costs $N_A = 1296 + 1024 + 256 = 2576$.

The nodes for the sorting network only need buffers of length 1 and only 1 multiplexer. They have basic costs $S_A = 1928$ and $S_S = 4104$.

The network nodes for the Fluent Machine have width 76 for the forward and the backward network. Thus they have basic costs $\tilde{N}_A = 3104$ and $\tilde{N}_S = 8808$.

If we reduce the network nodes of design $D1$ to design an EREW PRAM, we spare the costs for the comparator and reduce the width of the instruction queue by one bit. We then have basic costs $N'_A = 2320$ and $N'_S = 6192$.

A.3 Network Speed

In one network cycle the maximum delay path is the following: a packet has to be read out of the input buffer, its address has to be compared with another, it has to be selected by a multiplexer and it has to be stored in the input buffer of the following node. Reading the input buffer takes 3 gate delays, comparing addresses with an i bit carry lookahead adder takes about $2\log i + 2$ gate delays, selecting with a multiplexer takes 2 gate delays, storing in a buffer takes about 3 gate delays. With 32 bit addresses we have $3 + 2\log 32 + 2 + 2 + 3 = 20$ gate delays. Because we did not count driver delays and setup and hold times we take a network cycle time $\sigma_N = 25$ gate delays.

B Analysis of Benchmark B1

Let $G = (V, E)$ be an undirected graph with $v = |V|$, $V = \{0, \ldots, v - 1\}$ and $E \subseteq V \times V$ with $e = |E|$. Represent E by the adjacency matrix A given by $a_{jk} = 1$ if $(j, k) \in E$, 0 otherwise. A is symmetric because G is undirected. The connected components algorithm from [5] first computes the connectivity matrix C from the given adjacency matrix. C is given by $c_{jk} = 1$ if there exists a path in G from j to k, 0 otherwise. Then it constructs a matrix D given by $d_{jk} = k$ if $c_{jk} = 1$, 0 otherwise. Finally each vertex k is assigned to component l with $l = \min\{i | d_{ki} \neq 0\}$.

We assume that sending one word across a link of the hypercube takes only one step and that source and destination of this word are registers. The connectivity matrix is computed by $\log v$ times multiplying A with itself thus computing $C = A^v$. It turns out that multiplying 2 $v \times v$ matrices on a hypercube with n processors can be done in $(38\frac{v^3}{n} + 20)\log v + 5\frac{v^3}{n}$ steps. The computation of the connectivity matrix then needs approximately $(\log v)^2(38\frac{v^3}{n} + 20) + 5\frac{v^3}{n}\log v$

steps. The computation of matrix D takes approximately $7\frac{v^2}{n}$ steps, finding of minimums takes approximately $10\frac{v}{n}\log v$ steps. The total time t_{D6} then is approximately $40\frac{v^3}{n}(\log v)^2$.

[24] M. Paterson. Improved sorting networks with $O(\log N)$ depth. *Algorithmica*, 5:75–92, 1990.

[25] D. A. Patterson and C. H. Sequin. A VLSI RISC. *Comput.*, 15(9):8–21, 1982.

[26] A. G. Ranade. How to emulate shared memory. In *Proc. 28th IEEE FOCS*, pages 185–194, 1987.

[27] A. G. Ranade, S. N. Bhatt, and S. L. Johnson. The Fluent Abstract Machine. In *Proc. 5th MIT Conference on Advanced Research in VLSI*, pages 71–93, 1988.

[28] J. Reif and L. G. Valiant. A logarithmic time sort for linear size networks. *J. ACM*, 34(1):60–76, January 1987.

[29] Y. Shiloach and U. Vishkin. An $O(\log n)$ parallel connectivity algorithm. *J. Algorithms*, 3:57–67, 1982.

[30] B. Smith. A pipelined shared resource MIMD computer. In *Proc. 1978 Int. Conf. on Parallel Processing*, pages 6–8. IEEE, 1978.

[31] E. Upfal. Efficient schemes for parallel communication. In *Proc. ACM Symp. on Principles of Distr. Comp.*, pages 55–59, August 1982.

[32] L. G. Valiant. A scheme for fast parallel communication. *SIAM J. Comput.*, 11:350–361, 1982.

[33] L. G. Valiant. A bridging model for parallel computation. *Commun. ACM*, 33(8):103–111, 1990.

[34] L. G. Valiant. General purpose parallel architectures. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. A*, pages 943–971. Elsevier, 1990.

[35] I. Wegener. *The Complexity of Boolean Functions*. Teubner, 1987.

Appendices

A   Parameter Values

A.1   Processor Costs

The ALU mainly consists of a 32 bit wallace tree multiplier, a barrel shifter and a carry lookahead adder (see [35]). The multiplication is performed by adding 32 terms of length 1 to 32 bits. Each bit of each term is computed by an AND gate. The AND gates have basic costs $32 \times 16 \times 2 = 1024$.

4 terms of length i can be reduced to 2 terms of length $i + 1$ with an i bit 4–2–Adder consisting of $2i$ full adders. Thus we get a first stage consisting of a 32 bit 4–2–adder for the longest terms, up to a 4 bit 4–2–adder for the shortest terms. In total the first stage contains $32 + 28 + \ldots + 4 = 144$ bit 4–2–adders. The second stage contains $28 + 20 + 12 + 4 = 64$, the third stage $24 + 8 = 32$ and the last stage 16 bit 4–2–adders. The wallace tree then consists of $144 + 64 + 32 + 16 = 256$ bit 4–2–adders containing 512 full adders. As we saw in section 1.1 a full adder has basic costs 18. The basic costs for the wallace tree then are 9216.

The carry lookahead adder finishing the multiplication consists of 2×32 components to compute generate and propagate signals. Each component consists of 4 AND and OR gates. Additionaly we need for each bit 2 EXOR gates to compute the sum bits and 1 AND and 1 OR gate to produce the generate and propagate signal for that bit. The carry lookahead adder has basic costs $64 \times 4 \times 2 + 32 \times (2 \times 2 + 2 \times 6) = 1024$.

The barrel shifter consists of 5 stages of multiplexers. Because we allow rotations and buffering in carry each stage needs 33 multiplexers with 3 inputs. These are built of 2 multiplexers with 2 inputs. A multiplexer with two inputs consists of 2 AND gates, 1 OR gate and 1 inverter, having costs 7. The total basic costs of the barrel shifter now are $5 \times 33 \times 2 \times 7 = 2310$. The basic costs of the ALU then are $A = 1024 + 9216 + 1024 + 2310 = 13392$.

A register file with 16 registers 32 bit wide has basic costs $F = 16 \times 32 \times 12 = 6144$.

A.2   Costs of the network

Own simulations [3] show that network nodes only need buffers of length 2. Packets on the way from processors to memory modules are 76 bits wide (32 bit address, 32 bit data, 12 bit control). In the backward network 32 bits for transported data are sufficient. Each network node needs the following hardware: 4 registers with 76 bits each, 4 registers with 32 bits each, $2c \log n$ 3 bit registers with routing informations for the backward network, 2 multiplexers with 76 bits and 2 multiplexers with 32 bits. Additionaly we need a comparator and an adder to test identical addresses and to select the smaller one.

The registers have basic costs $N_S = (4 \times 76 + 4 \times 32 + 2 \times 3 \times 7 \times 3) \times 12 = 6696$. The multiplexers have basic costs $(2 \times 76 + 2 \times 32) \times 6 = 1296$. The adder has basic costs 1024 as computed above. The comparator consists of 1 EXOR gate and 1 OR gate for each of the 32 bits and thus has basic costs $32 \times (6 + 2) = 256$. The arithmetic of a network node then has total basic

cost–effectiveness of PRAM's and DMM's. The results are surprisingly favourable for PRAM's. In reality things are somewhat worse, e.g. because of connectors and wires. Nevertheless we are applying for funds to build a prototype of design $D1$ with $n = 128$ physical processors. Some details about the planned prototype can be found in [3].

Acknowledgements

References

[1] F. Abolhassan, A. Bingert, and J. Keller. Vergleich des ALLIANT FX/2816 mit einer realisierbaren PRAM. Manuscript, Universität des Saarlandes, 1991.

[2] F. Abolhassan and J. Keller. Detailed analysis of PRAM machines. Manuscript, July 1990.

[3] F. Abolhassan, J. Keller, and W. J. Paul. On physical realizations of the theoretical PRAM model. FB 14 Informatik, SFB-Report 21/1990, Universität des Saarlandes, December 1990.

[4] M. Ajtai, J. Komlós, and E. Szemerédi. An $O(n \log n)$ sorting network. In *Proc. 15th ACM STOC*, pages 1–9, New York, 1983. ACM.

[5] S. G. Akl. *The Design and Analysis of Parallel Algorithms*. Prentice–Hall, 1989.

[6] H. Alt, T. Hagerup, K. Mehlhorn, and F. P. Preparata. Deterministic simulation of idealized parallel computers on more realistic ones. *SIAM J. Comput.*, 16(5):808–835, October 1987.

[7] K. Batcher. Sorting networks and their applications. In *AFIPS Spring Joint Computer Conference, Vol. 32*, pages 307–314, Reston, Va., 1968. AFIPS Press.

[8] G. Bilardi and K. T. Herley. Deterministic simulations of PRAMs on bounded degree networks. In *Proc. 26th Annual Allerton Conference on Communication, Control and Computation*, September 1988.

[9] Y. Chang and J. Simon. Continuous routing and batch routing on the hypercube. In *Proc. 5th ACM Symp. on Principles of Distr. Comp.*, pages 272–281, 1986.

[10] R. Cole. Parallel merge sort. *SIAM J. Comput.*, 17(4):770–785, August 1988.

[11] M. Dietzfelbinger and F. Meyer auf der Heide. A new universal class of hash functions and dynamic hashing in real time. Reihe Informatik Bericht Nr. 67, Universität–GH Paderborn, April 1990.

[12] A. Formella. Effizienz numerischer Rechnerarchitekturen. Manuscript, Universität des Saarlandes, 1991.

[13] A. Gibbons and W. Rytter. *Efficient Parallel Algorithms*. Cambridge University Press, 1988.

[14] T. Hagerup. Optimal planar algorithms on planar graphs. *Inf. Comput.*, 84:71–96, 1990.

[15] T. Hagerup, A. Schmitt, and H. Seidl. FORK: A high–level–language for PRAMs. In *Proc. PARLE91*, 1991.

[16] R. Hockney and C. Jesshope. *Parallel Computers 2*. Adam Hilger, Bristol and Philadelphia, 1988.

[17] A. R. Karlin and E. Upfal. Parallel hashing: An efficient implementation of shared memory. *J. ACM*, 35(4):876–892, October 1988.

[18] R. M. Karp and V. L. Ramachandran. A survey of parallel algorithms for shared–memory machines. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. A*, pages 869–941. Elsevier, 1990.

[19] F. Leighton and C. E. Leiserson. Theory of parallel and VLSI computation. Lecture notes, research seminar series, Massachusetts Institute of Technology, May 1990.

[20] F. Leighton, B. Maggs, and S. Rao. Universal packet routing algorithms. In *Proc. 29th IEEE FOCS*, pages 256–269, 1988.

[21] K. Mehlhorn and U. Vishkin. Randomized and deterministic simulations of PRAMs by parallel machines with restricted granularity of parallel memories. *Acta Inf.*, 21:339–374, 1984.

[22] Motorola, Inc. ASIC Division, Chandler, Arizona. *Motorola High Density CMOS Array Design Manual*, July 1989.

[23] S. M. Müller and W. J. Paul. Towards a formal theory of computer architecture. In *Proceedings of PARCELLA 90, Advances in Parallel Computing*. North–Holland, 1990.

6.2 Simulation of PRAMs by DMMs

The worst case for a DMM is a benchmark where any known algorithm for a DMM is less cost–effective than the step–by–step simulation of a PRAM.

Theorem 4 *Let B be a benchmark that fulfils the above assumptions and that is parallelizable with efficiency ϵ. Then*

$$R \geq L \geq \frac{1}{438}.$$

Proof: Let the sequential runtime of B be T. B needs $T/\epsilon cn \log n$ steps on a PRAM $D1$ with $cn \log n$ processors. Because each step takes $c \log n$ processor cycles, $t_{D1} = 50T/\epsilon n$.

Let $D5$ be a DMM with $n \log n$ processors interconnected as a hypercube. $D5$ has costs $c_{D5} = c_{\tilde{P}} n \log n = 12944 n \log n$ because we ignore network costs. In order to simulate one step of $D1$ on $D5$ we adapt RANADE's routing scheme in software. Because successing phases can overlap we use a link in forward manner for phases 1,3,5 and in backward manner for phases 2,4,6. Processors alternately execute one step of phase i and one of phase $i + 1$. Because of this toggling the routing scheme needs at most twice as many routing steps as RANADE's scheme. The number of machine instructions to perform one routing step is 24:

# steps	comment
6	read address, data, mode of both inputs
1	compare the addresses
1	jump if equal (combining)
1	jump if less (left packet is to send)
1	compare address with routing mask
1	jump if equal (routing to left output)
1	test whether successing queue is full
1	jump if full
3	write address, data and mode
2	append direction queue if mode==read
1	mark input queue not full
1	test whether other successing queue full
1	jump if full
3	write address,data,ghost
$\sum 24$	Total

If we assume that RANADE's scheme needs $11 \log n$ steps the new scheme needs $11 \log n \times 24 \times 2 = 528 \log n$ instructions. If we further assume that one instruction only takes one processor cycle, the total time to simulate one PRAM step is at most $S \log n$ processor cycles

for $S = 528$. $D5$ simulates a PRAM with $n \log n$ processors. Therefore B needs $T/\epsilon n \log n$ steps on $D5$ and $t_{D5} = 50ST/\epsilon n$. We now can compute R:

$$R = \frac{c_{D1}}{c_{D5}} \cdot \frac{t_{D1}}{t_{D5}} \geq \frac{15598 n \log n + 10179 n}{12944 n \log n} \cdot \frac{1}{S} = \frac{1}{438} = L$$

∎

If we add floating point arithmetic L changes to $\frac{1}{2640}$.

6.3 Examples

In order to show that the bounds on R are tight we present two examples matching the bounds. The first example $B0$ is multiplication of two $s \times s$–matrices. We use design $D1$ with n physical processors and a distributed memory machine $D5$ with n processors interconnected as a $\sqrt{n} \times \sqrt{n}$ torus. Each processor of the torus then holds a $\frac{s}{\sqrt{n}} \times \frac{s}{\sqrt{n}}$–submatrix of both matrices. This example comes very close to the worst case described in section 6.1 and therefore R approximately matches the upper bound.

The second example $B1$ is computing the connected components of an undirected graph with v nodes and e edges. For the PRAM we use an algorithm of [29] in a form presented in [14]. Its runtime is $O(\log v)$ steps on a PRAM with $2e$ (virtual) processors. The formal explanation and the proofs for correctness and runtime can be found in [29]. On a PRAM with $n < v$ physical processors we have $t_{D1} = (300\frac{e}{n} + 108\frac{v}{n}) \log v$ as analyzed in [3] and c_{D1} as computed in equation 1. For the distributed memory machine we could use an algorithm from [5] that runs on a hypercube. Its runtime is $O(\log^2 v)$ on v^3 processors. For a hypercube $D6$ with $n < v$ processors we would have $c_{D6} = 12944n$ as computed in section 6.1, $t_{D6} = 40\frac{v^3}{n}(\log v)^2$ as sketched in appendix B. Using the fact that $e \leq \frac{1}{2}v^2$ leads to $R \approx 5\frac{\log n}{v \log v}$. This would imply $R < L$ and therefore a simulation of the PRAM algorithm is more cost–effective.

7 Conclusions

We have used the framework from [23] which allows to treat computer architecture as a formal optimization problem and to deal quantitatively with hardware/software tradeoffs. In this framework we have improved the price/performance ratio of RANADE's Fluent Machine by a factor of 6. We have determined when combining should be done in hardware (namely always for practical purpose). We have compared the

$$c_{D2}t_{D2} \leq c_{D1}t_{D1}$$
$$\overset{(4)}{\Rightarrow} \quad \alpha \leq \frac{\frac{c_{D1}}{c_{D2}} - 1}{t_{sim} - 1}$$

If we assume in favour of $D2$ that $n' = c'n \log n$ then with equations 1, 2, 3 we get $\alpha \leq 0.117(\log n)^{-2}$.
■

For moderate n however, the exact value is even smaller.

5.4 Consequences

We mentioned in section 1 that PRAMs are classified in theory as EREW, CREW and CRCW PRAMs. Relations among these classes are given in [13, 18]. A further class of ERCW PRAMs is not considered there.

Definition 8 *A machine model A is said to be hierarchically weaker than B ($A \preceq B$) if each problem that can be solved on model A in time T and P processors can also be solved on model B in time $O(T)$ and $O(P)$ processors.*

Obviously $EREW \preceq CREW \preceq CRCW$.

Theorem 2 *If we change our CRCW design $D1$ to an EREW design $D2$, an ERCW design $D3$ and an CREW design $D4$ we get the relation*

$$c_{D2} < c_{D3} < c_{D4} = c_{D1}.$$

Thus if a PRAM supports combining in the way we described in section 3.1 it is not worthwhile to consider CREW PRAMs but it might be useful to examine the role of ERCW PRAMs in the hierarchy.

Proof: (of theorem 2)

We get $D3$ from $D1$ by reducing the width of the direction queues with the same argument as in subsection 5.2. This shows $c_{D3} < c_{D1}$. We cannot skip the comparators because we still have to detect concurrent writes. This shows $c_{D2} < c_{D3}$. For $D4$ we cannot skip the comparators because we have to detect concurrent reads. We cannot reduce the width of the direction queues because of the same argument. This shows $c_{D4} = c_{D1}$. ■

Theorem 2 shows that $D4$ is identical to $D1$ and that for any PRAM program B $t_{D1} = t_{D4}$. Thus $D4$ has the same TDC as $D1$ but $D1 \npreceq D4$.

6 PRAMs vs. Distributed Memory Machines

PRAMs have always been thought to be uncompetitive to Distributed Memory Machines (DMM) because some problems do not need the global memory. In order to compare our PRAM $D1$ with a DMM $D5$ one has to compute $R = \text{TDC}(D1, B)/\text{TDC}(D5, B)$. We are interested in how much more cost-effective DMMs can be than PRAMs and vice versa. Therefore we search for bounds U and L with $L \leq R \leq U$ independently of B and of the particular DMM. It will turn out that for reasonable values of n a DMM cannot be much more cost-effective than a PRAM but vice versa a PRAM can be much more cost-effective than a DMM.

6.1 Simulation of DMMs by PRAMs

Assume a benchmark that does not use the global memory but can be run on a distributed memory machine with simple hardwired communication. This is the worst case that can happen when comparing PRAMs and DMMs. We formulate an upper bound as theorem 3.

Theorem 3 *Assume we have a benchmark B as has just been described that has enough parallelism to be computed on a distributed memory machine with efficiency ϵ close to 1. We consider a DMM $D5$ with n processors and communication given by a graph of small degree with n nodes and our PRAM $D1$. Then we get*

$$R \leq U \leq 1.21 \log n + 0.79.$$

Proof: The distributed memory machine with n processors has costs $c_{D5} = nc_{\tilde{P}} = 12944n$. We only count processor costs $c_{\tilde{P}}$ and ignore network costs although this is unfair towards the PRAM. Suppose that B needs T steps on a sequential machine. The DMM needs $T/\epsilon n$ steps. We assume in favour of the DMM that the benchmark B can be pipelined perfectly and thus one step takes only one cycle. Thus one has $t_{D5} = 50T/\epsilon n$.

The PRAM has costs c_{D1} as computed in equation 1 and needs $\frac{T}{\epsilon cn \log n}$ steps each taking $c \log n$ processor cycles. Thus $D1$ needs $T/\epsilon n$ cycles and therefore $t_{D1} = 50T/\epsilon n$. We then get

$$R = \frac{c_{D1}}{c_{D5}} \cdot \frac{t_{D1}}{t_{D5}} = \frac{15598n \log n + 10179n}{12944n} \cdot 1$$
$$= 1.21 \log n + 0.79 = U.$$

■

For reasonable values of n, e.g. $n \leq 2^{16}$, the quotient is less than 20. If we would add floating point arithmetic to the ALU as usual in existing parallel machines, the parameter A increases to $A' \approx 100000$ [12] and the quotient decreases dramatically to $0.2 \log n + 0.97$. For $n \leq 2^{16}$ the quotient is smaller than 4.2. If cost of memory is considered too, things change further in favour of the PRAM.

9

Suppose processor P_i wants to access variable V_j. Then it writes (i, j) to location i in global memory (we assume that locations 0 to $n - 1$ are not accessed by the PRAM program). The contents of locations 0 to $n - 1$ get sorted now by j. Duplicates which represent concurrent accesses are replaced by dummy accesses $(i, -j)$. P_i reads the content (i', j') of location i and accesses $V_{j'}$ if $j' \geq 0$. Then P_i writes the result of a READ access to location i'. The processors with eliminated duplicates duplicate now the results. At last P_i reads the result of its own access from location i and assigns it to variable V_j.

The most time consuming part of the simulation is the sort of the tuples (i, j). The sort can be parallized by using all n processors to sort the n tuples. Because a sequential sort by comparison of n elements needs time $\Omega(n \log n)$, an optimal parallel algorithm using all n processors should need parallel time $\Theta(\log n)$. Optimal sorting algorithms are described in [4, 10, 24], a randomized one is given in [28]. The constant factor in their runtime however is quite large. We will use BATCHER's *bitonic sort* [7], a parallel sorting algorithm with small constant that needs time $O(\log^2 n)$ to sort n elements using n processors. The bitonic sorting network can be defined recursively as in definition 7.

Definition 7 *B_2 and S_2 are identical circuits sorting two numbers. B_n is a circuit that merges two bitonic sequences each of length $\frac{n}{2}$ to one bitonic sequence of length n. The bitonic sorting network for n numbers is a circuit S_n. For one of these circuits S, $\tilde{S}$ denotes the circuit with reversed order of outputs.*
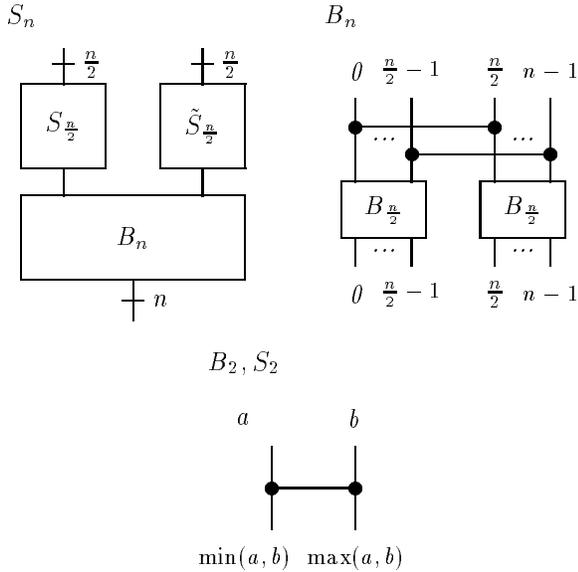
The bitonic sorter can be formulated as a program.

```
for pnum := 0 to n − 1 pardo
  for i := 1 to log n do
    for k := i − 1 to 0 do
      if bit k of pnum = 0 then
        if bit i of pnum = 0 then
          A[pnum] := min(A[pnum], A[pnum + 2^k])
        else
          A[pnum] := max(A[pnum], A[pnum + 2^k])
        fi
      else
        if bit i of pnum = 1 then
          A[pnum] := min(A[pnum − 2^k], A[pnum])
        else
          A[pnum] := max(A[pnum − 2^k], A[pnum])
        fi
      fi
    od
  od
od;
```

Figure 6: Bitonic Sort Algorithm

The program needs n processors that simulate in step i the n comparators B_2 in depth i of the circuit. The algorithm looks as shown in figure 6.

We assume that the compiler for our benchmark can recognize all instructions in which concurrent access can occur and that only these instructions are simulated in the way described above. We further assume that the compiler knows the number of processors that are working at this time. Now the compiler can generate code for the bitonic sort without using loops or subroutine calls. This makes it much faster. An assembler program would need $9.5 \left(\log n'\right)^2 + 10.5 \log n'$ instructions for the bitonic sort as described above. n' is the smallest power of two larger than the number of processors. In our design $D2$ $n' \geq c'n \log n$. The complete simulation of one step then takes

$$t_{sim} = \frac{19}{2} \left(\log n'\right)^2 + \frac{47}{2} \log n' + 46. \qquad (3)$$

instructions. The complete assembler program can be found in [2]. Now we will prove theorem 1 using the results of the previous subsections.

Proof: (indirect) Let B be a benchmark that needs time t_{D1} on $D1$. On $D2$ it will need time

$$t_{D2} = t_{D1} \left(\alpha t_{sim} + (1 - \alpha) 1\right). \qquad (4)$$

$$\text{TDC}(D2, B) \quad \leq \quad \text{TDC}(D1, B)$$

8

A	F	N_A	N_S	S_A	S_S
13572	6144	2576	6696	1928	4104

Table3: Actual Parameters

for its arithmetic part and N_S for its SRAM, the basic costs of a sorting node S_A and S_S.

Then we have costs $c_N = \rho_A N_A + \rho_S N_S$ for a network node and $c_S = \rho_A S_A + \rho_S S_S$ for a sorting node.

The improved machine consists of n physical processors, of $2\frac{c}{4}n \log n$ sorting nodes and of $n \log n$ network nodes. It has total costs

$$c_{D1} = nc_P + c_S \frac{2c}{4}n \log n + c_N n \log n.$$

The exact numbers for A, F, N_A, N_S, S_A, S_S are shown in table 3, the computation can be found in appendix A. The result is

$$c_{D1} = 10179n + 15598n \log n. \qquad (1)$$

The Fluent Machine's network nodes have slightly larger basic costs $\tilde{N}_A = 3104, \tilde{N}_S = 8808$ because RANADE's routing algorithm needs full routing information in forward and backward network. $c_{\tilde{N}}$ is computed in analogy to c_N. The costs of the Fluent Machine then are

$$c_{D0} = c_{\tilde{P}}n \log n + c_{\tilde{N}}n \log n = 19235.4n \log n.$$

For $n = 128$ the Fluent Machine is 1.128 times more expensive than our improved machine.

4.2 Speed of the Machine

In section A.3 we compute the maximal delay path in network nodes. We get a minimal cycle time of $\sigma_N = 25$ gate delays for the network and sorting nodes. For a particular processor design in [2] we computed a minimal cycle time $\sigma_P = 50$ gate delays, which comes from access times to the register file. In current VLSI technology with gate delays of $2ns$ we get cycle times of $100ns$ and $50ns$.

One step of the improved machine takes $c \log n$ processor cycles which is $\nu = c \log n \times 100ns$. RANADE reports in [27] simulation results such that one step of the Fluent machine takes $11 \log n$ network cycles which is $\nu' = 11 \log n \times 50ns$.

The improved machine then has a power of $\frac{cn \log n}{\nu}$ Instructions per second. For $n = 128$ we get 1280 MIPS. The corresponding value for the Fluent Machine is 232 MIPS. Thus the improved machine 5.5 times faster and 6.2 times more cost–effective than the Fluent Machine.

5 CRCW vs. EREW

5.1 Main Result

We investigate the question whether combining should be done by hardware *(hardwired combining)* or whether concurrent accesses should be simulated by software. We will prove the following theorem 1.

Theorem 1 *Let $D1$ be a CRCW PRAM as described in section 3.3 which supports combining by hardware. Let $D2$ be an EREW PRAM as described in section 5.2 on which each concurrent access is simulated by software as described in section 5.3. If a benchmark B that needs t_{D1} steps consists of αt_{D1} concurrent accesses with $0 \le \alpha \le 1$ then*

$$\mathtt{TDC}(D1,B) < \mathtt{TDC}(D2,B) \quad for \quad \alpha > 0.117 \frac{1}{(\log n)^2}.$$

This means: if a benchmark that needs t_{D1} steps consists of more than $0.117(\log n)^{-2}t_{D1}$ concurrent accesses it is better to run it on a CRCW PRAM instead of simulating it on an EREW PRAM.

5.2 Design of an EREW PRAM

To determine $\mathtt{TDC}(D2,B)$ it is necessary to sketch the design of an EREW PRAM $D2$. We get $D2$ from $D1$ by skipping all hardware that supports combining. These are the sorting networks in phases 1 and 6 of the routing and the comparators in the network switches which detect that combining is necessary. Additionally one can reduce the width of the direction queues in the switches to two bits because only four cases remain: 'in_i to out_j' where $i, j \in \{0, 1\}$. Removing the sorting networks reduces routing time and c can be decreased to $c' = 1.5$. The costs for the new processors ar $c_{P'} = \rho_A A + \rho_L c' \log nF = 10179 + 2857 \log n$. The costs for network and sorting nodes decrease from N_A to $N'_A = 2320$ and N_S to $N'_S = 6192$ as shown in appendix A. The total costs for $D2$ are

$$c_{D2} = c_{P'}n + c_{N'}n \log n = 10179n + 7389.4n \log n. \quad (2)$$

The cycle time of $D2$ is exactly the same as of $D1$, one step of $D2$ takes $c' logn$ processor cycles.

5.3 Simulation of CRCW on EREW

KARP and RAMACHANDRAN show in [18] how to simulate a CRCW PRAM on an EREW PRAM. They use the following method to simulate one step in which concurrent accesses can happen:

7

ing arrays is discussed in section 3.4. VALIANT calls this *parallel slackness* [33].

Definition 6 *A* round *in machine $D1$ is the time interval from the moment when the first vP injects its packet into the network to the moment when the last vP injects its packet into the network.*

At the end of a round there are on the one hand still packets of this round in the network, on the other hand the processors have to proceed (and thus must start the next round) to return these packets. CHANG and SIMON prove in [9] that this works and that the latency still is $O(\log n)$. The remaining problem how to separate the different rounds can easily be solved. After the last vP has injected its packet into the network, an *End of Round Packet (EOR)* is inserted. This is a packet with a destination larger than memory size m. Because the packets leave each node sorted by destinations, it has to wait in a network switch until another EOR enters this switch across its other input. It can be proved easily that this is sufficient to separate rounds.

3.4 Delayed `LOAD` and Sorting

One problem to be solved is that virtual processors that execute a `LOAD` instruction have to wait until the network returns the answer to their `READ` packets. Simulations indicate, that for $c = 6$ this works most of the time (see [3]). But this is quite large in comparison to $\log n$. We partially overcome this by using delayed `LOAD` instructions as in [25]. We require an answer to a `READ` packet being available not in the next instruction but in the next but one. Investigations show that insertion of additional 'dummy' instructions happens very rarely [25]. But if a program needs any dummy instructions, they can be easily inserted by the compiler. This reduces c to 3 without significantly slowing down the machine.

The sorting arrays should have length $c \log n$ too. But breaking a round in z parts is an alternative. This reduces the lengths to $\frac{c}{z} \log n$ but could slow down the machine's speed. Simulations show [3] that $z = 4$ is the maximum value that does not slow down speed if we double the sorting networks. Therefore we choose this value.

In order to examine the exact constants for runtime and costs in machine $D1$ by the method sketched in section 1.1 we have to model the processor for this machine. In [27] nothing special about it is mentioned except that the use of RISC processors is proposed.

3.5 A Processor

We use a processor similar to the Berkeley RISC processor [25]. Instead of register windows we have the register sets of the virtual processors. The processor has a `LOAD-STORE` architecture, i.e. that `COMPUTE` instructions only work on registers and immediate constants and that memory access only happens on `LOAD` and `STORE` instructions. The `COMPUTE` instructions involve adding, multiplying, shifts and bit oriented operations. All instructions need the same amount of time. In order to get a pipeline of depth $c \log n$, the ALU depth is increased artificially.

Because of the `LOAD-STORE` architecture the same multiplier can be used for multiplications in `COMPUTE` instructions and for hashing global addresses with a linear hash function in `LOAD` and `STORE` instructions. This means that hashing does not require much special hardware.

A more detailed description of the processor can be found in [3].

4 Cost and Speed

4.1 Cost of the machine

We compute the costs of the improved Fluent Machine with the method introduced in section 1.1. We will ignore control logic because it occupies only a fraction of at most 10 percent of the total costs. This would change if we would use CISC processors.

The RISC processor of section 3 mainly consists of an ALU and a register file. The ALU consists of a 32 bit WALLACE tree multiplier, a barrel shifter and a carry lookahead adder [35]. The register file of the Fluent Machine consists of 16 registers each 32 bits wide, the one in the improved machine consists of $c \log n \times 16$ registers each 32 bits wide. Let the basic costs of the ALU be A and the basic costs of the Fluent Machine's register file be F.

If we use the packing factors of table 2 we have costs $c_P = \rho_A A + \rho_L c \log n F$ for the processor of our design $D1$ and $c_{\bar{P}} = \rho_A A + \rho_S F$ for the Fluent Machine's processor.

Simulations [3] indicate that network nodes need buffers of length 2. A node consists of 2 buffers and 2 multiplexers on the way from processors to memory, 2 buffers and 2 multiplexers on the way back, a direction queue of length $2c \log n$ and a comparator and a subtractor to compare addresses. Sorting nodes only need buffers of length 1 and 1 multiplexer for each direction. Let the basic costs of a network node be N_A
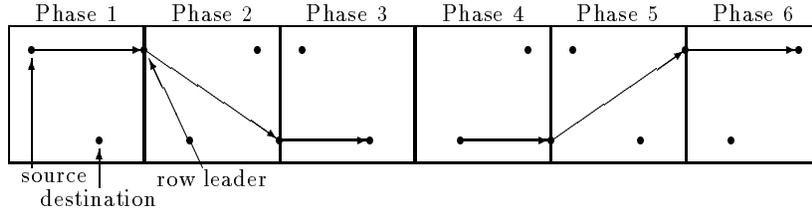
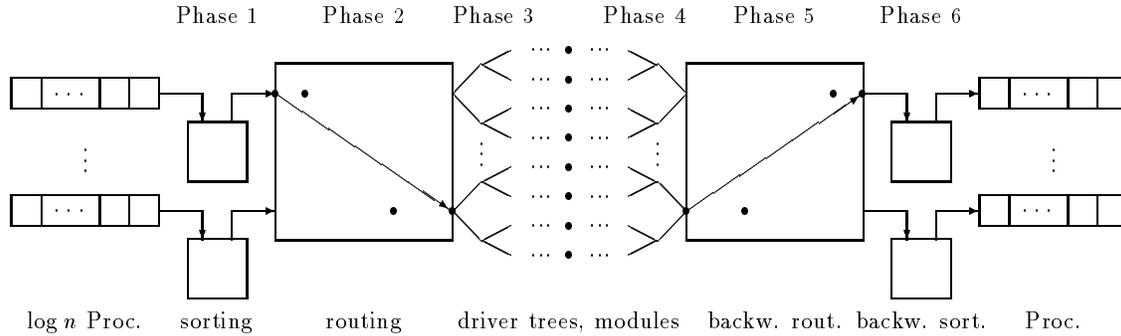Figure 3: 6 phase routing of the Fluent Machine

Figure 4: 6 phase Routing in the New Machine

ets. This cannot be avoided. But we can reduce the cost of the idle hardware by replacing the $\log n$ processors of a row by only one physical processor (pP) which simulates the original $\log n$ processors as virtual processors (vP). Another advantage of this concept is that we can increase the total number of PRAM processors by simulating $X = c \log n$ (with $c > 1$) vP's in a single pP. The simulation of the virtual processors by the physical processor is done by the principle of *pipelining*. This principle is well known from vector computers and was also used in the first MIMD computer marketed commercially, the Denelcor HEP [16, 30]. A closely related concept is *Bulk Synchronous Parallism* in [34].

In vector processors the execution of several instructions is overlapped by sharing the ALU. Figure 5 shows how pipelining is used in our design. Here the ALU needs x cycles. A single instruction in this example needs $x + 4$ cycles. Execution of t instructions needs $t + x + 3$ cycles. Without pipelining they need $t(x + 4)$ cycles.

Instead of accelerating several instructions of a vector processor with a pipeline, we use pipelining for overlapped execution of one instruction for all X vP's that are simulated in one physical processor. To simulate X vP's we increase the depth of our ALU ar-

Time	1	2	3	4	5	6		x+3	x+4
Stage									
Fetch	I1	I2	I3						
Decode		I1	I2	I3					
Load arguments			I1	I2	I3				
Compute cycle 1				I1	I2	I3			
$\vdots$									
Compute cycle x								I1	I2
Store results									I1

Figure 5: Pipelining in the Processor

tificially to $x = X - 4$. The virtual processors are represented in the physical processor simply by their own register sets. We save the costs of $X - 1$ ALU's.

The depth δ of this pipeline serves to hide network latency. This latency is proved to be $c \log n$ for some c with high probability [26]. Thus, if $\delta = c \log n$, then normally no vP has to wait for a returned packet. This c increases the number of vP's and the network conguestion. But network latency only grows slowly with increasing c. Thus there exists an optimal c. The exact value and its influence on the length of the sort-

5

row. Node $\langle col, row \rangle$, $col < \log n$ *is connected to node* $\langle col + 1, row \rangle$ *and to node* $\langle col + 1, row \oplus 2^{col} \rangle$, *where* $\oplus$ *denotes the bitwise exclusive or.*

Each network node contains a processor, a memory module of the shared memory and the routing switch. If a processor $\langle col, row \rangle$ wants to access a variable V_x it generates a packet of the form (destination,type,data) where destination is the tuple $(h(x), l(x))$ and type is READ or WRITE. This packet is injected into the network and sent to node $h(x) = \langle row', col' \rangle$ and back (if its type is READ) with the following six phase deterministic packet routing algorithm.

1. The packet is sent to node $\langle \log n, row \rangle$. On the way to column $\log n$ all packets injected into a row are sorted by their destinations.

2. The message is routed along the unique path from $\langle \log n, row \rangle$ to $\langle 0, row' \rangle$. The routing algorithm used is given in [26].

3. The packet is directed to node $\langle col', row' \rangle$ and there the memory access takes place.

4. – 6. The packet is sent the same way back to $\langle col, row \rangle$.

Figure 3 shows the phases performed on a network consisting of 6 butterflies. RANADE realizes these six phases with two butterfly networks where column i of the first network corresponds to column $\log n - i$ of the second one. Phases 1,3,5 use the first network, phases 2,4,6 use the second network. Thus the Fluent Machine consists of $n \log n$ nodes each containing one processor, one memory module and 2 network switches.

The reason for sorting in phase 1 is given in section 3.2.

3.2 Combining

In a CRCW PRAM several (possibly all) processors could access the same variable V_j at the same time. Let

$$S_j = \{P_i | P_i \text{ reads } V_j \text{ in current step}\},$$
$$PAC_j = \{pac_i | P_i \in S_j \text{ sends } pac_i \text{ into network}\}.$$

We talk only of READ accesses because WRITE accesses can be treated in a similar way with the simplification that they do not return an answer to the processor.

If all packets in PAC_j reach memory module $h(j)$, the module congestion c_m equals $|PAC_j|$. In the worst

case this could be n. Because the routing algorithms require module congestion $O(\log n)$ (see [3, 20]) the number of packets in PAC_j that reach $h(j)$ has to be reduced in the following way: The paths of the packets in PAC_j form a tree. However there is no need to send more than one packet along any branch of this tree. If a packet $pac_i \in PAC_j$ simply waits at each tree node until a packet $pac_l \in PAC_j$ appears along the other incoming edge (unless the node 'knows' that all future packets of the current step must originate from processors $P \notin S_j$), then the two packets can be merged and one forwarded along the tree. This merging is called *combining*.

In order to decide whether two incoming packets $pac_i \in S_x, pac_l \in S_y$ have to be combined, a network node has to compare the destinations $g(x)$ and $g(y)$.

How can a network node know that no more packets will arrive in the future? RANADE gives in [26] the following solution: sort the packets during phase 1 by their destinations and then maintain for each node the packets that leave the node sorted.

3.3 Improvements

Definition 5 *A* round *is the time interval from the moment when the first of all* $n \log n$ *packets is injected into the network to the moment when the last packet is returned to its processor again with the answer of a* READ *access.*

In RANADE's algorithm the next round can only be started when the actual round is finshed completely. This means that overlapping of several rounds (*pipelining*) is not possible in the Fluent Machine. This is the first disadvantage that we want to eliminate. This could be reached by using 6 physical butterfly networks as shown in figure 3. But the networks for phases 1 and 6 can be realized by n sorting arrays of length $\log n$ as described in [3, 19] and networks for phases 3 and 4 can be realized by driver trees respective OR trees. Both solutions have smaller costs than butterfly networks and are not slower. The sorting arrays only have one input and require that all $\log n$ processors of a row inject their packets sequentially into this input.

This leads to the following construction as shown in figure 4. The $\log n$ processors of a row inject their packets into the sorting array sequentially, the sorted packets are routed like in RANADE's phase 2, the packets are directed to the right modules via driver trees. Then the packets go all the way back to their processors.

The second disadvantage is that the processors spend most of the time waiting for returning pack-

4

ally compares two families of machines, the members of which are only different in size. Their costs and the runtime of the benchmark depend on the number of processors. To compare the families we take corresponding "representatives" of them. These could be members of the two families that have identical costs or that have equal processor numbers.

2 The PRAM Model

Definition 3 *An n–PRAM (parallel random access machine) is a parallel register machine with n processors $P_0, \ldots, P_{n-1}$, their local memories and a shared memory of size m. In each step each processor can work as a separate register machine or can access a cell of the shared memory. The processors work synchronously.*

We consider the following kinds of PRAMs:

- **EREW:** *(exclusive read exclusive write)* It is not possible to read or write a memory cell simultaneously with several processors.

- **CREW:** *(concurrent read exclusive write)* It is only possible to read a cell simultaneously.

- **CRCW:** *(concurrent read concurrent write)* Processors can read and write a cell simultaneously. Concurrent write forces to define which one of the concurrent processors will win. Usually three possibilities are studied:

  - **arbitrary:** One processor wins, but it is not known in advance which one wins.

  - **common:** All processors must write identical data, thus it does not matter which one wins.

  - **priority:** The processor with the largest or lowest index wins.

The last model is the most powerful. Overviews about algorithms for the different models can be found in [5, 13, 18].

One simulates an n–PRAM on a multi–computer machine (MIMD) by distributing the shared memory uniformly among memory modules $M_0, \ldots, M_{n-1}$ each of size $\frac{m}{n}$. The processors are connected by an interconnection network. Processor P_i communicates with module M_j on this network if P_i wants to access a variable that is stored in M_j $(i \neq j)$. To communicate means that P_i sends a packet to M_j specifying the required variable. In case of a `LOAD` instruction M_j sends the answer back to P_i.

$l(x)$	$h(x)$	$g(x)$

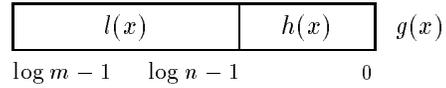$\log m - 1 \quad \log n - 1 \qquad\qquad 0$

Figure 2: Binary representation of Hash function $g(x)$

In order to map the variables used $V_0, \ldots, V_{r-1}, r \leq m$ onto the memory modules one uses a hash function $g : \{0, \ldots, m-1\} \rightarrow \{0, \ldots, m-1\}$ which is actually a tuple (h, l) of functions $h : \{0, \ldots, m-1\} \rightarrow \{0, \ldots, n-1\}$ and $l : \{0, \ldots, m-1\} \rightarrow \{0, \ldots, \frac{m}{n}-1\}$. The function h specifies the module, l specifies the location within the module. One gets h and l from g in the following way: $h(x) = g(x) \bmod n$, $l(x) = g(x)$ div n. If one considers the binary representation of g(x) with length $\log m$, the last $\log n$ bits give the module $h(x)$, the upper $(\log m - \log n)$ bits give the location $l(x)$ within the module. The binary representation of $g(x)$ looks as shown in figure 2.

Hash functions that distribute variables provably well are examined in [3]. An example are polynomials. Simulations [3, 27] indicate that for practical use particular linear hash functions g of the type $g(x) = ax \bmod m$ with greatest common divisor $\gcd(a, m) = 1$, $0 \leq a \leq m - 1$ are good enough. The advantages of the function $g(x)$ are its bijectivity and the short evaluation time.

The communication between processors and memory modules can be handled in several ways, e.g. as in [20, 26]. We base our work on RANADE's Fluent Machine as described in section 3.1.

3 The Machine D1

We first give a short summary of the Fluent Machine which is precicely described in [3, 26, 27]. Then we present some improvements that lead to our design $D1$.

3.1 The Fluent Machine

The Fluent Abstract Machine simulates a CRCW priority PRAM with $n \log n$ processors. The processors are interconnected by an $n \log n$ butterfly network as given in Definition 4.

Definition 4 *The* butterfly network *of degree 2 consists of $n(1 + \log n)$ network nodes. Each node is assigned a unique number $\langle col, row \rangle$ where $0 \leq col \leq \log n, 0 \leq row \leq n - 1$. $\langle col, row \rangle$ can be viewed as the concatenation of the binary representations of col and*

3

	INV	AND, OR	EXOR	1 bit Reg.
cost	1	2	6	12
delay	1	1	3	5

Table1: Basic cost and delay functions

Structure	Parameter	Value
Logic	ρ	1
Arithmetic	ρ_A	0.75
small SRAM	ρ_S	0.45
large SRAM	ρ_L	0.31

Table2: Packing Factors

a formalism from [23] which permits to compare cost–effectiveness of architectures. It will turn out that the reengineered version of the Fluent Machine is more than 5 times more cost–effective than the original machine and that it is surprisingly cost–effective when compared to distributed memory machines.

In section 1.1 we define the formalism to compare machines. Chapter 2 describes the theoretical PRAM model. Chapter 3 contains the description of the Fluent machine and the reengineered version. In chapter 4 we analyze both machines and compare them in the formalism given in section 1.1. In chapter 5 we show that it is worthwhile to support concurrent accesses by hardware. In chapter 6 we compare PRAMs and distributed memory machines.

1.1 Comparison of Machines

Definition 1 *Let D be a design of a machine with cost c_D. Let B be a program with runtime t_D on design D. B is called* **benchmark**. *We call $c_D t_D$ the* **time depending cost function** TDC *of design D with benchmark B.*

A motivation for the TDC is the well–known price/performance ratio, if we take performance as the reciprocal value of runtime at constant work B.

We determine c_D and t_D of a machine by specifying the whole machine by circuits and switching networks. Each type of gates has basic cost and delay given by functions *cost* and *delay*. Examples are shown in table 1. The cost of a circuit is the sum of the basic costs of its gates multiplied with *packing factors* which are examples of *technology parameters*. They represent the fact that structures such as logic, arithmetic and static RAM can be packed more or less densely. Typical parameters for different technologies can be derived from chip producers' statements about placement results. We will use particular parameters derived from [22] which are shown in table 2. The cost of a machine is the sum of the costs of all switching networks, main memory is not counted.

We take a carry–chain adder for 8–bit numbers as an example. It consists of 8 fulladders. A fulladder consists of two halfadders and an OR gate. A halfadder consists of an AND gate and an EXOR gate. We have

8 OR gates, 16 AND gates and 16 EXOR gates in total. The adder is an arithmetic unit and thus has a packing factor of 0.75. The cost of the adder is $0.75(8 cost(\texttt{OR}) + 16 cost(\texttt{AND}) + 16 cost(\texttt{EXOR})) = 108$.

We compute the execution times of the machine instructions (ignoring delays on wires) by searching for the maximum delay of all paths in all circuits. The delay of a path is the sum of the gate delays on this path plus a short time to load a register at the end of the path. This is a lower bound for the cycle time. The execution time of a machine command is the cycle time multiplied with the number of cycles the command needs (if all cycles have equal length).

In our example the longest path is the following one: in the first fulladder from input a_{in} or b_{in} to $carry_{out}$, in the 2^{nd} to the 7^{th} fulladder from $carry_{in}$ to $carry_{out}$, in the 8^{th} fulladder from $carry_{in}$ to sum_{out}. If the $carry_{in}$ of a fulladder goes to the 2^{nd} halfadder, our path meets an EXOR, an AND and an OR in the 1^{st} fulladder, an AND and an OR in the 2^{nd} to the 7^{th} fulladder and an EXOR in the 8^{th} fulladder. The total delay is $T_{total} = 7 delay(\texttt{AND}) + 7 delay(\texttt{OR}) + 2 delay(\texttt{EXOR}) = 20$.

We formulate benchmarks in PASCAL with the **pardo** construct [13] as parallel extention. This is sufficient for an analysis. But implementation of this language would be difficult. A better solution is given by the language FORK [15].

We determine the runtime of a benchmark B by compiling it by hand and analyzing the machine code. Depending on the CPU architecture the result is something like the number of LOAD, STORE and COMPUTE commands. For each group we multiply its number of commands with its execution time, then we sum over the groups. The result is the runtime t_D in gate delays. If pipelining is allowed, things become messier, but can still be handled.

Definition 2 *If two designs $D0$ and $D1$ have costs c_{D0} and c_{D1} and a benchmark B has runtime t_{D0} on $D0$ and t_{D1} on $D1$ then $D0$ is called better on B than $D1$ if and only if* TDC$(D0, B) < $ TDC$(D1, B)$.

If one compares scalable parallel machines, one re-

2

On the Cost–Effectiveness of PRAMs[*]

Ferri Abolhassan   Jörg Keller   Wolfgang J. Paul

Institute for Computer Architecture and Parallelism

Computer Science Dept., Universität des Saarlandes

W–6600 Saarbrücken, Germany

Abstract

We introduce a formalism which allows to treat computer architecture as a formal optimization problem. We apply this to the design of shared memory parallel machines. Present computers of this type support the programming model of a shared memory. But simultaneous access to the shared memory by several processors is in many situations processed sequentially. Asymptotically good solutions for this problem are offered by theoretical computer science. We modify these constructions under engineering aspects and improve the price/performance ratio by roughly a factor of 6. The resulting machine has surprisingly good price/performance ratio even if compared with distributed memory machines. For almost all access patterns of all processors into the shared memory, access is as fast as the access of only a single processor.

Figure 1: Concurrent Write on ALLIANT FX/2816

1 Introduction

Commercially available parallel machines can be classified as *distributed memory machines* or *shared memory machines*. Exchange of data between different processors is done in the first class of machines by explicit *message passing*. In the second class programs on different processors simply access variables in a *common address space*. Thus one gets a more comfortable programming model.

One is tempted to suspect big differences between the hardware architectures of the two classes, but this is actually not so. Processors of present shared memory machines tend to have local memories as well as large caches, and the exchange of cache lines between processors can be viewed as an automated way of message passing. As a consequence of this implementation on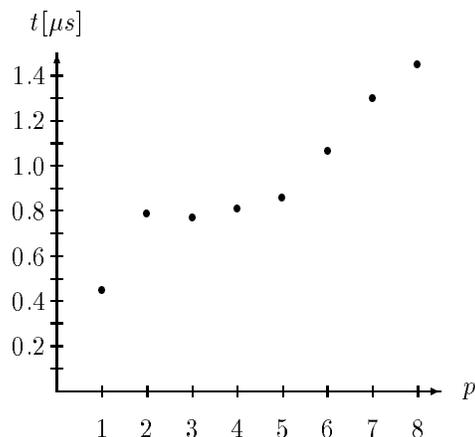e gets a large variation of the memory access time depending on the access patterns of the processors. In fact a single concurrent write of all say p processors of a parallel machine to the same memory location might very well be slower than p accesses of a single processor to its local memory. As an example figure 1 shows the time of a concurrent write by $p = 1, \ldots, 8$ processors to the same memory location in an ALLIANT FX/2816 [1]. Thus present shared memory machines support only the programming model but not the timing behaviour of a true shared memory.

Parallel machines which support both the programming model and the timing behaviour of true shared memory are called PRAMs in the theoretical literature. The problem of simulating PRAMs by more technically feasible models has been extensively studied [6, 8, 11, 17, 20, 21, 27, 31, 32, 34]. The construction from [27], called the Fluent Machine, is considered as the most efficient simulation.

In this paper we will describe the design of a reengineered version of the Fluent Machine. We will review

---