SIAM J. COMPUT.
Vol. 16, No. 1, February 1987

© 1987 Society for Industrial and Applied Mathematics
008

THE COMPLEXITY OF PARALLEL SORTING*

FRIEDHELM MEYER AUF DER HEIDE† AND AVI WIGDERSON‡

Abstract. The model we consider is the (concurrent-write, PRIORITY) PRAM. It has n synchronous processors, which communicate via an infinite shared memory. When several processors simultaneously write to the same cell, the one with the largest index succeeds. We allow the processors arbitrary computational power.

Our main result is that sorting n integers requires $\Omega(\sqrt{\log n})$ steps in this strong model. This bound is proved in two stages. First, using a novel Ramsey theoretic argument, we "reduce" sorting on a PRAM to sorting on a parallel merge tree. This tree is a generalization of Valiant's parallel comparison tree from [V] in which at every step n pairs of (previously ordered) sets are merged (rather then n pairs of elements compared). The second stage is proving the lower bound for such trees.

The Ramsey theoretic technique, together with known methods for bounding the "degree" of the computation, can be used to unify and generalize previous lower bounds for PRAM's. For example, we can show that the computation of any symmetric polynomial (e.g. the sum or product) on n integers requires exactly $\log_2 n$ steps.

Key words. parallel computation, PRAM, sorting, lower bounds, Ramsey theory

AMS(MOS) subject classification. 68Q25

1. Introduction. The interest in parallel sorting is obvious. In the last few years we have witnessed a multitude of upper bounds for parallel sorting on various models. These results culminated in the remarkable $O(n \log n)$ sorting network of Ajtai, Komlos and Szemeredi [AKS], and its beautiful adaptation to bounded degree n-node networks by Leighton [L].

Surprisingly, there are essentially no lower bound results for sorting. The reason seems to be that for many parallel models such bounds follow either from simple considerations or from parallel lower bounds on simpler functions. Here are a few examples.

1. Comparison trees and algebraic decision trees. The $\Omega(n \log n)$ sequential lower bound for sorting (Paul and Simon [PS]) implies an $\Omega(\log n)$ lower bound on the parallel versions of these trees with n processors.
2. Bounded degree n-node networks. An $\Omega(\log n)$ lower bound follows from diameter considerations.
3. Polynomial size, unbounded fan-in circuits. The $\Omega(\sqrt{\log n})$ on depth for computing the parity function, due to Ajtai from [A], yields a similar bound for sorting.
4. Exclusive-write PRAM. The $\Omega(\log n)$ lower bound for computing the Or function, due to Cook, Dwork and Reischuk from [CDR] implies a similar bound for sorting.

In all the parallel models mentioned above, there is a restriction on either the computational power of the individual processors, or on the nature of communication between them. Here we pose no such restrictions; we consider a concurrent-write PRAM, in which n processors of arbitrary computational power communicate via an infinite shared memory. Moreover, we use the strongest convention of resolving write conflicts; the processor with the largest index succeeds.

The only nontrivial lower bound on this model appears in [FMRW]. It is a (tight) $\Omega(\log \log n)$ bound on finding the maximum of n integers. As is almost traditional

* Received by the editors May 8, 1985; accepted for publication (in revised form) February 18, 1986.
† J. W. Goethe Universität, Frankfurt, West Germany.
‡ Mathematical Sciences Research Institute, Berkeley, California 94720.

with lower bounds for models i
depend on the input in a compl
and find some structure. Such
parallel comparison tree is ess
Then Valiant's $\Omega(\log \log n)$ lo

The problem with the R
variables are fixed to constar
proceeds, and hence it can ne

We prove an $\Omega(\sqrt{\log n})$ c
argument in which variables a
ordering is given to the algori
we show that essentially all a
(chains) of variables. It is natu
at every node n pairs of chain
tree is simply Valiant's paralle

We prove a general lowe
combine it with the dependenc
our $\Omega(\sqrt{\log n})$ lower bound.
improved to $\Omega(\log n)$.

As mentioned, our Ramse
ing of input variables to the alg
information of this type as pos
argument becomes much simp
sorted to begin with. Using this
"degree" of the computation fi
proof of the following result.
integers (e.g. the sum or the pr

This result was independer
in [IM]. It unifies and general
computational power of proces
interprocessor communication

2. Formal definitions and s
paper are the PRAM and the pa
of the PRAM.

A (PRIORITY) PRAM M
infinitely many shared memor
integers N. We say that M con
has the value a_i of the ith inp
to "0", and after T steps the f

Every step of the computa
by all processors. In the write p
In the read phase, each processo
and this write value depend on t
on the complexity of computin
at one step, its contents will be
them.

For the definition of the
Let V be a set of elements $\{x_1, x$
a partial order $<_Q$ in the natu

© 1987 Society for Industrial and Applied Mathematics
008

LEL SORTING*

ND AVI WIGDERSON‡

RIORITY) PRAM. It has n synchronous When several processors simultaneously ow the processors arbitrary computational

steps in this strong model. This bound is nent, we "reduce" sorting on a PRAM to liant's parallel comparison tree from [V] nerged (rather then n pairs of elements trees.

hods for bounding the "degree" of the ounds for PRAM's. For example, we can sum or product) on n integers requires

ds, Ramsey theory

s obvious. In the last few years rallel sorting on various models. sorting network of Ajtai, Komlos ounded degree n-node networks

d results for sorting. The reason nds follow either from simple pler functions. Here are a few

The $\Omega(n \log n)$ sequential lower ies an $\Omega(\log n)$ lower bound on essors.

n) lower bound follows from

e $\Omega(\sqrt{\log n})$ on depth for com- A], yields a similar bound for

nd for computing the Or func- CDR] implies a similar bound

e is a restriction on either the n the nature of communication e consider a concurrent-write nal power communicate via an t convention of resolving write

ears in [FMRW]. It is a (tight) egers. As is almost traditional

(in revised form) February 18, 1986.

94720.

with lower bounds for models in which the computation and communication behaviour depend on the input in a complicated way, Ramsey theory is used to "clear the smoke" and find some structure. Such an argument showed that, for finding the maximum, a parallel comparison tree is essentially (i.e. for a large input set) as good as a PRAM. Then Valiant's $\Omega(\log \log n)$ lower bound from [V] on such trees is applied.

The problem with the Ramsey theoretic argument in [FMRW] is that input variables are fixed to constants at a doubly exponential rate as the computation proceeds, and hence it can never yield a lower bound bigger than $\Omega(\log \log n)$.

We prove an $\Omega(\sqrt{\log n})$ on sorting n integers. We use a new Ramsey theoretic argument in which variables are not fixed, but instead some information about their ordering is given to the algorithm as the computation proceeds. With this argument we show that essentially all a processor can do in a step is to merge two ordered sets (chains) of variables. It is natural now to define a parallel (m, n)-merge tree, in which at every node n pairs of chains of size m each are merged. (Note that for $m = 1$ this tree is simply Valiant's parallel comparison tree.)

We prove a general lower bound, parametrized by m, for such trees. Then we combine it with the dependency of m on the time of the PRAM computation to obtain our $\Omega(\sqrt{\log n})$ lower bound. We believe that our lower bound on PRAM's can be improved to $\Omega(\log n)$.

As mentioned, our Ramsey theoretic argument gives information about the ordering of input variables to the algorithm. For sorting, we had to take care to give as little information of this type as possible. For symmetric functions, however, applying this argument becomes much simpler, since we can assume that the input variables are sorted to begin with. Using this idea, combined with known techniques to making the "degree" of the computation finite (albeit the infinite address space), we get a simple proof of the following result. The computation of any symmetric polynomial of n integers (e.g. the sum or the product) requires exactly $\log_2 n$ steps.

This result was independently obtained (at least for addition) by Israeli and Moran in [IM]. It unifies and generalizes similar bounds on models that restrict either the computational power of processors (Meyer auf der Heide and Reischuk [MR]) or the interprocessor communication (Parberry [P]).

2. Formal definitions and statements of results. The two important models in this paper are the PRAM and the parallel merge tree. We start with the (standard) definition of the PRAM.

A (PRIORITY) PRAM M consists of n processors (RAM's) $P_1, P_2, \cdots, P_n$, and infinitely many shared memory cells (which we abbreviate "cells"), indexed by the integers N. We say that M computes a function $f: N^n \to N^m$ in T steps, if initially P_i has the value a_i of the ith input variable in its local memory, all cells are initialized to "0", and after T steps the first m cells contain the m values of $f(a_1, a_2, \cdots, a_n)$.

Every step of the computation consists of two phases, synchronously performed by all processors. In the write phase, each processor writes some value into some cell. In the read phase, each processor reads some cells to its local memory. (These addresses and this write value depend on the processor's local memory, but we place no restriction on the complexity of computing them.) If a cell is written into by several processors at one step, its contents will be the value written by the highest index processor among them.

For the definition of the merge tree, we need some notation for partial orders. Let V be a set of elements $\{x_1, x_2, \cdots, x_n\}$. A directed acyclic graph $Q = (V, E)$ defines a partial order $<_Q$ in the natural way. $\bar{Q} = (V, \bar{E})$ is the *transitive closure* of Q. A set

$U \subset V$ is a *chain* (or a *total order*) in Q if the elements of U belong to a directed path in Q. A set U is an *antichain* in Q if U is an independent set in $\bar{Q}$.

Let $U_1, U_2 \subset V$ and $Q_1 = (U_1, E_1)$, $Q_2 = (U_2, E_2)$ be partial orders. Q_1 and Q_2 are *consistent* if there is no pair $a, b \in V$ such that $(a, b) \in \overline{E_1}$ and $(b, a) \in \overline{E_2}$. A family of partial orders on subsets of V, $Q_1, Q_2, \cdots, Q_k$ is *consistent* if every pair is consistent. The *union* of two consistent orders Q_1 and Q_2 is $Q_1 \cup Q_2 = (V, E_1 \cup E_2)$. Similarly we define $\bigcup_{i=1}^{k} Q_i$.

We say that $Q' = (V, E')$ is an *extension* of $Q = (V, E)$ iff $\bar{E} \subset \bar{E}'$. If Q' is a total order, we say that it is a *linear extension* of Q.

Let $C_1, C_2 \subset V$ be two chains in $Q = (V, E)$. A *merge* of C_1 and C_2 in Q is any total order $Q' = (C_1 \cup C_2, E')$, consistent with Q.

Now we are ready to define the model. An (m, n)-*merge tree* (for sorting) is a rooted tree with labels on the nodes and on the edges. The label of an internal node v is a partial order Q on V, together with n pairs of chains in Q, (C_{i1}, C_{i2}), $i = 1, 2, \cdots, n$, satisfying $|C_{ij}| \leqq m$ for all $i = 1, 2, \cdots, n$, $j = 1, 2$. Every branch e out of v is labeled with a different consistent family of merges C_i of C_{i1} and C_{i2}, $i = 1, 2, \cdots, n$. If u is the child of v connected by e, then the partial order in u is $Q \cup \bigcup_{i=1}^{n} C_i$. The root is labeled with the empty partial order on V. The cost of an (m, n)-merge tree H, denoted $c(H)$, is the length of the longest root-leaf path in H. H sorts n numbers, if each leaf is labeled with a total order. The cost of sorting on (m, n)-merge trees is $c(m, n)$, the minimum cost $c(H)$ of an (m, n)-merge tree H which sorts n numbers.

The main result of this paper is the following.

MAIN THEOREM. *Any* PRAM *needs* $\Omega(\sqrt{\log n})$ *steps to sort n integers.*

The main Theorem follows directly from the Theorems 1 and 2 below.

THEOREM 1. *Let* M *be a* PRAM *sorting n integers in T steps. Then there is a* $(2^T, n)$*-merge tree which sorts n numbers in T steps.*

THEOREM 2.

$$c(m, n) = \Omega\left(\frac{\log n}{\log (m \log n)}\right).$$

As a corollary to the proof method of Theorem 1, we get a general tight bound for a family of symmetric functions. For $S \subset N$ let $S_*^n = \{(a_1, \cdots, a_n) \in S^n \mid a_i \neq a_j \text{ for } i \neq j\}$. A function $f: N^n \to N$ is *strongly nonconstant*, if for each infinite $S \subset N$, f restricted to S_*^n is nonconstant.

THEOREM 3. *A* PRAM *needs exactly* $\log_2 n$ *steps to compute a symmetric, strongly nonconstant function.*

Example. Every nonconstant, symmetric polynomial on n variables (e.g. the sum or the product) is a symmetric, strongly nonconstant function.

3. Simulating PRAM's by merge trees.

THEOREM 1. *Let M be a* PRAM *sorting n integers in T steps. Then there is a* $(2^T, n)$*-merge tree which sorts n elements in T steps.*

We simulate M by a computation tree. We want to maintain at each node v of this tree in depth t the following property: Let I be the set of inputs arriving at v. Then, for inputs from I, each processor P_i only knows variables with indices from some set $X_i \subset [n]$ at time t. ($[n] = \{1, 2, \cdots, n\}$.) The exact meaning of this is that for inputs from I, the configuration of P_i after t steps (considered as a function in I) only depends on the variables with indices from X_i. In the sequel we shall refer to them as the variables from X_i for short.

Clearly, because of its arbitrary computational power, P_i can now sort the variables it knows in one step and can proceed dependent on their order. We shall see that this

is essentially the only property
if we restrict the input set suital
processors know are fixed, the
following sense. For each $i \in [n$
processor P_{j_i} at some fixed time
in this step, namely just those wl
before, the behaviour of P_i in the
of the two ordered sets of variab

In the sequel we only cons
that for a set $S \subset N$ $S_*^n = \{(a_1, \cdot$
and π_i be a total order on X_i. L
that the π_i's are consistent. Th
arriving at a node of the
$\{(a_1, \cdots, a_n) \in S_*^n \mid a_p < a_q$ for al

The following lemma is the

MAIN LEMMA I. *Suppose*
knows variables from X_i*. Then th*
after step t, for inputs from $I' =$

Before we prove this lemma
we define inductively a $(2', n)$-
when H has depth T, it sorts n
keep extra information at the nc

The set of inputs at the ro
denotes the trivial order on one el
in the lemma holds for $t = 0$.

Now let $t \in [T]$ and assume
t for inputs from S_*^n, such that
the form $I = I(X, \pi, S)$, and fo
time t. Furthermore we maintair
set for all nodes in depth t.

For each node v in depth t

Let $I = I(X, \pi, S)$ be the se
main Lemma I. Now we define
$\pi' = \bigcup_{i=1}^{n} \pi_i'$ of the sets in $X' = ($
arriving at v' is $I' = (X', \pi', S')$.
for inputs from I', P_i only know
defined sets of inputs arriving at
this construction for all nodes in

Now suppose that we have
that H sorts n numbers, i.e. to p
at a leaf of H the order on $[n]$
inputs from I each P_i only knov
between the different possible to
some inputs. Thus H is a merge

Before we prove main Lemn
theoretic results. For an infinit
$\{(a_1, \cdots, a_n) \in S^n \mid a_i < a_j$ if $i <_Q$
natural order on $[n]$, we write S

We apply the following "c
from [ER] (see also [GRS]).

ts of U belong to a directed path
endent set in $\bar{Q}$.
be partial orders. Q_1 and Q_2 are
$\in \overline{E_1}$ and $(b, a) \in \overline{E_2}$. A family of
sistent if every pair is consistent.
$\cup Q_2 = (V, E_1 \cup E_2)$. Similarly we

(V, E) iff $\bar{E} \subset \bar{E}'$. If Q' is a total

merge of C_1 and C_2 in Q is any

n)-merge tree (for sorting) is a
s. The label of an internal node
is in $Q, (C_{i1}, C_{i2}), i = 1, 2, \cdots, n,$
ry branch e out of v is labeled
and C_{i2}, $i = 1, 2, \cdots, n$. If u is
in u is $Q \cup \cup_{i=1}^{n} C_i$. The root is
an (m, n)-merge tree H, denoted
H sorts n numbers, if each leaf
$n)$-merge trees is $c(m, n)$, the
sorts n numbers.

teps to sort n integers.
orems 1 and 2 below.
ers in T steps. Then there is a

).

, we get a general tight bound
$= \{(a_1, \cdots, a_n) \in S^n \mid a_i \neq a_j$ for
if for each infinite $S \subset N$, f

o compute a symmetric, strongly

al on n variables (e.g. the sum
nction.

rs in T steps. Then there is a

o maintain at each node v of
ne set of inputs arriving at v.
s variables with indices from
act meaning of this is that for
dered as a function in I) only
sequel we shall refer to them

, P_i can now sort the variables
r order. We shall see that this

is essentially the only property of the input which can influence the behaviour of P_i, if we restrict the input set suitably. Namely, if the orders of the sets of variables the processors know are fixed, the next "communication pattern" is fixed, too, in the following sense. For each $i \in [n]$, P_i reads the value that was written by some fixed processor P_{j_i} at some fixed time t_i. Thus it is fixed which new variables P_i gets to know in this step, namely just those which P_{j_i} knew before. As they also were already sorted before, the behaviour of P_i in the next step only depends on the outcome of the merging of the two ordered sets of variables it now knows. Thus M behaves like a merge tree.

In the sequel we only consider inputs which consist of distinct numbers. Recall that for a set $S \subset N$ $S_*^n = \{(a_1, \cdots, a_n) \in S^n \mid a_i \neq a_j \text{ for } i \neq j\}$. For $i \in [n]$ let $X_i \subset [n]$ and π_i be a total order on X_i. Let $X = (X_1, \cdots, X_n)$, $\pi = \cup_{i=1}^{n} \pi_i$. We always assume that the π_i's are consistent. Then, following the above intuition, the set of inputs arriving at a node of the merge tree should be of the form $I(X, \pi, S) = \{(a_1, \cdots, a_n) \in S_*^n \mid a_p < a_q \text{ for all } p, q \text{ such that } p <_\pi q\}$.

The following lemma is the heart of the proof of the theorem.

MAIN LEMMA I. *Suppose that at time t, for inputs from $I = I(X, \pi, S)$, P_i only knows variables from X_i. Then there are $j_1, \cdots, j_n \in [n]$ and an infinite $S' \subset S$, such that after step t, for inputs from $I' = I(X, \pi, S')$, P_i only knows variables from $X_i \cup X_{j_i}$.*

Before we prove this lemma we first conclude Theorem 1 from it. For this purpose we define inductively a $(2^t, n)$-merge tree H of depth t. We shall show that finally, when H has depth T, it sorts n numbers. For the inductive construction to work, we keep extra information at the nodes.

The set of inputs at the root is $N_*^n = I((\{1\}, \cdots \{n\}), (\phi, \cdots, \phi), N)$, where ϕ denotes the trivial order on one element. At this time P_i only knows x_i. So the assumption in the lemma holds for $t = 0$.

Now let $t \in [T]$ and assume that we have constructed a $(2^t, n)$-merge tree of depth t for inputs from S_*^n, such that the set of inputs arriving at a node v in depth t is of the form $I = I(X, \pi, S)$, and for inputs from I, P_i only knows variables from X_i at time t. Furthermore we maintain that although X and π depend on v, S is the same set for all nodes in depth t.

For each node v in depth t successively perform the construction below.

Let $I = I(X, \pi, S)$ be the set of inputs arriving at v, and $S', j_1, \cdots, j_n$ be as in the main Lemma I. Now we define a son v' of v for each consistent tuple of mergings $\pi' = \cup_{i=1}^{n} \pi_i'$ of the sets in $X' = (X_1', \cdots, X_n')$ with $X_i' = X_i \cup X_{j_i}$. Then the set of inputs arriving at v' is $I' = (X', \pi', S')$. As $I' \subset I(X, \pi, S')$, we know by main Lemma I that, for inputs from I', P_i only knows variables from X_i' at time $t+1$. Now for the already defined sets of inputs arriving at nodes replace S by S'. Finally, after having performed this construction for all nodes in depth t, we get a $(2^{t+1}, n)$-merge tree of depth $t+1$.

Now suppose that we have constructed H up to depth T. We finally have to verify that H sorts n numbers, i.e. to prove that for each set of inputs $I = I(X, \pi, S)$ arriving at a leaf of H the order on $[n]$ induced by π is total. Suppose it is not. Then, as for inputs from I each P_i only knows variables from X_i at time T, M cannot distinguish between the different possible total orders and would compute the wrong output for some inputs. Thus H is a merge tree as demanded in the theorem.

Before we prove main Lemma I, we introduce some notations and useful Ramsey theoretic results. For an infinite set $S \subset N$ and a total order Q on $[n]$ let $S_Q^n = \{(a_1, \cdots, a_n) \in S^n \mid a_i < a_j \text{ if } i <_Q j\}$. Such a set is called a fixed order type. If Q is the natural order on $[n]$, we write $S_<^n$ for S_Q^n.

We apply the following "canonical" Ramsey theorem due to Erdös and Rado from [ER] (see also [GRS]).

THEOREM [ER] (Erdös-Rado Theorem). *Let $f: S_Q^n \to N$. Then there is $\tilde{S} \subset S$, $\tilde{S}$ infinite, such that $f' = f|_{\tilde{S}_Q^n}$ is 1-1 on the variables it depends on. Precisely, there is a $J \subset [n]$ such that $f'(a_1, \cdots, a_n) \neq f'(b_1, \cdots, b_n)$ if and only if $a_i \neq b_i$ for some $i \in J$. In particular, if f has a finite range, f' is constant.*

LEMMA 1. *Let $f: S_<^n \to N$ and $g: S_<^{n'} \to N$ be 1-1 functions. Then there is $\tilde{S} \subset S$, $\tilde{S}$ infinite, such that, restricted to $\tilde{S}_<^n$, f and g either have disjoint ranges or are identical. Precisely, either $f(\tilde{S}_<^n) \cap g(\tilde{S}_<^{n'}) = \emptyset$ or $n = n'$ and $f|_{\tilde{S}^n} \equiv g|_{\tilde{S}^n}$.*

Proof. Assume without loss of generality that $n \geqq n'$. Add dummy variables such that also g is defined on $S_<^n$, but only depends on the first n' variables. We first consider the 2-coloring c on $S_<^n$ with $c(\bar{a}) = 1$ if $f(\bar{a}) = g(\bar{a})$, and $c(\bar{a}) = 0$ otherwise. By the Erdös-Rado Theorem there is $\bar{S} \subset S$, $\bar{S}$ infinite, such that c is constant on $\bar{S}_<^n$. If $c \equiv 1$, then $f|_{\bar{S}^n} \equiv g|_{\bar{S}^n}$ and we are done. If $c \equiv 0$, then let G be the directed graph on $\bar{S}_<^n$ with $(\bar{a}, \bar{b}) \in E(G)$ if $f(\bar{a}) = g(\bar{b})$. G has no self-loops because $c \equiv 0$ on $\bar{S}_<^n$. G has indegree 1, because f is 1-1. Therefore it is easy to see that the underlying undirected graph is 3-colorable. Color it with 3 colors. By the Erdös-Rado Theorem there is an infinite $\tilde{S} \subset \bar{S}$ such that $\tilde{S}_<^n$ is monocromatic. Thus, $f(\tilde{S}_<^n) \cap g(\tilde{S}_<^n) = \emptyset$. Q.E.D.

Proof of Main Lemma I. Let $I = I(X, \pi, S)$ be such that P_i only knows variables from X_i at time t. Consider what P_i has done until time t. At each time $d \in [t]$, it wrote some value v_i^d to some cell w_i^d. Furthermore P_i read cell r_i at time t. These values are functions of the input set N^n. But, as P_i only knows variables from X_i, they only depend on the variables from X_i. Our goal now is to restrict the input set I to a set I' in such a way that, for inputs from I', P_i reads what some P_j wrote at some time d, only if r_i and w_j^d are the "same functions", and are applied to the "same arguments".

We refer to the functions w_i^d and r_i as address functions. The *clean form* of such a function f is derived by fixing all variables f does not depend on to arbitrary constants. Thus a clean form of a function always depends on all its variables. Let f' be the clean form of some address function f. If f was used by P_i (i.e. is r_i or w_i^d for some d), then it only depends on some (not necessary all) variables from X_i. As they are totally ordered according to π_i, f' is defined on a set of fixed order type. Thus we may apply the Erdös-Rado theorem successively to the clean forms of all address functions. The result is an infinite set $\tilde{S} \subset S$ such that, on $\tilde{I} = I \cap \tilde{S}^n$, all address functions are 1-1 on the variables they depend on.

From now on we assume that the domain of all address functions is $\tilde{I}$. Note that the clean form of the address functions can now depend on fewer variables than before we applied the Erdös-Rado theorem.

We know that the clean form of an address function is defined on a set of fixed order type. The function derived from it by reordering these variables such that it is defined on $\tilde{S}_<^n$ is called the *standard form* of f.

As we now know that the standard forms of address functions are 1-1 functions and are defined on $S_<^n$, we may apply Lemma 1 successively to all pairs of them. As a result we get an infinite set $\bar{S} \subset \tilde{S}$ with the following property (*). Assume from now on that the domain of all address functions is $I' = \tilde{I} \cap \bar{S}^n$.

(*) *Two address functions either have disjoint ranges or have the same standard form.*

Now let f and g be two address functions. Recall that they are defined on I'.

LEMMA 2. *f and g are either identical or $f(\bar{a}) \neq g(\bar{a})$ for all $\bar{a} \in I'$.*

Proof. Suppose $f(\bar{a}) = g(\bar{a})$ for some $\bar{a} \in I'$. Then, by (*), they have the same standard form h. Thus they are identical if they depend on the same variables. Assume that f and g depend on different sets of variables. Let $\bar{a}'$ and $\bar{a}''$ be the subvectors of the above $\bar{a}$ of those variables f and g depend on, increasingly ordered. These vectors

are different, because they contain
On the other hand, we know that
$h(\bar{a}'')$, which contradicts the supp

Now consider what P_i reads
for all j and d), then we know by
PRIORITY write conflict resoluti
time d, where (d, j) is lexicograp
depends on variables from X_j, aft
inputs from I', which proves the

4. A lower bound for sorting

of our main theorem by showing

THEOREM 2.

$$c(m$$

Proof of Theorem 2. It is s
$m \geqq 9 \log n$. Let H be any (m, n)-
exhibit a long path from the root
that we constructed the last nod
linear extensions. We shall look
we must show that there is a wa
few as possible transitive implic

To this end, we shall define
a handle on the quantities menti
path we construct, its partial or
nice class.

We now describe the "nice"
l, a. A (k, l, a)-*graph* is obtained
some k of the vertices, and repl
in and out neighbours).

A (k, l, a)-*partial order* is a
closure $\bar{G}$ of some (k, l, a) grap

Let $Q = (V, E)$ be a (k, l, a)
them $S_1, S_2, \cdots, S_k$. Clearly, to
S_i's. Consider a merge operation
Clearly, we may assume that C
those antichains is known). Mor
and $i = 1, 2, \cdots, k$. Hence, me
comparisons, at most one in eac

MAIN LEMMA II. *Let v be
(k, l, a)-partial order, with $a \leqq n$
partial order Q_u is contained in
$a + (3n/m) + km^4$.*

Before we prove main Lem
the empty partial order at the r
t times with $t = \log n/5 \log m$. V
(k, l, a)-partial order with $k \leqq n$
and because of the choice of t,
for this choice of t, $k \leqq n/4$, $a \leqq$

are different, because they contain values of different variables, and the a_i's are distinct. On the other hand, we know that $h(\bar{a}') = f(\bar{a})$, $h(\bar{a}'') = g(\bar{a})$. But as h is 1-1, $h(\bar{a}') \neq h(\bar{a}'')$, which contradicts the supposition we started with. Q.E.D.

Now consider what P_i reads in step t. If r_i was never used for writing (i.e. $r_i \neq w_j^d$ for all j and d), then we know by Lemma 2 that P_i reads 0. Otherwise, because of the PRIORITY write conflict resolution and Lemma 2, P_i reads v_j^d, the value P_j wrote at time d, where (d, j) is lexicographically maximal such that $w_j^d \equiv r_i$. Thus, as v_j^d only depends on variables from X_j, after this step P_i only knows variables from $X_i \cup X_j$ for inputs from I', which proves the lemma. Q.E.D.

4. A lower bound for sorting on merge trees. In this section we complete the proof of our main theorem by showing the demanded lower bound on (m, n)-merge trees.

THEOREM 2.

$$c(m, n) = \Omega\left(\frac{\log n}{\log (m \log n)}\right).$$

Proof of Theorem 2. It is sufficient to prove that $c(m, n) = \Omega(\log n / \log m)$ for $m \geq 9 \log n$. Let H be any (m, n)-merge tree which sorts n numbers. We shall inductively exhibit a long path from the root in H. Intuitively, the argument is as follows. Suppose that we constructed the last node in the path, v, s.t. the partial order in v has many linear extensions. We shall look for a child u of v with the same property. To do that, we must show that there is a way of merging the n pairs of chains given at v, s.t. as few as possible transitive implications are added to the partial order.

To this end, we shall define a "nice" class of partial orders, in which we can get a handle on the quantities mentioned above. We will show that for each node in the path we construct, its partial order has an extension which is a partial order in the nice class.

We now describe the "nice" partial orders. Let $n = kl + a$ for positive integers k, l, a. A (k, l, a)-*graph* is obtained as follows. Take a chain on $k + a$ vertices. Then choose some k of the vertices, and replace each by l copies of itself (each having the same in and out neighbours).

A (k, l, a)-*partial order* is a partial order Q s.t. $\bar{Q}$ is isomorphic to the transitive closure $\bar{G}$ of some (k, l, a) graph G.

Let $Q = (V, E)$ be a (k, l, a)-partial order. It has k antichains, each of size l. Call them $S_1, S_2, \cdots, S_k$. Clearly, to "sort" Q, it is necessary and sufficient to sort all the S_i's. Consider a merge operation on Q. It involves two chains, C_1 and $C_2 (|C_1|, |C_2| \leq m)$. Clearly, we may assume that $C_1, C_2 \subset \bigcup_{i=1}^{k} S_i$ (as the rank of the elements outside those antichains is known). Moreover, by the structure of Q, $|C_j \cap S_i| \leq 1$ for all $j = 1, 2$ and $i = 1, 2, \cdots, k$. Hence, merging C_1 and C_2 reduces to performing at most m comparisons, at most one in each S_i.

MAIN LEMMA II. *Let v be a node in H whose partial order Q_v is contained in a (k, l, a)-partial order, with $a \leq n/2$ (hence $kl \geq n/2$). Then there is a child u of v whose partial order Q_u is contained in an (k', l', a') partial order, with $k' \leq km^4$ and $a' \leq a + (3n/m) + km^4$.*

Before we prove main Lemma II, let us see how it implies Theorem 2. Clearly, the empty partial order at the root of H is a $(1, n, 0)$ partial order. Apply the lemma t times with $t = \log n / 5 \log m$. We reach a node whose partial order is contained in a (k, l, a)-partial order with $k \leq m^{4t}$ and $a \leq 3tn/m + m^{4t}$. (Note that since $m \geq 9 \log n$, and because of the choice of t, the assumption in the lemma holds at every step.) But for this choice of t, $k \leq n/4$, $a \leq n/2$, and therefore $l \geq 2$. So this partial order cannot

be a total order, and this node at depth t cannot be a leaf. Hence, $c(H) = \Omega(\log n/\log m)$.

Before we proceed to prove Lemma 2, we need the following technical lemma.

LEMMA 3. *Let $G=(V, E)$ be an undirected graph, and b a positive number. Then one can remove a set $V' \subset V$ of vertices, with $|V'| \leqq (2|V|/b)+(2|E|b/|V|)$, s.t. the remaining graph on $V-V'$ can be colored with $(2|E|b/|V|)$ colors, with each color class of the same size.*

This lemma is based on the following result, due to Hajnal and Szemeredi from [HS].

THEOREM [HS]. *Any graph $G=(V, E)$ with maximum degree Δ can be colored with $\Delta+1$ colors s.t. each color class has size $(|V|/\Delta+1)$ or $(|V|/\Delta+1)+1$.*

Proof of Lemma 3. Remove the $2|V|/b$ vertices of highest degree from G. Then the maximum degree is at most $(2|E|b/|V|)-1$. Apply the previous theorem, and then remove one vertex from each of the larger color classes (at most $2|E|b/|V|$), to make them all of equal size.

Proof of Main Lemma II. Assume without loss of generality that Q_v is a (k, l, a)-partial order, and let $S_1, S_2, \cdots, S_k$ be its antichains of size l. Consider the n pairs of chains, (C_{i_1}, C_{i_2}), $i=1, 2, \cdots, n$, that are merged at v. Each such pair gives rise to at most m comparisons. So we have a total of nm comparisons, distributed among the S_j's. By averaging, at most k/m of the S_j's have more than nm^2/k comparisons. For each such S_j, arbitrarily choose a total order (hence resolving all comparisons within it). We obtain an extension Q_1 of Q_v which is a (k_1, l, a_1)-partial order, with $k_1 \leqq k$ and $a_1 \leqq a+(k/m)l \leqq a+(n/m)$.

So each of the remaining S_j's has at most nm^2/k comparisons. Think of these comparisons as edges in an undirected graph with vertex set S_j. Now we use Lemma 3 on each of these graphs. From each S_j remove $\leqq 2l/m+(2(nm^2/k)m/l)$ vertices as in the lemma, make them all (say) bigger than the rest of S_j, and arbitrarily totally order them. This resolves all comparisons involving these vertices. The result is an extension Q_2 of Q_1, which is a (k_2, l, a_2) partial order, with $k_2=k_1 \leqq k$ and $a_2 \leqq a_1+k((2l/m)+(2nm^3/kl)) \leqq a_1+2n/m+4km^3 \leqq a+3n/m+km^4$.

Finally, the remaining graph on each S_j can be colored with $\leqq 2nm^3/kl \leqq 4m^3 \leqq m^4$ colors each color class of equal size. Totally order the *color classes* arbitrarily, thus resolving the remaining comparisons in each S_j. The result is a (k', l', a') partial order Q', with $k' \leqq km^4$ and $a' \leqq a+3n/m+km^4$. Since we resolved all comparisons in a consistent way, there must be a child u of v s.t. Q' is an extension of Q_u.

5. A general lower bound for a family of symmetric functions. In this section we prove Theorem 3. Recall that for $S \subset N$ we defined $S^n_* = \{(a_1, \cdots, a_n) \in S^n \mid a_i \neq a_j \text{ for } i \neq j\}$. A function $f: N^n \to N$ is *strongly nonconstant*, if for each infinite $S \subset N$, f restricted to S^n_* is nonconstant.

THEOREM 3. *A* PRAM *needs exactly $\log_2 n$ steps to compute a symmetric, strongly nonconstant function.*

Proof. The upper bound is obvious because of the computational power of the processors: in $\log_2 n$ steps one processor can get to know all variables and then compute the function in one step.

In order to show the lower bound, we shall again apply the Erdös–Rado theorem to find an infinite $S \subset N$ such that M is oblivious to inputs from $S^n_<$.

We first describe what we mean by oblivious. Let $\bar{x} \in N^n$ be an input for M. Then it is well defined which processor writes to or reads from which memory cell at a given time, if M is started with $\bar{x}$. Thus we can define the following *communication pattern*

for $\bar{x}$: For each $t \in [T]$, $i \in [n]$, P_i
$d \in [t]$ and $j \in [n]$ depend only o
the communication patterns of all
is often proved in similar forms i

LEMMA 4. *Let $I \subset N^n$ such th*
M computing f is oblivious to in
compute f.

LEMMA 5. *If M computes f:*
oblivious to inputs from $S^n_<$.

The lower bound in Theorem
the definition of strongly noncon
functions, if such a function is no

Proof of Lemma 5. We restri
and the number of steps M exec
communication patterns. Thus, by
such that all inputs from $S^n_<$ have
for inputs from $S^n_<$. Q.E.D.

[A] M. AJTAI, Σ^1_1-*formulae on fini*

[AKS] M. AJTAI, J. KOMLOS AND
(1983), pp. 1-19.

[CDR] S. COOK, C. DWORK AND
access machines without sim

[ER] P. ERDÖS AND R. RADO, *A c*

[FMRW] F. FICH, F. MEYER AUF
three $\cdots$ infinity: Lower bou
to appear.

[GRS] R. L. GRAHAM, B. L. ROTH
York, 1980.

[HS] A. HAJNAL AND E. SZEMER
its Applications, 2 (1970),

[IM] A. ISRAELI AND S. MORAN

[L] T. LEIGHTON, *Tight bounds*
DC, 1984, pp. 71-80.

[MR] F. MEYER AUF DER HEIDE
large hardware and unboun

[P] I. PARBERRY, *A complexity*

[PS] W. J. PAUL AND J. SIMON, *L*
und Algorithmik, Zürich,

[V] L. VALIANT, *Parallelism in c*

for $\bar{x}$: For each $t \in [T]$, $i \in [n]$, P_i reads at time t what P_j has written at time d, where $d \in [t]$ and $j \in [n]$ depend only on i and t. M is oblivious to inputs from $I \subset N^n$, if the communication patterns of all inputs from I are the same. The following lemma is often proved in similar forms in literature, e.g. in [MR], [CDR].

LEMMA 4. *Let $I \subset N^n$ such that $f: I \to N$ depends on all its variables. If the* PRAM *M computing f is oblivious to inputs from I, then M needs at least $\log_2 n$ steps to compute f.*

LEMMA 5. *If M computes $f: N^n \to N$ then there is an infinite $S \subset N$, such that M is oblivious to inputs from $S^n_<$.*

The lower bound in Theorem 3 now follows directly from the above lemmas and the definition of strongly nonconstant functions. As we only deal with symmetric functions, if such a function is nonconstant on S^n_*, then it is so even on $S^n_<$.

Proof of Lemma 5. We restrict f to inputs from $N^n_<$. As the number of processors and the number of steps M executes are finite, there are only finitely many different communication patterns. Thus, by the Erdös-Rado theorem, we find an infinite $S \subset N$ such that all inputs from $S^n_<$ have the same communication pattern, i.e. M is oblivious for inputs from $S^n_<$. Q.E.D.

REFERENCES

[A] M. AJTAI, *Σ^1_1-formulae on finite structures*, Ann. Pure Appl. Logic, 24 (1983), pp. 1-48.

[AKS] M. AJTAI, J. KOMLOS AND E. SZEMEREDI, *Sorting in c log n parallel steps*, Combinatorica, 3 (1983), pp. 1-19.

[CDR] S. COOK, C. DWORK AND R. REISCHUK, *Upper and lower time bounds for parallel random access machines without simultaneous writes*, preprint, 1983.

[ER] P. ERDÖS AND R. RADO, *A combinatorial theorem*, J. London Math Soc., 25 (1950), pp. 249-255.

[FMRW] F. FICH, F. MEYER AUF DER HEIDE, P. RAGDE AND A. WIGDERSON, *One, two, three $\cdots$ infinity: Lower bounds for parallel computation*, 16th ACM STOC, Providence, 1985, to appear.

[GRS] R. L. GRAHAM, B. L. ROTHSCHILD AND J. H. SPENCER, *Ramsey Theory*, John Wiley, New York, 1980.

[HS] A. HAJNAL AND E. SZEMEREDI, *Proof of a conjecture of P. Erdös*, Combinatorial Theory and its Applications, 2 (1970), pp. 601-623.

[IM] A. ISRAELI AND S. MORAN, private communication.

[L] T. LEIGHTON, *Tight bounds on the complexity of parallel sorting*, 16th ACM STOC, Washington DC, 1984, pp. 71-80.

[MR] F. MEYER AUF DER HEIDE AND R. REISCHUK, *On the limits to speed up parallel machines by large hardware and unbounded communication*, 25th IEEE FOCS, Miami, 1984, pp. 56-64.

[P] I. PARBERRY, *A complexity theory of parallel computation*, Ph.D. dissertation, Warwick, 1984.

[PS] W. J. PAUL AND J. SIMON, *Decision trees and random access machines*, Symposium ueber Logik und Algorithmik, Zürich, 1980, pp. 331-339.

[V] L. VALIANT, *Parallelism in comparison problems*, this Journal, 4 (1975), pp. 348-355.