Parallel Algorithms in STAPL:
Sorting and the Selection Problem

Anna Tikhonova

Gabriel Tanase   Olga Tkachyshyn

Nancy M. Amato   Lawrence Rauchwerger

Parasol Lab, Dept. of Computer Science, Texas A&M University

{atikhono,gabrielt,olgat,amato,rwerger}@cs.tamu.edu

July 19, 2005

Abstract

Parallel and distributed processing provide massive computational power demanded by modern computing. The development of properly designed parallel libraries is crucial for the advancement of parallel processing and moving parallel computing into the mainstream. The Standard Template Adaptive Parallel Library (STAPL) is a parallel library designed as a superset of the (sequential) ANSI C++ Standard Template Library (STL). STAPL provides routines that are easily interchangeable with their sequential counterparts and allows users to execute programs on uni- or multiprocessor systems that utilize shared or distributed memory, insulating less experienced users from managing parallelism and, at the same time, allowing more sophisticated users sufficient control to achieve higher performance gains.

Our research is focused on designing and implementing a set of parallel sorting and selection algorithms in STAPL. This paper will use the Selection Problem (Nth Element Algorithm) as an example of a parallel algorithm and discuss the way it is implemented in the STAPL environment. Our goal is to design the most appropriate algorithms for STAPL parallel library that will provide good efficiency without sacrificing generality and portability.