

## **FPGA and GPU Accelerators for Linux Clusters**

Robert F. Lucas, Dan M. Davis and James R. Moore  
Information Sciences Institute, University of Southern California  
4676 Admiralty Way, Suite 1001  
Marina del Rey, CA 90292  
{rflucas, ddavis, jjmoore } @isi.edu  
Telephone: 310 822-1511

### **Abstract**

The Test and Evaluation (T&E) community has been making great advances in using Modeling and Simulation (M&S) in their work. They would be even better served had they ready access to higher resolution, quicker turn-around times, more elements, and richer behavioral characteristics in their physics-based and entity-level simulations. As rapidly as it has been enabled to accomplish superior results, the T&E environment is still constrained by computational limits. High Performance Computing (HPC) can ameliorate those constraints. The use of Linux Clusters is one path to higher performance; the use of Field Programmable Gate Arrays (FPGAs) and Graphics Processing Units (GPU) as accelerators are two others. Merging these paths together holds even more promise. The authors report their experiences with the new HPCMP-provided 512 CPU (1024 core), GPU-enhanced Linux Cluster for the Joint Forces Command's Joint Experimentation Directorate (J9). They further relate their work on FPGAs, used as computational accelerators that bring with them the reprogrammable efficiency that are a complement to the GPUs powerful floating-point efficacy. Basic concepts are laid out that underlie the use of FPGAs and GPUs as accelerators for intelligent agent, entity-level simulations and for multi-frontal attacks on sparse systems of linear equations. These two disparate fields will be used to show the broad range of capability improvements projected by the authors for FPGAs and GPUs. They discuss the use of the two accelerators in tandem as well. The simulation needs of the T&E community, the ability of FPGA- and GPU-enhanced clusters to respond to T&E needs, and the careful analysis of the intersection of these are explicitly discussed. Existing configurations and potential configurations of clusters are addressed and the potential increase in performance are identified and justified. Anticipated problems and solutions will all be reported objectively, as guides to the T&E community. The paths to reliable and timely capability enhancement will be fully explicated. Early characterization runs of a single CPU with GPU-enhanced extensions are reported.