

Educational Extensions of Large-Scale Simulations Enabled by High Performance Computing

Dan M. Davis
Information Sciences Institute, USC
Marina del Rey, California
ddavis@isi.edu

Laurel K. Davis
Apple Tree Services
Culver City, California
laurel.davis@appletreeservices.com

Summer L. Allen
Science Applications Int'l. Corp.
Alexandria, Virginia
sallen@objectsciences.com

ABSTRACT

Large-scale intelligent agent simulations, enabled by high performance computing (HPC), have been effectively used by the Department of Defense for experimentation and analysis. The authors analyze their experiences in these and related areas, then present data and conclusions to support new applications of proven pedagogies to broaden the value of these capabilities across the areas of training and education. Over more than a decade, HPC has shown the ability to enable otherwise unattainable sizes of intelligent agent simulations, growing from small unit, to battlefield, to theater of war, and, finally, to global-scale operations. The techniques necessary to achieve these levels were imported and adapted from early supercomputing research in basic science projects at major universities. Among the insights from that research were the reductions of validity and utility suffered when constrained samples of the subject phenomena were simulated. This paper extends that concept into the discipline of education and demonstrates the putative desirability of having large-scale capabilities in the educational environment as well. The authors describe the available technologies for large-scale simulations, review the successes of experimentation and analysis enabled by those technologies, and outline the many opportunities for implementation in education. They then focus on early experimentation using distributed HPC to aid in technical and non-technical education for all age cohorts. They lay out a roadmap for future development and for assessments of applicability of their techniques by others who should benefit from such capabilities. Cost/benefit analyses are invoked to assist the potential users in making valid evaluations of the applicability of these proven techniques to their own uses. The development of an interactive educational module is outlined, described and lessons learned are reported. A test on a trans-continental meta-computing platform will be reported from the viewpoint of both HPC performance and educational efficacy.

ABOUT THE AUTHORS

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 Caltech, he managed Synthetic Forces Express, a major simulation project. Prior to that, he was a Software Engineer on the All Source Analysis System project at the Jet Propulsion Laboratory and worked on a classified project at Martin Marietta, Denver. An active duty Marine Cryptologist, he recently retired as a Commander, USNR, Cryptologic Specialty. He has served as the Chairman of the Coalition of Academic Supercomputing Centers and the Coalition for Academic Scientific Computation. He received a B.A. and a J.D., both from the University of Colorado in Boulder.

Laurel K. Davis is the Director of Apple Tree Services, an independent educational consulting and research organization in Culver City, California. She is an experienced classroom educator who has served in several public schools in the Los Angeles basin. Her current activities include teacher training, consulting with parents about the transition of their children from elementary into middle school environments, creation of materials to address non-academic skills and techniques to facilitate and enable scholarly pursuits, and leadership training and assessment. She has additional experience in commercial and private settings as a programmer and web site designer/manager. She received a B.A. in Communications and a M.Ed., both from the University of California, Los Angeles.

Summer L. Allen is the Division Director at the Science Applications International Corporation (SAIC) in Alexandria, Virginia. A former Naval Cryptologic analyst and linguist, she has specializes in the application of technology to the challenges posed by the Intelligence and Security Command (INSCOM). She has pursued the implications of cognitive research on the transfer of intelligence information to the warfighter. Ms. Allen earned a B.A. degree in Psychology from the University of Hawai'i.

Educational Extensions of Large-Scale Simulations Enabled by High Performance Computing

Dan M. Davis
Information Sciences Institute, USC
Marina del Rey, California
ddavis@isi.edu

Laurel K. Davis
Apple Tree Services
Culver City, California
laurel.davis@appletreeservices.com

Summer L. Allen
Science Applications Int'l. Corp.
Alexandria, Virginia
sallen@objectsciences.com

Introduction

This paper will speak to the application of distributed High Performance Computing (HPC) to the field of education, especially education of Department of Defense (DoD) personnel. It will be argued that the explosion in need for education can best, or perhaps only, be served by the expansion of the use of the capabilities of HPC. The DoD's needs for education and the recent advances in HPC will be delineated, followed by a discussion of the future uses of HPC technology to redress those needs.

Background

Computer scientists have developed an effective response to one of the constraining constants with which all physical scientists are familiar: the speed of light. This constraint, exacerbated by the laws of thermodynamics, seems to present a limit to the speed and power of digital computers using a single central processing unit (CPU). Early in the 1980s, a group of visionary scientists at the California Institute of Technology began to postulate that the restrictions articulated by Gene Amdahl (Amdahl, 1967) could be overcome. This led to the development of a series of parallel computers, ranging from the tens to hundreds of processors, all working in parallel on a single problem. Now thousands of processors in one computing platform are common (Dongarra, 2006).

Especially following the advent of the Intel Delta, a 512 node parallel computer, techniques were conceived, developed and implemented that indeed did prove many problems were practicably resolvable on "massively" parallel machines. One after another of the problems of our age were successfully attacked using the power of this admittedly expensive configuration. Thomas Sterling and Don Becker advanced a low-cost approach to parallel computing, often called a Beowulf, which entailed the collecting of commercial off the shelf (COTS) personal computers and connecting them with a low cost inter-node communication network (Sterling, 1997). The academic research community flocked to these

technologies and the brightest researchers produced truly amazing results. However, programming these machines is a non-trivial task.

Meanwhile, the DoD was making use of collections of computers, more informally connected, to simulate large-scale battlefields. They further had a manifest interest in being able to simulate more than one million vehicles, all with sophisticated behaviors, operating on a global-scale, variable resolution terrain database. This was driven by the government's needs to accommodate new high technology, network centric warfare (Cebrowski, 1998) and simulate more complex human functions (Ceranowicz, 2004) in situations with critical human participants (Sanne, 1999). The U.S. DoD began a series of experiments to model and simulate urban environments and asymmetric warfare. In support of their mission, analysts needed to conduct interactive experiments with entity-level simulations, using programs like Joint Semi-Automated Forces (JSAT) and associated simulation federates (Ceranowicz, 2002).

Not only was this difficult, but it needed to be done at a scale and level of resolution adequate for modeling the complexities of military operations in urban situations. All of this led to the government analysts' requirement of simulations of at least 1,000,000 entities, such as vehicles, non-combatants and warfighters, all on a global-scale terrain database with high-resolution insets. The Joint Forces Command (JFCOM) experimenters were using large numbers of Linux PCs, interconnected in the standard LAN configuration. They found that communications constraints limited the simulations to tens of thousands of vehicles, about one hundredth of the vehicles that they needed. The Joint Experimentation on Scalable Parallel Processors (JESPP) project was a paradigm example of the successful application of computational science and parallel computing to this type of situation (Lucas, 2003).

Developing these simulations so that the DoD could involve humans, i.e. Human-in-the-Loop (HITL), additionally augmented the DoD's ability to assess true impacts of a new system by including personnel and

procedures. (Ben-Ari, 1998) This called for several new methods to model human behavior (Hill, 2000) and take account of Political Military Economic Social Infrastructure and Information (PMESII) (Chaturvedi, 2000), both of which require significant additional compute power. Software implementations stressing efficient inter-node communications were necessary to achieve the desired scalability.

Battlespace simulations and distributed computing is based on the earlier work funded by DARPA in the mid-nineties and done by Paul Messina at Caltech (Messina, 1997). The Synthetic Forces Express project (SF Express) began to explore the utility of SPPs as a solution to the communications bottlenecks that were then being experienced by ModSAF. The SF Express project was charged with proving the capability of scalable communications architectures on multiple SPPs in order to simulate 50,000 vehicles: an order-of-magnitude increase over earlier major simulation, Synthetic Theater of War–Europe (STOW-E).

More than 100,000 individually simulated vehicles were simulated in March of 1998 by the SF Express project. The runs used several different types of SPPs at nine separate sites spanning seven time zones. These sites were linked by a variety of wide-area networks. (Brunett, 1997) All used 64 bit processors, just now being incorporated in PCs. The meta-computing platform was very effective, despite the heterogeneousness with respect to “endian-ness,” shared/distributed memory, processor design, switch topology, and other parameters.




Figure 1 Plan View display from a SAF

The typical interface for this type of simulation is a map visualization, as shown in Figure 1. The menu on the left allows the operator to lay down forces on the

map, give them orders, modify their behaviors, and adjust their movements. The independent agents are fully capable of many autonomous operations, such as route finding, obstacle adjustment, small unit tactics implementation and reactions to enemy activity. A new program further allows the operator to vary the parameters of the independent agent entities, *e.g.* the top speed of a simulated aircraft, the weapons on an APC, the survivability of a tank, *etc.*

A few workstations on a local area network (LAN) can support simulations of a few thousand entities in a typical JSAF run. For this size simulation, a simple broadcast of all data to all nodes is sufficient. Each node discards data that is not of interest to it. The Run Time Infrastructure (RTI) controls this activity. When the simulation is extended to tens of thousands of entities and scores of workstations, broadcast is not sufficient. UDP multicast has proven to be a good solution to this issue, in lieu of the simple broadcast. In this case, each simulator receives only the data to which it has subscribed, *i.e.* in which it has a stated interest. This leaves enough compute capacity for data management and visualization, *e.g.* a three dimensional, rendered view of the urban area, shown in Figure 2.



Figure 2 3D Rendered display from a SAF

Large Scale Forces Modeling and Simulation

Recent experiments within the Joint Forces Command's Experimentation Directorate, J9, demonstrate the feasibility of forces modeling and simulation applications in a large field of play with fine-grained resolution. Simulating such battle spaces requires large computational resources, often distributed across multiple sites. The ongoing Joint

Urban Operations (JUO) experiments utilize the JSAT application suite and RTI-s to scale to over 300 federates distributed across the continental United States and Hawaii (Ceranowicz, 2002). The JUO exercise has shown the scalability of the JSAT/RTI-s infrastructure and of interest-based, router-managed communication and its ability to serve over 500 human participants with ease.

EDUCATION AND TRAINING: GENERAL DEFINITIONS

A narrowly focused program that leads to high proficiency in a specific skill is generally regarded as training. It prepares the service member for one particular job or activity but is not designed to provide either a broad perspective nor a range of approaches.

Education is designed to allow the service member to attack problems not envisioned as part of their MOS and to prepare for career advancement. Education fosters innovative and general approaches to problems and strives to enhance productivity, career satisfaction and process improvements. (Moore, 1998)

To a very large degree, the American armed forces have been famed for this aspect of their capabilities, even drawing the respect of Field Marshall Rommel (Young, 1994). An education prepares the service member to deal with and solve a broad range of problems, and to choose which issues are important and which are not.

When on officer selection boards, one of the authors (D. Davis) always asked candidates what book had been important to them and what book they were currently reading. That question was widely regarded as an important discriminator by other members of the board, supporting the thesis that education is vital to the DoD.

DIFFERENTIATE TRAINING & ED WITH RELATION TO TECHNOLOGY

Research in cognitive science and education indicates that meaningful learning which can be transferred to multiple settings occurs when high-quality instruction is integrated with frequent opportunities to 1) apply abstract knowledge in multiple authentic contexts (Gick and Holyoak, 1983; Cognition and Technology Group at Vanderbilt, 1997; Bransford et al, 1998), 2) receive feedback on the validity and success of the application (Bransford et al, 2000), and 3) re-engage in

instruction with refined knowledge targets. However, due to time, budget, training and material constraints, the majority of learners do work passively from a textbook or in a setting of diminished rigor (Bransford et al, 2000; Scales, 2004) and rarely engage in the content in the depth described above for the critical understanding that will allow them to actually use the knowledge in varied contexts.

Advances in HPC technology provide a welcome opportunity to engage learners in real-time authentic contexts most relevant to their desired training outcomes. The DoD would accelerate the learning curve and raise the competency of trainees by providing immediate, repeated and user- or variable-influenced simulation experiences in which learners must synthesize and apply their developing content knowledge. Use of HPC would also provide an infrastructure for incorporating the newly honed knowledge of service members who had been in the field and could return to the instructional setting and inform the training process further.

DoD Directives

DoD Instruction 1322.25 (DoD, 1997) calls for:

“E3.1.1.1. Provide for the academic, technical, intellectual, personal and professional development of Service members, thereby contributing to the readiness of the Armed Forces and the quality of life of Service members and their families.

E3.1.1.2. Increase Service members' opportunities for advancement and leadership by reinforcing their academic skills and occupational competencies with new skills and knowledge.

E3.1.1.3. Lead to a credential, such as a certificate, diploma, or college degree, signifying satisfactory completion of the educational program.”

This Instruction directs the provision of voluntary education opportunities for all DoD personnel, but does not suggest how this can be accomplished when faced with the exigencies of military life.

COLLABORATIVE LEARNING PRINCIPLES

An optimal way to facilitate and intensify engagement with knowledge and the quality and frequency of feedback provided is through collaborative learning environments. The recent explosion in social networking technologies indicates the potential for increased engagement of students through use of high-quality virtual communities. Not only does this allow for access to “distributed cognition,” (Vye et al.,

1998); it promotes superior problem-solving and cognitive development in ways that individual exercises may not (Evans, 1989; Newstead and Evans, 1995; Kobayashi, 1994). Collaborative learning environments enabled by HPC would allow for the participation of experts and novices in complex contexts in tandem, thus providing opportunities for conflict, internalization (absorption of knowledge by novices exposed to experts) and self-explanation (attempts by experts to convey knowledge to novices), all of which serve to expose misconceptions, deepen understanding and lower affective filters that may inhibit learning (Dillenbourg and Schneider, 1994).

One criticism of collaborative learning environments is that peer-centered learning is not always as effective as knowledge-centered learning because experts who have a mastery over the knowledge are not necessarily as knowledgeable about strategies for helping novices understand those concepts. (Shulman, 1986). Scalable HPC-supported collaborative learning would help to address this liability in collaborative learning. HPC can provide the communications bandwidth and computational power to manipulate not only the variables of the learning contexts themselves, but the variable factors of the learners in the group (Davis and Davis, 2003). This would deliver not only an environment that would be modifiable from a content perspective, but an environment that could be modified both by expert students and skilled instructors to meet the pedagogical needs of all learners in the group.

DoD Interest

The DoD has many requirements for educating its members using collaborative learning. One of those reasons is that current global engagement requires service members who are equipped with exceptional cultural awareness and an intuitive sense for the nature and character of war. Some argue that higher-level military colleges and schools fail to meet the learning needs of the services (Scales, 2004). Gen. Scales points out that very few military leaders are fortunate enough to be selected to attend institutions that really expand their educations. He asserts those selected are chosen based solely on job performance rather than for the excellence of their intellect. He further is critical of the trend to diminish the experience in the depth and rigor. The authors agree that the central elements necessary to gain a deeper understanding of the issues of consequence to the nation are often slighted.

Education necessarily must be a life-long process. Every service member, regardless of grade or specialty, should be given unfettered and continuous access to

the best and most inclusive educational programs. It is common knowledge that those who take advantage of the opportunity to learn should receive recognition and professional reward commensurate with the quality of that learning. The authors assert that HPC has provided the capability of expanded distance learning to allow the learning process, as conveyed by the best teachers, to be amplified and proliferated, such that every service member can learn to his or her own capacity and motivation. Without HPC capabilities, only a few could experience the best in education, with the concomitant unacceptably poor results (Thernstrom, 2004).

JIOC Distributed System Needs

U.S. Strategic Command's Joint Information Operations Center (JIOC) is responsible for the integration of Information Operations (IO) into military plans and operations across the spectrum of conflict. This inherently involves collaboration with dimension of heterogeneity along several axes: unit, language, mission, tradition, etc.

Training for collaborative activities is a major goal of the Joint Forces Command. At the JIOC, an even longer term, broader process is indicated. Over periods of decades, personnel who will be active participants in JIOC activities in the case of exigent operations, will have to develop and hone collaborative skills. The cost of physically moving personnel to participate in this long-term education process would be prohibitive. Time conflicts between the personnel to be engaged in times of emergencies, may effectively preclude any significant joint educational activities. The distributed platform provided to JFCOM may hold promise for resolving this conundrum, both in its distributed availability and its scalable ability to provide any size of problem area whenever needed.

The newly coined term, Informational Operations, involves actions taken to affect adversary information and information systems while defending one's own systems. IO should impact all phases of an operation, span the entirety of military operations, and be evident at every level of war. This is vital to the modern commander's capability to mount and sustain operations by maintaining the level of information superiority required for decisive joint operations.

EXPERIENCE WITH HPC ON TRANS-CONTINENTAL ENVIRONMENT

One of the great strengths of the JFCOM experimental design is its use of distributed assets and service to dispersed users. The experiments themselves are housed and controlled by the JFCOM out of its experimental bay near Suffolk, Virginia. Environments and data are managed remotely out of Fort Belvoir in Northern Virginia. The civilian “culture” entities are laid down and managed by a team a continent away in San Diego, at the SPAWAR center on Point Loma. The two 128 node, 256 processor Linux clusters that are provided by the HPCMP, are located in Maui at the Maui High Performance Computing Center (MHPCC) and at Wright Patterson Air Force Base at the Aeronautical Systems Center Major Shared Resource Center (ASC-MSRC) in Ohio.

Communications between the sites are provided by the Defense Research and Engineering Network (DREN). Experiments for Urban Resolve have been both unclassified and classified. The work that has been undertaken for CENTCOM also requires encryption of the communications’ links. The Linux operating system is common across the net (usually Fedora, but also Red Hat Enterprise) and most of the programs are written in C++, with a smattering of Java.



Figure 3 Notional JFCOM Net Map

Note that there are geographical separation issues, Maui being on the order of five thousand miles from Suffolk, as indicated in the notional diagram in Figure 3. This precludes easily and economically meeting with the entire staff. Further, with a five time zone span, the operational synchronization is difficult, most

especially in the summer when the mainland sites go to daylight savings time, while Hawai‘i does not, thereby creating a six-hour difference. However, this experience would be critical in a DoD-wide distance educational effort, as such an initiative might well be literally worldwide. While efficient routing reduces latencies, the speed-of-light-latency is a phenomenon with which future user will have to contend.

VISIONS OF POTENTIAL OUTCOMES (with and without large-scale technology)

Education module

Issues of current concern to the DoD educator include the difficulties that arise as the target student population becomes more geographically dispersed, culturally heterogeneous, and beset by disruptive and unpredictable schedules. Further, increasing evidence (Thernstrom, 2004) shows that there is a paucity of good educators, which is driving the search for ways to leverage the truly gifted instructors to maximize the exposure of students to them. One of the areas that holds promise is the burgeoning discipline of distance education. In this environment live educators can provide materials for students around the world. The University of Southern California has significant experience in this field, granting more Engineering Masters degrees via distance education than any other University. They have specialized in serving the defense industry, an environment not dissimilar to the DoD in its dispersion of personnel and frequent disruptions of student availability.

New visualization techniques and interactive computer instruction show promise for opening up education to wider populations and enabling a Digitized One-on-One Teaching Method for teaching a plethora of concepts more effectively. A concept being explored at ISI is the concept of using animated avatars (Johnson, 2003) or a multitude of short live MPEG clips (Davis & Davis, 2003) to provide the student with the interactive learning experience. High Performance Computing can provide the richer MPEG experience to a dispersed and heterogeneous student body in the same way that it has provided a valid battlespace environment to the analysts at JFCOM. There is an active debate about the use of animated avatars vice using “live” instructors, captured by video technology and saved as MPEGs. A consistent complaint about the avatars is that they seem un-real and are adequate for gaming, but not to establish the trust inherent in a classic educational setting.

It is believed that this expanded audience and improved interactivity will be most effective if it allows for cultural and intellectual differences in each learner. Potential learners suffer from restraints due to geographical separation, erratic work schedules, and the restricted access to necessary materials. Static and unidirectional computer and internet presentations of this material may not fully exploit the opportunities to introduce new subjects, foster increased sophistication in analysis, and motivate learners to appreciate the subject matter as a tool and as a professionally enhancing experience.

As noted above, experience indicates that the learner becomes more involved if exposed to an interactive environment. There is evidence that true interactivity, both in the interface and in the presentation methodology, will further enhance learning and retention. The key question in this area is the efficacy of the combination of virtual environments and interactive teaching techniques in enhancing DoD-focused education. Two other critical issues are the architecture of the compute/networking assets needed to economically serve such a program and the impact such a program might have on DoD students' achievement and attitudes.

It can be argued that the advances in computer visualization programs, high performance computing, and high bandwidth networking can be used in reducing the barriers of access. It can also help correct the lack of effective evaluation of learning across a widely distributed and diverse population.

One easily monitored test bed for an initial implementation of these new technologies is the entire United States Department of Defense. Rigorous pre- and post-training evaluations will give meaningful and reliable data on which to analyze the impact of the program and to assist in identifying the need for further improvements. The program envisioned would automatically monitor progress of the participants.

Hypothetical

Easy Access by Population of Interest

For a new method of instruction to have any impact, it must reach a significant segment of the target population. The target population of this study will initially be the technical professionals of the DoD, but eventually the entire United States DoD population. Teacher training and motivation, as well as the acceptance of the program as "their own," are necessary and seen as a major component of this project. The interactive training program itself must

reduce as many barriers for the learner as possible or it may dissuade that learner from participating.

In diverse populations represented in the DoD, trying to teach concepts using a slavishly uniform pedagogical approach across these populations may exacerbate this difficulty. Ignoring either learning styles of different individuals or traditional views of different cultural groups may just perpetuate and emphasize these cultural differences. The geographical dispersion of the populations will also reinforce the isolation of some of the groups. The goal is to afford each person a more individualized experience in learning about the DoD mandated topic. Distributed HPC has proven it can reliably supply the compute power needed to meet these challenges.

Verifiable Measurement of Learning

Quantitatively verifiable results are necessary for useful research. A well-designed implementation plan would carefully establish what the desired results are in terms of useful knowledge and scientific sophistication. That factor will mandate rigorous interviewing of military and civil education leaders in the DoD. It would then be possible to identify the goals and re-visit the same group to insure the proper selection and articulation of these goals. The identified goals will then be quantified in a way to enable measurement and test/retest validity could be established within the target population. Progress toward these goals must be assessed and the authors have designed a three-step process for this.

- Pre-instructional assessment at the initiation of the training
- Periodic assessment as an integral part of the training
- Comprehensive assessment at the end of the training

These assessments will be focused on new or improved capabilities in the DoD member. To insure that these changes are appropriate, constant feedback from military leaders will be sought and heeded. Sensitivity to cultural, moral, and motivational issues will be a central focus of not only content, but of the assessment methodology as well. However, the results sought shall comport with recognized scientific standards and be centered on:

- Understanding basic concepts
- Articulating the key concepts central to understanding
- Synthesizing facts, military organization needs, and unit structure understanding

- Converting plans into appropriate plans of action

Figure 4 shows a flow chart depicting the way one type of interaction might be effectively offered to various service members. HPC could provide sufficient power to analyze strengths and weaknesses and address them as appropriate.




Figure 4

An early flow chart showing a putative design for an instructional module

Real

So, just what has been the experience of the DoD with trans-continentially distributed High Performance Computing and what does that experience tell the DoD educational professional about this new capability? The short answer is clearly that it has been a very good experience (Wagenbreth, 20005). The stability and the utility of the large Linux clusters at Wright Patterson AFB and on Maui have elicited the unsolicited admiration of both the technical and the analytical professionals at JFCOM. These experiments have attracted attention from DDR&E, DDS&T, Senator Hillary Clinton, Former House Speaker Gingrich, and nearly half a dozen current members of the House. The message is that large-scale computing on Linux clusters is a sufficiently mature technology that any educational issue that would benefit from more powerful computing can look with confidence to this paradigm.

The JSAF experiments at JFCOM call upon the clusters at ASC-MSRC and MHPCC to run, “flat out” for a week at a time, all the while collecting data at a rate greater than 10 GB per hour (Yao, 2005). This computational power and data analysis capability

- Citing appropriate facts to substantiate a course of action

should easily be implemented to address educational issues as well.

How to adapt, adopt, and implement

While the above treatment of the subject lays out the capabilities provided to the JFCOM experimenters for analytical purposes, the central topic of this paper is education. It is the authors’ opinion that virtually all of the cognitive, pedagogical and organizational issues raised earlier in this paper can be well served, better served or, in some cases, only served by distributed HPC. Traditional educational approaches all have disabling impediments that are exaggerated by the service environment of the armed forces. The current capabilities in the FMS community have both shown the way and provided a ready magazine of legacy codes that are tested and available to meet the needs of education for the service member. Some adaptation will be required, but that has proven in the past to be straightforward (Lucas, 2003). More critical is the rigor with which the educational modules are designed and the consistency and efficacy of the feedback mechanisms.

In any case, experience and logic dictate that it is often safer to adopt existing code bases than to take a leap into the uncertain future of an entirely new program. A senior executive at IBM stated that their overwhelming experience was that when a new program was designed to replace an existing code, by the time the new code was ready, the old program had advanced far beyond the new one and continuation of the existing code was often directed (Herbert Schorr, private conversation with one of the authors, April, 2002).

Implementation of an educational program to meet the needs laid out herein would not be a trivial process, but there are both sufficient intellectual centers to drive the creative focus on DoD educational goals and experienced system integrators to produce a program that leveraged the motivational talents of the few truly gifted teachers and provided the DoD service members with a ubiquitous and asynchronous program of education to enable the member to participate wherever and whenever time is available. While the power of HPC may be necessary, retaining that power calls for scalability and that requires experienced computational scientists. This capability has been demonstrated over and transcontinental network that has an allocated bandwidth of 20 Mbps, but typically

uses less than 1Mbps and is tolerant of latencies as long as 500 milliseconds (Davis, 2005)

Conclusion

The authors started from what they consider to be a virtually unassailable set of premises:

- Simulation and training have benefited from high performance computing
- Education needs newly developed capabilities to keep up with modern demands
- Users, educational assets and computing are often dispersed geographically
- Technology is available to deliver HPC power to educational systems

A second lesson, flowing from the first, was that the experience and insights from the High Energy Physics community were more germane than were the commercial transaction-processing programs or the Internet recreational data-search designs. The reason for this is assumed to be the closer relation to the types of data, the technical literacy of the users, the more uniform access to high-bandwidth, and the lack of the need for elaborate inter-user security, but very high external security. However, there are differences from the HEP community, *e.g.* the simulation community's broader and more dynamic analytical interest. The HEP community tends to work with output from a single, well-defined experiment, looking for a specific set of data. The military Forces Modeling and Simulation community had many user interests and changing goals.

The third lesson is that working with open-source, public licensed software has many advantages for the developer. The Linux community is active and involved. Source code is available for scrutiny, modification, and implementation. Hundreds of hours that would be expended in the procurement process are side-stepped. While the authors recognize the value added by major commercial vendors and their support staffs, the JFCOM experience indicates that open source software should be seriously considered.

Finally, it is the authors' experience that, at this time, there is no substitute for absolutely world-class parallel architects. While the great bulk of the code to be developed is easily within the capabilities of journeymen programmers, the overall design can be fatally crippled if it is not optimized for parallel processing. That requires finding experienced personnel at major HPC centers. That this is true would not be surprising at all to the venerable and venerated Fred Brooks, (Brooks, 1995) who

emphasized the value of extraordinary programmers as being worth as much as ten other programmers. The authors therefore conclude that a prudent implementer of a distributed educational system would be ill-advised to proceed without an experienced and successful parallel architect, typically a Ph.D.

In closing, it should be restated that High Performance Computing brings FMS new simulation capabilities, which brings new floods of information, which leads to the need for large scale and sophisticated data analysis, for which HPC facilities are optimized. Education is in dire need of all of these, especially considering the constraints of the DoD target group. The combination of careful planning and openness to others skills are a *sin qua non* of success in the provision of Education to the service members of the armed forces.

ACKNOWLEDGEMENTS

The authors wish to acknowledge the members of the ISI JESPP team who have contributed to this paper through their efforts and their intellectual stimulation. Much of the success reported here came from the Joint Experimentation on Scalable Parallel Processor project, initiated, directed and funded by the Joint Forces Command and to a very large degree conducted on the compute assets of the Maui High performance Computing Center, ASC-MSRC at Wright Patterson Air Force Base and other members of the High Performance Computing Modernization Program. This material is based on research sponsored by the Air Force Research Laboratory under agreement number FA8750-05-2-0204. 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 the Air Force Research Laboratory or the U.S. Government.

REFERENCES

Amdahl, G.M. (1967) Validity of the single-processor approach to achieving large scale computing capabilities. In *AFIPS Conference Proceedings* vol. 30. AFIPS Press, Reston, Va., , pp. 483-485.

Ben-Ari, E. (1998). Mastering Soldiers: Conflict Emotions and the Enemy in an Israeli Military Unit. New Directions in Anthropology, V. 10. Oxford: Berghahn Books.

Brooks, F., (1995), The Mythical Man-Month: Essays on Software Engineering, Addison Wesley, New York, New York

Brunett, S., Davis, D., Gottschalk, T., and Messina, P., (1998), Implementing Distributed Synthetic Forces Simulations in Metacomputing Environments, Seventh Heterogeneous Computing Workshop, Orlando, Florida

Cebrowski, A.K., & Garstka, J.J., (1998), Network Centric Warfare: Its Origin and Future, Naval Institute Proceedings, 124/1, 28-35.

Ceranowicz, A., Torpey, M., Hellinstine, W., Evans, J. & Hines, J. (2002) Reflections on Building the Joint Experimental Federation. *Proceedings of the Interservice/Industry Training, Simulation and Education Conference*, Orlando, FL.

Chaturvedi, A., Gupta, M., Mehta, S., & Yue, W., (2000), Agent-based Simulation Approach to Information Warfare in the SEAS Environment, 33rd Annual Hawaii International Conference on System Sciences (HICSS-33), 4-7 January.

Davis, D., Lucas, R., Amburn, P., & Gottschalk, T., (2005), Joint Experimentation on Scalable Parallel Processors, ITEA Journal of Test and Evaluation, Vol. 26, Number 2, Fairfax, Virginia.

Davis, D.M. & Davis, L.K. (2003), An Interactive-teaching Virtual Museum: Implementation of a New Digital One-on-One Pedagogy for K-6 Students from a Multi-Cultural Society, Unpublished USC Whitepaper, Available at http://www.isi.edu/~ddavis/DanzFiles/K-12Educ/IVM/IVM_Paper2Col.doc

DOD, (1997), DoDI 1322.25, *Voluntary Education Programs*, U. S. Department of Defense, Washington, D.C.

Bransford, J.D., L. Zech, D. Schwartz, B. Barron, N.J. Vye, and Cognition and Technology Group at Vanderbilt. (1998), Designs for environments that invite and sustain mathematical thinking. In *Symbolizing, Communicating and Mathematizing: Perspectives on Discourse, Tools and Instructional Design*, P.Cobb, ed. Mahwah, NJ: Erlbaum.

Bransford, J.D., A.L. Brown, and R.R. Cocking, eds. (2000), Learning and Transfer. In *How People Learn: Brain, Mind, Experience, and School*, Washington, D.C.: National Academy Press.

Cognition and Technology Group at Vanderbilt. (1997), *The Jasper Project: Lessons in Curriculum, Assessment and Professional Development*. Mahwah, NJ: Erlbaum.

Dillenbourg, P., and Schneider, D. (1994), Collaborative Learning in the Internet. *Proceedings, Fourth International Conference on Computer-Assisted Instruction*, Taiwan. S10-6 to S10-13.

Dongarra, J., Meuer, H., & Strohmaier, E. (2006), TOP500 Supercomputer Sites 22th Edition, Proceedings, SC06, Seattle, Washington, down loaded url: <http://www.top500.org/list/2005/11/100>

Evans, J. St. B.T. (1989), *Bias in Human Reasoning*, Hillsdale, NJ: Erlbaum.

Gick, M.L. and K.J. Holyoak. (1983), Schema introduction and analogical transfer. *Cognitive Psychology* 15:1-38.

Gottschalk, T. & Amburn, P., (2005, publication pending), Extending The MeshRouter Framework for Distributed Simulations. *Proceedings of the Interservice / Industry Training, Simulation and Education Conference*, Orlando, FL.

Graham, S., Snir, M., and Patterson, C., Editors, (2004) *Getting up to Speed, the Future of Supercomputing*, The National Academy Press, Washington, D.C.

Graebener, R. & Rafuse, G., Miller, R., & Yao, K-T., (2003) The Road to Successful Joint Experimentation Starts at the Data Collection Trail. *Proceedings of the Interservice / Industry Training, Simulation and Education Conference*, Orlando, FL.

Johnson, L., & Beale, C., (2003), Iterative Evaluation of Intelligent Game in Language Education, In the *Proceedings of the Artificial Intelligence in Education*, Sydney, Australia

Kobayashi, Y. (1994) Conceptual acquisition and change through social interaction, *Human Development* 37:233-241.

Lucas, R. & Davis, D. (2004) Joint Experimentation in Scalable Parallel Processors, *Proceedings of the Interservice / Industry Training, Simulation and Education Conference*, Orlando, FL.

Messina, P., Brunett, S., Davis, D., Gottschalk, T., Cerkendall, D., & Seigel, H., (1997) "Distributed Interactive Simulation for Synthetic Forces," in the Proceedings of the 11th International Parallel Processing Symposium, Geneva, Switzerland

Moore, J., (1998), Education versus Training, *Journal of Chemical Education*, Vol. 75 No. 2

Newstead, S.E., and J. St. B. T. Evans, eds. (1995), *Perspectives on Thinking and Reasoning: Essays in Honour of Peter Wason*. Hillsdale, NJ: Erlbaum.

Sanne, J. (1999). Creating Safety in Air Traffic Control. Unpublished doctoral dissertation, Institute of Tema Research, Linköping University, S-581 83 Linköping, Sweden.

Scales, R., (2004) Statement of Major General Robert Scales, USA (ret.) Testifying before the House Armed Services Committee on July 15, 2004.

Shulman, L., (1986) Those who understand: Knowledge growth in teaching. *Educational Researcher* 15(2):4-14.

Sterling, T., Salmon, J., Becker, D.J. & Savarese, D.F., (1999) *How to Build a Beowulf A Guide to the Implementation and application of PC Clusters*. MIT Press, Cambridge, MA,

Thernstrom, A., & Thernstrom, S., (2004), *No Excuses: Closing the Racial Gap in Learning*, Simon and Schuster, New York, New York

Vye, N.J., S.R. Goldman, C. Hmelo, J.F. Voss, S. Williams, and Cognition and Technology Group at Vanderbilt. (1998), Complex mathematical problem solving by individuals and dyads. *Cognition and Instruction* 15 (4).

Wagenbreth, G., Davis, D., Gottschalk, T., Lucas, R. & Yao, K-T. (2005, publication pending), Operational Experience, Distributed Simulations, Data Management and Analysis , *Proceedings of the 2005 Winter Simulation Conference*, Orlando, FL.

Yao, K-T. & Wagenbreth, G., (2005, publication pending), Simulation Data Grid: Joint Experimentation Data Management and Analysis, *Proceedings of the Interservice / Industry Training, Simulation and Education Conference*, Orlando, FL.

Young, D., (1994), Rommel the Desert Fox, Amereon Ltd (December 1994), London, UK