Petascale Computing for Military Operations

David R. Pratt, PhD
Phil Amburn, PhD
Science Applications International Corporation (SAIC)
12901 Science Drive
Orlando, FL 32817
407 243-3308
prattda@saic.com

Robert F. Lucas, PhD
Dan M. Davis
Information Sciences Institute
University of Southern California,
4676 Admiralty Drive
Suite 1001
Marina del Rey, CA 90292

Keywords:
Military Simulation, HPCS, Petascale

ABSTRACT: *The traditional combat models we have employed to date can no longer represent current military operations. The reasons for this are threefold: limited scale, insufficient fidelity, and inadequate combat focus. Ironically, all three share a common root cause, namely, lack of processing ability. The legacy codes in use have had to make compromises in order to operate within the distributed processing environments for which they were developed. At the same time the models have been proliferating, military operations have become significantly more integrated, thus increasing the gap between simulations and operations. While the reduced cost of determining the war fighting impacts of various resource allocations is a major benefit of Forces Modeling and Simulation (FMS) on a High Productivity Computer Systems (HPCS) level resource, the most significant benefit is closing the gap between simulations and operations. More realistic training, experimentation, analysis, and planning will lead to a 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.*

There is now a confluence of events that provide a dramatic opportunity for the use of HPCS for the Department of Defense (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 exciting new capabilities to the warfighter by combining HPC-based functional, physical, logical, and behavioral models of battlespace components and effects in a human-in-the-loop application. But it needs to be done in a disciplined manner. 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 high fidelity first order physics models or proxy interfaces to operational systems. In doing so, we are replacing the areas that traditionally have been most simplified by the computational and network limitations of the distributed processing model with those elements most needed to emulate current military operations.

1 Introduction

Arguably, we are in the mist 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 the change in culture that is transforming the services from four quasi-independent organizations into a single integrated force. This transformation is no where near complete, nor will it be in the near term. However, when combined with world events, 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, 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.

2 Objective

At the most basic level, the objective of this effort can be synopsized by ***"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 a several major elements come to the forefront:

- Dynamic environment modeling – The ability to model the variation of the light, climate, wind, cultural and vegetation changes 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 and exploit the terabytes 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 to link all of these capabilities together at a high level of fidelity to provide the object system.

3 Current System Limitations

This begs the question: why hasn't this been done? The simplistic answer is that it can't be done on current systems. In this section we will look at some of the reasons why.

3.1 Stovepiped 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 (buzzword: Net Centric, Effects Based Operations) are a fairly recent additions. As a result, the legacy models emphasize the physical platforms vice the organizational and cognitive structures.

3.2 Single processor lack of computational power

With the platform centric approach, the majority of the models were required to run on a uniprocessor 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. To do this there was a gross simplification or elimination of physical reality. An example of this is the lack of shadows. The use of variable light conditions is one of the key elements in small unit tactics. A more esoteric example is the authors know of no simulation that model non-human animals. Yet in real life, these animals are a fact of life and can have an effect on the operation.

3.3 Network of Workstations (NoW) limitations

To address the computational limitations of the uniprocessors, the logical evolution was to the NoW model of distributed processing. This evolved into the DIS and HLA communities. However, this development was not with out limitations of its own. Notable among these was need for a world consistency mechanism, making sure every body had a consistent view of the world. This added a simulation induced 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 functionality basis, weapons and environmental servers come to mind. But there 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 an increased the ability to address large problems, abet with increased level of complexity in terms of both management and computational infrastructure, both of which limit the scalability of the approach.

3.4 Clusters move bottleneck

The advent of the Linux cluster has simplified the management of events and reduces the versioning problems associated with the NoW paradigm. Event the increased bandwidth of the networks internal to the cluster has reduced the latency, but not eliminated it. More importantly, it has not changed the fundamental structure of the models. Thus, many of the same bottlenecks and limitations still exist. Just the relative size of the bottlenecks has been altered.

3.5 Co-processors not flexible enough

For those of us who have been in the field for long enough, the advent of the Graphics Processing Unit (GPU) is a case of déjà vu all over again. We have seen the Host/IG construct give way to first the integrated Host-IG and later to the PC graphics solutions. Likewise, the math co-processor is a thing of the distant past. That is not to say that there are not advantages to using co-processors like GPUs, particularly in the math intensive processing problems. Rather, it is not the solution to fundamental architectural limitations.

4 Characteristics of an HPCS

<<Bob / Dan, I need your help on this>>

5 Application of FMS to HPCS

Current military operations can no longer be repressed 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. Thus, increasing the gap between simulations and operations. 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 is by 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 battle space, this can only be done on a HPCS class resource.

6 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 HPC-based functional, physical, logical, and behavioral models of battlespace components and effects in a HITL application. But it needs to do it is a disciplined manner. As shown in the figure, but 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. The figure below shows a

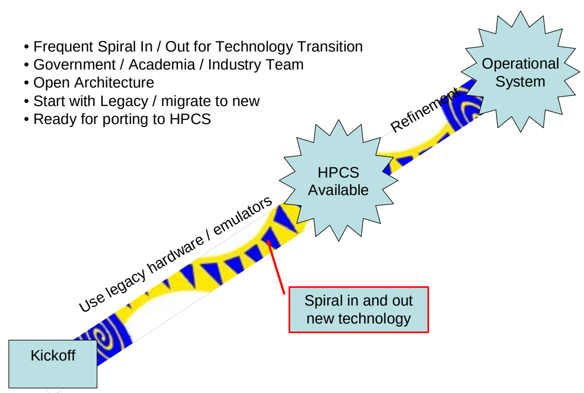

Figure 1. Implementation Path

7 Using the OOS Architecture as a baseline

Conceptually, the approach is shown in the figure below. 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 high fidelity first order physics models or proxy interfaces to operational systems. In doing this, we are replacing the areas that traditionally have been most simplified by the computational and network limitations of the distributed processing model with those elements most needed to emulate current military operations.

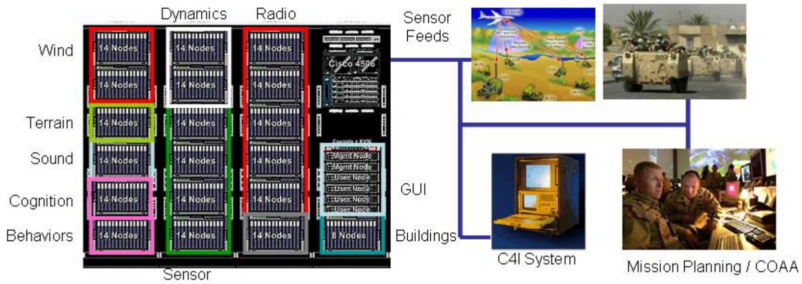

Figure 2. Notional System Architecture

8 Enabled new functionality

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

- Communications – modeling the electromagnetic spectrum and network protocols will provide the level of detail to identify dead spots and the bottlenecks in the communications networks that hinder 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 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 millions of entities needed to realistically populate urban environment
- Data Management – where the current approach exhibits a constrained data collection and management capability, with very limited / non existent during event analysis and limited 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 – Thought the co-evolution of friendly and opposing force tactics, the on scene commanders will 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: FCS, 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.

9 Benefits of the project to the DoD

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

9.1 Fidelity / Scale.

Given a large shared memory system we can eliminate / drastically reduce the latency between the elements of the system. In doing this we can partition the simulation differently. We can have 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 greater consistency of the environment. This will allow for lifecycle simulation to include the inclusion of long time spanning functions such as logistics, cultural, and Effects Based Operations.

9.2 Real-time COAA.

Incorporating the real world sensor and C4I data in real time and used as a basis for updating the planned scenario, we can run multiple excursions (50+) much faster than real time to provide a more relevant COAA. Change of plans (Turkey denying the use of based to launch war (time and location)) options and implications (just in time logistics) - exposition of second and third order effects.

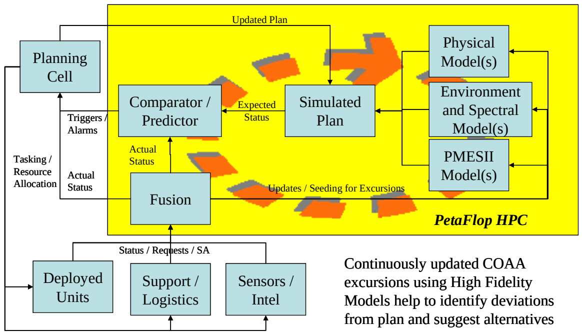

Figure 3. Real-Time COAA

9.3 Integrated T&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 straight, level, mid altitude, constraint speed flight. But what happens if somebody is shooting at the guy dropping the bomb? How does the fact that the pucker factor has been maxed change the characteristics?

9.4 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 all civilians, commanders etc. With significantly more computational power we can insert different model to more closely represent the variability of the battlespace.

10 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 SIP, CFD, CEA, CCM, and 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 that FMS provides that the effects of acquisition decisions, training methodologies and tactics development can be evaluated -- in a benign environment. While a wing drop in an F-18 can be studied in CFD simulation, FMS provides the framework to understand the quantity and tactics that make the F-18 an effective combat platform. To provide the next level of insight and operationally relevant environment, to include cultural, environmental, and physical of sufficient fidelity to the war fighter in the needed timeframes, several orders of magnitude increase in performance is needed.

Petaflop computing capabilities can support a quantum leap in the fidelity of the simulated battlefield with high fidelity emulation of the physical and electromagnetic environments, with behavioral, organizational, and cultural simulation and real world sensors and events to 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 a FMS on an HPCS level resource program, the most significant benefit is by 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 battle space, this can only be done on a HPCS class resource.

11 Acknowledgement

"This publication was made possible through support provided by DoD HPCMP PET activities through Mississippi State University under contract No. GS04T01BFC0060. The opinions expressed herein are those of the author(s) and do not necessarily reflect the views of the DoD or Mississippi State University."

12 Author's Biographies

DAVID R. PRATT is currently the Chief Scientist (Fellow) for SAIC's Training and Simulation Solution Business unit. As a vice president for technology, his responsibilities include developing and fostering leading- edge information technology and M&S technologies. He also serves as the Forces Modeling and Simulation point of contact for DoD's High Performance Computing Modernization Program (HPCMP). He received a Master of Science degree and a Ph.D. in computer science from the Naval Postgraduate School and a bachelor of science in electrical engineering, Duke University.

PHILIP AMBURN currently works for Science Applications International Corporation (SAIC) as the Programming Environment and Training On Site for

Forces Modeling and Simulation, with an office at Wright-Patterson Air Force Base (AFB), Ohio. His research interests include constructive and virtual simulation, interactive three-dimensional (3-D) graphics and visualization. He retired from the U.S. Air Force in the rank of lieutenant colonel. Amburn received a bachelor of science degree in physics from Kansas State Teachers College; a master of science degree in computer science from the Air Force Institute of Technology; and a Ph.D. in computer science from the University of North Carolina, Chapel Hill.

ROBERT F. LUCAS is the Director of the Computational Sciences Division of the University of Southern California's Information Sciences Institute (ISI). There he manages research in computer architecture, VLSI, compilers and other software tools. He has been the principal investigator on the JESPP project since its inception in the Spring of 2002. Prior to joining ISI, he was the Head of the High Performance Computing Research Department for the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory and the Deputy Director of DARPA's Information Technology Office. Dr. Lucas received his BS, MS, and PhD degrees in Electrical Engineering from Stanford University in 1980, 1983, and 1988 respectively.

DAN M. DAVIS is the Director, JESPP Project, Information Sciences Institute (ISI), University of Southern California, and has been active in large-scale distributed simulations for the DoD. While he was the Assistant Director of the Center for Advanced Computing Research at the Caltech, he managed Synthetic Forces Express, a multi-year simulation project. He has served as the Chairman of the Coalition of Academic Supercomputing Centers and the Coalition for Academic Scientific Computation. He was part of the University of Hawai‘i team that won the Maui High Performance Computing Center contract in May of 2001. He received a B.A. and a J.D., both from the University of Colorado in Boulder.