THE DEPARTMENT OF DEFENSE HIGH LEVEL ARCHITECTURE

Judith S. Dahmann

Richard M. Fujimoto

Richard M. Weatherly

Defense Modeling and Simulation Office
1901 North Beauregard Street
Alexandria, VA 22311, U.S.A.

College of Computing
Georgia Institute of Technology
Atlanta, GA 30332-0280, U.S.A.

The MITRE Corporation
7525 Colshire Drive
McLean, VA 22102-3481, U.S.A.

ABSTRACT

The High Level Architecture (HLA) provides the specification of a common technical architecture for use across all classes of simulations in the US Department of Defense. It provides the structural basis for simulation interoperability. The baseline definition of the HLA includes (1) the HLA Rules, (2) the HLA Interface Specification, and (3) the HLA Object Model Template. This paper describes the motivations and processes used to develop the High Level Architecture and provides a technical description of key elements of the architecture and supporting software. Services defined in the interface specification for providing time management (TM) and data distribution management (DDM) for distributed simulations are described.

1 INTRODUCTION

The Defense Modeling and Simulation Office (DMSO), is addressing the continuing need for interoperability among new and existing simulations within the U.S. Department of Defense through its High Level Architecture (HLA) initiative. The HLA seeks to generalize and build upon the results from the Distributed Interactive Simulation (DIS) world and related efforts such as the Aggregate Level Simulation Protocol (ALSP) (Wilson and Weatherly 1994). A goal of the HLA activity was to develop and recommend an architecture to the Executive Council for Modeling and Simulation (EXCIMS) before the end of calendar year 1996. This was achieved, and the EXCIMS in turn, after appropriate review, recommended the architecture to the Under Secretary of Defense (Acquisition and Technology) for approval and standardization. The baseline HLA definition was approved as the standard technical architecture for all U.S. DoD simulations on 10 September 1996. This tutorial provides an overview of the HLA, updating and expanding upon (Dahmann 1997). Additional information concerning the HLA concept and the DMSO Master Plan is available at <http://www.dmsomil>.

The High Level Architecture (HLA) provides a common architecture supporting reuse and interoperation of simulations across the U.S. Department of Defense. The HLA is based on the premise that no one simulation can satisfy all uses and users. An individual simulation or set of simulations developed for one purpose can be applied to another application under the HLA concept of the federation: a composable set of interacting simulations. The intent of the HLA is to provide a structure which will support reuse of capabilities available in different simulations, ultimately reducing the cost and time required to create a synthetic environment for a new purpose.

The HLA is intended to have wide applicability, across the full range of defense simulation applications including training, analysis, and engineering functions, at a variety of levels of resolution. These widely differing application areas include a variety of requirements which had to be considered in the development and evolution of the HLA.

The definition of architecture used in this effort -- "major functional elements, interfaces, and design rules, pertaining as feasible to all simulation applications, and providing a common framework within which specific system architectures can be defined" -- is one that is commonly accepted and is consistent with the IEEE definition of architecture for computer simulations. For the purpose of this effort, the emphasis is on the development of a high level architecture which pertains as widely as possible to all simulation areas and which will provide a framework for the development of specific system architectures. The HLA does not prescribe a specific implementation, nor does it mandate the use of any particular set of software or programming language. Over time, as technological advances become available, new and different implementations will be possible within the framework of the HLA.

2 THE HLA DEVELOPMENT PROCESS

The HLA was developed by the US Department of Defense (DoD) based on a process involving government, academia, and industry. In FY94, the Defense Advanced Research Projects Agency (DARPA) awarded three industrial contracts for the definition of a high level architecture. The results of these three contractor efforts were received in January 1995, and a core government and academic team combined the best of those inputs with additional insights from other DoD modeling and simulation (M&S) projects to arrive at the initial definition of the HLA. This definition was presented to the Defense Modeling and Simulation Office (DMSO)-sponsored Architecture Management Group (AMG) on 31 March 1995.

The AMG developed the architecture based on cooperative efforts to apply the HLA in a series of prototypes designed to ensure that it addressed the broad set of application requirements for DoD simulations. The result was the baseline HLA definition, completed in August 1996 and approved as the standard technical architecture for all US DoD simulations on 10 September 1996. The AMG continues to evolve the HLA as needed.

The AMG is made up of representatives of major U.S. DoD simulation programs. Programs were selected because they represent the principal simulation requirements, ranging from analysis to training to test and evaluation, from detailed engineering representation to campaign-level warfighting to man-in-the-loop interactive vehicle simulators. Members, along with their supporting government and industrial teams, participate in regular meetings every 6 to 8 weeks to address the key HLA issues, conduct cooperative prototyping efforts with other AMG members (prototype federations and experiments) and participate in technical working groups to evolve crosscutting technical aspects of the HLA.

It is also very important that the HLA be integrated into the broader technical community. Work is underway in a partnership with the Simulation Interoperability Workshop (SIW) and the Simulation Interoperability Standards Organization (SISO). The DIS Workshop (the precursor to the SISO and SIW) was represented on the AMG during the HLA baseline development period and has committed to establish industry standards to support the HLA. This effort was initiated at the 14th DIS Workshop in March 1996.

3 TECHNICAL ARCHITECTURE

3.1 Motivations

The HLA design is based on certain assumptions. First, the HLA is based on the premise that no one simulation can solve all the U.S. DoD functional needs for modeling and simulation. The needs of the users are too diverse. The technical complexity of needed implementations is beyond what has been shown to be possible today or is reasonably likely in the foreseeable future to be handled in a single simulation. Further, with changing user needs, it is just not possible to anticipate how simulations will be used in the future or in what combinations. Perhaps equally as important is the changing state of technology; the HLA will need to allow for leveraging new technical innovations.

Since it is not possible to anticipate all the uses of M&S in the U.S. DoD, it is important to think in terms of multiple simulations which can be reused in a variety of ways. This means that as simulations are developed, they must be constructed so that they can be easily brought together with other simulations, to support new and different applications. While the HLA attempts to impose as few constraints as possible on the design of individual simulations, given the drive for reuse, the HLA recognizes that simulations themselves will need to be constructed so that they can easily be extended to support new features needed by unanticipated applications.

These assumptions have affected the HLA design in several ways. Clearly, the architecture itself must have modular components with well-defined functionality and interfaces. Further, the HLA has separated the functionality needed for individual simulations from the infrastructure required for interoperability among simulations.

3.2 Functional Overview

Figure 1 shows how an HLA federation is partitioned into its major functional components. The first key component are the simulations themselves, or more generally, the federates. A federate can be a computer simulation, a manned simulator, a supporting utility (such as a viewer or data collector), or even an interface to a live player or instrumented range. All object representation is in the federates. The HLA imposes no constraints on what is represented in the federates or how it is represented, but it does require that all federates incorporate specified capabilities to allow the objects in the simulation to interact with objects in other

Figure 1: Functional View of an HLA Federation

simulations through the exchange of data supported by services in the Runtime Infrastructure (RTI).

The second functional component is the RTI. The RTI is, in effect, a distributed operating system for the federation. The RTI provides a set of services that support the simulations in carrying out these federate-to-federate interactions and federation management support functions. These services will be discussed later. All interactions among the federates flow through the RTI.

The third functional component is the runtime interface. The HLA runtime interface specification provides a standard way for federates to interact with the RTI, to invoke the RTI services to support runtime interactions among federates, and to respond to requests from the RTI. This interface is implementation independent and is independent of the specific object models and data exchange requirements of any federation.

Two other general capabilities of simulation systems are supported by the architecture. First, the HLA supports the passive collection of simulation data and monitoring of simulation activities. In the HLA, these tools act in the same way as simulations and interact with the RTI using the HLA interface.

Second, the HLA supports interfaces to live participants, such as instrumented platforms or live command and control (C2) systems. Live participants interact with the simulated world through something that acts like a simulation from the point of view of the HLA, that feeds a representation of the live world into

the simulated world and that projects data from the simulated world back to the live system.

The HLA is formally defined by three components: the object model template (Defense Modeling and Simulation Office 1996), the interface specification (Defense Modeling and Simulation Office 1996), and the HLA rules (Defense Modeling and Simulation Office 1996).

3.3 HLA Object Models

HLA object models are descriptions of the essential sharable elements of the simulation or federation in 'object' terms. The HLA is directed towards interoperability; hence in the HLA, object models are intended to focus on description of the critical aspects of simulations and federations which are shared across a federation. The HLA puts no constraints on the content of the object models. The HLA does require that each federate and federation document its object model using a standard object model template. These templates are intended to be the means for open information sharing across the community to facilitate reuse of simulations. These completed templates will be openly available, and tools are being developed to allow for automated search and reasoning about object model template data, to further facilitate cost-effective information exchange and reuse.

The HLA specifies two types of object models: the HLA Federation Object Model (FOM) and the HLA

Simulation Object Model (SOM). The HLA FOM describes the set of objects, attributes and interactions which are shared across a federation. The HLA SOM describes the simulation (federate) in terms of the types of objects, attributes and interactions it can offer to future federations. The SOM is distinct from internal design information; rather it provides information on the capabilities of a simulation to exchange information as part of a federation. The SOM is essentially a contract by the simulation defining the types of information it can make available in future federations. The availability of the SOM facilitates the assessment of the appropriateness of the federate for participation in a federation.

While the HLA does not define the contents of a SOM or FOM, it does require that a common documentation approach be used. Both the HLA FOM and SOM are documented using a standard form called the HLA Object Model Template (OMT).

3.4 HLA Interface Specification

The HLA interface specification describes the runtime services provided to the federates by the RTI, and by the federates to the RTI. There are six classes of services. *Federation management* services offer basic functions required to create and operate a federation. *Declaration management* services provide a means for federates to declare what information the federate will provide and require during a federation execution. *Object management* services provide creation, deletion, identification and other services at the object level. State updates and interactions result in messages being transmitted to other federates that have indicated interest in receiving this information. *Ownership management* supports the dynamic transfer of ownership of object/attributes during an execution. This is necessary because at any instant, only one federate, termed the “owner,” is allowed to update the value of a particular attribute. *Time management* services support synchronization of runtime simulation data exchange. Finally, *data distribution management* services support the efficient routing of data among federates during the course of a federation execution. The HLA interface specification defines the way these services are accessed, both functionally and in a programmer's interface. The time management and data distribution management services are described next in greater detail.

3.4.1 Time Management

Time management is concerned with the mechanisms for controlling the advancement of each federate in simulation time (referred to as the federation time axis in the HLA). In general, time advances must be coordinated with object management services so that

information is delivered to federates in a timely and ordered fashion, e.g., to satisfy requirements for reproducing causal behavior in the system being modeled. Specifically, a federate may specify that the messages it sends should be delivered to federates in receive order or time stamp order (TSO). Advances in simulation time may or may not be paced to advances in wallclock time.

Federates will normally use one of the following approaches to time management:

- *paced, independent time advances.* Federate time advances are *paced* to occur in synchrony with wall clock time (or scaled wall clock time, derived as a linear function of wall clock time). The federate autonomously advances its own time without coordinating such advances with the RTI. “Human-in-the-loop” training federates and “hardware-in-the-loop” test and evaluation federates will typically utilize this approach.
- *paced, coordinated time advances.* Federate time advances are paced to occur in synchrony with (scaled) wall clock time, but time advances are coordinated with the RTI to ensure that before-and-after relationships in the physical system are properly reproduced. Analytic simulation models embedded in environments including humans and/or live elements or physical devices typically utilize this approach.
- *unpaced, coordinated time advances.* Federate time advances are not paced to occur in synchrony with wall clock time, and the federate coordinates time advances to ensure before-and-after relationships are properly modeled. Analytic simulation models intended to execute “as-fast-as-possible” typically utilize this approach.

It is anticipated that some federations may include combinations of federates using different approaches to time management (e.g., coordinated and independent time advance federates may be utilized in a single federation). The time management services are intended to support federations that include federates with different ordering and delivery requirements.

The time management services defined in the HLA are primarily concerned with coordinated time advance federates. There are three key components of the time management services:

1. *Logical time:* Logical time is synonymous with federation time for coordinated time advance federates. Temporal relationships among events are specified by the logical time values (time stamps) assigned to the events. In general, at any instant during the execution, different coordinated time federates may be at different logical times.

2. *Mechanisms to advance logical time:* Services are provided for each federate to advance its own logical time. Federates must explicitly request logical time advances, and the advance does not occur until the RTI issues a grant. If advances in logical time must be paced to be in synchrony with wallclock time, such pacing must be done within the federate(s).
3. *Message ordering and synchronization:* The RTI will only grant an advance to logical time T when it can guarantee all time stamp ordered (TSO) messages with time stamp less than T (or in some cases less than or equal to T) have been delivered to the federate. This guarantee enables the federate to simulate the behavior of the entities it represents up to logical time T without concern for receiving new events with time stamp less than T .

Each federate using logical time must specify a non-negative *lookahead* value. If a federate's current logical time is T , all TSO messages generated by that federate must have a time stamp value of at least T plus the federate's lookahead. This constraint increases the amount of concurrent execution in the distributed simulation.

In order to determine which TSO messages can be delivered to a federate F without violating the guarantee that messages are delivered in time stamp order, the RTI must internally compute a lower bound on the time stamp (LBTS) of future messages that will be later received that are destined for F . Only TSO messages with time stamp less than or equal to F 's LBTS value are eligible for delivery. At any instant, the RTI may hold one or more messages with time stamp less than LBTS in its internal queues. The minimum among the federate's LBTS value and messages stored in the RTI's queues gives a lower bound on the time stamp of any TSO messages that will be delivered to the federate in the future. LBTS is computed within the RTI using a conservative synchronization protocol, e.g., see (Fujimoto 1990).

The HLA provides several services to support the inclusion of optimistic federates, e.g., federates synchronized using the Time Warp mechanism (Jefferson 1985). The LBTS value computed by the RTI enables the federate to compute Global Virtual Time (GVT). A service for retracting previously scheduled events is provided, thereby enabling the implementation of anti-messages. A service is provided to allow optimistic federates to flush the time stamp ordered message queue, enabling the optimistic processing of events. An important property of the time management services is they ensure that optimistic events (events that may later be canceled) are not delivered to federates unless they explicitly request them via the flush queue primitive, thereby enabling conservative federates to operate in the

same federation with optimistic simulations without being concerned with having to implement rollback and message cancellation mechanisms. At the same time, optimistic federates that provide this capability can freely exchange optimistic events.

A federate is said to be *time regulating* (or equivalently, time regulation is enabled) if it can send time stamp ordered (TSO) messages. Information such as the current logical time and lookahead of time regulating federates are used by the RTI in LBTS computations. TSO messages sent by a federate that is *not* time regulating are automatically converted to receive ordered by the RTI. Similarly, a federate is said to be *time constrained* (time constrained is enabled) if it can receive time stamp ordered (TSO) messages. TSO messages received by a federate that is not time constrained are automatically converted to receive ordered by the RTI.

3.4.2 Declaration Management and Data Distribution Management

Data management (DM) and data distribution management (DDM) in the HLA are used to specify which federates should receive messages for each state update and interaction (Morse 1996). DM services *Publish* and *Subscribe* allow a federate to update and receive updates to object attributes based solely on object class. The RTI uses information provided in publish/subscribe calls to set up filters that direct data among federates that need them. The object management's (OM) *Update Attribute Values* service call notifies the RTI that one or more attributes have been modified. For example, an object might subscribe to the attribute 'location' of all tanks in the battlefield. However, such filtering will only be appropriate for relatively small federations. DDM services provide more powerful data distribution services enabling value-based filtering (Defense Modeling and Simulation Office 1996). For example a tank might want to receive data from other tanks only if they are in its visible range. Note that since the RTI does not know the meaning of object attributes it cannot provide range-based filtering in this example. Federates must agree on a filtering strategy to do this.

The fundamental concept used in the HLA to support value-based DDM is the *routing space*. A routing space is a normalized (coordinate values range from 0.0 to 1.0) multidimensional coordinate system in which federates indicate interest in receiving or providing updates via *subscription* and *update regions*. Subscription and update regions are rectangular (in N dimensions) and are specified by indicating *extents*, with one extent for each dimension. Each extent indicates the portion of that dimension covered by the region. For example, the

extents $[0.0,0.5][0.0,1.0]$ specify the left half of a two dimensional routing space. Federates express their interest in receiving updates from other federates through subscription regions defined over a routing space, and are called subscribers (to a specific attribute). On the other hand, federates that are providing updates define their update regions over the routing space, and are called publishers (of a specific attribute). The RTI detects when subscription and update regions associated with an attribute overlap, and provides data transfer from the publisher to all subscribers for all updates to that attribute.

Figure 2 shows a two-dimensional routing space with one update U1 and two subscription regions S2 and S3 that belong to federates F1, F2 and F3 respectively. Since U1 and S2 overlap, the RTI will transmit updates to attributes by F1 that are associated with this update region to subscriber F2, but not to F3. Regions can be changed dynamically by invoking the *Modify Region* service.

Figure 2: Update Region U1 and Subscription Regions S2 and S3 in a Two Dimensional Routing Space

3.5 HLA Rules

Finally, the HLA rules summarize the key principles behind the HLA. These are divided into two groups: federation and federate rules. Federations, or sets of interacting simulations or federates, are required to have a FOM in the OMT format. During runtime, all object representation takes place in the federates (not the RTI) with only one federate owning any given attribute of an instance of an object at any time. Information exchange among the federates takes place via the RTI using the HLA interface specification.

Additional rules apply to individual federates. Under the HLA, each federate must document their public information in their SOM using the OMT. Based on

the information included in their SOM, federates must import and export information, transfer object attribute ownership, updates attributes and utilize the time management services of the RTI when managing local time.

4 HLA PROTOTYPING

The purpose of the HLA baseline development period was to take the initial HLA definition to develop the needed specifications and tools to actually use the HLA to support prototype applications. This prototyping was directed toward addressing a set of common issues (Defense Modeling and Simulation Office 1996) and to evolve the HLA to meet the broad set of needs. These issues include the technical viability of a single interface specification to support the wide application base of the HLA and the technical feasibility of building an RTI with the necessary range of tools needed by those applications, the impact of the HLA on simulation internal development, the utility of the object model concept and formats throughout the HLA life cycle, the ability to specify common testing methods, the ability to operate HLA federations in a secure mode and, finally, the applicability of the HLA to address the wide range of US DoD simulation applications.

There were four major prototype implementations of federations ("proto-federations") using the HLA to support the HLA baseline development process. The simulations incorporated in these proto-federations were made available, along with their support technical development teams, by AMG members, as part of the HLA development process. Each proto-federation addressed a different corner of the simulation application domain, and as such, generally typifies one way simulations will be used in an HLA context.

Several overall points are warranted before the proto-federations themselves are described. The purpose of these prototypes was to develop the HLA through hands-on experience, not to solve a specific real-world problem. In several cases the scenarios for the federations had to be stretched to support the mix of federates available from the AMG. Second, all four prototypes used the same architecture, the same interface specification, the same object model template and, in fact, the same RTI software implementation.

For HLA baseline development purposes it was decided to develop one prototype implementation of the RTI and use that implementation with all the proto-federations. This was seen as both the most cost-effective approach and a strategy that would allow for the examination of the desirability of having not only a common interface specification, but also one common implementation which could serve the entire community. The primary objective of the RTI

prototyping was to provide the other proto-federations with the needed functionality for them to implement and use the HLA for their application needs. The RTI prototype also provided the mechanism to assess the functional and performance requirements for RTI service provision as input to future RTI development and transition. The RTI software is being re-engineered to incorporate the full set of services defined in the HLA Baseline Definition.

The first proto-federation was called the Platform Prototype federation. This was a federation of real-time platform level simulators and simulations which are current users of IEEE DIS 2.0. The key issues addressed by this proto-federation were the ability of the HLA to provide the capabilities now supported by DIS and to address the issues associated with adapting a current DIS-enabled system to work under HLA.

The second federation, called the Joint Training Prototype, was a federation of discrete event simulations. The key issues for this prototype were the ability to coordinate multiple discrete event simulations using the RTI time management services, the ability to hand-off ownership of attributes in a federation and options for providing dynamic environmental effects across a federation.

The third prototype federation was the Analysis Proto-federation. This proto-federation is comprised of closed form analytic applications; hence, runtime efficiency and replicability are important here, since these simulations will eventually require multiple executions for analysis purposes. The interest in this proto-federation was in the time management and data management services.

The Engineering Prototype federation examined use of the HLA to support simulations to support system acquisition using validated, engineering-level simulations. This proto-federation was particularly interested in object ownership management which affords the option of handing off computation of certain effects during select parts of a scenario to selected simulation facilities and in performance requirements to support hardware-in-the-loop applications.

These prototypes were implemented to apply the HLA to applications which typify its ultimate uses to support its definition. The results of the prototypes were fed into the specification and testing procedures through the working groups. The experiences of the individual federate developments support assessment of the effect of the HLA on simulations and offer an experience base for future users of the HLA.

Finally, the experience of developing these proto-federations provided the basis for the development of a general process view of the HLA. While HLA development focuses on the specification of the technical

components of the architecture, it has been recognized that it is very important to understand the process which supports the use of the architecture for different applications. The baseline definition of the HLA was developed based on a set of proto-federations. Using the experience of these proto-federations as a driver, a general process for the development and execution of applications using the HLA is being formulated as a recommended practice to accompany the HLA definition (Defense Modeling and Simulation Office 1996). Understanding this process is viewed as a very important element in the development and application of the HLA.

5 HLA SUPPORT FOR SIMULATION REUSE AND INTEROPERABILITY

Simulation interoperability is defined as "the ability of a ... simulation to provide services to, and accept services from, other ... simulations, and to use the services so exchanged to enable them to operate effectively together." (U.S. Department of Defense 1994) This definition of interoperability reflects the overall objective of the HLA that different simulations be able to effectively share data towards a common goal.

As the definition indicates, there are two elements in interoperation: effective data sharing and consistent data interpretation. The HLA requires that federates build the functionality required to interface with the RTI and exchange data with other federates via the HLA specified interfaces. This enables federates to participate in federations, to exchange data, and coordinate their operations with a federation. The HLA also requires all federates and federations to document characteristics of their object representations relevant to other potential users of the federate or federation. This documentation, in the form of completed object model templates, facilitates the information exchange needed for a consistent interpretation of shared data.

Universal interoperability (the ability of any simulation to interoperate with any other simulation, regardless of original purpose or technical implementation) is not feasible with today's technology. Realistically, interoperability will be attainable in degrees, with the required level of interoperability determined by the needs of the federation user. Where interoperability deals with the logical exchange of information between distinct federates during runtime, reuse refers to the adaptation of components (e.g., ideas, whole simulations, lines of code) during the development of a new simulation. Reuse is assisted by having well-defined modular components (the federates) which share the common understanding of what is the meaning of the objects contained within them.

The HLA rules, interface specification, and object model template provide minimum essential tools for interoperability. They ensure that the basic capability for information exchange is in place, they establish the mechanisms for runtime data transfer, and they provide the means to identify appropriate simulations for different purposes. Beyond this, additional interoperability requires additional consistency in the internal representation of the simulations themselves. The extent to which this consistency is required for any particular application depends on the characteristics of that application. While the HLA in itself is insufficient to guarantee interoperability, it provides the technical framework for simulation and federation developers to achieve the degree of interoperability needed to achieve their objectives.

6 HLA TECHNICAL LIBRARY

DMSO has established an on-line HLA Technical Library, accessible through the DMSO Home Page on the World Wide Web (<http://www.dmsomil>), which contains not only the HLA Baseline Definition documents, but a wealth of supporting documentation as well. This includes test procedures, design details, an HLA security architecture, examples of the use of the HLA, and an extensive archive of briefings and documents, including all of the AMG proceedings throughout the HLA prototyping phase and published HLA papers. The HLA Technical Library will continue to grow and evolve as the HLA itself is carried forward. It is intended to be an open source for use by implementors of the HLA.

REFERENCES

- Dahmann, J. S. 1997. High level architecture for simulation. *Proceedings of the First International Workshop on Distributed Interactive Simulation and Real-Time Applications*, 9-14.
- Defense Modeling and Simulation Office 1996. Data distribution and management design document, V0.2. Washington D.C. December (<http://www.dmsomil>).
- Defense Modeling and Simulation Office 1996. High level architecture baseline development plan, V1.7. Washington D.C. April 1 (<http://www.dmsomil>).
- Defense Modeling and Simulation Office 1996. High level architecture federation and execution process model, V1.0. Washington D.C. September 6 (<http://www.dmsomil>).
- Defense Modeling and Simulation Office 1996. High level architecture interface specification, V.1.0. Washington D.C. August 15 (<http://www.dmsomil>).

- Defense Modeling and Simulation Office 1996. High level architecture object model template. Washington D.C. August 15 (<http://www.dmsomil>).
- Defense Modeling and Simulation Office 1996. High level architecture rules, V1.0. Washington D.C. August 15 (<http://www.dmsomil>).
- Fujimoto, R. M. 1990. Parallel discrete event simulation. *Communications of the ACM* 33(10): 30-53.
- Jefferson, D. 1985. Virtual time. *ACM Transactions on Programming Languages and Systems* 7(3): 404-425.
- Morse, K. 1996. Interest management in large scale distributed simulations, University of California, Irvine.
- U.S. Department of Defense 1994. Department of Defense Directive on modeling and simulation management (5000.59). Washington D.C. January 4 .
- Wilson, A. L., and R. M. Weatherly 1994. The Aggregate Level Simulation Protocol: An evolving system. *Proceedings of the 1994 Winter Simulation Conference*, 781-787.

AUTHOR BIOGRAPHIES

JUDITH S. DAHMANN is the Chief Scientist for the U.S. Defense Modeling and Simulation Office. She has lead the development of the High Level Architecture.

RICHARD M. FUJIMOTO is a professor with the College of Computing at the Georgia Institute of Technology. He served as the technical lead in defining the time management services of the HLA.

RICHARD M. WEATHERLY is the Chief Engineer for the MITRE Corporation's Information Systems and Technology Division. He wrote the first version of the HLA Interface Specification and lead the RTI 0.X, F.0, and 1.0 software development teams.