How to Share Memory in a Distributed System

Eli Upfal
Avi Wigderson

Abstract

The power of shared-memory in models of parallel computation is studied, and a novel distributed data structure that eliminates the need for shared memory without significantly increasing the run time of the parallel computation is described. More specifically, it is shown how a complete network of processors can deterministically simulate one PRAM step in $O(log n(log log n)^2)$ time when both models use n processors and the size of the PRAM's shared memory is polynomial in n. (The best previously known upper bound was the trivial $O(n)$). It is established that this upper bound is nearly optimal, and it is proved that an on-line simulation of $T PRAM$ steps by a complete network of processors requires $\Omega(T(log n/log log n))$ time.

A simple consequence of the upper bound is that an Ultracomputer (the currently feasible general purpose parallel machine) can simulate one step of a PRAM (the most convenient parallel model to program) in $O((log n)^2 log log n)$ steps.