Technology Interoperation in ATM Networks: The REFORM System

Panos Georgatsos, Algosystems
Dimitris Makris, National Technical University of Athens
David Griffin, University College London
George Pavlou, University of Surrey
Stelios Sartzetakis, Institute of Computer Science-Foundation for Research and Technology-Hellas
Yves T'Joens, Alcatel Corporate Research Center
Daniel Ranc, Institut National des Telecommunications

ABSTRACT This article addresses the applicability and interoperation of standards and emerging technologies for the operation and management of ATM networks. The issue is tackled from a practical perspective based on experiences gained through the design, realization, and experimentation of a system developed by the ACTS AC208 REFORM project. Assuming an ATM-based network infrastructure offering a range of services with distinct QoS guarantees, the REFORM system encompasses the required functions for ensuring cost-effective network survivability and availability; fast-responding reliable fault detection and self-healing mechanisms, distributed dynamic routing functions with inherent load balancing capabilities, efficient VP layer design, and dynamic network reconfiguration functions. A number of standards and emerging technologies were used for designing and realizing the wide spectrum of functionality incorporated within the REFORM system, including ITU-T OAM I.610 and Q.2931, ATM Forum UNI 3.0 and PNNI v. 1, OMG CORBA and Component Model, TINA NRA and ISO/OSI, and ITU-T TMN. Based on the experience gained, the article discusses and draws conclusions on the applicability, coexistence, and interoperation of the adopted technologies. It is shown that these technologies can coexist, through careful design, to the benefit of network design and operation.

A range of technologies exist which aim at meeting the challenges of today's telecommunications networks and services. These technologies are continually evolving, and new ones are emerging to address the needs arising from the proliferation of current and envisioned telecommunications services. They cover different aspects of network operation and management; from physical to network layer, and from network to service and business management.

A unified technological framework for the analysis, specification, design, deployment, operation, and management of telecommunications networks and services is an attractive goal but seems unfeasible at the moment. Telecommunications is not a subject to be resolved within the realm of one technology. By its very nature, the field is a convergence point of a number of sciences and disciplines, and it is only to be expected that a variety of complimentary and sometimes competing technologies may apply. Interoperability of different technologies has always been and will continue to be a key issue.

In this article we focus on the issue of applicability, coexistence, and interoperability of technologies for network operation and management. We discuss the rationale behind the selection of suitable technologies and present our feedback, based on practical experiences, on their use, integration, and interoperation. The experiences were gained through the design, realization, and experimentation of a system developed in the course of the ACTS AC208 REFORM project [1].

Considering an ATM-based network infrastructure offering multiple services with distinct QoS guarantees, REFORM has specified, designed, and realized a system for ensuring network resilience to changing traffic and topological conditions. The system encompasses a wide spectrum of functionality, spanning both the control and management planes of network operation. It has been based on International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) Q.2931 and ATM Forum (ATM: asynchronous transfer mode) user–network interface (UNI) 3.0 signaling mechanisms, ITU-T I.610-OAM functions, ATM Forum private network–network interface (PNNI) v. 1 distributed routing functions, the Object Management Group (OMG) Component Model for the design of the network layer, Telecommunications Information Network Architecture (TINA), International Organization for Standards open systems interconnection (ISO/OSI), ITU-T telecommunications management network (TMN) principles and architectures for network management design and OMG Common Object Request Broker Architecture (CORBA) for providing distributed interactions between system components. We assess in this article the applicability of these technologies, and demonstrate that they can coexist and interoperate to the benefit of both the design and operation of the network.

The article is organized as follows. The next section introduces the REFORM system. The article then lists the design requirements. We discuss the rationale behind the selection of the technologies for meeting the design requirements, and present their assessment regarding their applicability, integration, and interoperation. The last section contains the conclusions, also highlighting a vision of their evolution.

THE REFORM SYSTEM

Adopting the viewpoint of an ATM-based network provider offering quality of service (QoS)-based, switched, on-demand connectivity services, the REFORM project [1] aims to design, implement, and test a prototype system which provides the necessary means and functions for ensuring network *resilience* within acceptable levels, and in a cost-effective manner. Network resilience encompasses two aspects: *network avail-*

ability, the capability of a network to accept new connections; and *network survivability*, the capability of a network to gracefully recover the service of existing connections under fault conditions. The REFORM system (Fig. 1) covers the configuration, performance and fault management functional areas and comprises two distinct parts [2].

The *control plane system* (embedded in the network elements, NEs) hosts the required connection-oriented network layer functionality in addition to the REFORM-specific fast-responding fault detection, self-healing, and QoS-based dynamic routing functions. This part has been integrated with existing off-the-shelf NEs to form the *REFORM Nodes* which are managed by the REFORM management system (Fig. 1).

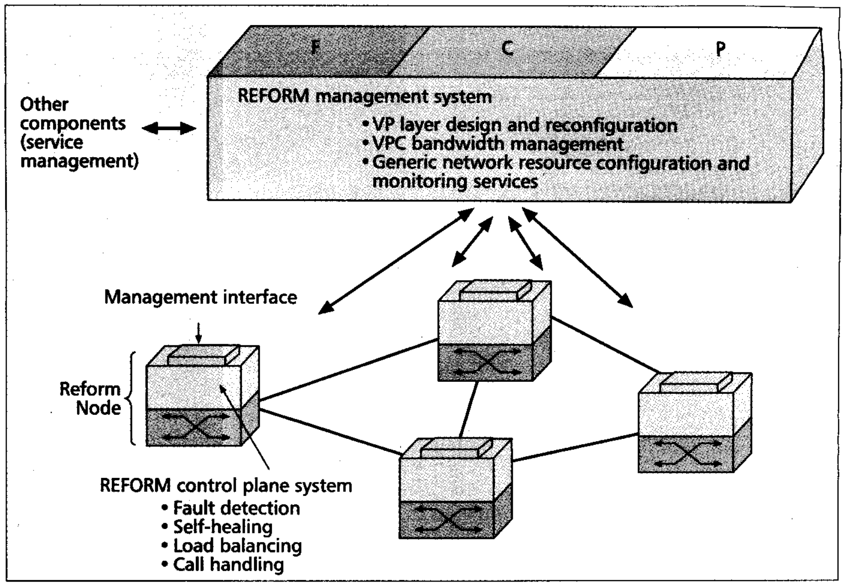

■ **Figure 1.** *Overall REFORM system view.*

The *management plane* system is concerned with the initial configuration and ongoing dynamic management of the ATM virtual path (VP) layer. Specifically, it hosts dynamic virtual path connection (VPC) bandwidth management, VP layer design and dynamic reconfiguration, fault management (filtering and correlation) functions, as well as generic configuration and network resource monitoring functions.

The entire REFORM system operates through a hierarchical model at different levels of time-scale and abstraction. Figure 2 depicts the components in both the management and control plane systems.

Design Issues and Requirements

In addition to general design requirements (scalability, portability, compatibility, interoperability, reusability, modularity, expandability, performance), the following particular aspects were taken into account during the design phase and therefore influenced the selection of the adopted technologies.

Network Element Design — As requirements become ever more demanding and technologies evolve, network terminals and nodes have been transformed from simple network termination and switching equipment to complex systems incorporating an increasing amount of intelligence. As network functionality moves from hardware- to software-based implementations, NE maintenance and evolution are becoming increasingly dependent on the configuration of embedded software. NEs should be designed so that new capabilities may be incorporated and existing ones modified in a cost-effective manner with a minimum of downtime.

Integration and Interoperation of Network Management Functions — The REFORM management system integrates the tasks of network planning and dimensioning with dynamic configuration, and fault and performance management. The network planning and dimension functions lie at the heart of the REFORM system, facilitating cost-effective network operation and enabling the implementation of the operator's business policy regarding service provisioning for the range of services being offered. Dynamic configuration, and fault and performance management are required to continuously optimize the performance of the network according to actual usage levels [3]. Traditionally, these areas have been handled by disjoint systems. Network planning, traffic estimation, and resource dimensioning have largely been offline activities; configuration management systems have been used to support the installation of new equipment and to provision specific customer resources; performance management systems have been limited to performance monitoring tasks, and fault management systems have concentrated mainly on alarm filtering and analysis to support human maintenance teams.

Active Management and Control Plane Interactions for an Intelligent Network Infrastructure — In our view, network management is much more than a data collection exercise for supporting configuration, fault, and performance reports to be subsequently fed to network operators to determine the next course of action. REFORM considers network management functions as built-in, automated, and intelligent facilities which respond to changes in network conditions as and even before they happen. This is necessary considering the complexity of the network environment and the policies required for provisioning advanced services. A human-oriented model of decision making may no longer always be viable for ensuring the cost-effective management of advanced, multiservice networks. Network management functions need to actively interact with the network, exploiting and complementing the built-in capabilities of the NEs, and therefore should be seen as an extension of embedded NE functionality.

Adopted Technologies and Their Assessment

Taking into account the functional objectives and requirements presented in the previous sections, the REFORM system was designed and realized based on a number of different technologies (Fig. 3). The overall design was based on an implementation- and vendor-independent specification — the REFORM functional model [2, 3] — rather than "blindly" following the design principles advocated by the adoption of a particular technology. It is believed that the successful integration of the various technologies in REFORM is attributable to this approach. The following sections discuss the application of the adopted technologies, and our assessment of their applicability and interoperation based on experiences gained through the implementation and experimentation activities.

I.610 OPERATIONS AND MAINTENANCE FUNCTIONS

ITU-T Recommendation I.610 defines a set of operations and principles for maintenance and administration of broadband integrated services digital network (B-ISDN) networks. Cells of a predefined special type are exchanged between network nodes for detecting failures, propagating failure notifications to neighboring nodes, verifying connectivity with loopback cells, and checking connection continuity. I.610 is a mature standard already implemented by a significant number of vendors. Since REFORM offers network survivability at the VP layer, I.610-based operations and maintenance (OAM) F4 continuity check[1] cells were used to monitor protected VPCs.

According to I.610, continuity check cells are transmitted every second, and VPC sink points will raise alarms if they fail to detect cells within 3.5 s. The REFORM requirements [2] state that VCC restoration should take place in the order of 500 msecs, therefore prohibiting the use of the I.610 recommendation as is. Appropriate modifications were made to cater for the restoration time requirements. Specifically, a modified OAM F4 flow was implemented, where cells could be transmitted every 25 ms. By increasing the frequency of continuity cells it was verified through field trials that faults can be reliably detected in milliseconds, leading to restoration times on the order of hundreds of milliseconds, and that this can be achieved with a tolerable overhead [4].

ITU-T Q.2931 AND ATM UNI 3.0

REFORM used networks built from multivendor NEs in its trials. As of today, implementations of the same signaling protocol by different vendors rarely fully interwork and cooperate. Furthermore, the other REFORM control plane components were required to cooperate with the signaling and call handling components for collecting statistics and providing routing information. This level of interaction with embedded capabilities is not feasible with today's commercially available NEs, because access to the internal functionality of the network equipment is very limited. To overcome these problems, the notion of the *REFORM Node* was introduced. The REFORM Node (Fig. 4) incorporates the control plane functions of the REFORM system, allowing vendor-independent interaction with the NEs, and uniform signaling and routing functions throughout the network. This is achieved by isolating vendor-specific technology from the system components through an intermediate adaptation layer to the specific NE. An existing ATM Forum UNI 3.0 signaling stack implementation was modified in order to offer uniform signaling capabilities at both UNI and network–network interface (NNI) reference points of the REFORM Node. However, it should be noted that any NNI protocol (e.g., B-ISUP, PNNI) could have been used. Appropriate modifications according to the REFORM model for QoS provisioning had to be made. This was expected, since standards for QoS provisioning are still under development. Through our prototype implementation it was verified that the adopted signaling technologies may interoperate at the required level with our intelligent dynamic routing and resource management functions (next section), by virtue of the REFORM Node design.

PNNI

Although specified to be deployed in private ATM networks, PNNI v. 1.0 [5] is often used in today's public ATM networks. The PNNI specification can be seen as containing two parts. The first, and most interesting to REFORM, is the routing protocol; the second is the signaling protocol. The PNNI routing protocol was used as the basis for dynamic routing in REFORM. The REFORM QoS-based routing algorithm [2] runs over the PNNI topology information distribution mechanisms.

Experimentation showed the validity of the concepts underlying the PNNI-based dynamic routing scheme in

[1] *According to I-610 "...the continuity check is the only in-service mechanism able to continuously detect for ATM layer defects (as opposed to physical layer defects) in real-time."*

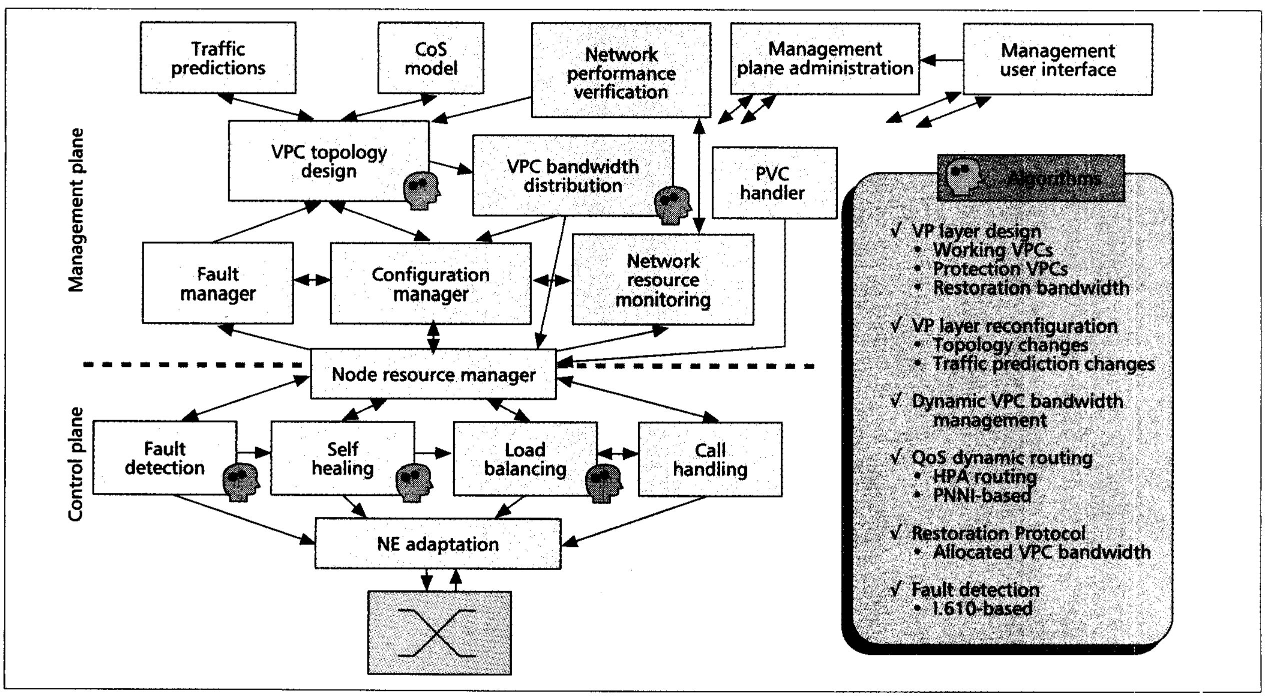

■ **Figure 2.** *REFORM system architecture.*

REFORM [2]. Furthermore, the interoperation of these dynamic routing functions with the self-healing and resource management functions was verified. The PNNI routing protocol specifications constitute a rich platform for building intelligent routing schemes. PNNI routing protocols may coexist with other signaling systems, which support source node or hop-by-hop routing.

The link state information need not be constrained to the metrics proposed in the PNNI specifications, but any other metric in the routing algorithms may rely on in a particular domain. It is debatable whether the specified metrics are appropriate for QoS-based routing, considering the granularity with which this information is renewed (at the cell level). Other metrics, such as those used in the REFORM routing algorithm [2], may be used. After all, the semantics of the link state metrics need only be applicable within a particular domain and therefore need not be standardized across all domains.

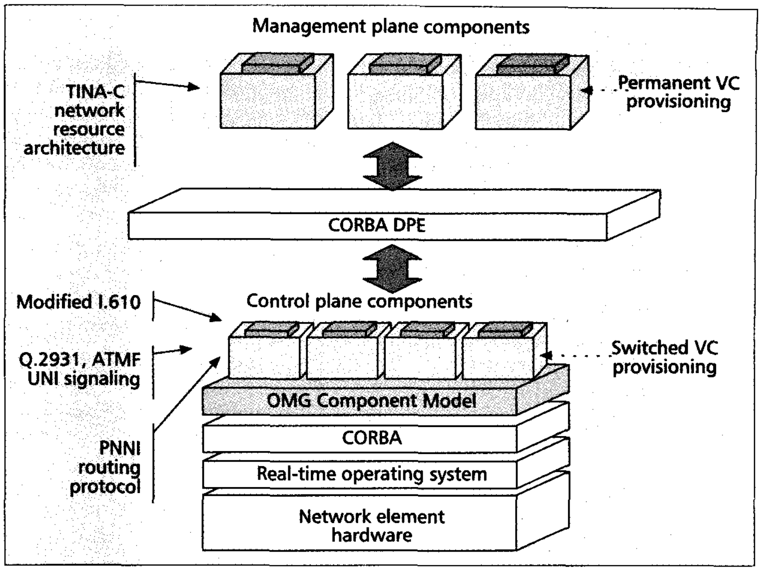

■ **Figure 3.** *Technologies in the REFORM system.*

The PNNI specifications imply a totally dynamic routing scheme through the automatic discovery of routing information (connectivity and, subsequently, routes to reachable destinations). This is clearly a useful feature; however, some routing schemes, such as the one adopted by REFORM, may be of a hybrid nature, combining the merits of centralized and distributed routing. The PNNI specifications should not be seen as dictating the particular routing scheme to be applied, but rather as a framework enabling distributed exchange of topology information.

In the REFORM system, a hybrid routing scheme is adopted (Fig. 5). The centralized part (residing in the management plane) defines the admissible routes per source–destination and distinct class of service (CoS) so that the QoS requirements of the CoSs and certain networkwide cost effectiveness criteria are met. The distributed part is responsible for influencing the routing decisions (from the many possible predefined admissible routes), according to actual traffic conditions, with the purpose of driving the network toward a load balanced state. This hybrid routing scheme maintains the merits of dynamic routing, but at the same time harnesses routing dynamicity to oper-

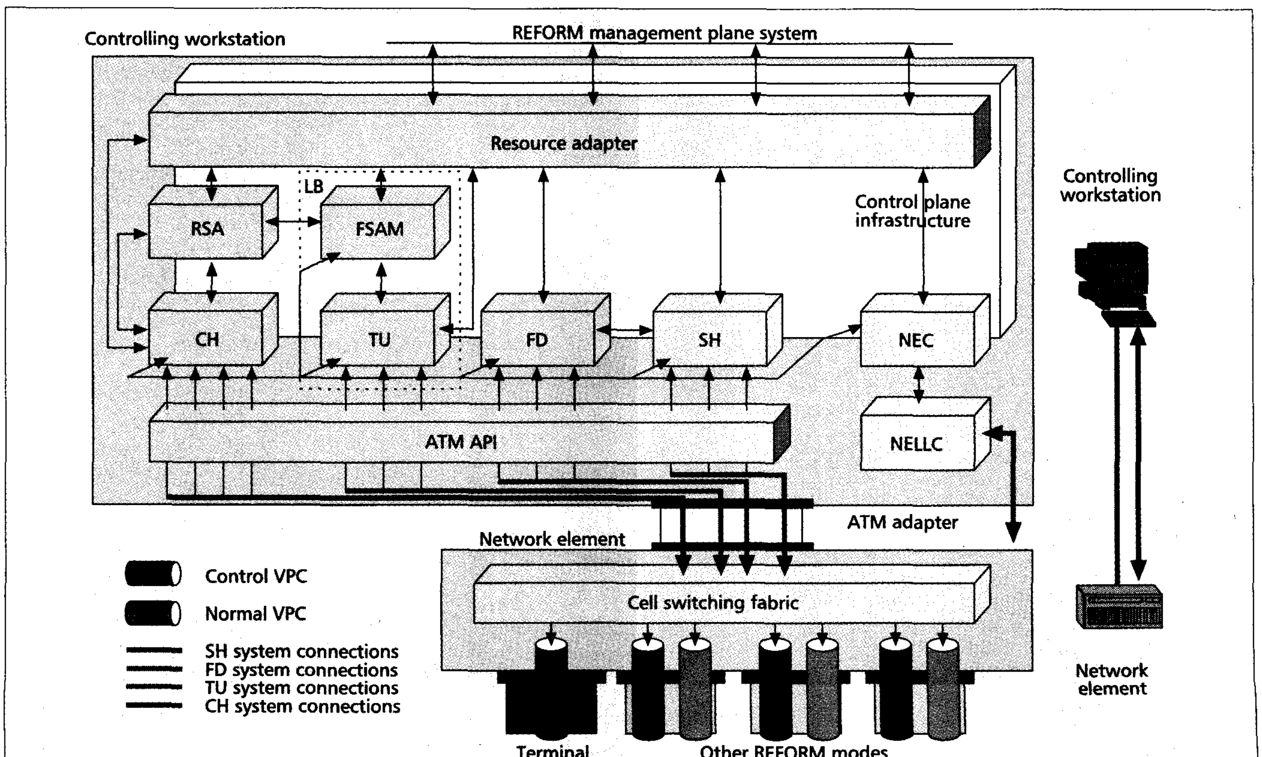

■ **Figure 4.** *The REFORM Node.*

ate within overall network operational policies with respect to QoS provisioning. Experimentation has shown that the distributed PNNI routing protocols can interoperate with overlying network management systems for QoS-based routing.

THE OMG COMPONENT MODEL

The main objective of the OMG Component Model [6] is to cut development costs by reducing the complexity and duration of the software development process. The innovation of this model is that two components may interact without requiring fully standardized interfaces. Each component obeys certain design patterns, and from these patterns the state and communication characteristics of the component may be discovered.

The Component Model has been used in the implementation and integration of the REFORM control plane system. This choice was made to achieve the desired levels of reusability, interoperation, flexibility, and expandability during the implementation, integration, and testing phases of the system. Specifically, the REFORM control plane system is viewed as a "container" of a number of building blocks — components — each fulfilling a specific functional aspect of a network node. Existing functionality can be modified or new components introduced to the embedded control system with a minimum of effort and disruption to system operation. Use of the Component Model allowed the developers to focus on functional aspects rather than on specific details of interprocess communications in the underlying operating systems, and this greatly simplified the testing and integration phases. Our experience has shown that the Component Model technology can be used to facilitate software reuse in development and deployment of telecommunication systems. It advocates a market of off-the-shelf components for upgrading and expanding the capabilities of existing systems with value-added, customized intelligence.

OMG CORBA

CORBA [7] was designed by the OMG with the goal of providing a truly open object distribution environment. The CORBA architecture is composed of a software bus, the object request broker (ORB), on which clients and servers are located. The clients and servers share a common definition of interfaces specified in the Interface Definition Language (IDL), which is independent of implementation language. In REFORM, CORBA was used as the ultimate integration means to allow transparent communication not only between distributed objects, but also between objects located in the same machine (e.g., two components in the same network node). In the same way CORBA insulates distributed applications from network details, it also abstracts implementations from operating system peculiarities.

We are in no doubt that future management frameworks will be based on distributed object technologies, with CORBA being a prime candidate. It is also likely that CORBA can be used in the control plane to support open, object-oriented network layer functionality, although this requires lightweight operations and mappings of its protocols over specific network technologies such as ATM adaptation layer 5 (AAL5). The use of the OMG Component Model in the REFORM control plane was a step in this direction.

The key benefits of CORBA experienced in REFORM are the fast development cycle and hiding of heterogeneity through mappings to different programming languages and operating systems. Another advantage is that since interfaces are specified in IDL, it is easier to integrate components which are implemented in different programming languages; for example, graphical user interfaces (GUIs) may be developed in Java, low-level system components in C/C++. The disadvantages perceived in REFORM are relatively slow operation, increased memory requirements, and instability of the existing implementations. More robust and streamlined implementations are necessary if CORBA is to be adopted as a key enabling technology by the telecommunications industry.

TINA

The TINA [8] initiative aims to provide a framework for all telecommunications software, encompassing components ranging from connection establishment through network and service management to service delivery and operation. TINA adopts the idea that new advanced services may be deployed directly in a distributed processing environment (DPE), and may be designed and implemented according to object-oriented principles and distributed processing techniques. REFORM relied on the TINA Network Resource Architecture (NRA) while designing its management system. The NRA specification covers the principles and concepts for managing networks and is strongly based on the ITU's TMN architecture [9]. The NRA includes the Network Resource Information Model (NRIM), which abstracts the communications resources forming the network infrastructure in a technology-independent model.

The REFORM system is one of the first attempts to validate the TINA NRA in a practical network management application encompassing configuration, performance, and fault management aspects. A conclusion to be drawn is that the TINA NRA and its underlying modeling techniques proved valuable in

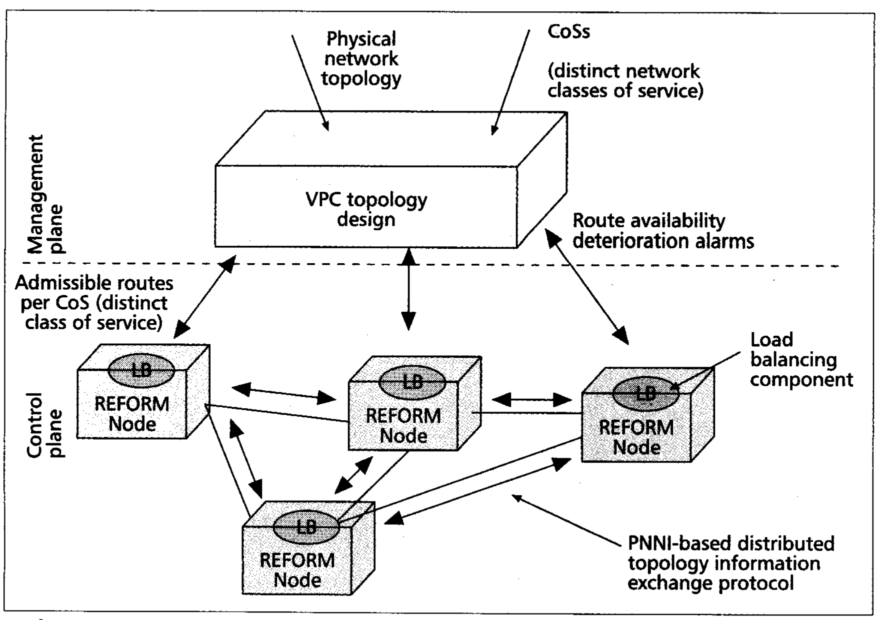

■ **Figure 5.** *Routing in the REFORM system; a hybrid scheme involving centralized and distributed logic.*

realizing the REFORM management system. The aspects of distribution and hierarchy inherent in the specifications combined with the distributed nature of the underlying CORBA environment are desirable when building management systems since they provide a means to address scalability issues.

Certain enhancements in the TINA NRA emerged, which mainly relate to QoS-provisioning functional aspects in multiservice networks [2]. Facilities for network planning, resource dimensioning, QoS provisioning, and dynamic routing management should be evident in the NRA in both the information and computational models. These can now be regarded as generic network functions, being widely recognized in the literature and standardization bodies. With minimum adaptations the REFORM system can interoperate with a TINA-based connection management system, providing added value, resilience, and QoS guarantees to TINA services. It should be stressed that by virtue of its design, the REFORM system can also operate at the same time with signaling-based systems, providing the same added values. This demonstrates that in an integrated service environment the network resources can be managed regardless of the technology used to carry service requests into the network.

TMN AND TINA COEXISTENCE

The ITU-T introduced TMN [9] as a means of provisioning network management systems. TMN relies on ISO/OSI systems management concepts and functions for the communication of management information, and modeling of network and service resources at various levels of abstraction. The TMN is built from a hierarchical system of managers and agents, and can be regarded as a separate network, logically distinct from the network being managed.

An interesting feature of REFORM work is in the mixing of TMN and TINA characteristics in the architecture of the management plane. In many ways, the design is fully inspired by the TMN architectural framework but realized based on TINA design and deployment principles, using CORBA as the underlying DPE. Thus, the REFORM management architecture is a concrete example of TMN-to-TINA migration and coexistence.

The approach relies on a management broker (MB) object which is used to provide access to objects in a CORBA environment through a generic TMN-like interface that supports object navigation through scoping and filtering assertions. The advantage of accessing managed objects through the MB is increased economy in network traffic and sophisticated querying capabilities in a fashion similar to the TMN "culture."

An additional use of TMN technology is in the network resource monitoring component. Monitoring activities in NEs are effected through CORBA-based versions of the OSI workload monitoring and measuring summarization systems management functions. These functions have proved essential for the provisioning of monitoring and measurement tasks in a generic fashion.

RESULTS, CONCLUSIONS, AND A VISION OF EVOLUTION

The REFORM system has been realized, demonstrated, and assessed in functional and performance terms through field trials in networks comprising a range of commercially available ATM switches. The field trials [1] were conducted on several testbeds: the ACTS EXPERT testbed in Basel, Switzerland, and on ATM networks in Norway, Greece, and Japan. The results prove the validity and effectiveness of the underlying functional architecture and the associated technology choices [2].

The REFORM system has shown that different technologies, even with different cultures, can indeed coexist and interoperate, demonstrating benefits in network design and operation. It is important to stress that this was achieved from a clear system design, which began from the system's functional requirements rather than from the doctrines implied by a particular technological framework. Emerging technologies should not treat existing ones as outdated but as their heritage. What matters is the intelligent operation and management of networks, which can be achieved through a range of technologies as appropriate for different aspects of the system, not through the adoption of a single technology to cover everything.

A number of the adopted technologies were not used "as is"; enhancements were made, especially to cater for QoS provisioning according to the specific model and approach developed by the project. However, this is to be expected considering that QoS provisioning is a subject still under investigation.

Our experience indicates that today's existing and emerging technologies, individually and integrated together, have the potential to form the means for cost-effective, scalable QoS provisioning in the near future. If this is the goal, it is believed that more consensus building is required at a higher functional level than the resolution of interoperability issues at detailed technical levels.

Considering that network intelligence is largely realized in software (either embedded in network elements or running in dedicated workstations), software engineering and information processing technologies inevitably play a key role in telecommunications. It is believed that distributed object technologies such as CORBA will form the middleware for executing the required network intelligence. The REFORM system uses CORBA not only for network management, but also for embedded NE functions. CORBA offers the "glue" for connecting network-layer components within NEs, components within the network management system, and components between NEs and the management system. The system performance assessment [2] proves the validity of this approach. A number of optimizations, however, are still required to tailor these technologies to the particular needs of the telecommunications environment.

Considering the ever-increasing need to enhance network flexibility and intelligence, a well-defined, interoperable cut between network software and hardware seems inevitable. This occurred in the computer industry and is also happening in the telecommunications industry. Recognizing the benefits, this view of the future is in line with current networking trends as reported in the literature (open network control, programmable and active networks, nomadic computing). In this context, networks will be built from basic hardware elements (multiprotocol input/output, switching devices, termination modules) on top of which software components will provide the required network intelligence at the operational (e.g., connection management) and management (e.g., monitoring, billing) levels. Network intelligence may also be directly provided by and to users through *programmable interfaces*.

In such a telecommunications environment, it is envisaged that network intelligence will not be provided by static subsystems interoperating through standardized interfaces or appropriate gateways. Instead, network intelligence will be provided through programmable interfaces by means of *easily adaptable components* encompassing specific intelligent capabilities (Fig. 6). These components may be provided by third parties specializing in particular aspects of network intelligence (e.g., connection management, network planning, routing, or billing). In this context, consensus rather than standardization is required with respect to the capabilities of the components,

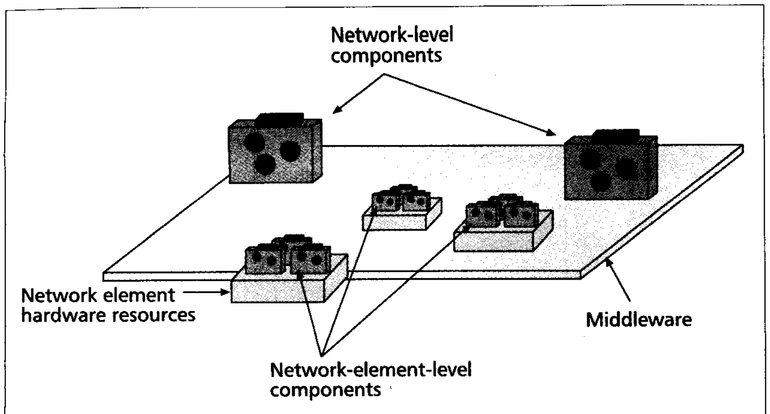

■ **Figure 6.** *Vision of the future: network intelligence provided in the form of easily adaptable components.*

formed by market needs. In this environment, the current distributed object-oriented technologies may also need to evolve. A number of emerging technologies (agents, Component Model) or combinations of them may prove suitable candidates to sponsor the client/server-based object-oriented technologies of today toward a future environment built on easily adaptable components.

The REFORM system entails aspects of the envisaged evolution. The REFORM control and management plane systems build on CORBA to deliver the required network intelligence, which interacts with today's commercially available NEs, treating them as input/output switching devices. Furthermore, the system was built as a collection of components (mimicking their provision by third parties) following the principles of Component Model technology. In this way, NEs and the systems managing them may be automatically equipped on demand with the behavioral capabilities appropriate for their role in the network. The REFORM system realization shows that the envisaged evolution toward flexible telecommunications systems built from off-the-shelf intelligent software and hardware components is feasible in the near future.

Acknowledgments

This article describes work undertaken in the context of the research project AC208 REFORM as part of the ACTS program [10], which is 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. Special thanks go to Costas Stathopoulos, Angelos Kypriotis of Algosystems, Dimitris Manikis, Eleni Mykoniatis of National Technical University of Athens, Keigo Makino of NTT, Fotis Kitsos, Magda Chantzaki of Institute of Computer Science-Foundation of Research and Technology-Hellas, Matthieu Verdier, and Theodore Michalereas of University College London who have undertaken the implementation of the REFORM system.

References

[1] The REFORM Project, on line information, http://ermis.algo.com.gr/acts/reform
[2] The REFORM Project Deliverables: http://ermis.algo.com.gr/acts/reform/deliver.htm
[3] Y. T'Joens, P. Georgatsos, and L. Georgiadis, "Network Reliability in ATM based IBCN: A functional Description of the REFORM System," *Proc. DRCN '98*, Brugge, Belgium, May 1998.
[4] D. Manikis *et al.*, "Performance Evaluation of Self-Healing Functionality in ATM Networks: The REFORM Perspective," *Proc. INTERWORKING '98*, May 1998, pp. 581–90.
[5] ATM Forum, "Private Network-Network Interface Specification v1.0 (PNNI 1.0)," af-pnni-0055.000, Mar. 1996.
[6] "The Common Object Request Broker Architecture and Specification (CORBA)," OMG, 1991.
[7] "CORBA Components," Joint initial submission, Nov. 1997.
[8] TINA-C, "Overall Concepts and Principles of TINA," doc. TB_MDC.018_1.0_94, Feb. 1995.
[9] ITU-T Rec, M.3010, "Principles of a telecommunications management network."
[10] The ACTS program online information: http://www.infowin.org/ACTS

Biographies

Panos Georgatsos (pgeorgat@algo.com.gr) received a B.Sc. degree in mathematics from the National University of Athens, Greece, in 1985, and a Ph.D. degree in computer science from Bradford University, United Kingdom, in 1989. He is currently working in Algosystems S.A., Athens, Greece, responsible for the R&D Group in Telecommunications. His research interests include service quality management, network routing, planning, resource dimensioning, analytic modeling, simulation, and architectures for distributed systems.

Dimitris Makris (dmak@telecom.ntua.gr) received a B.Sc. degree in electrical engineering from National Technical University of Athens (NTUA) in 1995. Since then, he has joined the Telecommunications Laboratory of NTUA and is now finishing his Ph.D. thesis. His research interests are in the area of broadband and mobile communication networks, real-time embedded systems, and distributed computing. He is a member of the Technical Chamber of Greece.

David Griffin (D.Griffin@ee.ucl.ac.uk) received a B.Sc. degree in electronic, computer, and systems engineering from Loughborough University, United Kingdom, in 1988. He joined GEC Plessey Telecommunications Ltd., and in 1993 worked for ICS-FORTH, Heraklion, Greece. He joined University London College in 1996 as a research fellow working on a number of EU research projects in the area of network and service management of broadband networks.

George Pavlou (G.Pavlou@ee.surrey.ac.uk) received his Diploma in electrical and mechanical engineering from NTUA in 1982, and his M.Sc. and Ph.D. degrees in computer science from University College London in 1986 and 1998, respectively. He has been leading research efforts in the area of management of broadband networks and services. Since 1998 he has been elected professor in telecommunications at the University of Surrey.

Stelios Sartzetakis (stelios@ics.forth.gr) received his B.Sc. degree in mathematics from Aristotelian University of Thessaloniki in 1983, and his M.Sc. degree in systems and computer engineering from Carleton University, Ottawa, Canada, in 1986. He joined ICS-FORTH in 1988, where he has been responsible for FORTH's telecommunications infrastructure at large, principally in the creation of FORTHnet. He is responsible for research projects in broadband telecommunications networks and services management.

Yves T'Joens (tjoensy@rc.bel.alcatel.be) received his M.Sc. degree in mechanical engineering from the University of Gent, Belgium, in 1992. He further received an M.Sc. degree in technology from the University of Manchester, Institute of Science and Technology, United Kingdom, and became an aeronautical engineer in 1993 (Universities of Gent and Leuven). Since 1994 he has been working in Alcatel Broadband Switching Division, and since 1996 in the Alcatel Corporate Research Center, presently heading the Internet access research activities.

Daniel Ranc (ranc@hugo.int-evry.fr) received an M.Sc. degree in computer science from the University of Paris. During 1987–1998 he worked at Alcatel Alsthom Recherche, the corporate research laboratory of the Alcatel Alsthom group. He has worked in leading edge object orientation technologies applied to Alcatel's TMN software. In 1998 he joined Institut National des Telecommunications, Evry Cedex, France. His interests are in the area of distributed systems.