



Basic Research in Computer Science

BRICS RS-95-10

Breslauer & Dubhashi: Transf. Comparison Model Lower Bounds to the PRAM

Transforming Comparison Model Lower Bounds to the PRAM

Dany Breslauer
Devdatt P. Dubhashi

BRICS Report Series

RS-95-10

ISSN 0909-0878

February 1995

**Copyright © 1995, BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

**See back inner page for a list of recent publications in the BRICS
Report Series. Copies may be obtained by contacting:**

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK - 8000 Aarhus C
Denmark
Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

**BRICS publications are in general accessible through WWW and
anonymous FTP:**

**`http://www.brics.dk/
ftp ftp.brics.dk (cd pub/BRICS)`**

Transforming Comparison Model Lower Bounds to the Parallel-Random-Access-Machine*

Dany Breslauer Devdatt P. Dubhashi

BRICS[†]
Department of Computer Science
University of Aarhus
DK-8000 Aarhus C, Denmark

Abstract

This note provides general transformations of lower bounds in Valiant’s parallel comparison decision tree model to lower bounds in the priority concurrent-read concurrent-write parallel-random-access-machine model. The proofs rely on standard Ramsey–theoretic arguments that simplify the structure of the computation by restricting the input domain. The transformation of comparison model lower bounds, which are usually easier to obtain, to the parallel-random-access-machine, unifies some known lower bounds and gives new lower bounds for several problems.

1 Introduction

Valiant’s parallel comparison decision tree model [24] is very attractive for studying parallel algorithms and lower bounds for *order invariant problems* whose solution depends on equality or order relations between the input variables. One of the major drawbacks of this model, however, is that the information obtained by determining the relations between input variables becomes “common knowledge” and the model fails to capture the difficulty in communicating information between various processing units that run in parallel. The more realistic parallel-random-access-machine model (PRAM) captures some issues of communication between the different processing units, what makes it a more natural model to describe parallel algorithms.

The two models, however, are not comparable in general. While some problems in the comparison model, e.g. finding the maximum [24], have similar algorithms in the PRAM model [23], other problems, e.g. finding the median [1, 5], have only much slower PRAM algorithms [7] with any polynomial number of processors. On the other hand, there exist problems, e.g. element distinctness, that have slow comparison model algorithms [6] and constant-time PRAM algorithms on integer input domains.

[†]Basic Research in Computer Science, Centre of the Danish National Research Foundation.

*The research reported in this paper was partially supported by ESPRIT Basic Research Action Program of the EC under contract #7141 (ALCOM II).

Since comparison model lower bounds are often easier to obtain than PRAM lower bounds, it can sometimes be useful to translate comparison model lower bounds into PRAM lower bounds. Clearly, if a PRAM algorithm can only access its input by determining the relations between the input variables, then the comparison model lower bounds will hold for the PRAM. However, this assumption prevents PRAM algorithms from using their powerful capabilities. Moreover, solutions to problems that are defined in the equality-comparison model sometimes benefit from the introduction of an arbitrary order on the input domain so that order comparisons can be used; e.g. element distinctness. Since the input variables on the PRAM are usually assumed to be integers, the input domain is naturally ordered. This makes lower bounds in the equality-comparison model inapplicable to the PRAM model.

This note gives two general translations of lower bounds in the order-comparison model to lower bounds in the priority CRCW-PRAM model:

1. Any comparison model lower bound can be converted into a corresponding lower bound in the priority CRCW-PRAM with bounded memory. By bounded memory we mean that the memory size is not permitted to grow as a function of the input domain size.
2. Any comparison model lower bound that holds if the input variables are known to be all *distinct* can be converted into a corresponding lower bound in the priority CRCW-PRAM with infinite memory.

The proof techniques used are standard multi-variable Ramsey–theoretic arguments that were developed by several authors for studying specific problems [9, 19]. The main idea is that one can restrict the original input domain in such a way that the processors must communicate in a manner that depends only on the relative order between the input variables. This implies that the PRAM can only determine the relations between input variables that were communicated to a given processor and not by the communication pattern itself. We then apply comparison model lower bounds to obtain lower bounds on the PRAM.

The transformation of the comparison model lower bounds provides a unified way to obtain lower bounds for the PRAM. It generalises previous results and provides PRAM lower bounds for problems that had only comparison model lower bounds. Some of the lower bounds obtainable by the transformation, for input of size n on a p processor PRAM, are:

1. Sorting requires $\Omega(n/p + \log_{\lceil (p/n) \log_{n+1} \rceil} n)$ time [9, 19].
2. Element distinctness requires $\Omega(n/p + \log_{\lceil (p/n) \log_{n+1} \rceil} n)$ time if the memory size is bounded [9, 14, 21].
3. Finding the maximum and merging require $\Omega(n/p + \log \log_{\lceil p/n+1 \rceil} n)$ time [17, 22].
4. String matching and some related problems on strings require, assuming bounded memory size, $\Omega(n/p + \log \log_{\lceil p/n+1 \rceil} n)$ time.

5. Finding an approximate maximum, namely, an element whose rank belongs in the top ϵn ranks, requires $\Omega(n/p + \log \log_{\lceil p/n+1 \rceil} (1/\epsilon) + \log^* n - \log^*(p/n))$ time¹, for $1/n \leq \epsilon \leq 1/2$.

The paper is organised as follows. §2 and §3 review Valiant’s parallel comparison-decision-tree and the parallel-random-access-machine models. §4 gives the general PRAM lower bounds and §5 shows how these lower bounds are applied to specific problems. Conclusions and open problems are given in §6.

2 Parallel comparison models

The input variables x_1, \dots, x_n are chosen from some infinite totally ordered domain \mathcal{D} . Denote by \mathcal{D}^k the set of all k -tuples of elements of \mathcal{D} , by \mathcal{D}_*^k the set of all k -tuples of \mathcal{D} with no two equal elements and by $\mathcal{D}_{<}^k$ the set of all increasing k -tuples of \mathcal{D} .

Following the notation of [20], two tuples x_1, \dots, x_k and y_1, \dots, y_k are said to be *order equivalent* if $x_i < x_j \Leftrightarrow y_i < y_j$, for all $i, j = 1, \dots, k$. (And hence $x_i = x_j \Leftrightarrow y_i = y_j$.) The equivalence class containing x_1, \dots, x_k is called the *order type* of x_1, \dots, x_k . A decision problem \mathcal{P} on the variables x_1, \dots, x_n partitions the inputs from \mathcal{D}^n into classes $\mathcal{P}_1, \dots, \mathcal{P}_q$. \mathcal{P} is said to be *order invariant* if order equivalent tuples are always in the same class.

A comparison between two variables $x_i :: x_j$ determines if $x_i < x_j$, $x_i = x_j$ or $x_i > x_j$. We will also consider sometimes the additional relations $x_i \neq x_j$, without order information. These inequality relations, which may not be established by comparisons, might instead be given *a priori*, as part of the definition of a problem. Such *a priori* restrictions are useful for problems that are defined on partial domains; e.g. in the merging problem the two lists to be merged are assumed to be sorted. Restricting the input so that all variables are distinct, i.e. $x_i \neq x_j$, for $i \neq j$, will be of particular interest in this paper. We refer to this restriction as the *distinctness assumption*.

An order invariant problem $\mathcal{P}(x_1, \dots, x_n)$ can be solved by comparing pairs of input variables until all input tuples satisfying the relations that were established are in the same class \mathcal{P}_i ; e.g. in the problem of finding the maximum it suffices to discover that some $x_i \geq x_j$, for $j = 1, \dots, n$, without caring about the relative order between the other variables. Clearly, sorting is the hardest problem in this sense since it determines the exact relations between all input variables and thus the order type of the input.

Valiant’s parallel comparison decision tree model [24] proceeds in rounds in which up to p pairwise comparisons of input variables are made simultaneously. According to the outcome of the comparisons, and the relations established in previous rounds, the comparison model algorithm decides which variables to compare in the next round, or it may decide to terminate with an answer. We denote by $\mathcal{C}_{\mathcal{P}}(n, p)$ the depth of the shallowest comparison decision tree that

¹Define $\log^{(0)} n = n$, $\log^{(i)} n = \log \log^{(i-1)} n$ and $\log^* n = \min\{i \mid \log^{(i)} n \leq 1\}$. In this paper $\log n = \max\{0, \log_2 n\}$.

solves the problem $\mathcal{P}(x_1, \dots, x_n)$ using p comparisons in each round. Clearly, comparing all $\binom{n}{2}$ pairs of variables gives complete information about their order type, and thus, for any problem \mathcal{P} , $\mathcal{C}_{\mathcal{P}}(n, \binom{n}{2}) \leq 1$.

Boppana [9], following Meyer auf der Heide and Wigderson [19], defines a similar comparison decision tree model that we call the *merging-comparison decision tree*. In the p processor merging-comparison model, each processor knows a certain subset of the input variables and their order type (initially these sets are empty). In every round, according to the partial order that is formed by the order types of the subsets of variables known by all processors, the merging-comparison model decides whether to terminate with an answer or to continue, letting each processor to *merge* its set of variable either with the set of variables known by some other processor at the end of the previous round or with a single input variable. We denote by $\mathcal{M}_{\mathcal{P}}(n, p)$ the depth of the shallowest p processor merging-comparison decision tree for the problem $\mathcal{P}(x_1, \dots, x_n)$. The following lemma relates lower bounds in the parallel comparison model to the merging-comparison model.

Lemma 2.1 *Let*

$$\hat{\mathcal{C}}_{\mathcal{P}}(n, p) = \max \left\{ t \mid \mathcal{C}_{\mathcal{P}}(n, 2^{2^{(t-1)}}p) \geq t \right\}.$$

Then, $2 \cdot \mathcal{C}_{\mathcal{P}}(n, p) \geq \mathcal{M}_{\mathcal{P}}(n, p) \geq \hat{\mathcal{C}}_{\mathcal{P}}(n, p)$.

Proof: Clearly, every comparison model round with p comparison can be simulated by at most two rounds of a p processor merging-comparison model, establishing that $2 \cdot \mathcal{C}_{\mathcal{P}}(n, p) \geq \mathcal{M}_{\mathcal{P}}(n, p)$. Inductively, the number of variables known by each processor after h rounds is at most 2^{h-1} . The relations established by merging two sets of at most 2^{h-1} variables can be determined by performing $2^{2^{(h-1)}}$ comparisons. Hence, a p processor merging-comparison model can be simulated by a comparison model that makes $2^{2^{(h-1)}}p$ comparisons in round number $h + 1$. If there are at most t rounds, this can be overestimated by $2^{2^{(t-1)}}p$ comparisons in each round. If $\mathcal{C}_{\mathcal{P}}(n, 2^{2^{(t-1)}}p) \geq t$, then even with this larger number of comparisons the solution of the problem \mathcal{P} require at least t rounds. \square

We say that a comparison model lower bound $\mathcal{C}_{\mathcal{P}}(n, p)$ is *resilient* if $\alpha \cdot \hat{\mathcal{C}}_{\mathcal{P}}(n, p) + \beta \geq \mathcal{C}_{\mathcal{P}}(n, p)$ for some constants α and β and all $n, p \geq 1$. When this holds, we write $\mathcal{C}_{\mathcal{P}}(n, p) = \Theta(\hat{\mathcal{C}}_{\mathcal{P}}(n, p))$. Resilient comparison model lower bounds translate to the same lower bounds in the merging-comparison model (up to constants).

3 The Parallel-Random-Access-Machine

In this paper we consider a powerful version of the *priority concurrent-read concurrent-write parallel-random-access-machine (CRCW-PRAM)*. The model consists of p synchronous processors that communicate via a shared memory with cells of *unlimited size*. Processors are allowed to read and write simultaneously at the same memory location; write conflicts are resolved by accepting the value that is written by the processor with the highest preassigned priority.

We assume that the execution of a PRAM program proceeds in rounds. Each round consists of a computation phase in which every processor can make *any* computation on the information it has obtained before, followed by a write phase and then by a read phase. Note that these assumptions result in an extremely powerful model that can compute *any function* in $O(n)$ steps using one processor or in $O(\log n)$ steps using $n/\log n$ processors. Hence, lower bounds in this model emphasise the limits of the interprocessor communication. We say that the PRAM solves an order invariant problem \mathcal{P} , if for each pair of order types in different \mathcal{P} -equivalence classes, there exists at least one processor that is able to distinguish between them.

Let $R_{i,t}$, for $i \in [p]$, denote the *read access* function of processor number i at round t ; i.e. at round t , processor i reads the memory cell that whose address is given by $R_{i,t}$. Similarly, let $W_{i,t}$, for $i \in [p]$, denote the *write access* function of processor i at round t ; i.e. at round t , processor i writes the value $X_{i,t}$ into the memory cell whose address is $W_{i,t}$. $R_{i,t}$, $W_{i,t}$ and $X_{i,t}$ are functions of the *state* of processor i at round t . The lower bound argument given next show that by restricting the input domain, it is possible to simplify the interaction between processors so that the state of each processor depends only on the input variables it “knows”.

4 Lower bounds

In this section we show that, under certain assumption, a PRAM algorithm for an order invariant problem can be simulated by a merging-comparison decision tree on some restricted input domain. This allows us to transform lower bounds from the merging-comparison model to lower bounds in the PRAM model. The arguments are essentially the same as those used by Meyer auf der Heide and Wigderson [19] and Boppana [9]; we observe that these arguments are more generally applicable than to the problems considered in those papers and even when input variables are allowed to be equal. In §4.1 we summarize the Ramsey-theoretic components of the proofs. In §4.2 we give the lower bounds for PRAM algorithms with bounded memory and in §4.3 for PRAM algorithms with unbounded memory, under the distinctness assumption.

4.1 Ramsey theory

Let \mathcal{D} be an infinite totally ordered set. We say that a function f is a \mathcal{D} *fixed order type* function if f is defined on tuples from \mathcal{D}^k of a fixed order type, for some $k \geq 0$. The *standard form* of f is obtained by removing all variables that f does not depend on and ordering the remaining variables in increasing order according to their order type. More precisely, the standard form of f is obtained by removing all but one representative of equal variables (since the domain of f has a fixed order type, equal variables are equal on all the domain); removing all variables that f does not depend on; and reordering the remaining variables in increasing order. Hence the standard form of f is defined on the domain $\mathcal{D}^l_{<}$, for some $l \leq k$.

Given a function f that is defined on some subset of \mathcal{D}^k , we denote by $f|_{\mathcal{E}}$, for $\mathcal{E} \subseteq \mathcal{D}$, the restriction of f to \mathcal{E}^k . Similarly, if \mathcal{F} is a family of functions, we use the notation $\mathcal{F}|_{\mathcal{E}} := \{f|_{\mathcal{E}} \mid f \in \mathcal{F}\}$. The following lemma can be derived from the ‘‘canonical’’ Ramsey theorem due to Erdős and Rado [15]. (See also §5.5 in [18].)

Lemma 4.1 *Let \mathcal{F} be a finite collection of \mathcal{D} fixed order type functions. Then, there exists an infinite subset $\mathcal{E} \subseteq \mathcal{D}$, such that the standard form of every function in $\mathcal{F}|_{\mathcal{E}}$ is injective and every pair of functions in $\mathcal{F}|_{\mathcal{E}}$ either have identical standard forms or disjoint ranges. In particular, if the range of the functions in \mathcal{F} is finite, then the functions in $\mathcal{F}|_{\mathcal{E}}$ are constant.*

4.2 Bounded memory

In this section we assume that the memory size m is fixed. Hence all read and write access functions have the finite range $[m]$. We will construct a merging-comparison decision tree that will simulate the computation of the PRAM.

The construction proceeds step by step. We will maintain the set of variable indices $V_{i,t} \subseteq [n]$ known by processor $i \in [p]$ at round t and an infinite set $S_t \subseteq \mathbb{N}$, such that the input tuples will be restricted to S_t^n . Initially $V_{i,0} = \emptyset$ and $S_0 = \mathbb{N}$, and for $t \geq 1$, $V_{i,t-1} \subseteq V_{i,t}$ and $S_t \subseteq S_{t-1}$, for $i \in [p]$. Also, we shall maintain that for inputs in S_t , the state of the processor i in the original PRAM, and therefore $R_{i,t}$, $W_{i,t}$ and $X_{i,t}$, are functions of the variables whose indices are in $V_{i,t}$.

Suppose by induction that we have described the behavior of the merging-comparison model up to round t . Thus, at the beginning of round t , each processor $i \in [p]$ knows the variables whose indices are in $V_{i,t}$ and their order type. Consider the collection of access functions $\mathcal{F}|_{S_t}$ consisting of the access functions $R_{i,t}$ and $W_{i,t'}$, for $i \in [p]$ and $0 \leq t' \leq t$. Since we are at a specific node of the merging-comparison decision tree in which the order types of $V_{i,t}$ are fixed, these access functions are S_t fixed order type functions. By Lemma 4.1, there exists an infinite subset $S_{t+1} \subseteq S_t$, such that the access functions in $\mathcal{F}|_{S_{t+1}}$ are constant. (If a processor does not write, it does not write on all the restricted inputs.)

Let $c = R_{i,t}|_{S_{t+1}}$ be a read access function and let $c = W_{i',t'}|_{S_{t+1}}$ be the write access function (provided it exists), that corresponds to what $R_{i,t}$ actually reads (lexicographically maximal $\langle t', i' \rangle$, corresponding to the most recent write by highest priority processor). If such $W_{i',t'}$ exists, then processor i reads the value $X_{i',t'}$ which is also a function of $V_{i',t'}$; we can assume that it actually reads $V_{i',t'} \subseteq V_{i',t}$ and computes $X_{i',t'}$ by itself. Define $V'_{i,t}$ to be $V_{i',t}$ if such $W_{i',t'}$ exists, otherwise $V'_{i,t} = \{c'\}$ if memory cell number c initially contains the input variable $x_{c'}$. If neither of these conditions hold, then $V'_{i,t} := \emptyset$. The merging-comparison model merges in round t the set of variables $V_{i,t}$ known by processor $i \in [p]$ with $V'_{i,t}$. This completely defines the behavior of the merging-comparison model in round t . Since the PRAM interprocessor communication pattern is fixed, the state of the processors can depend only on the variables they know.

Suppose the simulating merging-comparison algorithm cannot solve the problem \mathcal{P} in T rounds. Then there are two order types in different \mathcal{P} -equivalence classes that it cannot distinguish from each other. Since the domain S_T is large enough, for each processor i , there are input tuples of these two order types that agree on the variables in $V_{i,T}$. But by the invariant maintained, for inputs in S_T , the state of each processor in the original PRAM can only depend on the variables it knows, namely $V_{i,T}$. Hence, the original PRAM cannot distinguish between the inputs either. We conclude:

Theorem 4.2 *If a p processor merging-comparison model requires $\mathcal{M}_{\mathcal{P}}(n, p)$ rounds to solve the problem $\mathcal{P}(x_1, \dots, x_n)$, then a bounded memory PRAM must take $\mathcal{M}_{\mathcal{P}}(n, p)$ time as well.*

4.3 Unbounded memory, distinct inputs

In this section we assume that the input variables are distinct. Similarly to the previous section, we will construct inductively a merging-comparison decision tree that will simulate the computation of the PRAM.

Suppose by induction that we have described the behavior of the merging-comparison model up to round t and let $\mathcal{F}|_{S_t}$ be the collection of the access functions $R_{i,t}$ and $W_{i,t'}$, for $i \in [p]$ and $0 \leq t' \leq t$. Since the order types of $V_{i,t}$ are fixed, these access functions are S_t fixed order type functions and by Lemma 4.1, there exists an infinite subset $S_{t+1} \subseteq S_t$, such that the standard forms of the access functions in $\mathcal{F}|_{S_{t+1}}$ are injective and that every pair of functions in $\mathcal{F}|_{S_{t+1}}$ either have identical standard forms or disjoint ranges.

Since the input variables are distinct, identical standard forms of access functions in $\mathcal{F}|_{S_{t+1}}$ are either functions of exactly the same variables, in which case they are *always* equal, or have *always* disjoint ranges. Hence, given the standard form of each read access function $R_{i,t}|_{S_{t+1}}$ one can determine precisely the unique write access function $W_{i',t'}|_{S_{t+1}}$ that corresponds to what $R_{i,t}$ actually reads. As in the previous section, this defines $V_{i,t}'$ and therefore, the complete behavior of the merging-comparison decision tree. This leads to the following theorem:

Theorem 4.3 *If the p processor merging-comparison model requires $\mathcal{M}_{\mathcal{P}}(n, p)$ rounds to solve the problem $\mathcal{P}(x_1, \dots, x_n)$ under the distinctness assumption, then a PRAM must take $\mathcal{M}_{\mathcal{P}}(n, p)$ time as well.*

5 Applications

We give some applications of the general lower bounds from the previous section. Some of these bounds were known previously, but we present them in a unified fashion. We also give some new bounds.

In our discussion below we assume that the input has size n and the number of processors p satisfies $1 \leq p \leq \binom{n}{2}$. Since all the problems we consider depend on most of their input variables, any PRAM algorithm for these problems must read $\Omega(n)$ variables and thus must take at least $\Omega(n/p)$ time.

5.1 Element distinctness and sorting

The $\Omega(\log_{\lceil p/n+1 \rceil} n)$ comparison model lower bound for sorting [6] gives only an $\Omega(\sqrt{\log n})$ lower bound in the n processor merging-comparison model. However, there are better direct lower bounds for sorting in the merging-comparison model.

Lemma 5.1 (*Boppana [9]*) *Sorting in the p processors merging-comparison model requires $\Omega(\log_{\lceil (p/n)\log n+1 \rceil} n)$ rounds.*

Since the merging-comparison model lower bound holds also under the distinctness assumption, it translates to an $\Omega(\log_{\lceil (p/n)\log n+1 \rceil} n)$ time lower bound for sorting on the PRAM [9, 19]. It is straightforward to establish that sorting comparison model lower bounds hold for the element distinctness problem. Hence, the merging-comparison model lower bound translates to an $\Omega(\log_{\lceil (p/n)\log n+1 \rceil} n)$ time lower bound for the element distinctness problem in the PRAM with bounded memory [9, 14, 21].

5.2 Finding the maximum and related problems

Several problems have $\Omega(\log \log_{\lceil p/n+1 \rceil} n)$ rounds lower bounds in the parallel comparison model. The list includes:

- finding the maximum [24];
- merging two lists of equal length [10];
- string-matching [11];
- two-dimensional array-matching [13];
- testing if a string is square-free [4] and
- finding initial palindromes in a string [12].

The following lemma, whose proof is similar to Lemma 5.3 below, shows that these lower bounds can be transformed into $\Omega(\log \log_{\lceil p/n+1 \rceil} n)$ lower bounds for the PRAM model.

Lemma 5.2 *The $\Omega(\log \log_{\lceil p/n+1 \rceil} n)$ lower bounds are resilient.*

The comparison model lower bounds for finding the maximum and merging hold under the distinctness assumption. Hence, these lower bounds can be transformed into PRAM lower bounds [17, 22]. The comparison model lower bounds for string-matching and the related problems mentioned above do not hold under the distinctness assumption, and therefore, these lower bounds translate only to lower bounds in the PRAM with bounded memory.

5.3 Finding an approximate maximum

Alon and Azar [2, 3] give tight lower and upper bounds on the comparison complexity of several approximation problems. We consider as an example the problem of finding an approximate maximum (AM). Namely, an element whose rank belongs in the top ϵn ranks, $1/n \leq \epsilon \leq 1/2$. Alon and Azar prove that under the distinctness assumption,

$$C_{\text{AM}}(n, p) = \Omega(\log \log_{\lceil p/n+1 \rceil} (1/\epsilon) + \log^* n - \log^*(p/n)).$$

We prove next that this lower bound is resilient. Hence the same lower bound holds in the merging-comparison and the PRAM models.

Lemma 5.3 *The $\Omega(\log \log_{\lceil p/n+1 \rceil} (1/\epsilon) + \log^* n - \log^*(p/n))$ lower bound is resilient.*

Proof: We prove that $\log \log_{\lceil p/n+1 \rceil} (1/\epsilon) + \log^* n - \log^*(p/n)$ is resilient (constants do not matter). If $\log \log_{\lceil p/n+1 \rceil} (1/\epsilon) \leq \log^* n - \log^*(p/n) \leq \log^* n$, then

$$\log^* n - \log^*(p/n) = \Theta(\log^* n - \log^*(2^{\log^* n} p/n)).$$

If $\log^* n - \log^*(p/n) \leq \log \log_{\lceil p/n+1 \rceil} (1/\epsilon) \leq \log \log(1/\epsilon)$, we proceed with two cases. If $1 \leq p \leq n \log(1/\epsilon)$, then since $\log^* n \leq 1/\epsilon$, we get that for large enough n ,

$$\log \log(1/\epsilon) \leq 2 \log \log_{\lceil \log^2(1/\epsilon)+1 \rceil} (1/\epsilon) \leq 2 \log \log_{\lceil (p/n) \log(1/\epsilon)+1 \rceil} (1/\epsilon).$$

If $n \log(1/\epsilon) \leq p \leq \binom{n}{2}$, then

$$\log \log_{\lceil (p/n)+1 \rceil} (1/\epsilon) \leq \log \log_{\lceil (p/n) \log(1/\epsilon)+1 \rceil} (1/\epsilon) + 1.$$

Putting the inequalities above together, we establish that

$$\log \log_{\lceil p/n+1 \rceil} (1/\epsilon) = \Theta(\log \log_{\lceil (p/n) \log(1/\epsilon)+1 \rceil} (1/\epsilon)). \quad \square$$

6 Conclusions

By using a finite version of the Erdős–Rado Theorem [16], it is possible to replace the infinite Ramsey–theoretic arguments in §4.1 by finite ones. However, the lower bounds obtained by the general transformation would still require that the input domain is huge. In some cases, direct lower bounds that were given for specific problems require an input domain that is much smaller [8, 14, 17]. It would be of interest to extend the lower bounds given here to smaller domains.

References

- [1] M. Ajtai, J. Komlós, W.L. Steiger, and Szemerédi. Optimal Parallel Selection Has Complexity $O(\log \log n)$. *J. Comput. System Sci.*, 38:125–133, 1989.
- [2] N. Alon and Y. Azar. Finding an approximate maximum. *SIAM J. Comput.*, 18(2):258–267, 1989.
- [3] N. Alon and Y. Azar. Parallel Comparison Algorithms for Approximation Problems. *Combinatorica*, 11(2):97–122, 1991.
- [4] A. Apostolico and D. Breslauer. An Optimal $O(\log \log n)$ Time Parallel Algorithm for Detecting all Squares in a String. *SIAM J. Comput.*, to appear.
- [5] Y. Azar and N. Pippenger. Parallel Selection. *Discrete Applied Mathematics*, 27:49–58, 1990.
- [6] Y. Azar and U. Vishkin. Tight Comparison Bounds on the Complexity of Parallel Sorting. *SIAM J. Comput.*, 16(3):458–464, 1987.
- [7] P. Beame and J. Hastad. Optimal bound for decision problems on the CRCW-PRAM. *J. Assoc. Comput. Mach.*, 36(3):643–670, 1989.
- [8] O. Berkman, Y. Matias, and P. Ragde. Triply-Logarithmic Upper and Lower Bounds for Minimum, Range Minima, and Related Problems with Integer Inputs. In *Proc. 3rd Workshop on Algorithms and Data Structures*, number 709 in Lecture Notes in Computer Science. Springer-Verlag, Berlin, Germany, 1993.
- [9] R. Boppana. Optimal separation between concurrent-write parallel machines. In *Proc. 21st ACM Symp. on Theory of Computing*, pages 320–326, 1989.
- [10] A. Borodin and J.E. Hopcroft. Routing, merging and sorting on parallel models of comparison. *J. Comput. System Sci.*, 30:130–145, 1985.
- [11] D. Breslauer and Z. Galil. A Lower Bound for Parallel String Matching. *SIAM J. Comput.*, 21(5):856–862, 1992.
- [12] D. Breslauer and Z. Galil. Finding all Periods and Initial Palindromes of a String in Parallel. *Algorithmica*, to appear.
- [13] R. Cole, Z. Galil, R. Hariharan, S. Muthukrishnan, and K. Park. Parallel Two Dimensional Witness Computation. Manuscript, 1993.
- [14] J. Edmonds. Lower Bounds With Smaller Domain Size On Concurrent Write Parallel Machines. In *Proc. 3rd ACM Symp. on Parallel Algorithms and Architectures*, pages 322–331, 1991.

- [15] P. Erdős and R. Rado. A Combinatorial Theorem. *J. London Math. Soc.*, 25:249–255, 1950.
- [16] P. Erdős and R. Rado. Combinatorial theorems on classifications of subsets of a given set. *Proc. London Math. Soc.*, 2:417–439, 1952.
- [17] F.E. Fich, F. Meyer auf der Heide, and A. Wigderson. Lower bounds for parallel random-access machines with unbounded shared memory. *Advances in Computing Research*, 4:1–15, 1987.
- [18] R.L. Graham, B.L. Rothschild, and J.H. Spencer. *Ramsey Theory (Second Edition)*. John Wiley & Sons, New York, NY, U.S.A., 1990.
- [19] F. Meyer auf der Heide and A. Wigderson. The complexity of parallel sorting. *SIAM J. Comput.*, 16(1):100–107, 1987.
- [20] S. Moran, M. Snir, and U. Manber. Applications of Ramsey’s Theorem to Decision Tree Complexity. *J. Assoc. Comput. Mach.*, 32(4):938–949, 1985.
- [21] P. Ragde, W. Steiger, E. Szemerédi, and A. Wigderson. The parallel complexity of element distinctness is $\Omega(\sqrt{\log n})$. *SIAM J. Disc. Math.*, 1(3):399–410, 1988.
- [22] B. Schieber and U. Vishkin. Finding all nearest neighbors for convex polygons in parallel: a new lower bound technique and a matching algorithm. *Discrete Applied Mathematics*, 29:97–111, 1990.
- [23] Y. Shiloach and U. Vishkin. Finding the maximum, merging and sorting in a parallel computation model. *J. Algorithms*, 2:88–102, 1981.
- [24] L.G. Valiant. Parallelism in comparison models. *SIAM J. Comput.*, 4:348–355, 1975.

Recent Publications in the BRICS Report Series

- RS-95-10 Dany Breslauer and Devdatt P. Dubhashi. *Transforming Comparison Model Lower Bounds to the Parallel-Random-Access-Machine*. February 1995. 11 pp.
- RS-95-9 Lars R. Knudsen. *Partial and Higher Order Differentials and Applications to the DES*. February 1995. 24 pp.
- RS-95-8 Ole I. Hougaard, Michael I. Schwartzbach, and Hosein Askari. *Type Inference of Turbo Pascal*. February 1995. 19 pp.
- RS-95-7 David A. Basin and Nils Klarlund. *Hardware Verification using Monadic Second-Order Logic*. January 1995. 13 pp.
- RS-95-6 Igor Walukiewicz. *A Complete Deductive System for the μ -Calculus*. January 1995. 39 pp.
- RS-95-5 Luca Aceto and Anna Ingólfssdóttir. *A Complete Equational Axiomatization for Prefix Iteration with Silent Steps*. January 1995. 27 pp.
- RS-95-4 Mogens Nielsen and Glynn Winskel. *Petri Nets and Bisimulations*. January 1995. 36 pp. To appear in TCS.
- RS-95-3 Anna Ingólfssdóttir. *A Semantic Theory for Value-Passing Processes, Late Approach, Part I: A Denotational Model and Its Complete Axiomatization*. January 1995. 37 pp.
- RS-95-2 François Laroussinie, Kim G. Larsen, and Carsten Weise. *From Timed Automata to Logic - and Back*. January 1995. 21 pp.
- RS-95-1 Gudmund Skovbjerg Frandsen, Thore Husfeldt, Peter Bro Miltersen, Theis Rauhe, and Søren Skyum. *Dynamic Algorithms for the Dyck Languages*. January 1995. 21 pp.
- RS-94-48 Jens Chr. Godskesen and Kim G. Larsen. *Synthesizing Distinguishing Formulae for Real Time Systems*. December 1994. 21 pp.
- RS-94-47 Kim G. Larsen, Bernhard Steffen, and Carsten Weise. *A Constraint Oriented Proof Methodology based on Modal Transition Systems*. December 1994. 13 pp.