Time-Optimal Simulations of Networks by Universal Parallel Computers

Friedhelm Meyer auf der Heide*
Rolf Wanka
Informatik II, Universität Dortmund
D - 4600 Dortmund 50, Fed. Rep. of Germany

ABSTRACT: For technological reasons, in a realistic parallel computer the processors have to communicate via a communication network with bounded degree. Thus the question for a "good" communication network comes up. In this paper we present such a network, a universal parallel computer (UPC) with the following properties:

(i) It has optimal time-loss, namely $O(\log c)$ for simulating networks of degree c. (We also prove the lower bound $\Omega(\log c)$ for the time-loss.)

(ii) We introduce the broadcast-capability (how many processors can be reached by one processor in i steps?) and demonstrate its influence on the number of processors needed for a simulation of a network with n processors. E.g. for broadcast-capability $O(c^i)$ (e.g. networks with degree c), $O(n^{1+\varepsilon} \log n)$ processors are needed ($\varepsilon > 0$ arbitrary) whereas $O(n \cdot \text{polylog}(n))$ processors suffice for networks with polynomial broadcast-capability (e.g. k-dimensional grids).

(iii) The UPC is potentially infinite and has multi-user capabilities, i.e. it can be arbitrarily partitioned into finite UPC's each with the above efficiency.

This construction generalizes a UPC described in [MadH2], where, given a fixed degree c, for each n a UPC M_0 is constructed which needs $O(n^{1+\varepsilon} \log n)$ processors to achieve constant time-loss for simulating networks with n processors and degree c.

1. Introduction

I. The problem

Recently in complexity theory the following basic concept for parallel computation has been developed: p processors cooperate in the execution of parallel algorithms by interchanging data via a common mechanism for communication.

The most powerful instance of this concept is the model of the parallel random access machine (PRAM) where the processors have read and write access to a shared memory consisting of an infinite number of storage locations. This model is widely used to study various problems of parallel computation. Unfortunately it is not realistic because it ignores the current technological constraints of realizing communication among the processors.

In this paper we study the more realistic instance of the basic concept where networks of processors are investigated. In this model pairs of processors are interconnected by wires along which they can communicate. The interconnection pattern which we call the *communication graph* must be fixed.

* Supported in part by DFG-grants Me 872/1-2 and We 1066/2-1

A network is unrealistic if the number of wires going out of a single processor becomes too large. Therefore we demand that the degree of a processor is bounded by a small integer which is independent of the size, i.e. the number of processors in the network. A network with this property is called *realistic.* All networks we propose in this paper have degree less than 8.

In literature, many different communication networks are proposed for special purposes, e.g. the *Hypercube* [Pease], the *Butterfly network* [Waksman], or the *Shuffle-Exchange network* [Stone] for sorting, permuting and fast Fourier transform, Leighton's modification [Leighton] of the *AKS*-network [AKS] for sorting , multi-dimensional grids for interpolation problems, matrix operations, or graph problems (compare e.g. [Hambrusch],[Atallah,Kosaraju]).

On the other hand it is not conceivable to build a new network for each new application. Thus we have to look for versatile, i.e. for *universal parallel computers (UPCs).* A UPC is a realistic network that can simulate all other networks from a given class U using "few" resources (time, number of processors). Let us make this concept more precise.

Let U be a class of networks. A realistic network M_0 is *universal* for U (is a UPC for U), if it can simulate each network from U. The time-loss of a simulation is the factor by which the simulation takes longer than the simulated machine would take. As usual, we do not take into account the time needed for preprocessing M_0 for the network to be simulated.

II. Known results about universal parallel computers

There are two well-known propositions for UPC's: the *Shuffle-Exchange Network* due to Stone [Stone], and the *Cube-Connected Cycles Network* due to Preparata and Vuillemin [Prep,Vuil].

Both networks are universal for $\mathcal{U}(n)$, the class of all networks with n processors. For networks from the subclass $\mathcal{U}(n,c)$ of $\mathcal{U}(n)$ of networks of degree c, the above UPCs have time-loss $O(\min\{c \cdot \log n, \mathrm{sort}(n)\})$, where $\mathrm{sort}(n)$ denotes the time necessary to sort n items on the respective network. $\mathrm{sort}(n)$ can be chosen as $O((\log n)^2)$ using Batcher's Sort [Batcher] or $O(\log n)$ using the probabilistic sort due to Reif and Valiant [Reif,Valiant].

Galil and Paul have in [Galil,Paul] extended this approach to potentially infinite UPC's. They also take the preprocessing mentioned above into account to define a model for purposes of uniform complexity theory.

In [MadH3] a method is presented to equip the above UPC from [Paul,Galil] with "multi-user properties", i.e. to modify it in such a way that, after removing the first m processors, the remaining network is isomorphic to M_0, and thus can be used in the same way by other users for further simulations.

In [MadH2] it is shown that a UPC for $\mathcal{U}(n,c)$, c constant, exists with constant time-loss. It has $O(n^{1+\varepsilon} \log n)$ processors, for arbitrary $\varepsilon > 0$. It is not known whether this UPC is best possible for those with constant time-loss. The best lower bound from [MadH1] shows the trade-off

$$(\text{time-loss}) \star (\text{size}) = \Omega(n \log n / \log\log n)$$