SAND2014-XXXXR

LDRD PROJECT NUMBER: 158787

LDRD PROJECT TITLE: Scheduling Irregular Algorithms

PROJECT TEAM MEMBERS:
Erik Boman, 01426
Donald Nguyen, U. of Texas at Austin

ABSTRACT:

This LDRD project was a campus exec fellowship to fund (in part) Donald Nguyen's PhD research at UT-Austin. His work has focused on parallel programming models, and scheduling irregular algorithms on shared-memory systems using the Galois framework. Galois provides a simple but powerful way for users and applications to automatically obtain good parallel performance using certain supported data containers. The naïve user can write serial code, while advanced users can optimize performance by advanced features, such as specifying the scheduling policy. Galois was used to parallelize two sparse matrix reordering schemes: RCM and Sloan. Such reordering is important in high-performance computing to obtain better data locality and thus reduce run times.

INTRODUCTION:

Parallel computing is important since almost all computers now have multiple cores (processing units). High-performance computing uses massively parallel systems with hundreds of threads on a single node and up to tens of thousands of nodes. Future exascale systems may have billion-way parallelism. How to program and schedule tasks efficiently on such systems is a challenging problem. For data analytics, there are environments such as MapReduce/Hadoop, Giraph, GraphLab, PowerGraph, and many more. Typically, they provide convenience but poor performance. For scientific computing, performance is crucial so people code directly to low-level interfaces such as MPI, OpenMP, and CUDA, but programmer productivity suffers. Galois attempts to provide both ease-of-use and high performance.

DETAILED DESCRIPTION OF EXPERIMENT/METHOD:

Galois is a C++ library and runtime system. It provides a set of concurrent data structures, such as unordered sets and graphs. The idea is the programmer uses these data structures and lets Galois find concurrency and schedule the execution for maximum performance. Some work was spent improving the Galois system itself. The example of using Galois to do sparse matrix reordering is of independent interest. The goal in this application is to improve the performance of sparse matrix-vector (SpMV) by reordering the matrix. The RCM algorithm is known to reduce the bandwidth of the matrix, and thus improve locality. Typically, RCM reordering is performed in serial but the SpMV in parallel. The one-time cost of reordering is often greater than the benefit during SpMV; therefore, the reordering itself should be executed in parallel. RCM is based on breadth-first-search (BFS) , so it is possible to parallelize (though typically not done in HPC codes). This work showed that Galois could satisfactorily parallelize RCM so it is worth doing the reordering. See the paper "Parallelization of Reordering Algorithms for Bandwidth and Wavefront Reduction" (SC14) for details.

RESULTS:

The major results of this project are:

1. Donald Nguyen will graduate in December 2014.
2. The Galois system can potentially be useful in both data analysis (machine learning) and for sparse matrix computations in scientific computing.

Several papers were in part supported by this LDRD project:

1. *Konstantinos I. Karantasis, Andrew Lenharth, Donald Nguyen, Maria Garzaran, Keshav Pingali, "Parallelization of Reordering Algorithms for Bandwidth and Wavefront Reduction", in Proc. of SuperComputing 2014 (SC14), 2014.*
2. *Donald Nguyen, Andrew Lenharth, and Keshav Pingali, "Deterministic Galois: On-demand, portable and parameterless.", In Proceedings of International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS '14, 2014.*
3. *Donald Nguyen, Andrew Lenharth, and Keshav Pingali: A lightweight infrastructure for graph analytics. Proc. Of SOSP, pp. 456-471.*

DISCUSSION:

The Galois system is likely to have most impact on computing in data mining and machine learning. It is hard to have impact on high-performance computing, as developers are highly skilled and willing to hand-optimize codes for maximum performance. The RCM example is not convincing as there are several well-known ways to parallelize BFS. Also, there are other orderings (such as nested dissection methods) that have recently proven useful for SpMV. The numerical effects of reordering the matrix (e.g., impact on preconditioners) were not considered, but is a significant issue in scientific computing.

ANTICIPATED IMPACT:

There is no anticipated impact or planned follow-on work. There might be potential future impact since some Sandia staff are now aware of Galois through Nguyen's visits to the lab.

CONCLUSION:

This project has been fairly successful, with a paper accepted to SC14 but unfortunately, no impact on Sandia. The Galois system seems viable for data analytics but will struggle to impact the HPC community. First, there are other ways to parallelize breadth-first-search (BFS). Second , it is unclear how much gain parallel RCM gives us. Typically, matrix reordering is just a tiny fraction of the total work performed in a scientific simulation. Therefore, there has been little focus on parallel orderings.

LDRD
LABORATORY DIRECTED RESEARCH & DEVELOPMENT

Sandia National Laboratories

U.S. DEPARTMENT OF ENERGY