Electronic Transactions on Numerical Analysis.
Volume 21, pp. 66-80, 2005.
Copyright © 2005, Kent State University.
ISSN 1068-9613.

ETNA
Kent State University
etna@mcs.kent.edu

ADDRESSING THE STOCHASTIC NATURE OF SCIENTIFIC COMPUTATIONS VIA DYNAMIC LOOP SCHEDULING[*]

IOANA BANICESCU[†] AND RICOLINDO L. CARIÑO[‡]

Dedicated to Alan George on the occasion of his 60th birthday

Abstract. In general, scientific applications are large, computationally intensive, irregular and data-parallel. One of the main performance degradation factors when running scientific applications in parallel and distributed environments is load imbalance that may rise from a wide range of characteristics, due to problem, algorithmic, and systemic factors. Moreover, due to the stochastic behavior of many scientific applications, the load imbalance can unexpectedly occur at runtime, leading to difficulties in making decision about optimal partitioning, allocation and scheduling. In recent years, since parallel loops are prevalent and a major source of parallelism in scientific computations, many dynamic loop scheduling algorithms have been developed to improve performance via load balancing. These algorithms address the unpredictable behavior of computations, and many of these are based on probabilistic analyses. These have proven to be effective in improving performance of important applications used in scientific computing by addressing the entire spectrum of factors that lead to load imbalance. This paper proposes the use of a dynamic loop scheduling approach for repartitioning in scientific applications that contain computationally intensive parallel loops. A new partition is computed and data is redistributed during the execution of a parallel loop. This approach addresses load imbalance factors as they arise during a computational step, in contrast to the traditional load balancing approach of determining a new partition and migrating data after a computational step. Preliminary performance tests of an application that implements the proposed approach validate the feasibility and the effectiveness of the approach.

Key words. dynamic load balancing, data repartitioning, loop scheduling

AMS subject classifications. 68W10

[*] Received August 19, 2004. Accepted for publication October 3, 2005. Recommended by S. Oliveira. This work was partially supported by the following National Science Foundation grants: CAREER #9984465, ITR/ACS #0081303, ITR/ACS #0085969, #0132618, and #0082979.

[†] Dept. of Computer Science and Engineering, Mississippi State University, PO Box 9637, Mississippi State MS 39762 (`ioana@cse.msstate.edu`). Also with the Center for Computational Sciences–ERC, Mississippi State University.

[‡] Center for Computational Sciences–ERC, Mississippi State University, PO Box 9627, Mississippi State MS 39762 (`rlc@erc.msstate.edu`).