IA^3 – SC'11 Workshop on Irregular Applications: Architectures & Algorithms

Grand Hyatt Favorita
Seattle, WA
November 13, 2011

9:00 – 9:10 Welcome & introduction
John Feo (Pacific Northwest National Laboratory)

9:10 – 9:50 Keynote 1
"Scalable Algorithms for Massive-scale Graph Analytics"
Fabrizio Petrini (IBM TJ Watson)

Abstract: Graph algorithms are increasingly important for biology, transportation, business intelligence and a wide range of commercial workloads, but they are ill-suited for machine architectures that are designed for 3D physics simulations. Most graph algorithms stress to the limit various architectural aspects of a conventional machine. The memory access pattern is irregular, with little spatial locality and data re-use; the amount of computation per loaded byte is very low, typically involving bit manipulation; pointer-chasing is often the norm. Similarly, the generated network traffic is composed of small packets that are sent to random destinations with a very high messaging rate.

With our recent winning Graph500 submissions in November 2010 and June 2011 (http://www.graph500.org), we have achieved a processing rate of 18.5 billion edges per second on a scale 38 problem (2^38 vertices and 2^42 edges, requiring several tens of terabytes of memory) on 131,072 BlueGeneP processing cores, demonstrating that it is possible to scale irregular graph problems to large node configurations.

In this talk we will describe our recent work, which includes two main tracks: (1) the development of an arsenal of run-time mechanisms to support graph algorithms, such as the implementation of lightweight messaging protocols (based on puts and gets) on the native HW, and mechanisms for efficient thread load-balancing and (2) the design of novel algorithmic optimizations to efficiently map graph algorithms on large-scale machines.

9:50 – 10:30 Keynote 2
"Easy and Efficient Graph Analysis with A Domain-Specific Language"
Kunle Olukotun (Stanford University)

Abstract: The increasing importance of graph-data based applications is fueling the need for highly efficient and parallel implementations of graph analysis algorithms. However, these implementations are difficult to get right and are usually tied to a particular hardware architecture. In this talk I will introduce Green-Marl, a domain-specific language (DSL) for graph analysis. Green-Marl allows application developers to use high level graph abstractions in a way that implicitly exposes the data-level parallelism inherent in their algorithms. The Green-Marl compiler can use graph-specfic knowledge to optimize these high-level descriptions and produce efficient code for a variety of parallel architectures. These architectures might include, multicore, GPU, cluster and

specialized graph architectures.

10:30 – 11:00 Coffee Break

11:00 – 12:15 Paper Session 1: Algorithms for Irregular Applications
Chair: *Simone Secchi (Pacific Northwest National Laboratory)*

1. "An OpenMP Algorithm and Implementation for Clustering Biological Graphs"
   Timothy Chapman (UC Santa Cruz) and Ananth Kalyanaraman (Washington State University)

   ***Abstract**:* Graph algorithms on parallel architectures present an interesting case study for irregular applications. Among the graph algorithms popular in scientific computing, graph clustering or community detection has numerous applications in computational biology. However, this operation also poses serious computational challenges because of irregular memory access patterns, large memory requirements, and their dependence on other auxiliary (also irregular) data structures to supplement processing. In this paper, we address the problem of graph clustering on shared memory machines. We present a new OpenMP-based parallel algorithm called pClust-sm, which uses adjacency lists, hash tables and union-find data structures in parallel. The algorithm improves both the asymptotic runtime and memory complexities of a previous serial implementation. Preliminary results show that this algorithm can scale up to 8 threads (cores) of a shared memory machine on a real world metagenomics input graph with 1.2M vertices and 100M edges. More importantly, the new implementation drastically reduces the time to solution from the order of several hours to just over 4 minutes, and in addition, it enhances the problem size reach by at least one order of magnitude.

2. "Parallel Bipartite Maximum Matching: Worth the Effort?"
   Ariful Azad (Purdue University), Mahantesh Halappanavar (Pacific Northwest National Laboratory), Florin Dobrian (Conviva Inc.), and Alex Pothen (Purdue University)

   Abstract: Bipartite maximum matching, also known as the linear assignment problem, is a fundamental combinatorial problem with numerous applications in science and engineering. In particular, its application to compute the block triangular form of a sparse matrix is our motivation. Extensive experiments of key matching algorithms in serial have shown an inherent weakness of these algorithm: their sensitivity to the order of processing. Since a fair amount of randomness is involved in the execution of any parallel algorithm, the issue of sensitivity raises a serious doubt: are parallel matching algorithms really worth the effort? This question gains relevance due the serial nature combined with irregular memory access patterns inherent to the matching algorithms. In this paper, we answer this question in the affirmative. Our answer is based on efficient parallel implementations of two key algorithms (Hopcroft-Karp and Pothen-Fan) and their variants, and extensive experiments on three multithreaded platforms (a 128-processor Cray XMT, a 48-core AMD Opteron, and a 32-core Intel Nehalem) using a carefully chosen set of input drawn from real-world

applications as well as generated synthetically.

3. "Parallel Simulation of Dendritic Growth On Unstructured Grids"
   Andreas Schäfer (Friedrich-Alexander-Universität Erlangen-Nürnberg), Julian Hammer (Friedrich-Alexander-Universität Erlangen-Nürnberg), Dietmar Fey (Friedrich-Alexander-Universität Erlangen-Nürnberg).

   Abstract: In this paper we present our findings from parallelizing a material science application which simulates dendritic growth in molten metal alloys. The simulation itself is based on an iterative meshfree model. This means that the underlying graph is not a regular grid, as in the case of stencil codes, but a graph whose nodes are scattered randomly across the 2D plain. Because the cells depend on neighbors in a relatively large radius, the model's calculations are tightly coupled and the code turned out to be communication bound. We present two different approaches for the parallelization: one specifically written for this application, and one which uses LibGeoDecomp, a stencil code library. Benchmarks show that the stencil code library performs much better than expected, despite not being designed for this use case.

12:15 – 13:45 Lunch Break

13:45 – 15:00 Paper Session 2: Architectures for Irregular Applications
Chair: *Antonino Tumeo (Pacific Northwest National Laboratory)*

1. "Implementation of a hierarchical N-Body Simulator using the OmpSs Programming Model"
   Miquel Pericàs (Barcelona Supercomputing Center), Yoav Etsion (Barcelona Supercomputing Center), Xavier Martorell (Barcelona Supercomputing Center)

   Abstract: Many HPC algorithms are highly irregular. They have input-dependent control flow and operate on pointer-based data structures such as trees, graphs, or linked lists. This irregularity makes it challenging to parallelize such algorithms in order to efficiently run them on modern HPC systems.
   In this paper we study the architectural and programming bottlenecks of the OmpSs task-based programming model when implementing irregular applications. We select a sequential N-body simulation code and describe its parallelization using OmpSs. We then analyze the code, focusing on scalability and load balancing. We conclude that, in general, task-based programming models are well suited to the exploitation of irregular parallelism. Nevertheless, in order to avoid the overheads associated with manually managing the load balancing, the hardware and runtime will need to collectively support much finer-grained tasks.

2. "Exploring Irregular Memory Accesses on FPGAs."
   Robert Halstead (UC Riverside), Jason Villarreal (Jacquard Computing) and Walid Najjar (UC Riverside).

   Abstract: Algorithms that exhibit irregular memory access patterns are known to show poor performance on multiprocessor architectures, particularly when

memory access latency is variable. Many common data structures, including graphs, trees, and linked-lists, exhibit these irregular memory access patterns. While FPGA-based code accelerators have been successful on applications with regular memory access patterns, they have not been further explored for irregular memory access patterns. Multithreading has been shown to be an effective technique in masking long latencies. We describe the compiler generation of concurrent hardware threads for FPGAs with the objective of masking the memory latency caused by irregular memory access patterns. We extend the ROCCC compiler to generate customized state information for each dynamically generated thread.

3. "A Memory Accelerator with Gather Functions for Bandwidth-bound Irregular Applications"
Noboru Tanabe (Toshiba), Boonyasitpichai Nuttapon (Tokyo University of Agriculture and Technology), Hironori Nakajo (Tokyo University of Agriculture and Technology), Yuka Ogawa (Nara Women's University), Junko Kogou (Nara Women's University), Masami Takata (Nara Women's University) and Kazuki Joe (Nara Women's University).

Abstract: Compute intensive processing can be easily accelerated using processors with many cores such as GPUs. However, memory bandwidth limitation becomes serious year by year for memory bandwidth intensive applications such as sparse matrix vector multiplications (SpMV). SpMV is an important kernel for iterative solver and PageRank etc. In the case of irregularly structured matrices, the memory bandwidth requirement for SpMV increases by order of magnitude on conventional processors. In order to accelerate memory bandwidth intensive irregular applications, we have proposed a memory system with additional functions of scattering and gathering. For the preliminary evaluation of our proposed system, we assumed that the throughput of the memory system was sufficient. In this paper, we propose a memory system with scattering and gathering using many narrow memory channels. We evaluate the feasible throughput of the proposed memory system based on DDR3 DRAM with the modified DRAMsim2 simulator. In addition, we evaluate the performance of SpMV using our method for the proposed memory system and a GPU. We have confirmed the proposed memory system has good performance and good stability for matrix shape variation using fewer pins for external memory.

15:00 – 15:30 Coffee Break

15:30 – 17:00 Panel Session
Moderator: *Oreste Villa (Pacific Northwest National Laboratory)*

Panelists: *David A. Bader (Georgia Tech), Walid Najjar (UC Riverside), Alex Pothen (Purdue University), Michael Garland (NVIDIA), Udeepta Bordoloi (AMD), Fabrizio Petrini (IBM TJ Watson), Kunle Olukotun (Stanford University)*