The Future Uses for the GPGPU-Enhanced Cluster at JFCOM

Dan M. Davis, Robert F. Lucas, Gene Wagenbreth, and Ke-Thia Yao
Information Sciences Institute, Viterbi School of Engineering, University of Southern California
{ ddavis, rflucas, genew, kyao) @isi.edu

David W. Roberts and Charles Brewton
SimIS, Inc.
{ david.roberts, chuck }@simisinc.com

Abstract

One of the major concerns attached to the Dedicated High performance computing Project Investments (DHPIs) is the durability of the need for which they are awarded. Preferably, such need will long outlast the nominal observation period by HPCMP personnel of the use of the DHPI assets. The authors present the case of the 2007 award of the 256-Node, GPGPU-Enhanced Linux Cluster Joshua, for which they are partly responsible, at least in terms of research agenda. JFCOM has had a continuing need of High Performance Computing (HPC) since the inception of its Joint Concept Development and Experimentation Directorate (J9) and has been the beneficiary of support and awards from HPCMP. J9 uses battlespace simulation to fulfill its mission: the development of emerging joint concepts, the conduct of joint experimentation, and the coordination of DoD experimentation efforts in order to provide joint capabilities. The award of Joshua began with the configuration of the system and finally resulted in the transfer of that asset to JFCOM last year. During the waning days of 2009, J9 underwent a significant change in emphasis at the direction of Joint Forces Commander and the Director of Joint Experimentation, so the concerns set forth above were reawakened, at least in the minds of the authors. This paper sets out what these new directives and constraints were, discusses how the new needs differed from the old ones, analyzes how this impacted HPC requirements, presents how the "open and balanced" design of the system made it a effective tool for the new simulation emphasis and records how effectively Joshua is being used today. The authors conclude that this particular DHPI award is still an invaluable asset and hold that their approach to ensuring that it maintains its utility to the warfighter is applicable in other settings, by other HPCMP users.

1. Introduction

Mandates for DoD research closely follow the vagaries international events. The rapidly changing geopolitical environment imposes a rapidly changing research agenda. (CRA, 2010). The High Performance Computing (HPC) professionals are in the forefront of these groups being impacted by those changes. They are tasked with projecting the future of US weapons, systems, tactics and strategies. One of the principle vehicles for this capability is the High Performance Computing Modernization Program (Henry *et al*., 2003). This program has two major equipment acquisition initiatives: a set of general-use supercomputing centers for shared use by communities of researchers and a smaller sub-set of dedicated platforms for special programs that cannot effectively use the general purpose centers. To some degree, both initiatives face the same issue *vis-à-vis* the dynamics of DoD research directions: How should one structure acquisitions so as to optimize utility in the face of an uncertain future? This paper addresses the conceptualization, design, acquisition, use and subsequent change in mandate for one of the dedicated systems: U.S. Joint Forces Command's *Joshua*.

1.1. Joint Forces Command Mission and Requirements

The Joint Forces Command (JFCOM) has the mission of leading the transformation of the defense establishment of the United States and to enable the U.S. to exert broad-spectrum dominance as described in Joint Vision 2010 (CJCS, 1996) and 2020 (CJCS, 2000). The Joint Concept Development and Experimentation Directorate, J9, is JFCOM's research arm. Having a research activity lodged within an operation command is a unique situation. That leads to experiments in which warfighters in uniform are staffing the consoles during interactive, HPC-supported simulations.

Experiments to model and simulate the complexities of urban warfare use well-validated entity-level simulations, *e.g.* Joint Semi-Automated Forces (JSAF) and the Simulation of the Location and Attack of Mobile Enemy Missiles (SLAMEM). These need to be run at a scale and resolution adequate for modeling all the convolutions of urban combat.

There is a long lineage of entity-level synthetic battlefield codes. They consist of representations of terrain that are populated with intelligent-agents representing: friendly forces, enemy forces and civilian groups. Experience has shown that to produce their behaviors in at timely manner, significant compute resources are required (Messina, 1997). Much of this is true because a major computational load is imposed in the performance of line-of-sight calculations between the entities. The inherently onerous "n-squared" growth-characteristics of an all-see-all program have been identified previously (Brunett, 1998). If several thousand entities need to interact with each other in an urban setting with vegetation and buildings obscuring the lines of sight, inter-node communications chaos and failure are often observed. This situation has been successfully ameliorated, but only partially so, by the use of an innovative interest-managed communication's architecture (Barrett, 2004).

Because of this, JFCOM required an enhanced Linux cluster of adequate size, power, and configuration to support larger and more sophisticated simulations, especially when there is a requirement for more than 2,000,000 entities within high-resolution insets on a global-scale terrain database. The cluster has been used occasionally to interact with live exercises, but more often has been engaged interactively with users and experimenters while generating only virtual or constructive simulations. (Ceranowicz, 2005) It had to be robust to reliably support hundreds of personnel committed to the experiments and it had to be scalable to easily handle small activities as well as larger global-scale experiments with hundreds of live participants, many distributed trans-continentally, as shown in Figure 1 below.

1.2. Joint Futures Lab (JFL)

Creating a standing experimentation environment that can respond immediately to DoD time-critical needs for analysis is one of the major, but difficult to achieve, goals of the Joint Futures Lab in Suffolk. The JFL must operate in a distributed fashion over the Defense Research and Engineering Network (DREN), at a scale and level of resolution that allowed JFCOM and its partners to conduct experimentation on issues of concern to combatant commanders, who participated in the experiments themselves. This provided the requisite simulation federations, software, and networks, joined into one common computer infrastructure that was supporting experiments. Quantitative and qualitative analysis, flexible plug-and-play standards, and the opportunity for diverse organizations to participate in experiments were all mandated by DoD needs, and those general mandates remain.

1.3. Joint Advanced Training and Tactics Laboratory (JATTL)

Supporting mission rehearsal, training, operational testing, and analysis remain the goals of JATTL. The principle concerns of the JATTL were developing technologies that support the pre-computed products required for joint training and mission rehearsal. One of the new concepts, the Joint Rapid Distributed Database Development Capability, as well as others, were seen to be required during its execution. The latter include phenomenology such as environment, cultural assets, civilian populations, and other effects necessary to represent real operations. The JATTL was connected via both DREN and the National Lambda Rail (NLR) to over thirty Joint National Training Capability sites nationally.

1.4. JFCOM's JESPP

To enable JFCOM to meet these goals, a J9 team designed and developed a scalable simulation code that has been shown capable of modeling more than 1,000,000 entities. This effort is known as the Joint Experimentation on Scalable Parallel Processors (JESPP) project (Lucas, 2003.) This work was begun under an earlier DARPA/HPCMP project named SF Express. (Messina, 1997) The early JESPP experiments on the University of Southern California Linux cluster showed that the code was scalable well beyond the 1,000,000 entities actually simulated, given the availability of enough nodes (Wagenbreth, 2005).

Early on in the JESPP Project, JSAF, in its current instantiation, was successfully fielded and operated using JFCOM's HPCMP-provided compute assets hosted at ASC-MSRC, Wright Patterson AFB, and at the Maui High Performance Computing Center (MHPCC) in Hawai'i. The J9 team had been able to make those dedicated systems suitable and reliable for day-to-day use, both unclassified and classified. This prevented the disruption of HPCMP batch-scheduled jobs at the

Major Shared Resource Centers (MSRCs). The dedicated compute power provided additionally allows for the easy identification, collection, and analysis of the voluminous data from these experiments, enabled by the work of Dr. Ke-Thia Yao and his team (Yao, 2005).

A typical experiment would find the JFCOM personnel in Suffolk Virginia interfacing with a “Red Team” in Fort Belvoir Virginia, a civilian control group at SPAWAR San Diego, California, participants at Fort Knox Kentucky and Fort Leavenworth Kansas, all supported by the clusters on Maui and in Ohio. The use of interest-managed routers on the network has been successful in reducing inter-site traffic to low levels.

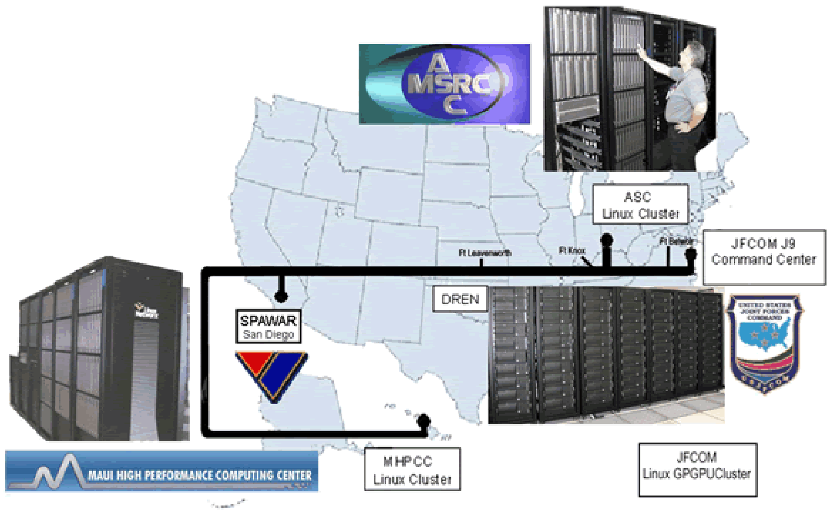

Figure 1

JFCOM’s Experimentation network: MHPCC, SPAWAR, ASC-MSRC, TEC/IDA and JFCOM

Even using these powerful computers, the JFCOM experimenters were constrained in a number of dimensions, *e.g.* number of entities, sophistication of behaviors, realism of various environmental phenomenology, etc. While the scalability of the code would have made the use of larger clusters feasible, a more effective, efficient, economical and elegant solution was sought. The end of the Moore’s Law speed and power improvements decreed that eventually JFCOM would need to seek new ways to deliver their services to increasingly sophisticated users with increasingly complex and geographically diverse conflict environments. One approach is the use of more powerful clusters, especially a GPU-enhanced cluster like *Joshua*.

1.5. Return on Investment

While precise quantification of the return on investment of this DHPI cluster may be elusive, it is possible to consider the cost savings of the JFCOM activities, which are, in turn, enabled by the HPCC assets via the DHPI program. In addition, there are incalculable benefits from battlespace simulation when one looks at the need to understand urban conflict, while realizing the impossibility of conducting full-scale military exercises in downtown U.S. cities. The same holds true, perhaps even more so, when the reaction to the loss of life of team members and allied forces needs to be felt in as realistic way as possible, while not even risking the loss of life and injury actually experienced in live combat military exercises.

The authors have seen a range of analyses, the cost of a “soldier day” in the field in domestic training has been estimated as high as $2,100. Using standard staffing tables, a Company exercise would costs on the order of $210K per day, a battalion $1M per day and a Division $20M per day. JFCOM commonly conducts exercises that are five to ten days in these sizes. When engaged in a virtual exercise, this is done with something on the order of one percent of the military personnel who would otherwise be engaged if the exercise were live. Not uncommonly JFCOM simulates large engagements in the constructive mode, utilizing only a few military Subject Matter Experts (SMEs). Clearly, one of the more difficult issues is to parse out how much is actually saved. The personnel who would be used in the live exercise are already incurring some of the costs. No matter how effective a simulation, some live training is needed. On the other hand, tank crews returning from the Gulf reported that the best training they had was the simulator training in the tank mock-ups. Further, putting a value on the cost to society of an accidental death during a live exercise, on an injury to a soldier or even on the wear and tear on very expensive military assets, is really beyond the scope of this paper, yet should resonate with all engaged in this analysis. As new threats emerge, as even more sophisticated equipment needs evaluation, as public sensitivities about military activities increase, even more burden will be born by the simulation community, JFCOM chief amongst them. The technology being researched, developed and implemented here will play a vital role in enabling that capability.

1.6. Broader Impacts for the HPCMP Community

Forces Modeling and Simulation (FMS) is a field often driving new computer science due to its virtually unlimited needs and open-ended requirements. Nothing short of a fully realistic experience, resting on a global model to obviate the edge effects of a restricted synthetic environment is sought. This field is unique to the DoD, compared to many of the other standard science disciplines, *e.g*. Computational Fluid Dynamics (CFD) and Weather. In a similar way, interactive computing is a new frontier being explored by the JESPP segment of FMS, coordinating with a few other user groups. Along these lines, the newly enhanced Linux Cluster capability has provided significant synergistic possibilities with other computational areas such as signals processing, visualization, advanced numerical analysis techniques, weather modeling and other disciplines or computational sciences such as SIP, CFD, and CSM.

2. Objective

Research objectives for this project were to develop a platform and system administrator cadre to provide 24x7x365 enhanced, distributed and scalable compute resources to enable joint warfighters at JFCOM as well as its partners, both U.S. Military Services and International Allies and to advance the technology of GPGPU utilization within the DoD. These enabled the users to develop, explore, test, and validate 21st century battlespace concepts in JFCOM J9's JFL. The specific immediate goal was to enhance global-scale, computer-generated support for experimentation by sustaining more than 2,000,000 entities on appropriate terrain, along with valid phenomenology.

3. Methodology

The use of existing DOD simulation codes on advanced Linux clusters operated by JFCOM was the implementation method chosen as being most efficient. This effort supplanted the previous JFCOM J9 DC clusters with a new cluster enhanced with 64-bit CPUs and nVidia 8800 graphics processing units (GPUs). Further, the authors were able to modify a few legacy codes. As noted above, the initial driver for the FMS use of accelerator-enhanced nodes was principally the faster processing of line-of-sight calculations. Envisioning other acceleration targets is easy: physics-based phenomenology, CFD plume dispersion, computational atmospheric chemistry, data analysis, *etc*. The computer was given the *Joshua*, the Hebrew commander.

Figure 2
Joshua

The first experiments were conducted on a smaller code set, to facilitate the programming and accelerate the experimentation. An arithmetic kernel from an MCAE "crash code" (Diniz, 2004) was used as vehicle for a basic "toy" problem. This early assessment of GPU acceleration focused on a subset of the large space of numerical algorithms, factoring large sparse symmetric indefinite matrices. Such problems often arise in Mechanical Computer Aided Engineering (MCAE) applications. It made use of the SGEMM (Single precision GEneral Matrix Multiply) algorithm (Whaley, 1998) from the BLAS (Basic Linear Algebra Subprograms) routines (Dongarra, 1993)
.

As an accelerator for computational hurdles such as sparse matrix factorization, the GPU was seen as a very attractive candidate. Previous generations of accelerators, such as those designed by Floating Point Systems (Charlesworth 1986) were for the relatively small market of scientific and engineering applications. Contrast this with GPUs that are designed to improve the end-user experience in mass-market arenas such as gaming. Meaningful speed-up was observed, on the order of an overall acceleration of 2X, using the GPUs. It was determined that the data transfer and interaction between the host and the GPU had to be reduced to an acceptable minimum.

4. Results

While great strides were made in the number of entities to be simulated, the research implications of this achievement, albeit expected, were a dramatic reaffirmation of the value of HPCs. Most notably, on 14 December 2007, the Joint Forces Command Personnel, under the leadership of Rich Williams simulated a full ten million CultureSAF entities on the Baghdad terrain database. This was accomplished as is visualized in Figure 3 below:

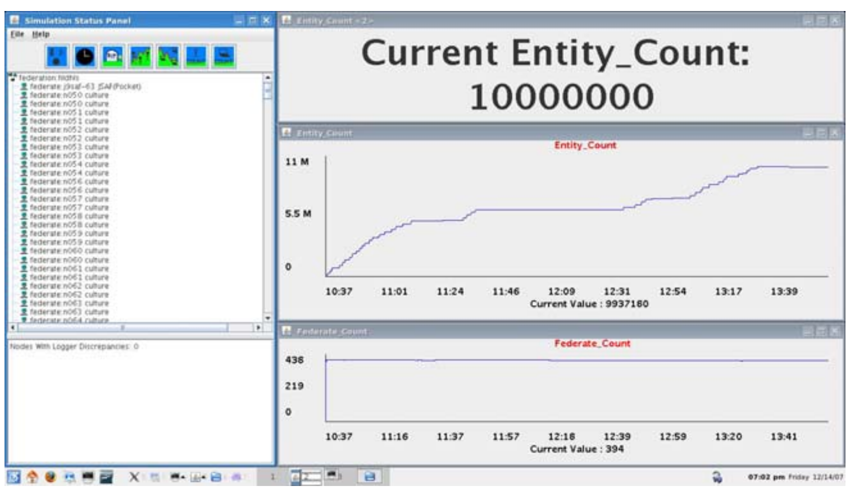

Figure 3
Graph of Entity Count Grow by Time on 14 Dec 07

This achievement had meaning in terms of the proposal that had been prepared by ISI and submitted to HPCMP by JFCOM, in which a major goal was the accomplishment of at least two million entities. A few early frustrations occurred due to the machine's not being configured for some of the experiments desired to help characterize the cluster's capabilities. Once again, the ten million entity run made the JFCOM, AFRL and ISI the largest agent-based, SAF implementation known, by at least two orders of magnitude, a size necessary to easily simulate any urban area with O(5M) inhabitants.

A summary of the results in the Mesh-Router tests showed that there were no problems with the mesh-router architecture or implementation. JFCOM operators and analysts confirmed that observed performance of the mesh-router was similar to performance of the tree router. Scalability was the desired parameter, so performance improvement was not expected nor required for success. There were no crashes or anomalies, a most critical factor.

The first effort was to create a baseline with the tree router which the mesh-router would match. The mesh-routers were incorporated in the latest RTI version. It typically took a long time to bring up a stable tree router configuration. Some of the problems the ISI team encountered were as follows:

- network problems
- clock skew between J9 and MHPCC inhibited Kerberos authentication
- MARCI problems with Runtime Initialization Data (RID) file and command line arguments that changed between RTI versions
- simulations that died immediately due to scenario input discrepancies
- cross-country network performance that was poor. This was solved by changing TCP to UDP.
- problem compiling loggers with new RTI. (an ongoing problem in this work was the use of upgraded compilers would fail until the code was re-written)
- applications were subscribing to too many entities, causing network problems.

USC generated a new connectivity map and mesh map by hand. JFCOM replaced the tree routers on Koa with mesh-routers using the hand created maps. Map generation for mesh-routers needed to be implemented in MARCI. JFCOM put down 100K culture entities and the system was stable and was left to run over several nights to test stability.

While the direct service to the warfighter is one of the hallmarks of HPCMP's initiatives for JFCOM, broader service to the DoD research community is also part of the overall set of goals. The uniformed service member sitting at a console linked to distributed and dedicated HPCMP assets is one recipient of the value of HPC, but another set is the researcher facing similar problems in their own environment as they serve other warfighters.

Some of the successful results of the distributed and dedicated HPCMP assets produced by the JFCOM team are the informational and instructional aspects of this work. It should be noted that, since the inception of the *Joshua* effort, the ISI/Caltech/JFCOM team has published thirteen papers, almost all at competitive, peer-reviewed conferences and journals. This represents both a recognition of the value of the work and a successful outreach effort to highlight the value of supercomputing and of the HPCMP. Further, the team was enabled through the HPCMP PET program to support the instructional efforts as well. Three courses were taught to DoD computer programmers, one each in 2007, 2008 and 2009. At sites in Suffolk Virginia, Marina del Rey California and San Diego California, nearly 80 programmers attended a "user-oriented" practicum, led by Gene Wagenbreth of ISI. These sessions received some of the highest student evaluations seen by the PET program personnel.

Figure 4
Course in GPGPU Programming - JFCOM, Oct 2007

5. History of HPCMP at JFCOM

The current USC/Caltech team has been working with HPCMP for more than a decade now (Messina, 1997). They and JFCOM have been one of the prime exemplars of the power HPC brings to Forces Modeling and Simulation. Varying machines and several centers have participated in this long effort. The insights provided have been lauded by General Officers and other warfighters. This mode of operations required an entirely new paradigm for computer use. These experiments were conducted by a cadre of several tens of operators, with literally hundreds of military participants and at half a dozen different locations, all distributed trans-continentally. Clearly this activity would not admit of batch operations and the interactive nature of the scenarios required low latencies, *i.e.* < 500 milliseconds.

An example event would be one in which the TEC site at Fort Belvoir, Virginia, had 30+ workstations and Saber, a quad-CPU machine with four TeraBytes of disk space that were used for after-event storage. The SPAWAR site at San Diego, California, had 20+ workstations. The J9 Joint Futures Lab at Suffolk, Virginia, had 50+ workstations and a 16-node mini-cluster. The ASC Wright Patterson Air Force Base at Dayton, Ohio, had the *Glenn* cluster with 128 dual CPU nodes. The MHPCC site at Maui, Hawaii, had the *Koa* cluster with 128 dual CPU nodes. After the acceptance of *Joshua*, the GPGPU-enhanced cluster in Suffolk, that machine joined the JFCOM experimentation network.

These experiments usually ran five days a week, ten hours a day. Simulators might run all night, but with little activity and usually with logging disabled. Depending on availability and requirements, some combination of *Glenn, Koa* and *Joshua* were used. Up to two hundred thousand clutter entities were simulated on the large clusters. (In this simulation, civilian entities are termed clutter, in that they serve to mask military entities.) Several thousand non-clutter entities were simulated on the other sites. A single node on the large clusters simulated 1000-2000 clutter entities.

Figure 5

DC Clusters *Glenn* at MSRC ASC and *Koa* at MHPCC

6. Changing Needs of the DoD

When the authors started work in this area, the major concern was the flood of Soviet tanks through the Fulda Gap in Germany (Messina, 1997). Shortly after that, it became apparent that there was a change in the nature of the threat to smaller, more urban, more constrained conflict (Boot, 2003). These wars were characterized with small but agonizing losses for the U.S., a heightened scrutiny of warfighter actions, and a critical public.

Table 1. Wars or Actions the U..S. Has Fought in Last Two Decades

Year	War/Action	Troops Committed	Duration	U.S. Casualties
1941	WW II	16,100,000	5 Years	416,000
1950	Korea	1,800,000	4 Years	40,000
1965	Vietnam	3,400,000	10 Years	55,000
1989	Panama	28,000	1 Month	23
1999	Desert Storm	500,000	< 1 Year	294
1995 - Present	Bosnia	10,000	15 Yeas & Continuing	85
2001 - Present	Afghanistan	60,000	8 Years & Continuing	951
2003 - Present	Iraq	250,000	7 Years & Continuing	4,400

That changed caused the DoD to refocus its battlespace simulations, leaving terrain models of the fields of Europe and moving to the urban areas of the Middle East, while abandoning hundreds of thousands of combatant agent-based-models for millions of civilian inhabitants surrounding a few hundred combatants. The basic concept remained the same: Create the battlespace of the future in which a warfighter can realistically engage the enemy with capabilities that are not yet available for test or in areas prohibitively expensive to recreate in a live simulation, .e.g. an entire city populated with demographically correct inhabitants.

7. New Mandates for JFCOM

None of this has taken place in a governmental vacuum. Since the HPC-for-battlespace simulation concept was born in the early 1990's, there have been three different U.S. Presidents and Congress has changed hands several times. Each group brings with it new emphases and that impact spending decisions. Also, external events drive other changes in expenditure and different mandates.

In the case of JFCOM, there are also periodic changes in command which impact the types of research emphasized and funded. This impact has recently been felt with a new initiative to make J9 more responsive to expressed request from other Operational Commands, albeit, less open to creating new technologies for others to use, before they ask for them. This has had the effect of curtailing the range of speculative or entrepreneurial research. The question is raised: How well did the conceptualization of the current DHPI platform anticipate the possible need to accede to such a change in direction?

8. Applicability of DC and DHPI Experience

In the authors' opinion, this is a matter of some moment for the DoD. In order to protect the research goals of the HPCMP and to foster the trust of several "stake-holders," the HPCMP would be best served if decisions in makes about the applicability of the technology being developed easily admitted of radical changes in command direction and Congressional mandates. There are typically skeptics of technology in any armed force and in any government, even the most technically oriented ones. (Neitzel, 2007) The authors recognize the fine line that must be walked between being too conservative in attempting new technology and in being rash to the point of possible allegations of waste. While not articulated, these considerations were necessarily kept in mind in the selection of the technology for both the Distributed Computing (DC) and Dedicated High Performance computer Project Investment applications.

These programs were set up by HPCMP to be "... modest-sized HPC systems awarded to technically sound, mission critical projects that cannot be performed at HPC centers due to special operational requirements (e.g., classification level, real-time response, hardware-in-the-loop, embedded implementations, use of emerging technologies)." (HPCMP, 2010) In the case of JFCOM, the need for real time, interactive and classified use mitigated against the use of the Major Shared Resource Centers (MSRCs).

With a history of dedicated hardware, this was an exemplar of the desirability of having a sufficiently general purpose design that virtually any change in direction, mandate or emphasis, would find the DHPI cluster still regarded, not only as a useful tool, but a critical part of the infrastructure being developed to meet the changing needs.

9. The Future Uses Envisioned

The uses envisioned for the restructured J9 will require a dedicated high performance computing asset. During the Spring of 2010, J9 conducted the JIPS HITL (Joint Integrated Persistent Surveillance - Human-In-The-Loop) experiment. This project was executed under the new mandates from JFCOM and conformed to the "respond to direct requests" mandate. In any case, it made full use of JSAF and SLAMEM, two legacy codes proven to work well on the cluster *Joshua*, but previously known to run into scaling problems on "PCs-on-a-LAN" configurations.

The next big event coming up is currently scheduled to use another of the legacy codes. Experience has shown that it too can be easily ported to the nodes on *Joshua*, thereby maintaining that platform as a critical component for the test. The designers of *Joshua* had just such an event in mind when they were looking at potential specifications for the platform.

Also presented at this 2010 UGC meeting is a paper extolling the use of the cluster for the implementation of physics-based simulations within the agent-based framework provided by the HLA and JSAF. Long seen by many as one of the grand challenges of FMS, the ability to simultaneously make use of the two simulation disciplines simultaneously in a real-time, interactive synthetic battlespace environment may be enabled by the inclusion of GPGPUs in *Joshua*.. The GPGPUs are well suited to do their Single Instruction Multiple Data (SIMD) calculations on the physic-based models that rely predominantly on Matrix Multiply and FFT calculations, all without adversely impacting the Central Processing Unit (CPU) processing of the Agent Based Modeling (ABM) portion of the environment or the communications fabric of the cluster.

10. Conclusions

The award of *Joshua* has met all of its major goals. It has been a productive and critical part of the J9 operations for several years. Its technology was in the vanguard of the use of GPUs as heterogeneous accelerators. The new technology has not proven to be disruptive in any way, as the computer has been successfully used on the mandate for its use that existed when it was procured. Further, the general design of the asset has shown an additional quality in that it has gone through a significant change in research direction without losing any of its value or reducing the criticality of its contribution.

Acknowledgments

Thanks are due to the excellent staffs at JFCOM, ASC-MSRC and MHPCC and our colleagues at NVIDIA. Some of this material is based on research sponsored by the Air Force Research Laboratory under agreement number FA8750-05-2-0204. Other work is based on research sponsored by the U.S. Joint Forces Command via a contract with the Lockheed Martin Corporation and SimIS, Inc. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation thereon. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of these organizations.

References

Barrett, B. & Gottschalk, T.D., (2004), Advanced Message Routing for Scalable Distributed Simulations, in the Proceedings of the Interservice/Industry Training Simulation and Education Conference, Conference, Orlando, Florida

Boot, M. (2003), The Savage Wars Of Peace: Small Wars And The Rise Of American Power, Basic Books, New York, New York

Brunett, S., & Gottschalk, T.D., (1998), A Large-scale Meta-computing Framework for the ModSAF Real-time Simulation, Parallel Computing:, V24:1873-1900, Amsterdam

Buck, I., (2007), GPU Computing: Programming a Massively Parallel Processor, International Symposium on Code Generation and Optimization, San José, California

Ceranowicz, A. & Torpey, M., (2005), Adapting to Urban Warfare, Journal of Defense Modeling and Simulation, 2:1, January 2005, San Diego, Ca

Charlesworth, A., & Gustafson, J., (1986), Introducing Replicated VLSI to Supercomputing: the FPS-164/MAX Scientific Computer, in IEEE Computer, 19:3, pp 10-23, March 1986

CJCS, (1996), Joint Vision 2010, Director for Strategic Plans and Policy, J5: Strategy Di-vision, Washington, D.C.: Government Printing Office

CJCS, (2000), Joint Vision 2020, Director for Strategic Plans and Policy, J5: Strategy Di-vision, Washington, D.C.: Government Printing Office

CJWC, (1997), Concept for Future Joint Operations, Commander, Joint Warfighting Center, Fort Monroe, VA.

CRA, (2010), The Change at DARPA, Computing Research Policy Blog, retrieved from the internet on 13 May 2010 from: http://www.cra.org/govaffairs/blog/2010/03/the-change-at-darpa/

Diniz, P., Lee, Y.-J., Hall, M. & Lucas, R., (2004), A Case Study Using Empirical Optimization for a large, Engineering Application, in the Proceedings of the 18th International Parallel and Distributed Processing Symposium, April 2004, pp: 200-208

Dongarra, J., (1993), Linear algebra libraries for high-performance computers: a personal perspective, Parallel & Distributed Technology: Systems & Applications, IEEE, Feb. 1993, Volume: 1, Issue: 1, pp: 17 - 24

Fatahalian, K., Sugerman, .J. & Hanrahan, P., (2004), Understanding the efficiency of GPU algorithms for matrix-matrix multiplication, Workshop on Graphics Hardware, Eurographics/SIGGRAPH

Frigo., J., Palmer, D., Gokhale, M., and M. Popkin-Paine, M. (2003), Gamma-ray pulsar detection using reconfigurable computing hardware, 11th Annual IEEE Symposium on Field-Programmable Custom Computing Machines (FCCM),.

Henry, C., Pfeiffer-Vega, S. & Gabor, B., (2003), The High Performance Computing Modernization Program, *CHIPS — the Department of the Navy Information Technology Magazine*, Charleston, South Carolina, Spring 2003

HPCMP, (2010), Resource Management: DHPI's, retrieved from the HPCMP web page on the Internet on 13 May 2010 from: http://www.hpcmo.hpc.mil/cms2/index.php/dphiinformation

Lastra, A., Lastra, M. Lin, and D. Minocha, (2004), ACM Workshop on General Purpose Computations on Graphics Processors.

Linderman, R. W., Linderman, M. H., and Lin, C-S., (2005), FPGA Acceleration of Information Management Services, 2005 MAPLD International Conference, Washington, DC

Lucas, R.F., Wagenbreth, G., Tran, J.J., & Davis, D. M., (2007), Multifrontal Computations on GPUs, unpublished ISI White Paper, on line at: www.isi.edu/~ddavis/JESPP/2007_Papers/SC07/mf2_gpu_v0.19a-nms.doc

Lucas, R., & Davis, D., (2003), Joint Experimentation on Scalable Parallel Processors, 2003 I/ITSEC Conference, Orlando, FL

Joint Pub 1-02, (2000), Department of Defense Dictionary of Military and Associated Terms, Chairman of the Joint Chiefs of Staff, Washington, D.C.

Messina, P. C., Brunett, S., Davis, D. M., Gottschalk, T. D., (1997) Distributed Interactive Simulation for Synthetic Forces, In Mapping and Scheduling Systems, International Parallel Processing Symposium, Geneva

MSCO, (2009), Modeling and Simulation Coordination Office, definition of LVC Simulations, Alexandria, VA

Neitzen, S., (2007), Tapping Hitler's Generals: Transcripts of Secret Conversations, 1942-1945, page 102, line 30, Frontline Books, St Paul Minnesota

Sumanaweera, T. and Liu D., (2005), Medical Image Reconstruction with the FFT, in GPU Gems 2, M. Pharr, Ed. Boston: Addison-Wesley

Tran, J.J., Lucas, R.R., Yao, K-T., Davis, D.M. and Wagenbreth, G., Yao, K-T., & Bakeman, D.J., (2008), A High Performance Route-Planning Technique for Dense Urban Simulations, WSC05-2008 I/ITSEC Conference, Orlando, FL

Yao, K-T., Ward, C. & Wagenbreth, G., (2006), Agile Data Logging and Analysis, 2003 I/ITSEC Conference, Orlando, FL