

# The OSIMIS Platform: Making OSI Management Simple

*George Pavlou, Kevin McCarthy, Saleem Bhatti, Graham Knight*  
*Department of Computer Science, University College London, Gower*  
*Street, London, WC1E 6BT, UK*  
*tel: +44 71 380 7215 fax: +44 71 387 1397*  
*e-mail: g.pavlou k.mccarthy s.bhatti g.knight @cs.ucl.ac.uk*

## Abstract

The OSIMIS (OSI Management Information Service) platform provides the foundation for the quick, efficient and easy construction of complex management systems. It is an object-oriented development environment in C++ [1] based on the OSI Management Model [2] that hides the underlying protocol complexity (CMIS/P [3][4]) and harnesses the power and expressiveness of the associated information model [5] through simple to use Application Program Interfaces (APIs). OSIMIS combines the thoroughness of the OSI models and protocols with advanced distributed systems concepts pioneered by ODP to provide a highly dynamic distributed information store. It also combines seamlessly the OSI management power with the large installed base of Internet SNMP [6] capable network elements. OSIMIS supports particularly well a hierarchical management organisation through hybrid manager-agent applications and may embrace a number of diverse technologies through proxy systems. This paper explains the OSIMIS components, architecture, philosophy and direction.

## Keywords

Network, Systems, Application Management, Distributed Systems, Platform, API

## 1 INTRODUCTION AND OVERVIEW

OSIMIS is an object-oriented management platform based on the OSI model [2] and implemented mainly in C++ [1]. It provides an environment for the development of management applications which hides the details of the underlying management service through object-oriented Application Program Interfaces (APIs) and allows designers and implementors to concentrate on the intelligence to be built into management applications rather than the mechanics of management service/protocol access. The manager-agent model and the notion of managed objects as abstractions of real resources are used but the separation between managing and managed systems is not strong in engineering terms: a management application can be in both roles and this is particularly true in situations where a management system is decomposed according to a hierarchical logical layered approach.

In fact, OSIMIS was designed from the beginning with the intent to support the integration of existing systems with either proprietary management facilities or different management models. Different methods for the interaction with real managed resources are supported, encompassing loosely coupled resources as it is the case with subordinate agents and management hierarchies.

The fact that the OSI model was chosen as the basic management model facilitates the integration of other models, the latter usually being less powerful, as is the case with the Internet SNMP [6] and the emerging OMG CORBA based technologies [7]. OSIMIS provides already a generic application gateway between CMIS and SNMP [8][9] while a similar approach for integrating OSI management and the OMG CORBA framework may be pursued in the future. Using OSI management as an end-to-end integrating model is also in line with the NMF OMNIPoint [10] approach which suggests a multiple technology architecture for future heterogeneous environments.

OSIMIS uses the ISODE (ISO Development Environment) [11] as the underlying OSI communications mechanism but it may also be decoupled from it when the currently ongoing work on supporting the X/Open XOM/XMP/XDS [12] APIs is finalised. The advantage of the ISODE environment though is the provision of services like FTAM and a full implementation of the OSI Directory Service (X.500) which are essential in complex management environments. Also a number of underlying network technologies are supported, namely X.25, CLNP and also TCP/IP through the RFC1006 method. These constitute the majority of currently deployed networks while interoperation of applications across any of these is possible through Transport Service Bridging.

OSIMIS has been and is still being developed in a number of European research projects, namely the ESPRIT INCA, PROOF and MIDAS and the RACE NEMESYS and ICM. It has been used extensively in both research and commercial environments and has served as the management platform for a number of other ESPRIT and RACE projects in the TMN and distributed systems and service management areas. OSIMIS has been fully in the public domain until version 3.0 to show the potential of OSI management and serve as a benchmark implementation while it is still freely available to academic and research institutions for non-commercial use.

### **Components and Architecture**

OSIMIS as platform comprises the following types of support:

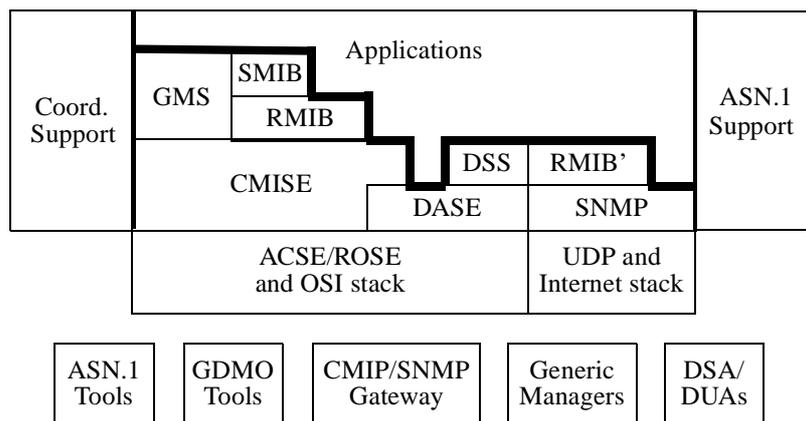
- high level object-oriented APIs realised as libraries
- tools as separate programs supporting the above APIs (compilers/translators)
- generic applications such as browsers, gateways, directory servers etc.
- specific useful management applications

Some of these services are supported by ISODE and these are:

- the OSI Transport (class 0), Session and Presentation protocols, including a lightweight version of the latter that may operate directly over the Internet TCP/IP
- the Association Control and Remote Operations Service Elements (ACSE and ROSE) as building blocks for higher level services
- the File Transfer Access and Management (FTAM) and Directory Access Service Element (DASE)
- a ASN.1 compiler with C language bindings (the pepsy tool)
- a Remote Operations stub generator (the rosy tool)
- a FTAM service for the UNIX operating system
- a full Directory Service implementation including an extensible Directory Service Agent (DSA) and a set of Directory User Agents (DUAs)
- a transport service bridge allowing interoperability of applications over different types of networks

OSIMIS is built as an environment using ISODE and is mostly implemented in the C++ programming language. The services it offers are:

- an implementation of CMIS/P using the ISODE ACSE, ROSE and ASN.1 tools
- an implementation of the Internet SNMP over the UNIX UDP implementation using the ISODE ASN.1 tools
- high-level ASN.1 support that encapsulates ASN.1 syntaxes in C++ objects
- an ASN.1 meta-compiler which uses the ISODE pepsy compiler to automate to a large extent the generation of syntax C++ objects
- a Coordination mechanism that allows to structure an application in a fully event-driven fashion and can be extended to interwork with similar mechanisms
- a Presentation Support service which is an extension of the coordination mechanism to interwork with X-Windows based mechanisms
- the Generic Managed System (GMS) which is an object-oriented OSI agent engine offering a high level API to implement new managed object classes, a library of generic attributes, notifications and objects and systems management functions
- a compiler for the OSI Guidelines for the Definition of Managed Objects (GDMO) [19] language which complements the GMS by producing C++ stub managed objects covering every syntactic aspect and leaving only behaviour to be implemented
- the Remote and Shadow MIB high level object-oriented manager APIs
- a Directory Support service offering application addressing and location transparency services
- a generic CMIS to SNMP application gateway driven by a translator between SNMP and OSI GDMO MIBs
- a set of generic manager applications (MIB browser and other)
- agents for the OSI version of the Internet TCP/IP MIB (native version) and the OSI transport protocol



**Figure 1** OSIMIS Layered Architecture and Generic Applications.

The OSIMIS services and architecture are shown in Figure 1. In the layered part, applications are programs while the rest are building blocks realised as libraries. The lower part shows the generic applications provided; from those the ASN.1 and GDMO tools are essential in providing off-line

support for the realisation of new MIBs. The thick line indicates all the APIs an application may use. In practice though most applications use only the Generic Managed System (GMS) and the Remote MIB (RMIB) APIs when acting in agent and manager role respectively, in addition to the Coordination and high-level ASN.1 support ones. The latter are used by other components in this layered architecture and are orthogonal to them, as such they are shown aside. Directory access for address resolution and the provision of location transparency may or may not be used, while the Directory Support Service (DSS) API provides more sophisticated searching, discovery and trading facilities.

## 2 THE ISO DEVELOPMENT ENVIRONMENT

The ISO Development Environment (ISODE) [12] is a platform for the development of OSI services and distributed systems. It provides an upper layer OSI stack that conforms fully to the relevant ISO/CCITT recommendations and includes tools for ASN.1 manipulation and remote operations stub generation. Two fundamental OSI applications also come with it, an extensible full Directory Service (X.500) [20] and File Transfer (FTAM) [21] implementations. ISODE is implemented in the C programming language [22] and runs on most versions of the UNIX operating system. ISODE does not provide any network and lower layer protocols e.g. X.25, CLNP, but relies on implementations for UNIX-based workstations which are accessible through the kernel interface. The upper layer protocols realised are the transport, session and presentation protocols of the OSI 7-layer model. Application layer Service Elements (ASEs) are also provided as building blocks for higher level services, these being the Association Control, Remote Operations and Reliable Transfer Service Elements (ACSE, ROSE and RTSE). These, in conjunction with the ASN.1 support, are used to implement higher level services. A special lightweight presentation layer is also provided that runs directly on top of TCP; this may be used for the CMOT (CMIP over TCP) [16] stack. In engineering terms, the ISODE stack is a set of libraries linked with applications using it.

ASN.1 manipulation is very important to OSI distributed applications. The ISODE approach for a programmatic interface (API) relies in a fundamental abstraction known as Presentation Element (PE). This is a generic C structure capable of describing in a recursive manner any ASN.1 data type. An ASN.1 compiler known as pepsy is provided with C language bindings, which produces concrete representations i.e. C structures corresponding to the ASN.1 types and also encode/decode routines that convert those to PEs and back. The presentation layer converts PEs to a data stream according to the encoding rules (e.g. BER) and the opposite. It should be noted that X/Open has defined an API for ASN.1 manipulation known as XOM [13] which, though similar in principle to that of ISODE, is syntactically very different. Translations between the two should be possible and such approaches are being investigated.

One of the most important concepts pioneered in ISODE is that of interworking over different lower layer protocol stacks which is realised through Transport Service bridging (TS-bridge) [15]. ISODE provides an implementation of the ISO Transport Protocol (TP) class 0 over X.25 or even over the Internet TCP/IP using the RFC1006 method [23], in which TCP is treated as a reliable network service. The ISODE session protocol may also run over the ISO TP class 4 and the Connectionless Network Protocol (CLNP). Transport service bridges, which are simple relaying applications similar to Interworking Units (ITUs), may be used to link subnetworks of all these different communities and to provide end-to-end interoperability hiding the heterogeneity of the

underlying network technology. The combinations mentioned before constitute the vast majority of currently deployed networks.

### 3 MANAGEMENT PROTOCOL AND O-O ABSTRACT SYNTAX SUPPORT

OSIMIS is based on the OSI management model as the means for end-to-end management and as such it implements the OSI Common Management Information Service/Protocol (CMIS/P). This is implemented as a C library and uses the ISODE ACSE and ROSE and its ASN.1 support. Every request and response CMIS primitive is realised through a procedure call. Indications and confirmations are realised through a single “wait” procedure call. Associations are represented as communication endpoints (file descriptors) and operating system calls e.g. the Berkeley UNIX `select(2)` can be used for multiplexing them to realise event-driven policies.

The OSIMIS CMIS API is known as the MSAP API, standing for Management Service Access Point. It was conceived much before standard APIs such as the X/Open XMP were specified and as such it does not conform to the latter. Having been designed specifically for CMIS and not for both CMIS and SNMP as the XMP one, it hides more information and may result in more efficient implementations. The reason this is a procedural C object-based and not a fully object-oriented implementation is to conform to the ISODE style, the trend of industry APIs and to be easily integrated in diverse environments. Higher-level object-oriented abstractions that encapsulate this functionality and add much more can be designed and built as explained in section 6.

OSIMIS offers as well an implementation of the Internet SNMPv1 and v2 which is used by the generic application gateway between the two. This uses the UNIX implementation of the Internet UDP and the ISODE ASN.1 support and is implemented in much the same fashion as CMIS/P, without conforming to XMP. Applications using CMIS need to manipulate ASN.1 types for the CMIS managed object attribute values, action, error parameters and notifications. The API for ASN.1 manipulation in ISODE is different to the X/Open XOM. Migration to XOM/XMP is possible through thin conversion layers so that the upper layer OSIMIS services are not affected. Such conversion infrastructure was first designed in the ESPRIT MIDAS while it has been implemented in the RACE ICM project.

Regarding ASN.1 manipulation, it is up to an application to encode and decode values as this adds to its dynamic nature by allowing late bindings of types to values and graceful handling of error conditions. From a distributed programming point of view this is unacceptable and OSIMIS provides a mechanism to support high-level object-oriented ASN.1 manipulation, shielding the programmer from details and enabling distributed programming using simply C++ objects as data types. This is achieved by using polymorphism to encapsulate behaviour in the data types determining how encoding and decoding should be performed through an ASN.1 meta-compiler which produces C++ classes for each type. Encode, decode, parse, print and compare methods are produced together with a get-next-element one for multi-valued types (ASN.1 SET OF or SEQUENCE OF). Finally, the very important ANY DEFINED BY construct is automatically supported through a table driven approach, mapping types to syntaxes. This high-level OO ASN.1 approach is used by higher level OSIMIS services such as the GMS, RMIB and SMIB ones.

## 4 APPLICATION COORDINATION SUPPORT

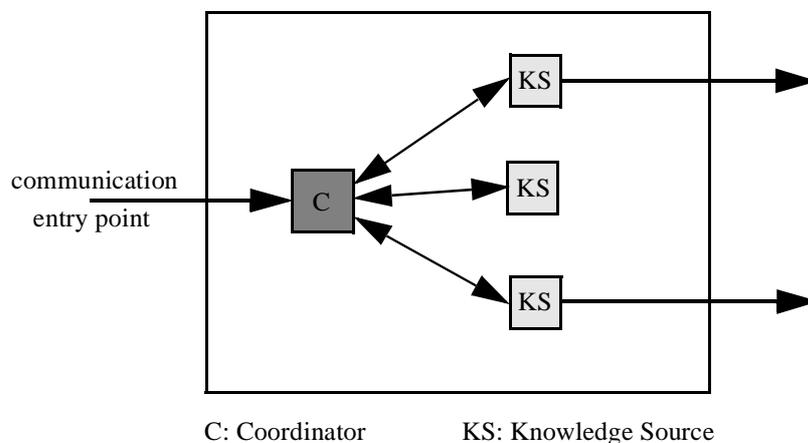
Management and, more generally, distributed applications have complex needs in terms of handling external input. Management applications have additional needs of internal alarm mechanisms for arranging periodic tasks in real time (polling etc.) Furthermore, some applications may need to be integrated with Graphical User Interface (GUI) technologies which have their own mechanisms for handling data from the keyboard and mouse. In this context, the term application assumes one process in operating systems terms.

There are in general different techniques to organise an application for handling both external and internal events. The organisation needs to be event driven though so that no resources are used when the system is idle. The two major techniques are:

- a. use a single-threaded execution paradigm
- b. use a multi-threaded one

In the first, external communications should follow an asynchronous model as waiting for a result of a remote operation in a synchronous fashion will block the whole system. Of course, a common mechanism is needed for all the external listening and demultiplexing of the incoming data and this is a part of what the OSIMIS Application Coordination Support provides. In the second, many threads of control can be executing simultaneously (in a pseudo-parallel fashion) within the same process, which means that blocking on an external result is allowed. This is the style of organisation used by distributed systems platforms as they are based on RPC which is inherently synchronous with respect to client objects performing remote operations to server objects.

The advantage of the first mechanism is that it is supported by most operating systems and, as such, it is lightweight and efficient while its drawbacks are that it introduces state for handling asynchronous remote operation results. The second mechanism allows more natural programming in a stateless fashion with respect to remote operations but it requires internal locking mechanisms and re-entrant code. In addition, such mechanisms are not yet commonly supported by operating systems and as such are not very efficient. An additional problem in organising a complex application concerns the handling of internal timer alarms: most operating systems do not “stack” them i.e. there can only be one alarm pending for each process. This means that a common mechanism is needed to ensure the correct usage of the underlying mechanism.



**Figure 2** The OSIMIS Process Coordination Support Model.

OSIMIS provides an object-oriented infrastructure in C++ which allows to organise an application in a fully event-driven fashion and a single-threaded execution paradigm, where every external or internal event is serialised and taken to completion on a “first-come-first-served” basis. This mechanism allows the easy integration of additional external sources of input or timer alarms and it is realised by two C++ classes: the Coordinator and the Knowledge Source (KS). There should always be one instance of the Coordinator or any derived class in the application while the Knowledge Source is an abstract class that allows to use the coordinator services and integrate external sources of input or timer alarms. All external events and timer alarms are controlled by the coordinator whose presence is transparent to implementors of specific KSs through the abstract KS interface. This model is depicted in Figure 3.

This coordination mechanism is designed in such a way as to allow integration with other systems’ ones. This is achieved through special coordinator derived classes which will interwork with a particular mechanism. This is achieved by still controlling the sources of input and timer alarms of the OSIMIS KSs but instead of performing the central listening, these are passed to the other system’s coordination mechanism which becomes the central one. This is needed for OSIMIS agents which wish to receive association requests since ISODE imposes its own listening mechanism which hides the Presentation Service Access Point (PSAP) on which new ACSE associations are accepted. A similar mechanism is needed for Graphical User Interface technologies which have their own coordination mechanisms. In this case, simply a new special coordinator class is needed for each of them. At the moment, the X-Windows Motif, the TCL/TK interpreted language and the InterViews graphical object library are integrated.

## 5 THE GENERIC MANAGED SYSTEM

The Generic Managed System (GMS) provides support for building agents that offer the full functionality of the OSI management model, including scoping, filtering, access control, linked replies and cancel-get. OSIMIS supports fully the Object Management [24], Event Reporting [25] and Log Control [26] Systems Management Functions (SMFs), the qualityofServiceAlarm notification of the Alarm Reporting one and partly the Access Control [27], Metric [28] and Summarization [29] objects. In conjunction with the GDMO compiler it offers a very high level API for the integration of new managed object classes where only semantic aspects (behaviour) need to be implemented. It also offers different methods of access to the associated real resources, including proxy mechanisms, based on the Coordination mechanism.

The Generic Managed System is built using the coordination and high level ASN.1 support infrastructure and most of its facilities are provided by three C++ classes which interact with each other:

- the CMISAgent, which provides OSI agent facilities
- the MO which is the abstract class providing generic managed object support
- the MOClassInfo which is a meta-class for a managed object class

The GMS library contains also generic attribute types such as counter, gauge, counterThreshold, gaugeThreshold and tideMark and specific attributes and objects as in the Definition of Management Information (DMI) [30], which relate to the SMFs. The object-oriented internal structure of a managed system built using the GMS in terms of interacting object instances is shown in Fig. 5.

## 5.1 The CMIS Agent

The CMISAgent is a specialised knowledge source, more specifically a specific IsodeAgent, as it has to accept management associations. There is always one only instance of this class for every application in agent role. Its functions are to accept or not associations according to authentication information, check the validity of operation parameters, find the base object for the operation, apply scoping and filtering, check if atomic synchronisation can be enforced, check access control rights and then apply the operation on the target managed object(s) and return the result(s)/error(s).

There is a very well defined interface between this class and the generic MO one which is at present synchronous only: a method call should always return with a result e.g. attribute values or error. This means that managed objects which mirror loosely coupled real resources and exercise an “access-upon-external-request” regime will have to access the real resource in a synchronous fashion which will result in the application blocking until the result is received. This is only a problem if another request is waiting to be served or if many objects are accessed in one request through scoping. Threads would be a solution but the first approach will be a GMS internal asynchronous API which is currently being designed. It is noted that the CMISAgent to MO interface is bidirectional as managed objects emit notifications which may be converted to event reports and passed to the agent.

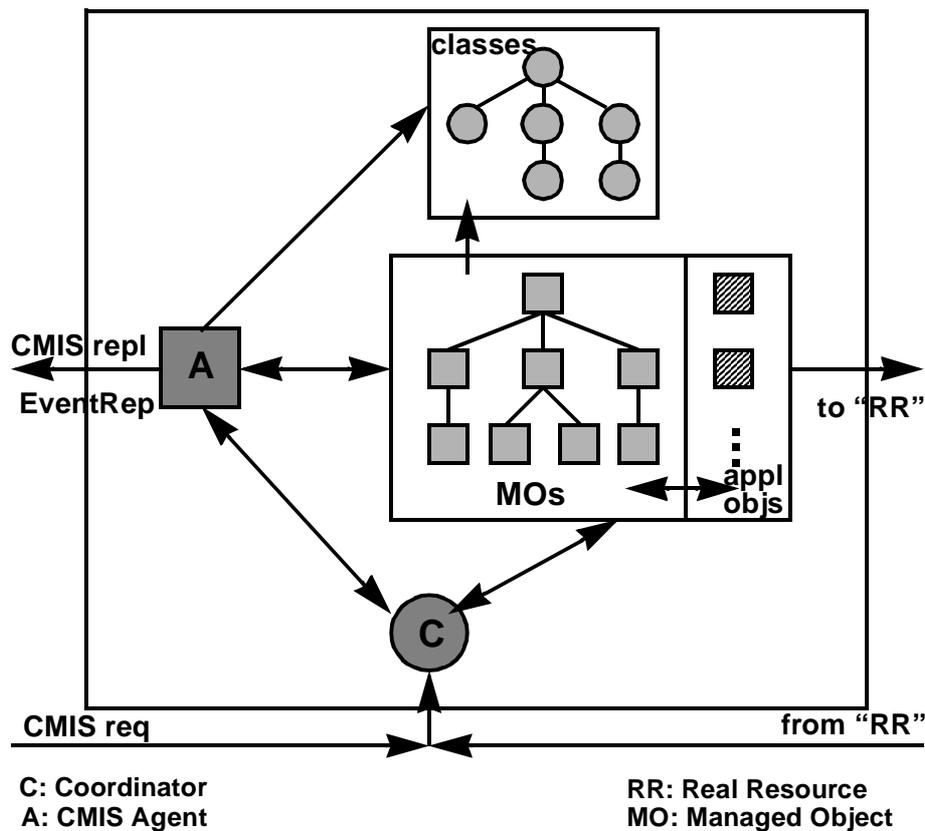


Figure 3 THE GMS Object-Oriented Architecture.

## 5.2 Managed Object Instances and Meta-Classes

Every specific managed object class needs access to information common to the class which is independent of all instances and common to all of them. This information concerns attributes, actions and notifications for the class, initial and default attribute values, “template” ASN.1 objects for manipulating action and notification values, integer tags associated to the object identifiers etc. This leads to the introduction of a common meta-class for all the managed object classes, the MOClassInfo. The inheritance tree is internally represented by instances of this class linked in a tree fashion as shown in the “classes” part of Figure 3.

Specific managed object classes are simply realised by equivalent C++ classes produced by the GDMO compiler and augmented manually with behaviour. Through access to meta-class information requests are first checked for correctness and authorisation before the behaviour code that interacts with the real resource is invoked. Behaviour is implemented through a set of polymorphic methods which may be redefined to model the associated real resource. Managed object instances are linked internally in a tree mirroring the containment relationships - see “MOs” part of Figure 3. Scoping becomes simply a tree search while special care is taken to make sure the tree reflects the state of the associated resources before scoping, filtering and other access operations. Filtering is provided through compare methods of the attributes which are simply the C++ syntax objects or derived classes when behaviour is coded at the attribute level.

## 5.3 Real Resource Access

There are three possible types of interaction between the managed object and the associated resource with respect to CMIS Get requests:

1. access upon external request
2. “cache-ahead” through periodic polling
3. update through asynchronous reports

The first one means that no activity is incurred when no manager accesses the agent but cannot support notifications. In the second requests are responded to quickly, especially with respect to loosely coupled resources, but timeliness of information may be slightly affected. Finally the third one is good but only if it can be tailored so that there is no unnecessary overhead when the agent is idle.

The GMS offers support for all methods through the coordination mechanism. When asynchronous reports from a resource are expected or asynchronous results to requests, it is likely that a separate object will be needed to demultiplex the incoming information and deliver it to the appropriate managed object. It should be noted here that an asynchronous interface to real resources driven by external CMIS requests is not currently supported as this requires an internal asynchronous interface between the agent and the managed objects. These objects are usually referred to as Internal Communications Controllers (ICCs) and are essentially specialised knowledge sources.

The internal organisation of a GMS-based agent in terms of the major interacting object instances is shown in Figure 5. Note that the instances shown are only the major ones defining the internal flow of control. The application’s intelligence may be realised through a whole lot of other objects which are application specific. Note also that the OSI stack is essentially encapsulated by the CMISAgent object.

## 5.4 Systems Management Functions

As already stated, OSIMIS supports the most important of the systems management functions. As far as the GMS is concerned, these functions are realised as special managed objects and generic attribute and notification types which can be simply instantiated or invoked. This is the case for example with the alarm reporting, metric and summarization objects. In other cases, the GMS knows the semantics of these classes and uses them accordingly e.g. in access control and event and log control. Notifications can be emitted through a special method call and all the subsequent notification processing is carried out by the GMS in a fashion transparent to application code. In the case of object management, the code generated by the GDMO compiler together with the GMS hide completely the emission of object creation and deletion notifications and the attribute change one when something is changed through CMIS.

Log control is realised simply through managed object persistency which is a general property of all OSIMIS managed objects. This is implemented using the GNU version of the UNIX DBM database management system and relies on object instance encoding using ASN.1 and the OSI Basic Encoding Rules to serialise the attribute values. Any object can be persistent so that its values are retained between different incarnations of an agent application. At start-up time, an agent looks for any logs or other persistent objects and simply arranges its management information tree accordingly.

## 5.5 Security

General standards in the area of security for OSI applications are only now being developed while the Objects and Attributes for Access Control Systems Management Function is not yet an International Standard. Nevertheless, systems based on OSI management have security needs and as such OSIMIS provides the following security services:

- peer entity authentication
- data origin authentication and stream integrity
- access control

These were developed in the ESPRIT MIDAS project to cater for the security of management of a large X.400 mail system [22] and will also be used in the RACE ICM project for inter-TMN security requirements on virtual private network applications. Peer entity authentication relies on public key encryption through RSA as in X.509. Data origin authentication is based on cryptographic checksums of CMIP PDUs calculated through the MD5 algorithm. Stream integrity is provided in a novel way that is based on a “well-known” invokeID sequence in ROSE PDUs. It should be noted that as CMIP does not make any provision for the carrying of integrity checksums, these are carried in the ROSE invokeID field. Finally access control is provided through an early implementation of the relevant SMF. Migration to the relevant full standards is envisaged when the latter will be completed.

## 6 GENERIC HIGH-LEVEL MANAGER SUPPORT

Programming manager applications using the CMIS API can be tedious. Higher object-oriented abstractions can be built on top of the CMIS services and such approaches were initially investigated in the RACE-I NEMESYS project while work in this area was taken much further in the

RACE-II ICM project.

The Remote MIB (RMIB) support service offers a higher level API which provides the abstraction of an association object. This handles association establishment and release, hides object identifiers through friendly names, hides ASN.1 manipulation using the high-level ASN.1 support, hides the complexity of CMIS distinguished names and filters through a string-based notation, assembles linked replies, provides a high level interface to event reporting which hides the manipulation of event discriminators and finally provides error handling at different levels. There is also a low level interface for applications that do not want this friendliness and the performance cost it entails but they still need the high-level mechanisms for event reporting and linked replies.

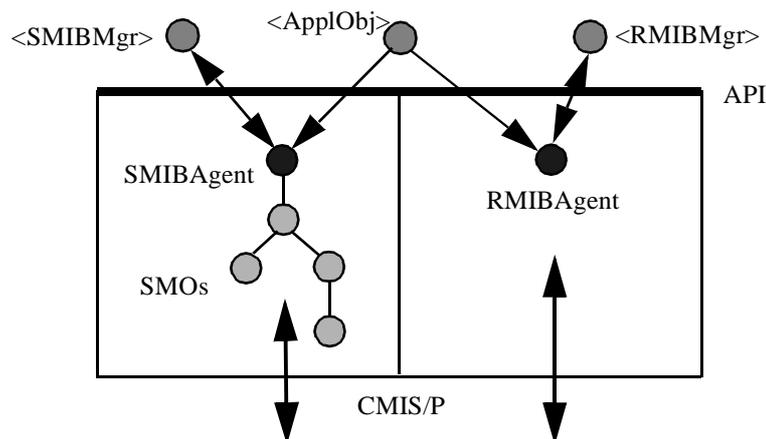
In the RMIB API there are two basic C++ classes involved: the RMIBAgent which is essentially the association object (a specialised KS in OSIMIS terms) and the RMIBManager abstract class which provides call-backs for asynchronous services offered by the RMIBAgent. While event reports are inherently asynchronous, manager to agent requests can be both: synchronous, in an RPC like fashion, or asynchronous. In the latter case linked replies could be all assembled first or passed to the specialised RMIBManager one by one. It should be noted that in the case of the synchronous API the whole application blocks until the results and/or errors are received while this is not the case with the asynchronous API. The introduction of threads or coroutines will obviate the use of the asynchronous API for reasons other than event reporting or a one-by-one delivery mechanism for linked replies.

While the RMIB infrastructure offers a much higher level facility than a raw CMIS API such as the OSIMIS MSAP one or X/Open's XOM/XMP, its nature is closely linked to that of CMIS apart from the fact that it hides the manipulation of event forwarding discriminators to effect event reporting. Though this facility is perfectly adequate for even complex managing applications as it offers the full CMIS power (scoping, filtering etc.), simpler higher-level approaches could be very useful for rapid prototyping.

One such facility is provided by the Shadow MIB (SMIB) support service, which offers the abstraction of objects in local address space, "shadowing" the real managed objects handled by remote agents. The real advantages of such an approach are twofold: first, the API could be less CMIS-like for accessing the local objects since parameters such as distinguished names, scoping etc. can be just replaced by pointers in local address space. Second, the existence of images of MOs as local shadow objects can be used to cache information and optimise access to the remote agents. The caching mechanism could be controlled by local application objects, tailoring it according to the nature of the application in hand in conjunction with shared management knowledge regarding the nature of the remote MIBs. Issues related to the nature of such an API are currently investigated in the ICM project []. The model and supporting C++ classes are very similar to the RMIB ones. The two models are illustrated in Figure 4.

Both the RMIB and SMIB support services are based on a compiled model while interpreted models are more suitable for quick prototyping, especially when similar mechanisms for Graphical User Interfaces are available. Such mechanisms currently exist e.g. the TCL/TK language/widget set or the SPOKE object-oriented environment and these are used in the RACE ICM project as technologies to support GUI construction.

Combining them to a CMIS-like interpreted scripting language can lead to a very versatile infrastructure for the rapid prototyping of applications with graphical user interfaces. Such languages are currently being investigated in the ICM and other projects.



**Figure 4** The Remote and Shadow MIB Manager Access Models.

## 7 DIRECTORY SUPPORT SERVICES AND DISTRIBUTION

Management applications need to address each other in a distributed environment. The OSI Directory Service (X.500) provides the means for storing information to make this possible. It essentially provides the a highly distributed hierarchical information store with high performance access characteristics, so it is suitable for storing information related to the location and capabilities of management applications.

The Directory Service model structures information in an object-oriented hierarchical fashion similar to that of OSI management. Information objects can be created, have their attributes accessed (read or written) and deleted. Access can involve complex assertions through filtering as with CMIP. This object-oriented information store can be highly distributed over physically separate entities known as Directory Service Agents (DSAs). These communicate with each other through a special protocol and requests for information a DSA does not hold can be “chained” to all the other DSAs until the information is found. This information can be accessed through Directory User Agents (DUAs) which talk to the local domain DSA through the Directory Access Protocol (DAP) while chaining guarantees the search of the global information store. This model is very powerful and closely resembles that of OSI management. From an information modelling perspective, the latter is a superset of the X.500 one and could be used to much the same effect. It is the chaining facility though that distinguishes the two and makes X.500 more suitable as a global information store. Current research tries to unify the two models, considering all the Management Information Trees of applications as extensions of the global Directory Information Tree.

Directory Services can be used for application addressing in two different styles: the first resolving Application Entity Titles (AETs) to Presentation Addresses (PSAPs) in a static fashion with the second introducing dynamic “location transparency” services as in distributed systems platforms. OSIMIS applications may equally well operate without the Directory Service: in this latter case, a flat file known as “isoentities” should exist at each site where management applications need addressing information. This maps statically AETs to PSAPs and should be administered in a fashion that guarantees global consistency, but this is exactly the purpose of services like X.500! In the first level of X.500 usage, the static information residing normally in the isoentities file is

converted into directory objects and stored in the directory (the ISODE QUIPU X.500 implementation provides a utility to do this off-line. This information becomes then globally accessible while central administration and consistency maintenance become fairly simple. This approach is adequate for fairly static environments where changes to the location of applications are infrequent. For more dynamic environments where distributed applications may often be moved for convenience, resilience, replication etc., a highly flexible solution is needed. This is provided in the form of location transparency services, wherever these are appropriate. It should be noted that these services may not be appropriate for the lowest management layer (Network Element), as the same application may exist at multiple sites.

Location transparency is implemented through special directory objects holding location, state and capability information of management applications. The latter register with it at start-up time and provide information of their location and capabilities while they de-register when they exit. Applications that wish to contact another one for which they know its logical name (AET), they contact the directory through a generic “broker” module they contain and may obtain one or more locations where this application runs. Further criteria e.g. location may be used to contact the right one. Another level of indirection can be used when it is not the name of an application known in advance but the name of a resource. A special directory information model has been devised that allows this mapping by following “pointers” i.e. Distinguished Names that provide this mapping. Complex assertions using the directory access filtering mechanism can be implemented to allow the specification of a set of criteria for the service or object sought.

## 8 APPLICATIONS

OSIMIS is a development environment; as such it encompasses libraries providing APIs that can be used to realise applications. Some of these are supported by stand-alone programs such as the ASN.1 and GDMO compilers. Generic management applications are also provided and there are two types of those: semantic-free manager ones that may operate on any MIB without changes and gateways to other management models. OSIMIS provides a set of generic managers, graphical or command-line based, which provide the full power of CMIS and a generic application gateway between CMIS/P and the Internet SNMP.

### 8.1 Generic Managers

There is a class of applications which are semantic-free and these are usually referred to as MIB browsers as they allow one to move around in a management information tree, retrieve and alter attribute values, perform actions and create and delete managed objects. OSIMIS provides a MIB browser with a Graphical User Interface based on the InterViews X-Windows C++ graphical object library. This allows to perform management operations and provides also a monitoring facility. It is going to be extended with the capability of receiving event reports and of monitoring objects through event reporting. This has recently been re-engineered in TCL/TK.

OSIMIS provides also a set of programs that operate from the command line and realise the full set of CMIS operations. These may be combined together in a “management shell”. There is also an event sink application that can be used to receive event reports according to specified criteria. Both the MIB browser and these command line programs owe their genericity to the generic CMIS facilities (empty local distinguished name { } for the top MIB object, the localClass facility

and scoping) and the manipulation of the ANY DEFINED BY ASN.1 syntax through the table driven approach described in section 3.

## 8.2 The Generic CMIS/SNMP Application Gateway

The current industry standard for network element management is the Internet SNMP, which is a simplified version of the OSI CMIP. The same holds for the relevant information models; the OSI is fully object-oriented while the SNMP supports a simple remote debugging paradigm. Generic application gateways between them are possible without any semantic loss for conversion from CMIS to SNMP as the latter's operations and information model are a subset of the OSI ones. Work for standards in this area has been driven by the Network Management Forum (NMF) while the RACE ICM project contributed actively to them and also built a generic application gateway based on OSIMIS [9].

This work involves a translator between Internet MIBs to equivalent GDMO ones and a special back-end for the GDMO compiler which will produce run-time support for the generic gateway. That way the handling of any current or future MIBs will be possible without the need to change a single line of code. It should be added that the generic gateway works with SNMP version 1 but it will be extended to cover SNMP version 2. The current approach for the gateway is stateless but the design is such that it allows the easy introduction of stateful optimisations in order to reduce SNMP managed systems access.

## 8.3 Other Potential Uses and Applications

Up to now it has been assumed that OSIMIS will be used as a management platform. Though this is the reason it was conceived, the power and generality of the OSI management model, on which OSIMIS is based, together with the abstractions OSIMIS provides and make it suitable for other potential usages outside the management systems realm.

The facilities OSIMIS provides make it suitable as a general vehicle for building distributed systems: the Action CMIS primitive is essentially a general method on an object which can have any parameters and return results or errors. This object is contained within an application, together possibly with other objects and its services may be accessed by objects in other applications in a distributed manner. The Directory Support Services enable one to find out which objects operate in a domain, what services they provide, which application contains them and finally use their services in a location transparent fashion. In combination to the high level object access mechanisms (RMIB / SMIB), it is possible to build distributed systems quickly and efficiently.

Realising a simple statistics server requires no more than 10 lines of C++ code in addition to its functionality, using the GMS and GDMO compiler. Accessing those objects in a particular environment using the location transparency and the RMIB support services requires about 20 lines of C++ code. Finally, the performance of OSIMIS-based systems is generally very high, a calculate square root action taking about 40 msecs on the local LAN using Sun SPARC Classics, with connection establishment taking about 150 msecs and connection release taking 20 msecs.

## 9 EPILOGUE

OSIMIS has proved the feasibility of OSI management and especially the suitability of its object-oriented concepts as the basis for higher-level abstractions which harness its power and hide its complexity. It has also shown that a management platform can be much more than a raw management protocol API together with sophisticated GUI support which is provided by most commercial offerings. In complex hierarchical management environments, object-oriented agent support similar to that of the GMS and the associated tools and functions is fundamental together with the ability to support the easy construction of proxy systems. Higher level manager support is also important to hide the complexity of CMIS services and allow the rapid but efficient systems realisation. OSIMIS has also shown that an RPC mechanism together with an Interface Definition Language (IDL) and a trader is not the only model to build distributed systems and, in fact, is not the best model for management systems. Broking services like the ones traders offer can be built using the OSI Directory (X.500) and Management (X.700) models while at the same time management applications conform fully to the ISO/CCITT management standards gaining from their rich event-based nature.

Finally, as we live in a multi-model and protocol world, it is unwise to assume existing investment will be simply thrown away to conform to a new model. In the management world there are two prevalent solutions, OSI and Internet, while a third one, the ODP-biased OMG CORBA framework threatens to become a de-facto industrial standard due to the fact it comes as a software platform rather than a pile of documents. All these three solutions will have to co-exist in the years to come and the basis for their integration through generic application gateways should be the most powerful of the three in order to ensure translation without loss of semantics. OSI is the most powerful one in terms of both protocol operations and information model expressiveness, justifying as such its choice as the basis of OSIMIS. Future work is envisaged in supporting gateways to CORBA systems and vice-versa, lightweight approaches to avoid the burden and size of OSI stacks through service relays, interpreted policy managed and management domains, sophisticated discovery facilities and a lot more.

### **Acknowledgements**

Many people have contributed to OSIMIS to be mentioned in this short space here. James Cowan of UCL though should be mentioned for the innovative design and implementation of the platform independent GDMO compiler, Thurain Tin, also of UCL, for the excellent RMIB infrastructure and Jim Reilly of VTT, Finland for the SNMP to GDMO information model translator that was produced over a week-end(!), the first version of the metric objects and general useful feedback. This work was carried out under the RACE ICM and NEMESYS and the ESPRIT MIDAS and PROOF European applied research projects.

## 10 REFERENCES

- [X701] ITU X.701, Information Technology - Open Systems Interconnection - Systems Management Overview, 7/91
- [M3010] ITU M.3010, Principles for a Telecommunications Management Network, Working Party IV, Report 28, 12/91
- [X710] ITU X.710, Information Technology - Open Systems Interconnection - Common Management Information Service Definition, Version 2, 7/91

- [XOpen] X/Open, OSI-Abstract-Data Manipulation and Management Protocols Specification, 1/92
- [Pav93a] Pavlou, G., S. Bhatti and G. Knight, OSIMIS User Manual Version 1.0 for System Version 3.0, 02/93
- [Pav93b] Pavlou G., The OSIMIS TMN Platform: Support for Multiple Technology Integrated Management Systems, Proceedings of the 1st RACE IS&N Conference, Paris, 11/93
- [Strou] Bjarne Stroustrup, The C++ Programming Language, Addison-Wesley, Reading, MA, 1986
- [X500] ITU X.500, Information Processing, Open Systems Interconnection - The Directory: Overview of Concepts, Models and Service, 1988
- [X720] ITU X.720, Information Technology - Structure of Management Information - Part 1: Management Information Model, 8/91
- [MUIB] Chapter 10, Experiment 3 Design, NEMESYS RACE Project 1005, ref. 05/DOW/SAR/DS/B/024/a1, 5/92
- [X730] CCITT Recommendation X.730 (ISO 10164-1) Information Technology - Open Systems Interconnection - Systems Management - Part 1: Object Management Function (for CCITT Applications), 10/91
- [X734] CCITT Recommendation X.734 (ISO 10164-5) Information Technology - Open Systems Interconnection - Systems Management - Part 5: Event Report Management Function, 8/91
- [X735] CCITT Recommendation X.735 (ISO 10164-6) Information Technology - Open Systems Interconnection - Systems Management - Part 6: Log Control Function, 6/91

## 11 BIOGRAPHIES

**George Pavlou** received his Diploma in Electrical, Mechanical and Production Engineering from the National Technical University of Athens in 1983 and his MSc in Computer Science from University College London in 1986. He has since worked in the Computer Science department at UCL mainly as a researcher but also as a lecturer. He is now a Senior Research Fellow and has been leading research efforts in the area of management for broadband networks, services and applications.

**Kevin McCarthy** received his B.Sc. in Mathematics and Computer Science from the University of Kent at Canterbury in 1986 and his M.Sc. in Data Communications, Networks and Distributed Systems from University College London in 1992. Since October 1992 he has been a member of the Research Staff in the Department of Computer Science, involved in research projects in the area of Directory Services and Broadband Network/Service Management.

**Saleem N. Bhatti** received his B.Eng.(Hons) in Electronic and Electrical Engineering in 1990 and his M.Sc. in Data Communication Networks and Distributed Systems in 1991, both from University College London. Since October 1991 he has been a member of the Research Staff in the Department of Computer Science, involved in various communications related projects. He has worked particularly on Network and Distributed Systems management.

**Graham Knight** graduated in Mathematics from the University of Southampton in 1969 and received his MSc in Computer Science from University College London in 1980. He has since worked in the Computer Science department at UCL as a researcher and teacher. He is now a Senior Lecturer and has led a number of research efforts in the department. These have been concerned mainly with two areas; network management and ISDN.

