Parallel Sorting With Limited Bandwidth

Micah Adler *
Computer Science Division
UC Berkeley
Berkeley, CA 94720

John W. Byers †
Computer Science Division
UC Berkeley
Berkeley, CA 94720

Richard M. Karp ‡
International Computer Science Institute
and
Computer Science Division
UC Berkeley
Berkeley, CA 94720

December 19, 1996

Abstract

We study the problem of sorting on a parallel computer with limited communication bandwidth. By using the recently proposed PRAM(m) model, where p processors communicate through a globally shared memory which can service m requests per unit time, we focus on the trade-off between the amount of local computation and the amount of inter-processor communication required for parallel sorting algorithms. We prove a lower bound of $\Omega(\frac{n \log m}{m})$ on the time to sort n numbers on the exclusive-read and queued-read variants of the PRAM(m). We show that Leighton's Columnsort can be used to give an asymptotically matching upper bound in the case where m grows as a fractional power of n. The bounds are of a surprising form, in that they have little dependence on the parameter p. This implies that attempting to distribute the workload across more processors while holding the problem size and the size of the shared memory fixed will not improve the optimal running time of sorting in this model. We also show that both the upper and the lower bound can be adapted to bridging models that address the issue of limited communication bandwidth: the LogP model and the BSP model. The lower bounds provide further convincing evidence that efficient parallel algorithms for sorting rely strongly on high communication bandwidth.

1 Introduction

A large body of theoretical research has concentrated on algorithms designed in the Parallel Random Access Machine (PRAM) model of computation. The PRAM allows processors to communicate with each other in unit time through a large globally shared memory, which leads to algorithms

*Supported by a Schlumberger Foundation Graduate Fellowship.
†Supported by a GAANN Graduate Fellowship.
‡Supported by NSF grant number CCR-9005448

that have a high degree of parallelism but perform a great deal of inter-processor communication, an inexpensive operation in the PRAM model. This leaves unresolved the question of how to design algorithms for machines which have limited inter-processor communication bandwidth.

Addressing this limitation has motivated the development of other models of parallel computation, representative of which are the BSP model [Val90a], the LogP model [CKP+93], and the PRAM(m) model [MNV94]. Provably efficient algorithms in the PRAM model are not necessarily the most efficient algorithms for these new models, so a host of problems need to be re-evaluated in this new framework. In this paper, we examine the problem of sorting in the context of parallel machines with limited communication bandwidth. We formalize the sorting problem as follows:

Definition 1 The Sorting Problem
Input: *n distinct keys $k_1 \ldots k_n$, with total order $k_{(1)} < k_{(2)} < \ldots < k_{(n)}$.*
Output at processor i: *A sorted list of keys:*

$$k_{(\frac{in}{p}+1)} \ldots k_{(\frac{(i}{p}+\frac{n}{p})}$$

We concentrate first on the complexity of sorting in the PRAM(m) model. In this variant of the classical PRAM model, p processors communicate through a globally shared memory consisting of m memory cells, and the entire input is provided to each processor in a globally shared read only memory (ROM). This model allows us to focus on the trade-off between the amount of information derived from local computation and the amount of information derived from inter-processor communication.

In the exclusive-read and queued-read variants of the PRAM(m) model (ER PRAM(m) and QR PRAM(m) respectively), we prove a lower bound for sorting n distinct keys of

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

The bound holds when $n > p^2$, which is the case of primary interest, since typical parallel applications involve problems where the input size is much larger than the number of processors. This lower bound does not rely on any restriction on the local computation of a processor, in contrast to sorting results which prove lower bounds on comparison based algorithms. The proof extends to both Monte Carlo and Las Vegas randomized algorithms, and to algorithms which employ concurrent writing.

In order to prove the lower bound, we introduce the *oracle model of computation*, a model that allows us to quantify a tradeoff between local computation and information received from other processors. In this model, which is defined more fully in Section 2.1, processors are not required to transmit any information. Rather, all inter-processor communication is simulated by an oracle that is assumed to know the entire input before computation begins. The oracle is of unlimited computational power, and thus can pre-compute any function of the inputs before computation begins. We prove that even in this setting, local computation is of such limited utility that the oracle must provide a large amount of information in order to enable the processors to solve the sorting problem efficiently.

When $m = O(n^\beta)$, for some $\beta < 1$, we show that a version of Columnsort [Lei85] has a running time that is bounded by

$$O\left(\frac{n}{m} \log n (1-\beta)^{-3.42}\right).$$

For $n \gg m$, the case of greatest interest, the final factor becomes a small constant and so in this setting the ratio between the upper and lower bounds is $\Theta(\frac{\log n}{\log m})$. The algorithm used in the upper bound runs with only a small constant factor slowdown on a machine model in which the input is distributed among all processors, and for which there is no globally shared ROM. This gives the algorithm more credibility from a practical standpoint.

We also show that our results can be generalized to two other models that incorporate limited communication throughput, the LogP model and the BSP model. In both of these models, the processors communicate using point-to-point messages, and a parameter g represents the minimum number of cycles between transmission of successive messages from a processor. To prove a lower bound for the ER PRAM(m) model, we show a lower bound on the number of bits which must be transmitted through the network in order to solve the sorting problem efficiently. Coupling this bound with model dependent lower bounds on the amount of time required to transmit a fixed number of bits in the BSP and LogP models, results in lower bounds for sorting in these two models. We defer definitions of the models and the exact form of the bounds to the section where those results are discussed.

Our results show that fast parallel algorithms to solve the sorting problem must rely on large amounts of communication. Furthermore, we have the surprising result that both our upper and lower bounds are unaffected by attempting to distribute the work across an unlimited number of processors, while holding fixed the problem size, the size of the shared memory, and the number of processors that actually output the result. Therefore, to increase the speed of parallel sorting on a machine with limited communication bandwidth, increasing bandwidth is more likely to improve the running time than increasing the number or computational power of processors.

The remainder of the paper is organized as follows. In the rest of Section 1, we briefly compare our results with previous work in parallel sorting. In Section 2, we provide a complete description of both the PRAM(m) model, and our lower bound tool, the oracle model of communication. Sections 3 and 4 provide proofs of our lower and bound for deterministic and randomized algorithms for sorting in the ER PRAM(m) and QR PRAM(m) models of computation. Section 5 provides the matching upper bound, and Section 6 briefly describes extensions of those proofs to the LogP and BSP models.

1.1 Previous Work

From the large body of research in the realm of parallel sorting algorithms, we discuss a few results which also focus on inter-processor communication requirements. Using Thompson's VLSI model [Tho80], Leighton, in [Lei85], proves a lower bound of $AT^2 = \Omega(n^2 \log^2 n)$ for sorting n keys of size $\Theta(\log n)$, where A is the area of a VLSI chip and T is the running time of the chip. His methods can be used to show bounds of the form $\Omega\left(\frac{n \log n}{m}\right)$ in a PRAM model with a globally shared memory of size m, but in which the input is evenly distributed across the p processors, rather than stored in a globally shared ROM. Indeed, an interesting question would be to determine whether we could apply our lower bound technique to a non-standard VLSI model in which the chip could receive each input in more than one location and at more than one time.

Other related work on parallel sorting includes [BC82], where Borodin and Cook prove that sorting requires $\text{TIME} \cdot \text{SPACE} = \Omega(\frac{n^2}{\log n})$. Aggarwal, Chandra and Snir show in [ACS90] that any parallel comparison-based algorithm that sorts n words requires $\Omega(\frac{n \log n}{p \log(\frac{n}{p})})$ communication steps. Also, the

same authors show in [ACS89] that sorting requires $\Omega(\frac{n \log n}{p} + l \log p)$ in a model where reading or writing a block of size b from memory takes time $l + b$.

When introducing the PRAM(m) model in [MNV94], Mansour, Nisan and Vishkin prove a lower bound of $\Omega(\frac{n}{\sqrt{mp}})$ for several problems, including sorting, in a concurrent read version of the PRAM(m), which implies the same bound in the ER PRAM(m). An easy upper bound can be obtained by using a variant of Cole's parallel merge sort [Col88] for the PRAM, which uses $O(n \log n)$ bits of memory and runs in $O(\log n)$ time. By letting each of the m bits of shared memory simulate $\frac{n \log n}{m}$ bits of the PRAM memory, we can run Cole's algorithm in time $O(\frac{n \log^2 n}{m})$ on the ER PRAM(m). Thus, the ratio between upper and lower bounds for sorting on the ER PRAM(m) was $\Theta(\sqrt{\frac{p}{m}} \log^2 n)$ prior to this work.

Subsequent to a preliminary version of this paper in [ABK95], Adler [Adl96] provides an algorithm for sorting in the CR PRAM(m) that is considerably faster than the lower bound for the ER PRAM(m) presented in this paper. Thus, that result together with the lower bound presented here imply that the CR PRAM(m) is strictly more powerful than the ER PRAM(m). [Adl96] also slightly improves the ER PRAM(m) upper bound to $O(\frac{n \log p}{m})$.

Related work on upper bounds includes [CS92], in which Cypher and Sanz allude to a recursive version of Columnsort, and introduce Cubesort, which can be used to obtain a running time of $O(\frac{n \log n}{m}(1-\beta)^{-2})25^{\log^* n - \log^*(n/m)}$ for sorting on the PRAM(m). However, this algorithm has substantial overhead and is considerably more involved then the one presented in this paper.

Finally, with respect to BSP algorithms for sorting, Gerbessiotis and Valiant [GV94] introduce a randomized algorithm for parallel sorting in the BSP model. Also, subsequent to an earlier version of this paper, the upper bound for sorting in the BSP model has been improved by both Goodrich [Goo96], as well as by Gerbessiotis and Siniolakis [GS96]. The form of these bounds is deferred to Section 5.2, where they are described in the context of our description of the BSP model.

2 The PRAM(m) Model

In this section, we define the PRAM(m) model, and then describe a theoretical tool derived from the PRAM(m) model, the oracle model of communication.

In a classical PRAM, p processors communicate by writing to and reading from a large globally shared memory in unit time. Access to the shared memories is slowed by such factors as long message send overheads [CKP+93], contention at memory banks, the fact that memory banks are much slower than processors [**?**], and bandwidth limitations of the network connecting processors to memory banks. In practice, the available per-processor bandwidth to shared memory is quite small. Similar difficulties exist in distributed memory parallel machines. The parameter m of the PRAM(m) model focuses attention on this bottleneck, by enforcing the condition that the shared memory can service only m requests per unit time, where $m < p$. We model this as a PRAM consisting of m shared memory cells, each consisting of 1 bit. We note that all the results in this paper can be extended to a model where each memory cell consists of a word of w bits.

The input of size n is provided to the PRAM(m) in a read-only shared memory (ROM) concurrently available to all of the processors. Conceptually, this is equivalent to having each processor begin with an identical copy of the input in its local memory. This restriction serves to concentrate our

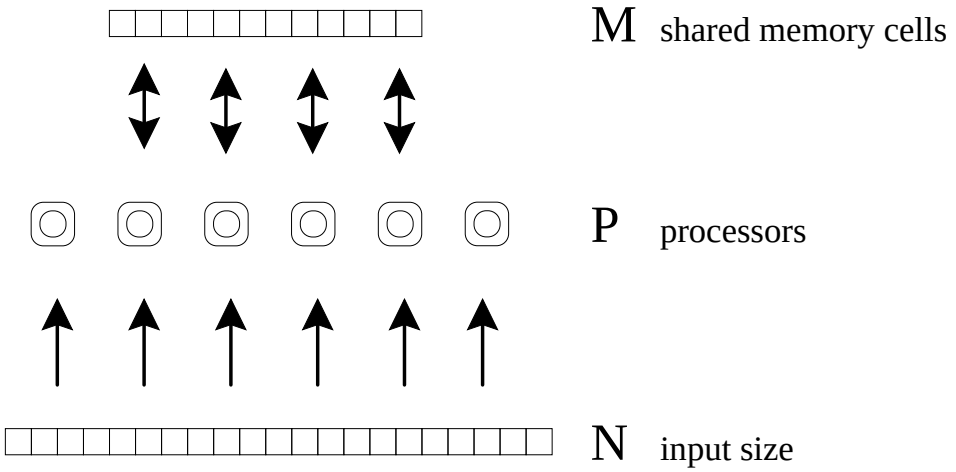

Figure 1: The PRAM(m) Model

lower bound efforts on the amount of communication required for actual computation, rather than the amount required to distribute the input. During each synchronized round of computation, every processor can perform one of four actions: it can read the contents of a ROM location, read the contents of a globally shared memory location, write to a globally shared memory location, or perform local computation.

As defined in [MNV94], the PRAM(m) model allows processors concurrent read, concurrent write access to the globally shared memory. In this paper, we consider exclusive read and queued read variants of the PRAM(m) model. In the ER PRAM(m), distinct processors i and j are forbidden from reading the same memory cell at the same time step. We also define the QR PRAM(m), where read contention at a memory cell is resolved as follows: each step of an algorithm completes in k time steps, where k is the maximum number of processors reading the same memory location. Finally, we define the asynchronous QR PRAM(m), where every memory cell services one request if any requests to that cell are pending, and all other requests are stored in a FIFO queue. We note that both queued contention resolution strategies are analogous to those devised for the standard PRAM [GMR94a]GMR94b. The contention resolution strategy for write access to the shared memory can be any of concurrent write, queued write or exclusive write. Our upper and lower bounds are not affected by this choice, and thus, we leave this component of the model unspecified.

The primary goal of this model is to examine the effectiveness of parallel computation given a sharp limitation on inter-processor communication. Since such a ROM may be unrealistic from a practical standpoint, upper bounds achieved in this model that rely on use of the ROM are only applicable to problems in which the entire input is initially known by all the processors.

2.1 The Oracle Model of Communication

It is often the case in parallel computing that the amount of computation required by a processor is greatly reduced by receiving results of computations performed by other processors. In order to quantify a tradeoff between local computation and information received from other processors, we define the oracle model of computation. This lower bound tool uses the principle that the combined information a processor receives from all other processors is no more useful than the information it can receive from a single processor with unlimited computational resources and access to all the

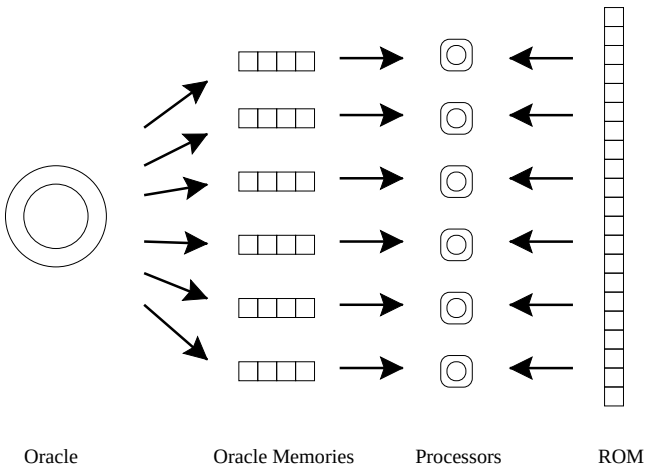

Figure 2: The flow of information in the oracle model

information the processors have.

In the oracle model, processors do not transmit any information. Rather, each processor only receives information from an oracle of unlimited computational power, and a read only memory (ROM) that contains the input. The oracle transmits information to the processors through p *oracle memories*. Processor i has read-only access to the ith oracle memory and not to any of the other oracle memories. We subject the oracle to the restriction that it compute and set all the values of the oracle memory before the processors begin computation. This restriction does not alter the power of the model; it only serves to simplify the analysis.

The processors access the input using a concurrently-readable ROM, which is identical to the ROM of the PRAM(m) model. During computation, at each synchronous time step every processor is allowed to perform one of three actions: it can read a bit from its oracle memory, read an input from the ROM, or perform local computation. The oracle knows the entire input and the programs executed by each of the processors. We are interested in the tradeoff between the maximum number of time steps required by any processor and the total number of bits read from the oracle memory.

We derive lower bounds in the oracle model for the number of bits written to the oracle memories. These bounds, combined with the restriction that at most m bits are read during any step of a PRAM(m) algorithm, provide lower bounds for the PRAM(m) model. More formally, consider algorithms A_e and A_o designed for the ER PRAM(m) and the oracle model, respectively. Let $r(A_e, i, p)$ denote the aggregate number of reads the p processors perform from the shared memory and the oracle on input i. Likewise, let $t(A_e, i, p)$ denote the number of time steps A_e runs on i with p processors. Define $r(A_o, i, p)$ and $t(A_o, i, p)$ similarly for algorithm A_o.

Definition 2 *For any positive integer c, an oracle algorithm A_o* exactly simulates *a PRAM(m) algorithm A_e if for all values of p and on all inputs i, $r(A_o, i, p) = r(A_e, i, p)$, $t(A_o, i, p) = t(A_e, i, p)$, and A_o computes the same output as A_e.*

Lemma 3 *Given any ER PRAM(m) algorithm A_e, there is an oracle algorithm A_o such that A_o exactly simulates A_e.*

Proof: Consider the execution of the ER PRAM(m) algorithm A_e on input i. Let $w(u, v)$ denote

the value of the bit processor u would read at time v. The oracle can compute $w(u,v)$ for all u,v instantaneously in advance of the simulation by using its unlimited resources. To perform the simulation, the oracle simply furnishes $w(u,v)$ in oracle memory u at time v for all u,v. The processors then execute their ER PRAM(m) algorithms, ignoring all write operations and reading from their oracle memory in place of reading bits from shared memory locations. The simulation has zero slowdown and the processors read the same total number of bits in both executions. ■

Note that the oracle model can also exactly simulate any QR PRAM(m) algorithm, as well as any asynchronous QR PRAM(m) model. Furthermore, a similar lemma shows that an oracle model in which processors have concurrent read access to a single, m bit oracle memory can exactly simulate any CR PRAM(m) algorithm, but proving strong lower bounds for sorting in this model remains open. Another use of the oracle model is to prove lower bounds about randomized algorithms. To allow the oracle to simulate the programs of processors in such a setting, we give the oracle access, at the start of an algorithm, to the random bits used by each processor.

3 Lower Bounds

In this section, we prove lower bounds for sorting algorithms in the oracle model by showing that even when all processors know the range of keys that need to be output by each processor, the task of locating those keys within the input is difficult. If we wish to sort n keys, and if all processors know the range of key values that will be output by each processor, but not the value of these keys, nor their location within the ROM, then the work remaining can be formalized as the *permutation routing problem*. We show that any lower bound for this problem implies an identical lower bound for sorting.

Definition 4 The Permutation Routing Problem
Input: *n memory locations, each containing a processor id, such that each processor id appears exactly $\frac{n}{p}$ times.*
Output at processor i: *a list of the locations where i appears.*

Lemma 5 *Any lower bound on the time to solve the permutation routing problem implies an identical lower bound for the sorting problem.*

Proof: It suffices to show that any algorithm that sorts n distinct keys can solve any instantiation of the permutation routing problem of size n. To do this, we derive from each location a key to be sorted, where the key is the concatenation of the processor id stored at that location and the location index within the ROM. Sorting these keys is sufficient to inform each processor i of the locations where i appears. ■

In order to prove our lower bounds for the permutation routing problem in different scenarios, we first prove a lower bound on the following problem in the oracle model, called the *processor d permutation routing problem*, defined as follows, where d is any processor id between 1 and p. The input is chosen uniformly at random from the set of all possible n element inputs to the permutation routing problem. Processor d is required to determine the list of locations in the ROM where d appears, but the remaining processors are not required to do anything. Processor d is allowed to perform $\frac{n \log m}{m}$ ROM queries. We are interested in the average number of oracle bits processor d

receives, where this average is taken over all possible n element inputs to the permutation routing problem.

Lemma 6 *In any deterministic algorithm solving the processor d permutation routing problem, where $n \geq p^2$, if the average over all inputs of the number of oracle bits received by processor d is at most $\frac{n \log m}{8p}$, and processor d never performs more than $\frac{n \log m}{m}$ ROM queries, then on an input chosen uniformly at random, processor d produces an incorrect result with probability at least $\frac{1}{4}$.*

The proof of this lemma is deferred to Section 3.1. We first discuss its implications.

Theorem 7 *For any deterministic ER PRAM(m) algorithm that solves the sorting problem where $n \geq p^2$, the average over all inputs of the time required by the algorithm is at least $\frac{n \log m}{8m}$. The same is true for any deterministic algorithm for the QR PRAM(m).*

Proof: We assume there is an algorithm A for the ER PRAM(m), where the average over all inputs of the time to perform A is at most $\frac{n \log m}{8m}$, and reach a contradiction. Let $b(A, C)$ be the total number of bits read from the shared memory when any algorithm A is executed on input C. Let $r(A, C, i)$ be the number of ROM locations read by processor i when A is executed on input C. The average of $b(A, C)$ over all inputs C is at most $\frac{n \log m}{8}$. Also, the average over all C of $\max_i r(A, C, i) \leq \frac{n \log m}{8m}$. By Lemmas 3 and 5, this implies the existence of an oracle algorithm A' for the permutation routing problem, where the average over all inputs C of $b(A', C) \leq \frac{n \log m}{8}$, and the average over all inputs of $\max_i r(A, C, i) \leq \frac{n \log m}{8m}$.

For any such A', there is some processor d such that the average over all inputs, of the number of oracle bits written to the oracle memory for processor d, is at most $\frac{n \log m}{8p}$, and the average over all inputs of $r(A', C, d)$ is at most $\frac{n \log m}{8m}$. By Markov's inequality, in A', the fraction of inputs C where $r(A', C, d) \geq \frac{n \log m}{m}$ ROM queries is at most $\frac{1}{8}$. Since algorithm A' is required to respond correctly on all inputs, even if processor d were restricted to only $\frac{n \log m}{m}$ ROM queries, processor d would still respond correctly on at least $\frac{7}{8}$ of all possible inputs. This contradicts Lemma 6, and thus there does not exist such an algorithm A for the ER PRAM(m). ■

For Las Vegas algorithms, or randomized strategies which are guaranteed to provide a correct solution with a bound only on the expected running time, we have a lower bound which follows from a direct application of Yao's lemma [Yao77].

Theorem 8 *For any Las Vegas algorithm A_v for the ER PRAM(m) where $n \geq p^2$, there is some input I for the sorting problem, such that A_v requires expected time at least $\frac{n \log m}{8m}$ to solve I. Also, the expected running time of any Las Vegas algorithm on an input chosen uniformly at random from the set of all inputs is at least $\frac{n \log m}{8m}$.*

Proof: Yao's lemma [Yao77] states that a lower bound of L for the expected running time of any deterministic algorithm acting on some distribution over the inputs implies that for any randomized algorithm there exists an input for which the expected running time is at least L. This combined with Theorem 7 directly implies the first claim of Theorem 8. The second claim follows from

Theorem 7 and the fact that any Las Vegas algorithm is actually a distribution over deterministic algorithms, and thus cannot fare better than the best deterministic algorithm. ■

For Monte Carlo algorithms, or randomized strategies with bounded running time which provide a correct solution with probability greater than $\frac{3}{4}$, we have the following lower bound.

Theorem 9 *For any Monte Carlo algorithm A_m for the ER PRAM(m) where $n \geq p^2$, there is some input I for the sorting problem, such that A_m requires time at least $\frac{n \log m}{8m}$ to solve I. Also, for the uniform distribution over all possible inputs, the running time of any Monte Carlo algorithm is at least $\frac{n \log m}{8m}$.*

We prove this theorem using an alternate formulation of Yao's lemma, provided in [Mac95]:

Lemma 10 *Let P_1 be the success probability of a T step randomized algorithm solving problem B, where the success probability is taken over the random choices made by the algorithm, and minimized over all possible inputs. Let P_2 be the success probability over a distribution $\mathcal{D}$ of inputs, maximized over all possible T step deterministic algorithms to solve B. Then, $P_1 \leq P_2$.*

Theorem 9 follows from a direct application of Lemma 10 to the following lemma.

Lemma 11 *For any deterministic ER PRAM(m) algorithm for the permutation routing problem, where $n > p^2$, that never uses more than $\frac{n \log m}{8m}$ time steps, when the input is chosen uniformly at random from the set of all possible inputs, the probability that every processor successfully produces the correct output is $\leq \frac{3}{4}$.*

Proof: We assume that there is such a deterministic algorithm A_d which produces the correct output with probability $> \frac{3}{4}$, and reach a contradiction. In such an algorithm, for every input C, $b(A_d, C) \leq \frac{n \log m}{8}$ and $\max_i r(A_d, C, i) \leq \frac{n \log m}{8m}$. By Lemmas 3 and 5, this implies the existence of an oracle algorithm A'_d for the permutation routing problem, where the total number of oracle bits is at most $\frac{n \log m}{8}$, and $\max_i r(A'_d, C, i) \leq \frac{n \log m}{8m}$. For any such A', there is some processor d such that the average over all inputs, of the number of oracle bits written to the oracle memory for processor d, is at most $\frac{n \log m}{8p}$, and $r(A'_d, C, d) \leq \frac{n \log m}{8m}$. However, this implies the existence of an algorithm for the processor d permutation routing problem, where the average over all inputs of the number of oracle bits written to the oracle memory for processor d is at most $\frac{n \log m}{8p}$, $r(A'_d, C, d) \leq \frac{n \log m}{8m}$, and yet processor d responds correctly with probability $> \frac{3}{4}$. This contradicts Lemma 6, and thus there does not exist such an algorithm A_d. ■

3.1 The Processor d Permutation Routing Problem

In this subsection, we prove Lemma 6, restated here for convenience.

Lemma 6 *In any deterministic algorithm solving the processor d permutation routing problem where $n > p^2$, if the average number of oracle bits received by processor d is at most $\frac{n \log m}{8p}$, and processor*

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 1 & 0 & 0 \end{bmatrix}$$

Figure 3: Permutation routing problem input: $n = 6$, $p = 3$.

d never performs more than $\frac{n \log m}{m}$ ROM queries, then on an input chosen uniformly at random, processor d produces an incorrect result with probability at least $\frac{1}{4}$.

For concreteness and simplicity, we represent an input to the permutation routing problem as a bit matrix B with n rows and p columns. If j appears in location i in the permutation routing problem instance, then $B_{ij} = 1$, otherwise $B_{ij} = 0$. Rows of B correspond to locations, and columns of B correspond to processors, so each column of B has exactly $\frac{n}{p}$ ones, and there is exactly one 1 in each row. A ROM query to location k reveals row k of this matrix. In order to correctly solve the processor d permutation routing problem, processor d must be able to specify column d of B exactly.

The proof of this lower bound employs the "little birdie" principle: giving a processor additional information never increases the complexity of the problem that the processor must solve. The side information that processor d is assumed to know a priori is a perturbed representation of the input specification B called a *hidden matrix*. Providing this matrix to processor d allows us to prove lower bounds on the amount of information subsequent ROM queries provide to processor d.

We examine algorithms that proceed as follows. A permutation routing problem instance B is chosen uniformly at random, and revealed to the oracle. Then, based on the choice of B, a hidden matrix H is chosen and revealed to processor d and the oracle. Based on B and H, the oracle provides processor d with some number of oracle bits. Processor d then executes its deterministic algorithm, and produces its output.

To construct the hidden matrix H, we first choose an $n \times p$ matrix G uniformly at random from binary matrices whose d^{th} column is identical to the d^{th} column of B, and the remaining columns each have exactly $\frac{n}{p(p-1)}$ 1's, such that every row with a 1 in column d has exactly one 1 in some other column, and all other rows have no 1's. The hidden matrix H is defined to be $H = B \oplus G$, where $\oplus$ denotes the bitwise XOR of the two matrices. This mapping evenly redistributes the 1's from column d of B across the other columns while leaving all other rows unchanged. A pictorial representation of an input matrix and a possible hidden matrix constructed from it are given in Figures 3 and 4.

We say that an input C is *consistent* with a hidden matrix H if there is a matrix G perturbing C as defined above such that $C \oplus G = H$. Let $\mathcal{C}(H)$ denote the set of inputs that are consistent with hidden matrix H. Since we are discussing deterministic algorithms, an input C coupled with a consistent hidden matrix H uniquely determines S, the setting of oracle memory d. We say that S is consistent with hidden matrix H, if there is $C \in \mathcal{C}(H)$ such that H and C determine S. Let

$$\begin{bmatrix} 0 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \\ 0 & 1 & 0 \end{bmatrix}$$

Figure 4: A hidden matrix for processor 1

$\mathcal{S}(H)$ be the set of all settings of the oracle bits S that are consistent with H. Also, we say that input C is consistent with both H and S if H and C determine S. Let $\mathcal{C}(H,S)$ be the set of all inputs C that are consistent with H and S. Finally, for a hidden matrix H and an input $C \in \mathcal{C}(H)$, let the indicator variable $R(H,C) = 1$ if processor d produces the correct output on C when given H, and 0 otherwise.

Claim 12 *In any algorithm for the processor d permutation routing problem, where processor d is provided with any hidden matrix H and oracle bits S, and processor d performs no more than $\frac{n \log m}{m}$ ROM queries,*

$$\sum_{C \in \mathcal{C}(H,S)} R(H,C) \leq Z, \text{ where } Z = \sum_{r=1}^{\frac{n}{p}} \binom{\frac{n \log m}{m}}{r}$$

Proof: After receiving the oracle bits and the hidden matrix, we can model processor i's actions by a decision tree, in which each node of the decision tree corresponds to a ROM query, and each leaf of the tree corresponds to a processor state achievable after performing at most $\frac{n \log m}{m}$ ROM queries. The number of distinct results that processor d can produce is at most the number of leaves in this tree. We show that for any non-redundant algorithm, i.e. an algorithm that only examines each ROM location once, the tree has at most Z leaves. Any redundant algorithm can be simulated by a non-redundant algorithm, and thus the number of distinct results processor d can produce is at most Z for all algorithms.

Suppose processor d reads the value of row i of C from the ROM. Then, either (a) row i of H is identical to row i of C, or (b) C has a 1 in column d of row i whereas H has a 1 in some other column. Since processor d knows H at the start of the algorithm, the decision tree has branching factor two. Since processor id d only appears in $\frac{n}{p}$ elements, at most $\frac{n}{p}$ of these ROM queries can discover elements where processor id d appears. Mapping successful discoveries to left branches and unsuccessful queries to right branches ensures that any algorithm which is non-redundant has a decision tree where any path from root to leaf can have at most $\frac{n}{p}$ left branches. The number of distinct leaves of the decision tree that can be reached by a path from root to leaf with k left branches is at most $\binom{\frac{n \log m}{m}}{k}$. Thus, the possible number of leaves in the decision tree is at most $\binom{\frac{n \log m}{m}}{0} + \binom{\frac{n \log m}{m}}{1} + \ldots + \binom{\frac{n \log m}{m}}{n/p}$. ■

For any algorithm for the processor d permutation routing problem, let $a(H)$ be the average number of oracle bits provided to processor d, where the average is taken over all inputs in $C(H)$.

Claim 13 *Consider any algorithm for the processor d permutation routing problem where $n > p^2$, where the little birdie provides processor d with a hidden matrix and where processor d performs at most $\frac{n \log m}{m}$ ROM queries. For any H, if $a(H) \leq \frac{n \log m}{4p}$, then on an input chosen uniformly at random from the set $C(H)$, the probability that processor d produces the correct output is $\leq \frac{1}{2}$.*

Proof: The total number of inputs in $C(H)$ on which processor d responds correctly is at most:

$$\sum_{C \in C(H)} R(H,C) = \sum_{S \in S(H)} \sum_{C \in C(H,S)} R(H,C),$$

and so the probability of a successful response is at most

$$\frac{1}{|C(H)|} \sum_{S \in S(H)} \sum_{C \in C(H,S)} R(H,C).$$

But, by Claim 12, this is at most Y, where

$$Y = \sum_{S \in S(H)} \min \left(\frac{Z}{|C(H)|}, \frac{|C(H,S)|}{|C(H)} \right).$$

We assume that Y $\geq \frac{1}{2}$, and reach a contradiction by showing that the average number of oracle bits written to oracle memory d is too high. Let $|S|$ denote the number of bits in oracle memory setting S. We minimize the average number of oracle bits used in conjunction with matrix H:

$$a(H) = \sum_{S \in S(H)} |S| \frac{|C(H,S)|}{|C(H)|},$$

subject to $Y \geq \frac{1}{2}$. First, note that $Y \geq \frac{1}{2}$ implies that $|S \in S(H)| \geq \frac{|C(H)|}{2Z}$. Next, we can assume that for any $S \in S(H)$, $|C(H,S)| \geq Z$, since given any valid solution, any sets $C(H,S)$ of smaller cardinality can be combined into sets of cardinality Z without changing the value of Y, and without increasing $a(H)$. Because there are at most 2^r settings of the oracle bits with $\leq r$ bits, there are at least $\frac{3|C(H)|}{8Z}$ oracle memory settings $S \in S(H)$ such that

$$|S| \geq \log \left(\frac{|C(H)|}{2Z} \right) - 1.$$

However, by counting the total contribution to $a(H)$ of these sets, we see that

$$a(H) \geq \frac{3}{8} \log \left(\frac{|C(H)|}{4Z} \right).$$

Using the fact that $n \geq p^2$, we have that for any j, $|C(H)|$ is

$$\left(\frac{\frac{n}{p-1}}{\frac{n}{p(p-1)}} \right)^{p-1}.$$

Using the inequality

$$\left(\frac{a}{b} \right)^b \leq \binom{a}{b} \leq \left(\frac{ae}{b} \right)^b,$$

and the fact $m \leq p$ implies that the sum expressed by Z is dominated by the final term, this gives us

$$a(H) \geq \frac{3n}{8p} \log m - o\left(\frac{n}{p} \log m\right),$$

which is a contradiction. Thus, the probability of a correct response from processor d is at most $\frac{1}{2}$. ■

Proof: (of Lemma 6) Since the number of hidden matrices consistent with a given input is invariant over the choice of input and the number of inputs consistent with a given hidden matrix is invariant over the choice of matrix, the average of $a(H)$ over all H is at most $\frac{n \log m}{8p}$. By Markov's inequality, $a(H) \leq \frac{n \log m}{4p}$ for at least $\frac{1}{2}$ of the hidden matrices H. We choose the input to the permutation routing problem as follows: first a hidden matrix H is chosen uniformly at random from the set of all hidden matrices that can be given to d. Then, an input C_i is chosen uniformly at random from $C(H)$. This roundabout means for choosing an input is equivalent to choosing an input uniformly at random from the set of all inputs. With probability at least $\frac{1}{2}$, $a(H) \leq \frac{n \log m}{4p}$. Given this, the probability that processor d responds correctly is at most $\frac{1}{2}$. Therefore, the probability that processor d responds correctly on an input chosen uniformly at random, when given the matrix H, is at most $\frac{3}{4}$. By the little birdie principle, the probability that processor d responds correctly on an input chosen uniformly at random, when not given the matrix H, is also at most $\frac{3}{4}$. ■

4 The Upper Bound

We show that a version of Leighton's Columnsort [Lei85] performs well in the ER PRAM(m). If w is the number of bits per key, then by introducing an additional factor of $\frac{w}{\log n}$ to the running time, this algorithm can be converted to one where all processors read only from distinct portions of the ROM of size $\frac{n}{p}$. The resulting algorithm can be run in a model where there is no ROM for the input, and the input is distributed across the processors' local memories.

Theorem 14 *There is an ER PRAM(m) algorithm for sorting n keys which runs in time $O(\frac{n \log n}{m}(1-\beta)^{-3.42}))$, provided $p \geq m$ and m grows no faster than some fixed function of n of the form n^β, $\beta < 1$.*

Proof: To provide a sorting algorithm for the ER PRAM(m), we use a recursive version of Columnsort, which we describe below. In Columnsort, the n keys are thought of as elements in a matrix M. The elements are sorted using seven phases, where each phase is one of three types: phases that sort the columns of the matrix, phases that perform an odd-even transposition sort along the rows of the matrix and phases that route a fixed permutation of the matrix elements, where each column routes an equal number of elements to every other column. The following simple description of Columnsort is provided in [Lei92]. In phases 1, 3, and 7, the columns are sorted into increasing order. In phase 5, odd columns are sorted into increasing order and even columns are sorted into decreasing order. In phase 2, the matrix is "transposed": the items are picked up in column major order and set down in row-major order (preserving the shape of the matrix). Phase 4 applies the reverse of the permutation applied in phase 2, and phase 6 performs two steps of odd-even transposition sort to each row. There is an additional requirement on the aspect ratio of M: if M is an $s \times r$ matrix, then s must be larger than r^2.

In the recursive version of Columnsort we provide here, each of the m memory cells is paired with a single processor. These m processors sort the n keys, and inform each of the $p - m$ remaining processors of the range of keys that they need to output. Note that since the algorithm makes effective use of only m processors, this algorithm is consistent with the observation that increasing communication throughput, as opposed to adding processors, is required for faster parallel sorting. We specify a call to recursive Columnsort in the ER PRAM(m) by two parameters: k, the number of keys to be sorted and a, the number of memory cell-processor pairs that are dedicated to this function call. We develop the following recurrence relation for the running time of recursive Columnsort, where the keys are contained in a ROM of size n:

$$\begin{aligned} RC(k,a) &= \mathrm{O}\left(\frac{k \log n}{a}\right), \text{ if } k \geq a^3 \\ RC(k,a) &= 4\, RC\left(k^{2/3}, ak^{-1/3}\right) + \mathrm{O}\left(\frac{k \log n}{a}\right), \text{ when } k > a. \end{aligned}$$

The base case for the recurrence, where $k \geq a^3$, works as follows. The keys are thought of as being entries in a matrix of keys of size $\frac{k}{a} \times a$; this matrix satisfies the aspect ratio requirement of Columnsort. Since we have a available memory cell-processor pairs at our disposal, we assign one pair to each column of $\frac{k}{a}$ keys of the matrix. We now show that we can implement each of the three types of phases of Columnsort on the PRAM(m) in time $\mathrm{O}(\frac{k \log n}{a})$. Sorting the columns can be performed by each of the a processors locally in time $O(\frac{k \log k}{a})$ by any of a variety of known serial algorithms.

Routing the fixed permutation on the matrix elements requires each processor to send an identical number of keys to every other processor, and thus can be done with a single pass through all the entries. A single element is routed by the source and destination processors using the shared memory location that is paired with the destination processor for $\log n$ consecutive steps, during which the key's address in the ROM is written to that memory location by the source processor, one bit at a time. Thus, the entire permutation can be routed in time $O(\frac{k \log n}{a})$. One phase of odd-even transposition sort can be performed by each of the a processors routing all the keys currently in their column to the processor that is assigned to the neighboring column. This can also be done in time $O(\frac{k \log n}{a})$.

For the case where $k < a^3$, we sort the keys as follows. The keys are thought of as the elements in a matrix M of size $k^{2/3} \times k^{1/3}$. This matrix satisfies the aspect ratio of Columnsort, and thus we use Columnsort to sort this matrix as well. We can still route the permutations of the matrix M and perform the phases of odd-even transposition sort in time $\mathrm{O}(\frac{k \log n}{a})$. This follows from the fact that for each permutation only k keys need to be routed through the shared memory and this operation can be performed without conflict while making use of each cell during each of the $\mathrm{O}(\frac{k \log n}{a})$ time steps. For the phases which sort the columns of the matrix, we employ parallel calls to recursive Columnsort, one call per column. Each column consists of $k^{2/3}$ keys, and we evenly distribute the a available memory cell-processor pairs across the columns. Thus, each of the 4 sorting phases takes time RC $\left(k^{2/3}, ak^{-1/3}\right)$. An example which graphically describes one level of this recursive procedure is given in Figure 5.

We start with a call to recursive Columnsort with n keys and m memory cell-processor pairs. After the n keys are sorted, the m processors can inform each of the remaining $p - m$ processors of the correct sorted list of $\frac{n}{p}$ keys to output in the same amount of time as is required to route one permutation. To analyze the running time of sorting n keys on the ER PRAM(m), we evaluate the

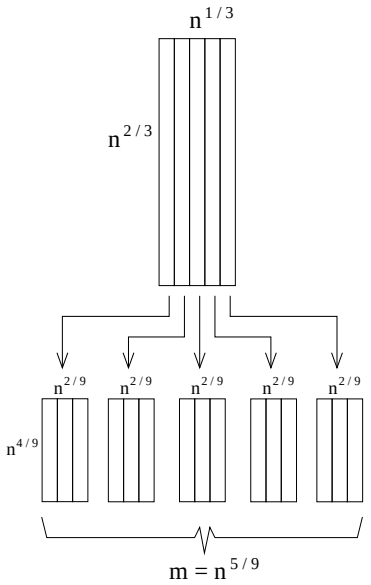

Figure 5: The recursive algorithm when $m = n^{5/9}$

recurrence RC (n, m). The running time of the algorithm is dominated by the time for sorting at the bottom of the recursion at level j: $O(\frac{n \log n}{m} 4^j)$. When m is a fixed function of n where $m = O(n^\beta)$ for $\beta < 1$, this j is the smallest integer that satisfies $n^{((2/3)^j)} \leq n^{1-\beta}$, and the following bound on 4^j follows directly:

$$4^j \leq (1-\beta)^{\frac{2}{\log \frac{2}{3}}} \approx (1-\beta)^{-3.42}$$

Substituting in to the formula above gives an upper bound on the running time of $O(\frac{n \log n}{m}(1-\beta)^{-3.42})$. ■

As m approaches n, this function becomes much larger than $\frac{n \log n}{m}$. But this is as expected, given the known lower bound for sorting of $\Omega(\log n)$ on any n processor PRAM, and the fact that our bound of the form $O(\frac{n \log n}{m})$ in the bit model becomes a bound of the form $O(\frac{n}{m})$ in a "word" model where each memory cell can contain one key.

5 Other Limited Bandwidth Models

We can use the techniques discussed for the PRAM(m) to derive bounds in other parallel models which address the issue of limited communication throughput. We give a brief discussion of the complexity of sorting in the LogP model [CKP+93] and the BSP model [Val90a].

5.1 The LogP model

In the LogP model, limited communication communication throughput in a parallel machine is enforced by requiring that each processor must wait for a gap of at least g cycles between the transmission of consecutive point-to-point messages. The three other LogP parameters are P, the number of processors, L, the latency of a message in the network, and o, the overhead (in cycles) to place a fixed-size message onto the network. Note that this model uses point-to-point messages

for communication, as opposed to the global shared memory used in the PRAM(m) model. We also make the assumption that the point-to-point messages, also called packets, have a maximum size w, measured in bits.

In this model, only P packets can be issued into the network each g time steps, and thus the throughput of the network is

$$m_L = \frac{wP}{g}$$

bits per machine cycle. We denote this expression for throughput by m_L to make plain its correspondence with m in the PRAM(m) model. In this paper we discuss the case where the entire input is known to each processor. The case where the input is distributed across the processors requires long keys to be sent in their entirety, rather than sending just the original index of the key. Otherwise, the algorithm, as well as the resulting bounds, are the same.

In order to prove a lower bound for sorting in the LogP model, we first show that any LogP model algorithm can be simulated in the oracle model.

Lemma 15 *If $w > \log g$, then given any LogP algorithm A_l that completes in time T, there is an oracle algorithm A_o that computes the same function as A_l, also in time T, and A_o writes at most $O(m_L T)$ bits to the oracle memories.*

Proof:

We partition the time steps of the LogP algorithm A_l into *epochs*, where each epoch consists of g consecutive time steps. Note that each processor receives at most one message during any epoch. We can simulate A_l with an oracle algorithm A_o, where epoch i in A_l is represented in A_o by the oracle utilizing bits $(i-1)(\lfloor \log g \rfloor + 1 + w) + 1$ through $i(\lfloor \log g \rfloor + 1 + w) + 1$ bits in oracle memory j. The first $\lfloor \log g \rfloor + 1$ of the bits for each epoch are used to inform processor j of whether or not a message arrives during that epoch, and in the case of an arrival, the exact time step of the arrival during the epoch. In the case of an arrival, the remaining w bits contain the message contents. In A_o, the processors execute their algorithms for A_l, ignoring steps where messages are sent, and reading from the oracle memory at the start of every epoch. The total number of bits read is $P(\lfloor \log g \rfloor + 1 + w)\lceil \frac{T}{g} \rceil$. When $w > \log g$, this is $O(m_L T)$. ■

We briefly point out why we require that the oracle give each processor the timing information provided by the $\lfloor \log g \rfloor + 1$ additional bits used for each epoch. The LogP model is an asynchronous model, and thus processors cannot use the timing information to ensure the correctness of the algorithm. However, for the purpose of running time analysis, it is assumed that each processor behaves synchronously. Thus, in the optimal algorithm, it is possible that a processor is able to use the timing information to achieve a better running time than an algorithm that does not make inferences based on this information. Note that any algorithm that does not use this timing information can be simulated in the oracle model using at most $m_L T$ bits, even in the case where $w \leq g$.

As in the ER PRAM(m) model, when m_L grows as a fractional power of n, the time required to sort n keys is asymptotically no less than the time required to route all n keys through the network, even in the case where every processor knows all the keys in advance. Let $T_s(n)$ be the optimal sequential time to sort the n keys.

Theorem 16 *In the LogP model, sorting n distinct keys requires expected time*

$$\Omega\left(\frac{T_s(n)}{P} + \frac{n \log m_L}{m_L} + L + o\right),$$

provided that $n > P^2$, $w > \log g$ and $T_s(n) \geq L$.

Proof: We first assume there exists an algorithm A_l for the LogP model, where the average over all inputs of the time to perform A_l is at most $\frac{n \log m_L}{16 m_L}$, and reach a contradiction. By Lemma 15, such an algorithm A_l implies the existence of an oracle model algorithm where the average number of oracle bits is at most $\frac{n \log m_L}{8}$, and the average of the maximum number of ROM queries by any processor is at most $\frac{n \log m_L}{8 m_L}$. However, as we saw in Theorem 7, this leads to a contradiction of Lemma 6, and thus there does not exist such an algorithm A_l for the ER PRAM(m). The lower bound then follows from the fact at least one transmission, which requires time at least $L + 2o$, is required in any algorithm that sorts in time faster than time $T_s(n)$, and that the time to sort on P processors is no faster than the optimal time to sort on one processor divided by P. ■

Theorem 17 *In the LogP model, sorting n keys known to all processors can be completed in time*

$$O(T_s(\frac{n}{P}) + \frac{n \log n}{m_L} + L + o),$$

provided that $P = O(n^\beta)$ for some constant $\beta < 1$, and that $w \leq \log n$.

Proof: We show that we can sort in time $O(T_s(\frac{n}{P}) + \frac{gn \log n}{wP}(1-\beta)^{-3.42} + L + o)$. When $P = O(n^\beta)$ for some $\beta < 1$, we again use the recursive version of Columnsort. The difference between the LogP algorithm and the PRAM(m) algorithm is that we must now always use all of the processors, since reducing the number of utilized processors means reducing the rate of information transmission, which in turn reduces the effective value of m_L. However, if we reexamine the PRAM(m) algorithm, then we see that we can use p memory cell-processor pairs by simulating $\frac{p}{m}$ memory cells with each of the existing m memory cells, and this leads to an algorithm with running time $O(\frac{n \log n}{p})\frac{p}{m}$, which uses every processor. This algorithm has the same running time as the one given in Section 4 in the case where $n \geq p^3$, but is slower for the PRAM(m) for the case where $n \approx p \gg m$, since it then requires more levels of recursion.

In the LogP model, however, we need to use every processor, and thus we require that $P = O(n^\beta)$ for some $\beta < 1$. Thus, we start the recursion with n keys, and P processors. We recurse until we have sets to sort of size $\frac{n}{P}$. This requires i levels of recursion, where i is the smallest integer such that $n^{[1-(2/3)^i]} \geq P$. We again show that each of the three kinds of steps of Columnsort can be performed in the stated time bound. Locally sorting each column requires time $O(T_s(\frac{n}{P}))$. The communication time to perform the odd-even transposition sort is at most twice that required to perform the permutation routing step. The permutation routing step requires each processor to transmit and receive a total of $\frac{n}{P}$ keys, each of length $\log n$. Provided that we are able to make use of all w bits allowed in each packet, this requires a total time of $O(\frac{gn \log n}{wP} + L + o)$.

We can make use of a constant fraction of the bits in any packet, provided that $\log n \geq w$. Actually, this requirement can be relaxed somewhat, since if processor i sends any keys to processor j, then i transmits at least x keys to j, where x is $\frac{n}{P}$ divided by the number of columns in the smallest matrices of the recursion. We see that $x \geq \frac{n^{[1-1/2(2/3)^i]}}{P}$, and thus the requirement that $w \leq \log n$, can be relaxed to $w \leq \log n \frac{n^{[1-1/2(2/3)^i]}}{P}$. ■

5.2 The BSP model

We now briefly describe analogous bounds for sorting in the bulk-synchronous parallel (BSP) model proposed by Valiant [Val90a][Val90b]. The model consists of a set of processors capable of transmitting point- to-point messages through a communication network and facilities for performing barrier synchronization across any subset of the processors. The three parameters of the model are P, the number of processors, L, the minimum number of local computation steps between successive synchronization operations and g, the ratio between the throughput of local computation to the throughput at which a processor may inject point-to-point messages into the network. As in the LogP model, g enforces a limit on the communication throughput available to each processor. We assume that each transmitted packet is at most w bits in size.

Computation in the BSP model proceeds in *supersteps*, wherein each processor in parallel executes a task consisting of some number of local computation steps, message transmissions and message receipts, subject to the constraints imposed by the parameter g. The superstep lasts for kL time steps, where k is the minimum integer such that all processors have completed their tasks before time kL. As in the LogP model, the communication throughput of a single processor in the BSP model is $m_B = \frac{wP}{g}$ bits per step of local computation.

Theorem 18 *In the BSP model, sorting n distinct keys requires time*

$$\Omega(\frac{T_s(n)}{P} + \frac{n \log m_B}{m_B} + L),$$

when $n > P^2$ and $T_s(n) \geq L$.

Proof: Using the technique from Lemma 15, we see that the oracle model is also capable of simulating any BSP algorithm. Thus, the existence of any algorithm that sorts faster than $\frac{n \log m_B}{16 m_B}$ again implies the existence of an algorithm that contradicts Lemma 6. ■

Theorem 19 *In the BSP model, sorting any n keys can be completed in time*

$$O(T_S(\frac{n}{P}) + \frac{n \log n}{m_B} + L),$$

provided that $P = O(n^\beta)$ for some constant $\beta < 1$, and that $w \leq \log n$.

Proof: We show that we can sort in time $O(T_S(\frac{n}{P}) + \frac{gn \log n}{wP} + L)$. The only significant difference between the BSP algorithm and the LogP algorithm is that the BSP algorithm is partitioned into supersteps that are separated by synchronization operations. We achieve this simply by synchronizing after any phase of Columnsort completes at any level of the recursion, noting that computation performed during any phase of Columnsort is not needed by another processor until the subsequent phase. The running time holds provided that $P = O(n^\beta)$ for some constant $\beta < 1$, since in this case there is only a constant number of phases. ■

This result for sorting in the BSP model compares with the previous best randomized methods of Gerbessiotis and Valiant [GV94] for the BSP model with the assumption that each packet that is transmitted consists of exactly one key. Their algorithms run in time $O(\frac{n \log n}{P} + gp^\epsilon + \frac{gn}{P} + L)$, with high probability, for any positive constant $\epsilon < 1$, and for $P \leq n^{1-\delta}$, where δ is a small

constant depending on ϵ. After the preliminary version of this paper appeared in [ABK95], work on this problem by Goodrich [Goo96] has tightened the bounds for sorting on the BSP, giving deterministic algorithms which run in time $O(\frac{n \log n}{P} + (L + \frac{gn}{P})(\log n / \log(n/P)))$ for all values of P, coupled with a matching lower bound. Other recent work by Gerbessiotis and Siniolakis [GS96] gives a deterministic algorithm for sorting on the BSP which runs in time $(1 + o(1))(\frac{n \log n}{P} + L) + O(\frac{gn}{P})$ for $P = n^{1-\epsilon}$, $0 < \epsilon < 1$, and uses 1-optimal local computation.

6 Conclusion

We have examined the problem of sorting on parallel models with limited communication bandwidth. Our main results include upper and lower bounds for sorting on exclusive and queued-read variants of the PRAM(m) model which are asymptotically optimal for many practical settings of the parameters and are otherwise asymptotically tight to within at most a logarithmic factor. The form of our bound is noteworthy in that it demonstrates that all efficient parallel algorithms for sorting in this limited bandwidth model depend on large amounts of inter-processor communication. The techniques used to develop the bounds also apply to the LogP model and the BSP model, bridging models which consider the effect of limited bandwidth on parallel computation. For all three models of computation considered, when $m = \Omega(n^\beta)$, the time to sort and the time to transmit all the keys through the shared memory (or the network) are asymptotically equivalent, even in the case where the entire input is known to each of the processors. Furthermore, as long as $n > p^2$, the bounds do not depend on the parameter p, so when attempting to improve the performance of parallel sorting on machines with limited communication bandwidth, increasing communication bandwidth is more likely to be beneficial than increasing the number of processors. The lower bound, however, does not apply to the concurrent read version of the PRAM(m) originally introduced by Mansour, Nisan and Vishkin in [MNV94], and thus the asymptotic complexity of sorting in this model remains an open question.

7 Acknowledgements

We would like to thank Ralph Werchner for his useful comments and suggestions on an earlier version of this paper.

References

[ABK95] M. Adler, J. Byers, R. Karp. Parallel Sorting with Limited Bandwidth *Proc. 7th ACM Symp. on Parallel Algorithms and Architectures*: pp. 129 - 136, 1995.

[ACS90] A. Aggarwal, A. Chandra, and M. Snir. Communication Complexity of PRAMs. *Theoretical Computer Science* 71: pp 3-28, 1990.

[ACS87] A. Aggarwal, A. Chandra and M. Snir. Hierarchical Memory with Block Transfer. In *Proc. 28th IEEE Symp. on Foundations of Computer Science*, pp. 204-216, 1987.

[ACS89] A. Aggarwal, A. Chandra and M. Snir. On Communication Latency in PRAM Computations. In *Proc. 1st ACM Symp. on Parallel Algorithms and Architectures*, 1989.

[Adl96] M. Adler. New Coding Techniques for Improved Bandwidth Utilization In *Proc.* 37^{th} *IEEE Symp. on Foundations of Computer Science*, pp. , 1996.

[AHU74] A. Aho, J. Hopcroft, and J. Ullman. *The Design and Analysis of Computer Algorithms.* Addison-Wesley: Reading, MA, 1974.

[AKS83] M. Ajtai, J. Komlós and E. Szemerédi. An $O(n \log n)$ sorting network. *Combinatorica* 3: pp. 1 - 19, 1983.

[BGMZ95] G. Blelloch, P. Gibbons, Y. Matias and M. Zagha. Accounting for Memory Bank Contention and Delay in High-Bandwidth Multiprocessors. In *Proc.* 7^{th} *ACM Symposium on Parallel Algorithms and Architectures*, pp. 84-94, 1995.

[BC82] A. Borodin and S. Cook. A Time-Space Tradeoff for Sorting on a General Sequential Model of Computation. *SIAM J. of Computing*, 11(2): pp. 287 - 297, 1982.

[Col88] R. Cole. Parallel Merge Sort. *SIAM J. of Computing*, 17(4): pp. 770 - 785, 1988.

[CD82] S. Cook, C. Dwork, and R. Reischuk. Upper and Lower Bounds for Parallel Random Access Machines Without Simultaneous Writes. *SIAM J. of Computing* 15: pp. 87-97, 1985.

[CKP+93] D. Culler, R. M. Karp, D. Patterson, A. Sahay, K. E. Schauser, E. Santos, R. Subramonian and T. von Eicken. LogP: Towards a Realistic Model of Parallel Computation. In *Proc. 4th ACM SIGPLAN Symp. on Principles and Practice of Parallel Programming*, pp. 1-12, January 1993.

[CS92] R. Cypher and J. Sanz. Cubesort: A Parallel Algorithm for Sorting N Data Items with S-Sorters. *Journal of Algorithms 13*: pp. 211-234, 1992.

[Dus94] A. Dusseau. Modeling Parallel Sorts with LogP on the CM-5. Technical Report: UCB/CSD-94-829, May 1994.

[GMR94a] P. Gibbons, Y. Matias, and V. Ramachandran. The QRQW PRAM: Accounting for Contention in Parallel Algorithms. In *Proc. 5th ACM-SIAM Symposium on Discrete Algorithms*, pp. 638-648, January 1994.

[GMR94b] P. Gibbons, Y. Matias, and V. Ramachandran. Efficient Low-Contention Parallel Algorithms. In *Proc. 6th ACM Symposium on Parallel Algorithms and Architectures*, pp. 236-247, June 1994.

[Goo96] M. Goodrich. Communication-Efficient Parallel Sorting *Proceedings of the* 28^{th} *Annual ACM Symposium on Theory of Computing*: pp. 247 - 256, 1996.

[GS96] A. Gerbessiotis, C. Siniolakis. Deterministic Sorting and Randomized Median Finding on the BSP Model. In Proc. 8th ACM Symposium on Parallel Algorithms and Architectures, pp. 1996.

[GV94] A. Gerbessiotis, L. Valiant. Direct Bulk-Synchronous Algorithms. In *Journal of Parallel and Distributed Computing*, 22:251 -267, 1994.

[KR90] R. M. Karp and V. Ramachandran. Parallel Algorithms for Shared-Memory Machines. In *Handbook of Theoretical Computer Science*, J. van Leeuwen, Ed., pp. 869-941. Elsevier Science Publishers: Amsterdam, The Netherlands, 1990.

[Lei85] T. Leighton. Tight Bounds on the Complexity of Parallel Sorting. *IEEE Trans. on Computers*, c-34(4): pp. 344-354, 1985.

[Lei92] T. Leighton. *Introduction to Parallel Algorithms and Architectures*, Morgan Kaugmann, 1992.

[Mac95] P. Mackenzie. Lower Bounds for Randomized Exclusive Write PRAMs. In *Proc. 7th ACM Symposium on Parallel Algorithms and Architectures*, pp. 254-263, July 1995.

[MNV94] Y. Mansour, N. Nisan and U. Vishkin. Trade-offs Between Communication Throughput and Parallel Time. In *Proceedings of the* 26^{th} *Annual ACM Symposium on Theory of Computing*: pp. 372-381, 1994.

[MV 84] . Mehlhorn and U. Vishkin. Randomized and Deterministic Simulations of PRAMs by Parallel Machines with Restricted Granularity of Parallel Memories. In *Acta Informatica*, 21:339-374, 1984.

[Tho80] C. Thompson. A Complexity Theory for VLSI. *PhD Thesis*. Carnegie-Mellon University, Pittsburgh, PA, 1980.

[Val90a] L. Valiant. A Bridging Model for Parallel Computation. *Communications of the ACM*, 33(8): pp 103-111, August 1990.

[Val90b] L. Valiant. General Purpose Parallel Architectures. In *Handbook of Theoretical Computer Science*, J. van Leeuwen, Ed., pp. 943-971. Elsevier Science Publishers: Amsterdam, The Netherlands, 1990.

[VW85] U. Vishkin and A. Wigderson. Trade-Offs between Depth and Width in Parallel Computation. *SIAM Journal of Computing*, 14(2): pp. 303 - 314, 1985.

[Yao77] A. Yao. Probabilistic Computations: Toward a Unified Measure of Complexity. In *Proc.* 18^{th} *IEEE Symp. on Foundations of Computer Science*, pp. 222-227, 1977.