Extending the High Level Architecture Paradigm to Economic Simulation

James A. Calpin **-** calpinj@mitre.org,
Marnie R. Salisbury **-** marnie@mitre.org,
John A. Vitkevich, Jr. **-** jvitkev@mitre.org,
David R. Woodward **-** woodward@mitre.org.

The MITRE Corporation
1820 Dolly Madison Blvd.
McLean VA 22102

Abstract

This paper describes the potential application of a US Department of Defense-developed simulation architecture to the domain of economic and financial simulation. This approach, known as the High Level Architecture (HLA) for Simulation Interoperability, provides the ability to link disparate simulations. This approach supports simulation interoperability and re-use, allowing for the development of composable sets of interacting simulations. The HLA approach has proven to be extremely successful in the military simulation arena, a domain that has a number of strong analogies to the economic and financial simulation domain. This paper provides an overview of the tenets of the HLA approach and proposes an extension of the tools and capabilities to the economic simulation domain.

1. Background

Currently, there appear to be strong analogies between the approaches used by the macroeconomic community in their modeling of national economies to those used by the military modeling and simulation (M&S) community in their representation of theater-level warfare. Both approaches can be characterized by simplified, low-resolution, aggregate representations of complex processes. Both domains have many different models purporting to address the same questions and phenomena, all with varying degrees of success. The validity of models in both domains is often open to question, and there is little agreement other than "something better is needed."

In the military M&S world, entity-level representation (where each battlefield system such as a tank or plane is represented individually) has been implemented to introduce a much higher level of fidelity in training simulations. The technologies associated with this technical approach have generally been referred to as the domain of advanced distributed simulation, or ADS. In select areas, these techniques and tools have been applied to the analytical domain. Standardized protocols (e.g., Distributed Interactive Simulation (DIS)) and technological innovations like the High Level Architecture (HLA) have steadily improved the efficacy of distributed simulation in the training and analysis domains. This approach has allowed a radical scaling upwards of the number of entities involved in the simulation. Such scaling, combined with constant advances in computational horsepower, have opened up the promise of representing the interactions of hundreds of thousands (perhaps millions) of highly complex entities.

We postulate that the principles of ADS in general, and in particular the HLA paradigm, may be transferable to the realm of macroeconomics. The similarities between the domains are rather striking. In both cases, the entities in question are all acting in accordance with fairly well-defined and relatively well understood rules of behavior. In both the military and economic domains, there are hierarchies of entities acting in concert, sometimes in opposition to each other, with each subject to random events and fluctuations. Entities act and react in an environment that is dynamic (sometimes highly so), and which has a direct impact on the decisions and behaviors being made. Analysts are typically interested in the aggregated results of the interactions of these entities over time; on the military side, such metrics include force exchange ratios, losses over time, and Forward Line of Own Troops (FLOT) movement, while in the macroeconomic arena these measures would typically include gross domestic product (GDP), unemployment rates, inflation rates, and productivity growth, and balance of international payments.

The overall objective of this paper is to describe the novel features of the HLA and how its capabilities might be extended to the economic simulation domain, particularly the branch associated with agent-based computational economics. By developing a framework for the use of HLA concepts, it is anticipated that the computational economics community will recognize the benefits of such an approach and will embrace and extend the framework developed in the defense sector, resulting in a textbook case of true technology transfer.

This paper is organized into three distinct sections. The first explains the underlying requirements and motivation that prompted the development of a capability like the HLA. The next section explicitly defines the components that comprise the HLA, while the last section addresses the potential extension of ADS and HLA principles into the computational economics domain.

2. Motivations for a High Level Architecture

From a military standpoint, simulations of all sorts can provide powerful tools to help maintain military readiness, plan operational missions, make optimal investment decisions, analyze force structure alternatives, and augment live-fire testing and training. Simulations (a general term including both pure-software, or "constructive," simulations and human-in-the-loop, or "virtual," simulators) are by definition abstractions of the real world.[1] Different user needs dictate different abstractions, while different entities, attributes and interactions must be represented at different levels of resolution and fidelity. These representations will, of necessity, be implemented in different computing environments and run on hardware platforms that range from personal computers to massively-parallel, high performance computers.

The genesis of the High Level Architecture can be traced back to the mid 1990s, when the Department of Defense surveyed the tremendous amount of resources that were being poured into different simulation initiatives within the Department. In an era of declining government resources devoted to defense, there was a strong motivation to seek across-the-board improvements in productivity and efficiency within DoD. The defense M&S domain was perceived to be particularly amenable to such actions.

Desires for improved efficiency did not change the fact that the DoD was faced with the need for many different simulations to support many different requirements. If the Department was indeed to use simulations cost-effectively, it needed the flexibility to reuse simulations to the maximum possible extent, building new representations only when existing simulations cannot provide the needed capabilities. In order to get the greatest return on investment for the simulations it did build and maintain, DoD needed to be able to team these representations together in different combinations ("federations") to satisfy a diverse and ever-evolving set of user needs. To allow maximum utility and flexibility, these simulation environments would be constructed from affordable, reusable components interoperating through an open systems architecture. From this set of requirements the concept of the HLA was born.

A number of generalizeable observations were made concerning these requirements and other relevant issues associated with military M&S, observations that were critical in guiding the development of the underlying rationale for the HLA. It is interesting to note the degree to which many (if not all) of these are also directly extensible or applicable to the computational economics domain, thereby clearly illustrating the domain-independent nature of the HLA requirements:

- No single, monolithic simulation can satisfy the needs of all users. Users differ in their interests and requirements for fidelity and detail.

[1] In order to promote a clearer understanding of the scope of this requirement, it is necessary to understand the difference between <u>model</u> and <u>simulation</u>. A model is "a physical, mathematical, or otherwise logical representation of a system, entity, phenomenon, or process;" whereas, a simulation is "a method for implementing a model over time" (Dahmann 1998). Models do not interoperate in parallel in a time-coordinated manner (except as internal components of simulations). It should be made clear that the HLA is largely applicable to simulations, rather than models such as spreadsheets or linear programs (although such models may in fact benefit from HLA interoperability, depending on the individual application and requirement).

- Simulation developers vary in their knowledge of domains to be simulated. No one set of developers is expert across all details even in one domain.
- No one can anticipate all the uses of simulation, and all the ways simulations can be usefully combined. Even if we were able to satisfy a comprehensive set of requirements in a domain of application (by building a "monolithic simulation"), the effort would fail at the outset in anticipating the requirements for such a system.
- The ability to include and incorporate future tools and technology must be allowed. Were there to be a set of perfect requirements for a given problem or domain, and had a perfect implementation been developed to address those requirements, information technology and associated capabilities would still advance. Moore's Law has held for the past 30 years and the continuing growth of computational power shows no end in sight.

These observations led key decisionmakers within the defense M&S community to conclude that a common framework for simulation was necessary and that the time was ripe for the development of an overarching architecture. Further study and analysis of the underlying requirements for such a framework led to the following conclusions:

- It should be possible to decompose a large simulation problem into smaller parts. The smaller parts will likely be easier to define, build correctly, and verify.
- It should be possible to combine the resulting pieces into a larger, correct, and useful simulation.
- It should be possible to combine those components with others to form a completely new simulation.
- Functions that are common to simulations should be separated from the source code of specific simulations and made available as services. A resulting generic infrastructure should be reusable from one simulation to the next.
- The interfaces between components and the generic infrastructure should insulate the components from changes in the underlying technology or approach used to construct the infrastructure.

In 1994, the DoD put together a Program Evaluation Team (PET) of computer scientists and engineers with a broad range of M&S experience. The PET studied the pros and cons of different architectures proposed by a set of contractor and academic groups. As no single proposal was satisfactory, the PET recommended that an Architecture Management Group (AMG) be formed to blend the best elements from the proposals and draft the definition and specifications for the HLA. In March 1995, the initial draft of the HLA was completed. The AMG initiated a diverse set of experimental prototypes to test the utility and completeness of the HLA. The lessons learned from the prototypes were used to refine the HLA specifications. In August 1996, the baseline definition of the HLA was established by the AMG.

In September 1996 the Undersecretary of Defense for Acquisition & Technology, Dr. Paul Kaminski, mandated that the HLA was to be the standard technical architecture for all DoD simulations. This policy decision essentially directed that all simulations currently planned or in use be modified to comply with the HLA requirements. The goal of this directive was to migrate all simulations to an HLA-compliant status by early 2001. Although the requirement to modify simulation code to meet HLA standards was viewed as onerous by some developers or sponsors, it has by and large been accepted and integrated into normal business processes within the defense M&S community. The downstream payoffs of composability, interoperability, and reusability serve as significant motivation for the conversion.

3. Description of the High Level Architecture

This section describes the HLA components, the supporting tools and services available, and the DoD efforts to standardize the HLA. The authors also include a brief discussion of known non-DoD uses of the HLA.

Before addressing the constituent elements that comprise the HLA, there are a number of terms that must be better defined. Nomenclature is as important in the simulation world as in other domains, and there are a number of key terms used repeatedly that readers must be familiar with to better understand the HLA approach. An enumeration of some of the more common and important terms is given below:

- **Federation**: a set of simulations, a common federation object model, and supporting RTI, that are used together to form a larger model or simulation
- **Federate**: a member of a federation; one simulation
  - Could represent one platform, like a cockpit simulator
  - Could represent an aggregate, like an entire national simulation of air traffic flow
- **Federation Execution**: a session of a federation executing together
- **Object**: An entity in the domain being simulated by a federation that
  - Is of interest to more than one federate
  - Is handled by the Runtime Infrastructure
- **Interaction**: a non-persistent, time-tagged event generated by one federate and received by others (through RTI)
- **Attribute**: A named datum (defined in Federation Object Model) associated with each instance of a class of objects
- **Parameter**: A named datum (defined in Federation Object Model) associated with each instance of a class of interactions
- **Simulation Object Model (SOM)**:
  - Describes objects, attributes and interactions in a particular simulation which *can* be used externally in a federation
- **Federation Object Model (FOM)**:
  - A description of all shared information (objects, attributes, and interactions) essential to a particular federation

Figure 1 is a depiction of an archetypical federation. The three circles indicate the three broad categories of federates which can be instantiated in myriad ways for a specific federation. A federation will typically use a subset of the RTI Services shown in the box.

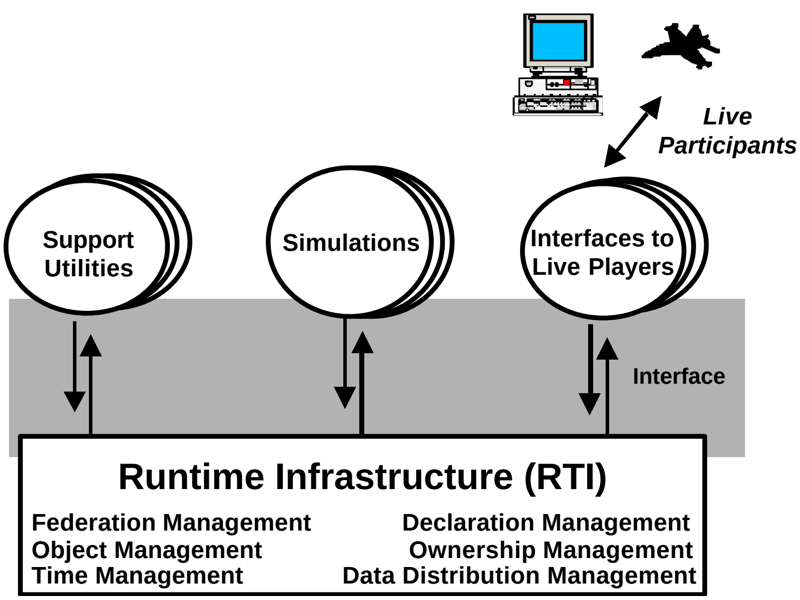

Figure 1. Archetypical Federation

3.1 HLA Components

The HLA as it exists today is essentially composed of three distinct elements:

- Ten **Rules** which define relationships among federation components
- An **Object Model Template** which specifies the form in which simulation elements are described
- An **Interface Specification** for a Runtime Infrastructure (RTI) which describes the way simulations interact during operation

These three elements map to the definition of a software architecture from Shaw and Garlan [cite] as follows: Abstractly, software architecture involves the description of elements from which systems are built, interactions among those elements, patterns that guide their composition, and constraints on those patterns.

HLA Rules

The HLA Rules are principles and conventions which must be followed to achieve proper interaction of federates during a federation execution. These describe the responsibilities of federates and federation designers. The HLA rules are divided into two groups consisting of 5 rules for HLA federations and 5 rules for HLA federates, as shown in Table 1.

Federation Rules	Federate Rules
Rule 1. Federations shall have a Federation Object Model (FOM), documented in accordance with the HLA Object Model Template (OMT).	Rule 6. Federates shall have an HLA Simulation Object Model (SOM), documented in accordance with the HLA OMT.
Rule 2. In a federation, all simulation-associated object instance representation shall be in the federates, not in the runtime infrastructure.	Rule 7. Federates shall be able to update and/or reflect any attributes, and send and/or receive interactions, as specified in their SOMs.
Rule 3. During a federation execution, all exchange of FOM data among federates shall occur via the RTI.	Rule 8. Federates shall be able to transfer and/or accept ownership of attributes dynamically during a federation execution, as specified in their SOMs.
Rule 4. During a federation execution, federates shall interact with the RTI in accordance with the HLA interface specification.	Rule 9. Federates shall be able to vary the conditions (e.g., thresholds) under which they provide updates of attributes, as specified in their SOMs.
Rule 5. During a federation execution, an instance attribute shall be owned by at most one federate at any time.	Rule 10. Federates shall be able to manage local time in a way that will allow them to coordinate data exchange with other members of a federation.

Table 1. HLA Rules

Object Model Template

The OMT is the prescribed common method for describing the entities to be simulated and interactions between entities in a federation. The OMT is an example of a meta-model, i.e., a way to describe models of information.

Within the HLA paradigm, object models can take one of two forms. A Federation Object Model (FOM) provides a specification of the exchange of public data among all of the participants in a HLA federation. An HLA Simulation Object Model (SOM), in contrast, provides a specification of the intrinsic capabilities that an individual simulation offers to federations. The HLA object model is key to achieving the stated goals of interoperability and reuse, and represents one of the fundamental tenets upon which the HLA has been defined.

Interface Specification

This is the specification of the interface between federates and the *Runtime Infrastructure* (RTI). The RTI is software that allows a federation to execute together. It is the interface between the RTI and federates that is standardized; an implementation of the RTI could take a variety of forms.

The RTI is software that conforms to the specification but is not itself part of the specification. It provides the software services, which are necessary to support an HLA compliant simulation.   These services fall into six categories: federation management, declaration management, object management, ownership management, time management, and data distribution management.

- *Federation Management:* Services to support the creation and deletion of federation executions, the membership of federation executions (i.e., federate join and federate resign), and federation-wide functions such as saves (checkpoints) and restores.

- *Declaration Management:* Services to allow a federate to declare its intent to produce (publish) and consume (subscribe to) data.

- *Object Management:* Services to support the actual exchange of data between producing and consuming federates.

- *Ownership Management:* Services to allow federates to transfer responsibility for selected attributes of objects to another federate.

- *Time Management:* Services to support the ordering of federation events in logical time between federates.

- *Data Distribution Management:* Services to allow federates to apply advanced techniques of data filtering to further refine the delivery of data to other federates.

3.2 Supporting Tools and Processes

Tools

It is critical to keep in mind that the HLA is an architecture, not software. However, to facilitate cost-effective implementation of the HLA, the Defense Modeling & Simulation Office has developed an initial suite of supporting software and is distributing it in the public domain. This software suite includes the following:

- HLA Runtime Infrastructure (RTI) software
- Object Model Development Tool
- Data Collector Tool
- Federation Management Tool
- Federation Verification Tool
- Federation Execution Planner's Workbook

The software and documentation is available from an on-line Software Distribution Center, accessible through the DMSO web site at http://hla.dmso.mil.

To foster the development of commercial software, all HLA specifications have been made public via the Internet, and HLA-based tools and development environments are already emerging in the commercial marketplace.

Processes

In addition to actual software, DMSO has developed a recommended practice for federation development, called the Federation Development and Execution Process (FEDEP). The FEDEP is a generalized process for building HLA federations; it is based on classic systems engineering practices. The FEDEP is a useful point of departure for project teams initiating an HLA federation project.

User Services

DMSO provides a number of services to the HLA user community including a robust up-to-date web site, a general on-line HLA Help Desk, an RTI Help Desk, and an HLA integrated training program. The training program consists to two types of training: lecture-style course modules offered in one and a half day packages called Regional Training Events and hands-on software development training with the RTI called the HLA Hands-On course. Information about the training events, the schedules, and the registration procedure is available through at the web address given above.

3.3 Standardization Efforts

The DoD has taken steps to gain recognition of the HLA as a standard by international and commercial standards organizations. In particular, the Object Management Group (OMG) has adopted the HLA Interface Specification as the OMG "Facility for Distributed Simulation Systems." OMG is a consortium of software vendors, users, government, and academia, that works to establish standards in distributed object computing.

Further, the Standards Activity of the Institute of Electrical and Electronics Engineers (IEEE) is currently reviewing and balloting the HLA Rules, OMT, and Interface Specification as the standard architecture for distributed simulation. Approval by the Standards Activity is expected by March 2000.

The extent to which the HLA has been embraced by these standards-making bodies is one indication of its general-purpose use. The HLA is *not* a DoD-specific technology.

4. Extensibility of the HLA Paradigm to Computational Economics

We believe that technical approaches and applications currently in use by the DoD M&S community can be transferred to the macroeconomic community to extend the state-of-the-practice in economic modeling. For example,

- The use of the HLA allows a macroeconomist to build up a complex simulation tool by joining together individual simulations that represent individual aspects of the overall economy. In other words, you no longer need one big monolithic simulation, built in a single programming language by a single development team.
- As the domain of computational economics grows and advances, a significant population of disparate models and simulations will develop. Many of these will actually be complementary, examining different aspects of a larger macroeconomic system. If the HLA were to become a *de facto* standard within the computational economics community, then the linking of "best of breed" simulations would become a reality, enabling significant capabilities.
- The use of entity-level modeling allows an economist to study the aggregate behavior of groups, communities, industry sectors, and nations by varying the behavior of the individuals that make up those units.
- The use of distributed simulation techniques allows an economist to achieve scalability by adding host computers to the hardware base, rather than moving to a larger, more expensive computer. [It is interesting to note that most economic modeling focuses on running simple agents on massively parallel processors. The DoD community has spent years pursuing a similar approach with little success.]

We base these premises on a number of assumptions, including

1. Economic systems can be modeled as the aggregation of a set of individual entities and their interactions.
2. The expertise on specific sectors of the economy and specific entities within the economy is dispersed among many organizations. These organizations do not normally work together to address larger, macroeconomic issues.

3. Different facets of the national economy have been simulated to some degree of fidelity by individual organizations, such as the Federal Reserve Board, the Council of Economic Advisers (CEA), or the National Bureau of Economic Research (NBER) .
4. There are many real-world results or circumstances that cannot be predicted easily or assessed using the state-of-the-practice modeling techniques.
5. The sheer size of the real world "system" compels the macroeconomic community to use models and computer-based simulations rather than live experiments.
6. The nature of the world economy is becoming more complex, thus putting an even greater burden on the equations and simple models that macroeconomists use to do their work.

It is worth noting that HLA compliance satisfies the most important condition for interoperability and reuse: a common, efficient technical means to join simulations together in federations, optionally including live players, and exchange information in a coherent manner. However, the HLA is not an interoperability "magic wand," that is, it will not automatically make every simulation suitable for federating with every other simulation nor guarantee a valid, meaningful exchange of information across the federation. For example, it would likely make little sense to federate a high-resolution agent-based financial market simulation with an individual consumer simulation; while in fact such a federation would be technically feasible, it would fail the relevancy test. Prudent, common sense planning is still required, but the HLA does provide the critical technical foundation for the required, a capability which represents a major advance in simulation technology.

HLA compliance delivers new functional capabilities and allows different organizations to produce/maintain a diverse set of products (e.g., simulations, live system interfaces, utilities, runtime infrastructures) which can be wisely used together in different combinations as user needs dictate. This yields reuse of individual products and allows simulations to bring in new capabilities without having to build them. This in turn equates to reductions in time, expense, and risk that justify the modest near-term costs of transitioning legacy systems to the HLA.

4.1 Current Approaches to Economic Modeling

We have undertaken an Internet survey of macroeconomic models and model research to see which might benefit from federation. Models are created for various purposes, forecasting, training through simulation, impact analysis, research on economic behavior, trade and balance of payments analysis, banking analysis, to name a few. Some that may have potential for federation benefits are bulleted below:

- Theoretical conceptual models - Most of the agent-based or object-oriented models seem to be in this category or the next. Clearly such models are a federation of agents or classes. Where a broader system is required for analysis, federation at the same or a higher level may be beneficial.
- Simulation models for training - For helping students understand how an economy works through hands on testing of policy instruments. For training purposes, it might be advantageous to begin with local or partial subsystems, and then progress to integrated or global models.
- Planning/forecasting models (data intensive) - In use in some developing countries, though traditional investment/sectoral planning is now giving way to creating an environment conducive to private sector growth through policy analysis and reform. Countries with highly open trade-oriented economies may require models that depend on the inclusion of other countries and regions. Development in small countries depends on the growth of export markets and tourism demand in larger countries. Large countries depend on each other and on the availability of primary products from small countries.
- Impact analysis models (also data intensive) – The impact of oil prices, military activity, aid flows, and economic reforms. As investigation of impacts focuses on subsystems, existing or new models could be added or deleted.
- National vs. Sub-National Models - Federating several national models may be useful, although there are also regional and global models out there. What may be more important is the potential for reuse and adaptation of existing models if a common operating environment is available.

- Federating national subsystem models – Banking system models, trade system models, etc. Banks, governmental agencies, and corporations have developed subsystem models. Creating links between them would have advantages analogous to trade and information links in real world systems.

4.2 Potential Benefits of HLA to Economic Modeling

Quantitative macroeconomic simulation models are designed to answer specific questions about the state of the economy to achieve particular goals of a sponsoring organization or research team. There are no standardized economic databases, sources of data, or data export features that are common to all models used today. Assuming that some workable strategy can be devised to federate existing models together into a single integrated network of models, what would the benefits be?

The authors believe that a number of important benefits can be achieved by federating several "as is" macroeconomic simulation models using the HLA approach. The primary benefit of federating several macroeconomic models is to provide a cross cutting view through an interacting network of models, which can broaden the view permitted by a single model and provide opportunities for modelers to gain new insights into economic phenomena. Other benefits arise from the implementation itself, namely, sharing a common set of data, creating meta-data that characterizes this broader view, and responding more rapidly to changing situations.

- <u>Broadening the view permitted by a single model</u> - No matter how all-encompassing, a single economic model that is developed and/or operated by a single organization or research group will have a particular focus that reflects the stated mission or purpose. Federating several models make it possible for the participating organizations to interact with each other, in a way that will enhance research perspectives through unprecedented opportunities for sharing data about the economy in "real time" from many sources. This can be accomplished in a way that preserves the distinctive character of each participating model federate.

- <u>Gain new insights into economic phenomena</u> - Federated models provide the capability to simulate interactions of greater complexity than previously possible. The simulation toolset enabled by HLA technology will increase the visibility into the dynamics of the economy at the macro level, and lead to avenues of enquiry that were previously not possible.

- <u>Sharing a common set of data</u> - When several different models use a common set of input data, there will be greater consistency in the results of computations; fewer data items to gather, track, and update; a wider consensus about the input values and assumptions; and a common understanding of what the inputs mean.

- <u>Creating meta-data that characterizes this broader view</u> - Summary level data describing the overall operational state, status of the inputs, status of the outputs, and other high level parameters can be extracted and tracked over time. This information is essential to understanding how well the federated system of models is working, particularly when simulating rapidly changing situations.

- <u>Responding more rapidly to changing situations</u> - The HLA infrastructure enables a robust simulation environment. It permits the rapid analysis of "what-if" questions, and can be used to explore sensitivity to changes when a number of federated models are simultaneously impacted by variations input data or outputs from other models.

- <u>Providing common operating environment for future economic model development</u> – Future sub-system modeling can be done with an awareness of the HLA infrastructure availability. This is likely to impact the design of future models. Existing models will have to be encapsulated or require a special interface, but new ones could be designed to plug into the common infrastructure.

4.3 Possible Strategy for Adopting the HLA Paradigm

Virtually none of the economic simulation models in use today were designed to be "federated" into an over-arching distributed simulation architecture. Despite the heterogeneous modeling approaches found among current economic models, a robust experience base and infrastructure already exists for applying HLA, and the current state-of-the-art makes it possible to integrate existing stand-alone computational economic models into a distributed network. If this were to be done, the candidate group of models to be federated must be carefully selected, so that the combined simulation environment produces distinct benefits that are greater that the sum of the individual models.

An alternative to federating whole models would be to use essentially the same data but translate or redesign the model to be more compatible with the HLA. If the models to be federated are relatively small, this may be easier than encapsulating or interfacing with an existing model. In other words, there may be alternative ways of migrating to a federation of simulations.

It is very important to establish the appropriate eligibility criteria for selecting candidate models before federation development takes place. These criteria should be based on the overall purpose of the distributed simulation task: What new questions will be answered or what special problems will be solved, by combining several economic models into an extended simulation environment? What advantage would federation have over creating an entirely new model (possibly using existing data) at a higher level?

After the purpose is established, it becomes necessary to identify the kind of data that must be passed between models. For each model to be federated, answers are needed to questions such as: What are the inputs, what are the outputs, and what data items must be shared?

Since all models in use today were not designed to be federated together using the HLA, it will be necessary to plan, design, and develop the appropriate modifications needed to interface each to the common infrastructure. Other implementation steps include federation planning, configuration of hardware, integration testing, and finally federation execution.

In the future, macroeconomic models can be designed and implemented with the HLA approach in mind. The design-to-HLA strategy can be achieved by incorporating standardized interfaces to the HLA and common data elements, and shared database designs, as common features of new macroeconomic simulation models.

5. Summary

This paper has addressed the High Level Architecture, the motivation behind its development, its basic precepts, and the potential for application in the computational economics community. As stated, the HLA is a simulation- and domain-independent approach that should work equally well in the economics domain as it has in the military domain. The HLA is also both technologically and procedurally mature, and is ready to be applied and adapted within simulation domains outside the strict confines of the military M&S community. Given the significant analogies between the economics and the combat simulation domains, it would seem that the field of computational economics is another logical candidate for the transition of the HLA to non-military applications.

We believe that it is feasible to select a group of existing macroeconomic models and create a federation of interacting models. By developing a framework for the use of ADS concepts, it is anticipated that the agent-based economics community will have the means to assess the benefits of such an approach. A first step would be to develop a prototype using existing or adapted subsystem models and test the potential benefits of this federation.

As the number of simulations achieving HLA compliance grows, the greater the opportunity for meaningful and productive interoperation. The adoption of the HLA as an IEEE specification is indicative of the acceptance of the approach within the broader simulation community; the computational economics community should leverage the experience and knowledge developed to date within the Department of Defense to extend and improve its own growing capabilities. Proponents or developers of simulations in this domain interested in learning more about the HLA are strongly encouraged to contact the authors or the Defense Modeling & Simulation Office.

6. References

A comprehensive HLA bibliography is available on-line at the McLeod Insitute for Simulation Science at Cal State-Chico, and can be found at http://www.ecst.csuchico.edu/~mcleod/. Further HLA documentation can be found at the Defense Modeling & Simulation Office's HLA Homepage, found at http://hla.dmso.mil/hla/papers. A list of the references cited in this paper is given below.

Dahmann, J. S., Fujimoto, R. M., Weatherly, Richard M. "The DoD High Level Architecture: An Update" *1998 Winter Simulation Conference*.

______, Kuhl, F., Weatherly, R. "Standards for Simulation: As Simple As Possible But Not Simpler. The High Level Architecture For Simulation" *Simulation* 71.6 (1998): 378-387.

McLeod Institute for Simulation Science. "Basic Concepts of the High Level Architecture (HLA)" http://www.ecst.csuchico.edu/~hla/LectureNotes/HLA_M1Part1.doc.

Shaw, Mary, and Garlan, David. *Software Architecture: Perspectives on Emerging Discipline*, Prentice-Hall, 1996.

7. Biographies

James A. Calpin is a Lead Staff Member in MITRE's Advanced Information Technology Division. He has a B.S. in Mathematics from the College of William & Mary, and M.S. in Systems Engineering from George Mason University, and is a doctoral candidate at The Institute of Public Policy at George Mason University. He has over 10 years experience in the defense analysis field, and is currently playing a key role in applying distributed simulation technologies to emerging analytical and experimentation-related requirements.

Marnie R. Salisbury is a Principal Engineer in MITRE's Advanced Information Technology Division. She has a B.S. in Computer Science from the Rochester Institute of Technology, and a M.S. in Systems Engineering from George Mason University. She has over 12 years experience in defense modeling and simulation, including 15 months as the Technical Director for the DARPA Synthetic Theater of War (STOW) program. She is currently the leader of the HLA cadre, a team assigned to provide intense technical support through partnerships with ongoing and emerging DoD federations.

John A. Vitkevich, Jr. is a Lead Economic/Business Analyst in MITRE's Economic & Decision Analysis Center (EDAC). He has a B.S. in Physics from Clarkson University, an A.M. in Physics from Harvard University, and a MBA from the Wharton School of the University of Pennsylvania. He has over 23 years experience at MITRE applying quantitative economic analysis, cost estimating, and modeling techniques to defense systems.

David R. Woodward is a Senior Economic/Business Analyst in MITRE's Economic & Decision Analysis Center (EDAC). He has a B.S. in Physics from the California Institute of Technology, and a Doctorate in Educational Planning and Economic Development from Harvard University. He has had 16 years experience working as an economic and financial analyst/planner in developing countries in the Pacific, Africa, and the Caribbean, and 10 years at experience at MITRE applying quantitative economic analysis techniques to a variety of MITRE initiatives. Recently Dr. Woodward took a one-year sabbatical/leave of absence from MITRE to join an economic reform project funded by the Asian Development Bank in Micronesia.