Trade-offs Between Depth and Width in Parallel Computation

Uri Vishkin
Avi Wigderson

Abstract

A new technique for proving lower bounds for parallel computation is introduced. This technique enables us to obtain for the first time, non-trivial tight lower bounds for shared-memory models of parallel computation that allow simultaneous read/write access to the same memory location. The size m of the common memory is called communication width or width in short.

For a wide variety of problems (including parity and majority) we show that the time complexity T (depth) and the communication width m are related by the trade-off curve $mT^2=\Omega(n)$, (where n is the size of the input). This bound is tight for every $m \leq n/log^2n$. We extend our technique to prove $mT^3=\Omega(n)$ trade-off for a class of “simpler” functions (including Boolean Or) on a weaker model that forbids simultaneous write access. This result improves the lower bound of Cook and Dwork [CD-82] when communication is limited.