

Trade-offs Between Depth and Width in Parallel Computation

Uzi 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, nontrivial tight lower bounds for shared-memory models of parallel computation that allow several processors to have simultaneous access to the same memory location. Specifically, we use a concurrent-read concurrent-write model of parallel computation. It has p processors, each has access to a common memory of size m (also called *communication width* or *width* in short). The input to the problem is located in an additional read-only portion of the common memory.

For a wide variety of problems (including parity, majority and summation) 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), *regardless* of the number of processors. Moreover, for every point on this curve with $m = O(n/\log^2 n)$ we give a matching upper bound with the *optimal* number of processors.

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. We also state and give a proof of a new result by Beame [B-83] that achieves a tight lower bound for the OR in this model, namely $mT^2 = \Omega(n)$. These results improve the lower bound of Cook and Dwork [CD-82] when communication is limited.