

ERCW PRAMs and Optical Communication

Philip D. MacKenzie*

Vijaya Ramachandran†

Dept. of Computer Sciences
University of Texas
Austin, TX 78712-1188

May 29, 1997

Abstract

This paper presents algorithms and lower bounds for several fundamental problems on the Exclusive Read, Concurrent Write Parallel Random Access Machine (ERCW PRAM) and some results for unbounded fan-in, bounded fan-out (or ‘BFO’) circuits. Our results for these two models are of importance because of the close relationship of the ERCW model to the OCPC model, a model of parallel computing based on dynamically reconfigurable optical networks, and of BFO circuits to the OCPC model with limited dynamic reconfiguration ability.

Topics: Parallel Algorithms, Theory of Parallel and Distributed Computing.

1 Introduction

In this paper we develop algorithms and lower bounds for fundamental problems on the Exclusive-Read Concurrent-Write (ERCW) Parallel Random-Access Machine (PRAM) model. The ERCW PRAM model has not received much attention, due in part to a general belief that concurrent writing does not add much power to a model without concurrent reading. We show that this is not always the case by presenting algorithms that solve problems on the ERCW PRAM much faster than they could be solved on the EREW PRAM. (See [41] for more details on the different PRAM models.)

We further motivate the ERCW PRAM by its relation to massively parallel computers with dynamically-reconfigurable optical networks. Specifically, we show that the ERCW PRAM (using

*Current address: Department of Mathematics and Computer Science, Boise State University, Boise, ID 83725. This research was supported by Texas Advanced Research Projects Grant 003658480. (philmac@cs.idbsu.edu)

†This research was supported in part by Texas Advanced Research Projects Grants 003658480 and 003658386, and NSF Grant CCR 90-23059. (vlr@cs.utexas.edu)

the ‘Tolerant’ protocol for resolving write conflicts) with n global memory cells and unlimited local memory is computationally equivalent to the OCPC (Optical Communication Parallel Computer) model [3, 28, 29, 31, 50] on n processors. (This is in contrast to the statement given in [3] that the OCPC model is equivalent to an EREW PRAM with n global memory cells.) In previous work, it was shown that the EREW PRAM could be simulated on the OCPC with some overhead per step [44, 18]. This overhead was a result of trying to simulate accesses to an arbitrary number of memory cells. We achieve better simulation results (and in fact, computational equivalence of models) by limiting the number of memory cells of the ERCW PRAM. The following example illustrates the benefit of our approach. Computing the global OR of n bits using the approach in [44] would require simulating the $\Theta(\log n)$ step EREW PRAM algorithm using an expected overhead of $\Theta(\log \log n)$ time per step, resulting in a total expected time of $\Theta(\log n \log \log n)$. However, in our direct approach, we simulate the constant time ERCW PRAM algorithm for global OR with only constant overhead, resulting in a constant time algorithm. (Note that since there is no ‘queue’ delay in optical communication networks, the ERCW PRAM is a better model for parallel machines with such networks than the recently proposed QRQW (or ERQW) model [30].)

Many results for the ERCW PRAM follow directly from results for the EREW PRAM or CRCW PRAM. For instance, the global OR of n bits can be found in constant time on an n processor ERCW PRAM, as on a CRCW PRAM, but broadcasting 1 bit to n processors requires $\Theta(\log n)$ steps, as on an EREW PRAM. The result for broadcasting implies that computing the prefix sums of n inputs and merging two lists of size n both require $\Theta(\log n)$ time also. However, some results obtained directly from EREW PRAM and CRCW PRAM results do not give tight bounds. For instance, the problem of computing the parity of n bits on the ERCW PRAM has a lower bound of $\Omega(\log n / \log \log n)$ from the result for the CRCW PRAM, and an upper bound of $O(\log n)$ from the EREW PRAM. Tight bounds are not known for the ERCW PRAM. Furthermore, tight bounds are not known for many other problems, including the problems of compaction and finding the maximum. In this paper, however, we make significant progress towards developing tighter bounds for these and other problems.

Here is a summary of our results for the ERCW PRAM. Many of these results depend on the write collision resolution protocol used, which we ignore here; these results are stated more precisely in the sections that follow. In the following, n is the size of the input, and all algorithms perform linear work except as noted. We present a k -compaction algorithm that runs in $O(\log \log n + \log k)$ time; a randomized algorithm for k -compaction that runs in $O(\log k)$ expected time; a randomized algorithm for approximate k -compaction that runs in $O(\log \log k)$ time, with failure probability $1/k$; an algorithm for finding the maximum of inputs in the range $[1, n]$ that runs in $O(\log \log n)$ time; an algorithm for chaining that runs in $O(\log \log n)$ time; an algorithm for integer chain-sorting (linear-size integers) that runs in $O(\log \log n)$ time; and an algorithm for integer sorting (polynomial-size integers) that runs in $O(\log n)$ using almost linear work.

We present two lower bounds results for the ERCW PRAM: a lower bound of $\Omega(\sqrt{\log \log n})$ time for solving compaction, and a lower bound of $\Omega(\sqrt{\log n})$ for finding the maximum of general inputs. (The former result, along with a similar result for the OCPC discovered independently by Goldberg and Jerrum, led to the $\Omega(\sqrt{\log \log n})$ lower bound on h -relation routing in Goldberg, Jerrum and MacKenzie [32].)

Finally, we consider unbounded fan-in, bounded fan-out (BFO) circuits. The computations on such circuits can be mapped optimally onto an ERCW PRAM as oblivious algorithms. This could be important if the ERCW PRAM is implemented using an optical network that has only limited

reconfiguration abilities. In an oblivious algorithm the communication patterns are fixed before the algorithm is run. Thus, each processor may have only a small set of other processors with which it needs to communicate, and this set is fixed before the algorithm is run. By designing oblivious algorithms, we may avoid some of the costs of reconfiguration.

We show that any BFO circuit for adding two n -bit integers, merging a bit into an n bit sorted sequence, sorting n bits, or computing the prefix sums or parity of n bits requires $\Omega(\log n)$ depth. Let $TH_{k,n}$ denote the threshold function which outputs 1 if and only if at least k of the inputs are equal to 1. We show that $TH_{k,n}$ can be computed by a linear size, $O(\log \log n + \log k)$ depth circuit, and that any BFO circuit which computes $TH_{k,n}$ requires $\Omega(\log \log n + \log k)$ depth.

The current interest in the OCPC model, the close relation between the OCPC model and the ERCW PRAM model, and the richness of results obtained so far on the OCPC, the ERCW PRAM, and the BFO circuit model, all indicate that these are important models of parallel computation which should be studied further.

The rest of this paper is organized as follows. In Section 2, we define the ERCW PRAM and discuss different write conflict protocols. Section 3, we describe the relationship of the ERCW PRAM to the OCPC model. Section 4 gives lower and upper bounds for compaction problems, and Section 5 gives lower and upper bounds for computing the maximum. In Section 6, we give algorithms for chaining and integer sorting. Section 7 gives lower and upper bounds for computing certain functions on unbounded fan-in, bounded fan-out circuits. Finally, in Section 8 we give relations between the different ERCW PRAM models.

2 Preliminaries

An Exclusive Read, Concurrent Write (ERCW) PRAM consists of a collection of processors, each with infinite local memory, which operate synchronously and communicate through a global memory. Each read or write to global memory takes one time step. Only one processor can read from any memory cell at any time step, but multiple processors may write to a memory cell in a single time step. Write conflicts are handled according to one of the following collision resolution protocols:

Priority The lowest numbered processor succeeds and writes its value to the cell;

Arbitrary An arbitrary processor succeeds and writes its value to the cell;

Common All processors must be writing the same value, which is written to the cell;

Collision A special collision symbol is written to the cell;

Tolerant The cell remains unchanged.

Nice Robust Either the cell remains unchanged, or an arbitrary processor succeeds and writes its value to the cell.

Robust An arbitrary value is written to the cell.

(Since the standard OCPC model uses the Tolerant protocol, we will be most concerned with developing ERCW PRAM algorithms using the Tolerant protocol. We define the OCPC model in Section 3.)

The ERCW(ack) PRAM is an ERCW PRAM with the added feature that a processor which successfully writes to a cell receives an acknowledgement. To retain the spirit of the Common model, we assume no processor receives an acknowledgement in the Common model. To retain the spirit of the Robust model, we assume that false “successful” writes could cause bogus acknowledgements to be sent.

Often we would like to separate the issues of using the global memory as storage for inputs and outputs, and using the global memory for communication. In these cases, we will assume that inputs and outputs are spread evenly among the local memories of the processors. For instance, given p processors and n inputs, we will assume each processor contains n/p inputs in its local memory. With this assumption, we will be free to design algorithms which use less than n cells of global memory.

In our algorithms we do not require that all processors learn the output of an algorithm, for this would force a trivial $\Omega(\log n)$ time lower bound on all our algorithms.

Lemma 2.1 *An n processor ERCW(ack) PRAM with m global memory cells can be simulated on a $\max\{n, m\}$ processor ERCW PRAM with $2m+n$ global memory cells with the same write conflict protocol (except Robust).*

Proof: On the Common model, the ERCW(ack) PRAM and ERCW PRAM are the same, so the simulation is trivial. Otherwise, let E_1 be the ERCW(ack) PRAM and let E_2 be the ERCW PRAM. The first m cells of E_2 will correspond to the m cells of E_1 , the second m cells of E_2 will be used for finding the processor that succeeds in writing to the corresponding cell of E_1 , and the last n cells will be used for writing acknowledgements to the successfully writing processors. We simulate a read of cell j by processor i of E_1 , by having processor i read cell j of E_2 . We simulate a write step as follows. First, every processor j ($0 \leq j \leq m-1$) writes n to cell $m+j$, and then every processor i ($0 \leq i \leq n-1$) writes 0 to cell $2m+i$. For any processor i of E_1 that writes some value v_i to any cell c_i , processor i of E_2 writes i to cell $m+c_i$. Then processor j ($1 \leq j \leq m$) reads cell $m+j$, and if the value read, say v , is not n or “collision”, processor j writes 1 to cell $2m+v$. Now, for each processor i of E_1 writing to some cell c_i , processor i of E_2 reads cell $2m+i$. If it reads a 1, then it writes v_i to cell c_i (with no contention). Note that we require E_2 to have the same write conflict protocol as E_1 , so that the processor that succeeds in writing to cell $m+j$ in E_2 (for some j) is the same as the processor that succeeds in writing to cell j in E_1 . (Note: The simulation for the Collision model is slightly different. We omit the details.) \square

The following lemma will be useful in designing Robust ERCW PRAM algorithms.

Lemma 2.2 *An n processor Nice Robust ERCW(ack) PRAM with m global memory cells can be simulated on a $\max\{n, m\}$ processor Robust ERCW PRAM with $m(n+2)$ global memory cells.*

Proof: Let E_1 be the ERCW(ack) PRAM and let E_2 be the ERCW PRAM. For each cell of E_1 , we associate 1 extra cell to test for the processor which is successful, and n extra cells to be used for writing acknowledgements. We simulate a read of cell j by processor i of E_1 , by having processor i read cell j of E_2 . We simulate a write step as follows. First, every processor j ($0 \leq j \leq m-1$) writes n to cell $m+j$. Then for any processor i of E_1 that writes some value v_i to any cell c_i , processor i of E_2 writes i to cell $m+c_i$, and writes 0 to cell $m(2+i)+c_i$. Then processor j ($1 \leq j \leq m$) reads cell $m+j$, and if the value read, say v , is such that $0 \leq v \leq n-1$, processor j

writes 1 to cell $m(2 + v) + j$. Now, for each processor i of E_1 writing to some cell c_i , processor i of E_2 reads cell $m(2 + i) + c_i$. If it reads a 1, then it writes v_i to cell c_i (with no contention). \square

3 Optical Communication and ERCW PRAMS

Here we describe the technology for optical communication, the OCPC model which is derived from this technology, and its relation to the ERCW PRAM.

3.1 Optical Communication Technology

There are two basic types of optical interconnection networks, fiber optic networks, and free-space optic networks. Research on routing in general fiber optic networks can be found in [1]. In this paper, however, we are only concerned with a specific type of fiber optic network, namely the *Passive optical star coupler* network [6, 7, 20, 35], which allows unit time communication between any pairs of processors. In this network all processors are connected via optical fibers to a passive optical star coupler, which broadcasts messages sent from one processor to all other processors. To allow more flexible communication, time division multiplexing (TDM) or wavelength division multiplexing (WDM) is used. For unit time communication, we must use WDM. For dynamic reconfiguration ability, we must have tunable transmitters and/or receivers. Currently, tunable transmitters and receivers are too slow to be practical.

The other type of optical interconnection network is a free-space optic network. In this network, we do not use wires, but instead use the directional property of light to send messages to the correct destination. Free-space optics has the potential to reduce space requirements and to alleviate many topological difficulties associated with more conventional routing. There are two types of free-space optic networks which achieve unit time communication, the beam spreading/masking network, and the beam steering network. The beam spreading/masking network [39, 48] (also called the crossbar or matrix multiplication network) uses an $n \times n$ array of switching elements with each row assigned to a transmitting processor and each column assigned to a receiving processor. A transmitting processor spreads its light encoded message out to the row of optical switches, and those switches either block the message or send it on to the designated receiving processor. This network has the disadvantages of having n^2 switches, $1/n$ total power transfer, and slow switching time. A beam steering network can either be made up of reconfigurable holograms [33, 5] or acousto-optic deflectors [39]. The typical reconfigurable hologram method assumes the processors are sitting on a board, and there is some holographic material above the board. To transmit a message, a reflecting hologram is written into the holographic substrate and the processor transmits a light encoded message to that hologram, which then steers it to the correct receiving processor. In the acousto-optic deflector method, it is assumed that each processor is connected to a two-dimensional deflector which can be programmed to steer a light encoded message to any other processor. In terms of speed of reconfiguration, the acousto-optic deflector seems to be the fastest, although it is still not as fast as an electronic switch.

3.2 OCPC model

One abstraction of the beam steering model (which could also be considered an abstraction of the passive optical star coupler with tunable transmitters) was first considered by Anderson and Miller [3], and has since been studied in [18, 22, 28, 29, 31, 32, 44, 50]. Various names for this model have been proposed, including *Local Memory PRAM*, *S*PRAM*, *OMC*, *OCPC*, and *OCPC*. We will use the term *OCPC*, denoting *Optical Communication Parallel Computer*.

An *OCPC* consists of a collection of processors, each with infinite local memory, which operate synchronously and communicate by transmitting messages to each other. At any step, a processor can transmit at most one message to another processor. The message will succeed in reaching the processor if it is the only message being sent to that processor at that step. Concurrent transmissions to the same processor will be handled according to one of the following standard collision resolution protocols: Priority, Arbitrary, Common, Collision, Tolerant and Robust [38]. (Note that the standard *OCPC* model uses the Tolerant protocol.)

The *OCPC(ack)* is an *OCPC* with the added feature that a processor which successfully transmits a message to another processor receives an acknowledgement as in the *ERCW(ack) PRAM* (see Section 2).

In Anderson and Miller [3] it is stated that an n processor *OCPC* is equivalent to an n processor *EREW PRAM* with n global memory locations. However this result is incorrect and we show in this section that an n processor *OCPC* is equivalent to an n processor *Tolerant ERCW PRAM* with n global memory locations. We also show additional relationships between *OCPC* and *ERCW PRAM* models, with and without acknowledgements.

Note that our simulations are not at all like those in [44, 18]. Our simulations are simple, straightforward and efficient, causing only constant slowdown.

Lemma 3.1 *A step of an n processor OCPC can be simulated in constant time on an n processor ERCW PRAM with n global memory cells with the same write conflict protocol.*

Proof: One step of the *OCPC* is simulated by a write and a read on the *ERCW PRAM*. We simulate an attempted transmission from processor i to processor j on the *OCPC* by processor i writing to cell j and processor j reading cell j . Since the write conflict resolution protocols are the same, processor j receives the same value on the *ERCW PRAM* as on the *OCPC*. \square

Note that in the next proof, the simulation of an *ERCW PRAM* on an *OCPC* requires that the *OCPC* has at least as many processors as the *ERCW PRAM* has global memory cells. This is not necessarily bad, since only *global* memory cells are counted, and for many *ERCW PRAM* algorithms, only $O(n)$ global memory cells are required.

Lemma 3.2 *A step of an n processor ERCW PRAM with m global memory cells can be simulated in constant time on a $\max\{n, m\}$ processor OCPC with the same write conflict protocol.*

Proof: Assign each processor of the *OCPC* to a memory cell and assign the first n processors also to the n processors of the *ERCW PRAM*. We simulate a write from processor i to cell j on the *ERCW PRAM* by processor i transmitting to processor j . Because the write conflict resolution protocols are the same, processor j receives the same value on the *OCPC* as is written to cell j on

the ERCW PRAM. Processor j stores the value transmitted to it. We simulate processor i reading cell j in the ERCW PRAM by processor i transmitting a read request to processor j and processor j transmitting the value stored there to processor i . Note that there can never be any transmission conflicts when simulating the read step. \square

Lemma 3.3 *A step of an n processor OCPC(ack) can be simulated in constant time on an n processor OCPC with the same write conflict protocol (except Robust).*

Proof: On the Common model, the OCPC and the OCPC(ack) models are equivalent, since no processor ever “succeeds” in a write. On the other models we simulate a transmission of a message M from processor i to processor j by the following procedure. Processor i first transmits i to processor j . If processor j receives a value v , then it transmits an acknowledgement to processor v . If processor i receives an acknowledgement, then it transmits M to processor j . This will be guaranteed to succeed without collision. \square

Lemma 3.4 *A step of an n processor ERCW(ack) PRAM with m global memory cells can be simulated in constant time on an on a $\max\{n, m\}$ processor OCPC with the same write conflict protocol (except Robust).*

Proof: By Lemma 3.3, it suffices to prove the theorem for the OCPC(ack) model. Assign each processor of the OCPC(ack) to a memory cell and assign the first n processors also to the n processors of the ERCW(ack) PRAM. A write step of the ERCW(ack) PRAM is simulated by a transmission step of the OCPC(ack). We simulate a write from processor i to cell j on the ERCW(ack) PRAM by processor i attempting a transmission to processor j . If processor i is successful in writing to the cell, then it will receive an acknowledgement. Because the OCPC(ack) is using the same write conflict resolution protocol, processor i ’s transmission will succeed and processor i will receive an acknowledgement. Processor j now stores the value transmitted to it. A read step of the ERCW(ack) PRAM is simulated by two transmission steps of the OCPC(ack). We simulate processor i reading cell j in the ERCW(ack) PRAM by processor i transmitting a read request to processor j and processor j transmitting the value stored there to processor i . Note that there can never be any transmission conflicts when simulating the read step. \square

4 Compaction problems

In this section, we study the problems of k -compaction and approximate k -compaction on the ERCW PRAM. The *k -compaction* problem takes an array of size n with k marked elements, and places the marked elements into an array of size k . The *approximate k -compaction* problem takes an array of size n with k marked elements, and places the marked elements into an array of size $O(k)$. Compaction and approximate compaction are important subproblems in processor reallocation and load balancing. For instance, if some processors have completed their tasks and some have not, we would like to be able to round up those tasks in a small area, so other processors can find them and help out; this can be achieved using approximate compaction. In addition, compaction and approximate compaction have been used as subroutines for many algorithms, including algorithms for space allocation, estimation, sorting, CRCW PRAM simulation, generation of random permutations, and computational geometry.

We will primarily deal with the k -compaction problem and our only result for approximate k -compaction is a randomized algorithm for the Nice Robust ERCW PRAM, which is described at the end of Section 3.2.1.

4.1 Lower Bounds

Here we will show that 2-compaction on the Robust, Nice Robust, Tolerant, Collision, or Common ERCW PRAM requires $\sqrt{(\log \log n)/2} - 1$ time; that k -compaction on the Arbitrary ERCW PRAM requires time k for $k \leq \sqrt{(\log \log n)/2} - 1$; that 4^k -compaction on the Priority ERCW PRAM requires time k for $k \leq \sqrt{(\log \log n)/2} - 1$; and that $2k$ -compaction on the Priority ERCW PRAM requires time k for $k \leq (\log \log \log n)/4 - 1$. We do not place any restrictions on the number of global memory cells, or the number of processors. (We assume that each of the first n global memory cells contains an input.)

Without loss of generality, we assume that each input is tagged by a pair (i, b) , where i is its index (from 1 to n) and b is 0 if the input is unmarked, 1 if the input is marked. Then we will show a lower bound on solving compaction for the simpler problem of performing compaction on the tags which are marked with 1s. This will obviously imply a lower bound for the general compaction problem.

Our lower bound proof will use the following definition and results.

Definition 4.1 *A sunflower is a collection of sets in which any element that is contained in two of the sets is contained in all of the sets.*

Theorem 4.1 (Erdős-Rado) [21] *Let r and s be positive integers and let \mathbf{F} be a family of sets such that every set in \mathbf{F} has at most r elements, and the number of sets in \mathbf{F} is greater than $r!(s-1)^r$. Then \mathbf{F} contains a sunflower of size s .*

Theorem 4.2 (Turan) [4] *A simple graph with n vertices and m edges has an independent set of size $n^2/(2m+n)$.*

Theorem 4.3 (Dilworth) [19] *Let X be a set and P be a partial ordering on X . If the largest antichain in X has k elements, then there is a set of k chains whose union is X .*

Corollary 4.1 [2] *Given a sequence of n distinct integers, there is either an increasing subsequence of size \sqrt{n} , or a decreasing subsequence of size \sqrt{n} . Furthermore, given k sequences of n integers (each sequence being a permutation of the n integers), there is a subset of size $n^{1/2^k}$ such that in each sequence, the subset is arranged in either increasing or decreasing order.*

Let 2COMP be an algorithm for 2-compaction on the Tolerant, Collision, or Common ERCW PRAM. For a given k , let COMP-A be an algorithm for k -compaction on the Arbitrary ERCW PRAM; let COMP-P- 4^k be an algorithm for 4^k -compaction on the Priority ERCW PRAM; and let COMP-P- $2k$ be an algorithm for $2k$ -compaction on the Priority ERCW PRAM. We will use an adversary argument for our lower bound proof. A step will consist of a write followed by a read.

At each step, the adversary will designate some of the inputs as unmarked (by setting their values to 0) or marked (by setting their values to 1). Let V_t be the set of indices of inputs which have not been designated by step t . These will be the *live inputs* in step t . Initially $V_0 = \{1, \dots, n\}$.

We will say that a *processor is affected* by a live input i in step t if there exists a 0-1 assignment to the values of the other live inputs for which the state of the processor in step t is different when the value of i is 0 and the value of i is 1. Similarly, a *cell is affected* by live input i in step t if there exists a 0-1 assignment to the values of the other live inputs such that the contents of the cell in step t are different when $i = 0$ and when $i = 1$.

We will assume that concurrent writes on the Robust ERCW PRAM are resolved using the Tolerant protocol. We will assume that concurrent writes on the Arbitrary ERCW PRAM are resolved by allowing the lowest numbered processor that is not affected by any live input to write, and if none, then simply the lowest numbered processor.

Let p_t be the maximum number of processors that could be affected by a single live input in step t . Let c_t be the maximum number of cells that could be affected by a single live input. Let $h_t = \max\{c_t, p_t\}$.

Lemma 4.1 *We can construct an adversary such that after step t of 2COMP, COMP-A, or COMP-P- 4^k , as long as $h_t \leq |V_t|^{1/(2^{h_t})}/8$, (1) $h_t \leq 4^t$; (2) $|V_t| \geq |V_{t-1}|^{1/h_{t-1}}/248h_{t-1}^2$; (3) each processor and cell is affected by at most one live input; (4) for COMP-A, at most t items have been designated as marked; and (5) for COMP-P- 4^k , at most 4^{t+1} items have been designated as marked. Also, we can construct an adversary such that after step t of COMP-P- $2k$, (1) $h_t \leq 4^t$; (2) $|V_t| \geq |V_{t-1}|^{1/4^{h_{t-1}}}/124h_{t-1}$; (3) each processor and cell is affected by at most one live input; and (4) at most $2t$ items have been designated as marked.*

Proof: We prove this by induction. First, $p_0 = 0 < 2^0$, $c_0 = 1 = 2^0$, each processor is affected by no inputs, and each cell is affected by at most one input; $h_0 = 1 \leq |V_0|/8$ if $n \geq 8$.

Now assume the lemma is true up to step t . Then we show how to make it hold for step $t + 1$. Let $q = h_t$. By the induction hypothesis, each processor and each cell is affected by at most one live input at the start of step $t + 1$. We will refer to this live input of a processor (or a cell) as *its live input*. Similarly, for a live input, *its processors* and *its cells* are the processors and cells that are affected by it. Let P be a processor that is affected by a live input at the start of step $t + 1$. Processor P *zero-writes* to a cell C in step $t + 1$ if P writes to C if its live input is unmarked. Processor P *one-writes* to a cell in step $t + 1$ if P writes to the cell if its live input is marked.

We will use the following strategy to push through the induction for step $t + 1$. We will describe a strategy for the adversary to fix the values of a subset of the live inputs in step $t + 1$ (thereby rendering these inputs to be no longer live) so that each cell and processor is affected by only a constant number of live inputs. To do this, the adversary fixes certain live inputs to value 0 so that every cell is affected by at most one live input whose processor zero-writes to it. Then the adversary does the same for one-writes – it fixes certain remaining live inputs to value 1 so that every cell is affected by at most one live input whose processor one-writes to it. So at this point, every cell is affected by at most 3 live inputs – its live input from step t , one live input for zero-writes, and one live input for one-writes. As a result, every processor is affected by at most 7 live inputs – its live input l from step t , and 3 live inputs each from the cell it reads in step $t + 1$ if $l = 0$ and the cell it reads in step $t + 1$ if $l = 1$. The adversary then uses Theorem 4.2 on a graph derived from

processors, cells and live inputs to extract a large independent set in the graph that corresponds to a subset of live inputs; all other live inputs are set to 0, and it is shown that this results in each processor and cell being affected by at most one live input in step $t + 1$. We fill in the details of this strategy below, where we also establish that this can be accomplished by fixing the values of a sufficiently small number of live inputs so that all four of the inductive assumptions will hold at the end of step $t + 1$.

We first consider zero-writes to a cell. We say that the adversary *zeros* a live input if it designates it as unmarked. Note that once an adversary zeros a live input, the input is not live anymore. We will make the adversary zero some of the live inputs so that each cell is affected by at most one live input whose processor zero-writes to it in this step. To do this, we describe a simple procedure for the adversary. Until each cell has at most one live input whose processor zero-writes to it, the adversary arbitrarily chooses a remaining live input l and for each cell to which l 's processors could write, the adversary zeros the other live inputs whose processors zero-write to the same cell – up to two per cell for Tolerant, Collision, or Common (one suffices for Common), and the one corresponding to the lowest numbered processor that zero-writes to that cell, for Arbitrary. Notice that once we zero live inputs whose processors zero-write to a cell according to the rule specified above, the value of the cell is fixed for this step, and no information about live inputs is written to it. Using this procedure, the adversary sets at most $2p_t$ live inputs for each live input l chosen. Thus we are left with $|V_t|/(2p_t + 1)$ live inputs. Let $m = |V_t|/(2q + 1)$. At this point, there are at least m live inputs, and for each cell there is at most one live input that could cause a processor to zero-write to it.

Now we deal with one-writes. We will fix the values of some live inputs so that every cell is affected by at most one live input during a one-write. For this, we will find a sufficiently large set of live inputs such that if one of these inputs is one-written to a cell by a processor, then either each other live input is also one-written to the same cell by a processor, or none are one-written to the same cell. This is equivalent to finding a sunflower in a group of sets, where each set contains the cells one-written to by processors which know a given live input. Thus we can apply Theorem 4.1 with $r = q$ and m as a lower bound on the number of sets. It follows that there must be a sunflower of size $(m/q!)^{1/q} \geq m^{1/q}/q$. We set all the live inputs that are not in the sunflower to zero. Let m' be the number of live inputs remaining. Then $m' \geq m^{1/q}/q$. At this point we have the property that each cell is either affected by at most one live input, or is affected by all live inputs on a one-write. We only need to consider the latter set of cells. Call this set of cells \mathcal{S} . On the Tolerant, Collision, and Common ERCW PRAM models, note that the cells in \mathcal{S} are fixed (because two of the live inputs must be marked, implying that there will be a write-conflict in each cell) and thus the cells are not affected by any live input on a one-write. For Arbitrary we need to do something slightly different. We mark a live input corresponding to a processor that one-writes to a cell in \mathcal{S} . Note that by the definition of \mathcal{S} , every other cell in \mathcal{S} is one-written to by a processor that is affected by that input. By our Arbitrary rule, all cells in \mathcal{S} will be fixed. For Priority we also need to do something different. Note that there are at most $p_t \leq q$ cells in \mathcal{S} . Then we proceed according to one of the following options. (Option 1 is used to obtain the bound for COMP-P- 4^k , and Option 2 is used to obtain the bound for COMP-P- $2k$.)

1. For each of the cells in \mathcal{S} we mark the live input corresponding to the lowest numbered processor that performs a one-write into that cell. Thus we fix at most q live inputs, and we fix each cell in \mathcal{S}
2. Consider a cell $v \in \mathcal{S}$, and rank each live input from 1 to m' in the same order as the processor

numbers of processors that one-write to v . For each other cell $v' \in \mathcal{S}$ consider the sequence formed by the ranks of live inputs (computed from cell v) in order of processor numbers of processors that one-write to v' . By Corollary 4.1, we can choose a subset of live inputs of size $m^{1/2^{q-1}}$ such that for each cell's sequence, the ranks (from the subset) are all present in increasing or decreasing order. We zero the other live inputs, and mark the live inputs corresponding to the lowest and highest ranks in this subset. Thus we fix each cell in \mathcal{S} .

Let m'' be the number of live inputs remaining. Note that $m'' = m' \geq |V_t|^{1/q}/4q - q \geq |V_t|^{1/q}/8q$, (using the fact that $q \leq |V_t|^{(1/2^q)}/8$) except for COMP-P-2k, in which case $m'' \geq m^{1/2^{q-1}} \geq |V_t|^{1/4^q}/4$.

At this point, each cell is affected by at most 3 live inputs – its live input from step t , and one live input each for a zero-write and a one-write in the current step. Now at most $c' = c_t + 2p_t$ cells are affected by a live input after this write step. Also, since each cell can be read by at most one processor, after the read step at most $p' = c' + p_t$ processors will be affected by a live input, and each processor could be affected by at most seven live inputs (its live input l from step t and 3 live inputs each from the cell it could read in the current step if $l = 0$ and if $l = 1$). The adversary will zero some inputs so that each cell and processor is only affected by one live input. To do this, we construct a graph in which the vertices are the live inputs and edges between vertices exist if the live inputs are known to the same cell or processor. There are m' vertices in this graph and the degree of each vertex is at most $2c' + 6p'$. Hence the number of edges in the graph is no more than $m'' \cdot (2c' + 6p')/2 < m'' \cdot (4c_t + 11p_t)$. By Turan's Theorem (Theorem 4.2), we can find an independent set of vertices of size

$$(m'')^2 / (m'' + 2m''(4c_t + 11p_t)) = m'' / (1 + 8c_t + 22p_t) \geq m'' / (30q + 1)$$

This is at least $|V_t|^{1/q} / (31q \cdot 8q) \geq |V_t|^{1/q} / 248q^2$, except for COMP-P-2k, in which case it is at least $|V_t|^{1/4^q} / (31q \cdot 4) \geq |V_t|^{1/4^q} / 124q$. The inputs corresponding to vertices not in the independent set are then zeroed by the adversary.

Now we can set $p_{t+1} = c_t + 3p_t \leq 4^{t+1}$, $c_{t+1} = c_t + 2p_t \leq 4^{t+1}$, V_{t+1} to be the remaining set of live inputs, where $|V_{t+1}| \geq |V_t|^{1/q} / 248q^2$ except for Option 2 for the Priority model, in which case $|V_{t+1}| \geq |V_t|^{1/4^q} / 124q$. Since $q = h_t$, the lemma follows. \square

Theorem 4.4 (1) *Solving 2-compaction on a Robust, Common, Collision, or Tolerant ERCW PRAM requires $\sqrt{(\log \log n)/2} - 1$ steps; (2) for $k \leq \sqrt{(\log \log n)/2} - 1$, solving k -compaction on an Arbitrary ERCW PRAM requires at least k steps; (3) for $k \leq \sqrt{(\log \log n)/2} - 1$, solving 4^k -compaction on a Priority ERCW PRAM requires at least k steps; and (4) for $k \leq (\log \log \log n)/4 - 1$, solving $2k$ -compaction on a Priority ERCW PRAM requires at least k steps.*

Proof: For the first three parts of the theorem, we can bound $|V_t|$ using Lemma 4.1 as follows.

$$|V_t| \geq \frac{|V_0| \prod_{0 \leq i \leq t} 1/4^i}{4^{2(t+1)} 248^t} \geq \frac{\prod_{0 \leq i \leq t} 1/4^i}{16(4000^t)} \geq \frac{n^{4^{-\sum_{0 \leq i \leq t} i}}}{16(4000^t)} \geq \frac{n^{4^{-t(t+1)/2}}}{16(4000^t)} \geq \frac{n^{2^{-(t+1)}}}{16(4000^t)}$$

provided t satisfies the constraint $h_t \leq |V_t|^{1/(2h_t)}/8$. We note that this condition is satisfied for all $t \leq \sqrt{(\log \log n)/2}$ for sufficiently large n . Thus after $T = \sqrt{(\log \log n)/2} - 1$ steps, we will have $|V_T| \geq \Omega(2^{\sqrt{\log n}} / 16(4000^{\sqrt{\log n}/2}))$. Now assume $k \leq T$. For large n , $|V_T| \geq \max\{5, k+2, 4^k+2\}$.

For the case of 2-compaction on the Robust, Nice Robust, Collision, Common, or Tolerant models, there will be 3 live inputs which do not affect either of the first two global memory cells. Thus the adversary could mark two of them so that the compaction fails.

For the case of k -compaction on the Arbitrary model, after $k - 1$ steps, at most $k - 1$ inputs have been designated as marked, and there will be at least 2 live inputs that do not affect any of the first k cells. Thus the adversary could mark (or unmark) one of them so that the compaction fails.

For the case of 4^k -compaction on the Priority model, after $k - 1$ steps, at most $1 + 4 + 4^2 + \dots + 4^{k-1} \leq 4^k - 1$ inputs have been designated as marked, and there will be at least 2 live inputs that do not affect any of the first 4^k cells. Thus the adversary could mark (or unmark) one of them so that the compaction fails.

For the last part of the theorem ($2k$ -compaction on the Priority model), we can bound $|V_t|$ using Lemma 4.1 as follows.

$$|V_t| \geq \frac{|V_0| \prod_{0 \leq i \leq t} 1/2^{4^i}}{4^{t+1} 124^t} \geq \frac{n^{2^{-4^t+1}}}{4(500^t)}$$

Thus after $T = (\log \log \log n)/4 - 1$ steps, we will have $|V_T| \geq \Omega(2\sqrt{\log n}/4(500^{\log \log \log n}))$. Now assume $k \leq T$. For large n , $|V_T| \geq 2k + 2$.

After $k - 1$ steps, at most $2(k - 1) < 2k - 1$ inputs have been designated as marked, and there will be 2 live inputs which do not affect any of the first $2k$ cells. Thus the adversary could mark (or unmark) one of them so that the compaction fails. \square

Since this lower bound holds for any number of processors and global memory cells, it also holds for the ERCW(ack) PRAM, the OCPC, and the OCPC(ack) models with the same write conflict resolution protocols.

4.2 Upper Bounds

First we note that there is a simple algorithm that consists of performing one concurrent write, which solves k -compaction in $O(k)$ time on an Arbitrary ERCW PRAM. However, this algorithm will not work unless some processor can succeed in each write. For the other write conflict resolution protocols we need a different approach.

We construct an algorithm which runs in $O(\log \log n + \log k)$ time on a Tolerant ERCW PRAM. This is an adaptation of an $O(\log k)$ time algorithm for k -compaction on the Robust CRCW PRAM given in Fich *et al.* [24]. If $k \geq n^{1/5}$, we use a standard EREW prefix sums algorithm to perform the compaction in $O(\log k)$ time. Otherwise, as in [24] we partition the input cells into groups of l cells, where

$$l = \begin{cases} 2k(k-1) & \text{if } k \leq \frac{\log n}{4 \log \log n} \\ \frac{(k-1) \log n}{3 \log \log n - 1} & \text{if } \frac{\log n}{4 \log \log n} < k \leq \log n, \text{ and} \\ \frac{(k-1) \log n}{3 \log k - 1} & \text{if } \log n < k < n^{1/5}. \end{cases}$$

We solve k -compaction within each group in $O(\log l) = O(\log k)$ time using the standard EREW algorithm.

Let $y_j = j$ if the j th group contains a non-zero entry, and let $y_j = 0$ otherwise. As in [24], if we solve the k -compaction problem for $y_1, \dots, y_{n/l}$, we can solve the original k -compaction problem in $O(\log l) = O(\log k)$ more steps. To solve the k -compaction problem on $y_1, \dots, y_{n/l}$, we proceed as in [24], and reduce the problem in $O(\log l)$ (i.e., $O(\log k)$) time on a Tolerant or Collision ERCW PRAM to k -compaction in an array of size $2ln^{(k-1)/l}$. If $k > \log n/4 \log \log n$, then $2ln^{(k-1)/l} = k^{O(1)}$, and clearly, we can solve this k -compaction problem in $O(\log k)$ time. If $k \leq \log n/4 \log \log n$, then in [24] the problem is solved in constant time using a CRCW technique, but on the ERCW PRAM we will, instead, solve the problem recursively on an array of size $2ln^{(k-1)/l} \leq 2ln^{1/2k}$. For $k \leq \log n/4 \log \log n$, $2l \leq n^{1/k}$, so we can bound the time of this recurrence by

$$T(n) \leq T(n^{3/(2k)}) + O(\log k) = O\left(\frac{\log \log n}{\log k}(\log k)\right) = O(\log \log n).$$

Hence, using the fact that the Tolerant protocol is a Nice Robust protocol, and using Lemma 2.2, we obtain the following theorem.

Theorem 4.5 *Let $t(n, k) = \log \log n + \log k$. The k -compaction problem can be solved in $O(t(n, k))$ time on an $n/t(n, k)$ processor Collision or Tolerant ERCW PRAM with $n/t(n, k)$ global memory cells, and on an $n/t(n, k)$ processor Robust ERCW PRAM with $O((n/t(n, k))^2)$ global memory cells. It can also be solved in $O(k)$ time on an n/k processor Arbitrary ERCW PRAM with 1 global memory cell.*

4.2.1 Randomized

We present two results for randomized algorithms for compaction. Both results are obtained by having processors hash into random locations in an array. We will assume the inputs are given in the local memories of the processors.

Our first result is an $O(\log k)$ expected time randomized algorithm for compaction on the Robust ERCW PRAM with n processors and n memory locations. If $k > n^{1/16}$, we use the standard $O(\log n) = O(\log k)$ time parallel prefix algorithm to perform the compaction. Otherwise, let A be an array of size k^4 . Clear this array, and let each processor representing some marked element write its processor number to a random location of A . We use an $O(\log k)$ time prefix operation to check if k processors succeeded without collision. If so, we compact them into the first k locations in the array by computing prefix sums, and inform the marked processors of their success. If any processor doesn't receive notice of success, (and so, by construction no processor receives notice of success) it simply retries the procedure. It is straightforward to see that the probability of failure decreases geometrically with the number of attempts. Then using Lemma 2.2, we obtain the following theorem.

Theorem 4.6 *An $n/\log k$ processor Robust OCPC or an $n/\log k$ processor Robust ERCW PRAM with no more than $n/\log k$ global memory cells can solve k -compaction in $O(\log k)$ expected time.*

Proof: We only need to show that the Nice Robust ERCW(ack) PRAM procedure runs in $O(\log k)$ expected time. The probability that there are any collisions when writing to the array of size k^4 is $\leq k/k^4$. The time for writing and the prefix sums computation is $c \log k$ for some constant c , and

thus the total expected time until the procedure is successful is given by

$$c \log k + (c \log k)/k^3 + (c \log k)/k^6 + \dots = c \log k \sum_{i=0}^{\infty} k^{-3i} = O(\log k).$$

□

Now we describe an $O(\log \log k)$ algorithm for approximate compaction on an n processor Nice Robust ERCW(ack) PRAM which works with probability $1 - \frac{1}{k}$ and uses only $O(k)$ global memory locations. Each processor with a marked element writes it to a random location in an array of size $8k$. If a processor receives an acknowledgement, it idles. If not, the processor writes its element into an array of size $4k$. This procedure continues for a total of $\log \log k$ steps as the array size reduces by half each time. Then we attempt for three steps to write the remaining elements into arrays of size k .

It is not difficult to see, using a Chernoff bound, that the number of remaining elements after step t is at most $\max\{k2^{-(2^t+t-1)}, k^{1/4}\}$ with probability $1 - te^{-k^{1/4}/4}$.

Lemma 4.2 *After step t , with probability $1 - te^{-k^{1/4}/4}$, the number of remaining elements will be at most $\max\{k2^{-(2^t+t-1)}, k^{1/4}\}$.*

Proof: By induction. The base case is trivial. For the induction step, the probability of having more than the number left will be at most $(t-1)e^{-k^{1/4}/4}$. Assume we have the required number left. We will randomly write these to an array of size $k2^{-t+4}$. The probability of a write collision for any given processor is $\leq k2^{-(2^t+t-1)}/k2^{-t+4}$, so the expected number remaining after step t will be

$$k2^{-(2^{t-1}+(t-1)-1)}(2^{t-4-(2^{t-1}+(t-1)-1)}) \leq k2^{-(2^t+t)}$$

If $k2^{-(2^t+t)} \geq k^{1/4}/2$, then by a Chernoff bound, the probability of having over twice this number, or $k2^{-(2^t+t-1)}$, is less than $e^{-k2^{-2^t+t+1}/4} \leq e^{-k^{1/4}/4}$. If $k2^{-(2^t+t)} < k^{1/4}/2$, then by a Chernoff bound, the probability of having over $k^{1/4}$ remaining is less than $e^{-k^{1/4}/4}$. Thus the total probability of failure at step t is less than $te^{-k^{1/4}/4}$. □

The probability of any element colliding in the last three steps is $k^{1/4}(1/k^{3/4})^3 \leq 1/k^2$. Since $(\log \log k)e^{-k^{1/4}/4} \leq 1/k^2$ for sufficiently large k , we can bound the total probability of not succeeding by $1/k$. Then using Lemma 2.1, we obtain the following theorem.

Theorem 4.7 *An $n/(\log \log k)$ processor Nice Robust OCPC or an $n/(\log \log k)$ processor Nice Robust ERCW PRAM with $n/(\log \log k)$ global memory cells can solve approximate k -compaction in time $O(\log \log k)$, with probability $1 - 1/k$.*

If we are on a model that can compute a Global OR in constant time (i.e. not the Robust or Nice Robust model), a single fixed processor can detect a failure and inform the “successful” processors in $O(\log k)$ time. To use this information in an algorithm that runs in $\Omega(\log k)$ time and uses compaction as a subroutine, the algorithm would run in stages of $O(\log k)$ steps, and a check for successful approximate compactions would be performed at the end of each stage. The expected asymptotic running time of this algorithm would then not be affected by the possible failure of the approximate compaction.

5 Maximum

Finding the maximum of n input elements requires $\Theta(\log n)$ time on an EREW or CREW, even when the values are restricted to be either 0 or 1 [17]. Finding the maximum of n inputs on a Priority CRCW with n processors requires $\Theta(\log \log n)$ time if the inputs come from a large range and $O(k)$ time if the inputs are restricted to the range $[1, n^k]$ [26]. In this section we will show that finding the maximum over an unrestricted range requires $\Omega(\sqrt{\log n})$ time on an ERCW PRAM. If input values are restricted to the range $[1, s]$, $s \leq n$ we will show that the maximum can be found in $O(\log \log s)$ time on the Common or Tolerant ERCW PRAM, and that $\Omega(\sqrt{\log \log s})$ time is required to find the maximum.

5.1 Lower Bounds

Our first lower bound will be for the case of unrestricted input range. Wlog, we will assume that all the inputs are distinct. Let MAX be an algorithm on the Priority ERCW PRAM which finds the maximum of n inputs stored one per processor in the first n processors. For concreteness, assume that the output of MAX is to be stored in the first global memory cell. Consider step t of MAX. Let $V_t \subseteq \{I_1, \dots, I_n\}$ be the set of inputs which could still be the maximum. These will be called the live inputs. Let $S_t \subseteq \{1, 2, 3, \dots\}$ be the possible values for the live inputs, as restricted by the adversary. Let $F_t = \{f_i | I_i \notin V_t\}$ be the adversary's assignment of values to fixed inputs. Let k_t be the maximum number of processors or cells which are affected by any given live input.

As the computation proceeds, the adversary fixes the values of certain inputs and maintains a set of allowed inputs, such that, after each step, each processor is affected by at most one live input. Initially $V_0 = \{I_1, \dots, I_n\}$, $S_0 = \{1, 2, 3, \dots\}$ and is infinite, and $F_0 = \emptyset$.

Lemma 5.1 *We can construct an adversary such that after step t of MAX, the following properties hold: (1) $V_t \subseteq V_{t-1}$ and $|V_t| \geq \frac{|V_{t-1}|}{3^{t+1}}$; (2) each processor and cell is affected by at most one input in V_t ; (3) $k_t \leq 3^t$; (4) $S_t \subseteq S_{t-1}$ and S_t is infinite; (5) $F_{t-1} \subseteq F_t \subseteq \{1, 2, 3, \dots\} - S_t$; and (6) an input in V_t affects at most k_t processors and k_t cells.*

Proof: Define a processor's read (write) function at step t to be a function which maps the live input this processor knows about at step t to the cell which it reads(writes).

The lemma is obviously true for step 0. Now assume the lemma is true for all steps up to $t - 1$. Consider step t of MAX. Each processor knows at most 1 live input. Depending on this live input, it will (possibly) write to a specific cell and (possibly) read from a specific cell. As in [25], we use Ramsey Theoretical arguments to restrict the possible values for inputs such that (1) each processor either writes for all values or for no values; (2) each processor either reads for all values or for no values; (3) each processor's read function is either constant or one-to-one; (4) each processor's write function is either constant or one-to-one; and (5) any two read and/or write functions from any step $t' \leq t$ are either identical or disjoint. Let $S'_t \subseteq S_{t-1}$ be the set of possible values for inputs after this restriction. The fact that S'_t is still infinite is shown in [25]. Now, because we are assuming the inputs are distinct, one-to-one read (or write) functions can not be used to by processors to gain any information. (Notice that in this case a processor will only read a cell if the input I it knows is a certain value, and the only possible writes to that cell occur from some processor which also only knows I , since no other input could be that value) Thus we can restrict

our attention to constant read and write functions. In this case, since the reading and writing cells are fixed, and since we are using the Priority model, a processor reading a location knows exactly which processor last wrote (successfully) to that cell, and thus which single live input it will learn about, if any.

Form a graph with the live inputs (V_{t-1}) as vertices and an edge between two vertices if a processor is affected by those two live inputs (one from the previous step, and one from the cell just read). Take the largest independent set in the graph and let V_t be the inputs associated with this independent set. Fix the smallest distinct values f_i from S'_t to variables $i \in V_{t-1} - V_t$ and remove them from S'_t to obtain S_t . Add these values to F_{t-1} to obtain F_t . By Turán's Theorem, we can choose a set V_t such that $|V_t| \geq \frac{|V_{t-1}|^2}{|V_{t-1}| + 2e}$, where e is the number of edges in the graph. The degree of this graph is at most $3k_{t-1}$. (Consider a live input I_j . At most k_{t-1} processors are originally affected by that live input, and possibly read a cell affected by another live input. Also, each of the k_{t-1} cells originally affected by I_j , plus the k_{t-1} cells written to by the processors affected by I_j could be read by other processors.) Thus the number of edges is at most $\frac{|V_{t-1}|3k_{t-1}}{2}$, so $|V_t| \geq \frac{|V_{t-1}|^2}{|V_{t-1}|3^{t+1}} \geq \frac{|V_{t-1}|^2}{|V_{t-1}|3^{t+1}} \geq \frac{|V_{t-1}|}{3^{t+1}}$. By the same reasoning as above, the number of cells affected by a given live input after step t is at most $2k_{t-1} < k_t$, and the number of processors affected by a given live input after step t is at most $3k_{t-1} < k_t$. \square

Theorem 5.1 *Finding the maximum of n inputs on a Priority ERCW PRAM requires $\Omega(\sqrt{\log n})$ communication steps.*

Proof: By Lemma 5.1, one can conclude that $|V_t| \geq n/3^{(t+1)+t+(t-1)+\dots+2} \geq n/3^{(t+2)(t+1)/2}$. Then at step $T = \sqrt{2 \log n / \log 3} - 2$, $V_T \geq 2$, and the first cell is affected by at most one of these inputs. Then the adversary can simply set the other input to be higher than the value stored in the first global memory cell. \square

Theorem 5.1 only holds when the inputs are drawn from a very large range. For the case of inputs restricted to a small range, we can prove the following lower bound.

Theorem 5.2 *Finding the maximum of n inputs drawn from the range $[1, s]$, for $s \leq n$, requires $\Omega(\sqrt{\log \log s})$ time on a Robust, Nice Robust, Tolerant, Collision, or Common ERCW PRAM.*

Proof: Consider an input array of size n which consists of all zeros except for two entries at locations $i, j \in [1, s]$, which contain the values i and j , respectively. If we have an algorithm to find the maximum in this n element input, then we solve the 2-compaction problem on this array with the same time bound, and thus the $\Omega(\sqrt{\log \log s})$ lower bound on 2-compaction applies to the problem of finding the maximum. \square

5.2 Upper Bounds

We first show a doubly logarithmic time algorithm for Rightmost One problem, and then show an algorithm for Maximum which is doubly logarithmic in the number of different values allowed for the input elements (i.e., the 'range'), up to a range of size n .

Theorem 5.3 *The rightmost one of n bits can be found on an $n/\log\log n$ processor Common, Tolerant, or Collision ERCW PRAM in $O(\log\log n)$ time, or on an n processor Priority ERCW PRAM in constant time.*

Proof Sketch: The algorithm for the Priority model is trivial. For the other models, we divide the array into subarrays of size \sqrt{n} and recursively find the rightmost subarray which contains a 1 and the rightmost one in each subarray. Note that on the CRCW PRAM, the recursion is unnecessary, since n processors can find the rightmost one in an array of size \sqrt{n} in constant time. \square

Theorem 5.4 *The maximum of n inputs in the range $[1, s]$ can be found on a $\max\{n, s\}/\log\log s$ processor Common, Tolerant, or Collision ERCW PRAM in $O(\log\log s)$ time.*

Proof Sketch: We create an array of size s , place 1's at positions in the array which correspond to input values, and find the rightmost one in $O(\log\log s)$ time. \square

Using an algorithm similar to one in [23], we obtain the following result for finding the maximum of binary inputs (i.e., the global OR) on a Robust ERCW PRAM.

Theorem 5.5 *An $n/\log\log n$ processor Robust ERCW PRAM can find the global OR of n bits in $\Theta(\log\log n)$ time with error probability $\frac{1}{n}$.*

6 Chaining and Integer Sorting

Our goal is to obtain a fast ERCW PRAM algorithm to sort integers from a polynomial range. To do this, we first develop algorithms for the Chaining problem, which takes an n -bit array as input and finds for each 1 in the input, the position of the nearest 1 to its left.

Theorem 6.1 *The Chaining problem on n bits can be solved on an $n/\log\log n$ processor Common, Tolerant, or Collision OCPC or ERCW PRAM in $O(\log\log n)$ time.*

Proof: First partition the input array into consecutive groups of $\log^2 n$ bits and solve the Nearest Ones problem in these groups computing prefix sums in $O(\log\log n)$ time. Now we assign a 1 to each group which contained a 1, and solve the Chaining problem on $n/\log^2 n$ bits. Once this is done, the processor associated with the leftmost 1 bit in each group can simply read the position of the rightmost 1 bit in the nearest group to the left which contains a 1, and write it to the output array.

To solve the Chaining problem on $n/\log^2 n$ bits, notice that in $O(\log\log n)$ steps, we can broadcast each bit to $\log n$ processors, so we have $\log n$ processors working for each bit. Imagine a complete binary tree formed over the $n/\log^2 n$ bits. For each bit, associate one of its associated processors with each of its ancestors. Now, in parallel for each node in the tree, solve the Rightmost One and Leftmost One problems for the subarray containing the elements at the leaves of the subtree rooted at the node. This computation can be performed in $O(\log\log n)$ time and linear work as shown in Section 5. Each processor assigned to a bit with value 1 will then know if its bit is the rightmost or leftmost at that node. For each bit b with value 1, use a standard prefix sums

computation over the processors associated with that bit to find the lowest node (closest to the leaves) for which b is not the leftmost bit. Then the processor for b assigned to that node can look at the left child x of that node to find the rightmost bit r with value 1 in the subarray for x . This is the nearest one to the left of b . There will be no read conflict, because at each node there is at most one bit in its subarray which is the leftmost bit with value 1. \square

The following theorem addresses the Chaining problem in the case when there is a processor associated with each non-zero element in the input.

Theorem 6.2 *Let $A[1..n]$ be an array of zeros and ones, with a processor associated with each $A[i] = 1$ (hence the number of processors is equal to the number of ones in the input). Let the priorities of the processors decrease with the position within A of the element to which a processor is associated. The Chaining problem on this input can be solved in $O(\log \log n)$ time on a PRIORITY ERCW(ack) PRAM.*

Proof: The following algorithm solves the problem within the stated bounds. We create an auxiliary array of size \sqrt{n} and divide the input array A into \sqrt{n} blocks of size \sqrt{n} . All processors assigned to elements in the i th block perform a concurrent write of their element's position within A into location i of the auxiliary array. The processors that succeed delete their entry in A and recursively solve the problem in the auxiliary array. The remaining processors recursively solve the problem within their blocks. The recursive solutions are then combined into a solution for the original problem in constant time. Since all of the recursive subproblems are of size \sqrt{n} , the overall algorithm runs in $O(\log \log n)$ time. \square

We can now perform a stable sort of n integers in the range $[0..n - 1]$ in $O(\log n)$ time with $n \log \log n / \log n$ processors and $O(n^2)$ space on a PRIORITY ERCW(ack) PRAM as follows. As in the CRCW algorithm of [36] we use an $n \times n$ array (which is assumed to be initialized to zero). For each index i in the input, if element i has value j then a 1 is written into position (i, j) of the array. We then solve the chaining problem on the $n \times n$ array (interpreted as a $1 \times n^2$ array) to obtain the sorted elements in a linked list. This portion of the algorithm runs in $O(\log \log n)$ time using n processors using the algorithm in the proof of Theorem 6.2. To obtain the sorted list in an array form, we perform list ranking to find the position of each element in the output array. Since the sort is stable this allows us to sort n integers in the range $[0..n^k - 1]$, for any constant k , within the same processor-time bounds. It also allows us to reduce the space requirement to $n^{1+\epsilon}$, for any constant $\epsilon > 0$, by viewing each value as the sum of powers of n^ϵ . This gives us the following theorem.

Theorem 6.3 *Integer chain-sorting can be performed on n integers in the range $[0..n - 1]$ in $O(\log \log n)$ time with n processors on a Priority ERCW(ack) PRAM. Integer sort into an array can be performed on n integers in the range $[0..n^k]$ in $O(\log n)$ time with $n \log \log n / \log n$ processors and $n^{1+\epsilon}$ space on a Priority ERCW(ack) PRAM.*

7 Unbounded Fan-in, Bounded Fan-out Circuits

Since fast dynamic reconfiguration between a large number of processors in optical networks does not yet seem to be technically feasible, we would like to find ways of reducing the need for it.

One way is to design *oblivious* algorithms. An oblivious algorithm for an OCPC is an algorithm in which, if a processor transmits a message during a step, the destination of that transmission is fixed before the algorithm is run. An oblivious algorithm for an ERCW PRAM is an algorithm in which, if a processor reads or writes to any cells during a step, the locations read and written to by that processor are fixed before the algorithm is run. In oblivious algorithms the pattern of transmissions is known prior to the start of the algorithm (i.e., it is not dependent on the inputs). Thus, each processor may have only a small set of other processors with which it needs to communicate, and this set is fixed before the algorithm is run. Therefore, we would be able to fix or preset the transmission elements, and we may avoid some of the reconfiguration costs.

A *Boolean circuit* is a directed acyclic graph. The nodes of in-degree 0 are called the inputs. The nodes of in-degree $k > 0$ are called *gates* and are labeled with AND, OR, or NOT. (Nodes labeled with NOT have must have in-degree 1.) The *fan-in* of a node is its in-degree and the *fan-out* of a node is its out-degree. One of the nodes is designated as the output node. The *size* of a circuit is the number of gates, and the *depth* of a circuit is the maximum length of a directed path from an input to the output. A *Boolean formula* is a Boolean circuit whose underlying graph is a tree (all nodes except the output node have out-degree 1).

A special type of oblivious algorithm is given by a *bounded fan-out Boolean circuit (BFO)*, which is a Boolean circuit in which the fan-out of each node is at most two. We assume standard definitions for circuits and formulas [10]. A BFO circuit with size s and depth d can be simulated in a straightforward way by an s processor, d step oblivious OCPC algorithm. Just as unbounded fan-in, unbounded fan-out circuits correspond closely to the CRCW PRAM [11], and the study of bounded fan-in circuits often sheds light on algorithms for the CREW and EREW PRAM, we believe that the study of BFO circuits should enhance the understanding of the ERCW PRAMs

We now give some results on solving some fundamental problems on BFO circuits.

7.1 Lower Bounds

Our first result shows how to transform a BFO circuit into something resembling a formula, so that we can obtain a lower bound the depth of the circuit using known lower bounds on formula size.

We will assume the size of a formula is the number of inputs in the circuit corresponding to the formula.

Theorem 7.1 *Let f be a Boolean function over n variables. If f is computed by a circuit of depth d with fan-out at most c (with one input corresponding to each variable), then there is a Boolean formula of size at most nc^d which computes f .*

Proof: Let C be a depth d circuit in which each gate has fan-out at most c . Let C' be the same circuit, but with every gate with some fan-out $c' > 1$ replaced by a gate with a single output leading into a “fan-out” gate which fans out the output to c' other gates. Then C' has depth at most $2d$. Now consider the following percolate operation. Assume a gate g has an output which enters a c' -fan-out gate. The percolate operation replaces this with a c' -fan-out gate at each input which fans out each input into c' duplicates of gate g . This has the effect of percolating the gate g up in the circuit.

We perform percolate operations on C' until all standard gates are above all fan-out gates.

Notice that we have not changed the result nor the depth of the circuit. Call this new circuit C'' . Notice that the standard gates of C'' all have one output, and thus correspond to a formula for f . Let F correspond to this formula, i.e., the circuit consisting of the standard gates of C'' , with the inputs corresponding to every input into a gate which is an actual input or an output from one of the fan-out gates. Since there are at most d levels of fan-out gates, and each of those gates has fan-out c , each input can be fanned out to at most c^d inputs of F . Thus there are at most nc^d inputs to F . \square

Corollary 7.1 *Any BFO circuit which computes parity requires $\Omega(\log n)$ depth.*

Proof: By Khrapchenko [42], any formula for parity must have size $\Omega(n^2)$. By the previous lemma, $nc^d = \Omega(n^2)$, and since c is a constant, $d = \Omega(\log n)$. \square

Let $TH_{k,n}$ denote the threshold function which outputs 1 if and only if at least k of the inputs are equal to 1.

Corollary 7.2 *Any BFO circuit which computes $TH_{k,n}$ requires $\Omega(\log k + \log \log n)$ depth.*

Proof: By Khrapchenko [42] any formula for $TH_{k,n}$ must have size $\Omega(k(n-k+1))$. By Krichevskii [43] any formula for $TH_{k,n}$ must have size $\Omega(n \log n)$. By the previous lemma, $nc^d = \max\{\Omega(k(n-k+1)), \Omega(n \log n)\}$, and since c is a constant, $d = \Omega(\log k + \log \log n)$. \square

We next consider the computation of multiple-valued Boolean functions.

Lemma 7.1 *Let $f : R^n \rightarrow R^m$ be a Boolean function. Consider the j th input variable for some $j, 1 \leq j \leq n$. Let O be a set of output variables with the property that for each $o \in O$ there is some n -bit input I such that the value of o is complemented when the j th bit in I is complemented. Then any bounded fan-out circuit that computes f will require depth $\Omega(\log |O|)$.*

Proof: The circuit must contain a path from the j th input node to each of the output nodes in O . Since the circuit has bounded fan-out, the lemma follows. \square

Corollary 7.3 *Any bounded fan-out circuit for adding two n -bit integers, merging a bit into an n -bit sorted sequence, sorting n bits, or computing the prefix sums of n bits requires $\Omega(\log n)$ depth.*

7.2 Upper bounds

There are well known bounded fan-in circuits with $O(\log n)$ depth and linear size for parity, addition, merging, sorting binary inputs, and prefix sums on binary inputs. By [40], these circuits can be converted into bounded fan-out circuits of the same size and depth. By Corollaries 7.1 and 7.3, these are optimal BFO circuits for these problems in terms of both size and depth.

Next we present a BFO circuit that computes the threshold function $TH_{k,n}$ in optimal size n and optimal depth $O(\log k + \log \log n)$. Our construction makes use of an optimal logarithmic depth circuit for computing (in binary) the sum of n bits [46] and two constructions for monotone formulas due to Valiant [49] and Friedman [27] which we sketch below.

The monotone formula construction of Valiant [49] shows that any monotone symmetric function on n variables can be written as a monotone formula of size $O(n^{5.3})$. Implicit in the construction of this formula is a monotone BFO circuit of size $O(n^{5.3})$ and depth $O(\log n)$.

The monotone formula construction of Friedman [27] shows that $TH_{k,n}$ can be written as a monotone formula of size $O(k^{12.6}n \log n)$. This construction uses Valiant's construction on threshold functions with $4k^2$ inputs, and thus has depth $O(\log k)$. The threshold function developed by Friedman has the form

$$TH_k(y_1, \dots, y_n) = \bigvee_{j=1}^{k^4 \log n} TH_k \left(\bigvee_{i \in A_1^j} y_i, \bigvee_{i \in A_2^j} y_i, \dots, \bigvee_{i \in A_{4k^2}^j} y_i \right),$$

where for each j , $A_1^j, \dots, A_{4k^2}^j$ is a partition of the n inputs. Thus each of the n inputs must be fanned out to $k^4 \log n$ of Valiant's threshold circuits. This can be done in $O(\log k + \log \log n)$ depth and $O(nk^4 \log n)$ size. The total size of all of the Valiant circuits are then $k^4 \log n$ times $O((4k^2)^{5.3})$.

We use these results for our circuit as follows. First we place the n inputs into groups of size $k^{15} \log n$ and use the addition circuits to find the sum of the number of ones. This takes linear size in each group, and thus linear size overall. The depth required is $O(\log k + \log \log n)$. Then using a standard comparison circuit, each group can check to see if it has more than k 1's. The output to this circuit goes into a final OR gate which determines the final outcome. Along with this, each circuit fans out each of its first $\lceil \log k \rceil$ outputs into the appropriate number of ones, so that we will have at most $2k$ outputs from each group, with the number of ones output equal to the number of ones in the group (assuming the number of ones is at most k). The outputs of all the groups can now be input into the Friedman circuit. Since there are $2kn/(k^{15} \log n)$ inputs, the size of the Friedman circuit will be $O(n)$. The depth will be $O(\log k + \log \log n)$. The output to the Friedman circuit is then ORed with the output of the comparator circuit at each of the groups. If any group had more than k inputs, then the output of the total circuit will be 1. If not, then the output from each group will be the correct number of ones in the group, and the output of the total circuit will be the output of Friedman's circuit, which will be 1 if and only if the number of ones in the input is at least k .

This circuit implies the following theorem

Theorem 7.2 *There is a size $O(n)$, depth $O(\log k + \log \log n)$ BFO circuit which computes $TH_{k,n}$.*

8 Relations between ERCW Models

We now discuss the relative powers of the different write conflict resolution protocols on the ERCW(ack) PRAM. Many of our results parallel those on the CRCW PRAM. Using the results from Section 3, some of these results can be generalized to the ERCW PRAM model and the OCPC model.

We use the notation $\text{Protocol}(m)$ to denote a collision protocol on an ERCW(ack) PRAM with m global memory cells. If we let $X \leq Y$ mean “ X conflict resolution protocol can be simulated on Y conflict resolution protocol with constant slowdown”, then it is not hard to see that

$$\begin{aligned} \text{Robust}(m) &\leq \text{Collision}(m) \leq \text{Arbitrary}(m) \leq \text{Priority}(m), \text{ and} \\ \text{Robust}(m) &\leq \text{Nice Robust}(m) \leq \text{Tolerant}(m) \leq \text{Collision}(2m). \end{aligned}$$

(The last simulation simply associates an extra memory cell with each memory cell of the Tolerant ERCW(ack) PRAM, to test whether there will be a collision at that cell, so that the value of the cell is not overwritten if there is a collision.) In addition, $\text{Common}(m) \leq \text{Arbitrary}(m)$.

In the next two subsections, we describe some less obvious simulation results.

8.1 Separations

In our lower bounds, we will always assume the simulating machine has infinite memory. Boppana [9] showed that solving Element Distinctness on the Common CRCW PRAM (and thus the Common ERCW(ack) PRAM) requires $\Omega(\log n / \log \log n)$ time. However on any of the other ERCW(ack) PRAM models, Element Distinctness can be solved in constant time. The following theorem follows.

Theorem 8.1 *There is a separation of $\Omega(\log n / \log \log n)$ between the Common ERCW(ack) PRAM model and any other ERCW(ack) PRAM model.*

This separation is tight for the CRCW PRAM, but so far the best algorithm for Element Distinctness on the Common ERCW(ack) PRAM requires $\Omega(\log n)$ time.

Similarly, Grolmusz and Ragde [34] and Chaudhuri[12] provide separations between some other CRCW PRAM models, which can be easily transferred to the ERCW(ack) PRAM, yielding the following theorem.

Theorem 8.2 *There is a separation of $\Omega(\log \log \log n)$ between the Collision ERCW(ack) PRAM and Common ERCW(ack) PRAM models, a separations of $\Omega(\log \log n)$ between the Collision ERCW(ack) PRAM model and the Arbitrary ERCW(ack) PRAM model, and a separation of $\Omega(\log \log n)$ between the Tolerant ERCW(ack) PRAM model and the Collision ERCW(ack) PRAM model.*

8.2 Simulations

First we note that the simulation of a Priority(m) CRCW PRAM on an Arbitrary(mn) CRCW PRAM, given in Chlebus *et al.* [15] can be easily transferred to the ERCW(ack) PRAM, yielding the following theorem.

Theorem 8.3 *There is a simulation of a Priority(m) ERCW(ack) PRAM on an Arbitrary(mn) ERCW(ack) PRAM that runs in $O(\log \log n)$ steps.*

The rest of this subsection describes the simulation of an Arbitrary(m) ERCW(ack) PRAM on a Tolerant(mn) ERCW(ack) PRAM.

For the simulation we need to use a partition algorithm from [15], which is run with a subset of processors, and results in either one processor being marked, or at least one but at most half of the processors being marked. This partition algorithm uses $O(n)$ memory cells and takes $O(\log \log n)$ time. An additional feature of this algorithm is that each marked processor has an associated unmarked processor, and thus if k processors are initially assigned to each processor in the subset, then the algorithm can assign $2k$ processors to each marked processor, k from the marked processor,

and k from its associated unmarked processor. (This partition algorithm was written for the Collision CRCW PRAM, but can also be run on the Tolerant ERCW(ack) PRAM.)

In the first phase of our simulation, we divide the processors into groups according to the cells they write to, and run the partition algorithm $\sqrt{\log n / \log \log n}$ times, with each subsequent application on the marked processors from the previous application. Then if there is more than a single processor remaining in a group, each of the remaining processors will be assigned $2^{\sqrt{\log n / \log \log n}}$ processors.

In the second phase of our simulation, we use a technique from [16] to choose one of the remaining (marked) processors for each cell as follows. Let $k = 2^{\sqrt{\log n / \log \log n}}$. Assign each marked processor to a cell in an n element array according to its processor number. Form a k -ary tree T over this array. For each marked processor P do the following: Associate the k processors assigned to P with the k leaves in T with the same parent as P . Now let P write a 0 to its own cell. Then let P write a 1 to its own cell, and let the auxiliary processors associated with cells at positions greater than P also write 1 to their cells. Then let P read its location. P will only read a 1 if it is the first (lowest numbered) child of its parent which is writing. Say a marked processor “wins” this level if it succeeds in writing a 1. Assume P wins this level. Then P and its k assigned processors move up to the next level. Notice that on the ERCW(ack) PRAM, the marked processors that lose can’t inform their assigned processors that they’ve lost, so those processors will also move up to the next level. But they will simply mimic the assigned processors of the winner, and this will still allow only the first child of each parent to succeed in writing a 1. We continue this procedure through the $\log n / \log k$ levels, until exactly one winning processor remains. This processor then writes without contention to the appropriate cell, completing the simulation of the Arbitrary Write step.

The time for performing $\sqrt{\log n / \log \log n}$ partitions is $O(\sqrt{\log n \log \log n})$ and the time to proceed through $\log n / \log k$ levels of the tree, each one taking constant time, is also $O(\sqrt{\log n \log \log n})$. Thus the time for simulating a step is $O(\sqrt{\log n \log \log n})$. The following theorem follows.

Theorem 8.4 *There is a simulation of a Arbitrary(m) ERCW(ack) PRAM on an Tolerant(mn) ERCW(ack) PRAM that runs in $O(\sqrt{\log n \log \log n})$ steps.*

9 Conclusions

In this paper we have studied the Exclusive Read, Concurrent Write (ERCW) PRAM model. This model is of importance since we show a tight correspondence between an ERCW PRAM with a linear number of memory locations and the Optical Communication Parallel Computer (OCPC), a model of parallel computation that uses optical communication between processors and memory.

We have also presented results for bounded fan-out circuits (BFO’s). Algorithms designed for BFO’s will map on to the OCPC without the need for fully dynamic reconfiguration.

The OCPC model has been widely studied, and most of the prior work on this model has been devoted to implementing one step of parallel communication using a small number of steps on the optical communication medium. In contrast, we show in this paper that each step of the ERCW PRAM (with number of shared memory locations equal to the number of processors) can be implemented in constant time on the OCPC and vice versa. Thus the algorithms we present in this paper can be mapped on to the OCPC so that each step takes only a constant amount of time

thus obviating the need to map a step of parallel communication on to the OCPC communication medium.

In this paper, we have presented algorithms as well as lower bounds on the time needed to solve many fundamental problems on the ERCW PRAM and the OCPC. In many cases the bounds are not tight, and it would be interesting to close the gap between upper and lower bounds.

References

- [1] A. Aggarwal, A. Bar-Noy, D. Coppersmith, R. Ramaswami, B. Schieber, M. Sudan. Efficient Routing in Optical Networks. *Journal of the ACM*, 43(6):973–1001, 1996.
- [2] I. Anderson. *Combinatorics of Finite Sets*. Oxford Science Publications, 1987.
- [3] R. J. Anderson and G. L. Miller. Optical communication for pointer based algorithms. Technical Report CRI 88-14, University of Southern California, 1988.
- [4] C. Berge. *Graphs*. Elsevier Science Publishers B.V., 1985.
- [5] P. B. Berra, A. Ghafoor, M. Guiznani, S. J. Marcinkowski, and P. A. Mitkas. Optics and supercomputing. *Proceedings of the IEEE*, 77(12):1797–1815, 1989.
- [6] P. Berthomé, Th. Duboux, T. Hagerup, I. Newman, A. Schuster. Self-Simulation for the Passive Optical Star Model. *European Symp. on Algorithms*, Lecture Notes in Comp. Sci., 979:369–380, 1995.
- [7] P. Berthomé, A. Ferreira. On Broadcasting Schemes in Restricted Optical Passive Star Systems. *DIMACS Series in Discrete Math. and Comp. Sci.*, 21:19–29, 1995.
- [8] P. C. P. Bhatt, K. Diks, T. Hagerup, V. Prasad, T. Radzik, and S. Saxena. Improved deterministic parallel integer sorting. *Inform. and Comput.*, 94(1):29–47, September 1991.
- [9] R. B. Boppana. Optimal separations between concurrent-write parallel machines. In *Proc. 21st Symp. on Theory of Computing*, pages 320–326, 1989.
- [10] R. B. Boppana and M. Sipser. The complexity of finite functions. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, chapter 14, pages 757–804. MIT Press/Elsevier, 1990.
- [11] A. K. Chandra, L. J. Stockmeyer, and U. Vishkin. A complexity theory for unbounded fan-in parallelism. In *Proc. 23th Symp. on Found. of Comp. Sci.*, pages 1–13, 1982.
- [12] S. Chaudhuri. Separating the power of CRCW PRAMs. unpublished, 1993.
- [13] B. S. Chlebus. A parallel bucket sort. *Inform. Process. Lett.*, 27(2):57–61, February 1988.
- [14] B. S. Chlebus. Parallel iterated bucket sort. *Inform. Process. Lett.*, 31(4):181–183, May 1989.
- [15] B. S. Chlebus, K. Diks, T. Hagerup, and T. Radzik. Efficient simulations between CRCW PRAMs. In *Proc. 13th Symp. on Math. Found. of Comp. Sci.*, volume 324, pages 231–239. Springer Lecture Notes in Computer Science, 1988.

- [16] B. S. Chlebus, K. Diks, T. Hagerup, and T. Radzik. New simulations between CRCW PRAMs. In *Proc. 7th Intl. Conf. on Fundamentals of Comp. Theory*, volume 380, pages 95–104. Springer Lecture Notes in Computer Science, 1989.
- [17] S. Cook, C. Dwork, and R. Reischuk. Upper and lower time bounds for parallel random access machines without simultaneous writes. *SIAM J. Comput.*, 15(1):87–97, February 1986.
- [18] M. Dietzfelbinger and F. Meyer auf der Heide. Simple, efficient shared memory simulations. In *Proc. ACM Symp. on Para. Alg. and Arch.*, pages 110–119, 1993.
- [19] R. P. Dilworth. A decomposition theorem for partially ordered sets. *Ann. Math.*, 51:161–165, 1950.
- [20] P. Dowd. High performance interprocessor communication through optical wavelength division multiple access channels. In *Proc. 18th Symp. on Comp. Arch.*, pages 96–105, 1991.
- [21] P. Erdős and R. Rado. Intersection Theorems for Systems of Sets. *J. London Math. Soc.*, 35:85–90, 1960.
- [22] M. M. Eshaghian. Parallel algorithms for image processing on omc. *IEEE Trans. Comput.*, 40(7):827–833, 1991.
- [23] F. Fich, R. Impagliazzo, B. Kapron, V. King, and M. Kutylowski. Limits on the power of parallel random access machines with weak forms of write conflict resolution. In *Proc. of 10th Symp. on Theor. Aspects of Comp. Sci.*, page unknown, 1993.
- [24] F. Fich, M. Kowaluk, M. Kutylowski, K. Loryś, and P. Ragde. Retrieval of scattered information by EREW, CREW, and CRCW PRAMs. In *Proc. 3rd Scand. Workshop on Alg. Theory*, pages 30–41. Lec. Notes in Comp. Sci., Vol. 621, 1992.
- [25] F. E. Fich, F. Meyer auf der Heide, P. Ragde, and A. Wigderson. One, two, three ...infinity: Lower bounds for parallel computation. In *Proc. 17th Symp. on Theory of Computing*, pages 48–58, 1985.
- [26] F. E. Fich, P. Ragde, and A. Wigderson. Relations between concurrent-write models of parallel computation. *SIAM J. Comput.*, 17:606–627, 1988.
- [27] J. Friedman. Construct $O(n \log n)$ size monotone formulae for the k th threshold function of n boolean variables. *SIAM J. Comput.*, 15(3):641–654, 1986.
- [28] A. V. Gerbessiotis and L. G. Valiant. Direct bulk-synchronous parallel algorithms. In *Proc. Scandinavian Workshop on Algo. Theory*, 1992.
- [29] M. Geréb-Graus and T. Tsantilas. Efficient optical communication in parallel computers. In *Proc. ACM Symp. on Para. Alg. and Arch.*, pages 41–48, 1992.
- [30] P. B. Gibbons, Y. Matias, V. Ramachandran. The Queue-Read Queue-Write PRAM model: Accounting for contention in parallel algorithms. In *Proc. ACM-SIAM Symp. on Discrete Algs.* 1994, *SIAM J Comput.*, to appear.
- [31] L. A. Goldberg, M. Jerrum, T. Leighton, and S. Rao. A doubly logarithmic communication algorithm for the completely connected optical communication parallel computer. In *Proc. ACM Symp. on Para. Alg. and Arch.*, pages 300–309, 1993.

- [32] L. A. Goldberg, M. Jerrum, and P. D. MacKenzie. A lower bound for routing on a completely connected optical communication parallel computer. accepted to SPAA, 1994.
- [33] J. W. Goodman, F. Leonberger, S. Y. Kung, and R. A. Athale. Optical interconnections for vlsi systems. *Proceedings of the IEEE*, 72(7):850–866, 1984.
- [34] V. Grolmusz and P. Ragde. Incomparability in parallel computation. In *Proc. 28th Symp. on Found. of Comp. Sci.*, pages 89–98, 1987.
- [35] I. M. I. Habbab, M. Kavehrad, and C. E. W. Sundberg. Protocols for very high-speed optical fiber local area networks using a passive star topology. *J. Lightwave Tech.*, LT-5:1782–1793, 1987.
- [36] T. Hagerup. Towards optimal parallel bucket sorting. *Inform. and Comp.*, 75:39–51, 1987.
- [37] T. Hagerup. Constant-time parallel integer sorting. In *Proc. 23rd ACM Symp. on Theory of Computing*, pages 299–306, 1991.
- [38] T. Hagerup and T. Radzik. Every robust CRCW PRAM can efficiently simulate a Priority PRAM. In *Proc. 2nd ACM Symp. on Para. Alg. and Arch.*, pages 117–124, 1990.
- [39] A. Hartmann and S. Redfield. Design sketches for optical crossbar switches intended for large-scale parallel processing applications. *Optical Eng.*, 28(4):315–327, 1989.
- [40] H. J. Hoover, M. M. Klawe, and N. J. Pippenger. Bounding fan-out in logical networks. *J. Assoc. Comput. Mach.*, 31(1):13–18, 1984.
- [41] R. M. Karp and V. Ramachandran. Parallel algorithms for shared-memory machines. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, chapter 17, pages 869–941. MIT Press/Elsevier, 1990.
- [42] V. M. Khrapchenko. A method for determining lower bounds for the complexity of Π -schemes. *Mat. Zametki*, 10(1):83–92, 1971. (in Russian); English translation in: *Math. Notes* **10**(1) (1971) 474–479.
- [43] R. E. Krichevskii. Complexity of contact circuits realizing a function of logical algebra. *Dokl. Akad. Nauk SSSR*, 151(4):803–806, 1963. (in Russian); English translation in: *Soviet Phys. Dokl.* **8**(8) (1964) 770–772.
- [44] P. D. MacKenzie, C. G. Plaxton, and R. Rajaraman. On contention resolution protocols and associated probabilistic phenomena. In *Proc. 26th ACM Symp. on Theory of Computing*, pages 153–162, 1994.
- [45] P. D. MacKenzie and Q. F. Stout. Ultra-fast expected time parallel algorithms. In *2nd ACM-SIAM Symp. on Disc. Alg.*, pages 414–423, 1991. submitted to Journal of Algorithms.
- [46] D. E. Muller and F. P. Preparata. Bounds to complexities of networks for sorting and for switching. *J. Assoc. Comput. Mach.*, 22(2):195–201, April 1975.
- [47] S. Rajasekaran and J. H. Reif. Optimal and sublogarithmic time randomized parallel sorting algorithms. *SIAM J. Comput.*, 18(3):594–607, June 1989.

- [48] A. A. Sawchuk, B. K. Jenkins, and C. S. Raghavendra. Optical crossbar networks. *IEEE Computer*, 20(6):50–60, 1987.
- [49] L. G. Valiant. Short monotone formulae for the majority function. *J. Algorithms*, 5:363–366, 1984.
- [50] L. G. Valiant. General purpose parallel architectures. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, chapter 18, pages 945–971. MIT Press/Elsevier, 1990.