SAND2012-6752P

** Division 1000 News Notes Submission Form **

Please complete all sections and obtain approval following Center review and approval process. After receiving approval of content and concurrence with sensitivity marking, please email the final version to Tanci Wilson at: tlwilso@sandia.gov.

Technical POC:

Name: Kyle Wheeler **Org:** 01423

Phone #: 844-0394 **Email:** kbwheel@sandia.gov

Date (m/d/yy): 8/13/2012

[If Appropriate, please include website references (URLs), photos, event dates, and media references.]

Topic Area: Technical Accomplishments

If "Other" is selected, please describe:

Program Funding Source(s) [eg. LDRD, DARPA, NW Program]: DARPA

Manager Approval: Ron Brightwell **Date:** 8/13/2012

Title: **Flexible Preconditions – A Model for Efficient Macro-Dataflow Execution**

Summary (provide paragraph): We have developed a new task creation strategy for macro-dataflow models, called flexible-preconditions. Our approach addresses task dependences in a way that combines the strengths of the two current approaches: the expressiveness and convenience of eager task creation and the performance of dependence-delayed task spawns, when a set of preconditions is known a priori. This work focuses on the performance of parallel programming models based on macro-dataflow, in which applications are composed of tasks and inter-task dependences. Data-flow models usually choose between specifying the task dependences before task creation, as preconditions, or during task execution, when they are actually needed (eager execution). We have created a programming and runtime environment that is sufficiently flexible to support both eager and delayed execution by means of preconditions. We have implemented the flexible-preconditions approach for an existing dataflow model and studied the implications that the choice between eager and lazy dependences has on the memory footprint, performance, and expressiveness of such models. Our results show that the use of preconditions when all dependences are known ahead of time can lead to a performance improvement of up to 38% (Figure 1 below) and a reduced memory footprint of up to 62% (Figure 2 below). In addition, on an application where the complete set of per-task dependences is determined after the task is spawned, we obtain 25% increase in performance by using flexible-preconditions.

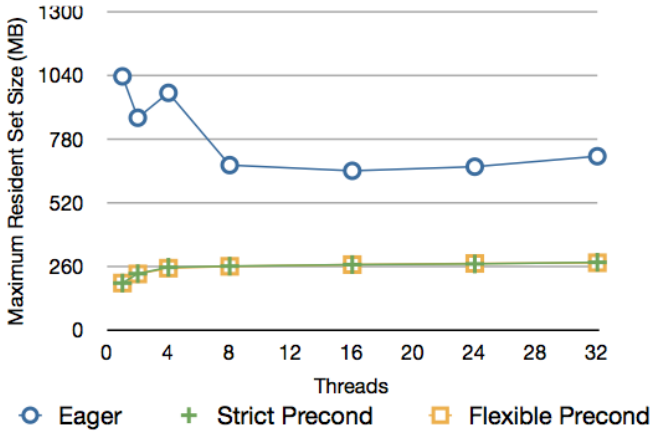

Figure 1: Performance results of a Cholesky factorization that show a 38% difference in performance between the eager approach and the strict and flexible preconditions approaches.

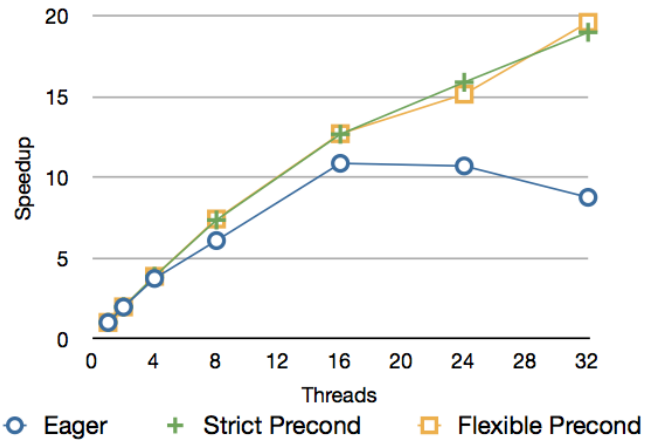

Figure 2: Memory size of Cholesky factorization that shows a reduction of 62% for flexible preconditions compared to eager execution.

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.