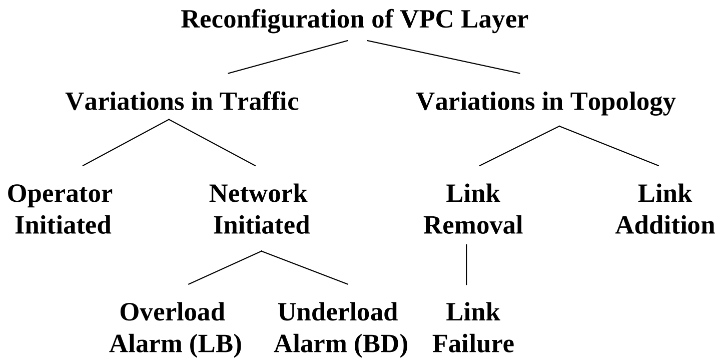

Figure 2 VPC Layer Reconfiguration overview

For each of these individual events shown in Figure 2, algorithms have been developed within the project, so as to smoothly migrate from the present VPC layer towards an improved VPC layer, without the interruption of a single VC connection.
Essentially, any variation in traffic is assumed to be reflected in an updated traffic estimation matrix. By mathematically subtracting the original traffic estimation matrix from the updated traffic estimation matrix, a new matrix is created that highlights the variation in traffic for every source/ destination/CoS triplet. First, for all source/destination pairs, where traffic is decreased, the bandwidth associated with this decrease is proportionally released over the VPCs that constitute the routes between the selected source and destinations. Next, the network topology consisting of all links where some unallocated (and regained) bandwidth exists, together with the sparse traffic estimation increase matrix is fed to the design algorithms discussed in chapter 3 (route design). This way, a 'delta' VP network is created that can be easily superposed on the existing one, that is to say, preferentially, existing VPCs will see their bandwidth allocated increase, while sometimes new VPCs and routes may be introduced. The main goal however remains balancing estimated traffic over the physical topology. The reader is referred to [D14] for a full description of the further mechanisms.

8. Conclusions

In this paper, we have presented an overview of the architecture, functional model, and results obtained in the course of the ACTS-REFORM project.
This paper has presented the specifications and experimental results of the REFORM system. It has been described how advanced network planning and adaptive network operation are integrated with a high speed protection switching scheme in order to optimize in a cost effective way the provisioning of reliable VC switched services, given an operator investment in physical network topology.

Acknowledgements

The work described in this paper would not have been possible without the active participation of all partners of the REFORM project (see [1]). This project has partly been funded by the European Commission ACTS programme.

References

[1] http://www.algo.com.gr/acts/reform
[D1] "REFORM system requirements and analysis", ACTS-REFORM, January 1997
[D4] "Final REFORM system specifications and architecture", July 1998
[D14] "The self healing mechanisms, load balancing and dynamic routing capabilities and algorithms of the REFORM system", ACTS-REFORM, September 1999