

Application of Advanced Multi-Core Processor Technologies to Oceanographic Research

Mark R. Abbott
104 COAS Administration Building
Oregon State University
Corvallis, OR 97331-5503
phone: (541) 737-5195 fax: (541) 737-2064 email: mark@coas.oregonstate.edu

Grant Number: N000140610342
<http://www.coas.oregonstate.edu>

LONG-TERM GOALS

Improve our ability to sense and predict ocean processes, utilizing state-of-the-art information processing architectures.

OBJECTIVES

New processor architectures (multi-core, multi-threaded) hold the promise of delivering enormous amounts of compute power in a small form factor and with low power requirements. However, new programming models are required to realize this potential. Our objectives are to deploy signal processing algorithms onto the IBM Cell Broadband Engine architecture, and to assess the Cell B/E potential as a processing engine in an autonomous vehicle.

APPROACH

We established a partnership with IBM to have early access to Cell BE technology, including future releases along its evolutionary path. We partnered with OASIS, Inc. (Lexington, MA) to develop signal processing code for the Cell BE. We analyzed the algorithm and software, subdividing tasks, and allocating these tasks to the programming elements. Managing the communications and scheduling of these tasks is another challenge with the Cell BE. Our primary contacts at OASIS, Inc. are Phil Abbott and Vince Premus. Will Dillon (OSU) is a PhD graduate student in computer science working on this project, under the supervision of Mike Bailey (OSU Engineering).

WORK COMPLETED

We acquired and installed IBM Cell BE rack consisting of two blades, with 2 Cells per blade system. Each Cell consists of one PowerPC chip and eight Synergistic Processing Elements (SPE). We installed the Rapid Mind development environment. After initial testing and system integration, we ported the acoustic signal processing application from OASIS, Inc. to the Cell BE. We have tested this application and begun to optimize it for the Cell BE.

Target detection, tracking and identification is performed in three discrete steps, each with unique design specifications; the first of which is "Data-Independent Beamforming" (Van Veen and Buckley, 1988). This step is a frequency-domain beamforming operation across n (64, currently) beams. The

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE 30 SEP 2007		2. REPORT TYPE		3. DATES COVERED 00-00-2007 to 00-00-2007	
4. TITLE AND SUBTITLE Application Of Advanced Multi-Core Processor Technologies To Oceanographic Research				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Oregon State University,104 COAS Administration Building,Corvallis,OR,97331				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES code 1 only					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 4	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

first and second derivatives of the result are found and identify unique audio sources, which are added to a “target queue.” The second stage is “Statistically-Optimum Beamforming,” and works by collecting the maximum auditory information from the beam in the direction of the desired target and a “Null” is placed in the strongest source of interference. Finally, this signal is used to perform a machine learning operation with the intent of identifying the type of target: man-made, natural, etc.

The progress made toward a completed system is twofold. First, the architecture and system design is nearly complete, and design documents for each subsystem are in the process of being drafted. Second, the initial programming for the first stage is complete. Each subsystem runs as a standard Unix task and subsystems communicate through network sockets. The data-independent beamforming task runs as a network server, and all audio data is sent along a network link. Note that the “network link” may be virtual and constrained to the on-board computer.

RESULTS

At present, the data-independent beamformer is heavily instrumented to collect run times for various operations. Optimization of this code base for the Cell BE has not yet begun; however, baseline timing runs have been compiled on a 2.4 GHz. Core 2 Duo running Mac OS X. The most meaningful run time is the time taken to perform the FFT on input data, in this case 4.8 GFLOPS which works out to be about 600 μ s. The initial FFT benchmarks on the QS20 (IBM Blade center containing Cell BE processors) are in the 2.3 GFLOPS range and through extrapolation: about 1.2 ms. On first look, the performance of the Cell BE seems sub-optimal. However, it bears noting that we are performing 16 one-dimensional FFTs, one for each hydrophone channel, and the FFT benchmark is performing a single FFT. It is reasonable to assume that the performance of the FFT step may increase by up to 8 times when all 8 SPEs are used on their own data, which could yield 18.4 GFLOPS. This would result in a theoretical FFT computation time of 200 μ s. This assumption seems reasonable when compared to the best observed FFT benchmark on the QS20, which achieved 39 GFLOPS on a three-dimensional FFT. It is important to note that these values are significant because the ratio of time spent to complete computation of a data set must be less the time to collect it. In this case, the FFTs are performed upon 2048 samples, which represents 0.6 seconds. Thus computation is complete in 3/10,000 of the time taken to collect data, a demonstration of how much computation capability remains after the first step of processing.

IMPACT/APPLICATIONS

The evolution of computing power has followed twin paths of higher frequency processors and symmetric multiprocessing. Although Moore’s Law has held for the past several decades, we have reached the limits of heat dissipation that are required for higher speed chips that are built at higher densities. Design tradeoff decisions must be made. For example, some chip designs are moving to multiple core architectures within in a single chip, with each core running multiple processing threads. As with all multi-processor architectures, communications and latency are key determinants of overall processing throughput. These cores behave as programmable I/O devices or attached co-processors. Each of these co-processors can be optimized for complex functions such as signal processing, video and audio, etc. The next generation of computer processing architectures are now appearing in both technical and home computing systems. The challenge is that programming models must undergo fundamental changes in order to exploit these new multi-core architectures.

The new generation of chips, such as the IBM/Toshiba/Sony Cell BE, relies on task-specific processing elements to achieve the necessary throughput. Moreover, future versions will incorporate an increasing number task-specific cores, capable of running tens of processing threads concurrently. These tasks today are traditionally found in Field Programmable Gate Arrays (FPGA) and Graphics Processing Units (GPU) that can be customized by the user to perform specific tasks, but are hard to program.

In sense, this architecture does not look much different than architectures from 20 years ago where a single computer had a separate, dedicated array processing board. In fact, many of the issues are the same, including understanding I/O from the specialized hardware, integrating optimized software libraries for the special purpose hardware within the processing flow, etc. The challenge now is that there is greater complexity and flexibility available for user applications. The advantage is that these capabilities are incorporated into a single “system on a chip,” expanding the deployment and packaging opportunities far beyond what can be done with typical systems today. With the eventual appearance of low-power versions of these architectures, the ability to deploy enormous amounts of computational power into a wide variety of platforms could transform how we sense and respond to the environment.

RELATED PROJECTS

None.

REFERENCES

Van Veen, B.D., and K.M. Buckley, 1988, Beamforming: A Versatile Approach to Spatial Filtering, IEEE ASSP Magazine, vol. 5, pp. 4-24