Unified Fault, Resource Management and Control in ATM-based IBCN

S.Sartzetakis, P.Georgatsos, G.Konstantoulakis, G.Pavlou, D.Griffin

ICS-FORTH/GR, AlgoSystems/GR, NTUA/GR, UCL/UK

Abstract: In this paper we present the initial specification of a revolutionary system that covers both the control and management planes of network operation, with emphasis on fault, performance, and configuration management. The system incorporates rapid and reliable ATM network layer self-healing mechanisms, intelligent load balancing, dynamic routing and resource management functions all interworking together with the overall goal to ensure cost-effective network performance and availability under normal and fault situations. Restoration mechanisms in the ATM network layer are integrated with the control and management plane functionality, aiming at providing an integral and network-wide treatment to the problem of fault recovery. The system is designed and is being implemented adopting the emerging TINA framework and using a CORBA-based distributed processing environment. The validity and effectiveness of the system are assessed and demonstrated through international ATM field trials.

Keywords: Integrated ATM control and management. Fault management and network survivability. Dynamic routing, self-healing, OAM. TINA-C, CORBA DPE.

1. Introduction

ATM networks will be key components of the future Global Information Infrastructure. Compared with today's segmented telecommunications and computer networks, the Information Infrastructure will become an integrated distributed processing environment that will support delivery, processing and integration of any medium (video, image, text and audio) or mix of them (multimedia). The information market allows for outsourcing of specialised skills by providing the virtual space for people to work in, even without leaving their home. This evolutionary process may have very soon a fundamental effect on the way people are working, especially in the information sector—already today, the Internet illustrates some of these sociological changes that are beginning to take place.

Gradually the effects will spread into other sectors; new technology will allow the efficient decentralisation of support services, such as medicine and education, towards local communities. There will be an opportunity to make the support services customised to the needs of individuals with special needs, e.g. continuous medical monitoring and tele-diagnosis. Because of the integral significance of these services to society, the basis of the Information Infrastructure must be very robust: the reliability and availability requirements will be very high. In addition, the existence of advanced services with guaranteed quality and the open market competition for the delivery of those services dictate that, the Information Infrastructure must have the means to provide cost-effective performance at all times.

In order to minimise the cost of building and maintaining wide area ATM networks, efficient and cost-effective resource management techniques, are needed as that cost increases dramatically due to over-allocation of resources for reliability and availability. This requires building a network that has the intelligence to monitor, control, and enforce its own QoS agreements, guaranteeing the high availability expected. There is an urgent need for efficient resource routing and failure management mechanisms.

Research in resource and routing management has led to a number of optimising algorithms. However, these studies tackle the problem of network availability and performance in isolation, not considering the interactions between the control and management functions. They do not take explicitly into account the diversity on performance and bandwidth requirements of the many service classes supported by the networks. Furthermore, there is today no integrated perception of resource control, routing and alarm management in the perspective of network reliability and availability.

Recognising the need for combining the functional capabilities of the control and management planes to the benefit of network operation, this paper introduces the REFORM (REsource Failure and restORation Management) [1] system with the purpose to ensure network performance and availability within acceptable levels under normal and fault conditions. The specifications of this system both from the functional and architectural aspects is the main theme of this paper.

Our system integrates control and management plane functionality for providing a complete, network-wide solution to the problem of network availability under normal and fault conditions. Specifically, control plane functions such as route selection, Operation Administration and Maintenance (OAM), and self-healing mechanisms are integrated with higher level network-wide routing and resource management functions, with the purpose to ensure network-wide performance and availability.

We treat the problem of network availability in ATM-based IBC networks in an integral manner. The term "integral" refers to the ability to:

- cover the complete failure management cycle, that is prior, post failure and failure normalisation phases;
- involve both control and management plane functionality.

While research has already been conducted in the area, it is the first time that a complete system is introduced that provides:

- **an integrated approach to the problem of network performance and availability**, considering the whole failure management cycle and integrating the OAM and control mechanisms with the network management functions in the performance and fault management areas;
- **intelligent load balancing, OAM restoration, dynamic routing and spare resource management** mechanisms all **interworking together**, taking explicitly into account the multi-service nature of the network environment.

The REFORM system has been designed along the TINA architectural lines, using a CORBA-based distributed processing environment with all the advantages this brings and which are explained later in this paper. While the distinction between the control and management planes is maintained at a conceptual level, this distinction is relaxed during the system design and implementation phases. Part of control and management plane functionality will be modelled and implemented based on a DPE platform, using a uniform set of DPE services. The REFORM system is one of the first systems to provide such integrated solutions.

In the following section we provide briefly some background information we consider helpful in understanding the scope of our system. Section 3 is presenting the REFORM system specifications, and 4 the architecture. Following that we describe a realistic environment scenario that is used to assess the REFORM system as whole. Finally Section 6 concludes the paper with a summary and some future work.

2. Background information, state-of-the-art.

Considering networks as multiple-resource, multiple access systems, an obvious problem to arise is how access to the resources is granted. **Resource management** refers to the necessary means and functions for resolving such contention problems with the overall objective to obtain solutions satisfying certain cost-effective criteria while constrained by general performance requirements. ATM networks by their very asynchronous multiplexing nature and the flexibility of the VPC layer that offer, provide several degrees of freedom in traffic multiplexing. This flexibility although eases the task of traffic admission and multiplexing with potential benefits, it might be proved disastrous to network performance if not managed properly. The resource management problem in general is related to the routing management problem, complementing each other. Routes in ATM networks are defined in terms of VPCs and VPCs have been defined in order to support routing.

Routing functionality is spread over the control and management planes of network operation and across the network elements. Route selection functionality is part of the call control functionality running at the network switches, whereas route definition functionality is regarded as management plane functionality. Different sets of route may be defined for different connection types according to their performance requirements.

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. Routing policies and algorithms should therefore cater for traffic changes. Moreover, they should adapt gracefully to topological changes avoiding routing through network damaged areas. Moreover routing functions should harmonically co-exist with the resource management functions in all layers of network operation with the overall goal to optimise network performance. Different routing policies may be defined, differing in the placement of the associated functionality, the type of algorithms they use, the degree of adaptivity they offer and the type of information they utilise.

Successful **fault management**, being one of the major services, that the ATM network management must provide even end-to-end to the customers, requires that the information is gathered from the network elements in a standardised way. While taking in account the work of standardisation actors in this field, and recognising that the standard layout provides an efficient way to reliably inform the user of faults occurring in the transmission network, the belief is strong that additional functionality is needed in order to insure a higher level of network reliability and resiliency.

Practical experience on existing networks shows that some fault scenarios involve alarm burst phenomena (e.g. uncontrolled pouring-in of numerous redundant alarm notifications) which tend to overload the Operation System, in a situation where availability and synthetic information are both required in order to take management decisions as efficiently and accurately as possible.

Network robustness can be achieved with the deployment of autonomously operating **restoration mechanisms**. These restoration mechanisms will reside in different network levels and different sub-networks and will inter-work with each other through escalation schemes. The autonomously operating restoration mechanisms, provided by the network will interact with the network management plane, in order to achieve the goal of network robustness. The network management system can simply monitor the restoration process or it can also actively participate in this process.

Fast control plane restoration (self-healing) mechanisms as well as medium or longer term restoration mechanisms, in the management plane must be accommodated. The management will use the hooks provided by OAM in order to complement and optimise the effect and operation of the self-healing actions taken at the control plane.

There are two possible self-healing techniques that can be used for restoration management. The first is the **pre-assignment of VPs**, that are characterised as back-up VPs and the second is a flooding based dynamic route search scheme. The pre-assignment of the back up VPs yields several benefits compared to the flooding based dynamic route search scheme used in most existing self healing schemes. Its primary advantages are restoration rapidity and realisation of path restoration between path termination nodes. In this scheme there is the need of the management of back up VPs as well as of original VPs. In particular this restoration mechanism includes a message transmission system which is based on OAM flows, a spare resource management algorithm which is based on iterative algorithms in order to achieve optimum utilisation of network resources and also a well defined failure management cycle which can be realised as components of a management system.

The ATM-based networks operation and maintenance is organised in a layered fashion. Five hierarchical **OAM** layers and the relative OAM information flows have been defined (F5-1: Virtual Channel, Virtual Path, Transmission Path, Digital Section, Regenerator Section). The two first OAM layers (F4-F5) are in the ATM layer and the rest three (F1-F3) in the physical layer. In the various network locations it is not necessary for all the OAM layers to be present.

The OAM functions at the ATM layer are independent of the underlying transmission system (ATM/SDH/PDH). F4 and F5 OAM cells must be multiplexed into the user cell stream. In the de-multiplexing direction, the cells must be routed to the OAM processor unit. The OAM flows F4 and F5 offer a unique possibility to evaluate network's performance on the actual connections in use. The in-service performance measurements offer realistic information on the connections in use by an application, while the OAM flows load on a connection can be as low as 0.1% of the user load. This load does not seem to affect the quality of service expected by the user.

The importance of the *integration* between *management system and OAM* has been identified by ITU. In its recommendation I.610 proposes an interaction scheme between the TMN and OAM. This scheme involves both X & Q interfaces for the entire path starting from the customer premises, passing the access local exchange up to the core network.

3. REFORM system scope, approach, and functional specifications

Adopting a network operator's viewpoint and by taking into account the broadband and multi-service nature of IBC network environment, the scope of the REFORM system is to provide the necessary means and functions for ensuring network performance and availability within acceptable levels under normal and fault situations. *Ensuring network performance* means the network will be able to maintain the performance of the existing connections within acceptable levels. *Ensuring network availability* means the network will be able to set-up connections guaranteeing their performance requirements. To achieve this, the REFORM system offers:

- fast fault detection, alarm indication and report functions;
- rapid and reliable network self-healing mechanisms for resource and service restoration, spread across both the control and management planes of the network operation;
- efficient dynamic and/or static resource and routing management schemes and algorithms with inherent load balancing functionality, coping with fault conditions; and
- efficient resource migration algorithms.

All these will interwork together with the overall goal to ensure the:

- cost-effective network performance and availability in normal conditions; and
- cost-effective, reliable and robust network recovery in the performance of existing connections as well as of the availability for new connections, from fault situations.

The term **self-healing**, as used above, refers to the ability of the network to reconfigure itself around failures quickly and gracefully with the goal to restore service availability within acceptable levels for existing as well as for future connections. Self-healing implies resource restoration by means of a distributed mechanism, as opposed to centralised schemes. Self-healing mechanisms mainly reside in the control plane (one node), while higher level decisions can be escalated to the management plane. The REFORM system provides self-healing mechanisms at both the control and management planes. The latter is provided as part of suitable distributed routing and resource management schemes, dynamically adapting in traffic cases and in cases of resource failures. Centralised protection is also provided through management activities for defining suitable sets of routes and for allocating spare resources, in a cost-effective manner. These activities are adaptive in a quasi-static form regarding network resource failures, redesigning new sets of routes or reassigning spare resources. When failures occur and the lower-level self-healing capabilities of the network cannot succeed in restoring network availability under the current traffic conditions these activities may also be triggered. The REFORM system integrates all these functions through a carefully designed hierarchical architecture, where a CORBA DPE platform provides an abstract interface to discriminate high level management functions from low level control ones. Aspects of this architecture are discussed in Section 4.

To facilitate hierarchy, and thus functional decomposition, we adopted the distinction between the management and control planes in network operation. After all, such a distinction is necessary since it maps to the different domains and actors' interfaces currently being formed in the telecommunications arena; those of network equipment manufacturers, network operators and management system providers.

In general, the control plane is required for the operation of the network, including the base functionality of signalling, Call Control (CC), and Operation Administration and Maintenance (OAM). The management plane is required for the optimal operation of the control plane, by tuning and appropriately managing various operational parameters in the control plane.

With respect to fault management, the approach adopted is that the control plane encompasses the functionality of fault detection and fast-responding self-healing mechanisms through the OAM capabilities that it offers. The management plane encompasses resource and routing management functionality for tuning and optimising the operation and effect of the control plane actions as well as for guaranteeing cost-effective allocation and usage of the spare resources in the network. Re-routing is performed to avoid damaged network areas while at the same time relieving the load from the back-up network resources. If such activities fail to restore network performance and availability given the current traffic conditions, other -centralised- management activities will be activated (e.g. for defining new sets of routes, service migration). While the control plane functionality provides a fast reaction to a fault situation, the management plane provides the required short, medium or longer term reactions, depending on the severity of the failure and the actual traffic conditions, complementing therefore the effect of the self-healing actions taken at the control plane.

We broaden therefore the scope of restoration management beyond the level of self-healing mechanisms for service restoration at a local level. By incorporating effective resource and routing management functionality -suitably coping with fault conditions- our system restores service availability for existing and new connections, at a network-wide level.

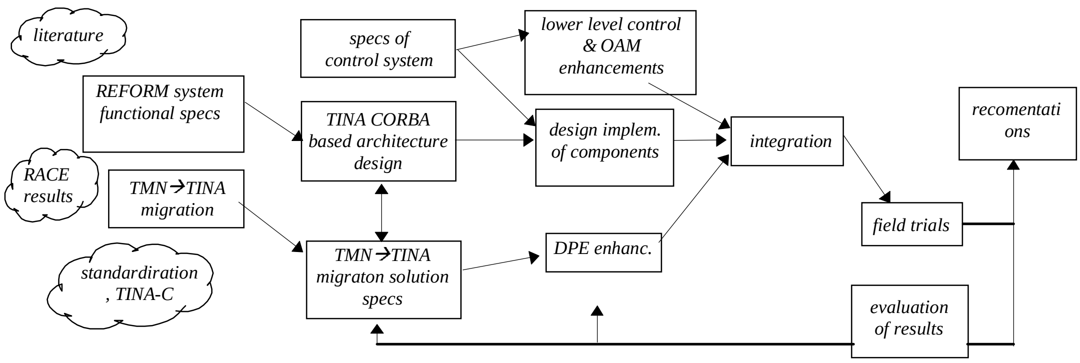

Figure 1.: The REFORM Approach.

Figure 1. shows the REFORM approach. We build upon the results of the RACE II projects *ICM, IMMUNE* and *TRIBUNE* [XX],[XX],[14], and on other proven results in the literature [9], [11]. Specifically, REFORM makes use of the *VPCRM (VPC and Routing Management) architecture* , the specified load balancing[10], VPC bandwidth management and route design algorithms and principles[11] developed in the ICM project. The validity and effectiveness of these results have been verified through extensive experimentation in real and simulated network environments [10], [9]. The

REFORM system also utilises the results of the IMMUNE project regarding self-healing networks and restoration management. The results of the TRIBUNE project regarding control plane and OAM functionality set the basis for the required lower-level control functionality. Finally, state-of-the-art studies in the literature like [7],[8], set the ground of the self-healing , spare capacity allocation and reconstruction algorithms employed in the REFORM system.

3.1. System Functional Specifications

We are concerned with failures caused in or escalated to the ATM network layer. For example, when a physical link failure has occurred and cannot be recovered by the restoration mechanisms in the physical layer(s), the fault is escalated to the ATM layer. This type of failure will cause a large number of alarms, as many as the VPLs defined on this link. REFORM therefore will provide for alarm correlation functionality in order to determine the original link failure. Faults in the ATM layer are also concerned with equipment related failures i.e. VP cross-connect, VP TP (Termination Point) failures. The failures will be identified through the standardised OAM alarms. Restoration mechanisms and escalation procedures in the network levels below the ATM layer are not addressed.

The REFORM system incorporates rapid and reliable ATM layer self-healing mechanisms, intelligent load balancing, dynamic routing and resource management functions all interworking together with the overall goal to ensure cost-effective network performance and availability under normal and fault situations. Cost-effective solutions are pursued, on one hand by means of considering suitable objective functions in the optimisation problems corresponding to the route definition, spare resource allocation and management algorithms, and on the other hand by considering different survivability of service classes.

In the rest of the section we outline the scope of the functionality of our system per each of the network phases with respect to failures. We concentrate on describing how control and management functionality is integrated to fulfil the objectives of our system.

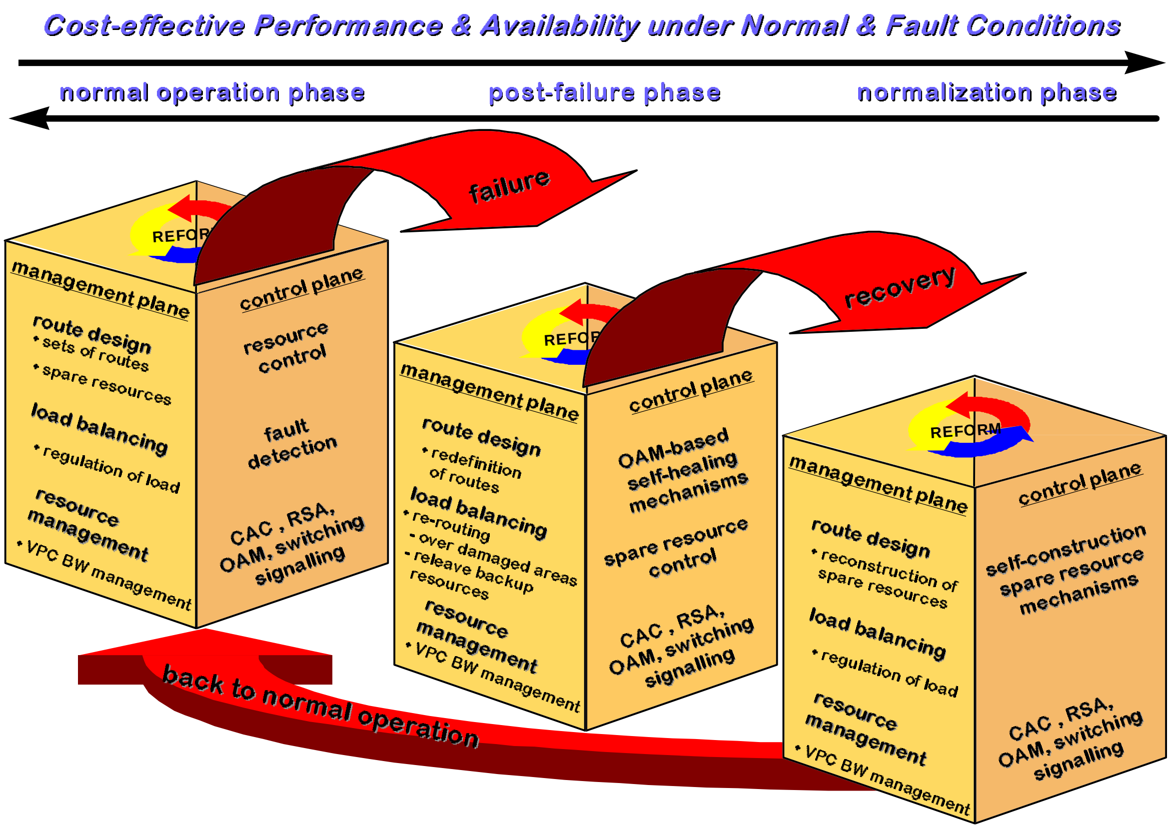

Figure 2.: The REFORM system during the network operation phases.

During **normal operation** of the network, the system incorporates intelligent routing, load balancing and resource management functions, which -taking into account network-wide traffic and actual usage conditions- with appropriate management actions tune the operation of the control plane. All this functionality is provided through a suitable hierarchy corresponding to different levels of abstraction and time-scales. The overall aim during this phase is to guarantee effective utilisation of network resources and cost-effective network availability within acceptable levels, while ensuring that the performance requirements of the different service classes are met. The *route design functions* are responsible for the design of appropriate admissible sets of routes per connection type, based on suitably defined sets of VPCs, taking into account the different performance requirements of the different connection types that the network supports. The route design functions are also responsible for the allocation of appropriate spare resources in the network according to certain cost-effective objectives and an overall protection strategy to provide as many options as possible for bypassing faults. They should provide for adaptability to cater for network traffic changes. The *load balancing functions* (which may distributed or centralised) operate on the defined sets of routes trying to influence the routing decisions taken in the network switches, by conveying to them network-wide information. Through their activities load distribution may be regulated over the whole network. Note that balanced networks have the potential benefit of minimising disruptions of existing connections in case of network eventualities (node, link failures). The control plane encompasses the base functionality of CC, including routing and CAC, signalling and OAM. The routing functionality in the control plane is responsible for taking the actual routing decisions by means of suitable route selection algorithms. It operates on the routing data (sets of admissible routes, route selection parameters) designed by the overlaid management functions. The control plane routing algorithms may also perform load balancing at a local level (i.e. at the vicinity of a node). The resource management functions are responsible for the management of the bandwidth of the VPCs according to actual network conditions. The system employs such functions both in the control and management plane, differing in terms of time-scale and abstractions. They try to ensure effective usage of VPCs and to avoid congestion at the cell and the connection level.

During the **post-failure phase** the system employs fast-responding self-healing functions in the control plane through the OAM facilities that it offers. The system follows a self-healing scheme based on pre-assigned back-up resources, rather than based on flooding messages. The standardised OAM flows are used to provide a fast and reliable means for inter-node message exchange. In addition , the adaptive, distributed and/or centralised routing, load balancing and resource management functionality of the system, with their inherent self-healing capabilities, will be activated for tuning the operation and complementing the effect of the control plane actions, to a network level. The overall goal during this phase, is to ensure a rapid recovery of network performance and availability, taken into account the actual traffic conditions. The entire system functionality is properly integrated through a suitable restoration escalation procedure through the levels of the REFORM hierarchy. The system combines the merits of distributed and centralised control schemes during the restoration process. Once a failure has been detected, the control plane will react through its OAM-based self-healing procedures and the associated automatic re-routing mechanisms. Spare resource control is also required for resolving competition as spare resources are normally shared. The load balancing and resource management activities in the management plane will also be in effect, through their own self-healing capabilities for providing a first level network-wide treatment to the restoration procedure. The load balancing functions will influence routing decisions so that future connections avoid routes traversing damaged paths, while at the same time taking care not to overload the back-up routes. The resource management activities will try to ensure that sufficient bandwidth on the back-up resources can be allocated. All these activities will take a network-wide view and the actual network traffic conditions. The system monitors network performance and availability with a wide perspective during the restoration process, in order to assess the effectiveness of the actions taken so far. If network performance is found unacceptable, the centralised routing functions will be activated for assigning new sets of routes and VPCs as necessary. The system allows for different *survivability classes* for prioritising access to network resources. If required, the restoration process may also be extended to higher levels for performing appropriate service migration functions, an extension of the system currently under study.

The **normalisation phase** is concerned with the reconstruction of routes and spare resources to cater for the predicted traffic and to build in flexibility to allow a rapid response to subsequent failures. The overall goal is to enter the normal operation mode by causing the minimum of disruptions to existing connections. The route design functions will determine new sets of routes and spare resources according to a general strategy satisfying certain cost-effective criteria and taking into account network usage predictions. Different policies for spare resource reconstruction, like for example distributed self-reconstruction mechanisms are under investigation. The load balancing and resource management functions operate as in the normal operation case taking into account the newly resulted configuration. The control plane OAM and routing functionality will also be activated according to the reconstruction directives given by the management plane.

The detailed description of a functional model decomposing the above functionality into individual components and specifying the information exchanges between them is beyond the scope of this paper. This can be found in [12].

4. System Architecture

The key aspect of the REFORM architectural approach is that we follow the emerging TINA framework with the purpose of evaluating its suitability and assessing its potential benefits when applied to network management, while comparing it as well to established technologies (OSI-based TMN, Internet SNMP). We use a commercial CORBA-based Distributed Processing Environment (DPE) for designing and implementing both the control and management plane functionality. Existing TINA-C architectures in the area of fault and resource configuration management are taken into account. Also existing TMN architectures for configuration and performance management, such as the ICM Virtual Path Connection and Routing Management, constitute the starting point for the management plane; these are enhanced and modified in the light of the TINA-C developments and ported on the CORBA-based DPE. They are complemented by new functionality, not yet addressed by TINA and not available in existing TMN architectures, which is under investigation in REFORM.

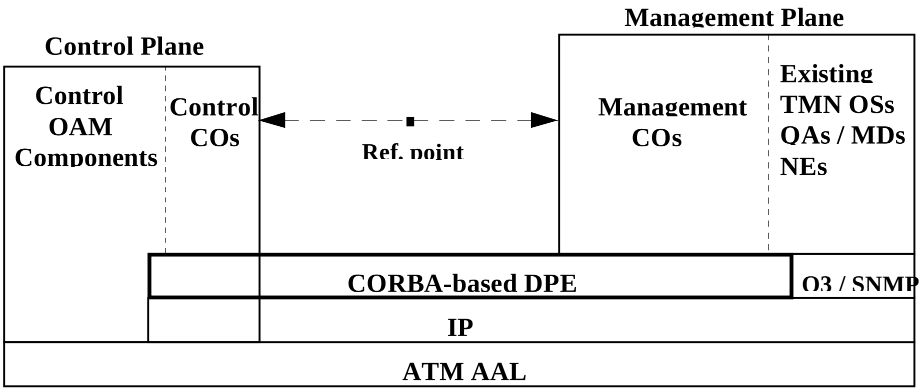

Figure 3.: Overall REFORM system architecture.

The overall system architecture is shown on Figure 3. There exist two major domains, related to management and control plane respectively. The major underlying assumption is the use of TINA-C engineering principles in the form of a CORBA-based DPE.

The detailed architecture of the management plane takes into account state of the art work in this area i.e. in TINA-C and the RACE-II ICM project. The management plane functionality is specified as Computational Objects (COs) for fault, performance and resource configuration management. Existing TMN applications from previous work - Operations Systems (OSs), Q-Adaptors (QAs) - are integrated in the REFORM system. Also, Network Elements (NEs) are accessed through TMN Q3 or SNMP interfaces. The integration of those elements and applications in the TINA-based REFORM environment raises issues of migration from TMN to TINA, which are discussed below. The architecture of the management domain has both the hierarchical aspects, as dictated by the TMN layered architecture which is adopted by TINA, and also the flat peer-to-peer aspects which are facilitated by the use of a DPE-based computational decomposition. Federation of activity between management domains is considered as well.

The control plane functionality will be supported by relevant control logic and OAM transport protocols operating directly over the ATM adaptation layer. The state of the art in terms of projects such as IMMUNE and TRIBUNE together with the relevant ongoing ITU-T work is the starting point. The integration between the control and management plane domains will be accomplished through the interaction of computational objects over the DPE. As such, the control plane domain includes computational objects for the purpose of interfacing to the management plane. Such interactions are expected to be bi-directional.

All functions and operational components of our system in the leftmost box in Figure 3. are regarded as local inside a node, and will be implemented according to the particular capabilities of the network elements. We take into account the low level i/fs that provide connection with the h/w subsystem. Thus, direct message exchange will be feasible. Further OAM functionality and protocol is under study for the layer above ATM layer. In particular, point-to-multi-point OAM functionality is studied for the AAL layer. A multiplexing scheme for the upstream will be adopted in order to support the point-to-multi-point video distribution application that will be used in the trials. A method for multiplexing of OAM flows in the upstream direction will be defined and an approach similar to the multiplexing of flows in AAL 3/4 will be followed (multiplexing in the same VP/VC using ids).

Overall, there are three types of transport in the REFORM system:

- the ATM AAL, which supports the control plane interactions; the direct use of the ATM AAL ensures high-speed interactions but the relevant presentation and distribution facilities are minimal;
- the TMN Q3 using RFC1006/TCP/IP over the ATM AAL; this supports management plane interactions for the existing TMN applications; and
- the CORBA-based DPE using TCP/IP over the ATM ALL; this supports the majority of the management plane interactions and the interactions between the control and management domains.

There is also a possibility to use the CORBA-based DPE directly over the ATM AAL. This depends on developments on CORBA mappings in the lifetime of REFORM.

In order to be able to test our system under stress and worst case conditions all OAM protocols are integrated within the testing systems we use for our experimentation. The test tool is attached to the switches and monitors the data streams, OAM flows and collects statistical information. On the other hand we are able to generate fault in the OAM flows and collect statistical information. We are also able to generate faults in the OAM streams in order to cause predefined error profiles.

Migration from TMN to TINA

The TINA-C project has defined a generic telecommunication systems architecture, as well as specification methods, the scope of which include both services and the management of services, network and computing infrastructure in an integrated fashion. REFORM tackles aspects such as how to integrate or migrate towards the TINA environment, using in particular the DPE engineering model and the problem decomposition from the enterprise, information and computational viewpoints.

More specifically, REFORM will use concepts and architectural elements from TINA related to its Management Architecture and in particular to Fault and Resource Configuration Management. The network resources will be represented as described in the TINA technology-independent Network Resource Information Model (NRIM) **[3].** Existing TMN architectures for ATM resource management are mapped onto TINA concepts and will be expanded during the evolution of our system. The resulting architecture will be fed back to TINA.

Given the fact that network elements will support TMN Q3 or SNMP management interfaces while existing TMN management applications (OSs) with Q interfaces may be reused, issues in migrating these components to the TINA environment will have to be tackled. In addition, existing TMN architectures for ATM resource management (e.g. ICM) will be reworked according to the new methodologies and the architectural input from TINA-C in these areas.

In the key area of integrating elements or applications with Q3 interfaces, the X/Open Joint Inter-Domain Management (XoJIDM) task force has been investigating the similarities and differences between GDMO/ASN.1 and CORBA IDL object models and proposes generic translations in both directions. It should be mentioned that elements with SNMP interfaces may also be seen as Q3-capable elements through generic translation: the ICM project has contributed to the relevant NMF specifications and has developed a generic Internet Q-Adaptor (IQA) which is available to REFORM.

Given the fact that fault and performance management need fine-grain event reporting / logging and sophisticated query and retrieval facilities, REFORM will assess the suitability of OMG Common Object Services (COS) to support such functionality. Specialised facilities in addition to the OMG common services will also be investigated, such as management brokers that enable multiple object access and sophisticated information retrieval, sophisticated event reporting and logging facilities, metric monitoring and summarisation facilities etc. In general, the full expressive power of the OSI Management / TMN framework needs to be available over the CORBA-based DPE.

5. An Experimentation Scenario

Our work is centred around a number of trials and experiments to validate the system in a realistic environment. The goal of this practical work is to assess the REFORM system as a whole, taking into account the contributions made by the self-healing mechanisms, load balancing, routing and resource management algorithms. The work will attempt to validate the developed functionality and the specific algorithms designed for each aspect of the system to show that the proposed mechanisms are effective in resolving fault conditions and restoring network performance with an adequate responsiveness so that user perceived quality of service is disrupted as little as possible.

The field trials will be based on European National Host (NH) platforms [2]. The NHs are consolidated advanced platforms comprising communications infrastructure, services and generic applications. The NHs are made available to projects interested in performing operational trials or experiments of leading edge applications, services and management that involve real networks, services and users.

As an example of a typical, but simplified, experimental scenario, the following describes a fault in the delivery of video services and the associated activity of the REFORM system to resolve the problem.

Consider a video distribution service consisting of a video server which distributes a video stream to the users of the service on demand. (Figure 4.). Faults can affect this service in a number of ways. For example: a user may be unable to access the video distribution server to request a new video stream; an existing video stream may be interrupted; or the video server itself may collapse. In the following a fault caused by a link failure is discussed.

On failure there are a number of ways of restoring the service. These can be at many levels of the REFORM architecture: from the reception of OAM alarms and the invocation of fast, localised self-healing mechanisms through to network wide, off-line redesign of the network routes available to the video services. The goal of our system is to enable these mechanisms to work together in an integrated way. By the co-operation of algorithms and functions in each of the architectural levels the network can be made operational as quickly as possible through the immediate response of localised algorithms embedded in the network elements while the higher level management algorithms restore efficiency in the network as a whole over a longer time period.

Figure 4 shows a simplified view of a sample network configuration formed by the EXPERT testbed (Swiss NH), the TRIBUNE (Dutch NH) and the Norwegian NH. A video distribution server is attached to the LATEX node of the Swiss NH, there is a single user group attached to the NT2 node of the same network and a second user group located in the Dutch NH.

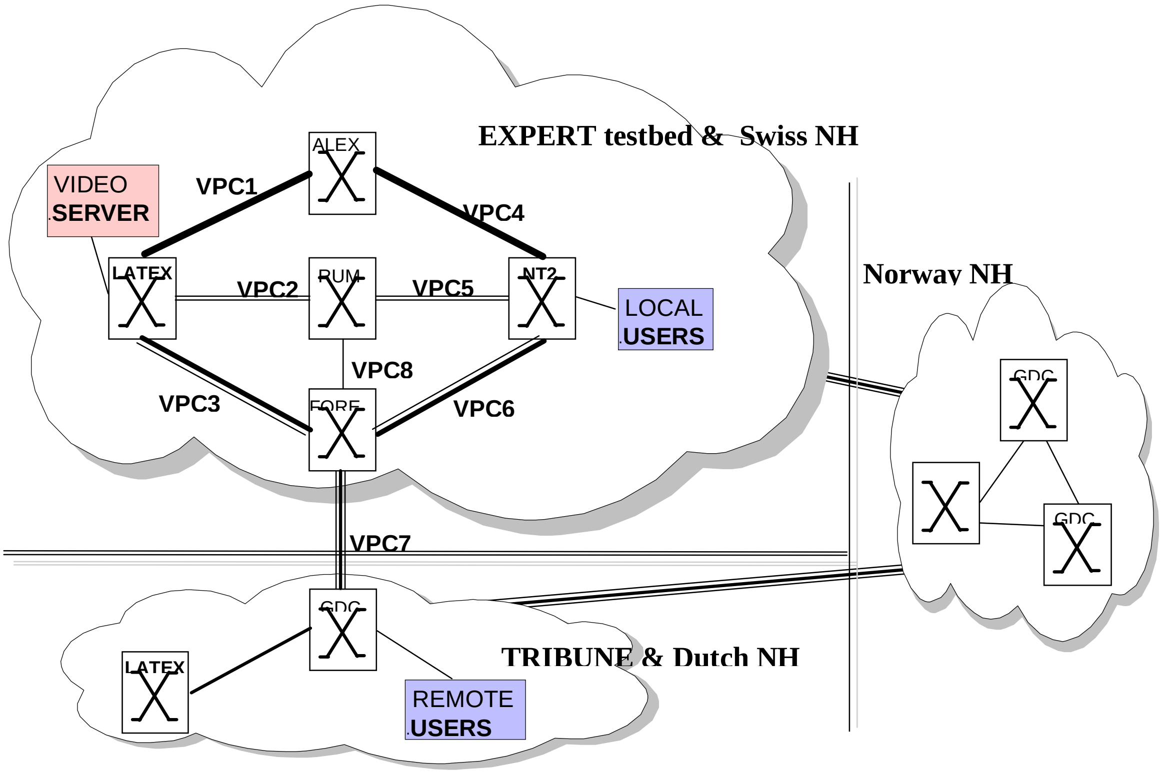

Figure 4.: The NH set-up for the REFORM experiments.

Video streams are transmitted via VCCs which are set up on demand between the users and the server and routed over previously installed VPCs. In this example, two possible routes have been assigned to video traffic between the video server and the local users on the Swiss NH: over VPC1 and VPC4; or over VPC2 and VPC5. Similarly, two routes have been assigned to the remote users at the Dutch NH: over VPC3 and VPC 7; and over VPC2, VPC8 and VPC7.

The fact that two routes have been assigned for each source-destination pair is implemented in the network by alternative route selection entries in the VC switches of the network nodes. When a connection request is made, Route Selection Algorithms (RSAs) determine which route should be used [10]. The route selection entries are defined initially by the route design algorithms in the upper levels of the REFORM architecture (note that this is an example of the co-operation between the management and control levels of the REFORM architecture).

During normal operation of the network, the RSAs select appropriate routes for new connections according to their algorithms [10]. The load balancing functions manage the route selection criteria used by the RSAs to balance the load over the network as evenly as possible by altering the selection priorities associated with each route selection entry.

In this scenario we now introduce a fault in the link between the LATEX and ALEX nodes in the Swiss NH, so putting VPC1 out of service and disrupting the video streams which were allocated to VCCs routed over this VPC.

On failure of this link, the LATEX node is notified by a control plane signal (OAM alarm) that the link is out of service. On receipt of this signal, OAM-based restoration protocols will attempt to restore the failed link. Because this may take a large amount of time in the case of physical failures requiring human intervention, the local control plane self-healing algorithms will prohibit the RSA from selecting this VPC for future connections. As a secondary response the control plane algorithms will try to obtain as much bandwidth as possible for the alternative route (VPC2 and VPC5) from the common pool bandwidth which is continually updated by the higher level bandwidth management functions which redistribute capacity to existing routes and VPCs according to measured and predicted traffic levels.

As new connections are established they are routed over the only remaining route and it is likely that it will become overloaded leading to connection rejections. Before this situation arises, the Load Balancing component, in the management plane, will register a high traffic load compared to the network wide average and attempt to resolve this problem. However, there are no alternative routes available from the LATEX node, and Load Balancing is unable to resolve the potential problem.

At this point, Load Balancing informs its superior component in the REFORM management hierarchy, Route Design, of an unbalanced network which it is unable to resolve [9]. Route Design then attempts to redesign the VPC network and the routing plan to cope with the new physical topology. In this case potential new routes are considered (VPC2, VPC8 and VPC6; or VPC3 and VPC6) according to existing and predicted traffic and the quality constraints of the video service (for example there may be a maximum number of links and nodes video traffic may be routed over due to maximum allowable delay figures) and a new route is established if possible. The management decisions are implemented in the network by updating route selection tables in the RSAs at the corresponding nodes. As time progresses Load Balancing will update the route selection priorities for the new and the already existing VPCs, to balance the load in the overall network and in the control plane the RSAs will make all new routing decisions according to the new routing table entries.

The above is an outline of the functionality of the REFORM system highlighting the co-operation required in an integrated system covering the control and management planes. The overall problem of fault, resource management and control is handled by a number of layers. Each layer has a well defined limit of responsibility with operational parameters defined and managed by the layer above.

The experimental scenarios should demonstrate the applicability of the REFORM system to a set of representative networks and operational conditions. Future scenarios will be designed to demonstrate more complex situations with large numbers of routes sharing multiple resources. These scenarios will be designed to show the robustness and scalability of the REFORM solutions.

6. Conclusions, Future work

We introduced an integrated framework for ensuring network performance and availability in multi-class ATM networks under normal and fault conditions. We integrated control and management plane functionality, covering the complete failure management cycle (prior, post, normalisation). We combine in one system:

- rapid and reliable network self-healing mechanisms, spread across the control and management layers of the network operation;
- intelligent routing coping with fault conditions, taking into account the multi-service environment;
- intelligent load balancing functionality with inherent self-healing capabilities;

- cost-effective spare resource allocation and appropriate management schemes

We provide a number of generic and specific software components to provide solutions to the problems of fault, and routing, management for ATM networks. These components can be exploited by the existing network operators, or provided to new network operators as part of a packaged solution.

Our goal is also to make substantial contributions to TINA-C especially in the fields of fault and performance management as well as to the definition of the migration path from currently used management frameworks, like TMN, to emerging ones, like TINA. Since we adopted the emerging TINA framework based on a commercial CORBA-type platform, we will evaluate, validate and demonstrate through experimentation in real network environment the suitability of such approaches to network management.

We plan to carry out field trials promoting in parallel the use of advanced networks and services as they are provided by NH environment incorporating real user communities in Europe. This way we demonstrate how the functional capabilities of the developed system can improve performance and reliability of end-to-end connections.

Much of the above work is still on-going, and experimentation is planed for this and the next years. Up to date information is widely available on line [1].

Acknowledgements

REFORM is a research project (July 96 - June 99) of the ACTS programme partially funded by the Commission of the European Union. The authors wish to thank all their project colleagues who contributed in many ways to the formation of the ideas described here. Without them this work and our co-operation for these three years would never have started.

7. References

[1] On-line and always up to date information on the REFORM system is available under: http://www.algo.com.gr/acts/reform

[2] On-line infomation about ACTS programme and NH activities: http://www.uk.infowin.org/ACTS/

[3] Overall Principles and Concepts of TINA, TINA-C Document.

Fault Management and Resource Configuration Management, TINA-C Document.

Connection Management Architecture, TINA-C Document.

Management Architecture, TINA-C Document.

Resource Configuration Architecture, TINA-C Document.

Network Resource Information Model Specification, TINA-C Document.

[4] NMF/X/Open: XoJIDM (X/Open Joint Inter-Domain Management) task force; Proposal "Inter-domain management: specification translation" (from ASN.1 to IDL);

[5] TINA-C: Answers to "Distributed Management Facilities" concerning application of CORBA in TMN Architecture as proposed by TINA-C; Critical assessment to the proposed TINA architectures for resource and fault management.

[6] ITU-T Recommendation M. 3010, Principles for a Telecommunications Management Network, Geneva, October 1992.

ITU-T Recommendation M. 3200, TMN Management Services: Overview, Geneve, October 1992.

ITU-T Recommendation M. 3100 Generic IM

ITU-T Transport IM, functional model

ITU-T SG4 "TMN Architecture and Principles, Management Services, Generic Object Model"

ITU-T SG11 "TMN Protocols and Interface Models"

[7] IEEE JSAC Vol12/1 Jan1994 : special issue.

[8] IEEE Communications magazine: Special issue of September 1995.

[9] D.Griffin, P.Georgatsos "A TMN system for VPC and routing management in ATM networks", Integrated Network Management IV, Proc. of 4th. ISINM 1995, ed. A.S.Sethi, et al., Chapman & Hall, UK, 1995.

[10] P.Georgatsos, D.Griffin "A Management System for Load Balancing through Adaptive Routing in Multi-Service ATM Networks", INFOCOM 1996.

[11] D.Griffin, P.Georgatsos "A General Framework for Routing Management in Multi-Service ATM Networks", submitted to ISINM'97.

[12] REFORM Deliverable D2 "Initial specifications and architecture of the REFORM system", due to June 1997.

[13] "Integrated Communications Management of Broadband Networks", Crete University Press, ed. D.Griffin, ISBN 960 524 006 8, July 1996.

[14] D.Manikis et al. "Signalling Performance Functionality for Multimedia Services in an ATM Integrated Environment", submitted for presentation in IEEE Infocom '97.