DIS AND THE TRANSITION TO OSI: A COEXISTENCE OF STANDARDS

Margaret L. Loper, David T. Shen
University of Central Florida
Institute for Simulation and Training

M. Bassiouni
Department of Computer Science
University of Central Florida

ABSTRACT

Distributed Interactive Simulation (DIS) is an emerging simulation system requiring state-of-the art communication services and protocols. The communication services identified by the DIS Communication Architecture and Security Subgroup (CASS) fall into two classes, interim and long-term. DIS has two choices for Commercial-Off-The-Shelf communication protocols, Internet or Open System Interconnection (OSI) standards. Both protocol suites meet interim requirements; however, neither protocol suite can meet the full range of multicast requirements. Work has begun in national standards organizations to develop these protocols, and this work is based on the OSI/GOSIP architecture. The DIS interim architecture will be based on Internet standards but, if DIS is to comply with the GOSIP mandate, a strategy for the transition to OSI must be devised. This paper develops a transition plan which addresses GOSIP compliance and a strategy for meeting long-range multicast requirements.

ABOUT THE AUTHORS

Margaret Loper received her masters degree in Computer Engineering from the University of Central Florida in 1991. She is currently a Systems Engineer for the Institute for Simulation and Training, Orlando. Ms. Loper has been actively involved in research on simulation networking, OSI protocols, and multicast communications. She is an active member of the DIS Communication Architecture and Security Subgroup and is also a member of ANSI working groups on OSI Upper Layers and Network & Transport Layers.

M. A. Bassiouni received his Ph.D. degree in Computer Science from Pennsylvania State University in 1982. He is currently an Associate Professor of Computer Science at the University of Central Florida, Orlando. He has been actively involved in research on computer networks, distributed systems, concurrency control, and relational databases. Dr. Bassiouni is a member of IEEE and the IEEE Computer and Communications Societies, the Association for Computing Machinery, and the American Society for Information Science.

David Shen received his masters degree in Computer Engineering from the University of Central Florida in 1991. He is currently an Associate Engineer at the Institute for Simulation and Training. For the past two years, he has been involved in research related to computer communication protocols applied to DIS.

DIS AND THE TRANSITION TO OSI: A COEXISTENCE OF STANDARDS

Margaret L. Loper, David T. Shen
University of Central Florida
Institute for Simulation and Training

M. Bassiouni
Department of Computer Science
University of Central Florida

1. INTRODUCTION

Distributed Interactive Simulation (DIS) is an emerging simulation system requiring state-of-the-art communication services and protocols. To reduce cost and facilitate the interoperability of dissimilar simulations, industry communication standards will be adopted to maximize the use of Commercial-Off-The-Shelf (COTS) products and to maximize the base of practical technical knowledge. While the majority of communication services required by DIS can be satisfied with current COTS protocols, there are requirements which cannot be satisfied by these protocols. Consequently, the DIS service requirements fall into two categories: interim services, which are required to support immediate DIS experiments, demonstrations, and tests, and the customization or long-range services, which require development. The Communication Architecture and Security Subgroup (CASS) of the DIS workshops is recommending a phased approach for the evolution of the architecture.

The DIS communication architecture is composed of a suite of protocols which satisfy the established service requirements defined in [1]. For example, the bulk transfer service requirement will be satisfied by a file transfer protocol, for which DIS can choose either Internet or Open Systems Interconnection (OSI) COTS products. Both the Internet and OSI standards are composed of a large number of protocols, not all of which are required by DIS. For the remainder of this paper, the term "base stack" will be used to designate the subset of Internet or OSI protocols required by DIS for operation. The base stack includes only those protocols required to satisfy the service requirements (See Table 1).

Table 1 Summary of DIS Service Requirements and Base Stack Protocols

Required Services	Base Stack
• Network Management	• Network Management & Terminal
• Bulk Transfer	• File Transfer
• Reliable Unicast	• Connection-Oriented Transport
• Unreliable Unicast	• Connectionless Transport
• Unreliable Multicast	• Network Layer Multicast
• Reliable Multicast	• Transport Layer Multicast
• Seamless Local/Global Community	• Connectionless Network
• Security	• Multi-Level Security
• Synchronization	• Time

As mentioned previously, not all requirements can be met with COTS products. One of these services is multicast. Consequently, multicast will have to be phased into the architecture over a period of years as services and protocols are adopted by the standards bodies. As a result, the base stack for the interim architecture will differ from the base stack for the long-range architecture.

The evolution of the DIS communication architecture will need to occur in phases to allow sufficient time for users and implementers to gain experience with current requirements before new services are introduced. This strategy will be accomplished with a Transition Plan which synchronizes interim and long-range services with the phases of the architecture. The Transition Plan should not only identify the services and base stacks for each phase, but should also postulate the time frame in which the transitions should occur.

In order to understand the Transition Plan, we must first lay the groundwork. We have already introduced the service requirements and identified the communication protocols which satisfy each requirement. Next, we will examine COTS protocols, and, based on product availability, determine the long-range requirements. All this information will be coalesced into a Transition Plan which will be based on the starting point recommended by CASS. This Transition Plan will define a strategy to evolve the communication architecture from interim to long-range services which emphasizes the coexistence of standards.

2. THE STANDARDS

At the January 1992 meeting of the CASS, proposals for the DIS interim architecture were discussed. Three proposals were submitted: one based on Internet standards, one based on OSI standards, and one based on the Navy's Survivable Adaptable Fiber Optic Embedded Network (SAFENET) architecture. Because the SAFENET architecture is composed of predominantly OSI protocols, the discussion focused on Internet and OSI standards. From a technical viewpoint, both the Internet and OSI standards met the service requirements; however, questions were raised about the maturity of OSI.

Maturity can be defined by the following characteristics: Protocol Maturity, Product Availability, Product Maturity, Cost, Implementations, and Risk. Protocol maturity is measured by the number of years since the protocol was published. As a protocol becomes stable, less revisions are necessary to the standard, and the protocol and its associated products become mature. The number of products available from vendors is the measure of protocol availability. Product maturity must be based on protocol maturity and is consequently affected by revisions to the standards. The cost of products is linked to protocol availability. The fewer the products on the market, the more expensive they are. The number of implementations and risk are also related. As the number of implementations grows, risk is reduced. This is a result of gaining experience with the protocols and lessons learned on product interoperability.

2.1 Internet Protocol Suite

The Internet Protocol Suite (IPS) is a family of protocols based on the Transmission Control Protocol/Internet Protocol (TCP/IP) standards. The IPS base stack was published in the mid 1970's. Products are available from most vendors for both workstation and PC platforms. The cost of the IPS base stack varies and is usually included in the cost of the hardware. The IPS is the defacto standard for computer networking and boasts numerous implementations, most notably the global Internet. There is no risk associated with products and only minimal risk associated with in-house implementations. Due to the large installed base of IPS and the twenty years of development, these protocols and their corresponding products are very mature.

One interesting note on the maturity of this protocol suite:

It is so mature it suffers from old age.

The Network Layer protocol IP is running out of address space, which will be depleted within the next two years [5]. Several proposals are on the table to solve the problem. Most proposals are merely "patches" which will extend the life of IP; they are not long term solutions to the real problem. The most widely accepted proposal is to use the OSI Network layer addressing scheme and shift to the OSI Connectionless Network Protocol (CLNP). With the exception of the addressing structure, IP and CLNP are practically the same protocol. They provide the same functionality and are virtually indistinguishable.

2.2 Open Systems Interconnection

The other option for protocol interoperability is to comply with the Government Open Systems Interconnection Profile (GOSIP) mandate which has been in effect since August 1990. GOSIP is the U.S. Government program for adoption of OSI standards across all Federal agencies. DIS will benefit from the OSI/GOSIP architecture through: reduced cost, increased interoperability (both nationally and internationally), and increased application-level functionality [2]. The Institute for Simulation and Training (IST) has developed the DIS protocol standard with the goal of using the GOSIP protocols. Unfortunately, GOSIP has not reached the level of maturity of the Internet protocol suite and consequently, many view GOSIP compliance as a long-term goal.

The OSI base stack was published in the mid 1980's. Based purely on the number of years, the IPS base stack is more mature; however, many of the OSI protocols are based on their Internet predecessors and therefore gain stability from lessons learned. Current documentation shows approximately 450 OSI products from 80 suppliers. From a recent survey of major computer vendors, users have a choice of an IPS or OSI stack for workstations. Several OSI products, such as Network Management, are still in development by vendors. Product maturity is hard to measure, but due to the limited installed base of OSI products, maturity is not near the level of the IPS. The cost of OSI products is higher than that of the IPS for several reasons. First, the development of the IPS was funded in large part by federal agencies through research grants. Therefore, vendors did not have to spend their own money to mature the protocols and products. In contrast, OSI is being developed by industry. Consequently, the capital expended in the development of both the protocols and products is passed on to the customer. From a limited survey, OSI base stacks range from zero to several thousand dollars. However, one major computer vendor ships all computer systems with dual (Internet and OSI) stacks.

Although OSI cannot boast implementations as numerous as IPS, OSI is slowly growing and is even being integrated into the global Internet. The National Science Foundation Network (NSFnet) backbone has offered national CLNP[1] service since August 1990. There are approximately 25 regional networks which are part of this "OSI over the Internet" testbed, including: Energy Sciences Network (ESnet), NASA Network (NASAnet), Southeastern Universities Research Association Network (SURAnet), and New England Academic and Research Network (NEARnet). These regional networks route both Internet and OSI traffic. There is also a world X.400 (OSI electronic mail) backbone connecting the U.S., Europe, and Pacific Rim. In addition, several new major government procurements are specifying OSI/GOSIP communication services. These procurements include the Department of the Treasury, the Department of Energy, and the Department of Agriculture.

There is risk associated with OSI products due to limited experience with the protocols. However, the integration of OSI into the Internet will help reduce the risk by exposing industry, academia and government to the protocols and products. In fact, the implementation of OSI routing protocols running in NSFnet is now available free to the public. As OSI implementations are tested and made available to the public, the risk associated with OSI will diminish.

3. THE REAL PROBLEM

As discussed previously, the DIS service requirements fall into two categories: interim services, which are required to support immediate DIS experiments and long-range services, which require development.

1 CLNP is analogous to the Internet IP protocol.

The interim services are those which can be satisfied with current COTS products. These includes network management, bulk transfer, security[2], reliable unicast, unreliable unicast, unreliable broadcast, seamless local/global communication, and synchronization. In fact, these are all basic communication services used by many types of applications.

The real problem with any application is the development or customization needed to meet unique requirements. For DIS, one such customization is multicast. Large scale DIS exercises will require multicast to selectively transmit information among simulators. Multicast will reduce the amount of PDU traffic a simulator must process by sorting out information which is of no interest to it. To allow hundreds of thousands of entities to simultaneously participate in simulation exercises [6], multicast will be required. DIS desires a full range of multicast capabilities which cannot be satisfied with any current COTS protocol.

Today there are only three possibilities for multicast: IP Multicast (IPMC), Stream-II (ST-II), and the Xpress Transfer Protocol (XTP). IP multicast, part of the IPS family, is the only commercial multicast product but has limited availability. Further, IPMC in its current form does not meet the requirements of DIS: *It is not a real-time protocol.* For this reason, a modified IPMC has been proposed. This proposal would require significant development cost and would introduce significant risk since there is no prototype implementation on which to base the protocol. The development would require a **minimum** of a year plus an additional two years to introduce it into the global Internet. Only after substantial use by the Internet community would there be commercially available products. The down side to this proposal is that IP, and consequently IPMC, will migrate to a successor (possibly OSI's CLNP) in the same three year time frame it would take to develop this protocol.

The Internet ST-II protocol is also part of the IPS family but is considered an experimental protocol. It is not commercially available, but it is the only protocol which has been proved to work for DIS applications (i.e., multicast <u>and</u> meets real time requirements). ST-II is the multicast protocol used by SIMNET and has been successfully applied to large scale exercises such as WAREX '90. The long-haul testbed for simulation applications, the Defense Simulation Internet (DSI), currently supports IP and ST-II traffic.

2 Secuity may be an exception; the CASS is currently investigating this requirement.

This compatibility with a permanent infrastructure is an added benefit for ST-II.

The last protocol, XTP, is commercially available but offered by only two vendors. Although this protocol offers both multicast and high speed, it has not been proven for DIS applications and has never been proven for wide area networks. While advertised as high-speed, it has yet to significantly out perform IPS' TCP or OSI's TP4. XTP is neither an IPS nor an OSI standard.

Although the ST-II and XTP protocols are candidates to meet interim multicast needs of DIS, they do not meet the long-term requirements. Consequently,

Multicast must be developed for DIS.

To do so, the DIS community has several options: develop a near-term version of multicast, such as the proposal to modify IPMC, or develop multicast for the long-range architecture (i.e., the GOSIP compliant stack). However, given the similarity of IP and CLNP, it would be possible to develop a multicast solution which is relatively "protocol independent" and can evolve to GOSIP as the architecture does. There is no obvious justification to develop unique solutions for each phase of the architecture and considerable reason not to.

4. TRANSITION PLAN

The approach to defining a Transition Plan consists of three steps: 1) establish goals for the communication architecture; 2) select a COTS protocol suite as a starting point; and 3) based on the Applications and standards evolution, establish a time frame for the transition of the architecture to long-range goals. The goals of the communication architecture are based on the research necessary to satisfy long-range requirements. This research will emphasize the coexistence of standards.

The goals of the Transition Plan are:

- To identify research necessary to satisfy long-range requirements,
- To avoid re-implementation at each phase,
- To reduce risk and cost.

4.1 Research

For the architecture to evolve gracefully, research on long-range requirements must continue. The U.S. Army Simulation, Training, and Instrumentation Command (STRICOM) has started work in the American National Standards Institute (ANSI), the International Organization for Standardization (ISO), and the International Telephone and Telegraph Consultative Committee (CCITT) to develop a full range of multicast functionalities. Initially, the work has concentrated in three main areas: the *Multipeer Addendum* to the OSI Reference Model, to provide an overall framework for multicast services; the *Transport Layer*, to provide end-to-end reliability; and the *Network Layer*, to provide a basic multicast data transfer facility. This work supports an overall strategy for multicast in OSI.

DIS will require a Network Layer multicast service in all phases of the architecture. The Network layer multicast being proposed in ANSI/ISO is based on a modular architecture which allows multiple types of multicast service to coexist with one another comfortably and economically [3]. The result are modules which can be combined in a variety of ways to create a wide range of multicast services. These modules are called Layer Functional Modules (LFMs) and are based on the Extended Application Layer Structure (XALS) concepts originally developed for the upper layers. The LFMs can be used as "building blocks" to provide the various functions from basic multicast data transfer to resource reservation to routing.

The LFM approach is fundamentally pluralistic and side steps the emotional issue of "my multicast" or "your multicast" by letting everyone have "their multicast". Any technique can be cast easily and simply into this structure. One benefit of this architecture is its "protocol independent" nature. While the basic multicast data forwarding will be provided by CLNP, interim DIS prototyping could initially build experimental LFMs around IP. Over time, the experimental LFMs can be replaced with standards. This would allow a phased evolution of multicast for prototyping, experimentation, and standardization. Prototyping the LFMs and experimenting with LFMs individually in the interim architecture avoids a complete re-implementation for the OSI/GOSIP architecture.

4.2 Three Phase Approach

To realize a graceful evolution of the architecture, the CASS is recommending a three phase approach. Phase 0, also known as the interim architecture, mandates COTS products with no development or customization. The protocols which make up this phase are basic communication services and will provide an infrastructure for proof-of-concept communication experiments. Due to the maturity of IPS products, Phase 0 will use Internet standards. Since the multicast requirement cannot be satisfied with a COTS product, CASS is leaving the choice "open" rather than requiring a particular protocol. The Phase 0 architecture is a proof-of-concept of the DIS OSI communication infrastructure. In addition to the basic communication protocols identified in the Phase 0 architecture, Phase 1 requires additional multicast capabilities.

The last phase of the architecture, Phase 2, is the GOSIP compliant architecture, based upon lessons learned in Phase 1, added functionality, and final versions of OSI/GOSIP multicast protocols. The phases and corresponding base stacks are defined in the following sections.

Phase 0: Internet Stack - The *base stack* for the Phase 0 architecture consists of the basic communication services. For consistency with subsequent phases of the architecture, Phase 0 will be described in terms of the OSI seven layer model. At the Application layer, five protocols are specified: DIS Application Protocol, Simple Network Management Protocol (SNMP), Telnet (a terminal protocol), File Transfer Protocol (FTP), and Network Time Protocol (NTP). SNMP and Telnet will be used to meet the Network Management service requirement; SNMP will provide network monitoring while Telnet will be used to establish terminal sessions for remote debugging and network management. FTP will be used to satisfy the file transfer requirement by providing a bulk transfer service (i.e., retrieval of databases). The NTP will be used to meet the synchronization requirement.

At the Transport Layer, the architecture is based on the User Datagram Protocol (UDP) for unreliable unicast (or datagram) service and the Transmission Control Protocol (TCP) for reliable data transfer. At the Network Layer, the architecture specifies the Internet Protocol (IP) for seamless local/global communication. As discussed previously, the DIS requirement for multicast cannot be met with any COTS product. There are two candidates which can meet interim requirements: ST-II and XTP. ST-II is the better choice because it has been proved to work for DIS applications and it also provides a better migration to the LFMs being proposed in ANSI/ISO. Prototyping the LFMs should begin in this Phase using IP as the basic data forwarding service. All experience gained with LFMs at this stage of development will be directly transferable to subsequent phases of the architecture. This is not true for experience gained with XTP. At this time CASS is leaving the selection of the multicast protocol to the individuals requiring it for an exercise; however, this Transition Plan recommends the use of ST-II. The interim services can be aligned with the Communication Classes defined in [1] (See Table 2).

As part of the standard operation of IP-over-Ethernet, the mapping between IP-addresses and the corresponding local Ethernet addresses is handled by the Address Resolution Protocol (ARP). The Phase 0 architecture is based on WANs interconnecting Ethernet (IEEE 802.3) LANs with a local broadcast capability at each site.

The standards for Phase 0 are listed below. (The Internet standards are called Request For Comments.)

SNMP	Simple Network Management Protocol (RFC 1157)
Telnet	(Terminal Protocol) (RFC 854)
FTP	File Transfer Protocol (RFC 959)
NTP	Network Time Protocol (RFC 1119)
UDP	User Datagram Protocol (RFC 768)
TCP	Transmission Control Protocol (RFC 793)
IP	Internet Protocol (RFC 791)
ST-II	Stream-II (RFC 1190)
LFMs	Layer Functional Modules: Group Management, Routing, Data Forwarding, Addressing, and Quality of Service
ARP	Address Resolution Protocol (RFC 826)
RARP	A Reverse Address Resolution Protocol (RFC 903)
CSMA/CD	Carrier Sense Multiple Access with Collision Detection (IEEE 802.3)

Phase 1: OSI Stack - The proposed Phase 1 architecture incorporates the OSI seven layer stack to facilitate the migration to a full GOSIP compliant network. The *base stack* for Phase 1 combines the OSI basic communication services with the required multicast services. The Application Layer specifies

Table 2. Phase 0 Service Characterization

	Class 1: Unreliable Multicast	Class 2: Unreliable Unicast	Class 3: Reliable Unicast
Application	DIS PDUs	DIS PDUs SNMP NTP	DIS PDUs FTP Telnet
Transport	UDP	UDP	TCP
Network	IP ST-II LFMs	IP	IP

five protocols: DIS Application Protocol (DIS-AP), Group Management Protocol (DIS-GMP), Common Management Information Protocol (CMIP), Virtual Terminal Protocol (VTP), and File Transfer Access and Management (FTAM). Using XALS, the DIS-AP will allow PDUs to select the required Transport services (i.e., the appropriate communication class). The DIS-GMP will provide the capability to specify group membership management, group initiation, and group communication termination. To satisfy the network management requirement, CMIP and VTP will be used. CMIP provides network management and monitoring, and VTP will be used where ever terminal sessions are needed. FTAM will be used to satisfy the file transfer requirement by providing a bulk transfer service (i.e., retrieval of databases). The synchronization requirement is being developed within the OSI program of work.

The Session and Presentation Layers and Association Control Service Element (ACSE) will be implemented using the OSI Skinny Stack [7] approach and will incorporate extensions to ACSE (A^2CSE).

At the Transport Layer, the architecture is based on the Connectionless Transport Protocol (CLTP) for datagram service and the Class 4 Transport Protocol (TP4) for reliable data transfer. At the Network Layer, the architecture specifies the Connectionless Network Protocol (CLNP) for seamless local/global communication. DIS will need a long-lived Network Layer multicast protocol; LFMs will be used in Phase 1 to satisfy the multicast requirement. The basic data transfer will be provided by CLNP, while the DIS-specific functions, such as resource reservation and multicast distribution, will be provided by individual LFMs. Since LFMs were prototyped in Phase 0, the transition to Phase 1 will be much cleaner and will not require total re-implementation. The transition will require only the substitution of CLNP for IP for carrying data. The OSI services can be aligned with the Communication Classes defined in [1] (See Table 3).

The End System to Intermediate System (ES-IS) protocol provides the equivalent function as the Internet ARP protocol. The architecture will successfully operate over any type of communication subnetwork environment that meets minimum performance requirements (e.g., IEEE 802.3 or FDDI).

The OSI protocols are defined below:

DIS-API	DIS Application Protocol Interface (undefined)
DIS-GMP	DIS Group Management Protocol (undefined)
A^2CSE	ASO-Association Control Service Element (undefined)
XALS	Extended Application Layer Structure (ISO 9545/DAM1)
CMIP	Common Management Information Protocol (ISO 9596)
VTP	Virtual Terminal Protocol (ISO 9041)
FTAM	File Transfer Access and Management (ISO 8571)
CLTP	Connectionless Transport Protocol (ISO 8602)
TP4	Transport Protocol Class 4 (ISO 8073)
CLNP	Connectionless Network Protocol (ISO 8473)
LFMs	Layer Functional Modules: Group Management, Routing, Data Forwarding, Addressing, and Quality of Service (undefined)
ES-IS	End System to Intermediate System Routing Protocol (ISO 9542)
CSMA/CD	Carrier Sense Multiple Access with Collision Detection (ISO 8802/3)
FDDI	Fiber Distributed Data Interface (ISO 9314)

Phase 2: GOSIP Stack - The proposed Phase 2 architecture incorporates the future OSI multicast protocols into the GOSIP compliant network. The additional protocols required in Phase 2 will be the reliable multicast protocol. Although DIS does not currently require a reliable multicast protocol, it is

Table 3
Phase 1 Service Characterization

	Class 1: Unreliable Multicast	Class 2: Unreliable Unicast	Class 3: Reliable Unicast
Application	DIS-AP DIS-GMP	DIS-AP CMIP	DIS-AP CMIP FTAM VTP
Transport	CLTP	CLTP	TP4
Network	CLNP LFMs	CLNP	CLNP

desired for the long-term architecture. This reliability will likely be provided by a Transport Layer protocol called TP5 [4]. The Phase 2 service characterization adds the future Communication Class 4 (See Table 4).

4.3 A Timeline

To make the Transition Plan a viable strategy, it is important to develop a timeline which identifies the transitions from one phase to the next. But how do we define these time periods? The approach used in this Plan is based on the Applications using DIS and the projected standards evolution required to satisfy long-range requirements.

There are seven new procurements or upgrades to existing systems which refer to the DIS standard: CCTT, BFTT, MAIS, TCTS, TACTS, JACTS, and NTC. Most of these projects are in initial planning stages and do not have critical design reviews until the 1994-1995 time frame; actual deployment will take place even later. Until that time, exercises will consist mostly of homogeneous simulators or small groups of heterogeneous simulations (e.g., the *I/ITSC demonstration). This will be the period when* DIS is demonstrated and tested, as well as developed (e.g., versions 2.0 and 3.0 of the standard). During the next 3-4 years, DIS will begin to mature much the way traditional communication protocols do. Therefore, we can postulate that this is our interim time frame.

DIS will have to wait three to four years to get the full range of Network and Transport multicast services it requires. Since the goal of the DIS communication architecture is to move toward OSI/GOSIP, this is where the development work should focus. As discussed previously, the multicast work is being designed to be relatively protocol suite independent. This will allow DIS to begin experimentation with multicast in the interim architecture without losing the experience gained when the architecture evolves to GOSIP compliance. The LFM approach will allow DIS to move quickly from IP to CLNP and will not put DIS in a bind if the ISO standards process gets bogged down. If this strategy is sidetracked by developing unique interim solutions, it will be at great expense to the government and to the DIS community and will prolong the lack of marketability for the services being developed.

Since the Phase 0 (IPS) architecture has a very large experience base in industry, it will require minimal effort to learn and develop. However, the DIS community should not take this to mean that all their communication problems are solved. To the contrary, as discussed in Section 2.1, the IPS is only an interim solution itself since the internetworking protocol IP will have to migrate to a successor in the same interim time frame. The current disarray in the Internet community over the successor to IP makes the stable CLNP and associated protocols a much more attractive target to migrate to.

By synthesizing our knowledge of the Applications and standards, it possible to develop a timeline (See Figure 1) for the evolution of the DIS communication architecture. The line starts with the beginning of the multicast work in 1991 and ends in 1997 with the projected dates of the BFTT and TCTS contracts. Shown on the line are the projected dates of the known Applications which will use DIS along with the corresponding maximum number of entities expected for each program. Also shown is the recommendation of the Phase 0 architecture (3/92) and projected dates for LFM prototypes (9/93), the Phase 1 architecture (3/95), and the Phase 2 architecture (3/96).

<u>Table 4</u>
Phase 2 Service Characterization

	Class 1: Unreliable Multicast	Class 2: Unreliable Unicast	Class 3: Reliable Unicast	Class 4: Reliable Multicast
Application	DIS-AP DIS-GMP	DIS-AP CMIP	DIS-AP CMIP FTAM VTP	DIS-AP
Transport	CLTP	CLTP	TP4	TP5
Network	CLNP LFMs	CLNP	CLNP	CLNP LFMs

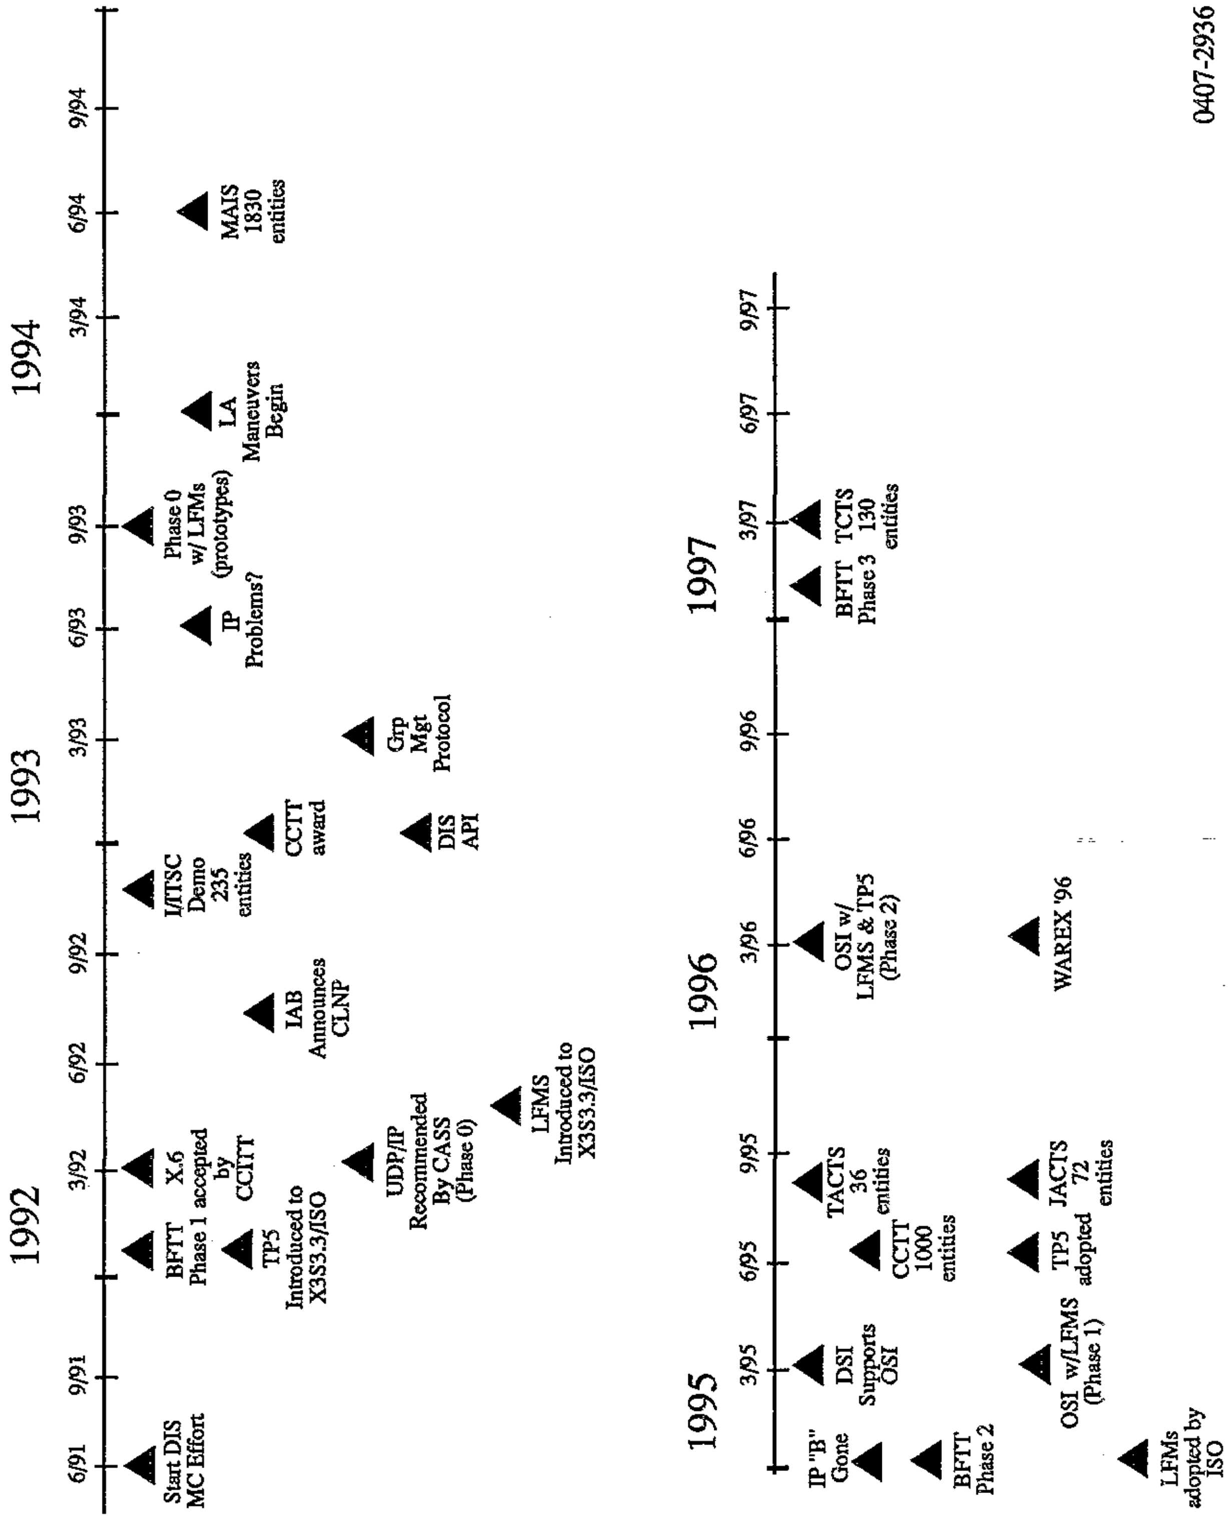

Figure 1 Timeline

A prediction was made at the March 1992 CASS meeting for when DSI would support OSI traffic. This date (3/95) is also included on the line and strengthens the postulated interim time frame.

5. CONCLUSIONS

Today, there is no one architecture (Internet or OSI) which meets the total needs of DIS. This means that the DIS community will have to develop an architecture to meet its requirements. This development will need to occur in phases to allow sufficient time for users and implementers to gain experience with current requirements before introducing new services. The phases CASS is recommending provides a good structure for the DIS transition to GOSIP compliance. The Phases provide DIS with a starting place (Phase 0), an interim OSI stack (Phase 1), and a final OSI/GOSIP stack (Phase 2). The basis for starting with the IPS base stack is strengthened by the fact that the DSI currently supports only IP and ST-II traffic. However, DSI will begin routing OSI traffic in the next three years (again, our postulated interim time frame). The DIS users must take advantage of this time to gain OSI experience and begin evolving the architecture before it is required for interoperability.

The transition to OSI/GOSIP will not happen overnight. If DIS desires GOSIP compliance, we as a community must take the initiative to push the evolution of the protocols which meet our requirements. Until the time that all services are met by one protocol suite, the communication standards (IPS and OSI) will have to coexist. This means that while development of services and protocols are occurring on the long-term architecture, proof-of-concept testing should occur on the interim architecture. The multicast LFMs are a perfect example. This approach will reduce the risk of integrating new services and provide a graceful evolution of the architecture DIS requires.

The approach presented here is flexible enough to accommodate existing and future communications technologies, while providing a graceful growth path for the maturation of the architecture. Now is the time for the DIS community to think about the range of problems it faces and formulate a strategy for achieving the goals within the required time frame. We should not blindly react to interim solutions which, at best, patch the problem and require re-development at a subsequent phase of the architecture. By following this strategy, DIS can focus attention on the long-range answers we need.

ACKNOWLEDGEMENTS

This work has been supported by the U.S. Army STRICOM under contract N61339-91-C-0103.

The authors would like to thank the following individuals for providing the Internet and OSI information found in this paper: Lyman Chapin, BBN Communications; Bob Cooney, Naval Computer and Telecommunications Station; John Day, BBN Communications; Richard desJardins, The GOSIP Institute; Debbie Deutsch, BBN Systems and Technologies; Susan Hares, Merit Network; Al Kerecman, USACECOM; Dave Marlow, Naval Warfare Systems Center; Jim Moulton, Open Network Solutions; and Dave Oran, Digital Equipment Corporation.

REFERENCES

[1] *Communication Architecture for DIS, Draft Standard,* March 1992.

[2] "The Department of Defense Open Systems Interconnection (OSI) Implementation Strategy", DCA/DCEC, May 1988.

[3] "Layer Management Enhancements for Network-Layer Multicast: Architecture", D. Deutsch and J. Day, March 27, 1992, X3S3.3/92-128.

[4] "Proposed Contribution to SC 6 on Multicast Transport Protocol", J. Moulton, March 26, 1992, X3S3.3/92-142.

[5] "Internet Architect Gives Long-Term View", *Network World,* May 18, 1992.

[6] *Strawman Distributed Interactive Simulation Architecture Description Document, Volume I: Summary Description,* Loral System Company, ADST Program Office, March 31, 1992.

[7] "The OSI Skinny Stack", P. Furniss and K. Ashley, University of London Computing Centre, March 25, 1992.