



# **Report of the International Workshop on Distributed Systems: Operations & Management**

**Sponsored by**  
IFIP TG6 WG6 on Network Management  
with participation of IEEE Communication Society/CNOM

**Co-Hosted by**  
MetaAccess, Inc. and  
University of California at Santa Barbara  
October 15-16, 1991

**Editor**  
Shri K. Goyal  
GTE Laboratories Incorporated

## **Workshop General Format Organization and Summary**

This Second International Workshop on Distributed Systems, co-hosted MetaAccess, Inc. and the University of California at Santa Barbara, was enthusiastically attended by over seventy participants. The workshop consisted of six sessions distributed over almost two days. The breaks between sessions allowed sufficient time for technical interchanges and making new acquaintances. Each session was composed of three or four papers, or participation by panelists. The sessions were structured such as to allow time for integration of issues and results from the session participants and, of course, the audience at the end.

The Program Committee responsible for organization of the workshop included:

Divyakant Agrawal, University of California at Santa Barbara Salah Aidarous, Bell Northern Research Roberta Cohen, AT&T Bell Laboratories Deborah Estrin, University of Southern California Rodney Goodman, California Institute of Technology Shri Goyal, GTE Laboratories Christian Huitema, INRIA Naseem Khan, CODEX Corporation Kim Kappel, Georgia Institute of Technology Kris Krishnan, The MITRE Corporation Kenneth Lutz, Bell Communications Research Branislav Meandzija, MetaAccess, Inc. (Program Chair) Keith McCloghrie, Hughes LAN Systems Dan Stokesberry, National Institute of Standards and Technology Liba Svobodova, IBM Zürich Research Laboratory Wolfgang Zimmer, GMD-FIRST Berlin Douglas Zuckerman, AT&T Bell Laboratories

The local Organization Committee included:

Peter Allen, College of Engineering, UCSB, Graphics Mary Olson, Dept. of Computer Science, UCSB, Organization Co-Chair Marla Zimba, MetaAccess, Inc., Organization Co-Chair Through their tremendous coordinated efforts, the conference ran very smoothly. All the participants received handouts ahead of the presentation. There was plenty of good food and parking spaces. The first day of the workshop ended with a meeting of IFIP WG6.6, followed by a lively reception at the University of California's Faculty Club. The second day included a sailboat cruise with dinner in a bit cold weather to make it even more exciting.

The conference opened with introductions by Kimberly Kappel, the IFIP WG6.6 Co -Chair, Douglas Zuckerman, IEEE CNOM Chair, and Branislav Meandzija, Program Chair and co-Host. This was followed by a welcome address by Roger Wood, Associate Dean, College of Engineering, UCSB. The workshop was followed by a full-day Program Committee meeting for NOMS'92.

Overall, the workshop offered a great opportunity to share views on Distributed System Operations and Management with the leading world experts and opened doors for an enthusiastic environment for future co-operation.

The following sections summarize the presentations and discussions of workshop papers, prepared by respective session chairs.

### **SESSION 1 : Enterprise Management**

Session Chair: Kimberly Kappel, Georgia Institute of Technology

This session was concerned with issues relating to the management of enterprise networks. Two of the three speakers discussed architectures for distributed management of networks. The third speaker's topic was different in that it discussed the impact of network management policy.

#### **Paper 1: The Design of a Management Delegation Engine; Goldszmidt & Yemini, Columbia University**

German Goldszmidt of Columbia University presented the work that he has done in conjunction with Professor Yechiam Yemini. This presentation discussed the design and prototype implementation of a framework for network management, called the Manager Agent Delegation (MAD) model, which allows dynamic delegation of management function between manager and agents within a distributed network. Delegated management processes are defined and delegated to the MAD agents, which then become hierarchical managers by receiving and executing delegated programs that monitor and control managed objects.

One of the key concepts of this architecture is that manager must have the ability to dynamically delegate management authority if bottlenecks are to be avoided. The MAD agent is a hybrid combination of a hierarchical manager and a proxy agent, and, therefore, can act in either role as required. One of the motivations for this approach is to allow the management authority to be executed in close proximity to the device being managed. From a performance perspective, delegation lowers management costs and results in more effective and reliable management.

A prototype of MAD has been implemented in C using Sun OS4.1.

#### **Paper 2: A Methodology for Deriving a Network Management Specification from a Network Management Policy, Putter & Roos, University of Pretoria**

Phillip Putter of the University of Pretoria presented the work that he is doing with J.D. Roos on deriving a network management specification from the network management policy. They contend that the process of defining the policy for network management within a particular environment will result in the specification of required applications, definition of required managed objects, determination of the management domains, and understanding of the patterns of interaction. The methodology which they propose relies on the work of the International Standards Organization in the Open Distributed Processing (ODP) efforts.

A central contention within this presentation is that the general management style of an organization is determined by the management policy which the organization subscribes to. The policy includes such issues as objectives and goals of the organization. In order to determine policy, it is necessary to consider the network to be managed from several viewpoints. The viewpoints identified by the ODP efforts are considered by Putter and Roos, and are considered multiple abstractions of a single system.

For example, the specification of the enterprise viewpoint involves the user view of network management applications and results in a specification of the behavior of the network management applications and their characteristics. In order to illustrate the mapping of policy to system specification, Putter used the example of a policy which requires 99.9% network availability and an analysis of each of the ODP viewpoints (enterprise, information, computation, engineering and technology) in order to determine system requirements.

**Paper 3: An Architecture for Distributed, Cooperative, Integrated Network Control, Sethi & Lobo, University of Delaware**

Andrea Lobo of University of Delaware presented the work that she is doing with Professor Sethi on an architecture for distributed, cooperative, and integrated network control. They propose an architecture which separates the network layer control functions from the operational aspects of the protocols. The central component of the architecture is an integrated network control agent (INCA) which resides in each node of the network and is responsible for network control and management. The INCAs only perform reactive, operational control on their own initiative and must work together, cooperating to manage the subnetwork. This architecture meets the goals of scalability, separation of control and operation, and integration of network control functions.

Each INCA is composed of a monitoring agent, a communication agent, and a control agent. The function of the monitoring agent is to observe local variables and perform some reactive control. The communication agent provides services to the control agent and enforces policies that specify how INCAs can interact. It also acts as a database for the data collected by the monitoring agents. The control agent implements the control functions using the services of the other two agents.

Lobo and Sethi are currently designing INCA components for congestion control and connection management applications. They plan to apply their work to high speed networks.

The discussions which followed the three presentations were initiated by the generic issue of how the work being done by each of the presenters enabled enterprise management. The architectural discussions have obvious implications in enabling management of distributed networks, but seemed to be more focused on management of a local nature. There was discussion about the necessity of developing architectures which address the problem of managing enterprise networks. The work of Putter is intended to determine the structure of network management systems which meet the policy requirements of a given enterprise. Other topics that were brought up during the ensuing discussion include the difference between control, operations and management. It was agreed that these terms are frequently used equivalently and that there needs to be more concise and rigorous use of terms, particularly in groups that include people with differing backgrounds (telecommunications vs. data communications vs. computing). Waltham, MA

**SESSION 2: Data Modeling and Organization of Management Information**

Session Chair: Roberta S. Cohen, AT&T Bell Laboratories, USA

Four speakers addressed aspects of Data Modeling and Organization of Management Information in Distributed Systems.

**Paper 1: Information Base Modeling for Open Heterogeneous Network Management, Sibilla, Faure, Desprats & Valderruten, IRIT/SIERA, FRANCE**

Michelle Sibilla of IRIT/SIERA presented an object-oriented model for a network management architecture based on an open repository of ASN-1 type object class models. The design and implementation were performed through consistent application of the same set of principles which enabled the team to propose the automation of the entire process. The process uses Non First Normal Form - NF 2 by Check and Pistor and can be implemented on any Object-Oriented Data Base Management System.

**Paper 2: Converting MIB Descriptions to MIB Implementations Hegering, Abeck & Boehnke, University of Munich, GERMANY**

Thorsten Boehnke of the University of Munich presented an implementation analysis of Management Information Bases. He pointed out that Managed Objects can not be seen as atomic units as they are composed of fragments which are generally distributed over distinct parts of the real resource. The criteria for building fragments are influenced by the abilities and features of the real resource and by the requirements of the network management application. This knowledge should be described in implementation-relevant amendments to the Managed Object description. Thorsten outlined a possible approach using the mapping types example.

**Paper 3: Directory Information Base: A Proposal for Manageable Configuration Management in a Distributed System Strich & Okulski, SSDS, Inc., USA**

Michael Okulski of SSDS Inc. reported that the X.500 Directory System supports a globally distributed low frequency write, high frequency read that may be used as a repository of configuration management information for decentralized processing systems. The general approach from a configuration management perspective is that the Directory System may contain all the information needed to produce 'system maps' or "wiring diagrams" from various perspectives. This information is defined by the application (or user) that uses the data, and therefore may act as an extension to data that may already exist in the system. Additional management systems are then needed to overlay these maps with current data regarding fault conditions, operational status, system usage, and performance. This approach allows system resources that are not being directly managed (e.g. phones, people, applications, etc.) to be accounted for regarding the management of system configurations.

**Paper 4: Object-Oriented Distributed Technology in Network Management D. Morse, MPR Teltech, Ltd., Canada**

Daryl Morse of MPR Teltech, Ltd. introduced a new way of thinking which will be required if the benefits of OSI Management standards are to be realized. While OSI Management standards offer the potential for greatly increasing the scope of the traditional domain of network management systems, the implementation of such systems is currently neither practical, nor economical. Open Distributed Processing (ODP) standards promise to address many of today's problems, but as ODP will not be available for a number of years, a short-term solution will be necessary. Such a solution can be achieved by applying new software architecture and implementation technology to the network management domain. Specifically, if OSI Management standards and object-oriented distributed systems technology are merged, a flexible, extensible framework for the implementation of network management systems can be produced.

**SESSION 3: Distributed Systems Implementation Issues**

Session Chair: Morris Sloman, Imperial College

Four speakers addressed the implementation issues.

**Paper 1: Application and Network Management in Open Distributed Processing Environments A Schill University of Karlsruhe Martina Zitterbart, IBM TJ Watson Research Center**

This paper discussed a distributed object oriented language approach to managing Open Distributed Processing (ODP) applications and the transport layer extensions needed to support different application requirements.

The following extensions to the ODP viewpoint model are suggested:

**Computation:** a declarative language, with support for hierarchical composition to define the initial application structure in terms of instance creation and object interconnections. Changes are also defined declaratively using a similar language. Abstract behavioral object interaction specification is also needed.

**Information:** object interaction specification must be augmented with detailed descriptions of transmitted data types.

**Engineering:** This needs fine grained object placement with automatic mobility control with migration decisions supported by goal oriented policies. Measuring and analysis tools are needed to support this.

**Technology:** a declarative notation similar to that for the computation viewpoint is needed to describe hardware and network structures. A run-time system must be provided to support dynamic reconfiguration.

**Enterprise:** a requirements description language to define business requirements from an end user's view.

Transport Protocols

The information exchange defined in the information viewpoint determines the transport service characteristics. The paper surveys new transport classes required for high performance applications and classifies service characteristics in terms of isochronous/non -isochronous, one or two way interactions, real-time, throughput and error rates. The protocol components needed to configure the transport class are also defined.

## **Paper 2: OSF Distributed Management Environment (DME) Matthias Autrata - OSF Munich**

This paper briefly summarized the DME environment which follows the trend of being object oriented. The key requirements to be supported by the DME include:

- 1) Consistency based on a common graphical user interface and consistency of syntax and semantics of the Application Programming interface for management.
- 2) Interoperability through use of CMIP and SNMP protocols. Interoperability with proprietary management systems will be via gateways.
- 3) Scalability for managing large networks through modular architecture, optional components, delegation of management activities to agents. They identify 3 levels of management units:

Single nodes managed as a federation of individual systems  
Cells - groups of systems administered as one domain. Management delegated to an agent within a cell

- 4) Fully distributed DME functionality
- 5) Remote Procedure call for communication and management operations.
- 6) Ability to support Policies and Roles
- 7) Security for management operations

The DME architecture provides the services needed for management:

Management User Interface Services

Applications and application Services such as license management, software management, printing services & host management

Management services for domains and policies.

Object services include: registering and locating objects, management request brokers for access both SNMP and CMIP, event management, object servers for executing management operations. The DME development kit provides facilities for compiling object definitions, event templates, dialog scripting language and associated user interface toolkit.

## **Paper 3: Managing Replicated Data Efficiently in Distributed Systems Diwa Agrawal & A. El Abadi of UCSB.**

This paper reviewed several approaches to partially mitigate the costs associated with replicated data management. These were variations of the basic quorum protocol which requires a read quorum for reading and a write quorum for writing the replicas.

- i) Views protocol optimizes read access by permitting read to a single copy
- ii) Copies of data can be organized into a logical structure such as a tree and which permits read-one write-all operations in case of no failure.
- iii) A non-blocking quorum protocol which uses an underlying propagation mechanism to update replicas and reduce latency.

**Paper 4: Archetype Distribution Management Environment, Branislav Meandzija, MetaAccess, Santa Barbara**

This presentation described the Archetype Distribution Computer Aided Software Engineering environment based on the Common LISP Object System that integrates the design, specification, implementation and operations and management of distributed software.

Distribution paradigms such as client-server, requester-responder, and distribution mechanisms such as remote procedure calls, are all part of the Archetype vocabulary. The Archetype Distribution environment is an Xwindows-based environment that combines the methodology with the ISO Open Systems Interconnection model and enables the automatic creation of distributed systems and distributed system components that conform to OSI standards (including managed objects). Archetype Distribution is based on an Xwindows environment. Software generated using these tools can be linked to graphical images to allow run time monitoring of protocol drivers and analysis of protocols.

The archetype method defines 2 different specification languages:

- i) An abstract natural language for rapid object oriented design of networking and distributed system software
- ii) An execution language which is data driven and concurrent for implementing protocol engines.

The Archetype Distribution provides the user with a variety of capabilities including the graphic design of managed objects (including OSI objects) and the dynamic linking of the dynamically generated custom management GUI to automatically generated communications programs and distributed applications based on TCP/IP and/or OSI.

Archetype Distribution also includes a unique monitoring capability for network operations and management. That capability provides an animated display of communications processes through run-time display of communications process internals exhibiting logical and physical protocol relationships. Managers can view network "hotspots" at run-time through color or audio options; view the code processed at specific nodes at run-time through recorder or step-through monitor functions; and edit code specific to individual modules or nodes during run-time.

**SESSION 4: Distributed Operations Architectures for Telecommunications**

Session Chair: Douglas N. Zuckerman, AT&T Bell Laboratories, USA

The four speakers addressed aspects of distributed operations architectures for telecommunications.

**Paper 1: An Operations Architecture for the Future P. Boissonneault & G. Sharp, Bell Canada, S. Aidarous, Bell Northern Research, CANADA**

Glen Sharp of Bell Canada addressed "An Operations Architecture for the Future". Traditionally, the bulk of operations, administration and maintenance (OAM) processing has been in the operating telecommunications company (OTC) domain. It has been handled mostly through operations support systems (OSSs) in an OSS client-NE server method of operation, primarily because network elements (NEs) were restricted by their inability to perform advanced OAM processing functions. In addition, interfaces between the operations network and the telecommunications network were poor. This has resulted in a complex operations environment characterized by a large number of technology and vendor dependent task-oriented OSSs which are costly to maintain, have their own islands of computing, lack connectivity and have redundant data.

Recently, OTCs and NE vendors have become motivated to explore opportunities for integrating OAM capabilities across the operations and the telecommunications networks and to explore opportunities for real-time operations and customer self-service. This shift of focus from OSS-oriented solutions to an integrated OSS/NE strategy was triggered by the introduction of operations controllers into the architecture of next-generation intelligent network elements (INEs). The opportunities are driven by several factors, including:

- Current OTC initiatives to reduce operations expense, to expedite new service deployment, and to generate more revenues.
- Next-generation NE technologies, which allow flexible deployment options and increased interworking with OSSs.
- Increasing customer requirements for improved response times, enhanced services and more control on their services.

Opportunities for an integrated operations/technology strategy must take into account the different OAM requirements of those involved with the telecommunications networks. Glen proposed an OAM architecture framework with an integrated view encompassing both OSSs and INEs (whether owned by an OTC or a private network operator), with appropriate apportioning of functions across OSSs and INEs.

**Paper 2: Operations Model and Architecture for Evolving Communications Networks A. R. Johnston and K. Ludwiczak, AT&T Bell Laboratories, USA**

Alan Johnston of AT&T Bell Laboratories addressed an "Operations Model and Architecture for Evolving Telecommunications Networks". Recent work on the evolution of telecommunications networks has led to new, highly-integrated network architectures. Integration is understood at several levels including physical integration of transport and switching (NEs) for cost containment, functional integration within NEs for improved performance, and operational integration both among NEs and between NEs and Operations Systems (OSS) for network simplification. For example, a service node could generally serve as the network's central point of control and service logic within the larger Telecommunications Management Network (TMN), and one access node could focus on the residential and smaller business markets and another access node could focus on the larger business market.

Broadly viewed, Operations, Administration, Maintenance and Provisioning (OAM&P) functions will increasingly move into NEs making them self-sufficient. External OSs will be left with the responsibility to provide greater consolidation of information and improved access management capabilities to, for example, deliver the right information to the right person at the right time. For telecommunications networks, such a distributed management and control architecture is a new phenomena. Alan's presentation examined this environment by modeling the TMN using task-based operations processes to understand the economic basis for distributing functions and data between NEs and OSs. Next he considered specific management capabilities in the NEs and showed that the resulting network can simplify and streamline the management processes. And finally, he briefly considered how the TMN can be used to deliver services faster and more efficiently.

**Paper 3: Operations Challenges for Large Voice/Data Hybrid Networks J. D. Igleheart and R. N. S. Rathore, Bell Communications Research, USA**

Ram Rathore of Bell Communications Research addressed "Operations Challenges for Large Voice/Data Hybrid Networks with Distributed Intelligence". The Local Exchange Carrier (LEC) networks are evolving to meet the needs of competitive business enterprise networks optimized for point-to-point voice and data transmission. The key trends are:

- Deployment of Digital Cross-connect Systems (DCSs) at fiber facility nodes with subnetwork management capabilities for circuits ranging from sub-DS0 to DS3 bandwidth. The subnetwork management capabilities include remote provisioning and reconfiguration, remote alarm monitoring, and self-healing maintenance features.
- Increasing deployment of Intelligent Channel Banks (ICBs) with remote circuit provisioning features. The ICBs typically support circuits from sub-DS0 to DS1 level.
- Deployment of Transport Resource Managers (TRMs) with multiplexing, switching, circuit termination, and subnetwork management capabilities optimized for point-to-point voice and data transmission. These TRM capabilities bundled together represent a cost effective means for providing customer network management and control desired by the business customers of the LECs. The TRMs

offer dynamic bandwidth allocation for sub-DS0 to DS1 level with potential future enhancement to DS3 and SONET (Synchronous Optical NETwork) rates.

Ram focused on the operations issues and challenges for TRM networks in this evolving hybrid network of the LECs. His view was that the economic necessities of the transition period dictate:

- Harnessing the intelligence of individual network elements and subnetwork management systems for service cost reduction and revenue-producing new service features.
- Developing interworking criteria for both signal compatibility and network management messages and commands.
- Potential evolution of a “manager of subnetwork managers” for streamlining service activation and assurance across the large hybrid networks.

While the problem domain is too complex for developing a single solution, Ram discussed some of the directions and heuristics for moving forward in this challenging arena.

#### **Paper 4: Aspects of Distributed Management in the Transport Network F. Kaplan and T. Sasada, NEC America**

Ted Sasada of NEC America addressed “Control Aspects of Distributed Network Management”. Operations systems have traditionally managed transport equipment through direct links to the equipment. To accommodate communication over noisy, low speed data lines, primitive protocols were used. The network management model was represented by a central control element directly connected to each piece of equipment. With improvements in the data networks, faster, more sophisticated protocols were deployed to provide more precise network management. The model was expanded to replace the point to point management links with a telecommunications management network (TMN). Although more precise control capability exists using the TMN, the intelligence built into the network equipment plays only a minor role in the network management. Allowing intelligent network elements to share in the overall management of the network presents the operations systems designers with a series of technical issues that must be addressed:

- How do we provide sufficient bandwidth to accommodate the control protocols?
- How much processing capability can we assume at each control element?
- How do we insure that delays do not exceed established limits?

Assuming the underlying network transport facility conforms to the evolving SONET standard, Ted stated that the embedded operations channels defined by SONET provide potential solutions for the stated issues. In particular, the Data Communication Channels (DCC) allow the implementation of network management protocols to pass information between the SONET intelligent network elements. Controlling the distributed network system requires choosing from among the potential solutions an optimal solution. Ted concluded by describing the results obtained through an experimental architecture for such control.

In the ensuing discussion, some additional issues were raised, specifically:

- Though three-layer architectures seem better suited than two layer architectures in meeting enterprise-driven needs for large networks, most current implementations do not contain the “mid-layer manager” which would be below the “manager of managers”. A suggested reason for this was the newness of the approach.
- All management layers are interdependent; for example, the upper layers need to account for the rapidly increasing intelligence in network elements and customer premises equipment (CPE). Traditional OSS functions are migrating to NEs, and vice versa.
- As an expanding set of complex services are provided by telecommunications networks, standard approaches are needed to help in their modeling, for example Object Oriented.



## **Session 5: Quality of Service Requirements and Architecture Panel**

Panel Chair: Deborah Estrin, University of California

Five members participated in the panel. Following the position presentation, there was an enthusiastic discussion with audience participation.

### **Participant 1: Management and Control of Resources in Broadband Networks with Quality of Service Guarantees, Aurel Lazar and Giovanni Pacifici, Columbia University**

Prof. Lazar addressed the problem of providing service guarantees from the perspective of scheduling and admission control and analyzed the benefits and difficulties associated with distributing resource allocation algorithms. He introduced the concept of an "admissible load region", which defines the supportable mixes of his three traffic classes. He concluded with a brief discussion of his Integrated Reference Model for broadband networks.

### **Participant 2: End System Considerations for Quality of Service, David Feldmeier, Bellcore**

Whereas Lazar concentrated on the switch and network functions, Feldmeier addressed end system, in particular transport-level, issues encountered when providing service guarantees. He gave examples of virtual stream multiplexing and sequential message segmentation/reassembly as inappropriate design decisions that could make it difficult to meet service guarantees.

### **Participant 3: BERKOM Project: A case study of advanced quality of service management, Wolfgang Zimmer, GMD FIRST**

The FINE (First Innovative Network Environment) service manager developed as part of the BERKOM project provides both connection oriented and connectionless service in a broadband environment. Both types of service provide fixed or variable size service data units, with an optional data-part descriptor to facilitate efficient implementations. The connectionless service includes a global time to live value in each data unit. The connection-oriented service allows many parameters to be negotiated at setup time: service data units, flow control, use of TTL and selectable error handling. The flow control options are free flow with loss probability, rate flow control, and rate flow based on credit mechanism. A 4-level priority scheme may also be used.

### **Participant 4: Internetwork Perspectives on QOS, Deborah Estrin, University of Southern California**

Estrin addresses QOS as an internet issue because the QOS experienced by the end application is a result of behavior all along the data path, which may cross several individual networks. These networks may differ in link bandwidth, resource management policies and mechanisms, routing, policy/access control restrictions, charging, etc. Depending upon the max-stream-data rate to link-BW ratio, some networks along a path may need to manage and account for individual flows, and others may manage traffic classes. At the internet level the resource management and routing mechanisms should not impose or prevent either approach in order to accommodate a wide range of traffic across a wide range of networks.

### **Participant 5: Managing QOS in Distributed Networks: A User Orientation, Roberta Cohen, AT&T Bell Laboratories.**

Cohen presented QOS from a customer perspective. The critical parameters are: responsiveness, transmission quality, availability, and reliability. Although each aspect of user orientation requires some amount of end-to-end measurement, user orientation is inevitably a local view of the network's performance. Service Level Agreements (SLAs) are used by network providers and users to define performance and quality expectations and relate those expectations to observable network characteristics. In this context, requests for QOS parameter guarantees, negotiation and manipulation of QOS parameters must represent the logical, end-to-end flow of data and must represent SLA-relevant measures for speed, accuracy, availability, and reliability.

## **SESSION 6: Management of Security and Security of Management**

Session Chair: Kenneth J. Lutz, Bell Communications Research, Inc.

Network security is growing rapidly, both in its importance and in its scope. Many aspects of computing and telecommunications networks need to be managed, such as capacity and performance, and security management needs to be added to this list. Besides managing security, however, another area of crucial importance is the security of management, particularly as more end users are allowed to exercise some form of control of the network, including network management. This session addressed some specialized aspects of both management of security and security of management.

Special thanks go to Liba Svobodova of the IBM Zurich Research Laboratory for organizing this session.

**Paper 1: Decentralized Management of Security in Distributed Systems Ravi Sandhu, George Mason University, USA**

In the first presentation, Sandhu stated that security and trust are fundamental issues that must be considered in the management and operation of distributed systems. Without rigid centralized control, distributed systems must give each component system the ability to specify the security policies for the resources it controls. It is important, therefore, to be able to analyze the composition and interaction of security policies globally. Safety analysis is required to determine whether or not the composition of multiple autonomous security policies gives an acceptable system-wide policy. Safety analysis was shown to be undecidable under very weak assumptions by Harrison, Russo, and Ullman in their model, commonly known as HRU. Sandhu stated that efficient safety analysis results have been recently obtained by introducing strong typing into the HRU model [1]. More research still needs to be done on this issue.

**Paper 2: Recent Developments in SNMP Security Keith McCloghrie, Hughes LAN Systems, Inc., USA**

The second presentation, by McCloghrie, began with a review of SNMP security [2], then discussed four recent developments. These recent developments do not change the SNMP Protocol Data Unit(PDU), just the SNMP message "wrapper," which identifies both the source and destination as SNMP parties. (1) The digest algorithm used for authentication can continue to be MD4 unless the frequency of key changes becomes too often. Then it is likely that MD5, which is more costly, will have to be used. (2) Access control needs to be able to be configured with a granularity at the instance level, which was cumbersome as specified because each instance of each attribute needed to be configured. A change has been to add a "mask" value so that one access control entry can apply to the same instance of multiple attributes. (3) Several minor changes to the ASN.1 encodings of the new SNMP message format have been proposed; the changes that allow for more optimized coding are likely to be accepted. (4) Since each agent needs to know about a number of SNMP parties, six initial Object Identifiers have been proposed. These are sufficient to provide secure SNMP communication which can then be used to configure more parties, as required.

**Paper 3: Access Control through a Domain Service - Specification and Implementation, Jonathan D. Moffett, Morris S. Sloman, and Kevin Twidle, Imperial College, UK**

The third presentation, by Twidle, described how the domain concept [3] can be used to introduce access control and to structure and manage object names within a distributed system. The domain concept copes with large systems containing millions of objects and thousands of users by grouping objects hierarchically into subdomains such that management policies can be defined for the domains. Access control policies are specified for the domains, and the member objects inherit those policies. Twidle reported that a prototype has been implemented with a centralized domain name server and node managers on every physical node. There are limited access rules that use access control lists in each domain. Work is continuing on a fully functional system and on tools for specification, evaluation, and analysis.

During the discussion, the subject of theoretical foundations of security was raised. There was agreement that, as the foundation is built, more is being learned and that SNMP provides a good beginning. Another issue discussed was at what level should security be enforced. In SNMP, security is enforced at the application level; but in OSI, it is at all levels. The consensus was that it is needed at the appropriate level relative to the object being secured. A third issue was the lost-key problem, which will become more acute

as distributed systems proliferate. SNMP requires that keys first be distributed manually, then changed using SNMP. If a key is lost, manual redistribution is required.