

Operational Use of a Large GPGPU-Enhanced Linux Cluster

Dan M. Davis, Robert F. Lucas and Gene Wagenbreth

Information Sciences Institute/Univ of So Calif.

Marina del Rey, California

{ ddavis, rflucas, genew }@isi.edu

ABSTRACT

This paper will focus on the effective use of General Purpose Graphics Processing Units (GPGPUs) in a 256-node Linux cluster. By using the computational power of the GPGPUs, the user wanted to attack the computational spikes that otherwise was thought to have required a cluster that would have been twice as large, expensive and power hungry. The Authors will brief how the current and future return on investment for the JFCOM DHPI GPGPU-enhanced Cluster joshua is a paradigm exemplar of leveraging technology to accomplish goals for orders of magnitude less funding and to analyze systems in social environments that cannot be disrupted in live exercises. By emulating forces that would cost tens of millions of dollars to equip and deploy and by simulating urban areas that are not open to U.S. DoD exercises, JFCOM can economically, safely and securely test new sensors, systems and strategies. HPCMP has achieved a preeminent position of professional leadership in the field of GPGPU-computing, showing the technical merit of the project. The computational merit of the project is clearly demonstrated in the achievement of the provision of adequate compute products to support several on-going exercises, one of which will be briefed. The authors carefully studied two sub-routine algorithms, Line of Sight and Route Finding. Stability of joshua in an operational setting will be explicated to show current progress. Appropriateness of requested resources in this case were a dead-on match, as joshua has exceeded the goal of 2,000,000 SAF entities by achieving 10,000,000. The authors conclude that the GPGPU approach saved purchase costs, enabled large city simulations, save energy costs and delivered simulations reliably for the users.

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 through two decades. As Assistant Director of the Center for Advanced Computing Research at Caltech, he managed Synthetic Forces Express, a major simulation project. Earlier, he was a Software Engineer on the All Source Analysis System project at JPL and did classified research at Martin Marietta. He was the principal author on the design and proposal of the 256 node, GPGPU-Enhanced cluster at JFCOM. An active duty Marine Cryptologist, he retired as a Commander, USNR. He served as the Chairman of the Coalition of Academic Supercomputing Centers. He received a B.A. and a J.D., both from the University of Colorado in Boulder.

Robert F. Lucas is the Director of the Computational Sciences Division of the University of Southern California's Information Sciences Institute (ISI) and a Research Associate Professor in the Department of Computer Science at the same University. There he manages research in computer architecture, VLSI, compilers and other software tools. He has been the principal investigator on the JESPP project since its inception in 2002. Prior to joining ISI, he was the Head of the High Performance Computing Research Department for the National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory, the Deputy Director of DARPA's Information Technology Office, and a member of the research staff of the Institute for Defense Analysis's Center for Computing Sciences. From 1979 to 1984 he was a member of the Technical Staff of the Hughes Aircraft Company. Dr. Lucas received his BS, MS, and PhD degrees in Electrical Engineering from Stanford University in 1980, 1983, and 1988 respectively.

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 tools for distributed and shared memory parallelization of Fortran programs. He has also been active in benchmarking, optimization and porting of software for private industry and government labs. He has programmed on CRAY, SGI, Hitachi, Fujitsu, NEC, networked PCs, networked workstations, IBM SP2, as well as conventional machines. He received a BS in Math/Computer Science from the University of Illinois in 1971.