

# **Providing Scalable Simulations, Environments and Data Structures for Theatre of War Analysis and Training**

Dan M. Davis & Dr. Robert F. Lucas  
Information Sciences Institute  
University of Southern California  
4676 Admiralty Way, Suite 1001  
Marina del Rey, California 90230, USA  
3104488434 & 310 4489449  
{ ddavis & rflucas } @isi.edu

Garth D. Baer  
Student, School of Law  
The George Washington University  
2000 H St., N.W., Washington, D.C., 20052  
202 9946261  
gdbaer@gwu.edu

Dr. Thomas D. Gottschalk  
Senior Research Scientist and Lecturer in Physics  
Center for Advanced Computing Research and Dept of Physics  
California Institute of Technology  
1200 E. California Blvd, Pasadena, California, 91125, USA  
626 3976671  
tdg@cacr.caltech.edu

## **ABSTRACT**

There are growing imperatives for realistic, valid, and insightful distributed simulations using global-scale, high-resolution battlespace environments. These can be well served by advances in cluster computing and high performance networks, but only if appropriate communications strategies enable scalability and fault-tolerance. Lessons learned from experiences with U.S. Joint Forces Command's JSAF simulations will cover innovations, data management techniques, and future development goals.

## Introduction

Arguably, we are in the midst of the greatest transformation of the military services since 1949 when the current Department of Defense (DoD) was stood up. Subsequent laws, notably the Goldwater Nichols DoD Reorganization Act of 1986, laid the foundation, but it has been recent events and change in culture that is transforming the services from four quasi-independent organizations into a single integrated force. This transformation is not nearly complete, nor will it be in the near term; however, when combined with current world events and DOD transformation efforts, it does provide a unique moment in time to make a quantum leap in our ability to have the simulated world more accurately reflect that of the real world, and thus, make those simulations more useful to the warfighters immersed in their operational reality. This paper discusses several of the approaches we can take to support both the transformation and our warfighters.

## Objective

At the most basic level, the objective of this effort can be summarized by the following: "Provide a large-scale Contemporary Operating Environment (COE) simulation to provide real time, operationally relevant, Course of Action Analysis (COAA)".

As we dissect this goal into its component parts, we see several major elements come to the forefront:

- Dynamic environment modeling – The ability to model light, climate, wind, cultural and vegetation variations over time
- Individual, organizational, cultural modeling – The ability to have the entities in the simulation not only reflect the cultural reality and diversity of a given location, but also the variations of individuals within that population
- Infrastructure / effects-based model – The ability to model the economic, manufacturing, logistics, and utility networks and the cascading effects of their interconnectivities and interrelationships
- Weapon, communication, sensor system modeling – The ability to use accurate first order physics based models in a fully task loaded and interconnected environment
- Learning / evolutionary systems – The ability for friendly, adversarial, and neutral forces to learn and adapt over time
- Data management and mining – The ability to manage, analyze, and exploit the petabytes of data that will be both needed and generated by these kinds of simulations in relevant timescales

There is nothing on this list that has not been done either in isolation, on a small scale, or in a simplified version. What has not been done is the linking of all of these capabilities together at a high level of fidelity to provide the object system.

## Current System Limitations

This discussion begs the question: Why hasn't this been done? The simplistic answer is that it cannot be done on current systems, even on current high performance computers. In this section we will look at some of the reasons why.

### Stove-piped Platform Centric Acquisition

The vast majority of the legacy codes were developed by programs or services for their particular applications. As a result, many have a natural bias to their sponsor's figures of merit. It is only natural for a model developed by the air defense community to emphasize the importance and effectiveness of air defense over precision strike. Additionally, the concepts of information sharing and interconnectivity (embodied in such buzzwords as net centric, effects based operations) are fairly recent additions. As a result, the legacy models emphasize the physical platforms vice organizational and cognitive structures.

### Single Processor Lack of Computational Power

With the platform-centric approach, the majority of the models were designed to run on a uni-processor device. Many times, it was more than one platform per processor. For example a recent simulation system was required to model a "Battalion level exercise" on a single processor. This necessitated considerable simplification often resulting in the elimination of physical reality; for example, the lack of shadows (the use of variable light conditions is one of the key elements in small unit tactics). As a more esoteric example, the authors know of no simulation system that models nonhuman animals. Yet in real life, these animals are a fact of life and can have an effect on the operation.

#### Network of Workstations (NoW) Limitations

To address the computational limitations of uni-processors, the logical evolution was to the NoW model for distributed processing. This approach motivated the development of DIS and HLA message passing standards for distributed simulations. However, this development was not without limitations of its own. Notable among these was the need for a world consistency mechanism; that is, making sure everyone had a consistent view of the world. This added simulation induced communication overhead and level of complexity to the systems. Additionally, the delays and bottlenecks associated with network transport drove the architecture of the systems. Thus, the elements that interacted most often tended to be grouped together, ideally on a processing element. This reinforced the platform centric view of the world. There were a few notable exceptions to this model that were built along a functional basis—weapons and environmental servers come to mind—but these were exceptions to the rule. The federations of models, with different terrain models and processing algorithms, often had "fair fight" issues. One system had a simulation induced advantage over another one. That is not to say the processing model was bad. It was the best we had available at the time and it increased our ability to address large problems, albeit with increased level of complexity in terms of both management and computational infrastructure, both of which limit the scalability of the approach.

#### Clusters Move Bottleneck

The advent of the Linux cluster has simplified the management of events and reduced the executable version inconsistency problems associated with the NoW paradigm. The physical proximity of processors and the increased bandwidth of the networks internal to the cluster have reduced communication bottleneck, but not eliminated it. More importantly, clusters have not changed the fundamental structure of the models. Thus, the same bottlenecks and limitations still exist; only the relative size of the bottlenecks has been altered.

#### Co-Processors Not Flexible Enough

For those of us who have been in the field for some time, the advent of the graphics processing unit (GPU) with substantial processing capability is a case of déjà vu. We have seen the host/IG construct give way to first the integrated hostIG and later to the PC graphics solutions. Likewise, the math coprocessor is now a thing of the distant past. That is not to say that there are not advantages to using coprocessors like GPUs, particularly in math intensive processing problems, but this approach is not the solution to fundamental architectural limitations.

### **Characteristics of an HPCS**

The transition to commodity processor based HPC systems in the mid1990s led to a new generation of large-scale systems composed of SMP processing nodes. These systems both maximize peak performance while minimizing purchase price. The ubiquitous Beowulf PC cluster is the best example. Such systems have been successfully used to support a number of military and other applications, including FMS as described above.

Unfortunately, HPC systems composed of clusters of shared memory processors (SMPs) have distributed address spaces, which greatly complicates the human cost of developing software for them. Worse, the high-latency of accessing remote data on such systems precludes their effective use for many tightly-coupled problems of National import. Even where they have been successfully ported, HPC applications typically only achieve a small fraction of the peak performance their users had hoped for.

To overcome the above problems in today's HPC systems, DARPA and its government partners, including NSA, NRO, DOE, NASA, and NSF, have begun a new high productivity computing systems (HPCS) program. The goals of the HPCS program are extremely ambitious. While maintaining the steady growth in memory volume and peak performance that users have come to expect, HPCS systems also address some of the shortcomings that have limited the utility of SMP clusters for DoD applications. Now in its second phase, HPCS is working with Cray, IBM, and Sun, and all of their proposed HPCS systems will feature global address spaces, memory latency hiding mechanisms, and other novel hardware to make it possible for them to support a growing range of DoD users. HPCS systems will also come with new programming languages and tools that will make these features accessible to software developers. When deployed at the end of the decade, HPCS systems will make the world's most powerful computing platforms accessible to a broad range of DoD applications, something that hasn't happened since the heyday of the Cray vector mainframes, over a decade ago. These systems will run DoD applications more efficiently than today's SMP clusters, and it will be easier and more cost-effective to program them.

## **Application of FMS to HPCS**

Current military operations can no longer be expressed by the traditional combat models we have employed to date. The reason for this is threefold: lack of scale, lack of fidelity, and combat focus. Ironically, all three of these share the common root cause – lack of processing ability. The legacy codes in use have had to compromise in the above areas to fit with the distributed processing environments in which they were developed. At the same time the models have been proliferating, military operations have become significantly more integrated. Consequently, the gap between simulations and operations is increasing. While the monetary impacts of being able to determine the war fighting impacts of resource allocations is a major benefit of Forces Modeling and Simulation (FMS) on an HPCS level resource, the most significant benefit may become the closing of the gap between simulations and operations. More realistic training, experimentation, analysis, and planning will lead to reduction in casualties and an increase in mission effectiveness. With the complexity of the modern and future battle space, this can only be done on a HPCS class resource.

## **Path to Implementation**

There is now a confluence of events that provide a dramatic opportunity for the use of HPCS for the DoD. Investments by US JFCOM, PEO STRI, DARPA and HPCMO are coupling HPC resources with operational needs to support the radical transformation of the US military. HPCS level resources can provide radically new capabilities to the warfighter by combining HPCbased functional, physical, logical, and behavioral models of battlespace components and effects in human-in-the-loop (HITL) applications. But using HPCS resources for HITL simulations must be done in a disciplined manner. As shown in Figure 1, by teaming academia and industry, we will extend baseline operationally relevant models that provide the architecture, but do not represent the fidelity or have the scale needed for emulation of real world operation. By working with the government managers of these models, the technology insertion and transition path is assured. Figure 1 shows a logical path that takes advantage of current technology in preparation for the availability of HPCS class resources.

Figure 1. Implementation Path

## **Using the OOS Architecture as a Baseline**

Figure 2 illustrates the conceptual approach. By taking advantage of the convergence of the processing memory space capabilities of the HPCS resources, the component nature of emerging simulations, such as OneSAF Objective System (OOS), and the location transparency provided by the long haul networks, we propose to replace the selected component elements of OOS architecture with either highfidelity first order physics models or proxy interfaces to operational systems. In doing this, we are enhancing those areas that

traditionally have been most simplified by the computational and network limitations of the distributed processing model by including those elements most needed to emulate current military operations.

## **Enabled New Functionality**

New simulations, taking advantage of HPCS, should address the following areas to increase the fidelity of battlefield simulation:

- Communications – modeling both the electromagnetic spectrum and network protocols will provide the level of detail necessary to identify those dead spots and bottlenecks in the communications networks that hinder both friendly and opposition operations
- Environmental Fidelity – where improvements are needed in dynamic terrain, weather, physics-based models for localized events such as sand storms, plume dispersal and the electromagnetic spectrum
- Human-Like Behaviors – significant improvement is needed here and is central to adequate simulation of US and allied soldiers, civilians and terrorists in urban environments. This must include variation of activity and learning – neither capability exists in today's fielded simulations.
- Wide Area Physically Accurate Sensor Networks – expanded much beyond simple table / probabilistic based sensor model with a limited number of target / sensor pairs. Certainly, space-based and networked sensors need to scan the millions of entities needed to realistically populate urban environments.
- Data management – the current approach exhibits limited data collection and management capability, minimal real time analysis during events, and an almost complete absence of computational steering. Data sets are growing substantially, with a 100X increase in just two years and with all indications that it will get worse.
- Knowledge Extraction – we have little or no learning from the current operational or experimental events. The volume of data overwhelms existing analysis tools. There is a significant need for visualization, data mining, and automated extraction tools to make data approachable by analysts.
- Real Time Course of Action Analysis – through modeling the co-evolution of friendly and opposing force tactics, the on scene commanders should be able to rapidly determine the optimal tactics to achieve the mission with the minimal casualties and resource expenditures.

A simulated battlefield environment with significantly higher fidelity will allow the DoD to use M&S to consider co-evolution of materiel and doctrine in programs such as: Future Combat System (FCS), Joint Strike Fighter (JSF), and the general areas of DT&E and OT&E support. This would also allow the unit of employment to be part of a Corps level deployment in a populated and politically dynamic area. This type of environment includes 10k+ Friendly forces, 20k+ Regular and irregular opposition, and 1M+ Civilians.

## **Benefits of the Project to the DoD**

A dramatically improved simulation of the battlefield can affect the following:

### **Fidelity / Scale**

Given a large, shared-memory system, we can eliminate or drastically reduce the latency between the elements of the system. In so doing, we can dramatically and fundamentally change the partitioning of simulations. We can use a functional decomposition of the simulation vice the traditional entity based decomposition. This will both simplify the inter-process communication allowing for higher fidelity and enable greater consistency of the environment, as well as allow for lifecycle simulation to include long time-spanning functions such as logistics, culture, and effects based operations.

### Real Time COAA.

We can consider incorporating the real world sensors and C4I data in real time and using simulation as a basis for updating the planned scenario. We can run multiple excursions (50+) much faster than real time, analyze the ensemble of possible future outcomes, and generate expectations allowing for relevant COAA. We will help the commander react to changes of plans (e.g., Turkey denying the use of bases to launch war) Real time COAA will enable evaluation of new options and their implications (justintime logistics) and the exposition). Adequate fidelity will expose second and third order effects that can not be examined today. Figure 3 show the proposed planning cycle that takes advantage of the future integrated simulation capability.

Figure 3. RealTime COAA

### Integrated DT&E and T&E (Lifecycle Simulation)

Using complex cues provided by the simulation system, we can have the component / system / System of Systems under test embedded in a realistic and task loaded environment. For example, the weapon might drop clear of the plane when it is in straight, level, mid altitude, constrained speed flight. But what happens if somebody is shooting at the aircraft dropping the bomb? How do fear, evasive action, and nearby explosions change the characteristics? These are the types of questions that can only be answered in contextually rich environment or in the real world.

### Cultural, Cognitive, and Physical Diversity Modeling.

Most, if not all, of our current models suffer from the aggregation effect. Everybody and everything acts and behaves the same way. All M60 tanks are alike, as are all civilians, commanders, etc. With significantly more computational power, we can insert different models to more closely represent the variability of the battlespace.

One issue that is not often modeled is the impact of civilian actions and infrastructure on military operations. Consider a "simple" thing like traffic patterns. As those of us who travel frequently know, there are different rules of road in different cities. Driving in the Midwest does not prepare you to rent a car in Boston. It is even more different overseas. These kinds of issues are critical if we are to understand and deal effective with asymmetric operations in the urban centers. Extending this further, the management and load of the civilian infrastructure and the cultural implication is a significant concern to leaders in the post combat phases. While these models exist in isolation, they are not linked in such a way to understand the implications of different course of action.

The modeling of environment as a weapon, such as blowing locks, dikes, levies, dams, and mountain sides, setting fire to oil wells, grassland, and forests, and other intentional significant environmental damage is something that must be considered as we enter the era of post Geneva Convention warfare. The modeling of these effects is well known phenomena, but they are computationally too expensive to run in real time. As we move forward in FMS we need to have these kinds of high-fidelity models on call, and the FMS on HPC must be able to dynamically re-task processors to simulate different cultural and physical phenomena as needed. We might take a hit for a short period of time in some cultural effects in order to flood an area as a dam is breached and then turn our attention to the culture again or re-task a processor to simulate a grass fire in another region of the exercise.

## Summary

Forces Modeling and Simulation, an emerging HPC discipline, is an operationally relevant computational technology area that integrates and provides a DoD relevant context for the capabilities and advances of many other computational technology areas to provide the high-fidelity simulated battlefield. Simulation capabilities and results from areas as diverse as Signal and Image Processing (SIP), Computational Fluid Dynamics (CFD), Computational Electromagnetics and Acoustics (CEA), Computational Chemistry and Materials Science

(CCM), and Computational Structural Mechanics (CSM) are needed to provide the high-fidelity simulated battlefields for acquisition, training, doctrine development, and test and evaluation.

It is only in the operationally relevant context provided by FMS that the effects of acquisition decisions, training methodologies, and tactics development can be evaluated – in a benign environment. While a wing drop in an F18 can be studied in CFD simulation, FMS provides the framework to understand the quantity and tactics that make the F18 an effective combat platform. To provide the next level of insight and operationally relevant environment, to include cultural, environmental, and physical simulations of sufficient fidelity to the war fighter in the needed timeframes, a performance increase of several orders of magnitude is needed. HPCS computing capabilities can support a quantum leap in the fidelity of the simulated battlefield. High fidelity emulation of the physical and electromagnetic environments, behavioral, organizational, and cultural simulation, and real world sensors and events will provide real-time seeding for mission planning / rehearsal.

While the monetary impacts of being able to determine the war fighting impacts of resource allocations is a major benefit of FMS on an HPCS class resource, the most significant benefit comes from closing the gap between simulations and operations. More realistic training, experimentation, analysis, and planning will lead to reduction in casualties and an increase in mission effectiveness. With the complexity of the modern and future battlespace, this can only be accomplished on a HPCS class resource.