

Fast and Parallelizable Ranking with Outliers from Pairwise Comparisons

S Im¹ and M Montazer Qaem²

¹ University of California, Merced, CA 95343, USA
sim3@ucmerced.edu

² University of California, Merced, CA 95343, USA
mmontazerqaem@ucmerced.edu

Abstract. In this paper, we initiate the study of the problem of ordering objects from their pairwise comparison results when allowed to discard up to a certain number of objects as outliers. More specifically, we seek to find an ordering under the popular Kendall tau distance measure, i.e., minimizing the number of pairwise comparison results that are inconsistent with the ordering, with some outliers removed. The presence of outliers challenges the assumption that a global consistent ordering exists and obscures the measure. This problem does not admit a polynomial time algorithm unless $NP \subseteq BPP$, and therefore, we develop approximation algorithms with provable guarantees for all inputs. Our algorithms have running time and memory usage that are almost linear in the input size. Further, they are readily adaptable to run on massively parallel platforms such as MapReduce or Spark.

Keywords: Rank aggregation · outliers · approximation · distributed algorithms.

1 Introduction

Ranking is a fundamental problem arising in various contexts, including web pages ranking, machine learning, data analytics, and social choice. It is of particular importance to order n given objects by aggregating pairwise comparison information which could be inconsistent. For example, if we are given $A \prec B$ (meaning that B is superior to A) and $B \prec C$, it would be natural to deduce a complete ordering, $A \prec B \prec C$. However, there exists no complete ordering creating no inconsistencies, if another pairwise comparison result $C \prec A$ is considered together. As a complete ordering/ranking is sought from partial orderings, this type of problems is called rank aggregation and has been studied extensively, largely in two settings: (i) to find a true ordering (as accurately as possible) that is assumed to exist when some inconsistencies are generated according to a certain distribution; and (ii) to find a ranking that is the closest to an arbitrarily given set of the partial orderings for a certain objective, with no stochastic assumptions.

This paper revisits a central rank aggregation problem in the second domain, with a new angle.

The Minimum Feedback Arc Set Tournament Problem (FAST). The input is a tournament.³ The goal is to delete the minimum number of edges in order to make the resulting graph acyclic. This is equivalent to finding a linear ordering of the vertices to incur the fewest 'backward' edges.

The FAST problem is well-studied: it does not admit a polynomial time algorithm unless $\text{NP} \subseteq \text{BPP}$ [3] and there are several constant approximations known for the problem [3,18,11]. We note that this is a special case of the more general problem (*Minimum Feedback Arc Set; FAS for short*) where the input graph is not necessarily complete.

1.1 Necessity of Another Measure for Inconsistencies

The FAST problem measures the quality of an ordering by the number of pairs that are inconsistent with the ordering; this measure was proposed by Kemeny [17] and is also known as the Kendall tau distance. Unfortunately, this measure fails to capture the 'locality' of inconsistencies, namely whether or not the inconsistencies are concentrated around a small number of objects. To see this, consider the two instances in Figures 1.




Fig. 1: The left graph consists of $\frac{n}{3}$ disjoint triangles (the rest of the arcs (v_i, v_j) for all $i < j$ are omitted). In the right graph, v_1, v_2, \dots, v_{n-1} has an ordering with no inconsistencies among them and v_n is a successor of other odd-indexed vertices while being a predecessor of even-indexed vertices.

It is easy to see no matter how we order the vertices in the left instance, we end up with having at least $n/3$ pairs that are inconsistent with the ordering, one from each triangle. Likewise, the number of inconsistent pairs in the right instance is at least $(n-1)/2$ in any ordering, one from each triangle $\{v_n, v_{2i-1}, v_{2i}\}$. Since the ordering, v_1, v_2, \dots, v_n , creates $O(n)$ inconsistent pairs in both examples, the optimal objective is $\Theta(n)$ for both. However, the two instances have inconsistencies of very different natures. In the first graph, the inconsistent pairs are scattered all over the graph. In contrast, in the second graph, the inconsistent pairs are concentrated on the single vertex v_n – the second input graph becomes acyclic when v_n is removed.

The above two examples raise a question if the Kendall tau distance alone is an effective measure in the presence of outliers. We attempt to order or rank objects because we believe that they are comparable to one another, and therefore, there exists an ordering well summarizing the whole set of pairwise comparisons. However, this belief is no longer justified if there are some outliers that do not belong to the 'same' category of comparable objects. In the second input

³ A directed graph $G = (V, E)$ is called a tournament if it is complete and directed. In other words, for any pair $u \neq v \in V$, either $(u, v) \in E$ or $(v, u) \in E$.

graph, perhaps, we should think of v_n as an outlier, as the input graph has no inconsistencies without it. Counting the number of inconsistent pairs could fail to capture the quality of the output ordering without removing outliers. This is the case particularly because we can only hope for approximate solutions and noises incurred by outliers could render multiplicative approximation guarantees not very interesting.

Motivated by this, we propose a new measure that factors in outlier vertices as well as inconsistent pairs/edges:

The Minimum Feedback Arc Set Tournament with Outliers Problem (FASTO). This is a generalization of FAST. As in FAST, we are given as input a tournament $G = (V, E)$, along with a pair of target numbers, (x^*, y^*) . A pair $(V' \subseteq V, E' \subseteq E)$ is a feasible solution if $(V \setminus V', E \setminus E')$ is a DAG – we refer to V' as the outlier set and E' as the backward edge set. Here, E' is a subset of edges between $V \setminus V'$. The solution quality is measured as $(\frac{|V'|}{x^*}, \frac{|E'|}{y^*})$, which are the number of outliers and backward edges incurred, respectively, relative to the target numbers, x^* and y^* . We can assume w.l.o.g. that $x^* > 0$ since otherwise FASTO becomes exactly FAST.

We will say that a solution is (α, β) -approximate if $|V'| \leq \alpha x^*$ and $|E'| \leq \beta y^*$. An algorithm is said to be a (α, β) -approximation if it always produces a (α, β) -approximate solution for all inputs. Here, it is implicitly assumed that there is a feasible solution (V', E') w.r.t. (x^*, y^*) , i.e., $|V'| \leq x^*$ and $|E'| \leq y^*$ – if not, the algorithm is allowed to produce any outputs. Equivalently, the problem can be viewed as follows: Given that we can remove up to αx^* vertices as outliers, how many edges do we need to flip their directions so as to make the input graph acyclic. But we assume that we are given a target pair (x^*, y^*) , as it makes our presentation cleaner.

1.2 Our Results and Contributions

Throughout this paper, we use n to denote the number of vertices in the input graph and $N = \Theta(n^2)$, which is the asymptotic input size. We use \tilde{O} or $\tilde{\Theta}$ to suppress logarithmic factors. Recall that x^* is the target number of outliers. While we propose several algorithms that are scalable and parallelizable, the following is our main theoretical result with performance guarantees for all inputs; the first case is more interesting but we also study the second case for completeness.

Theorem 1. *There is an approximation algorithm for FASTO with $\tilde{O}(N)$ running time and $\tilde{O}(N)$ memory usage that outputs a solution, with probability at least $1/2 - 1/n$, that is*

- $(O(1), O(1))$ -approximate when $x^* \leq \sqrt{n}$ (Section 2); and
- $(O(\log n), O(1))$ -approximate when $x^* > \sqrt{n}$ (Section 3).

Further, this algorithm can be adapted to massively parallel computing platforms so that they run in $O(\log n)$ rounds.

We note that the running time of our algorithm, which is almost linear in the input size, is essentially *optimal*. To see this, consider a simple instance that admits an ordering with one backward edge, with $(x^*, y^*) = (1, 0)$. Then, it is unavoidable to actually find the backward edge, which essentially requires to read all edges. While we do not know how to obtain a constant approximation for the case when $x^* > \sqrt{n}$ in the massively parallel computing setting, we can still get a relatively fast algorithm in the single machine setting. More precisely, we can obtain an $(O(1), O(1))$ -approximate solution for all instances, with probability at least $1/2 - 1/n$, using $\tilde{O}(\sqrt{N}x^{*2})$ running time and $\tilde{O}(\sqrt{N}x^{*2})$ memory. For the formal model of massively parallel computing platforms, see [6].

Below, we outline our contributions.

Proposing a New Metric for Ranking: Ranking with Some Outliers Removed. Outliers have been extensively studied in various fields, such as statistics, machine learning, and databases. Outliers were considered before together with ranking, but they were mostly focused on the evaluation of outliers themselves, e.g., ranking outliers [22]. Our work is distinguished, as we seek to find a clean ordering which otherwise could be obscured by outliers. Various combinatorial optimization problems have been studied in a spirit similar to ours, particularly for clustering problems [9,10,21,15]. For example, the k -means clustering problem can be extended to minimize the sum of squares of distances to the k centers from all points, except a certain number of outliers [16,15]. We feel that it is a contribution of conceptual importance to study ranking problems in this direction for the first time. We believe this new direction is worth further investigation; see Section 5 for future research directions.

Fast and Memory-efficient Algorithms with Provable Guarantees. Our work is inspired by Aboud’s work [1] on a closely related clustering problem – in the Correlational Clustering problem, there is an undirected complete graph where each edge is labeled as either ‘-’ or ‘+’. The goal is to partition the vertices so as to minimize the number of inconsistent pairs where an edge (u, v) labeled ‘+’ (resp., ‘-’) incurs a unit cost if the two vertices u and v are placed into different groups (resp., the same group). This problem is closely related to FAST, and there exist simple and elegant 3-approximate randomized algorithms, called KWIK-SORT, for both problems, which can be seen as a generalization of quicksort: a randomly chosen pivot recursively divide an instance into two sub-instances. In the case of FAST, one subproblem contains the predecessors of the pivot and the other the successors of the pivot. Likewise, in the correlational clustering case, the vertices are divided based on their respective affinities to the pivot.

Aboud considered an outlier version of the Correlational Clustering problem and gave an $(O(1), O(1))$ -approximation.⁴ Not surprisingly, one can adapt his result to obtain a $(O(1), O(1))$ -approximation for FASTO. Unfortunately, Aboud’s

⁴ More precisely, he considered a slightly more general version where each vertex may have a different cost when removed as an outlier.

algorithm uses memory and running time that are at least linear in the number of ‘bad’ triangles, which can be as large as $\Omega(n^3) = \Omega(N^{1.5})$. In our problem, FASTO, a bad triangle v_1, v_2, v_3 is a triangle that does not admit a consistent ranking, i.e., $(v_1, v_2), (v_2, v_3), (v_3, v_1) \in E$ or $(v_2, v_1), (v_3, v_2), (v_1, v_3) \in E$.

To develop a fast and memory-efficient algorithm, we combine sampling with Aboud’s algorithm. This combination is not trivial; for example, applying Aboud’s algorithm to a reduced-sized input via sampling does not work. At a high-level, Aboud’s algorithm uses a primal-dual approach. The approach sets up a linear programming (LP) and solves the LP by increasing the variables of the LP and its dual, where the constructed integer solution to the LP is used to identify outlier vertices. We have to adapt the LP, as we have to carefully handle the sampled points and argue their effects on potential backward edges. After all, we manage to reduce the memory usage and running time to $\tilde{O}(N)$ preserving the approximation factors up to constant factors⁵ under the assumption that the number of outliers is small, which we believe to be reasonable in practice.

Algorithms Adaptable to Massively Parallel Platforms. Finally, our algorithm can be easily adapted to run on massively parallel platforms such as MapReduce or Spark. On such platforms, each machine is assumed to have insufficient memory to store the whole input data, and therefore, multiple machines must be used. Aboud’s algorithm is not suitable for such platforms, as it uses significantly super-linear memory. In contrast, our algorithm uses sampling appropriately to reduce the input size while minimally sacrificing the approximation guarantees. More precisely, our algorithm for FASTO can be adapted to run in $O(\log n)$ rounds on the parallel platforms – the number of rounds is often one of the most important measures due to the huge communication overhead incurred in each round.

As a byproduct, for the first time we show how to convert KWIK-SORT for FAST to the distributed setting in $O(1)$ rounds (see Sections 2.1 and 2.4), which is interesting on its own. The algorithm recursively finds a pivot and divides a subset of vertices into two groups, thus obtaining $O(\log n)$ rounds is straightforward. But we observe that as long as the pivot is sampled uniformly at random from each subgroup, the algorithm’s performance guarantee continues to hold. Therefore, the algorithm still works even if we consider vertices in a random order as pivots – the sub-instance including the pivot is divided into two. Thus, we construct a decision tree from a prefix of the random vertex ordering, and using this decision tree, we partition the vertex set into multiple groups in a distributed way. This simple yet important observation also plays a key role in breaking bad triangles, thus reducing the memory usage of our algorithm for FASTO.

⁵ We show that our algorithm is (180, 180)-approximate, which can be improved arbitrarily close to (60, 60) if one is willing to accept a lower success probability. In contrast, Aboud’s algorithm can be adapted to be (18, 18)-approximate for FASTO; however as mentioned above, it uses considerably more memory and run time than ours.

1.3 Other Related Work

The only problem that studies ranking with the possibility of removing certain outlier vertices, to our best knowledge, is the Feedback Vertex Set problem (FVS). The FVS problem asks to remove a minimum subset of vertices to make the remaining graph acyclic. It is an easy observation that FVS is a special case of our problem with $y^* = 0$ if the input graph were an arbitrary directed graph, not just a tournament. The FVS problem is known to be NP-hard and the current best approximation for the problem has an approximation factor of $O(\log n \log \log n)$ [13]. Thus, if we allow the input graph to be an arbitrary directed graph for our problem FASTO, then we cannot hope for a better than $(O(\log n \log \log n), c)$ -approximation for any $c > 0$ unless we improve upon the best approximation for FVS. We are not aware of any other literature that considers rank aggregation with the possibility of removing outlier vertices, with the exception of the aforementioned Aboud’s work on a closely related clustering problem [1]. Due to the vast literature on the topic, we can only provide an inherently incomplete list of work on ranking without outliers being considered. There exist several approximation algorithms for FAST. As mentioned, Ailon et al. [3] give the randomized KWIK-SORT that is 3-approximate for the problem, which can be derandomized [25]. Also, the algorithm that orders vertices according to their in-degrees is known to be a 5-approximation [11]. Kenyon-Mathieu and Schudy [18] give a PTAS; a PTAS is a $(1 + \epsilon)$ -approximate polynomial-time algorithm for any fixed $\epsilon > 0$. The complementary maximization version (maximizing the number of forward edges in the linear ordering) was also studied, and PTASes are known for the objective [14,7]. For partial rankings, see [2] and pointers in the paper. Extension to online and distributed settings are studied in [26] but the performance guarantee is not against the optimal solution, but against a random ordering, which incurs $\Theta(n^2)$ backward edges. For another extensive literature on stochastic inputs, see [4,19,5,23,12,24] and pointers therein.

1.4 Notation and Organization

We interchangeably use $(u, v) \in E$ and $u \prec v$ – we will say that u is v ’s predecessor, or equivalently, v is u ’s successor. We use \tilde{O} or $\tilde{\Theta}$ to suppress logarithmic factors in the asymptotic quantities. We use n to denote the number of vertices in the input graph and $N = \Theta(n^2)$ to denote the asymptotic input size. The subgraph of G induced on V' is denoted $G[V']$. Let $[k] := \{1, 2, 3, \dots, k\}$.

In Section 2, we present our algorithm for FASTO when the target number of outliers is small, i.e., $x^* \leq \sqrt{n}$. The other case is handled in Section 3. In Section 4, we evaluate our algorithms and other heuristics we develop via experiments. In Section 5, we close with several interesting future research directions. Due to the space constraints, we will defer most analysis to the full version of this paper.

2 When the Target Number of Outliers x^* is Small

Our algorithm when $x^* \leq \sqrt{n}$ consists of three main steps:

Step 1: Partitioning via KWIK-SORT-SEQ on Sample. Each vertex in V is sampled independently with probability $\frac{1}{4x^*}$ and placed into S . Randomly permute S and run KWIK-SORT-SEQ on the ordered set S to construct a decision tree $\tau(S)$. Let $k = |S|$. Partition the other vertices, $V \setminus S$, into $k + 1$ groups, V_1, V_2, \dots, V_{k+1} , using $\tau(S)$.

Step 2: Identifying Outliers via Primal-Dual. Formulate Linear Programming (LP) relaxation and derive its dual. Solve the primal and dual LPs using a primal-dual method, which outputs the set of outliers to be chosen.

Step 3: Final Ordering. Run KWIK-SORT on the non-outlier vertices in each group V_i , $i \in \{1, 2, \dots, k + 1\}$.

In the following, we give a full description of all the steps of our algorithm; the last step is self-explanatory, and thus, is briefly discussed at the end of Section 2.2. Following the analysis of the algorithm, the extension to the distributed setting is discussed in the final subsection.

2.1 Step 1: Partitioning via KWIK-SORT-SEQ

Recall the original KWIK-SORT algorithm [3], which was shown to be a 3-approximation for FAST. The above KWIK-SORT algorithm is described in

Algorithm 1 KWIK-SORT ($G = (V, A)$)

- 1: If $V = \emptyset$, then return empty-list
 - 2: Set $V_L \rightarrow \emptyset, V_R \rightarrow \emptyset$
 - 3: Pick a random pivot $i \in V$ uniformly at random
 - 4: For all vertices $j \in V \setminus \{i\}$:
 - 5: If $(j, i) \in A$, then add j to V_L (place j on left side)
 - 6: Else add j to V_R (place j on right side).
 - 7: $G_L = G[V_L] \leftarrow$ the tournament induced on V_L .
 - 8: $G_R = G[V_R] \leftarrow$ the tournament induced on V_R .
 - 9: Return KWIK-SORT (G_L), i , KWIK-SORT (G_R)
(Concatenation of left recursion, i , and right recursion.)
-

a divide-and-conquer manner. As usual, there are multiple ways to serialize a parallel execution – for example, lines (7) and (8) can be reversed in the order of execution. So, as long as we ensure that a pivot is sampled from each subgraph *uniformly at random* for further partitioning, we can simulate the parallel execution keeping the approximation guarantee. Here, we introduce one specific simulation, KWIK-SORT-SEQ, which generates a random ordering of V , takes a pivot one by one from the random ordering, and gradually refines the partitioning. We show that this is indeed a valid way of simulating KWIK-SORT.

Lemma 1. *KWIK-SORT-SEQ is a legitimate way of simulating KWIK-SORT, keeping the approximation guarantee. Therefore, KWIK-SORT-SEQ is a 3-approximation for FAST.*

Algorithm 2 KWIK-SORT-SEQ ($G = (V, A)$)

-
- 1: $\pi(V) \leftarrow$ a random permutation of V
 - 2: $\mathcal{V} = \{V\}$
 - 3: For $i = 1$ to $n = |V|$:
 - 4: $\pi_i(V) \leftarrow$ i th vertex in the ordering $\pi(V)$
 - 5: Let $V' \in \mathcal{V}$ be such that $\pi_i(V) \in V'$
 - 6: $V'_L, V'_R \leftarrow \emptyset$
 - 7: For all vertices $j \in V' \setminus \{\pi_i(V)\}$:
 - 8: If $(j, i) \in A$, then add j to V'_L ; else add j to V'_R
 - 9: Put $V'_L, \{\pi_i(V)\}, V'_R$ in place of V' , in this order
 - 10: Return \mathcal{V} (Order vertices in the same order they appear in \mathcal{V}).
-


Due to the space constraints, we defer the proof to the full version of this paper. Note that a fixed random permutation $\pi(V)$ completely determines the final ordering of vertices. Likewise, a fixed length- i prefix of $\pi(V)$ completely determines the intermediate partitioning \mathcal{V} after $\pi_1(V), \pi_2(V), \dots, \pi_i(V)$ being applied, and the partitioning only refines as more pivots are applied. Thus, we can view this intermediate partitioning as the classification outcome of $V \setminus S$ via the decision tree $\tau(S)$ generated by KWIK-SORT-SEQ on the ordered set $S = \{\pi_1(V), \pi_2(V), \dots, \pi_i(V)\}$. See Fig. 2 for illustration.

The first step of our algorithm is essentially identical to what KWIK-SORT-SEQ does, except that our algorithm only needs a prefix of the random permutation, not the entire $\pi(V)$. It is an easy observation that sampling each vertex independently with the same probability and randomly permuting them is a valid way of getting a prefix of a random permutation. We note that we sample each vertex with probability $\frac{1}{4x^*}$, to avoid sampling any outlier (in the optimal solution) with a constant probability.

2.2 Step 2: Identifying Outliers

To set up our linear programming relaxation, we first need some definitions. We consider the ordered sample S and the induced groups V_1, V_2, \dots, V_{k+1} in the order they appear in the linear ordering produced by KWIK-SORT-SEQ performed in Step 1. For notational convenience, we reindex the sampled points so that they appear in the order of s_1, s_2, \dots, s_k . We classify edges into three categories: Let E_{in} be the edges within the groups, S_{back} the backward edges

Fig. 2: Illustration of the construction of $\tau(S)$ and partitioning of $V \setminus S$ via $\tau(S)$. In this example, the decision tree $\tau(S)$ is induced by a(n ordered) sample $S = \{s_1, s_2, s_3\}$, where $s_1 \prec s_2, s_3 \prec s_2$, and $s_1 \prec s_3$. If a vertex $v \in V \setminus S$ is such that $s_1 \prec v, v \prec s_2, s_3 \prec v$, then v is placed into G_3 .



with both end points in S , and E_{back} the backward edges $e = (u, v)$ such that $u \in V_i, v \in V_j$ for $i > j$; or $u = s_i, v \in V_j$ for $i \geq j$; or $v = s_i, u \in V_j$ for $i \leq j$. See Fig. 3.

Finally, we let T_{in} be the set of bad triangles with all vertices in the same group; recall that a bad triangle is a cycle of length 3. We are now ready to define an integer programming (IP) for a penalty version of FASTO, where a backward edge incurs a unit penalty and each outlier incurs $c := y^*/x^*$ units of penalty:

$$LP_{fasto}^{primal}(G) := \min \sum_{e \in E_{in}} y_e + \sum_{x \in V \setminus S} p_x \cdot c + \sum_{e \in E_{back}} z_e + |S_{back}| \quad (\text{PRIMAL})$$

$$s.t. \quad \sum_{e \subset t} y_e + \sum_{x \in t} p_x \geq 1 \quad \forall t \in T_{in} \quad (1)$$

$$p_u + p_v + z_e \geq 1 \quad \forall e = (u, v) \in E_{back} : \{u, v\} \cap S = \emptyset \quad (2)$$

$$p_u + z_e \geq 1 \quad \forall e = (u, v) \in E_{back} : v \in S \wedge u \in V \setminus S \quad (3)$$

$$p_v + z_e \geq 1 \quad \forall e = (u, v) \in E_{back} : u \in S \wedge v \in V \setminus S, \quad (4)$$

over variables $y_e \geq 0$ for all $e \in E_{in}$, $p_x \geq 0$ for all $x \in V \setminus S$, and $z_e \geq 0$ for all $e \in E_{back}$.




Fig. 3: Classification of edges. The rectangles shown represent groups V_1, V_2, V_3, V_4 from the left. Edges in S_{back} are shown as solid arcs and edges E_{back} as dotted arcs. Edges in E_{in} are those within groups and are omitted.

This IP has the following interpretation: an edge e in E_{in} (E_{back} , resp.) becomes backward if $y_e = 1$ ($z_e = 1$, resp.). Assuming that we will choose no sampled points as outliers, all edges in S_{back} will become backward. And each non-sampled point x incurs penalty c if it is chosen as an outlier when $p_x = 1$. Constraint (1) follows from the fact that for each bad triangle t , at least one edge e of t must become backward unless t is broken; a triangle gets broken when at least one of its vertices is chosen as outlier. The other constraints force each edge in E_{back} to become backward if its neither end point is chosen as outlier. A Linear Programming (LP) relaxation, which will be referred to as LP_{fasto}^{primal} , is obtained by allowing variables to have fractional values. Using the fact that KWIK-SORT is a 3-approximation, assuming that there is a feasible solution w.r.t. the target pair (x^*, y^*) , conditioned on the sample being disjoint from the feasible solution's outlier, we can argue that the expected optimal objective of LP_{fasto}^{primal} is at most $4y^*$ (Lemma 4).

To obtain an approximate integer solution to LP_{fasto}^{primal} , we will use the primal-dual method. Primal-dual is a common technique for designing approximation algorithms [27]. The dual LP is shown below.

$$LP^{dual} = \max \sum_{t \in T_{in}} \alpha_t + \sum_{e \in E_{back}} \beta_e + |S_{back}| \quad (\text{DUAL})$$

$$s.t. \quad \sum_{e \subset t} \alpha_t \leq 1 \quad \forall e \in E_{in} \quad (5)$$

$$\sum_{x \in t} \alpha_t + \sum_{x \in e} \beta_e \leq c \quad \forall x \notin S \quad (6)$$

$$\beta_e \leq 1 \quad \forall e \in E_{back} \quad (7)$$

To develop an algorithm based on a primal-dual method, we replace Constraint (6) with two sufficient conditions (8) and (9).

$$\sum_{x \in t} \alpha_t \leq \frac{3}{5}c \quad \forall x \in V \setminus S \quad (8)$$

$$\sum_{x \in e} \beta_e \leq \frac{2}{5}c \quad \forall x \in V \setminus S \quad (9)$$

Below, we give our algorithm, **Algorithm 3**, that sets the variables of our primal with the help of the above dual. Although the algorithm updates all variables, the only information we need to run the final Step 3 is the outlier set U , as Step 3 runs KWIK-SORT on each group with vertices U removed as outliers. But the other outputs will be useful for the analysis of our algorithm. Note that although the dual variables can have fractional values, the primal variables will only have integer values. Note that E' and E^2 represent the backward edges within the groups and across the groups, respectively; since our algorithm samples no outlier vertices in the optimal solution, all edges in S_{back} become backward and they are counted separately.

Algorithm 3 Primal-Dual Algorithm

- 1: *Initialization:* $p \leftarrow 0, y \leftarrow 0, z \leftarrow 0, \alpha \leftarrow 0, \beta \leftarrow 0, U \leftarrow \emptyset, E' \leftarrow \emptyset, E^2 \leftarrow \emptyset$. Initially, all α_t and β_e variables are active.
 - 2: **while** \exists active dual variables α_t or β_e **do**
 - 3: Uniformly increase active dual variables until Constraints (5), (7), (8) or (9) become tight.
 - 4: **for** each $e \in E_{in}$ s.t. $\sum_{e \subset t} \alpha_t = 1$, **do** add e to E' ; and inactivate α_t for all t s.t. $e \subset t$.
 - 5: **for** each $x \in V \setminus S$ s.t. $\sum_{x \in t} \alpha_t = \frac{3}{5}c$, **do** add x to U ; and inactivate α_t for all t s.t. $x \in t$
 - 6: **for** each $x \in V \setminus S$ s.t. $\sum_{x \in e} \beta_e = \frac{2}{5}c$, **do** add x to U ; and inactivate β_e for all e s.t. $x \in e$.
 - 7: **for** each $e \in E_{back}$ s.t. $\beta_e = 1$, **do** add e to E^2 ; and inactivate β_e .
 - 8: Remove from E' and E^2 all the edges e with $e \cap U \neq \emptyset$.
 - 9: **for** $e \in E'$ **do** $y_e \leftarrow 1$; **for** $\forall e \in E^2$ **do** $z_e \leftarrow 1$; and **for** $\forall x \in U$ **do** $p_x \leftarrow 1$
 - 10: **return** p, y, z, U, E', E^2
-

2.3 Sketch of the Analysis: Approximation Guarantees, Memory Usage, and Running Time

In this subsection, we only give a sketch of the analysis due to space constraints. In Step 1, we can show that the sample S is disjoint from the optimal solution's

outlier set, with a constant probability (at least $3/4$). Conditioned on that, as mentioned earlier, the expected optimal objective of LP_{fasto}^{primal} can be shown to be at most $4y^*$. The primal-dual method in Step 2 obtains an integer solution to LP_{fasto}^{primal} that is a constant approximate against the optimal LP objective, which is established by LP duality. Therefore, the outlier set U 's contribution to the LP_{fasto}^{primal} 's objective is $c|U| = (y^*/x^*) \cdot |U|$ and it is upper bounded by $O(1)y^*$. This shows our algorithm chooses at most $O(1)x^*$ outliers. We now turn our attention to upper bounding the number of backward edges. The primal solution to LP_{fasto}^{primal} gives the number of backward edges within groups 'covering' all the unbroken triangles, which upper bounds the minimum number of backward edges achievable within groups [3], and counts the number of other backward edges explicitly by z_ϵ and $|S_{back}|$. Since the latter is determined by the partial ordering produced by Step 1 and U and we run a 3-approximate KWIK-SORT on each group, we can also establish that the final number of backward edges output is $O(1)y^*$.

Now we discuss our algorithm's memory usage and running time. We show in Lemma 7 that each group size is $\tilde{O}(n/k)$ if $|S| = \Omega(\log n)$. This requires us to prove that a randomly chosen pivot partitions a problem into two subproblems of similar sizes with a constant probability, meaning that there is a large fraction of vertices that have similar numbers of predecessors and successors. Then, the total number of bad triangles within groups is $(\tilde{O}(n/k))^3 \cdot k = \tilde{O}(n^2) = \tilde{O}(N)$ when $k \simeq \frac{n}{x^*} \geq \sqrt{n}$, as desired. Since the number of variables in our algorithm, particularly in Step 2 is dominated by the number of bad triangles in consideration and edges, it immediately follows that the memory usage is $\tilde{O}(N)$. Further, one can increase dual variables by a factor of $(1 + \epsilon)$ in each iteration for an arbitrary constant precision parameter $\epsilon > 0$, starting from $1/n^2$. Using this one can show the number of iterations needed is $O(\log n)$. This immediately leads to the claim that the running time is $\tilde{O}(n^2) = \tilde{O}(N)$ when $x^* \leq \sqrt{n}$.

2.4 Extension to the Distributed Setting

Due to space constraints, in this subsection, we briefly discuss how we can adapt the algorithm to run in a distributed way. For formal theoretical models of the distributed setting we consider, see [6]. We assume that the input graph is stored across machines arbitrarily. Clearly, Step 1 of taking a sample S can be done in parallel. All edges between points in S are sent to a machine to construct the decision tree $\tau(S)$. The decision tree is broadcast to all machines to partition vertices in $k+1$ groups in a distributed way. If FAST were the problem considered, we would sample each vertex with probability $1/\sqrt{n}$ and move the subgraph induced on each group G_i to a machine and continue to run KWIK-SORT on the subgraph. Then, we can implement KWIK-SORT to run in $O(1)$ rounds assuming that each machine has $\tilde{\Omega}(n)$ memory. If FASTO is the problem considered, we can implement Step 2 in $O(\log n)$ rounds, as discussed in the previous subsection. Step 3, which is the execution of KWIK-SORT on each group, can be run in one round. Again, the only constraint is that each machine has $\tilde{\Omega}(n)$ memory for an arbitrary number of machines.

3 When the Target Number of Outliers x^* is Large

This section is devoted to proving Theorem 1 for the case $x^* > \sqrt{n}$. Our algorithm when $x^* \geq \sqrt{n}$ builds on the following key lemma.

Lemma 2. *Let $G' = (V, E')$ denote a subgraph of G where E' is the set of backward edges produced by KWIK-SORT-SEQ on $G = (V, E)$. There exists $V' \subseteq V$ with probability at least $1/2$ such that $|V'| \leq 260(\log n)x^*$ and $|E'[V \setminus V']| \leq 12y^*$.*

Using this lemma and a simpler primal-dual based algorithm, we can find a $(O(\log n), O(1))$ -approximate solution with probability at least $1/2$. Since we do not have to deal with bad triangles at all, the analysis of memory usage and running time and the extension to the distributed setting only become easier. The full algorithm and analysis are deferred to the full version of this paper.

4 Experiments

In this section, we perform experiments to evaluate our algorithm against synthetic data sets. All experiments were performed on Ubuntu 14.04 LTS with RAM size 15294 MB and CPU Intel(R) Xeon(R) CPU E5-2670 v2 @ 2.50GHz. We implemented the following four algorithms including ours for comparison. The last two (RSF and IOR) are new heuristics we developed in this paper, but with no approximation guarantees. We assume that all algorithms are given a ‘budget’ B on the number of outliers, which limits the number of vertices that can be chosen as outliers.

- Primal-Dual with Sampling (PDS): Our algorithm when x^* is small. We run the algorithm for all target pairs (x^*, y^*) of powers of two, where $x^* \in [0, B]$, $y^* \in [0, B']$, and choose the best solution with at most $1.5B$ outliers. Here, B' is the number of back edges output by KWIK-SORT. Note that we allow PDS to choose slightly more outliers although it may end up with choosing less since the only guarantee is that PDS chooses up to $O(x^*)$ outliers. The running time is summed up over all pairs.
- Aboud’s algorithm (ABD): The algorithm in [1] for Correlational Clustering is adapted to FASTO. As above, the best solution with at most $1.5B$ outliers is chosen over all the above target pairs. ABD is essentially PDS with $S = \emptyset$ in Step 1. The running time is again summed up over all pairs considered.
- Random Sample Filter (RSF): Take a random Sample S from V (for fair comparison, the same sample size is used as in PDS). Order S using KWIK-SORT and let $\pi(S)$ be the resulting order. For each $v \in V \setminus S$, let $b(v)$ be the minimum number of backward edges with v as one endpoint over all these $|S| + 1$ possible new permutations created by adding v to $\pi(S)$. Outputs B vertices $v \in V \setminus S$ with the highest $b(v)$ values as the outliers and order the remaining vertices by KWIK-SORT.

Algorithm 4 Iterative Outlier Removal (IOR)

-
- 1: $U = \emptyset$
 - 2: **while** $|U| < B$ **do**
 - 3: For each vertex $v \in V$, run KWIK-SORT on $V \setminus v$, which yields $|V|$ permutations. Among all the achieved permutations, consider the one with minimum number of backward edges. Let $v \in V$ be the missing vertex in this permutation. Add v to U ; and $V \leftarrow V \setminus \{v\}$.
 - 4: Output U as outliers and run KWIK-SORT on V/U .
-

- Iterative Outlier Removal (IOR): We iteratively remove the vertex without which KWIK-SORT outputs the least back edges. See Algorithm 4 for more details.

We mainly use two natural models to generate synthetic data sets. The first model, which we call the *uniform* model, assumes inconsistencies uniformly scattered over edges, in addition to randomly chosen outlier vertices. More precisely, the uniform model is described by a quadruple $\langle p, q, r, n \rangle$: For a tournament $G = (V, E)$ over n vertices with no inconsistencies, flip each edge with probability p and perturb each vertex v with probability q by flipping each edge incident on v with probability r – all independently. The second model, which essentially follows the *Bradley-Terry-Luce* model [8,20] and therefore we refer to as BTL, assumes that each vertex i has a certain (hidden) score $w_i > 0$. Then, each pair of vertices i and j has edge (i, j) with probability $\frac{w_j}{w_i + w_j}$; or edge (j, i) with the remaining probability. Intuitively, if edge (i, j) is present for every i, j such that $w_i < w_j$, we will have a tournament that is a DAG. However, some edges are flipped stochastically – in particular, edges between vertices with similar scores are more likely to be flipped. Assuming that the underlying vertex scores are a geometric sequence, we can compactly describe a BTL instance by a quadruple $\langle b, q, r, n \rangle$, where the score of the n vertices forms a geometric sequence of ratio b . Then, edges are first generated as described above, and vertices are perturbed using the parameters q and r as are done for the uniform model.

We first confirm that ABD is not very scalable.

Table 1: PDS with sample size $1/(2q) = 50$ vs. ABD for the uniform model $\langle p = 0.001, q = 0.01, r = 0.50, n \rangle$ and $B = nq$. Bad Δ s denotes the number of bad triangles in 1000 units.

pts (n)	outliers		back edges		bad Δ s		time (sec)	
	PDS	ABD	PDS	ABD	PDS	ABD	PDS	ABD
250	0	2	269	33	0	18	0.2	3.8
500	6	7	377	152	365	228	4.8	100.0
1000	9	N/A	1597	N/A	441	1032	20.0	600+

Even when $n = 1000$, ABD does not terminate in 600 seconds while our algorithm PDS does in 20 seconds. Our algorithm’s speed-up is explained by the significantly smaller number of bad triangles. For $n = 500$, PDS outputs factor 2 or 3 more back edges than ABD. But we were not able to compare the two algorithms for larger inputs because of ABD’s large run time.

Next, we compare PDS to RSF and IOR for inputs generated under the uniform and BTL models. Note that RSF and IOR choose exactly the same number of outliers, B .

Table 2: PDS with sample size $1/(2q) = 50$ vs RSF vs IOR for the uniform input ($p = 0.001, q = 0.01, r = 0.50, n$) and $B = nq$.

pts (n)	outliers		back edges (10^3)			time (sec)		
	RSF	PDS	RSF	IOR	PDS	RSF	IOR	PDS
500	5	6	1.1	0.98	0.4	0.1	0.1	4.7
1000	10	9	2.6	1.8	1.6	0.4	0.5	20.3
2000	20	16	7.5	6.1	2.5	1.1	3.8	145.4
4000	40	55	26.2	71.5	11.8	3.8	41.3	1132

Table 3: PDS with sample size $1/(2q) = 50$ vs RSF vs IOR for the BTL input with $q = 0.01, r = 0.50, (b, n) \in \{(4.2, 500), (2.3, 1000), (1.57, 2000), (1.266, 4000)\}$; here $B = nq$.

pts (n)	outliers		back edges (10^3)			time (sec)		
	RSF	PDS	RSF	IOR	PDS	RSF	IOR	PDS
500	5	1	0.3	0.0	0.3	0.1	0.1	1
1000	10	9	0.8	0.3	0.8	0.3	0.5	8.3
2000	20	26	7.9	7.6	1.6	0.9	2.3	73
4000	40	33	8.1	8.0	7.9	3.4	15.9	381.3

Our algorithm PDS outperforms the other two in terms of the number of back edges although it occasionally chooses slightly more outliers. However, PDS is considerably slower than RSF and IOR.

Finally, to showcase the major advantage of our algorithm PDS that it has performance guarantees for all inputs, in contrast to the two heuristics RSF and IOR, we consider certain specific instances. First, we observe that RSF significantly underperforms compared to PDS and IOR when the instance is constructed by choosing \sqrt{n} points at random and flipping edges among them. As before, note that RSF and IOR choose the same number of outliers, B , thus we only display RSF for the outliers column.

Table 4: PDS with sample size $\sqrt{n}/2$ vs RSF with sample size $\sqrt{n}/2$ vs IOR. Each vertex is perturbed with probability $1/\sqrt{n}$ – each edge between perturbed vertices is flipped with probability $1/2$. $B = \sqrt{n}$.

pts (n)	outliers		back edges			time (sec)		
	RSF	PDS	RSF	IOR	PDS	RSF	IOR	PDS
1000	31	31	226	0	0	0.2	1.2	6.4
2000	44	44	541	0	0	0.8	7.3	27.2
4000	63	63	1570	0	0	3.3	62.8	130.2
8000	89	89	2887	0	0	13.4	442.0	702.7

As shown in Table 4, when $n = 4000$, all algorithms choose exactly 63 outliers; but RSF produces 1570 back edges while the other two produce no back edges. For all cases when $n = 1000, 2000, 4000$ and 8000 , PDS and IOR create no back edges while RSF does a considerable number of back edges. Interestingly, IOR appears not to be very scalable. For $n = 8000$, IOR is only twice faster than PDS.

We continue to observe that IOR also quite underperforms compared to PDS for a certain class of instances. The instance we create is parameterized by t . The instance is constructed by combining $4t$ sub-tournaments, $G_1, G_2, G_3, \dots, G_{4t}$, which are identical. Each G_i has t vertices and admits a perfect ordering with one vertex removed therein – each edge in G_i incident on the vertex is flipped with probability $1/2$. We connect the sub-tournaments, so that for any $i < j$, all vertices in G_i are predecessors of all vertices in G_j . As shown in Table 5, when $n = 4096$, PDS creates no back edges while IOR creates 313 back edges; both chooses the same number of outliers, 128. Further, PDS is slower than IOR only by a factor of at most 3.

Table 5: PDS with sample size $t/2$ vs IOR

pts (n)	outliers		back edges		time (sec)	
	IOR	PDS	IOR	PDS	IOR	PDS
1024	64	62	87	7	2.4	21.4
2116	92	89	142	79	16.7	203.2
4096	128	128	313	0	129.5	305.5

5 Conclusions

In this paper, we studied how to order objects in the presence of outliers. In particular, we developed approximation algorithms that are nearly optimal in terms of running time and can be adapted to the distributed setting, along with potentially useful heuristics. There are many interesting future research directions. First, our algorithm may choose more than x^* outliers. It would be interesting if one can get an approximation algorithm that finds an ordering resulting in $O(1)y^*$ backward edges with strictly no more than x^* outliers. Second, we currently do not know if it is possible to obtain a $(O(1), O(1))$ -approximation algorithm whose running time is almost linear in the input size when $x^* > \sqrt{n}$. Finally, it

would be of significant interest to consider arbitrary directed graphs, not only tournaments, as input.

References

1. Aboud, A.: Correlation clustering with penalties and approximating the reordering buffer management problem. Master's thesis. The Technion Israel Institute of Technology (2008)
2. Ailon, N.: Aggregation of partial rankings, p -ratings and top- m lists. *Algorithmica* **57**(2), 284–300 (2010). <https://doi.org/10.1007/s00453-008-9211-1>, <https://doi.org/10.1007/s00453-008-9211-1>
3. Ailon, N., Charikar, M., Newman, A.: Aggregating inconsistent information: ranking and clustering. *Journal of the ACM* **55**(5), 23 (2008)
4. Altman, A., Tennenholtz, M.: Ranking systems: the pagerank axioms. In: ACM EC (2005)
5. Ammar, A., Shah, D.: Ranking: Compare, don't score. In: IEEE Allerton (2011)
6. Andoni, A., Nikolov, A., Onak, K., Yaroslavtsev, G.: Parallel algorithms for geometric graph problems. In: Symposium on Theory of Computing, STOC 2014, New York, NY, USA, May 31 - June 03, 2014. pp. 574–583 (2014). <https://doi.org/10.1145/2591796.2591805>, <http://doi.acm.org/10.1145/2591796.2591805>
7. Arora, S., Frieze, A., Kaplan, H.: A new rounding procedure for the assignment problem with applications to dense graph arrangement problems. *Math programming* **92**(1), 1–36 (2002)
8. Bradley, R.A., Terry, M.E.: Rank analysis of incomplete block designs: I. the method of paired comparisons. *Biometrika* **39**(3/4), 324–345 (1952)
9. Charikar, M., Khuller, S., Mount, D.M., Narasimhan, G.: Algorithms for facility location problems with outliers. In: ACM-SIAM SODA (2001)
10. Chen, K.: A constant factor approximation algorithm for k-median clustering with outliers. In: ACM-SIAM SODA (2008)
11. Coppersmith, D., Fleischer, L.K., Rurda, A.: Ordering by weighted number of wins gives a good ranking for weighted tournaments. *ACM Transactions on Algorithms* **6**(3), 55 (2010)
12. Duchi, J.C., Mackey, L.W., Jordan, M.I.: On the consistency of ranking algorithms. In: ICML. pp. 327–334 (2010)
13. Even, G., (Seffi) Naor, J., Schieber, B., Sudan, M.: Approximating minimum feedback sets and multicuts in directed graphs. *Algorithmica* **20**(2), 151–174 (Feb 1998). <https://doi.org/10.1007/PL00009191>, <https://doi.org/10.1007/PL00009191>
14. Frieze, A., Kannan, R.: Quick approximation to matrices and applications. *Combinatorica* **19**(2), 175–220 (1999)
15. Guha, S., Li, Y., Zhang, Q.: Distributed partial clustering. In: ACM SPAA (2017)
16. Gupta, S., Kumar, R., Lu, K., Moseley, B., Vassilvitskii, S.: Local search methods for k-means with outliers. *PVLDB* **10**(7), 757–768 (2017), <http://www.vldb.org/pvldb/vol10/p757-Lu.pdf>
17. Kemeny, J.G.: Mathematics without numbers. *Daedalus* **88**(4), 577–591 (1959)
18. Kenyon-Mathieu, C., Schudy, W.: How to rank with few errors. In: ACM STOC (2007)
19. Lu, T., Boutilier, C.: Learning mallows models with pairwise preferences. In: ICML (2011)

20. Luce, R.D.: Individual Choice Behavior a Theoretical Analysis. John Wiley and sons (1959)
21. Malkomes, G., Kusner, M.J., Chen, W., Weinberger, K.Q., Moseley, B.: Fast distributed k-center clustering with outliers on massive data. In: NIPS (2015)
22. Muller, E., Sánchez, P.I., Mülle, Y., Bohm, K.: Ranking outlier nodes in subspaces of attributed graphs. In: IEEE ICDEW (2013)
23. Negahban, S., Oh, S., Shah, D.: Rank centrality: Ranking from pairwise comparisons. *Operations Research* **65**(1), 266–287 (2016)
24. Rajkumar, A., Agarwal, S.: A statistical convergence perspective of algorithms for rank aggregation from pairwise data. In: ICML (2014)
25. Van Zuylen, A., Williamson, D.P.: Deterministic algorithms for rank aggregation and other ranking and clustering problems. In: WAOA. pp. 260–273. Springer (2007)
26. Wauthier, F., Jordan, M., Jojic, N.: Efficient ranking from pairwise comparisons. In: ICML (2013)
27. Williamson, D.P., Shmoys, D.B.: The Design of Approximation Algorithms. Cambridge University Press (2011), http://www.cambridge.org/de/knowledge/isbn/item5759340/?site_locale=de_DE

Appendix

A Analysis of Our Algorithm When $x^* \leq \sqrt{n}$

A.1 Analysis of Approximation Ratio

We first give the omitted proof of Lemma 1.

Proof. of Lemma 1. It suffices to show that for any $V' \in \mathcal{V}$, conditioned on $\pi_i(V) \in V'$, $\pi_i(V)$ is chosen from V' uniformly at random. Let \mathcal{V} be the current partition of V just before the i th pivot $\pi_i(V)$ is applied. Note that $\Pr[\pi_i(V) \in V'] = |V'|/(|V| - (i - 1))$, and for any $v' \in V'$, $\Pr[\pi_i(V) = v'] = 1/(|V| - (i - 1))$. Therefore, the conditional probability is $1/|V'|$, as desired. \square

For the sake of analysis, consider a fixed solution that is feasible w.r.t. the pair (x^*, y^*) , which we will refer to as the optimal solution. Let U^* denote the outlier set of the optimal solution. We begin our analysis with the following easy observation.

Lemma 3. $\Pr[S \cap U^* = \emptyset] \geq 3/4$.

Proof. Recall that each vertex v , particularly in U^* , is independently added to S with probability $\frac{1}{4x^*}$. Since $|U^*| \leq x^*$, via a union bound, we have $\Pr[S \cap U^* \neq \emptyset] \leq 1/4$. \square

For the sake of analysis, we will compare several stochastic or deterministic quantities that are subtly different from one another. We first summarize the quantities as follows for ease of reference. We will often condition some quantities on the sample S being S' where S' is an ordered subset of V . Let $[k + 1] := \{1, 2, \dots, k + 1\}$.

- LP_{fasto}^{primal} : The optimal objective of LP_{fasto}^{primal} ; for notional convenience, this may also refer to the LP itself.
- LP^{dual} : The optimal objective of LP^{dual} ; for notional convenience, this may also refer to the LP itself. (Note that $LP_{fasto}^{primal}|_{S=S'} = LP^{dual}|_{S=S'}$ for any S' due to the strong LP duality. Note that both quantities are deterministic for any fixed S' .)
- $COST_{fasto}^{pd}$: The objective of LP_{fasto}^{primal} for the solution output by the primal-dual in Step 2. Formally, $\sum_{e \in E'_{pd}} y_e^{pd} + \sum_{x \in U_{pd}} c \cdot p_x^{pd} + \sum_{e \in E_{pd}^2} z_e^{pd} + |S_{back}|$ (Here we add 'pd' to the subscript or superscript to note that this is the primal-dual's output.)
- $COST^{kwik}(G')$: the number of backward edges produced by KWIK-SORT on a tournament G' . (Note that this is a stochastic quantity because of the internal randomness of KWIK-SORT.)
- $COST_{fasto}^A$: Our algorithm's final cost. Recall that each backward edge incurs a unit cost and each outlier incurs $c = y^*/x^*$ units of cost. Thus, formally, $\sum_{i \in [k+1]} COST^{kwik}(G[G_i \setminus U_{pd}]) + \sum_{x \in U_{pd}} c \cdot p_x^{pd} + \sum_{e \in E_{pd}^2} z_e^{pd} + |S_{back}|$.

We observe that the optimal solution has a small objective for LP_{fasto}^{primal} when the sample is disjoint from U^* , meaning that its optimal objective is small.

Lemma 4. $\mathbb{E}[LP_{fasto}^{primal}|_{S \cap U^* = \emptyset}] \leq 4y^*$.

Proof. For the sake of analysis, we set the variables of LP_{fasto}^{primal} mimicking the optimal solution: $p_x = 1$ for all $x \in U^*$ and run KWIK-SORT-SEQ on $V \setminus U^*$. Here, the algorithm is given a random permutation of $V \setminus U^*$, which is generated in the following manner: Add each vertex from $V \setminus U^*$ to S_1 with probability $1/(4x^*)$ independently, and let $S_2 := V \setminus U^* \setminus S_1$; then, append a random permutation of S_2 to a random permutation of S_1 . This also gives an alternative way of obtaining S , conditioned on $S \cap U^* = \emptyset$, by setting $S = S_1$. Knowing that the optimal solution for FAST on $V \setminus U^*$ has at most y^* backward edges, the expected number of backward edges that KWIK-SORT-SEQ produces on $V \setminus U^*$ is at most $3y^*$. We set y, z, S_{back} according to these backward edges. Thus, the expected objective of LP_{fasto}^{primal} is at most $3y^* + cx^* \leq 4y^*$, since $c = y^*/x^*$. \square

Now we show that our primal-dual method yields a constant approximate solution to LP_{fasto}^{primal} against its optimum for any fixed S' . When we say S' is fixed, we mean that the ordering of vertices in S' is also fixed.

Lemma 5. For any S' , $COST_{fasto}^{pd}|_{S=S'} \leq 5 \cdot LP^{dual}|_{S=S'} = 5 \cdot LP_{fasto}^{primal}|_{S=S'}$.

Proof. Let $U_{pd} := U_1^{pd} \cup U_2^{pd}$ where U_1^{pd} and U_2^{pd} are the set of outliers that have been chosen in lines (5) and (6) of the algorithm, respectively. Note that:

1. If $e \in E'_{pd}$, then $\sum_{e \subset t} \alpha_t = 1$.
2. If $e \in E_{pd}^2$, then $\beta_e = 1$.
3. If $x \in U_{pd}$, then either $\sum_{x \in t} \alpha_t = \frac{3}{5}c$ or $\sum_{x \in e} \beta_e = \frac{2}{5}c$.

In the remaining proof, we drop the condition $S = S'$ for notational convenience. Then, we derive:

$$\begin{aligned}
& COST_{fasto}^{pd}(G) \\
&= \sum_{e \in E'_{pd}} y_e^{pd} + \sum_{x \in U_1^{pd}} c \cdot p_x^{pd} + \sum_{x \in U_2^{pd}} c \cdot p_x^{pd} + \sum_{e \in E_{pd}^2} z_e^{pd} + |S_{back}| \\
&= \sum_{e \in E'_{pd}} \left(\sum_{e \subset t} \alpha_t \right) \cdot y_e^{pd} + \sum_{x \in U_1^{pd}} \frac{5}{3} \cdot \left(\sum_{x \in t} \alpha_t \right) \cdot p_x^{pd} \\
&\quad + \sum_{x \in U_2^{pd}} \frac{5}{2} \cdot \left(\sum_{x \in e} \beta_e \right) \cdot p_x^{pd} + \sum_{e \in E_{pd}^2} \beta_e \cdot z_e^{pd} + |S_{back}| \\
&= \sum_{t \in T_{in}} \alpha_t \left(\sum_{e \in E'_{pd}: e \subset t} y_e^{pd} + \frac{5}{3} \cdot \sum_{x \in U_1^{pd}: x \in t} p_x^{pd} \right) \\
&\quad + \sum_{e \in E_{pd}^2} \beta_e \left(\frac{5}{2} \cdot \sum_{x \in U_2^{pd}: x \in e} p_x^{pd} + z_e^{pd} \right) + |S_{back}| \\
&\leq 5 \cdot \sum_{t \in T_{in}} \alpha_t + 5 \cdot \sum_{e \in E_{pd}^2} \beta_e + |S_{back}| \leq 5 \cdot LP^{dual}
\end{aligned}$$

The last inequality is due to α, β being a feasible solution to LP^{dual} . The second to last inequality follows from the fact that $\sum_{x \in t} p_x^{pd} \leq 3$ and $\sum_{x \in e} p_x^{pd} \leq 2$, and

1. $\sum_{e \in E'_{pd}: e \subset t} y_e^{pd} + \sum_{x \in U_1^{pd}: x \in t} p_x^{pd} \leq 3$; and
2. $\sum_{x \in U_2^{pd}: x \in e} p_x^{pd} + z_e^{pd} \leq 2$,

which hold true because of line (8) of the algorithm. This completes the proof, as we have $LP^{dual}|_{S=S'} = LP^{primal}_{fasto}|_{S=S'}$ due to the LP duality. \square

To complete our analysis, we need to upper bound the number of backward edges output by KWIK-SORT in the last Step 3.

Lemma 6. *For any S' ,*

$$\sum_{i \in [k+1]} \mathbb{E} \left[COST^{kwik}(G[G_i \setminus U]) \mid_{S=S'} \right] \leq 3 \sum_{e \in E'_{pd}} y_e^{pd} \mid_{S=S'}$$

Proof. The work of Ailon et. al. [3] establishes that KWIK-SORT for FAST is 3-approximate against the minimum fractional number of edges to cover all bad triangles. To formally state this, we need to define an LP for FAST. Let T' be the set of bad triangles in the input graph $G' = (V', E')$.

$$\min \sum_{e \in E'} y_e \quad s.t. \quad \sum_{e \subset t} y_e \geq 1 \quad \forall t \in T'; \quad \text{and } y_e \geq 0 \quad \forall e \in E'$$

It is shown in [3] that $\mathbb{E}[COST^{kwik}(G')]$ is at most 3 times the optimal objective of the above LP. Since the LP objective is to be minimized, it also holds

that $\mathbb{E}[COST^{kwik}(G')]$ is at most 3 times the objective of any feasible solution to the LP.

We now apply this fact to our setting. Namely, we observe that edges E'_{pd} cover all unbroken bad triangles within groups. To see this, consider any unbroken bad triangle $t \in T_{in}$. Since t is unbroken, it must be the case that $\sum_{x \in t} p_x^{pd} = 0$, so we have $\sum_{e \in t} y_e^{pd} \geq 1$. Since our primal-dual solution yields an integer solution to the primal LP, it means that $y_e^{pd} = 1$ for some $e \in t$. Because of the way y_e^{pd} is set, we have $\sum_{e \in t \cap E'_{pd}} y_e^{pd} \geq 1$. Therefore, the number of backward edges created by running KWIK-SORT within groups is at most $3 \sum_{e \in E'_{pd}} y_e^{pd}$.

The whole argument holds for any fixed ordered sample S' . \square

We are now ready to complete our analysis.

Theorem 2. *Our algorithm for FASTO outputs a $(O(1), O(1))$ -approximate solution with probability at least $1/2$.*

Proof. For any fixed S' , we have:

$$\begin{aligned}
& \mathbb{E}[COST_{fasto}^A |_{S=S'}] \\
&= \sum_{i \in [k+1]} \mathbb{E}[COST^{kwik}(G[G_i \setminus U]) |_{S=S'}] \\
&\quad + \sum_{x \in U_{pd}} c \cdot p_x^{pd} |_{S=S'} + \sum_{e \in E_{pd}^2} z_e^{pd} |_{S=S'} + |S_{back}| |_{S=S'} \\
&\leq 3 \sum_{e \in E'_{pd}} y_e^{pd} |_{S=S'} + \sum_{x \in U_{pd}} c \cdot p_x^{pd} |_{S=S'} + \sum_{e \in E_{pd}^2} z_e^{pd} |_{S=S'} \\
&\quad + |S_{back}| |_{S=S'} \quad [\text{Lemma 6}] \\
&\leq 3COST_{fasto}^{pd}(G) |_{S=S'} \leq 15LP_{fasto}^{primal}(G) |_{S=S'} \quad [\text{Lemma 5}]
\end{aligned}$$

This implies that $\mathbb{E}[COST_{fasto}^A |_{S \cap U^* = \emptyset}] \leq 60y^*$ due to Lemma 4. By Markov's inequality, we have $\Pr[COST_{fasto}^A \leq 180y^* |_{S \cap U^* = \emptyset}] \geq 2/3$. By deconditioning with Lemma 3, we have $\Pr[COST_{fasto}^A \leq 180y^*] \geq 1/2$. The theorem immediately follows since $COST_{fasto}^A$ is the number of backward edges, plus the number of outliers multiplied by (y^*/x^*) . \square

We note that the success probability claimed in Theorem 1 is slightly lower than $1/2$ to factor in the failure probability of $1/n$ that our algorithm is not memory efficient or slow, which will be shown in the subsequent section.

A.2 Analysis of Memory Usage and Running Time

Recall that in the beginning of Section 2.3 we analyzed our algorithm's memory usage and running time assuming that each group size is $\tilde{O}(n/k)$. This section is devoted to proving that the assumption holds true, which is formally stated as follows.

Lemma 7. *Consider the partition of the vertex set that results from running KWIK-SORT-SEQ on a sample S (see Step 1). If $|S| = k \geq 96 \log n$, then no group has a size of more than $\frac{n \log n}{k} = \tilde{O}(n/k)$, with probability at least $1 - 1/n$.*

The central idea in the proof of Lemma 7 is that a pivot chosen uniformly at random partitions the vertex set into two groups of similar sizes with a constant probability. To formally articulate the idea, we need the following definition.

Definition 1. *For every vertex $v \in V$, let $in(v)$ and $out(v)$ be the in-degree and out-degree of v , respectively. We define the set of balanced vertices, $B(V)$, high out-degree vertices, $HO(V)$, and high in-degree vertices, $HI(V)$ as follows: $B(V) = \{v \in V \mid \frac{1}{7} \leq \frac{in(v)}{out(v)} \leq 7\}$; $HO(V) = \{v \in V \mid 7 < \frac{out(v)}{in(v)}\}$; and $HI(V) = \{v \in V \mid 7 < \frac{in(v)}{out(v)}\}$.*

Lemma 8. $|B(V)| \geq \frac{n}{4} - 1$ for any tournament $G = (V, E)$.

Proof. For the sake of contradiction, suppose for some tournament $G = (V, E)$ we have $|B(V)| < \frac{n}{4} - 1$. We assume w.l.o.g. that $|HO(V)| \geq |HI(V)|$. Consider an arbitrary vertex $v \in HO(V)$. From the definition of high out-degree vertices, we have that $out(v) \geq \frac{7}{8} \cdot (n - 1)$. Let $out^{HO}(v)$ be the number of out-neighbors of v in $HO(V)$. Then, we can lower bound $out^{HO}(v) \geq out(v) - |B(V)| - |HI(V)| \geq \frac{7}{8} \cdot (n - 1) - (\frac{n}{4} - 1) - \frac{n}{8} \geq \frac{n}{8}$. What this means is that each $v \in HO(V)$ contributes to $\sum_{u \in HO(V)} in(u)$ by at least $n/8$, as there are at least $n/8$ vertices in $HO(V)$ connecting from v . By summing up over all $v \in HO(V)$, we have $\sum_{u \in HO(V)} in(u) \geq \frac{n}{8}|HO(V)|$. On the other hand, from the definition of the high-degree vertices, we know that $in(v) \leq \frac{n-1}{8}$ for every vertex $v \in HO(V)$, which implies $\sum_{u \in HO(V)} in(u) \leq \frac{n-1}{8} \cdot |HO(V)|$, which is a contradiction. \square

Let P_1, P_2, \dots, P_k denote the pivots in the order they are given to the algorithm. We will say that the algorithm applies P_j in the j th iteration. Let $G(i, j)$ be the sub-tournament/group including vertex i after picking j pivots; if i was chosen as a pivot in the j th iteration or before, then $G(i, j)$ has only one vertex, i . For notational convenience, we also let $G(i, j)$ denote the vertex set of the tournament. So, $|G(i, j)|$ denotes the number of vertices in the group including i after the j th iteration.

In the following, we show that the expected size of the group where a vertex i belongs decreases by a constant factor if a pivot is sampled from the group.

Lemma 9. $\mathbb{E}[|G(i, k + 1)| \mid P_{k+1} \in G(i, k)] \leq \frac{31}{32}|G(i, k)|$.

Proof. Consider the group $G(i, k)$. The group size never increases. If a pivot is sampled from $G(i, k)$ in the $k + 1$ th iteration, i.e., $P_{k+1} \in G(i, k)$, we know that $P_{k+1} \in B(G(i, k))$ with probability at least $1/4$ by Lemma 8. (Here, for better readability, we ignore -1 in the lemma statement, which has a negligible effect when n is sufficiently large.) When it happens, vertex i 's group size decreases by a factor of at least $1/8$. Thus, we derive, $\mathbb{E}[|G(i, k + 1)| \mid P_{k+1} \in G(i, k)] \leq \frac{1}{4} \cdot \frac{7}{8}|G(i, k)| + \frac{3}{4}|G(i, k)| \leq \frac{31}{32}|G(i, k)|$. \square

Lemma 10. For any $m > 0$ and $k \geq 0$,

$$\Pr[|G(i, k)| > m] \leq \frac{n}{m} \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right)^k.$$

Proof. If $|G(i, k-1)| \geq m$, then, from Lemma 9, we have,

$$\begin{aligned} & \mathbb{E}[|G(i, k)|] \\ & \leq \frac{|G(i, k-1)|}{n-k+1} \cdot \frac{31}{32} |G(i, k-1)| + \left(1 - \frac{|G(i, k-1)|}{n-k+1}\right) |G(i, k-1)| \\ & = |G(i, k-1)| \left(1 - \frac{1}{32} \cdot \frac{|G(i, k-1)|}{n-k+1}\right) \leq |G(i, k-1)| \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right). \end{aligned}$$

Note that $|G(i, k)| > m$ only if $|G(i, k-1)|, |G(i, k-2)|, \dots, |G(i, 0)| > m$. Thus, by a simple induction, we can show that

$$\begin{aligned} & \mathbb{E}\left[|G(i, k)| \mid |G(i, k-1)|, |G(i, k-2)|, \dots, |G(i, 0)| > m\right] \\ & \leq \mathbb{E}[|G(i, 0)|] \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right)^k = n \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right)^k \end{aligned}$$

By the Markov inequality, we have

$$\begin{aligned} \Pr\left[|G(i, k)| > m \mid |G(i, k-1)|, |G(i, k-2)|, \dots, |G(i, 0)| > m\right] \\ \leq \frac{n}{m} \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right)^k \end{aligned}$$

De-conditioning this yields the lemma. \square

We are now ready to complete the proof of Lemma 7. We apply Lemma 10 with $m = \frac{96n \log n}{k}$. Then, for any vertex v ,

$$\Pr[|G(i, k)| \geq m] \leq \frac{n}{m} \left(1 - \frac{1}{32} \cdot \frac{m}{n}\right)^k \leq n \left(1 - \frac{3 \log n}{k}\right)^k = 1/n^2.$$

In other words, a vertex v 's group size is greater than m with a very small probability, $1/n^2$. By applying a union bound over all vertices, we conclude that no group has a size greater than m , with probability at least $1 - 1/n$, as claimed.

B Our Algorithm When the Target Number of Outliers x^* is Large

This section is devoted to proving Theorem 1 for the case $x^* > \sqrt{n}$.

Assuming that E' stated in Lemma 2 exists, we can find V'' s.t. $|V''| \leq O(\log n)x^*$ and $|E'[V \setminus V'']| \leq O(1)y^*$. The algorithm is essentially a special case of the primal-dual method for the case when x^* is small. Assuming that Lemma 2 holds true, to prove Theorem 1 for the case $x^* > \sqrt{n}$, it suffices to solve the following problem: We are given as input an *undirected* graph $G = (V, E)$; G is

not necessarily complete. Our goal is to find $U \subseteq V$ such that $|U| \leq O(1)x^*$ and $|E[V \setminus U]| \leq O(1)y^*$ under the assumption that U^* exists such that $|U^*| \leq x^*$ and $|E[V \setminus U^*]| \leq y^*$. (Note that each undirected edge in E has a uniquely corresponding backward edge in E' and we overloaded x^* and y^* to simplify the notation.)

Consider the following LP for this problem:

$$\min \sum_{x \in V} p_x \cdot c + \sum_{e \in E} z_e \quad \text{s.t.} \quad p_u + p_v + z_e \geq 1 \quad \forall e = (u, v) \in E,$$

which is over $p_x \geq 0$ over all $x \in V$ and $z_e \geq 0$ for all $e \in E$. Here $c = y^*/x^*$ as before. Note that the optimal objective is at most $2y^*$.

Consider the dual:

$$\max \sum_{e \in E} \beta_e \quad \text{s.t.} \quad \sum_{e \in E: x \in e} \beta_e \leq c \quad \forall x \in V \quad (10)$$

$$\beta_e \leq 1 \quad \forall e \in E, \quad (11)$$

which is over $\beta_e \geq 0$ for all $e \in E$.

The primal-dual algorithm is very simple: Initially, $U = \emptyset$, and all $\beta_e = 0$ are active. We uniformly increase all active β_e . If $\sum_{e \in E: x \in e} \beta_e = c$, then we inactivate all β_e for all edges incident on x , and add x to U . If $\beta_e = 1$ for any e , we stop. We pay a unit cost for each edge e with $\beta_e = 1$.

So, our solution's LP primal objective is

$$\begin{aligned} & \sum_{x \in U} c + \sum_{e \in E: \beta_e = 1} 1 = \sum_{x \in U} \sum_{e \in E: x \in e} \beta_e + \sum_{e \in E: \beta_e = 1} \beta_e \\ & = \sum_{e \in E} \sum_{x \in U: x \in e} \beta_e + \sum_{e \in E: \beta_e = 1} \beta_e \leq \sum_{e \in E} 2\beta_e + \sum_{e \in E} \beta_e = 3 \sum_{e \in E} \beta_e \leq 6y^* \end{aligned}$$

Thus, we have $c|U| + |E[V \setminus U]| \leq 6y^*$. Therefore, we have $|U| \leq 6x^*$ and $|E[V \setminus U]| \leq 6y^*$, as desired. We note that one can easily improve these upper bounds to $4x^*$ and $4y^*$, but we do not optimize constants here.

The remainder of this section is devoted to proving Lemma 2.

Proof. (of Lemma 2) Let $\pi := \pi(V)$ be the random permutation given to KWIK-SORT-SEQ as pivot one by one. Fix this permutation π . If the pivot v fed into KWIK-SORT-SEQ is from U^* , then, for the sake of analysis, we delete all vertices in the group including v . Let $V^1(\pi)$ denote the set of remaining vertices and $E^1(\pi)$ the set of backward edges over $V^1(\pi)$ in the final ordering. Now consider another execution of KWIK-SORT-SEQ on $V \setminus U^*$ that ignores the optimal solution's outliers. That is, if KWIK-SORT-SEQ receives v from U^* as pivot, it ignores v . Let $E^2(\pi)$ denote the set of backward edges in the resulting final ordering of $V \setminus U^*$. Let us call these two executions the first and second, respectively.

We will show that

$$|E^1(\pi)| \leq |E^2(\pi)| \quad \text{for any } \pi; \quad \text{and} \quad (12)$$

$$\mathbb{E}_\pi |V \setminus V^1(\pi)| \leq (65 \log n)x^* \quad (13)$$

Then, we have $\mathbb{E}_\pi |E^1(\pi)| \leq \mathbb{E}_\pi |E^2(\pi)| \leq 3y^*$. By applying Markov inequality for both and a union bound, we have that $\Pr_\pi[|E^1(\pi)| \geq 12y^* \text{ or } |V \setminus V^1(\pi)| \geq (260 \log n)x^*] \leq 1/2$. By setting $V' = V \setminus V^1(\pi)$ and $E' = E^1(\pi)$, we have the lemma..

It now remains to prove the two claims (12) and (13). To see the easier claim (12), we just gradually change the second execution so that the two executions produce an identical output: when the second execution receives an optimal solution's outlier as pivot for the first time, it deletes all vertices in the group including the pivot, which only decreases the number of backward edges. We obtain the same ordering of $V \setminus U^*$ by repeating this for all vertices in U^* .

To prove the second claim (13), we take an equivalent parallel view of the first execution. That is, for the sake of analysis, we assume that in each iteration, every (non-singleton) group is partitioned into two or three. By Lemma 9, we know that each group's expected size decreases by a factor of at least $1/32$ in each iteration. It is an easy exercise to show that every group becomes a singleton group after $64 \log n$ iterations, with probability at least $1 - 1/n$. Let $\{G_i^\ell\}_i$ denote the collection of (disjoint) groups right after the ℓ th iteration. As before, we delete all vertices in G_i^ℓ if a pivot chosen from U^* partitions G_i^ℓ in the $\ell + 1$ th iteration. Then, the expected number of deleted vertices for the first $64 \log n$ iterations is at most, $\sum_{\ell=1}^{64 \log n} \sum_i \frac{|G_i^\ell \cap U^*|}{|G_i^\ell|} \cdot |G_i^\ell| \leq (64 \log n)x^*$. After that, a non-singleton group exists with probability at most $1/n$; thus, the expected number of vertices deleted afterwards is at most $n * (1/n) = 1$. Further, after all groups become singleton, at most x^* vertices are deleted. Thus, we have an upper bound of $64(\log n)x^* + 1 + x^* \leq 65(\log n)x^*$. \square