GPGPU Programming Courses: Getting the Word Out to the Test and Evaluation Community

Gene Wagenbreth, Dan M. Davis, & Robert F. Lucas
Information Sciences Institute, USC
Marina del Rey, California

ABSTRACT

Heterogeneous computing offers significant advantages for many disciplines that could beneficially use High Performance Computing. A demonstrably beneficial heterogeneity is the use of General Purpose Graphics Processing Units (GPGPUs) to accelerate algorithms, including those found in the Test and Evaluation (T&E) environments. The authors will share their extensive experience with GPGPUs and they will contrast programming ease with that of the Sony-Toshiba-IBM (STI) Cell chips. Experiences with Joshua, a 256 node GPGPU-Enhanced Linux Cluster at JFCOM will be presented. They will discuss the use of heterogeneous computing in the T&E community. Most importantly, though, they will discuss at length the creation, design, organization and presentation of three courses they have taught, introducing programmers to GPGPU programming. They will present sample materials and discuss lessons learned in order to assist potential GPGPU users to evaluate their own computational needs, training requirements, and logistics of conducting a course of their own.

AUTHORS

Gene Wagenbreth is a Systems Analyst for Parallel Processing at the Information Sciences Institute at the University of Southern California, doing research in the Computational Sciences Division. Prior positions have included Vice President and Chief Architect of Applied Parallel Research and Lead Programmer of Pacific Sierra Research, where he specialized in parallelization of Fortran programs. He is active in benchmarking, optimization and porting of software and has programmed on CRAY, SGI, Hitachi, Fujitsu, NEC, networked PCs, networked workstations, IBM SP2, as well as conventional machines. He received a B.S. in Math/Computer Science from the University of Illinois, Champagne-Urbana. genew@isi.edu

Dan M. Davis is the Director, JESPP Project, Information Sciences Institute (ISI), University of Southern California (USC), focusing on large-scale distributed DoD simulations. As the Assistant Director of the Center for Advanced Computing Research at Caltech, he managed Synthetic Forces Express, bringing HPC to DoD simulations. Prior experience includes work as a Software Engineer at the Jet Propulsion Laboratory and at a classified project at Martin Marietta. He saw duty in Vietnam as a Cryptologist in the USMC and he retired as a Commander, Cryptologic Specialty, U.S.N.R. He received B.A. and J.D. degrees, University of Colorado in Boulder. ddavis@isi.edu

Robert F. Lucas is the Director of the Computational Sciences Division of ISI at the USC. There he manages research in computer architecture, VLSI, compilers and other software tools. He's been the PI on the JESPP project since 2002. Prior to joining ISI, he was the Head of the High Performance Computing Research Department for NERSC at LBNL. Earlier he was the Deputy Director of DARPA's ITO, and a member of the research staff of IDA's Center for Computing Sciences. Dr. Lucas received B.S., M.S., and Ph.D. degrees in Electrical Engineering from Stanford University, Palo Alto, California. rflucas@isi.edu

GPGPU Programming Courses: Getting the Word Out to the Test and Evaluation Community

Introduction and Background

One of the most promising heterogeneous computing additions to computing power in the last decade has been the advent, development and programmability of General Purpose Graphics Processing Units (GPGPUs). The recognition that they can accelerate the execution of a wide range of algorithms and that their costs are kept low has made this a viable path. (Lastra *et al.*, 2004) The low cost has been made possible by mass marketing of this computational technology to the untold multitudes of "gamers" in the consumer market. (Davis, 2010)

The GPUs power potential is large and growing at a faster rate than the CPU for which they would act as accelerators. However, while it is true that performance may often be measured in peak GigaFLOPS (billions of Floating Point Operations Per Second), true utility is better measured by productivity that takes into account many other issues, *e.g.* programming ease, availability of trained programmers, maintainability of code, durability of technology, *etc.* (Kepner, 2004) To respond to some of these needs and to gain insights into their magnitude and impact, the authors have presented three courses and a graduate student symposium on the programming of GPGPUs for accelerated computation.

The first of these courses was in response to a request by the Joint Forces Command (JFCOM), whose managers were the custodians of the new Dedicated High Performance computer Investment (DHPI)-provided Linux Cluster, *Joshua*. It indicated JFCOM Joint Concept Development and Experimentation Directorate's (J9's) interest in and need for a class on the use of Graphics Processing Units (GPUs) as General Purpose GPU (GPGPU) accelerators. As these GPUs have incredible processing power, they were thought to be potentially useful in performing certain kinds of computation to enhance the processing of the Central Processing Unit (CPU) on each node of the cluster. JFCOM at that time had the largest GPGPU-enhanced cluster in the world, but it could not be effectively employed unless the managers and policy makers understood the implications of this new technology and the programmers and operators had a solid foundation for its use. Installation of the new DHPI machine was accomplished in August 2007 and operations were expected in September. This DHPI cluster was housed in computer facility at the JFCOM Joint Training Directorate (J7} in Suffolk Virginia and that directorate has expressed an interest in this class as well.

Offering this course enabled JFCOM and the FMS community to make rapid advances as they worked to take full advantage of the new DHPI Cluster. One of the unique aspects of the new Linux cluster was that there was a (then) state-of-the-art (GTS 8800) NVIDIA GPU on each node. JFCOM was interested in their being able to improve simulation performance for their battlefield experiments, a matter of significant interest to them for some time (Messina, 1997). They hoped to be able to do better on the new DHPI and that depended on their having the capability to make modifications of legacy simulation programs such as JSAF. This was anticipated to afford them the opportunity to take advantage of this particular heterogeneous High Performance Computing (HPC) architecture and to illuminate a path for implementation of other accelerators of significant interest, *e.g.* FPGAs, PIMs, Cell processors and Signal Processors. As JFCOM conducted the Noble Resolve experiments, the performance of the underlying codes

needed to be improved to achieve stated goals and there was a need to optimize the use of the new DHPI asset

The second course had similar underlying motivations, but the location and the target audience were different. This course was presented in the fall of 2008 in the ISI offices in Marina del Rey California, located some ten miles west of the main USC campus. Nevertheless, several of the students were USC undergraduates and many of the rest were USC researchers. Here the emphasis was much more on the theoretical underpinnings of the GPGPU concept and the class discussions took on a significantly more analytic tone.

The third course, again focusing more on the journeyman programmer, was conducted in La Jolla, at the San Diego Super Computer Center on the University of California campus there. The class participants here were evenly split between "working programmers" from the government research facilities on Point Loma and the "academic researchers" from SDSC and UC San Diego. This course was the first in which some of the students had NVIDIA "CUDA-capable" GPUs in their lap-top computers.

The fourth evolution was the presentation in March of 2010 of a graduate-level symposium at the University of Southern California in Los Angeles. This symposium was part of a "500-level" course in the Viterbi School of Engineering, Parallel Programming. The members of the audience here were all graduate students at that University and English was clearly a second language for the vast majority of the students. Their interest was more conceptual than the "work-a-day" programmers in the previous courses.

In total, some eighty students have completed the training offered. The common factor in all these courses was they were all taught by one of the authors, Gene Wagenbreth. They were each funded and sponsored by different agencies.

Objectives

The objective of these courses was for ISI to provide the necessary introduction to allow journeyman programmers to quickly make effective use of GPGPU acceleration in the codes on which they were working. Save for the graduate symposium, most of the "students" were experienced programmers working on battlefield simulations, DoD research programs, T&E professionals from the Navy, computer science students, and research managers.

ISI supported the creation and execution of all three GPU courses and the University of Southern California supported the symposium. These courses included the training material needed for course presentations and the student handouts needed for the execution of the programming after the course completion. Several students took advantage of the instructor's offer and called for assistance while programming their own code later.

Justification for GPGPU Acceleration

The simulation community has often been hampered by constraints in computing: not enough resolution, not enough entities, not enough behavioral variants. High Performance Computing (HPC) can ameliorate those constraints. The use of Linux Clusters is one path to higher performance; the use of Graphics Processing Units (GPU) as accelerators is another. These are often called General Purpose GPUs. Merging the cluster and GPGPU paths holds even more promise.

The ISI team members were the principal architects of a successful proposal to the High Performance Computing Modernization Program (HPCMP) for a new 512 CPU (1024

core), GPU-enhanced Linux Cluster, *Joshua*, for the Joint Forces Command's Joint Experimentation Directorate (J9). This cluster was awarded to J9 via the Dedicated HPC Project Investment program, DHPI and was configured in such a way to putatively utilize the GPUs to increase performance by a factor of two or more, with concomitant savings in cost, power and space.

Rationale for the Courses

Offering these courses were seen as a way to aid the U.S. DoD M&S community as they worked to take full advantage of the new heterogeneous opportunities. One of the unique aspects of the new JFCOM machine was that there was a state-of-the-art NVIDIA GPU in each node. These GPGPUs could be programmed using the new CUDA programming language (Compute Unified Device Architecture, a "C-like" language). DoD M&S had a need to improve simulation performance on new clusters and needed to make modifications to simulation programs such as JSAF, enabling them to take advantage of the newly heterogeneous HPC architecture. As the DoD M&S community has often conducted experiments focused on emergency responses in urban areas, performance needed to be improved to adequately model the complexities of these crowded areas and large populations. They needed to make the most of the new GPGPU-enhanced Linux cluster asset. Programming models, code examples and practice problems in CUDA were presented and implemented in the classes.

First Course - 2007

For the first class, the instruction was organized into two segments, one for a quick conceptual overview and the second for a practicum on programming for the end-users.

GPU 101 – Introduction and Overview (2 hour course) – no prerequisite

Target Audience, Intro Session:

- Technical staff seeking an introduction to GPUs
- Management needing an overview

Topics

- What is a GPU?
- How is it different from other devices?
- How does it work?
- How does it fit with the rest of the architecture?
- How can it be used as an accelerator or as a GPGPU?
- What kinds of algorithms work best on GPGPUs?
- What kinds of systems don't work well and cannot benefit from GPGPUs?
- What is the GPU market and projections?

GPU 102 – Practicum on Programming the GPU (20 Hour course) GPU 101 and hands on experience with C/C++ are prerequisite

Target Audience, Programmers:

- Technical staff seeking a basic knowledge reinforces by a practical application of how to program GPUs as part of heterogeneous processing model

Topics

- GPU Architecture
- Machine configurations
- GPU programming languages
- Software Development Kit (SDKs) (Linux / Windows)
- Floating and integer considerations
- Sample GPGPU applications
- How to structure your applications
- How to port your applications
- Benchmarking and Profiling
- How to optimize your application
- GPGPU Performance Tips

This course was given on the 23rd through the 25th of October 2007 in Suffolk, at the SAIC

offices there. The students were all from the JFCOM J9 Joint Concept Development and Experimentation Directorate, (J9). They were typically B.S. or M.S. graduates of state universities, but the course was attended by at least three Ph.D.s from elite universities. The setting was convenient to the students, being across the freeway from JFCOM where they usually worked.

Figure 1. SAIC Offices in Suffolk, Site of First GPGPU Course

In addition to Gene Wagenbreth, the course was taught by the Doctors Patrick Legresley and Paulius Micikevicius from NVIDIA, Dr. Phil Amburn from the High Performance Computing Modernization Program (HPCMP) and Mr. Bill Helfinstine from Lockheed Martin Corporation.

Figure 2. Dr. Patrick Legresley Briefs at the Large Screen Monitor

The course was structured to provide a significant amount of time to work on code of the students' own choosing and to provide immediate support and feedback to the programmers.

Figure 3. Dr. Paulius Micikevicius Looks Over the Shoulder of Dr. Andy Ceranowicz

This course was funded by the HPCMP via their PET (Professional Education and Training) program. That support was driven by Dr. David Pratt, one of their senior Forces Modeling and Simulation (FMS) research managers. He and Dr. Phil Amburn who provided on-site support for FMS activities at HPCMP centers, were deeply engaged in conceptualizing and designing the course. Dr. Amburn presented his view of the significance of GPGPU cluster programming to the course members and both Dr. Pratt and Dr. Amburn stayed to observe the entire course and provide assistance as required. HPCMP has an elaborate course evaluation and feedback mechanism which was briefed to the students. All of the presenting staff members were experienced presenters and lecturers.

It is the authors' opinion that careful selection of dynamic and organized presenters is a *sine qua no*n of a successful course. Both Dr. Pratt and Dr. Amburn provided real-time feed back

and suggestions to all of the presenters, and this mechanism was seen as beneficial in ensuring the high quality of the presentations and the continued maintenance of interest by the students in the course.

Figure 4. Course Sponsor Representatives: Dr.s Dave Pratt and Phil Amburn of HPCMP

Any good instructor will advise being very careful about picking a good setting for instruction. All four of the efforts covered here were in excellent settings. The had sufficient room for each student to use their own lap-top computers, sufficient outlets for power, internet connectivity for breaks, good lighting, computer visualization support (projected or large-screen monitor) and good ventilation. The authors' experience is that good ventilation is an oft over-looked, yet critical factor. Students become restless and lose interest in the material if the air gets stale.

Figure 5. Class Setting for First Course in Suffolk

A second critical factor is the material used. Most instruction relies on use of presentation slides and all four of these instances did as well. The instructors from NVIDIA came with a very large number of pre-drawn slides, which were very likely created by professional artists at NVIDIA

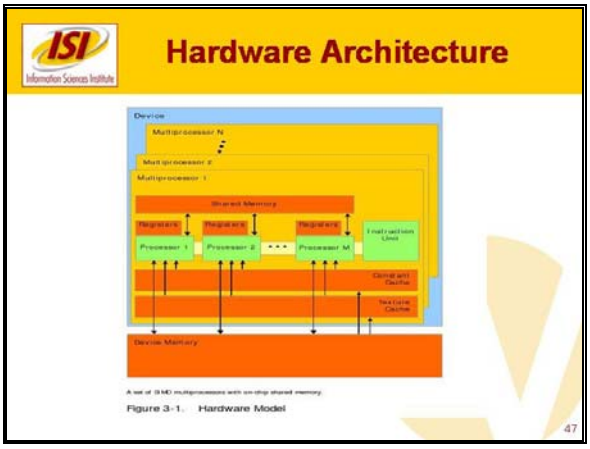

Figure 6. Slide showing NVIDIA Architecture

The critical factor here is to make sure that the slides are well designed to convey the message intended. One of the early issues encountered in the first course was that the slides were perhaps too generic for the rather focused needs of the audience. While teaching to a mixed audience must cover issues of interest to each participant, a more homogeneous audience will not respond well to instruction about issues for which they have not immediate need. The old communications admonition "Know your audience!" was reinforced by the experience in the first course. This also speaks in favor of locally developing a course tailored to the needs of the local user, rather than try to use a more generic course.

In order to fully comprehend the needs of J9, ISI considered their needs in terms of magnitude and flexibility had previously been impossible due to limitations of compute power. An earlier DC award of the two clusters at Maui and Wright-Patterson AFB

has enabled the development and implementation of a proven scalable code base capable of using thousands of nodes interactively. The JFCOM team continued to address community-wide issues such as: enhanced security for distributed autonomous processes, interactive HPC paradigms, use of advanced architectures, self-aware models, global terrain with high-resolution insets and physics-based phenomenology requisite for Joint Experimentation (Lucas, 2003). The ISI team was instrumental in advancing this research agenda. The need now was to ascertain how to make best use of the legacy codes and still optimize the benefits of heterogeneous computing.

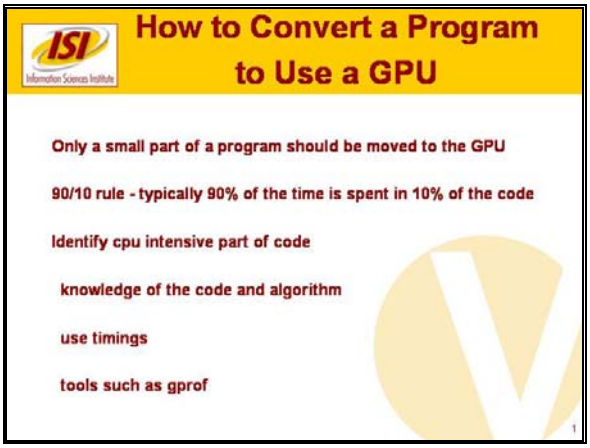

Figure 7. Slide Laying out Recommended Porting Approach

In wanting to address training, it was noted that the simulation community had often been hampered by constraints in computing: not enough resolution, not enough entities, not enough behavioral variants. High Performance Computing (HPC) was held to be able to ameliorate those constraints. The use of Linux Clusters was advanced as one path to higher performance; the use of Graphics Processing Units (GPU) as accelerators was another. These are called General Purpose GPUs (GPGPUs). Merging the cluster and GPGPU paths was seen to hold even more promise. The ISI team members were the principal architects of a successful proposal to the High Performance Computing Modernization Program (HPCMP) for a new 512 CPU (1024 core), GPU-enhanced Linux Cluster, *Joshua*, for J9. This cluster was configured in such a way to utilize the GPUs to increase performance by a factor of two or more, with concomitant savings in cost, power and space.

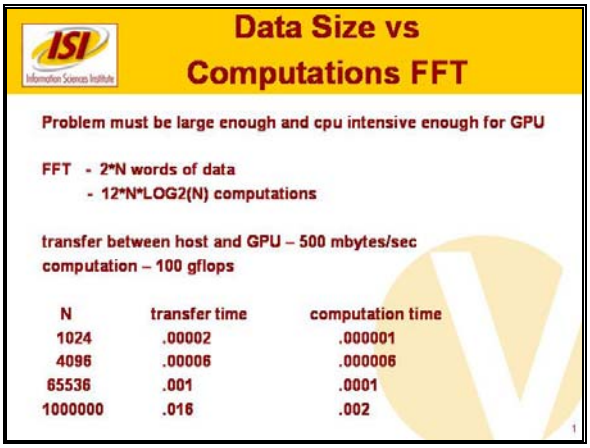

Figure 8. Slide Presenting Computation Parameters

The volume of material suggested is a matter of some discussion. The pace at which a presenter uses slides varies considerably. The degree to which users will rely on the slides as reference materials during later programming activities is similarly unknown. These courses were characterized by a truncating of the course length as time wore on and a reduction of the number of slides. Whether this is a concomitant relationship or just responding to the same pressure to present the data in more compact ways so as to be digestible by the already over-taxed programming community is open to interpretation.

Offering a course in GPGPU programming was seen as a way to aid JFCOM as they worked to take full advantage of the new JFCOM cluster. One of the unique aspects of the new machine was that there was a state-of-the-art NVIDIA GPU in each node. These GPGPUs could be programmed using the

new CUDA programming language (Compute Unified Device Architecture, a "C-like" language). The DoD computing community in general needed to improve simulation performance and to make modifications to simulation programs such as JSAF, enabling them to take advantage of heterogeneous HPC architectures. Programming models, code examples and practice problems in CUDA were developed, drafted, documented and presented and test codes were implemented in class.

Slides were delivered in three ways:

- Projected in class
- Handed out as materials
- Posted on a web site

Figure 9. Class Setting showing Large Display Monitors

In each of the classroom settings, the video was excellent. No problems with interfacing with the projection systems were encountered, avoiding the issue with which most people are all too familiar, someone standing at a podium getting frustrated with a non-compatible computer video-out problem while the rest of the crowd sits getting bored.

All of the authors having taught at length, they found their experience reinforced in any number of ways. One of these is the necessity of having good materials available. A white board with quick erase pens or a large pad of newsprint with markers is essential. Class discussions, unexpected topics, or interactive list creation all work better with the old hand-written communications.

This is another area where good planning and a dry run pay big dividends in an effective class. Assessing that effectiveness will be done real-time by good instructors, who can sense the interest of the class and can perceive the impact of the message. It would be advisable to select instructors who are good at eliciting class participation using traditional methods.

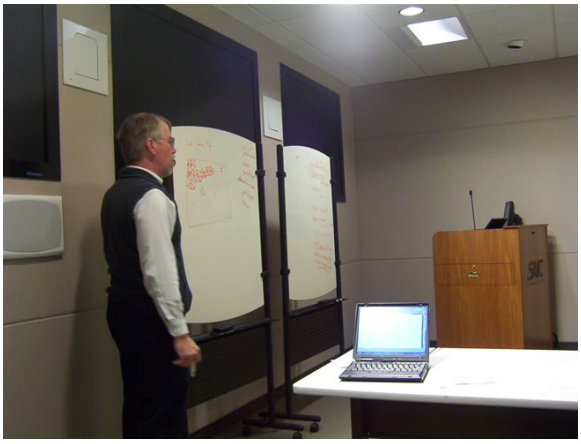

Figure 10. Gene Wagenbreth using Old Technology

Another issue in this and the latter two courses, was the provision of some computing platform that had a CUDA-capable GPU in it. Finding which ones are CUDA compatible is easily accomplished through the CUDA website (CUDA, 2010). The following (Table 1) is a representation of the time devoted to preparing for and delivering the first class. The time expended

after the course was largely in monitoring the HPCMP's formal course evaluation process. The 27 students who took the class gave this course one of the best evaluations in HPCMP history.

Table 1. Milestones

Date	Milestone	Responsibility and/or Objective(s)
Sept 2007	Kickoff (Via TelCon or VideoCon)	Collective TelCon meeting to formalize the objectives and timelines
23 – 25 Oct 2007	Training Course	To present a training course as specified in section 3.0.
16 Nov 2007	End of PoP	Completion of the project

Second Course - 2008

The success of the first course was followed by an immediate request by the researchers at ISI who had been excluded from the Suffolk course, largely by the lack of travel funds. HPCMP agreed to sponsor a second course and this one was scheduled for October of 2008 and was sited at the ISI building in Marina del Rey California, just north of Los Angeles International Airport (LAX).

Figure 11. Course 2 Site - Information sciences Institute in Marina del Rey

The setting was again quite conducive to the presentation of materials. The room reserved for the class was on the 11th floor, with windows overlooking the marina, but high above the distractions of traffic. Ample room and good ventilation make a good learning environment. The instructor was Gene Wagenbreth and the materials were:

- More focused
- Truncated
- More practicum oriented

Third Course - 2009

The putative success of the first two courses led to the request for a third, this time to be held in San Diego. It was scheduled for the Spring of 2009 and was held at the San Diego Supercomputer Center, which is actually just north of Lo Jolla California.

Figure 12. Course 3 Site - San Diego Supercomputer Center

The room selected was a classrooms setting with ample space, ventilation, and projection equipment. Approximately 24 people attended all or part of the presentations. Again, Gene Wagenbreth was the lone presenter. In general, this is not a recommended procedure. Even the most spirited presenter has trouble holding the attention of a class for days on end. Team teaching has its benefits, should more than one instructor by available.

Figure 13. Class Setting for Third Course, at SDSC

In San Diego, the projection equipment was excellent, but the setting did not have as good access to food during breaks as the other two sites had. This occasioned the thought that break and eating facilities also need be considered in selecting a site for training. Logistics may be a major driving factor in site selection in many cases.

Figure 14. Class Setting Showing Large Screen Monitor

As noted before, this course was made different because one of the students recognized that he had a CUDA-capable lap top computer, so was able to try out CUDA segments real time. This added a note of excitement to the proceedings that were a change in mood compared to earlier classes.

Fourth Course/Symposium - 2010

As mentioned before, this was a class presented to masters and doctoral graduate students at the University of Southern California. The class was made up of some 35 students, the majority of whom were marginally facile in English. Their intellectual competence was not in question, as USC's graduate school of engineering has ranked in the top ten nationally for more than a decade (US New and World Reports, 2010). This course was not taught as a practicum, but as a straight four-hour lecture. Nonetheless, several students raised questions or tried the computing examples, as was evidenced by comments and questions to the instructor and by reports given back to the professor.

Analysis and Lessons Learned

One issue that has been left unmentioned is the perplexing one of the desirability of formal training for computer science skill improvement. The authors in no way think this is a settled issue. One of the best, most innovative and productive programmers they know is a Caltech physicist who claims never to have had a class in computer science, but to be self taught, entirely from introductory programming language manuals. One of the authors was a mentee of a an Air Force General Officer who, as a company grade office, ran into a problem getting programmers to address a problem at one of the nation's best known test sites. He and two of his non-computing officer colleagues went and bought Kernighan and Ritchie and another book on C, taught themselves how to program, wrote their own program and solved the problem. This was so successful, that, at the time of the telling, some twenty years later, their code was still in use at that site.

That position is easily countered with the more universal experience of the venerable Fred Brooks who noted that, while solving a quick problem may be one thing, producing a

million-lines-of-code project takes more discipline and organization. (Brooks, 1995) Courses help the student get a better overview of what is commonly accepted practice, a broader view of what is available and a more collegial view of who is working in the field.

In each of the courses, the students were interested, committed and attentive. They quickly understood the programming concepts being advanced and they all saw and appreciated the similarities in the programming paradigms of C and of CUDA. Their ability to abstract, recognize and isolate appropriate targets for GPGPU processing remains slightly unknown and unknowable. Like parallel programming, the critical skill is not readily apparent and only time seems to be able to winnow out those who do not easily understand the complex calculus of the intersecting and coupled relationships.

Several students have contacted the authors for help after the classes, indicating there is some activity that was initiated or spurred on by the courses. Many of the original students from Suffolk were almost immediately re-assigned to other projects in an organizational move ostensibly unrelated to the class, so little progress on that front has been recorded.

Figure 15. Course in GPGPU Programming - JFCOM, Oct 2007

Another dimension is that of informing the community at large of the practicality of these approaches. While conference papers are one way, the course has a deeper and more pervasive impact on the community at large. If considered as an evangelical technique alone, the courses have proven to be a very good way to inject this technology into the simulation community.

Conclusions

The authors feel safe in concluding that the courses, as structured, achieved their immediate goals:

- Introduction of the new capability to users
- Establishing basic skills in CUDA
- Providing an overview of system analysis skills for heterogeneous programming

It seems clear that the courses were well received and appreciated. It similarly appears to be that a senior programmer or, better yet, a team of two or three senior programmers, could easily master CUDA programming, tailor the NVIDIA material for their own particular community, and effectively present a similar course to other programmers.

Whether this would be superior to just letting each individual programmer "boot-strap" themselves up via on-line and commercial materials is a less clear analysis. It would seem to the authors that this would depend on the size of the technical audience, the real heterogeneous or parallel programming experience held by the prospective instructors and other factors too numerous to name. In any case, the authors stand ready to discuss these or other issues that may be encountered by the prospective instructors.

Acknowledgements

This work was directed and funded by the Joint Experimentation Group at the Joint

Forces Command and the authors wish to thank Major General Woods, Anthony Cerri, Jim Blank, and the entire J9 staff. The authors especially would like to acknowledge the direction, support and counsel of Rae Dehncke who has been the program manger and guiding light for this project. The authors would like to acknowledge the excellent technical support provided by the members of the Computational Sciences Division of the Information Sciences Institute, Gene Wagenbreth and John Tran, as well as the guidance and assistance from members of the Scalable Systems Division there, the Dr.s Ke-Thia Yao and Robert Neches. None of this could be accomplished without the help of the JSAF team Dr. Andy Ceranowicz, Mark Torpey, Bill Hellfinstine and many others.

References

Brooks, F.P., (1995) *The Mythical Man-Month: Essays on Software Engineering*, Anniversary Edition (2nd Edition), Addison Wesley Professional, Reading Massachusetts

CUDA (2010), *CUDA GPUs*, retrieved form the internet via IE on 12 May 2010 from: http://www.nvidia.com/object/cuda_gpus.html

Davis, D.M., (2010), privately purchased a CUDA capable GPU at Fry's Electronics in Manhattan Beach California for $39 in January of 2010

Kepner, J., (2004), HPC Productivity: An Overarching View, *International Journal of High Performance Computing Applications*, Volume 18, Issue 4, Pages: 393 - 397

Kernighan, B.W. & Ritchie, D.M., (1998), *C Programming Language (2nd Edition)*, Prentice -Hall, Upper Saddle River New Jersey

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

Lucas, R.. F. & Davis, D. M., (2003) Joint Experimentation on Scalable Parallel Processors, in the *Proceedings of the Interservice/Industry Training, Simulation, and Education Conference*, 2003

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

U.S. News and World Reports, (2010), *The Best Engineering Schools*, retrieved from the internet on 12 May 2010 from: http://grad-schools.usnews.rankingsandreviews.com/best-graduate-schools/top-engineering-schools/rankings