

On the Complexity of List Ranking in the Parallel External Memory Model

Riko Jacob¹, Tobias Lieber¹, and Nodari Sitchinava²

¹ Institute for Theoretical Computer Science, ETH Zürich, Switzerland
{rjacob,lieberto}@inf.ethz.ch

² Department of Information and Computer Sciences, University of Hawaii, USA
nodari@hawaii.edu

Abstract. We study the problem of *list ranking* in the parallel external memory (PEM) model. We observe an interesting dual nature for the hardness of the problem due to limited information exchange among the processors about the structure of the list, on the one hand, and its close relationship to the problem of permuting data, which is known to be hard for the external memory models, on the other hand.

By carefully defining the power of the computational model, we prove a permuting lower bound in the PEM model. Furthermore, we present a stronger $\Omega(\log^2 N)$ lower bound for a special variant of the problem and for a specific range of the model parameters, which takes us a step closer toward proving a non-trivial lower bound for the list ranking problem in the bulk-synchronous parallel (BSP) and MapReduce models. Finally, we also present an algorithm that is tight for a larger range of parameters of the model than in prior work.

1 Introduction

Analysis of massive graphs representing social networks using distributed programming models, such as MapReduce and Hadoop, has renewed interests in distributed graph algorithms. In the classical RAM model, depth-first search traversal of the graph is the building block for many graph analysis solutions. However, no efficient depth-first search traversal is known in the parallel/distributed setting. Instead, list ranking serves as such a building block for parallel solutions to many problems on graphs.

The *list ranking* problem is defined as follows: given a linked list compute for each node the length of the path to the end of the list. In the classic RAM model, the list can be ranked in linear time by traversing the list. However, in the PRAM model (the parallel analog of the RAM model) it took almost a decade from the first solution by Wyllie [1] till it was solved optimally [2].

The problem is even more intriguing in the models that study block-wise access to memory. For example, in the *external memory* (EM) model of Aggarwal and Vitter [3] list ranking is closely related to the problem of permuting data in an array. The EM model studies the *input/output (I/O) complexity* – the number of transfers an algorithm has to perform between a disk that contains

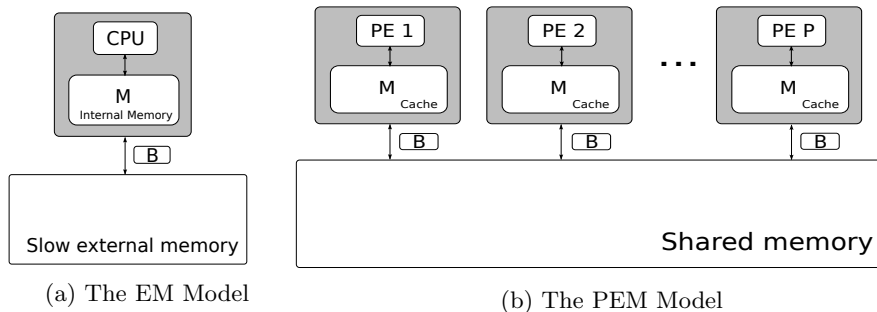


Fig. 1: The sequential and parallel external memory models.

the input and a fast internal memory of size M . Each transfer is performed in blocks of B contiguous elements. In this model, permuting and, consequently, list ranking require I/O complexity which is closely related to sorting [4], rather than the linear complexity required in the RAM model.

In the distributed models, such as bulk-synchronous parallel (BSP) [5] or MapReduce [6] models, the hardness of the list ranking problem is more poorly understood. These models consist of P processors, each with a private memory of size M . With no other data storage, typically $M = \Theta(N/P)$. The data is exchanged among the processors during the *communication rounds* and the number of such rounds defines the complexity metric of these models.

One established modeling of today's commercial data centers running MapReduce is to assume $P = \Theta(N^\epsilon)$ and $M = \Theta(N^{1-\epsilon})$ for a constant $0 < \epsilon < 1$. Since network bandwidth is usually the limiting factor of these models, $O(\log_M P) = O(1)$ communication rounds is the ultimate goal of computation on such models [7]. Indeed, if each processor is allowed to send up to $M = N/P$ items to any subset of processors, permuting of the input can be implemented in a single round, while sorting takes $O(\log_M P) = O(1)$ rounds [8,9]. On the other hand, the best known solution for list ranking is via the simulation results of Karloff et al. [7] by simulating the $O(\log N)$ time PRAM algorithm [2], yielding $O(\log P)$ rounds, which is strictly worse than both sorting and permuting. Up to now, no non-trivial lower bounds (i.e. stronger than $\Omega(\log_M P) = \Omega(1)$) are known in the BSP and MapReduce models.

In this paper we study lower bounds for the list ranking problem in the *parallel external memory (PEM)* model. The PEM model was introduced by Arge et al. [10] as a parallel extension of the EM model to capture the hierarchical memory organization of modern multicore processors. The model consists of P processing units, each containing a private cache of size M , and a shared (external) memory of conceptually unlimited size (see Figure 1b). The data is concurrently transferred between the processors' caches and shared memory in blocks of size B . The model measures the *parallel I/O complexity* – the number of parallel block transfers. From the discussion above, it appears that the hardness of list ranking stems from two factors: (1) limited speed of discovery of the

structure of the linked list due to limited information flow among the processors, and (2) close relationship of list ranking to the problem of permuting data. While only one of these challenges is captured by the distributed models or the sequential EM model, both of them are exhibited in the PEM model: the first one is captured by the distributed nature of the private caches of the model, and the second one has been shown by Greiner [11] who proves that permuting data in the PEM model takes asymptotically $\text{perm}_P(N, M, B) = \min \left\{ \frac{N}{P}, \frac{N}{PB} \overline{\log}_d \frac{N}{B} \right\}$ parallel I/Os, where $d = \max \left\{ 2, \min \left\{ \frac{M}{B}, \frac{N}{PB} \right\} \right\}$ and $\overline{\log}(x) = \max\{1, \log(x)\}$.

Part of the challenge of proving lower bounds (in any model) is restricting the model enough to be able to prove non-trivial bounds, while identifying the features of the model that emphasize the hardness of a particular problem.

An example of such restriction in the external memory models (both sequential and parallel) is the so-called *indivisibility assumption* [3]. The assumption states that each item is processed as a whole, and no information can be obtained from a part of the input, for example, by combining several items into one. To our knowledge, without the indivisibility assumption, it is not clear how to prove lower bounds in the PEM model exceeding the information-theoretic lower bounds of $\Omega(\log P)$ parallel I/Os [12,10].

1.1 Our Contributions

In this paper, we address the precise formulation of the power of the PEM model for the list ranking problem. In Section 2 we present the atomic PEM model which formalizes the indivisibility assumption in the PEM model. It can be viewed as the parallel analog of the model for proving permuting lower bounds in the sequential EM model [3]. We extend this basic model by allowing an algorithm to perform operations on the atoms that create new atoms. While we always keep the indivisibility of the atoms, the precise operations and the information the algorithm has about the content of the atom varies.

In the sequential EM model, Chiang et al. [4] sketch a lower bound for list ranking via a reduction to the *proximate neighbors problem*. However, since there is no equivalent to Brent’s scheduling principle [13] in the PEM model [11], the lower bound does not generalize to the PEM model.

Therefore, in Section 3 we derive a lower bound of $\Omega(\text{perm}_P(N, M, B))$ parallel I/Os for the proximate neighbor problem, and two problems which are related to the list ranking problem. Our lower bounds hold for both deterministic and randomized algorithms. In the process we provide an alternative proof for the proximate neighbor lower bound in the sequential external memory model, matching the result of Chiang et al. [4]. Those lower bounds essentially exploit the fact that the same problem can be represented as input in many different layouts.

The discussion in Section 1 about the dual nature of hardness of the list ranking problem might hint at the fact that the result above is only part of the picture and a stronger lower bound might be achievable. Part of the challenge in proving a stronger lower bound lies in the difficulty of combining the indivisibility assumption with the restrictions on how the structure of the linked list is shared

among the processors without giving the model too much power, thus, making the solutions trivial. We address this challenge by defining the *interval PEM* model and defining the *guided interval fusion (GIF)* problem (Section 4). We prove that GIF requires $\Omega(\log^2 N)$ parallel I/Os in the interval PEM. Our lower bound for GIF in the PEM model implies a $\Omega(\log N)$ lower bound for the number of rounds for GIF in the distributed models when $P = \Theta(M) = \Theta(\sqrt{N})$.

GIF captures the way how all currently known algorithms use information to solve list ranking in all parallel/distributed models. Therefore, if this lower bound could be broken for the list ranking problem, it will require completely new algorithmic techniques. Thus, our result brings us a step closer to proving the unconditional $\Omega(\log P)$ lower bound in the BSP and MapReduce models.

Finally, in Section 5 we improve the PEM list ranking algorithm of Arge et al. [14] to work efficiently for a larger range of parameters in the PEM model.

2 Modeling

We extend the description of the PEM model given in Section 1 to define the PEM model more precisely. Initially, the data resides in the shared memory. To process any element, it must be present in the corresponding processor's cache. The shared memory is partitioned into blocks of B contiguous elements and the transfer of data between the shared memory and caches is performed by transferring these blocks as units. Each transfer, an *input-output operation*, or simply *I/O*, can transfer one block of B items between the main memory and each processors' cache. Thus, up to P blocks can be transferred in each *parallel I/O* operation. The complexity measure of a PEM algorithm is the number of parallel I/Os that the algorithm performs.

Similar to the PRAM model, there are several policies in the PEM model, for handling simultaneous accesses by multiple processors to a block in shared memory. In this paper we consider the CREW PEM model, using a block wise concurrent read, exclusive write policy.

In order to prove lower bounds we make the definition of the model more precise by stating what an algorithm is able to do in each step. In particular, we assume that each element of the input is an indivisible unit of data, an *atom*, which consumes one memory cell in the cache or shared memory. Such atoms come into existence either as input atoms, or by an operation, as defined later, on two atoms. A program or an algorithm has limited knowledge about the content of an atom. In this paper an atom does not provide any information. Furthermore, the *atomic PEM* is limited to the following operations: an I/O operation reads or writes one block of up to B atoms in the shared memory, and atoms can be copied or deleted. Formal definitions of similar PEM machines can be found in [10,11].

For providing lower bounds for different problems, the concept of the atomic PEM is extended in later sections.

In the following, we distinguish between algorithms and programs in the following way: In an *algorithm* the control flow might depend on the input, i.e.,

there are conditional statements (and therefore loops). In contrast, a *program* has no conditional statements and is a sequence of valid instructions for a PEM model, independent of the input (atoms). For a given instance of a computational task, a program can be seen as an instantiation of an algorithm to which all results of conditional statements and index computations are presented beforehand. Note, that in the problems considered in Section 3 the copying and the deletion operation of the atomic PEM do not help at all, since a program can be stripped down to operations which operate on atoms which influence the final result.

3 Counting Lower Bounds to the List Ranking Problem

In this section we prove the lower bound for the list ranking problem by showing the lower bound to the proximate neighbors problem [4] and reducing it to the problems of semigroup evaluation, edge contraction and, finally, list ranking.

3.1 Proximate Neighbors Problem in PEM

Definition 1 ([4]). A block permutation describes the content of the shared memory of a PEM configuration as a set of at most B atoms for each block.

Definition 2. An instance of the proximate neighbors problem of size N consists of atoms x_i for $i \in [N]$. All atoms are labeled by a labeling function $\lambda: [N] \mapsto [\frac{N}{2}]$ with $|\lambda^{-1}(i)| = 2$. An output block permutation solves the problem if for every $i \in [\frac{N}{2}]$ the two neighboring atoms $\lambda^{-1}(i)$ are stored in the same block. The blocks in such an output may contain less than B atoms.

Lemma 3. Let A be a computational problem of size N for which an algorithm has to be capable of generating at least $(\frac{N}{eB})^{cN}$ block permutations, for a constant $c > 0$. Then in the CREW atomic PEM model with $P \leq \frac{N}{B}$ processors, at least half of the input instances of A require $\Omega(\text{perm}_P(N, M, B))$ parallel I/Os.

Proof. Straightforward generalization of the proof to Theorem 2.7 in [11]. \square

Theorem 4. At least half of the instances of the proximate neighbors problem of size N require $\Omega(\text{perm}_P(N, M, B))$ parallel I/Os in the CREW atomic PEM model with $P < N/B$ processors.

Proof. Any algorithm solving an instance of the proximate neighbors problem must be capable of generating at least $\frac{(N/2)!}{(B/2)^{\frac{N}{2}}}$ block permutations (see the full version of this paper [15] for a complete proof). The theorem follows from Lemma 3 and the observation that $\frac{(N/2)!}{(B/2)^{\frac{N}{2}}} \geq (\frac{N}{eB})^{\frac{N}{2}}$. \square

Note, that the bound holds even if a program has full access to the labeling function λ and thus is fully optimized for an input. The origin of the complexity of the problem rather is permuting the atoms to the output block permutation that solves the problem.

3.2 Semigroup Evaluation in the PEM Model

Consider the problem of evaluating a very simple type of expressions, namely that of a semigroup, in the PEM model.

Definition 5 (Semigroup Evaluation). *Let S be a semigroup with its associative binary operation $\cdot : S \times S \rightarrow S$. The semigroup evaluation problem is defined as evaluating the expression $\prod_{i=1}^N a_i$, with $a_i = x_{\pi(i)}$, for the array of input atoms $x_i \in S$ for $1 \leq i \leq N$ and where π is a permutation over $[N]$.*

To be able to solve the semigroup evaluation problem, algorithms must be able to apply the semigroup operation to atoms. Thus, we extend the atomic PEM model to the *semigroup PEM* model by the following additional operation: if two atoms x and y are in the cache of a processor, a new atom $z = x \cdot y$ can be created.

We say that a program is correct in the semigroup PEM if it computes the correct result for any input and any semigroup.

Theorem 6. *At least one instance of the semigroup evaluation problem of size N requires $\Omega(\text{perm}_P(N, M, B))$ parallel I/Os in the CREW semigroup PEM model with $P \leq \frac{N}{B}$ processors.*

Proof (Sketch). Let I_λ^P be an instance of the proximate neighbors problem over the input atoms $X = \{x_i | i \in [N]\}$ with its labeling function λ . We consider an instance I_π^S of the semigroup evaluation problem over the semigroup on the set X^2 with the semigroup operation $(a, b) \cdot (c, d) = (a, d)$, where $a, b, c, d \in X$. Furthermore, the instance I_π^S is defined over the input atoms $a_i = (x_i, x_i)$, with $1 \leq i \leq N$. The permutation π of I_π^S is one of the permutations such that for all $i \in [\frac{N}{2}]$, $\{\pi(2i-1), \pi(2i)\} = \lambda^{-1}(i)$ holds.

Then, the key idea is to write for each application of the semigroup operation $(a, b) \cdot (c, d)$ in a program solving I_π^S , the pair $\{b, c\}$ as a result for I_λ^P to the output. This would yield an efficient program for I_π^S , and therefore yields the lower bound by Theorem 4. The full argument can be found in [15]. \square

3.3 Atomic Edge Contraction in the PEM Model

Definition 7. *The input of the atomic edge contraction problem of size N consists of atoms x_i , $1 \leq i \leq N$, which represent directed edges e_i on a $(N+1)$ -vertex path between vertices s and t . Initially, the edges are located in arbitrary locations of the shared memory. The instance is solved if an atom representing the edge (s, t) is created and written to shared memory.*

To prove the lower bound for the atomic edge contraction problem, we extend the atomic PEM with an additional operation: two atoms representing a pair of edges (a, b) and (b, c) can be removed and replaced by a new atom representing a new edge (a, c) . We call the resulting model *edge-contracting PEM*.

Theorem 8. *There is at least one instance of the atomic edge contraction problem of size N which requires $\Omega(\text{perm}_P(N, M, B))$ parallel I/Os in the CREW edge-contracting PEM model with $P \leq \frac{N}{B}$ processors.*

Proof (Sketch). An instance I_π^S of the semigroup evaluation problem can be reduced to an instance $I^\mathcal{E}$ of the atomic edge contraction problem, by defining the atom $x_{\pi(i)}$ of $I^\mathcal{E}$, initially stored at location $\pi(i)$, as $e_{\pi(i)} = (\pi(i), \pi(i+1))$, where π is the permutation of I_π^S . The full argument can be found in [15]. \square

3.4 Randomization and Relation to the List Ranking Problem

Observe that the expected number of parallel I/Os of a randomized algorithm for an instance is a convex combination of the number of parallel I/Os of programs. Combining this observation with the $\Omega(\log P)$ lower bound of [12,10] mentioned in Section 1 we obtain:

Theorem 9. *For the proximate neighbors, semigroup evaluation, and the atomic edge contraction problems, there exists at least one instance that requires at least $\Omega(\text{perm}_P(N, M, B) + \log P)$ expected parallel I/Os by any randomized algorithm in the corresponding PEM model with $P \leq \frac{N}{B}$ processors.*

Although our semigroup PEM and edge-contracting PEM models might seem too restrictive at a first glance. To the best of our knowledge all current parallel solutions to list ranking utilize pointer hopping, which can be reduced to atomic edge contraction and thus the lower bound applies.

4 The Guided Interval Fusion Problem (GIF)

In this section we prove for the GIF problem, which is very similar to the atomic edge contraction problem, a lower bound of $\Omega(\log^2 N)$ in the PEM model with parameters $P = M$ and $B = M/2$ for inputs of size $N = PM = 2^x$ for some $x \in \mathbb{N}$. In contrast to the atomic edge contraction problem, in the GIF problem an algorithm is not granted unlimited access to the permutation π .

The chosen parameters of the PEM model complement the upper bounds of Section 5 at one specific point in the parameter range. Note that with careful modifications the $\Omega(\log^2 N)$ bound can even be proven for $N = M^{\frac{3}{2}+\epsilon}$.

Definition 10. *The interval PEM is an extension of the atomic PEM: Two atoms x and y representing closed intervals I_x and I_y , located in one cache, can be fused if $I_x \cap I_y \neq \emptyset$. Fusing creates a new atom z representing the interval $I_z = I_x \cup I_y$. We say z is derived from x if z is the result of zero or more fusing operations starting from atom x .*

Definition 11. *The guided interval fusion problem (GIF) is a game between an algorithm and an adversary, played on an interval PEM. The algorithm obtains a GIF instance \mathcal{G} in the first N cells of the shared memory, containing N uniquely*

named atoms x_i , $1 \leq i \leq N$. Each initial atom x_i represents the (invisible to the algorithm) closed interval $I_{x_i} = [k - 1, k]$ for $k = \pi(i)$ with $1 \leq k \leq N$.

The permutation π is gradually revealed by the adversary in form of boundaries $p = (i, j)$ meaning that the (initial) atoms x_i and x_j represent neighboring intervals ($\pi(j) = \pi(i) + 1$). We say that the boundary point p for x_i and x_j is revealed. The adversary must guarantee that at any time, for all existing atoms, at least one boundary is revealed. The game ends as soon as an atom representing $[0, N]$ exists.

Note the following: The algorithm may try to fuse two atoms, even though by the revealed boundaries this is not guaranteed to succeed. If this attempt is successful because their intervals share a point, a new atom is created and the algorithm solved a boundary. We call this phenomenon a *chance encounter*. Since the interval PEM extends the atomic PEM, copying of atoms is allowed.

For the lower bound, we assume that the algorithm is omniscient. More precisely, we assume there exists a central processing unit, with unlimited computational power, that has all presently available information on the location of atoms and what is known about boundaries. This unit can then decide on how atoms are moved and fused.

Thus, as soon as all boundary information is known to the algorithm, the instance is solvable with $\mathcal{O}(\log N)$ parallel I/Os: The central unit can virtually list rank the atoms, group the atoms by rank into P groups, and then by permuting move to every processor $\mathcal{O}(M)$ atoms which then can be fused with $\mathcal{O}(1)$ I/Os to the solving atom.

Hence, the careful revealing of the boundary information is crucial. To define the revealing process for GIF instances, the atoms and boundaries of a GIF instance \mathcal{G} of size N are related to a perfect binary tree $\mathcal{T}_{\mathcal{G}}$. The tree $\mathcal{T}_{\mathcal{G}}$ has N leaves, $N - 1$ internal nodes and every leaf is at distance $h = \log N$ from the root. More precisely, each leaf $i \in [N]$ corresponds to the atom representing the interval $[i - 1, i]$. And each internal vertex v_p corresponds to the boundary $p = (i, j)$ where i corresponds to the rightmost leaf of its left subtree, and j to the leftmost leaf of the right subtree. The levels of $\mathcal{T}_{\mathcal{G}}$ are numbered bottom up: the leaves have level 1 and the root vertex level $\log N$ (corresponding to the revealing order of boundaries).

The protocol for boundary announcement, shows for a random GIF instance \mathcal{G} , that a deterministic algorithm takes $\Omega(\log^2 N)$ parallel I/Os.

Definition 12. *The tree $\mathcal{T}_{\mathcal{G}}$ is the guide of \mathcal{G} if boundaries are revealed in the following way. Let x be an atom of \mathcal{G} representing the interval $I = [a, b]$. If neither of the boundaries a and b are revealed, the boundary whose node in $\mathcal{T}_{\mathcal{G}}$ has smaller level, is revealed. If both have the same level, a is revealed.*

Note that for the analysis it is irrelevant how to break ties (in situations when the two invisible boundaries have the same level). By the assumption that the algorithm is omniscient, when p is revealed, immediately all intervals having p as boundary know that they share this boundary. Thus, the guide ensures that at any time each atom knows at least one initial atom with which it can be fused.

A node $v \in \mathcal{T}_{\mathcal{G}}$ is called *solved*, if there is an atom of \mathcal{G} representing an interval that contains the intervals of the leaves of the subtree of v . The boundary p is only revealed by the guide if at least one child of v_p is solved.

An easy (omniscient) algorithm solving a GIF instance can be implemented: in each of $O(\log N)$ rounds, permute the atoms such that for every atom there is at least one atom, known to be fuseable, which resides in the same cache. Fuse all pairs of neighboring atoms, reducing the number of atoms by a factor of at least 2, revealing new boundaries. Repeat permuting and fusing until the instance is solved. Because the permuting step can be achieved with $\mathcal{O}(\log N)$ parallel I/Os, this algorithm finishes in $\mathcal{O}(\log^2 N)$ parallel I/Os. Note that solving a boundary resembles bridging out one element of an independent set in the classical list ranking scheme. Thus, most list ranking algorithms use information as presented to an algorithm solving a GIF instance \mathcal{G} guided by $\mathcal{T}_{\mathcal{G}}$. Hence, this natural way of solving a GIF instance can be understood as solving in every of the $\log N$ rounds a proximate neighbors instance, making the $\Omega(\log^2 N)$ lower bound reasonable.

In the following we prove the lower bound for GIF by choosing k and showing that $s = k - 1 = \mathcal{O}(\log N) = \mathcal{O}(\log M)$ *progress stages* are necessary to solve all nodes W of level k of $\mathcal{T}_{\mathcal{G}}$ (thus, $|W| = 2^{h+1-k}$). For each stage we show in Lemma 15 that it takes $\Omega(\log N)$ parallel I/Os to compute.

To measure the progress of a stage, configurations of interval PEM machines are used. The *configuration* C^t after the interval PEM machine performed t I/Os followed by fusing operations consists of sets of atoms. For each cache and each block of the shared memory, there is one set of atoms.

For $e \in W$, let T_e be the subtree of e in $\mathcal{T}_{\mathcal{G}}$, and \mathcal{B}_e be all boundaries in T_e . The progress measure towards solving e is the highest level of a solved boundary in \mathcal{B}_e . More precisely, T_e is *unsolved* on level i , if all boundaries of level i of \mathcal{B}_e are unsolved. Initially every T_e is unsolved on level 2. The solved level increases by one at a time if only revealed boundaries are solved, but chance encounters may increase it faster.

The execution of a deterministic algorithm A defines the following s progress stages: Let $s = \frac{2 \log M}{16}$ and $X = \frac{|W|}{s}$. In each stage $1 < i \leq s$, at least X elements increase their level to i . Over time, the number of elements that are unsolved on level i decreases, and we define t_i to be the last time where in C^{t_i} the number of elements of W that are unsolved on level $i + 1$ is at least $|W| - iX$. Further, let W_i be the elements (at least X of them) that in stage i get solved on level i or higher (in the time-frame from t_{i-1} to $t_i + 1$). We choose $k - 1 = \frac{h}{16}$ such that $X = \frac{2^{h+1-k}}{s} \geq 2^{\frac{15h}{16}} / s = M^{\frac{15}{8}} / s > M^{\frac{7}{4}}$ because $s = \frac{2 \log M}{16} < M^{\frac{1}{8}}$ for $M \in \mathbb{N}$.

In the beginning of stage i , for each $v \in W_i$ the level of v is at most $i - 1$, and hence all level i nodes are not announced to the algorithm. Let P_i be the set of boundaries for which progress is traced: For every $e \in W_i$, there is a node v_{p_e} of level i with boundary p_e that is solved first (brake ties arbitrarily). Then P_i consists of those boundaries. We define a_e and b_e to be the two level 1 atoms (original intervals) defining the boundary p_e . Then all intervals having boundary p_e are derived of a_e or b_e . Solving the boundary p_e means fusing

any interval derived of a_e with any interval derived of b_e . Furthermore a traced boundary is considered solved if in its interval (the one corresponding to an element of W) a chance encounter solves a boundary of level greater than i .

To trace the progress of the algorithm towards fusing the atoms of one stage, we define the graph $G_t^i = (V, E_t^i)$ from the configuration C^{t_i+t} . There is one vertex for each cache and each block of the shared memory (independent of t). There is an edge (self-loops allowed) $\{u, v\} \in E_t^i$ if for some $e \in W_i$ some atom derived of a_e is at u and some atom derived of b_e is at v or vice versa. The *multiplicity* of an edge counts the number of such e . The multiplicity of the graph is the maximal multiplicity of an edge.

Note that solving a node v_e requires that it counts as a self-loop somewhere. Hence the sum of the multiplicities of self-loops are an upper bound on the number of solved nodes, and for the stage to end, i.e., at time $t_{i+1} + 1$, the sum of the multiplicities of loops must be at least X . After each parallel I/O chance encounters may happen. Thus, the number of chance encounters is given by P times the multiplicity of self-loops at the beginning of the stage.

We say that two nodes of \mathcal{T}_G are indistinguishable if they are on the same level and exchanging them could still be consistent with the information given so far. Let l_e (derived of a_e) and r_e (derived of b_e) be the two children of $v_e \in P_i$. By definition, at time t_i both l_e and r_e are unsolved and hence v_e is not revealed.

Boundaries corresponding to nodes of level higher than k may be announced or solved (not only due to chance encounters). To account for that, we assume that all such boundaries between the intervals corresponding to W are solved. Hence the algorithm is aware of the leftmost and rightmost solved interval belonging to these boundaries, and this may extend to other intervals by revealed boundaries. Only the nodes of level i that correspond to this leftmost or rightmost interval might be identifiable to the algorithm, all other nodes of level i are indistinguishable. Because $i < k$, for all traced pairs a_e, b_e at least one of the elements belongs to this big set of indistinguishable nodes. We mark identifiable nodes. Hence, at stage i the algorithm has to solve a random matching of the traced pairs where all marked nodes are matched with unmarked ones (and unmarked ones might be matched with marked or unmarked ones).

The next lemma derives a high-probability upper bound on the multiplicity of a graph. All remaining proofs of this section are deferred to [15]. The proof of the following lemma uses the Hoeffding inequality.

Lemma 13. *Consider a deterministic GIF algorithm operating on a uniformly chosen permutation π defining \mathcal{T}_G for the GIF instance \mathcal{G} . Let $p(i, M)$ be the probability that G_0^i has multiplicity at most $M^{\frac{5}{8}}$ (where P, N, t_j , and k depend on parameter M). Then there is a M' such that for all $M \geq M'$ and for each $i \leq k$ it holds $p(i, M) \geq 1 - \frac{1}{M^2}$.*

Fundamental insights on identifying two pairs of a K_4 show that the progress achieved with one I/O can not be too large.

Lemma 14. *If the graph G_t^i has multiplicity at most m , then G_{t+1}^i has multiplicity at most $4m$.*

By the two previous lemmas we obtain the following result.

Lemma 15. *Let A be an algorithm solving an GIF instance \mathcal{G} guided by $\mathcal{T}_{\mathcal{G}}$ of height h traced at level $k - 1 = h/16$. Each progress stage $j < s = k - 1$ of A , assuming $t_j < \log^2 M$, takes time $t_j - t_{j-1} = \Omega(\log M)$.*

There are $\mathcal{O}(\log N)$ stages, each taking at least $\Omega(\log N)$ I/Os, yielding with a union bound over all G_0^i for all progress stages $i < s$:

Lemma 16. *Consider a deterministic GIF algorithm operating on a uniformly chosen permutation π defining $\mathcal{T}_{\mathcal{G}}$ for the GIF instance \mathcal{G} . Then there is a M' such that for all $M \geq M'$, solving \mathcal{G} in the interval PEM takes with high probability ($p > 1 - 1/M$) at least $\Omega(\log^2 N)$ parallel I/Os.*

By Yao's principle [16] this can be transferred to randomized algorithms:

Theorem 17. *The expected number of parallel I/Os to solve a GIF instance of size $N = PM$ on an interval PEM with $M = P$ is $\Omega(\log^2 N)$.*

A simple reduction yields:

Theorem 18. *Solving the GIF problem in the BSP or in the MapReduce model with $N = PM$ and $P = \Theta(M) = \Theta(\sqrt{N})$ takes $\Omega(\log N)$ communication rounds.*

GIF is an attempt to formulate how the known algorithms for list ranking distribute information by attaching it to atoms of the PEM model. Most known algorithms for list ranking use fusing of edges on an independent set of edge-pairs (bridging out edges). This means that every edge is used (if at all) either as first or second edge in the fusing. This choice of the algorithm is taken without complete information, and hence we take it as reasonable to replace it by an adversarial choice, leading to the definition of the guide of a GIF instance.

Additionally, the PEM lower bound shows that there is no efficient possibility to perform different stages (matchings) in parallel, showing that (unlike in the efficient PRAM sorting algorithms) no pipelining seems possible. At this stage, our lower bound is hence more a bound on a class of algorithms, and it remains a challenge to formulate precisely what this class is. Additionally, it would be nice to show lower bounds in a less restrictive setting.

5 Upper Bounds

Improvements in the analysis [11] of the PEM merge sort algorithm [10] yield:

Lemma 19 ([11]). *The I/O complexity of sorting N records with the PEM merge sort algorithm using $P \leq \frac{N}{B}$ processors is $\text{sort}_P(N, M, B) = \mathcal{O}\left(\frac{N}{PB} \log_d \frac{N}{B}\right)$ for $d = \max\{2, \min\{\frac{N}{PB}, \frac{M}{B}\}\}$.*

We use it to extend the parameter range for the randomized list ranking algorithm [14] from $P \leq \frac{N}{B^2}$, and $M = B^{\mathcal{O}(1)}$ to:

Theorem 20. *The expected number of parallel I/Os, needed to solve the list ranking problem of size N in the CREW PEM model with $P \leq \frac{N}{B}$ is*

$$\mathcal{O} \left(\text{sort}_P(N, M, B) + (\log P) \overline{\log} \frac{B}{\log P} \right)$$

which is for $B < \log P$ just $\text{sort}_P(N, M, B)$.

In order to make the standard recursive scheme work, two algorithms are used. By Lemma 19, the randomized algorithm of [14], which is based on [13], can be used whenever $N \geq P \min\{\log P, B\}$. This yields the $(\log P) \overline{\log} \frac{B}{\log P}$ term, if $B > \log P$. Otherwise ($N \leq P \min\{\log P, B\}$) a simulation of a work-optimal PRAM algorithm [2] is used. A careful implementation and analysis [15] yield Theorem 20.

References

1. Wyllie, J.: The Complexity of Parallel Computation. PhD thesis, Cornell University (1979)
2. Anderson, R.J., Miller, G.L.: Deterministic parallel list ranking. In: VLSI Algorithms and Architectures. Volume 319 of LNCS. Springer (1988) 81–90
3. Aggarwal, A., Vitter, J.S.: The input/output complexity of sorting and related problems. Commun. ACM **31**(9) (September 1988) 1116–1127
4. Chiang, Y.J., Goodrich, M., Grove, E., Tamassia, R., Vengroff, D.E., Vitter, J.S.: External-memory graph algorithms. In: Proceedings of SODA’95. (1995) 139–149
5. Valiant, L.G.: A bridging model for parallel computation. Commun. ACM **33**(8) (1990) 103–111
6. Dean, J., Ghemawat, S.: Mapreduce: Simplified data processing on large clusters. Commun. ACM **51**(1) (January 2008) 107–113
7. Karloff, H.J., Suri, S., Vassilvitskii, S.: A model of computation for mapreduce. In Charikar, M., ed.: SODA, SIAM (2010) 938–948
8. Goodrich, M.: Communication-efficient parallel sorting. SIAM J. Comput. **29**(2) (1999) 416–432
9. Goodrich, M., Sitchinava, N., Zhang, Q.: Sorting, searching, and simulation in the mapreduce framework. In: ISAAC. Volume 7074 of LNCS., Springer (2011) 374–383
10. Arge, L., Goodrich, M., Nelson, M., Sitchinava, N.: Fundamental parallel algorithms for private-cache chip multiprocessors. In: SPAA 2008. (2008) 197–206
11. Greiner, G.: Sparse Matrix Computations and their I/O Complexity. Dissertation, Technische Universität München, München (2012)
12. Karp, R.M., Ramachandran, V.: Handbook of theoretical computer science. (1990) 869–941
13. Vishkin, U.: Randomized speed-ups in parallel computation. In: STOC. (1984) 230–239
14. Arge, L., Goodrich, M., Sitchinava, N.: Parallel external memory graph algorithms. In: IPDPS, IEEE (2010) 1–11
15. Jacob, R., Lieber, T., Sitchinava, N.: On the complexity of list ranking in the parallel external memory model. CoRR **abs/1406.3279** (2014)
16. Yao, A.C.C.: Probabilistic computations: Toward a unified measure of complexity (extended abstract). In: FOCS, IEEE Computer Society (1977) 222–227