1. Introduction

1.1 The Telecommunications Management Network

The increasing sophistication of telecommunications infrastructures and, in particular, the introduction of broadband transmission and switching technologies require the use of sophisticated management facilities. These are important for harnessing the relevant multiplexing capabilities and for supporting the deployment of sophisticated telecommunications services with guaranteed Quality of Service (QoS) characteristics. The ITU-T has been developing the Telecommunications Management Network (TMN) [M3010] as the framework to support the *open* management of telecommunications networks and services. By management it is meant support for planning, provisioning, installing, operating, maintaining and administering these. The TMN corresponds to the *management* plane of the ITU-T Broadband Integrated Services Digital Network (B-ISDN) [BIDSN] reference model. It is an integral part for the operation of the network and services and complements the *control* and *user* B-ISDN planes.

The TMN is a logically separate network, overlaid over the telecommunications network being managed. It comprises a number of management applications, which are physically located “outside” the managed network. These communicate with each other and also access the network elements for the purpose of monitoring and controlling them. The management aspects of the network elements are considered as part of the TMN. A key aspect of the TMN model is its distributed hierarchical nature. Management applications are organised in layers of management responsibility, with each layer potentially comprising many distributed management applications. Higher layers build on the functionality provided by the underlying layers and offer increased encapsulation and abstraction. The TMN proposes the hierarchical separation of management responsibility into element, network, service and business management layers.

A key requirement for the TMN is *openness*. It should be possible for a network operator to buy network elements and management applications from various different suppliers and deploy them together into a TMN. This requires an agreed architecture with well-defined interfaces for managed elements and management applications. Additional key requirements include: the timely reaction to network events; the minimisation of load caused by management traffic; the graceful evolution to new generations of management functionality; the ability to cope with very large sizes of manageable entities; and the ability to cope with the problems caused by the wide geographical distribution of the relevant applications. In addition, good performance is desirable though the TMN does not have the same stringent real time requirements as the control plane functions.

The TMN started being developed by the ITU-T in the mid to late eighties and a key decision was to adopt emerging object-oriented concepts for management information modelling. Object-orientation has been a cultural achievement of information technology that supports better reusability, extensibility, genericity and system integrity [Rumb91]. The TMN needs a distributed object access technology as the basis of its interfaces which should also satisfy the key requirements identified above. The decision at the time was to adopt the emerging OSI Systems Management (OSI-SM) [X700] as the basis for the TMN information architecture and for the distributed object access across standardised interfaces.

OSI-SM was initially developed for the management of OSI data networks. It was the first OSI application to adopt a fully object-oriented approach. The OSI Directory [X500] was another OSI application that adopted some object-oriented principles but fell short of true object-orientation. In OSI-SM, clusters of managed objects model managed entities, being administered by applications in *agent* roles. Managed objects are specified in an abstract object-oriented information specification language, known as the Guidelines for the Definition of Managed Objects (GDMO) [X722]. Applications in *manager* roles access those objects in a distributed fashion in order to implement management policies. Distributed access is supported by the Common Management Information Service and Protocol (CMIS/P) [X710][X711], which was the first generic OSI protocol with “remote method call” semantics. A key aspect of the OSI-SM model is its support for multiple object access through sophisticated predicates and fine-grain control of notifications emitted by managed objects.

While TMN and OSI-SM are different, the former being the framework and the latter the supporting technology, they have typically been treated together, at least until recently. The development of both frameworks followed the ISO/ITU-T tradition, addressing the relevant functionality in an abstract fashion and disregarding software realisation aspects. The combined OSI-SM / TMN framework is mainly a *communications* concept, resulting in well-defined formats of interoperable messages but leaving unspecified aspects concerning the internal structure of applications.

The TMN framework [M3010] appears to be very complex. In addition, OSI-SM has been widely perceived as a technology difficult to implement, resource hungry and potentially resulting in poor performance [Rose91]. This is implicitly recognised in the TMN architecture which includes lightweight interfaces; these may be supported by "less capable" network elements. These lightweight interfaces need to be treated specially through mediation applications and this increases further the complexity of the overall framework.

In summary, the general perception of the TMN in the early nineties was as being too complex and ambitious, with its feasibility, implementability, performance and eventual deployment being widely doubted. In addition, despite the fact the TMN was supposed to be a large scale distributed system, there was no architectural provision for coping with the problems of distribution in terms of location transparency.

1.2 Thesis Motivation

The author developed an interest in OSI-SM as part of his work on the ESPRIT INCA[1] project in 1988. This involved the development of management applications for monitoring the activity of the OSI transport protocol, based on the emerging but incomplete OSI-SM standards [Knig89]. The author started following and contributing to the work of the British Standards Institution (BSI) on the protocol aspects of OSI-SM. The relevant standards reached a state of maturity in 1989 and the author implemented a complete version of CMIS/P under the ESPRIT PROOF project. When subsequently trying to implement management applications, it became clear that a generic approach was necessary in order to harness the complexity of the OSI-SM framework. The inherent object oriented aspects of GDMO could be exploited through mappings to object-oriented languages, hiding the underlying protocol aspects. The author subsequently designed and implemented a first version of a generic OSI-SM agent environment in early 1990 [Pav91a].

At the same time, the importance of the TMN was recognised in the RACE research framework, with four major projects addressing relevant technology aspects, NEMESYS, AIM, ADVANCE and GUIDELINE. The author started working in the NEMESYS project in late 1990 and became aware of a another framework and associated "culture", directly relevant to the realisation of distributed telecommunications software systems: the ISO/ITU-T Open Distributed Processing (ODP) [ODP]. ODP was influenced by the UK Alvey Advanced Networked Systems

[1] Details on the European projects the research work in this thesis was performed are given later, in section 1.4.

Architecture (ANSA) project [ANSA89a]. The latter realised early the need for an object-oriented distributed software framework instead of various application layer protocols. ODP and ANSA had been influential in the RACE research programme, including the NEMESYS project.

It was then that the author conceived the idea of a TMN object-oriented software platform: a development environment that would conform to the protocol aspects of the TMN framework for "on the wire" interoperability but would also adopt ODP concepts and provide easy to use Application Program Interfaces (APIs) for distributed object access. This environment should hide the complexity of the underlying OSI protocols but should provide access to the CMIS multiple object access and fine-grain event control features because of their importance for telecommunications management environments. In addition, the ODP access and location transparencies [X903] should be supported.

The relevant research work reused the author's experience on OSI upper layer protocols and applications and "married" it with ODP and ANSA concepts. The resulting TMN platform took a number of years to complete and is known as the OSI Management Information Service (OSIMIS) [Pav95b]. It was initially used in the NEMESYS project as the basis for first hierarchical TMN system with fully compliant TMN interfaces [Pav91b][Pav92b]. It was subsequently developed further and was used in a number of ESPRIT, RACE and ACTS projects. A substantial amount of research and development work took place in the RACE ICM project [ICM], between 1992 and 1995. It was then used to support one of the most complex TMN systems at the time [Grif96b][Reil96]. OSIMIS has also been used by numerous research institutions and companies, outside European research projects and has influenced a number of commercial developments. Finally, its software abstractions and APIs were input to standardisation work by the Network Management Forum (NMF).

Since the proposed TMN software platform was influenced by ODP concepts, an alternative approach would have been to consider the use of an ODP-based platform and add the necessary functionality for TMN environments. The author considered migrating the OSIMIS platform over ANSA in the first year of the ICM project. The relevant analysis, which is presented in this thesis, showed that this was not possible given a number of TMN requirements not met by ANSA. The Open Software Foundation's (OSF) Distributed Computing Environment (DCE) was considered as the second generation ODP-based platform. Examined carefully, this was a step sideways in comparison to ANSA and exhibited the same deficiencies when considered as the basis for TMN systems.

The Object Management Group's (OMG) Common Object Request Broker Architecture (CORBA) can be seen as the third generation ODP-based platform. This is the first true object-oriented platform and presents a credible alternative to OSI-SM. Since its inception, a number of researchers have been considering its use for telecommunications management, given its generality and superior distribution aspects. Particularly important work has been the comparison of the OSI-SM and OMG object models [Rutt94] and the identification of generic rules for the translation between object specifications in the two frameworks [JIDM95]. The author proposes a solution which maintains the full OSI-SM functionality over CORBA, based on the experience gained by designing and implementing the OSI-SM based platform. A part of this solution was validated through implementation in the ACTS VITAL and REFORM projects.

Throughout the period of this research work, the author has been contributing to the ITU-T TMN architecture group. As a result of the relevant design and implementation experience, the author has proposed a number of simplifications to the TMN architectural framework [Pav96a] which are presented in this thesis.

1.3 Problem Statement and Approach

Given the motivation and history behind the research work in this thesis, the actual thesis statement can be formulated as follows:

The TMN is a powerful and relatively simple framework. The choice of OSI-SM as its base technology satisfies key requirements while problems of distribution can be addressed through the use of the OSI Directory. The TMN inherent object-oriented information specification aspects can be exploited through an object-oriented realisation model that conforms to the "spirit" of the ODP framework. OMG CORBA is the first ODP-influenced technology that satisfies key TMN requirements and can be used to replace seamlessly OSI-SM and OSI Directory as its base technology.

The OSI-SM and TMN frameworks are presented first. This presentation serves both as state-of-the-art but goes further by providing insights to a number of design decisions and associating them to the underlying requirements and fundamental assumptions. A number of modifications and extensions to the TMN framework are presented next. The modifications propose the simplification of the overall framework and rely on the fact that it is later shown that full scale OSI-SM/TMN technology is feasible, performant, easy and economical to provide. The extensions introduce location transparency and discovery facilities through the use of the OSI

Directory, which is chosen because of its federated nature. Location transparency is a very important property of distributed systems.

The key properties of object-oriented distributed frameworks are identified next. These cover ODP aspects such as access and location transparency, but a systematic presentation of the ODP framework is attempted later. The proposed OSI-SM/TMN realisation model and the various incarnations of ODP-based technologies, i.e. ANSA, OSF DCE and OMG CORBA, will be measured against those properties.

It is then shown in detail how the inherent object-oriented information specification aspects of the OSI-SM / TMN framework can be exploited through the use of object-oriented design and implementation methodology that results in an easy-to-use, distributed management platform. Object-oriented mappings for the CMIS service [X710], the Abstract Syntax Notation One (ASN.1) “network data typing” language [X208] and the Guidelines for the Definition of Managed Objects (GDMO) [X723] are presented. The programming language of choice is C++ [Strau86] but the presented concepts are general enough to be used in other object-oriented languages such as Smalltalk [Gold83] or Java [Sun96]. A client or manager mapping to the Tcl scripting language [Oust94] is also presented; this could be used for applications with Graphical User Interfaces (GUIs). A performance analysis and evaluation of the resulting platform shows that it has good performance characteristics and relatively modest resource requirements. Finally, the resulting platform is measured against the properties of object-oriented distributed frameworks and it is shown how it can support TMN applications. Given the fact that the ultimate validation of the proposed framework was achieved through the research and development of TMN systems, work based on the proposed platform is presented in Appendix A.

The ODP framework is presented next, along with related technologies such as ANSA, OSF DCE and OMG CORBA. This is more than a state of the art description, considering the ODP framework and relevant technologies in the context of telecommunications management and examining the potential advantages and open issues. The OSI-SM model is positioned in the ODP framework and it is explained why ANSA or the OSF DCE cannot be used as TMN base technologies. OMG CORBA is described in some detail, since it is the first ODP-based technology that satisfies almost all the key properties of distributed object frameworks. A solution is then presented for mapping OSI-SM onto OMG CORBA without losing any of its expressive power. This solution is based on the design and implementation experience gained from the OSI-SM based platform and considers the issues behind a native CORBA-based telecommunications management environment. Through this approach it is possible to move gradually towards a

CORBA-based solution, protecting existing investment in OSI-SM based technology. A relevant performance analysis examines the potential advantages and drawbacks compared to the OSI-SM approach.

The nature of the research work in this thesis can be thought as having three different facets:

- The modifications and extensions to the TMN architecture is research work of *architectural* nature. It relies on experience from the design, implementation, performance and resource requirements of real TMN systems. It also brings in ODP concepts such as location transparency. These modifications and extensions result in a simpler and more powerful framework.
- The design and development of the OSI-SM-based TMN platform is research work that targets the *validation* of the OSI-SM / TMN framework through *design* and *implementation*. It is though far from simple implementation work because of the requirement to adhere to the ODP vision of an easy-to-use, distributed object access framework with object-oriented APIs. When this research work started, it was not at all clear how CMIS, ASN.1 and GDMO could be mapped onto object-oriented software concepts. The work in this thesis has contributed in this direction.
- Finally, the research work that considers the use of CORBA as the base technology for future TMN systems is *specification* work which attempts to bring together the OSI-SM and CORBA / ODP frameworks. It relies on the understanding of the requirements and capabilities of the two frameworks, both theoretically and practically. A part of the resulting specification has been *validated* through design and implementation.

It should be mentioned that although a performance analysis is presented for both the OSI-SM and CORBA based frameworks, this is *not* a performance type of thesis. The aim of OSI-SM performance analysis is to validate the proposed framework from a performance perspective and to demonstrate that the proposed software architecture has relatively good performance characteristics. The aim of CORBA performance analysis is to compare and contrast it to OSI-SM, examine if the migration towards CORBA brings a performance advantage and assess issues of scalability. In summary, this research work is mostly architectural, specification and design work, validated through implementation.

This research work has been undertaken over a long period of time. As a result, it is inevitable that there exist TMN platform products today which have some of the characteristics of the software architecture presented in this thesis. This has been also aided by the fact that the

OSIMIS environment has been in the public domain since 1991 in order to promote the OSI-SM / TMN concepts. At least three commercial TMN platforms are known to be based on it and possibly many more. On the other hand, important findings regarding this research work have been published early i.e. [Pav91a], [Pav92a], [Pav93a], [Pav93b]. Products with similar functionality appeared in the market place around 1995.

The research work concerning the use of CORBA took place in the last two and a half years, after commercial CORBA implementations with C++ programming language bindings became available. While the research community started working on this issue since 1994, there is not a solution yet that reproduces the full OSI-SM functionality in a native CORBA environment. Most research has concentrated in solutions that support gateways between CORBA and OSI-SM environments. The solution proposed by the author targets native CORBA-based TMN applications throughout a TMN system i.e. a “CORBA to the switch” approach.

The research work in this thesis has being mainly based on the following state of the art work:

- the TMN [M3010] and OSI-SM [X700] architectural frameworks;
- the ODP [ODP] architectural framework;
- the ANSA software platform concepts [ANSA89a][ANSA89b];
- the software abstractions in the ISODE OSI upper layer development environment [ISODE][Rose90];
- the object-oriented design methodologies and principles [Booch91][Rumb91]; and
- the OMG CORBA framework [CORBA], the comparison of the OMG and OSI-SM object models [Rutt94] and the relevant object model mappings [JIDM95].

1.4 The Environment and Evolution of this Research Work

Most of the research work in this thesis has been undertaken in the context of applied research projects funded by the Commission of the European Union (CEU). A brief description of those projects and the evolution of this research work are presented in this section.

Early work on the validation of the evolving OSI-SM framework was initially undertaken in the Integrated Communications Networking Architecture (INCA) project, under the European Specific Research in Information Technology (ESPRIT) framework in 1988. The INCA project was addressing the validation of the emerging OSI upper layer protocols and applications. A design and implementation of the incomplete CMIS/P specifications was produced, together with management applications for monitoring the activity of the OSI transport layer. This work is described in [Knig89] and did not target generic, reusable software model. A procedural implementation paradigm was followed, using the C programming language [Kern78].

This work continued in the ESPRIT Primary Rate ISDN OSI Office Facilities (PROOF) project, which was investigating the issues of gateways between primary rate ISDN and IP or X.25 data networks[2]. A prototype gateway was to be produced in the project and an associated OSI-SM agent had to be designed and developed. It became then clear that a different, generic style for the realisation of OSI-SM agents was necessary, exploiting the object-oriented aspects of GDMO and achieving software reusability and extensibility. The first embryonic ideas on the potential generic structure of OSI-SM agents are described in [Knig90]. The author subsequently undertook research work towards the object oriented decomposition of OSI-SM agents and a first design and implementation was produced in early 1990. This research work is described in [Pav91a] targeting a generic, reusable, object-oriented agent infrastructure. The C++ programming language [Strau86] was used since the first relatively stable, bug-free compilers had become available at that time.

In late 1990, the author started working in the Network Management using Expert Systems (NEMESYS) project of the Research in Advanced Communications in Europe (RACE) framework. The latter was addressing issues in broadband communications and telecommunications services, with TMN being a central theme addressed by a number of other projects. NEMESYS was investigating the applicability of Advanced Information Processing

(AIP) technologies to TMN and the use of object-orientation, distributed systems and ODP principles were central to the project. It was then that the idea of a distributed TMN platform was conceived. The generic agent environment was developed further but it was complemented by generic manager infrastructure and generic applications, including a generic Management Information Base (MIB) browser [Pav92a]. The first hierarchical TMN system with fully compliant TMN interfaces was designed and implemented using the resulting TMN software platform; the functionality of this system, which addresses performance management aspects of Asynchronous Transfer Mode (ATM) networks, is described in [Pav91b][Pav92b]. Towards the end of 1991, this platform became available to companies and research institutions in the RACE programme under the name OSIMIS-2.95. The strange version number was supposed to mean "just before version 3.0", though it took another two years until version 3.0 was released!

The NEMESYS project had two successors in the RACE-II programme, the Integrated Communications Management (ICM) and Pre-Pilot in Advanced Resource Management (PREPARE) projects. The former was addressing ATM configuration and performance management while the latter was concentrating in inter-domain service management; both projects used the TMN architectural framework. The initial idea in ICM was to follow a true ODP approach, adopting either ANSA [ANSA89a] or DCE [DCE] as the platform. Support for location transparency through trading in ANSA and through name services in DCE was perceived as an important feature. On the other hand, mapping the OSI-SM model onto either of those technologies is problematic, for reasons explained in Chapter 4 of this thesis. As a result, OSIMIS was adopted and developed further in that project, including the addition of location transparency through the OSI Directory [X500]. In the PREPARE project OSIMIS was used together with another four commercial platforms (!), two of which were based on OSIMIS.

OSIMIS was developed much further in the course of the ICM project. Various features were added, including: object-oriented ASN.1 and GDMO compilers [Pav96b]; location transparency through the OSI Directory [Stath93][Stath95][Pav96a]; a full version of object-oriented manager infrastructure [Pav94b]; Tcl-based scripting manager infrastructure [Tin95][Pav96d]; intelligent monitoring objects [Pav96c]; generic gateways to the Internet Simple Network Management Protocol (SNMP) [Pav93c][McCar95]; meta-management facilities [Sartz95]; lightweight private key based security services including authentication, integrity and confidentiality services [Bhat96]; and object-based access control [Pav96b]. In parallel to the ICM project, the ESPRIT

---

[2] Such gateways between PR ISDN and IP are a commercial reality today, used a lot by Internet Service Providers (ISPs).

Management in a Distributed Application and Service Environment (MIDAS) contributed also significantly to the OSIMIS development, including a public key based version of the security services [Bhat95] which preceded the lightweight version mentioned above. The OSIMIS environment is described in [Pav93a], [Pav95a] and [Pav96b].

OSIMIS-3.0 was released in early 1993 [Pav93b] and OSIMIS-4.0 in early 1995 [Pav95b]. These have been the two official releases that have been used by numerous research institutions and influenced a number of commercial products. The OSIMIS-4.1 version contains the lightweight security mechanisms but it was never publicly released.

The RACE framework was followed by the Advanced Communications Technologies and Services (ACTS) one, which started towards the end of 1995. By that time, a significant amount of TMN research work had already been accomplished through the ICM and PREPARE projects and through a third important project that addressed service management issues using ODP, the Pan-European Reference Configuration for IBC Service Management (PRISM) project. The ICM, PREPARE and PRISM projects have published their research results in books, [ICM], [PREPARE] and [PRISM] respectively. Given this culmination of the TMN research work in RACE, ACTS concentrated in research on integrated management and service control frameworks adhering to the emerging Telecommunications Information Networking Architecture (TINA) [TINA]. The author has been involved in the VITAL and REFORM projects.

The Validation of Integrated Telecommunications Architecture for the Long-term (VITAL) project attempts to validate the TINA framework. This is done through the design and implementation of an integrated architecture comprising both multimedia service control and ATM network management aspects. On the other hand, The Resource and Fault Restoration and Management (REFORM) project can be thought as the continuation of the ICM project. It adds fault management to the ATM configuration and performance management functions while it uses the TINA instead of the TMN architectural framework.

Since the base technology for TINA systems is a CORBA-based platform known as the Distributed Processing Environment (DPE), an important issue is how the TMN methodologies, principles, specifications and existing applications will be retained and reused. The key issue is how to map the OSI-SM / Directory model to CORBA; this issue has attracted significant attention from the research community. Relevant research work was undertaken by the author in the context of the VITAL and REFORM projects, resulting in a model that retains the OSI-SM power and expressiveness. This model has been validated through implementation [Pav97b][Pav97d] and is presented in Chapter 4 of this thesis. There have been two different implementations of this model: a gateway one, which allows CORBA client or manager objects to access OSI-SM agents; and a native one, in which TMN applications have native CORBA-based interfaces. These implementations were produced in the course of 1996 and 1997 respectively, using a combination of the OSIMIS infrastructure and a commercial CORBA platform.

Figure 1-1 shows the various research projects between 1989 and 1997. It should be noted that the VITAL and REFORM projects are still active. The various OSIMIS versions are also shown. The last two versions which include the CORBA-based gateway and native CORBA agents have not been released outside the ACTS research community.

It should be finally stated that OSIMIS is not only the work of the author. A number of researchers, at UCL and elsewhere, have contributed to it. Given the fact that this research work has been undertaken in the context of collaborative research projects, the role of other researchers to the contributions presented is made clear throughout the thesis. The large majority of the work presented though constitutes research work by the author.

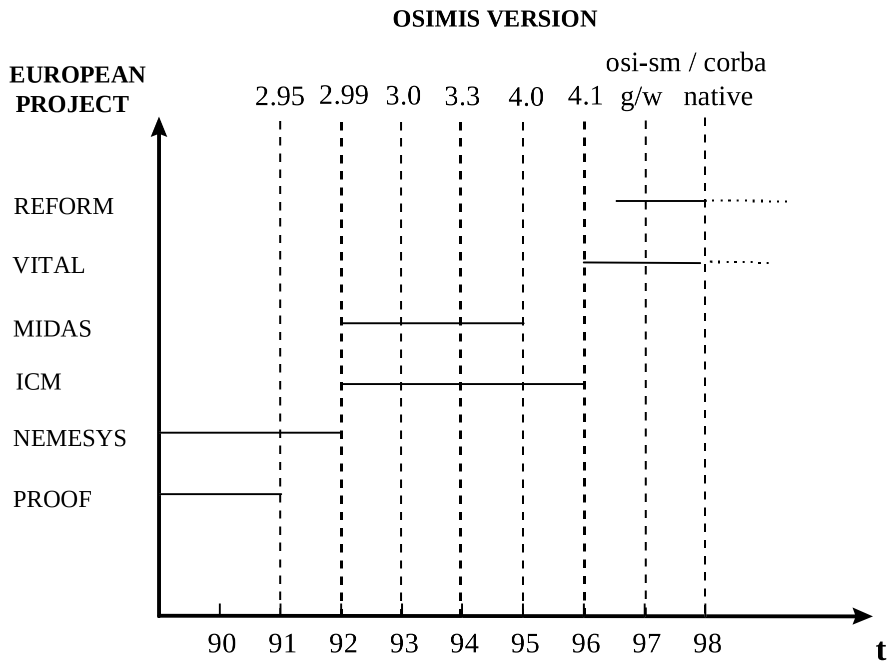

Figure 1-1 European Research Projects and OSIMIS Evolution

1.5 The Difficult Path to the Achievements

At the end of 1997 TMN systems have started to become a reality, especially in the areas of broadband transmission, switching and mobile network environments. The initial vision of the TMN has been translated into real systems eight to nine years later. The research work presented in this thesis has contributed to this direction. The road to this point though has not been easy. Initially, the overall framework was considered too complicated, using abstract object-oriented concepts with no concrete counterparts, unimplementable, resulting in poor performance and requiring excessive computing resources. Proponents of the Internet management approach [SNMP] argued that reliable connection-oriented transport is not suitable for management and object-orientation is an unnecessary complication. ODP proponents argued that the CMIS access model is complicated, difficult to implement and unnecessary. The fact that OSI protocols were used was also seen as an additional negative point. This negative perception was also exacerbated by the fact that early implementations of OSI upper layer protocols had been bad, resulting in cumbersome applications and poor performance.

Three real incidents are described below which demonstrate the general perception of OSI-SM and TMN in the first half of the nineties,. They also try to demonstrate the pioneering role of this research work in demystifying the OSI-SM / TMN framework and establishing its true properties and qualities.

Incident 1: January 1993

In early 1993 a new researcher joined the UCL network and service management group to work in one of the projects. He had previously worked for one of the commercial suppliers of OSI infrastructure and had participated in the development of their OSI-SM product. Since he was going to work with OSIMIS, the author gave to him a small demonstration that involved the manipulation of management information through a number of command line tools and also through a graphical MIB browser. The relevant agent and manager applications were running in different computers on the local network, using full TMN Q_3 protocol stacks over the Internet TCP/IP protocols [Q3], as described in Chapter 3.

That researcher asked a number of questions before the short demonstration but remained silent in the course of it. When the author asked him about his thoughts, he faced an expression of mixed suspiciousness and amazement, being asked back: "*Are you claiming these applications have full OSI-SM functionality and access each other in a fraction of a second*?". A second quick demonstration with the timing flag on for one of the programs revealed around 150 ms for establishing an association, and around 50ms for retrieving the root MIB object and its attributes, based on the computing power available at the time (detailed performance measurements are presented in section 3.7 of Chapter 3). He subsequently mentioned that the products he had encountered needed around 15 seconds (!) for establishing an association and more than one second for retrieving an object under similar conditions. A similar figure for association establishment is also mentioned in [Rose90].

Incident 2: April 1993

The first paper on generic OSI-SM agent infrastructures [Pav91a] and the subsequent early releases of OSIMIS had generated a lot of interest in the UCL work on OSI-SM and TMN. UCL was asked to exhibit OSIMIS in the 3rd IFIP/IEEE Integrated Management Symposium [IM-III] in San Francisco, during April 1993. The relevant costs were levied so that the exhibition could comprise a “product” that was more than a CMIS/P protocol stack, with real OSI-SM agents and generic managers. Since the author was also going to present a tutorial on “Implementing OSI-SM” [Pav93a], a demonstration was prepared. This involved the demonstration of a number of aspects of the OSI-SM model, including event management [X734], log control [X735], MIB browsing etc. The management interactions were shown both over the local network and also over the Internet, accessing agents at UCL. The latter supported the OSI version of TCP/IP MIB [Laba91b].

This demonstration proved to be immensely popular, since it provided a tangible proof that OSI-SM was after all both implementable and performant, in contrast to all the verbal attacks by the SNMP community during the conference. The OSI-SM booth was situated just next to the SNMP one, which included a number of vendors. This gave to the author the following idea: he “invited” SNMP vendors to conduct comparative tests by accessing SNMP agents at UCL over the Internet while the author would also conduct tests by accessing the equivalent OSI-SM agent at UCL through CMIS/P. This would allow to assess both success in retrieving complex information, such as routing tables, and also to compare relative performance. While the numerous visitors of the two booths got really excited, the SNMP vendors declined repeatedly the “invitation”. The author finally tried the experiment himself by using the generic SNMP manager that comes with the ISODE package [ISODE]. The result was that the OSI-SM based experiment succeeded every time, exhibiting much better performance characteristics due to intelligent remote object access through “scoping and filtering”.

Incident 3: September 1995

European research projects are typically audited every year. The final audit for the ICM project took place in September 1995 and included a demonstration after the project's presentation and questions by the auditors. The demonstration comprised the ICM TMN system, described in detail in [Gri96b] and [ICM], which provided ATM Virtual Path Connection and Routing Management (VPCRM) services and also ATM-based Virtual Private Network (VPN) services. This system was the most complex real TMN built at the time (and maybe to date), comprising twelve different types of TMN Operations Systems (OSs). Instances of those existed in different domains, providing end-to-end VPN services and intra-domain VPCRM services. The system had operated over real networks, but for the purpose of the demonstration it was operating over a simulated network of eight nodes. The processing power hungry simulator was running on one workstation, posing as a set of ATM switches with individual TMN Q_3 interfaces, while the operations systems were running on another workstation.

The demonstration was given in parallel to a slide presentation, which was explaining the various features of the system. The auditors could see those features and relate them to action on the screens of the two workstations. While the demonstration was going on, one of the auditors asked the author: "*Are you sure all the OSs in your TMN system have full Q_3 interfaces and do not simply interwork through an ad-hoc lightweight mechanism?*" After the affirmative reply, a second question followed: "*Can you then please show the trace of CMIS/P messages these applications send and receive?*". We had to stop the system and re-start it, configuring relevant output and redirecting it to log files so that we could observe those. The expression on the auditors' faces was little short of amazement since they had totally different ideas about the overhead of a fully compliant TMN OS.

In summary, TMN suffered initially from bad design and engineering. The fact early versions of OSIMIS were made publicly available early enough helped a new generation of better TMN products. Despite the relevant progress though, it is still non uncommon to see requirements for workstations with at least 96 Mb of memory and very powerful CPUs in order to run a TMN OS or do TMN system development. This is in contrast to the relatively lightweight OSIMIS approach, which is acknowledged in [Deri97] when comparing various management technologies and platforms.

1.6 Overview and Style of this Thesis

This thesis is organised in the following fashion. This chapter was the introduction to the thesis while Chapter 5 presents the summary and conclusions. Chapters 2, 3 and 4 constitute the core material of the thesis and can be thought as “super chapters”, each addressing a set of related topics. These topics are presented in the first level sections of those chapters and can be thought as “chapters within a chapter”. State of the art work is presented throughout this thesis, either in separate sub-sections or just before the author’s own work on a particular subject. For example, research work on CMIS APIs is presented just before the author’s own work on the subject.

Chapter 2 addresses the OSI-SM / TMN architectural aspects. It starts with two sub chapters on OSI-SM and TMN respectively, which serve mainly as of state of the art description but they also provide insights to a number of design decisions. The third sub-chapter proposes a number of modifications and extensions to the TMN model and architecture, supported by a relevant analysis. This constitutes the first research contribution of the thesis.

Chapter 3 addresses the realisation of the simplified and extended OSI-SM / TMN framework in an object-oriented fashion. The target to provide an easy-to-use distributed management platform that exhibits characteristics of object-oriented distributed processing frameworks. This chapter comprises a number of sub-chapters as follows. An introduction to object-oriented distributed systems is presented first, identifying their key properties. Issues behind realising the CMIS/P protocol are presented next, examining potential policies for the relevant API. Alternative lightweight mappings that avoid the use of a full OSI stack are also considered. The next sub-chapter examines issues behind treating ASN.1 data types in an object-oriented fashion. This leads to the next two sub-chapters that examine object-oriented mappings for managing and managed objects respectively. The issues behind synchronous and asynchronous remote execution models are then presented. A performance analysis and evaluation of the overall framework follows. Finally, the proposed framework is further validated against the desired properties of object-oriented distributed systems and the needs of TMN applications as identified in Chapter 2.

Chapter 4 considers the mapping of the OSI-SM / TMN model onto emerging distributed object frameworks that follow ODP principles, using OMG CORBA as the target framework and technology. An introduction to ODP is presented first. This serves as state of the art description but also positions OSI-SM in the ODP framework. Different incarnations of ODP-based technologies are presented next. These include ANSA, the OSF DCE and OMG CORBA, the latter considered in more detail. This is more than a state of the art description which positions those technologies in the context of telecommunications management and examines their characteristics against the identified properties of object-oriented distributed systems. The next sub-chapter examines in detail the issues behind using OMG CORBA instead of OSI-SM in TMN environments and proposes a solution that retains the expressive power of the latter. Finally, a performance analysis and evaluation of the CORBA-based framework is presented.

Chapter 5 brings together the various strands of the thesis and summarises the main findings. It explains the significance of this thesis and also identifies areas in which work could be developed further. A number of appendices follow at the end, together with the acronyms and bibliography.

It is also worth mentioning related issues that are *not* discussed in this thesis. Intrusive and inter-domain management is impossible without security. This thesis does not address security issues, though the author has been involved in relevant research work [Bhat96][Pav96b]. Another management paradigm is that of moving code rather than data between management applications. This was first proposed in [Yemi91] and is known as "management by delegation". The advent of languages such as Java [Sun96] and GUI technologies such as the World Wide Web (WWW) open new possibilities for management using mobile code. These will certainly have an impact on telecommunications management architectures in the long term but they are not examined in this thesis. Finally, another approach to management that has evolved over the last years is domain and policy based management [Slom89][X749]. This is somehow orthogonal to the work presented in this thesis and is not addressed here.

Considering the style of this thesis, a number of C++ object class specifications and code extracts are presented in Chapter 3 in order to demonstrate the various features of the proposed software infrastructure. These are fundamental since they demonstrate the relevant properties and user-friendliness, so they could have *not* been moved to appendices. In a similar fashion, CORBA design aspects are demonstrated through the use of specifications in the OMG Interface Definition Language (IDL) in Chapter 4. The Object Modelling Technique (OMT) [Rumb91] together with objects instance diagrams are used to demonstrate aspects of object-oriented design.

Finally, as the reader may have already noticed, "we" or "our" is used instead of "I" or "mine" when describing research work by the author throughout this thesis [3].

---

[3] This has nothing to do with the "royal *we*" but simply reflects the author's writing style.