The
ITEA Journal
of Test and Evaluation
Information Systems
Test & Evaluation
June/July 2007
Volume 28, Number 2

Featured Capability

**Arnold Engineering Development Center (AEDC)
Hypervelocity Wind Tunnel 9**

TECHNICAL PAPER ABSTRACTS:

Software Assurance Testing: A Market Survey

Steve Mackie
Wyle Laboratories, Aerospace Group, Arlington, Virginia

The consumer software industry today lacks any standards for software assurance. This market survey examines the readiness of the commercial software testing laboratory market to test software to a standard of assurance when such a standard is established. This article points out the relative immaturity of the consumer software market, lack of government oversight, and the lack of adequate laboratory accreditation programs to manage software testing laboratory accreditations. Despite these market shortfalls, there are encouraging signs of market maturity with the development of tailored personnel training and certification programs, automated testing tools, and the promotion of state-of-the-test-industry methods and best practices. According to the Committee on National Security Systems Instruction No. 4009, "National Information Assurance Glossary," Revised 2006, "Software Assurance is defined as the level of confidence that software is free from vulnerabilities, either intentionally designed into the software or accidentally inserted at any time during its lifecycle, and that the software functions in the intended manner."

Managing Large Distributed Data Sets Using the Storage Resource Broker

Dr. Reagan W. Moore
San Diego Supercomputer Center, La Jolla, California

The Storage Resource Broker (SRB) data grid enables the formation of shared collections that span administrative domains and multiple storage resources. A typical collection may have over a million files of both simulation and observational data. Uniform access mechanisms, uniform global names and uniform access controls are applied to the data, no matter where they are located. This makes it possible to automate data analyses across the distributed data, removing the need to aggregate data into a single location. The SRB is generic infrastructure that supports the requirements of multiple data management applications, ranging from data grids for shared collections, to digital libraries for data publication, to persistent archives for data preservation, to real-time data collection federation.

continued…

Using Sequential-Designed Experimentation to Minimize the Number of Research and Development Tests

Dr. Thomas F. Curry
Northrop Grumman Information Technology,
Colorado Springs, Colorado
Dr. Seung J. Lee
Defense Threat Reduction Agency,
Fort Belvoir, Virginia

Northrop Grumman provides advisory and assistance services to the Defense Threat Reduction Agency (DTRA). DTRA is interested in quantifying the attributes of secondary debris in rooms adjacent to munition bursts that result from weapons detonated in urban environments. If the detonation leaves the wall intact, the debris in the room is usually generated by bomb fragments that penetrate the wall, and by plaster or other surface material spalled off the back of the wall by the shock wave from the bomb. Little data exist on these phenomena, so tests need to be developed to collect information. Responses to the solicitation include a three-year test plan containing an arbitrary set of 421 tests. Subsequent analysis shows that a sequential-designed D-optimal experiment can obtain the required information in 74 (or fewer) tests. This 80-percent reduction in experimentation cost saves approximately $2,657,600.

U.S. Submarine Sonar Acoustic Rapid COTS Insertion (ARCI)

George Maris
Passive System Engineering Branch, Code 1542
Naval Sea Systems Command (NAVSEA)
Naval Undersea Warfare Center (NUWC), Division Newport, Rhode Island

By committing itself to the Acoustic Rapid COTS [Commercial Off-the-Shelf] Insertion (ARCI) process, the U.S. Navy has reduced its costs for developing and maintaining the sonar infrastructure for submarines. The Naval Sea Systems Command has demonstrated that COTS acquisition and a technology insertion program are cost-effective methods for achieving affordability and technical currency of ship hardware systems. In conjunction with the related advanced processing build software development effort, the ARCI program has enabled the submarine force to retain a position of underwater acoustic superiority, in spite of stealth advancements abroad and cost constraints at home. The ARCI program has simultaneously reduced not only the time it takes to introduce new hardware technology from 10 years to two, and but also has reduced the advance processing builds to a yearly process. This article reviews ARCI-related processes and the lessons learned—both positive and negative—about the program's application.

continued…

Test and Evaluation Opportunities for Information-Centric Operations

Matthew Bogusky, Kim Cain and Brian R. Gattoni
Test and Evaluation Management Center (TEMC),
Defense Information Systems Agency (DISA), Arlington, Virginia

This article discusses the need for an environment where development and testing of joint information technology capabilities can occur by leveraging Service and agency efforts into a singular federation of resources. Additionally, test and evaluation (T&E) agencies must establish policy for information sharing to minimize duplicative testing. The Defense Information Systems Agency is taking the initiative to answer the challenge of a federated virtual environment and identifying policy changes necessary to improve T&E execution.

Test on Demand: Continuous Testing Model for Agile Environments

Judith M. Hill
Test and Evaluation Management Center (TEMC),
Defense Information Systems Agency (DISA), Arlington, Virginia

To be an enabler of net-centric operations, test and evaluation must embrace the same Service-oriented architecture (SOA) framework used in the development environment. To do this will require comprehensive changes in the testing process and infrastructure. A new testing paradigm is required that is continuously available, fast and agile. Most important, testing services must be available on demand wherever and whenever development is taking place.

continued…

Modeling and Simulation Framework for Real-Time Services Over the U.S. Army's WIN-T Network

Michael A. Vincent and Joan M. Smith
U.S. Army Test and Evaluation Command,
Aberdeen Proving Ground, Maryland
Radhika R. Roy and Ira J. Hines
Science Applications International Corporation,
Abingdon, Maryland

A framework is provided of modeling and simulation criteria primarily for real-time services over the U.S. Army Warfighter Information Network-Tactical (WIN-T) network that may cover many geographical areas across the world. The modeling schemes for mobility, obstacles and routing of the tactical network, including variable topology and many widely spread mobile ad hoc networks, as well as the large-scale, fixed topology backbone network, are described in this article to analyze the performance specially related to real-time services.

Wireless Network Testing and Evaluation Using Real-Time Emulation

Dr. Sheetalkumar R. Doshi, Dr. Unghee Lee and Dr. Rajive L. Bagrodia
Scalable Network Technologies, Los Angeles, California

Full-scale physical testbeds used in network testing and evaluation of wireless network systems have certain limitations that do not allow testers and evaluators to obtain results in an easy way. Considerable time, effort and resource costs have to be invested to obtain comprehensive performance results on such testbeds. This article proposes an alternative approach to network testing and evaluation based on the real-time emulation capability of the QualNet simulator. This article provides an overview of the approach and gives illustrations to outline how a testbed based on real-time emulation can be set up. The effectiveness of such a testbed is highlighted by a sample use case, where performance evaluation results are presented for a wireless ad hoc network that is to be deployed over a range of scenarios.

PRESIDENT'S CORNER

Greetings and thank you for your continued support of ITEA. I am devoting this summer's column to reporting on the progress that we are making in improving ITEA, followed by my thoughts on a test capability whose time has come: embedded instrumentation.

We are continuing to make real progress in developing strategic initiatives to support ITEA's mission and goals. The broadened ITEA mission to reach out to the aerospace and transportation industries continues to manifest itself in such activities as expanded content for educational and workshop events that are relevant to these industries, annual symposia program tracks, leadership diversification into non-Department of Defense (DoD) sectors and other areas.

The ITEA Corporate Development, Chapter and Individual Membership, Publications and Communications committees have been very active in the past quarter developing short-and long-term objectives that hold the promise of exciting new programmatic activities. Included for 2007 are such projects as reconstituting committee membership to increase participation; revamping the ITEA web site and adding value to online services and benefits such as a job board, intern program and electronic *ITEA Journal* archives; providing additional support to chapter leaders and their constituencies through broader communication efforts, including an online chapter forum; exploring international outreach and engaging our current global partners to feel a part of the ITEA community; along with numerous other initiatives. I hope you'll plan to join me at the 2007 Annual ITEA International Symposium in Kauai, Hawaii, November 12-15, where I will be reporting on the success of many of these and other initiatives. I also encourage your involvement. If any of these projects are of personal or professional interest to you, please call or e-mail me to see how you can support their success.

Now, I would like to share my thoughts on the many advantages of embedded instrumentation, which provides a "cradle to grave" capability to gather data about the system on which it is installed. This capability, referred to as continuous test and evaluation (CT&E), supports the full life cycle support of the platform, to include developmental test (DT), operational test (OT), training, maintenance and live fire damage assessment. The Iraq War has taught us the benefit of data gathered during a conflict to aid in enhancing the effectiveness of a weapon system. Because we continue to develop separate instrumentation for test and training, embedded instrumentation would force us to develop a single instrumentation package, thus realizing development cost savings. Embedded instrumentation also saves retrofit costs incurred when we add and remove instrumentation for DT, then for OT, then for training.

Thomas J. Macdonald

Embedded instrumentation can utilize onboard systems. The Extended-Range Guided Munition intends to embed a datalink that would provide testers with data from the onboard Global Positioning System (GPS) navigation system. The miniaturized embedded instrumentation can be accommodated in the airframe without removing the warhead. Hence, every weapon shot will provide both live fire weapon effectiveness and missile system performance data. There will be no need (with the attendant cost) to expend inert munitions to test system performance. It is worth noting that, if the onboard GPS navigation receiver is properly chosen, the downlink data can be augmented with ground-based data to significantly enhance the inherent navigation accuracy, thus allowing the test community to use the data as a truth source.

There are many ongoing "non-intrusive" development projects that provide miniaturized instrumentation. They are referred to as "non-intrusive" and not "embedded" because the weapon developer makes the "to embed" decision, not the instrumentation development project. The issue here is that, in most cases, the weapon developer chooses not to embed. Why? Weapon development is time consuming and is often plagued with schedule slips and cost overruns. In this hectic environment, a weapon development program manager is unwilling to add further risk to the program by taking on the additional effort of providing this "add on." In addition, the cost savings of embedded instrumentation are not realized by the weapon developer. Nonetheless, these savings benefit DoD and, hence, the taxpayer, and these life cycle cost savings are significant.

The Under Secretary of Defense for Acquisition, Technology and Logistics (AT&L)—the department responsible for all DoD acquisitions, as well as for system testing—should consider making embedded instrumentation a required component of all weapon system acquisition/development projects. This top-level direction is essential to making this change happen. As always, your comments are welcomed and encouraged. Please e-mail me at: **comments@itea.org**.

"Born Joint"—Is The T&E Community Ready?

Dr. Steven J. Hutchison
Defense Information Systems Agency,
Arlington, Virginia

A quick Web search on the phrase "born joint" will show that this concept has been around for at least a decade. So, it may be a bit of a surprise that here, in 2007, I am asking if the test and evaluation (T&E) community is ready for systems that are born joint. And, just what do I mean by asking, "Is the T&E community ready?" More precisely put, my question asks if the Department of Defense (DoD) T&E community has the policy, infrastructure, methods and processes in place to execute Joint Operational Test and Evaluation, or JOT&E, for joint acquisition programs. My premise is simple: If a system is born joint, it should be tested joint. And my assessment today is that we are not ready to "test joint." That said, we do have many of the elements of joint testing being developed, but we do not have a plan to cause them to converge on realistic Joint OT&E. Because we do not have realistic Joint OT&E, the T&E community plays a diminished role in helping DoD achieve information superiority.

Dr. Steven J. Hutchison

"Information Systems Test & Evaluation" is the theme of this issue of the *ITEA Journal*. The T&E community has a vital role to play in helping DoD transform and field a network-centric joint force capable of achieving full spectrum dominance. Recall how Joint Vision 2020 (JV2020) described the role of information technologies: "The transformation of the joint force to full spectrum dominance rests upon information superiority as a key enabler…." *(See Figure 1, next page.)*

JV2020 depicts information superiority as encompassing all the elements of combat power. Our role as testers is to ensure that the information technologies procured by DoD function in the joint mission environment, provide the necessary capabilities, are supportable, and that their capabilities and limitations are understood by the warfighters.

True "jointness" occurs in the exchange of information and collaboration between forces with different missions and capabilities that achieves greater effects on the battlefield. This guest editorial is not about testing platforms or weapons used by multiple Services (an example being the high-mobility multipurpose wheeled vehicle, better known as the HMMWV); it is about our community's readiness to test and evaluate systems (or managed services) acquired *to improve information sharing* between forces in a joint mission environment.

When we look at our current processes for conducting OT&E in a joint mission environment, we see many initiatives and some progress, but no convergence toward true Joint OT&E. Joint OT&E is not a new issue; as far back as 1970, a Blue Ribbon Defense Panel concluded the following:

> *"Currently, there is no effective method for conducting OT&E which cuts across Service lines, although in most actual combat environments, the United States must conduct combined operations….Because of the lack of joint OT&E, it is not only very difficult to detect certain kinds of deficiencies and to predict combat capability in advance, but it is also difficult to make decisions relating to overall force composition."*[1]

Interestingly, this conclusion brought the Joint Test and Evaluation (JT&E) program to DoD. But JT&E is not Joint OT&E—at least not in the sense of OT&E as part of the acquisition process (more on this in the policy discussion to follow).

Just what *is* needed to conduct Joint OT&E today? Joint OT&E requires *policy, processes, methodology* and *infrastructure*. We have many of the ingredients needed to do this today; however, some are for soup, others are for a sandwich, and you know where that leaves us. What follows is a brief discussion of how to pull this together into a coherent capability.

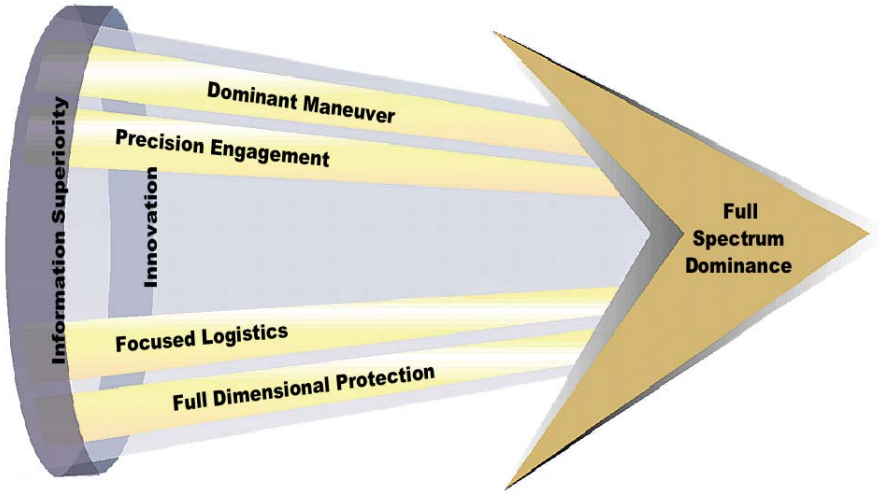

Figure 1. Full spectrum dominance

Policy

Current policy for conducting Joint OT&E for acquisition programs does not exist. In vague terms, responsibility for Joint OT&E lies with the Office of the Director, Operational Test and Evaluation (DOT&E). DoD Directive 5141.2 states: The DOT&E shall "coordinate Joint Operational Testing," and "Oversee Joint Operational Test and Evaluation programs…"[2] It is not clear if the directive intended Joint OT&E for acquisition programs, or the JT&E program as we have it today. The latter is alive and well; however, the former continues to be illusive. The key difference is that JT&Es focus on doctrine and tactics, techniques and procedures (TTPs) to fill gaps in interoperability between fielded systems.

To quote the JT&E Program Briefing: "Joint operational solutions are the principal products of JT&E projects—not hardware acquisition products."[3] Joint OT&E, as used here, is OT&E on materiel solutions managed by a joint program manager, for the purpose of informing the acquisition decision-making process.

Likewise, the DoD 5000 series does little to compel Joint OT&E. In fact, a quick scan of the DoDD 5000.1 and DoDI 5000.2 shows not one instance of the term "joint test," or "Joint OT&E," or similar phrases; yet there are numerous references to joint programs, joint concepts, capabilities, architectures, training and so forth. I find it interesting that we can write in the DoDI 5000.2: "…the user shall prepare the Capability Development Document (CDD) to…clarify how the program will lead to joint warfighting capability…" but not mandate that the initial (I)OT&E should be designed to confirm effectiveness and suitability of that joint warfighting capability.

These policies describing the joint user's role emerged in June 2003, in the Chairman of the Joint Chiefs of Staff Instruction (CJCSI) 3170.01C, Joint Capabilities Integration and Development System (JCIDS), replacing the Requirements Generation System. In doing so, the Chairman ended the stovepiped, Service-centric requirements definition process and began the process to define capabilities in joint terms.

So, we have directives and instructions describing responsibilities of the joint materiel developer, as well as the joint combat developer, yet we still do not have policy for the joint tester.

This was echoed in the Fiscal Year 2006-2011 Strategic Planning Guidance (SPG):

> *"Developing and fielding joint force capabilities requires adequate, realistic T&E in a joint operational context. To do this, the Department will provide new testing capabilities and institutionalize the evaluation of joint system effectiveness as part of new capabilities-based processes."*

In November 2004, in response to the SPG, DOT&E published the "Testing in a Joint Environment Roadmap." The Roadmap describes the shortcoming this way:

> *"Current Service T&E capabilities are world-class, but focus primarily on testing in a system-centric operational environment that does not fully reflect the complexity of the joint environment.…To develop and field capabilities 'born joint,' the Department needs a more robust, focused, and tightly coupled T&E capability that places testing in a joint environment and joint interoperability testing at the core of T&E activity, rather than as an extension to system-centric testing."*

The Roadmap reviewed 30 policy documents to identify change areas to accomplish the SPG task to institutionalize testing in a joint operational context.

Work is progressing, but we clearly need new policy underpinnings for Joint OT&E. Perhaps it is time to write a DoD Directive on Joint OT&E.

Processes

Our current processes for Joint OT&E are also not well-defined. Almost all OT&E performed by the Joint Interoperability Test Command (JITC) is joint. JITC is the Operational Test Agency (OTA) for numerous programs of record that serve the joint warfighting, intelligence, business and enterprise communities. Additionally, the Joint Systems Integration Command (JSIC) conducts joint testing for systems being acquired under Joint Forces Command's Limited Acquisition Authority.

Joint testing also occurs when OTAs work together under the processes documented by the Service OTAs in their memorandum of agreement (MOA) on Multi-Service OT&E (MOT&E).[4] However, MOT&E is usually conducted in a Service-centric fashion, using the lead Service's processes and reporting procedures. This MOA lays a good foundation for bringing the OTAs together in what could be a framework for Joint OT&E, but MOT&E, as defined today, is not Joint OT&E; nor is MOT&E required by policy.

An example of Joint OT&E is a recent operational event for the Joint Operational Planning and Execution System (JOPES), part of the Defense Information Systems Agency's (DISA's) Global Command and Control System-Joint (GCCS-J) system. The OT&E involved users from five Combatant Commands (COCOMS) and more than 10 subordinate commands across three continents (PACOM, SOCOM, CENTCOM, EUCOM and TRANSCOM).

For a strategic evaluation, the National Military Command Center (NMCC) also served as a participant, and all players received centralized support by the Joint Staff Support Center (JSSC). The COCOMs and their subordinate commands executed joint mission tasks for a period of 10 days, exchanging data and executing missions relative to real-world operations. JITC served as the OTA for this event. It was performed in an operationally realistic joint mission environment, with typical joint users executing joint mission tasks—in other words: Joint OT&E.

The Joint OT&E concept can be built on the MOT&E framework, but it must define the Joint Mission Environment (JME) and provide the methodology to be used if it is to be successful.

Methodology

We have been able to force jointness into the requirements process via JCIDS. But what is the forcing function for jointness in OT&E? We have a Roadmap, an emerging Joint Mission Environment Test Capability (JMETC) and an emerging Joint Test and Evaluation Methodology (JTEM). Even if we had something to force Joint OT&E, how would we execute a Joint OT&E? Who writes the joint test plan, and what does it look like? How do we go about getting commitment of joint test units? Who is the joint user representative? These questions pose significant challenges to our parochial test organizations.

Fortunately, some of these questions are being addressed today. We are seeing some development of Joint Mission Threads (JMTs) as a vehicle to test new capabilities. The JMTs are being developed under the Joint Forces Command (JFCOM) Joint Battle Management Command and Control (JBMC2) effort. To date, however, there is but one of the seven JMTs completed. Additionally, one of the chartered JT&Es, the JTEM, is working to "develop, test and evaluate methods and processes for defining and using a distributed live, virtual, constructive joint test environment to evaluate system performance and joint mission effectiveness." As of this writing, JTEM has already produced initial draft templates of test plans and other test artifacts.

JTEM is paving the way to answering some of the questions I have already posed. For example, once the joint capabilities have been identified, the Joint Test Agent will likely begin test planning and decomposing measures of effectiveness from critical operational issues. The well-understood framework of critical operational issues is being addressed in the joint world by JTEM as a critical joint issue (CJI).[5] JTEM proposes the following definition of a CJI:

> *"Those issues that can be used to measure joint program effectiveness. A critical joint issue (CJI) for test and evaluation should be carefully structured as a question addressing the key capability attributes described in Joint Capability Documentation (Initial Capabilities Document and Capability Development Document)….The critical joint issues should address the system of systems capability to perform joint operational tasks and/or the system of systems, system, or service attribute performance. CJIs are of primary importance to the decision authority in reaching a decision to allow the system of systems to advance into the next phase of development."*

This is precisely the kind of thinking that needs to take place to move from Service-centric OT&E to Joint OT&E. And these are exactly the kinds of challenges being confronted today by the Net Enabled Command Capability program. We must move ahead and resolve these issues.

Infrastructure

Defining the JME will be essential to adequate Joint OT&E. As a result of the "Testing in a Joint Environment Roadmap," a new program element was established in the Fiscal Year 2007 budget that funded (minimally) initial development of the JMETC. The JMETC is envisioned as an enabler for distributed T&E, engineering, training and experimentation by providing the necessary joint context.

As a persistent capability, it will reduce the need to establish unique testbeds for systems and the stovepiped testing that characterizes T&E today. To do this, it must develop and implement the infrastructure solution and provide the needed customer interface and support. On February 13, 2006, JMETC was transferred to the Test Resource Management Center (TRMC) in the Office of the Under Secretary of Defense (USD), Acquisition, Technology and Logistics (AT&L).

While we in DISA fully endorse the JMETC, getting the JMETC going is not without its issues. At the close of Fiscal Year 2006, the T&E Board of Directors, Executive Secretariat (BoD[ES]), wrote to express its concerns to the TRMC regarding the JMETC concept. The BoD(ES) wrote: "While the BoD(ES) endorses the [Secretary of Defense's] SECDEF's vision as expressed in the 'Testing in a Joint Environment Roadmap,' the BoD(ES) cannot endorse the current approach to satisfying that vision as embodied in the JMETC."[6]

Summary

So, are we ready to do Joint OT&E? No. However, we have some of the right ingredients. We need to make Joint OT&E a matter of policy. We need to embrace the intent of JMETC and fund it accordingly. We need to ensure the JTEM continues on its present course and is well-supported. We need to create environments in which joint warfighters can train, test and evaluate. But more importantly, we need leadership to make this vision happen. When we do this, then the T&E community will be a relevant part of the acquisition of information technologies that enhance the capabilities of the joint warfighter. ❑

DR. STEVEN J. HUTCHISON, a member of the Senior Executive Service, is the Test and Evaluation (T&E) executive for the Defense Information Systems Agency (DISA) and the director of the T&E Management Center (TEMC), Arlington, Virginia. He is responsible for developing and enforcing T&E policy and procedures; representing DISA to the Department of Defense (DoD) T&E community; and providing direct support to DISA programs for T&E of information technology (IT) capabilities. The TEMC coordinates and oversees investment funding for research, development, test and evaluation (RDT&E) and the IT testbed in the Major Range and Test Facility Base (MRTFB). Prior to his arrival at DISA, Dr. Hutchison served in the Office of the Director, Operational Test and Evaluation (DOT&E), Office of the Secretary of Defense, as a net-centric warfare systems analyst. While in DOT&E, he had oversight responsibilities for several of the major warfighting information systems in DoD, including the Global Command and Control System–Joint (GCCS-J), the Service variants of the Distributed Common Ground/Surface System, and the Net Enabled Command Capability (NECC).

Endnotes

[1]Report to the President and the Secretary of Defense on the Department of Defense by the Blue Ribbon Defense Panel, Department of Defense, Washington, DC, July 1, 1970.

[2]DoDD 5141.2, Director, Operational Test and Evaluation (DOT&E), May 25, 2000, paragraph 4.2.6, and paragraph 4.2.8.

[3]See http://www.jte.osd.mil/docs/JTE%20Pgm%20Briefweb-Jan%2007.ppt. The briefing goes on to say that: "Test products are new joint tactics, techniques, and procedures; architectures; and processes."

[4]JITC is not a signatory to the Multi-Service OT&E (MOT&E) memorandum of agreement.

[5]JTEM Lexicon Version 1.0, January 30, 2007. See https://www.jte.osd.mil/jtemctm.

[6]Memorandum for Director, Test Resource Management Center (TRMC), Subject: Joint Mission Environment Test Capability (JMETC) Requirements, dated August 31, 2006. Memorandum is signed by the Chairman of the Board of Directors, Executive Secretariat.

A Brief Look at FPGAs, GPUs and Cell Processors

Michael L. Stokes
Ohio Supercomputer Center, Columbus, Ohio

In recent months, there has been much discussion about new advancements in processor technology that promise huge performance returns for a small investment. Field programmable gate arrays (FPGAs), cell processors and graphics processor units (GPUs) are all the rage. So, one question is, "What are the salient features of this new technology?" But perhaps the most important questions, however, are: "What does all of this mean to the test and evaluation (T&E) community? Should these new technologies be adopted, put on the top shelf until the technology matures, or [should they be] written off altogether?"

FPGAs, invented around 1984 by Ross Freeman, Xilinx cofounder, have a relatively large number of programmable logic components with programmable interconnects. The logic components can be programmed to duplicate basic logic functions (AND, OR, XOR and NOT) or grouped to form simple mathematical functions such as integer addition. A complex circuit can be constructed from logic and memory components similar to a programmable breadboard. These components can be reprogrammed after the manufacturing process; thus, they are *field programmable*.

Programming an FPGA is accomplished using hardware description language, which is very similar to assembly language programming for general-purpose central processing units (CPUs). The software developer must be aware of hardware timings, interrupts, signals, thread synchronization and other interfaces at the hardware level. Several research efforts are underway to develop higher-level languages for programming FPGAs to make software development easier and faster.

Typical applications that use FPGAs are cryptography, specialized routers or network edge devices, and medical imaging. These applications can be characterized as those that perform a substantial amount of computation with a small amount of memory, such as the Fast Fourier Transform (FFT). Logic designs that are independent of one another execute in parallel (up to available memory); therefore, multiple FFTs, for example, execute in the same time to execute just one. In this sense, FPGAs are highly scalable. The down side of FPGAs is that limited onboard memory and input/output (IO) bandwidth restricts scalability. FPGAs also do not support floating point operations, thus (intrinsically) limiting the appeal for programming broader applications found in the T&E community.

A relative newcomer to general-purpose programming is the GPU *(see Figure 1)*. A GPU is that fancy graphics card that costs a few hundred dollars for a PC to make

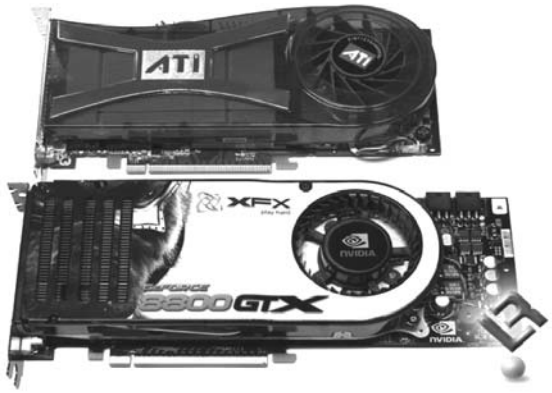

Figure 1. NVIDIA and ATI graphics processor units (GPUs)

video games vivid and realistic. The popular association of GPUs is with accelerating graphics, but the new architectures from manufactures such as NVIDIA Corporation and ATI are capable of performing general-purpose computing in addition to making animated monsters look more life-like. For a good overview of this exciting area, visit ***http://www.gpgpu.org***.

There are two approaches to consider when programming a GPU for general-purpose computing. The first is to pose the problem as a graphics problem and solve it using a graphics language such as OpenGL ***(see http://opengl.org)*** or DirectX ***(http://www.microsoft.com/directX)***. The second approach is to program the GPU directly. But be cautioned: This is not a conventional method of programming, so be prepared to think about the application in some new ways, because what is needed today to program GPUs may not be what is needed in the future.

The first approach to using GPUs is to recast the general-purpose application as a graphics problem, which is not always possible or simple to do. OpenGL or DirectX calls can be employed to define the geome-

try, then used to query the geometry to discover distances, intersections and other properties. The so-called *Line-of-sight* program has been solved this way using various techniques available in OpenGL. This approach is very nice when it works because GPUs are already optimized for OpenGL and DirectX in hardware. The programmer has to know little or nothing about the graphics accelerator or its capabilities. However, this model sometimes does not provide the programmer with enough flexibility. Graphics Language Shading Library is a relatively new capability in OpenGL (as of OpenGL 2.0) that allows additional instructions to be inserted in the vertex and pixel shader component of the OpenGL pipeline. While this offers substantial flexibility over the simple graphics API, the method still requires that the problem be cast in a geometric format. So, what if this is not possible?

One solution is the introduction of the Compute Unified Device Architecture (CUDA) ***(http://developer.nvidia.com/object/cuda.html)*** by NVIDIA. The CUDA environment is supported on the NVIDIA G8X and newer graphics adapters, and includes the CUDA toolkit for the Linux and Windows Operating Systems. This toolkit allows the user to access the GPU hardware directly.

The GPU hardware currently sports 16 multiprocessors. Each multiprocessor has a set of eight 32-bit subprocessors with a Single Instruction Multiple Data architecture shared instruction unit, or a total of 128 processors. Each multiprocessor has a set of 32-bit registers per processor, on-chip shared memory with fast access to the processors, a read-only constant cache memory, and a read-only texture cache. In addition, the device contains 768 MB of device memory with much slower access speeds than shared memory. Parallelism is achieved by associating up to 32 threads of execution in a group called a *warp,* of which up to 16 *warps* make up a *block.* Each multiprocessor executes one or more block(s) in parallel.

The peak performance from the NVIDIA 8800 is around 300+ GFLOPS for a unit that lists for around $600—a large performance-to-price ratio when compared to current-day high-performance computing (HPC). However, the GPU is not without its faults. In order to obtain the full throughput of the device, it is necessary to optimize onboard memory and keep the processors fully engaged, which can be difficult for some applications. There is an alternative to the GPU. It is called a cell processor *(see Figure 2)*, and it is one of the hottest options for accelerating HPC applications.

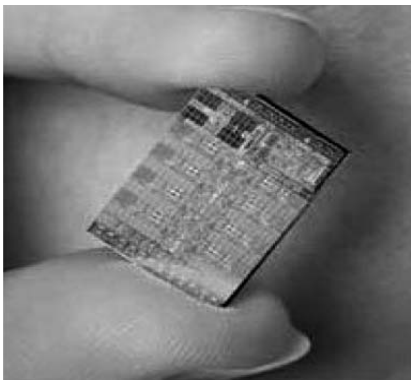

Figure 2. A cell with one Power Processor Element (PPE) and eight Synergistic Processor Elements (SPEs)

"Though sold as a game console, what will in fact enter the home is a Cell-based computer," according to Sony's Ken Kutaragi, in talking about Sony's PlayStation 3 (PS3). The Cell processor was designed from the very beginning to be a general-purpose computer, not as a gaming console, as often rumored. Most companies would have used specialized automated software in the design process, but IBM, the designer of the Cell processor (or simply Cell), decided to perform the design by hand, a choice costing millions more dollars but resulting in a more compact and power-efficient design. The resulting design is nothing less than stellar in the opinion of most industry observers due to its simplified design and power efficiency.

The architecture of the Cell is unique, although it shares a lot in common with older vector supercomputers (recall the Cray 1). The Cell can perform at rates similar to or better than GPUs, but it is much easier to program because it was designed for general-purpose application. A Cell is composed of a number of elements as shown in *Figure 3 (next page)*: one Power Processor Element (PPE), eight Synergistic Processor Elements (SPEs), one Element Interconnect Bus (EIB), one Direct Memory Access Controller (DMAC), two Rambus XDR memory controllers and one Rambus FlexIO interface.

The PPE primarily initiates and monitors jobs that run on the SPEs. The PPE runs the basic operating system and parts of the application, but compute-intensive components of the operating system and the applications are offloaded to the SPEs. The PPE is a 64-bit *Power Architecture* processor compatible with PowerPC binaries. The downside of this design is that one can expect poor performance when executing logic-heavy instructions, because it is an in-order processor (it does not pre-compute logic branches).

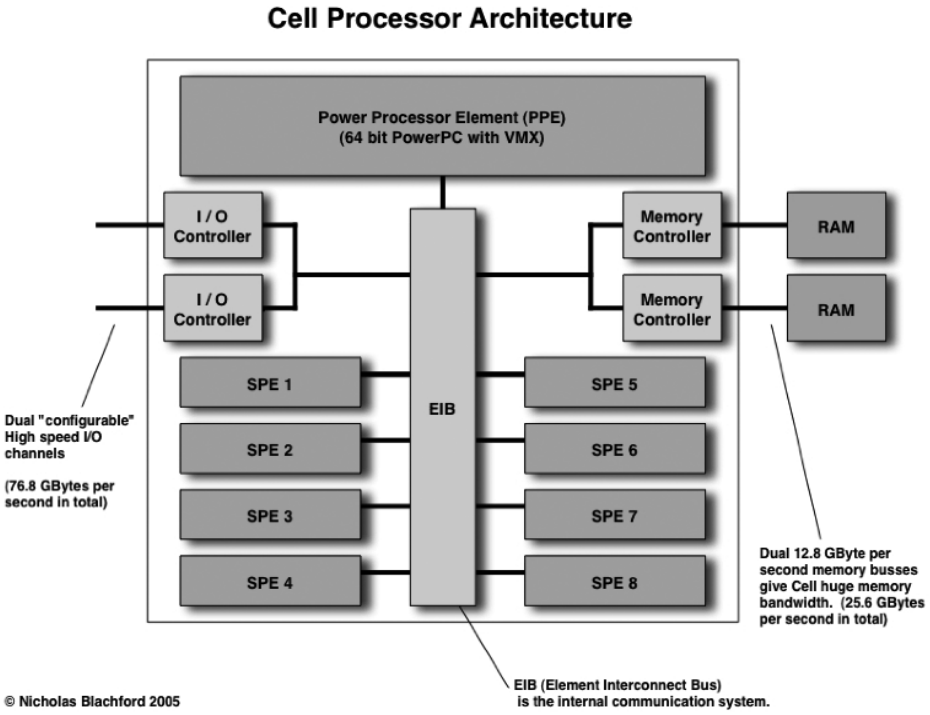

Figure 3. Cell processor architecture

SPEs, operating at clock speeds up to 4 GHz, contain 128 x 128 bit registers, along with 4 single precision floating point units with a maximum peak performance of 32 GFLOPS, 4 integer units capable of 32 GFLOPS, and a 256 K *local store* instead of a cache. Each SPE consumes less than 5 Watts at 4 GHz, making it a better performer than the current batch of GPUs (some of which recommend a 500 W or larger power supply). Like the PPE, SPEs are vector processors, namely, they perform multiple operations simultaneously with a single instruction. Each SPE is capable of 4 x 32 bit operations per cycle (8 if multiply-adds are included). The double precision calculations are Institute of Electrical and Electronics Engineers (IEEE) standard, whereas the single precision calculations are not, which is supposedly faster.

In the current implementation of the PS3, double precision calculations share the computational area of the single precision floating units. This saves a lot of room on the silicon, but it results in a large penalty in performance, or around 25 double precision GFLOPS at 4 Ghz. For a gaming box such as the PS3, double precision is not important, but IBM has hinted that future versions of the Cell might include a full-speed, double precision floating point unit that will perform at rates as high as integer performance, or 256 GLOPS per SPE. SPEs can execute instructions in parallel in synchronized loops, or can also be chained together, forming a stream processor through a process called pipelining. It is through pipelining that the Cell is capable of achieving near-peak performance.

In summary, as has always been the case, the choice of which acceleration architecture is best for any specific application has to be examined for each individual case. FPGAs, in general, are a good choice when seeking a hardware or embedded solution, while GPUs and Cells appear to be better choices for modeling and simulation and simulation and training applications. GPUs are well-suited for applications that can be cast in the form of a graphics solution and some general-purpose applications, while the Cell architecture seems to be the best choice when complex, general-purpose solutions are required. ❑

MICHAEL L. STOKES is an engineer that works for the Ohio Supercomputer Center, but resides at Redstone Technical Test Center (RTTC), Redstone Arsenal, Alabama. Funded by the High Performance Computing Modernization Program Office (HPCMPO) under the PET program, he works in the area of testing and evaluation. His current area of research is in the development of real-time ray tracers for multi- and hyperspectral synthetic scene injection. He received his doctorate degree in engineering science and mechanics at the University of Tennessee, Knoxville, in 1991. Comments and questions can be directed to him at ***michael.l.stokes@us.army.mil****, Room 27, Building 4500, Redstone Arsenal, Alabama 35898.*

Acknowledgment

Images/photographs in this article are copyrighted by, and printed with written permission from, Nicholas Blachford. The author gratefully acknowledges Blachford's contributions.