ASYNCHRONOUS GLOBAL OPTIMIZATION FOR MASSIVE-SCALE COMPUTING

By

Travis Desell

An Abstract of a Thesis Submitted to the Graduate

Faculty of Rensselaer Polytechnic Institute

in Partial Fulfillment of the

Requirements for the Degree of

DOCTOR OF PHILOSOPHY

Major Subject: Computer Science

The original of the complete thesis is on file
in the Rensselaer Polytechnic Institute Library

Examining Committee:

Carlos Varela, Thesis Adviser
Boleslaw Szymanski, Thesis Adviser
Malik Magdon-Ismail, Member
Heidi Newberg, Member
Dave Anderson, Member

Rensselaer Polytechnic Institute
Troy, New York

December 2009
(For Graduation December 2009)

ABSTRACT

As the rates of data acquisition and cost of model evaluation in scientific computing are far surpassing improvements in processor speed, the size of the computing environments required to effectively perform scientific research is increasing dramatically. As these computing environments increase in size, traditional global optimization methods, which are sequential in nature, fail to adequately address the challenges of scalability, fault tolerance and heterogeneity that using these computing systems entails. This thesis introduces asynchronous optimization strategies which while similar to their traditional synchronous counterparts, do not have explicit iterations or dependencies. This allows them to scale to hundreds of thousands of hosts while not being degraded by faults or heterogeneity. A framework for generic distributed optimization (FGDO) is presented, which separates the concerns of scientific model development, distributed computing and developing efficient optimization strategies; allowing researchers to develop these independently and utilize them interoperably through simple interfaces. FGDO has been used to run these asynchronous optimization methods using an astroinformatics problem which calculates models of the Milky Way galaxy on thousands of processors in RPI's BlueGene/L supercomputer and to run the MilkyWay@Home volunteer computing project, which currently consists of over 25,000 active computing hosts. A simulation environment was also implemented in FGDO, which allowed asynchronous optimization to be examined in a controlled setting with benchmark optimization problems. Results using the simulated environment show that the asynchronous optimization methods used scale to hundreds of thousands of computing hosts, while the traditional methods do not improve or even degrade as more computing hosts are added. Additionally, the asynchronous optimization methods are shown to be largely unaffected by increasing heterogeneity in the computing environment and also scale similarly in a computing environment modeled after MilkyWay@Home. This thesis presents strong evidence of the need for novel optimization methods for massive scale computing systems and provides effective initial work towards this goal.