Beyond Parallel Random-Access Machines

Bruce Maggs

The fundamental paradigm in theoretical computer science is to propose a formal model of computation, and then to analyze the complexity of different problems in the model. However, if the model fails to reflect constraints in the real world, these analyses have limited practical importance. The parallel random-access machine (PRAM) is by far the most popular model of parallel computation. Hundreds of algorithms have been designed for the PRAM. Yet, the PRAM ignores contraints on communication imposed by modern electronic packaging technology.

Recently, several networks have been shown to be *universal* in the sense that they can emulate the PRAM with at most polylogarithmic degradation in time [1, 2, 6, 8]. For example, with high probability, an N-processor butterfly network can emulate each step of an N-processor concurrent-read concurrent-write (CRCW) PRAM in $O(\log N)$ steps [6]. It is tempting to conclude that these results validate the PRAM as a practical model of parallel computation.

A problem with universal networks is that they may be difficult to build on a large scale. When a network is physically constRucted, Each packaging unit (chip, board, rack, cabinet) contains a subset of the processors of the network, and thus determines a cut of the network. For a universal network, the capacity (bandwidth) of every cut must be nearly as large as the number of processors on the smaller side of the cut; otherwise, it would be impossible to perform certain memory accesses in polylogarithmic time. Thus, when a universal network is built, the number of external connections (pins) on a packaging unit must be nearly as large as the number of processors in the unit. However, the number of pins on a packaging unit of an electronic system is typically much smaller than the number of components that the unit contains, and can be made larger only with great cost. Consequently, if all the pin constraints are met, a packaging unit cannot contain as many processors as might otherwise fit. Alternatively, if each packaging unit contains its full complement of processors, then pin limitations preclude the universal network from being assembled.

The impact of pin constraints can be modeled theoretically in the three-dimensional VLSI model [3, 4] where hardware cost is measured by volume and the pinboundedness of a region is measured by its surface area. In this model, the largest universal network that can fit in a given volume V has only about $V^{2/3}$ nodes. In the two-dimensional model [7], where pinboundedness is measured by perimeter, the bound is even worse. The largest universal network that can fit in area A has only about $\sqrt{A}$ nodes.

The difficulty of building universal networks suggests the exploration of alternative models of parallel computation. Leiserson and Maggs have proposed a model called the *distributed random-access machine* (DRAM) in which pin constraints are modeled explicitly and the communication requirements of an algorithm can be measured. However, the DRAM hasn't caught on with PRAM algorithm designers. Perhaps they believe that some fundamental change in electronic technology will eventually make the PRAM more practical. Most likely, they prefer the simplicity of the PRAM. Nevertheless, if we are to design practical algorithms for the machines that we can build today, we need models that reflect real world constraints.

References

[1] A. R. Karlin and E. Upfal, "Parallel hashing — an efficient implementation of shared memory," *Proceedings of the 18th Annual ACM Symposium on the Theory of Computing*, May 1986, pp. 160–168.

[2] F. T. Leighton, B. Maggs, and S. Rao, "Universal packet routing algorithms", *Proceedings of the 29th Annual Symposium on the Foundations of Computer Science*, IEEE, October 1988, pp. 256–271.

[3] F. T. Leighton and A. L. Rosenberg, "Three-dimensional circuit layouts," *SIAM J. Comput.*, Vol. 15, No. 3, August 1986.

[4] C. E. Leiserson, "Fat-trees: universal networks for hardware-efficient supercomputing," *IEEE Transactions on Computers*, Vol. C–34, No. 10, October 1985, pp. 892–901.

[5] C. E. Leiserson and B. M. Maggs, "Communication-efficient parallel graph algorithms for distributed random-access machines," *Algorithmica*, Vol. 3, pp. 53-77, 1988.

[6] A. G. Ranade, "How to emulate shared memory," *Proceedings of the 28th Annual Symposium on Foundations of Computer Science*, IEEE, October 1987, pp. 185–194.

[7] C. D. Thompson, *A Complexity Theory for VLSI*, Ph.D. thesis, Department of Computer Science, Carnegie-Mellon University, Pittsburgh, PA, 1980.

[8] L. G. Valiant and G. J. Brebner, "Universal schemes for parallel communication," *Proceedings of the 13th Annual ACM Symposium on the Theory of Computing*, May 1981, pp. 263–277.