# How to Rank with Few Errors 

# A PTAS for Weighted Feedback Arc Set on Tournaments* 

Claire Kenyon-Mathieu

Warren Schudy ${ }^{\dagger}$

Brown University<br>Computer Science<br>Providence, RI


#### Abstract

We present a polynomial time approximation scheme (PTAS) for the minimum feedback arc set problem on tournaments. A simple weighted generalization gives a PTAS for KemenyYoung rank aggregation.


## Categories and Subject Descriptors

F.2.2 [Analysis of Algorithms and Problem Complexity]: Nonnumerical Algorithms and Problems-Computations on discrete structures

## General Terms

Algorithms, Theory

## Keywords

Approximation algorithm, Feedback arc set, Kemeny-Young rank aggregation, Max acyclic subgraph, Polynomial-time approximation scheme, Tournament graphs

## 1. INTRODUCTION

Suppose you ran a chess tournament, everybody played everybody, and you wanted to use the results to rank everybody. Unless you were really lucky, the results would not be acyclic, so you could not just sort the players by who beat whom. A natural objective is to find a ranking that minimizes the number of upsets, where an upset is a pair of players where the player ranked lower on the ranking beats the player ranked higher. This is the minimum feedback arc set (FAS) problem on tournaments. Our main result is a polynomial time approximation scheme (PTAS) for this problem. A simple weighted generalization gives a PTAS for Kemeny-Young rank aggregation.

[^0][^1]rankings of objects of some sort, such as search engine outputs $[16,17]$, and desires to aggregate the input rankings into a single output ranking that is similar to all of the input rankings: it should have minimum average distance from the input rankings, for some notion of distance. This ancient problem was already studied in the context of voting by Borda [6] and Condorcet [10] in the 18th century, and has aroused renewed interest in learning [9]. A natural notion of distance is the number of pairs of vertices that are in different orders: this defines the Kemeny-Young rank aggregation problem [29, 30]. This problem is NP-hard, even with only 4 voters [17]. Constant factor approximation algorithms are known: choosing a random (or best) input ranking as the output gives a 2 -approximation; finding the optimal aggregate ranking using the footrule distance as the metric instead of the Kendall-Tau distance also gives a 2 approximation [16]. There is also a randomized $4 / 3$ approximation algorithm $[1,2]$. We improve on these results by designing a polynomial-time approximation scheme.

Theorem 1 (PTAS). There are randomized algorithms for minimum Feedback Arc Set on weighted tournaments and for Kemeny-Young rank aggregation that, given $\lambda>0$, output orderings with expected cost less than $(1+\lambda) O P T$. Let $\epsilon=\lambda b$. The running time is:

$$
O\left((1 / \epsilon) n^{6}+n^{2} 2^{\tilde{O}(1 / \epsilon)}+n 2^{2^{\tilde{O}(1 / \epsilon)}}\right)
$$

The algorithms can be derandomized at the cost of increasing the running time to $n^{2^{\tilde{O}(1 / \epsilon)}}$.

Our algorithm combines several existing algorithmic tools. It uses (a) existing constant factor approximation algorithms. In addition, it uses (b) existing polynomial-time approximation schemes for the complementary maximization problem (a.k.a. max acyclic subgraph tournaments), due to Arora, Frieze and Kaplan [4] and to Frieze and Kannan [22]. Moreover, our central algorithm, Algorithm 2, also uses (c) the divide-and-conquer paradigm. Of course, this has been previously used in this setting, in partitioning algorithms in general graphs [33] and in a Quicksort-type algorithm [1], but both of these constructions are top-down, whereas our algorithm is bottom-up, which gives it a very different flavor. Finally, it also relies on (d) a simple but useful technique, used already in 1961 by Slater and Alway [39]: iteratively improve the current ordering by taking a vertex out of the ordering and moving it back at a different position. (Such single vertex moves and variants have been studied at some length since that time $[25,44,16,17]$.)

It is fortunate that the FAS problem on tournaments has an approximation scheme. The closely related problem of correlation clustering on complete graphs ${ }^{1}$ is APX-hard [8]. The related problem of feedback vertex set is also hard to approximate even on tournaments, though it does have a 2.5-approximation algorithm [7].

We note that Amit Agarwal has informed us that he recently obtained similar results.

[^2]
## 2. ALGORITHM

In this section, we present our deterministic and randomized algorithms and reduce one to the other.

Our algorithms use the sampling-based additive approximation algorithms of [4, 22] as a subroutine. For simplicity, here we use the derandomized additive error algorithm. Section 4.1 discusses how to extend our analysis to use the randomized version, yielding the runtime in Theorem 1.

Theorem 2 (Small additive error). [4, 22] There is a randomized polynomial-time approximation scheme for maximum acyclic subgraph on weighted tournaments. Given $\beta, \eta>0$, the algorithm outputs, in polynomial time, [22]

$$
O\left(\left[n / \beta^{4}+2^{O\left(1 / \beta^{2}\right)}\right] \log (1 / \eta)\right)
$$

an ordering whose cost is, with probability at least $1-\eta$, less than $O P T+\beta n^{2}$. The algorithm can be derandomized into an algorithm which we denote by AddApprox, which increases the runtime to $n^{O\left(1 / \beta^{4}\right)}$.

We also use a simple type of local search. A total order $R$ can be specified by an ordering $\pi: V \rightarrow\{1, \ldots, n\}$ such that $\pi(x)$ is the position of vertex $x: \pi(x)<\pi(y)$ iff $R(x, y)$. A single vertex move, given a ordering $\pi$, a vertex $x$ and a position $i$, consists of taking $x$ out of $\pi$ and putting it back in position $i$.

All of our algorithms call the additive approximation algorithm with parameter

$$
\beta=\beta(\epsilon) \equiv 9^{-\frac{1}{\epsilon} \log _{3 / 2}\left(1 / \epsilon^{2}\right)} \epsilon^{3}=2^{-O(1 / \epsilon \log (1 / \epsilon))}
$$

Our deterministic PTAS is given in Algorithm 1. Recall from Problem 1 that $b$ is the lower bound on $w_{x y}+w_{y x}$ for every pair $\{x, y\}$. For ease of thinking, the reader may think of $b$ as being equal to 1 . Also recall that we are seeking a $1+\lambda=1+\epsilon / b$ approximation.

Our (somewhat faster) randomized PTAS is given in Algorithm 2. Curiously, if the constant factor approximation algorithm used is the sorting by indegree algorithms from [6, 12], the initial order $\pi^{l o c a l}$ computed in Algorithm 2 is precisely the order returned by the heuristic algorithm created by Slater and Alway [39] in 1961!

For the analysis, it is enough to study Algorithm 2, since the result for Algorithm 1 follows as a simple corollary:

```
Algorithm 1 Deterministic polynomial time approximation
scheme of Theorem 1 (PTAS)
Given: Fixed parameters \(\epsilon>0\) and \(b \in(0,1]\).
Input: A weighted tournament.
    Round weights to integer multiples of \(\epsilon / n^{2}\).
    \(\pi \leftarrow\) constant factor approximation [1, 12].
    While some move decreases cost, do it. Types of moves:
```

1. Single vertex moves. Choose vertex $x$ and rank $j$, take $x$ out of the ordering and put it back in so that its rank is $j$.
2. Additive approximation. Choose integers $i<$ $j$; let $U=\{$ vertices with rank in $[i, j]\}$. $\pi_{U}^{\prime} \leftarrow \operatorname{AddApprox}(U)$ with parameter $\beta(\epsilon)$. Replace the restriction $\pi_{U}$ of $\pi$ to $U$ by $\pi_{U}^{\prime}$.

Output: $\pi$.

```
Algorithm 2 Randomized polynomial time approximation
scheme (RPTAS) of Theorem 1. Vertical lines in left margin
denote differences from Algorithm 3.
Given: Fixed parameters \(\epsilon>0\) and \(b \in(0,1]\).
Input: A weighted tournament
    Round weights to integer multiples of \(\epsilon / n^{2}\).
    \(\pi \leftarrow\) constant factor approximation \([1,12]\).
    Apply single vertex moves to get a local optimum \(\pi^{l o c a l}\).
Output: \(\operatorname{Rec}(1, n)\)
\(\boldsymbol{\operatorname { R e c }}(i, j)=\)
    If \(\mathrm{i}=\mathrm{j}\) then Return i
    For any \(\ell, m\), let \(V_{\ell, m}=\left\{\right.\) vertices \(\left.x: \pi^{\text {local }}(x) \in[\ell, m]\right\}\)
    \(K \leftarrow\left\{k:\left|V_{i k}\right|,\left|V_{k+1, j}\right| \geq\left|V_{i j}\right| / 3\right\}\).
    Choose \(k\) uniformly at random from \(K\).
    \(\rho_{1} \leftarrow \operatorname{AddApprox}\left(V_{i j}\right)\) with parameter \(\beta(\epsilon)\)
    \(\rho_{2} \leftarrow\) Concatenate \(\boldsymbol{\operatorname { R e c }}(i, k)\) and \(\boldsymbol{\operatorname { R e c }}(k+1, j)\)
    Return \(\rho_{1}\) or \(\rho_{2}\), whichever has lower cost.
```

Lemma 3. If Algorithm 2 outputs an ordering whose cost is $\leq(1+\lambda)$ OPT with positive probability, then Algorithm 1 outputs an ordering with cost $\leq(1+\lambda) O P T$.

Proof. Let $\pi$ denote the output of Algorithm 1. Since we started with a constant factor approximation and only performed cost-improving moves, $\pi$ is still a constant factor approximation, of course. Execute Algorithm 2 starting with this ordering $\pi$. Since there are no cost-decreasing single-vertex moves, $\pi^{l o c a l}=\pi$. Since Algorithm AddApprox cannot improve any subinterval $U$ of $\pi$, the recursive process will always choose $\rho_{2}$ rather than $\rho_{1}$ : in the end, we obtain $\pi^{\text {out }}=\pi$ for every execution of our Algorithm 2. Therefore it must be that $\pi$ has cost $\leq(1+\lambda) \mathrm{OPT}$.

## 3. ANALYSIS: CORRECTNESS

### 3.1 Rounding analysis

The following lemma shows that rounding the weights, which is done to ensure polynomial-time termination of local search, is irrelevant for the correctness analysis.

Lemma 4 (Rounding). Let $C$ denote the cost function with the original weights, $\tilde{C}$ denote the cost function with the rounded weights, and $O \tilde{P} T=\min _{\rho} \tilde{C}(\rho)$ denote the value of the optimal solution to the problem with rounded weights. If $\tilde{C}(\pi) \leq(1+\lambda) O \tilde{P} T$, then $C(\pi) \leq(1+3 \lambda) O P T$.

Proof Sketch. The proof is split into two cases, depending on whether OPT $<b(1 / 2-\lambda)$.

Thanks to Lemma 4, without loss of generality we can assume that the input weights are integer multiples of $\epsilon / n^{2}$; we will make this assumption for the remainder of the paper.

### 3.2 Reduction to a virtual algorithm

Algorithm 2 is quite difficult to analyze directly, because it uses bottom-up recursion, and so we will instead analyze a different, top-down process, Algorithm 3. This is a virtual algorithm whose sole purpose is to guide the analysis: indeed, it takes as input parameter an optimal ordering $\pi^{*}$, so it cannot be a real algorithm!

For any subset $S$ of vertices, the Spearman's footrule distance between the two orderings used in the specification of

```
Algorithm 3 Virtual algorithm. Vertical lines in left mar-
gin denote differences from Algorithm 2.
Given: Fixed parameters \(\epsilon>0\) and \(b \in(0,1]\).
Input: Weighted tournament; optimal ordering \(\pi^{*}\)
    Round weights to integer multiples of \(\epsilon / n^{2}\).
    \(\pi \leftarrow\) constant factor approximation \([1,12]\).
    Apply single vertex moves to get a local optimum \(\pi^{\text {local }}\).
    Let \(\alpha=9^{-r}\), with \(r\) uniform in \(\left[0, \log _{9}\left(\epsilon^{3} / \beta\right)\right]\).
Output: \(\pi^{o u t}=\boldsymbol{\operatorname { R e c }}(1, \mathrm{n})\)
\(\operatorname{Rec}(\mathrm{i}, \mathrm{j})=\)
    If \(\mathrm{i}=\mathrm{j}\) then Return i
    For any \(\ell, m\), let \(V_{\ell, m}=\left\{\right.\) vertices \(\left.x: \pi^{\text {local }}(x) \in[\ell, m]\right\}\)
    \(K \leftarrow\left\{k:\left|V_{i k}\right|,\left|V_{k+1, j}\right| \geq\left|V_{i j}\right| / 3\right\}\).
    Choose \(k\) uniformly at random from \(K\).
    \(\rho_{1} \leftarrow \operatorname{AddApprox}\left(V_{i j}\right)\) with parameter \(\beta\)
    \(\rho_{2} \leftarrow\) Concatenate \(\boldsymbol{\operatorname { R e c }}(i, k)\) and \(\boldsymbol{\operatorname { R e c }}(k+1, j)\)
    If \(F_{V_{i, j}} \geq \alpha \epsilon^{2}\left|V_{i, j}\right|^{2}\)
        Return \(\rho_{1}\) ( \(V_{i j}\) is called a leaf)
    else
        Return \(\rho_{2}\) ( \(V_{i j}\) is called an internal node)
```

the stopping condition used in Algorithm 3 is defined by

$$
F_{S}=\sum_{x \in S}\left|\pi^{l o c a l}(x)-\pi^{*}(x)\right| .
$$

In tournament graphs, the footrule distance is related to the cost of the two orderings:

Lemma 5. $F_{V} \leq(2 / b)\left(C_{V}\left(\pi^{*}\right)+C_{V}\left(\pi^{\text {local }}\right)\right)$.
Proof. By [14]'s Theorem 2 relating Spearman's footrule and Kendall-Tau, we have:

$$
\begin{aligned}
& \sum_{j}\left|\pi^{l o c a l}(j)-\pi^{*}(j)\right| \\
& \quad \leq 2 \sum_{i, j} \mathbb{1}\left(\pi^{*}(i)>\pi^{*}(j)\right) \mathbb{1}\left(\pi^{\text {local }}(i)<\pi^{\text {local }}(j)\right) \\
& \leq \frac{2}{b} \sum_{i, j}\left[\mathbb{1}\left(\pi^{*}(i)>\pi^{*}(j)\right) w_{i j}+\mathbb{1}\left(\pi^{\text {local }}(i)<\pi^{\text {local }}(j)\right) w_{j i}\right] \\
& =(2 / b)\left(C\left(\pi^{*}\right)+C\left(\pi^{\text {local }}\right)\right) .
\end{aligned}
$$

where

$$
\mathbb{1}(p) \equiv\left\{\begin{array}{ll}
1 & \text { if } p \text { is true } \\
0 & \text { otherwise }
\end{array} .\right.
$$

The second inequality follows from $w_{i j}+w_{j i} \geq b$.
The following Lemma shows that it is sufficient to analyze Algorithm 3.

Lemma 6. If $\pi^{\text {out }}{ }_{2}$ denotes the output of the Algorithm 2 and $\pi^{\text {out }}$ denotes the output of the Algorithm 3, then:

$$
\forall x, \quad \operatorname{Pr}\left(\operatorname{cost}\left(\pi^{\text {out }}{ }_{2}\right) \leq x\right) \geq \operatorname{Pr}\left(\operatorname{cost}\left(\pi^{\text {out }}\right) \leq x\right)
$$

Proof. Couple the executions of the two algorithms so that they start with the same $\pi$, do the same sequence of single-vertex moves leading to the same $\pi^{\text {local }}$, and make the same choice of $k$ in their recursive divide-and-conquer process when they have the same value of $(i, j)$. The two divide-and-conquer processes have the same tree of recursive calls. By bottom-up induction on the tree nodes, the
ordering returned by the process used in Algorithm 2 has cost less than or equal to that of the ordering returned by the process used in Algorithm 3.

We now focus on the output $\pi^{\text {out }}$ of Algorithm 3, whose analysis hinges on the following main technical result, proven in Section 3.5. For any subset of vertices $S$, let

$$
C_{S}(\pi)=\sum_{\{x, y\} \subseteq S, \pi(x)>\pi(y)} w_{x y} .
$$

The objective is to minimize the overall cost $C_{V}(\pi)$, which we also denote as $C(\pi)$ (no subscript) for shorthand.

Theorem 7. For the orderings defined in Algorithm 3:

$$
\mathbf{E}\left[C\left(\pi^{o u t}\right)\right]-C\left(\pi^{*}\right) \leq O(\epsilon / b)\left(C\left(\pi^{l o c a l}\right)+C\left(\pi^{*}\right)\right)
$$

Since $\pi^{\text {local }}$ is an improvement over the constant factor approximation $\pi$, we still have $C\left(\pi^{\text {local }}\right) \leq C(\pi)=O\left(C\left(\pi^{*}\right)\right)$. Thus $\mathbf{E}\left[C\left(\pi^{o u t}\right)\right] \leq C\left(\pi^{*}\right)(1+O(\epsilon / b))$, hence Algorithm 3 is an approximation scheme. By Lemma 6, so is Algorithm 2, and by Lemma 3, so is Algorithm 1.

The running time analysis is deferred to Section 4.2, so all we need to do is analyze the virtual algorithm.

### 3.3 Intuition

We are now ready for some intuition for Algorithm 3. When the optimal cost $C\left(\pi^{*}\right)$ is large, the additive approximation algorithm is satisfactory. When $C\left(\pi^{*}\right)$ is small, $C\left(\pi^{\text {local }}\right)$ is also small (because it is a constant factor approximation). This implies that the footrule distance will also be small (Lemma 5), the stopping condition will not hold, and Algorithm 3 will recurse. It will perform some partition $V=V_{1, k} \cup V_{k+1, n}$, thereby irrevocably committing all vertices of $V_{1, k}$ to precede all vertices of $V_{k+1, n}$ in the output ordering. We need to show that the mistakes made in such dividing steps are not too bad.

However, since the footrule distance is small in that case, by definition of footrule, the positions of the vertices in $\pi^{\text {local }}$ are typically close to their positions in $\pi^{*}$. When $k$ is chosen at random, the vertices with rank $\{1, \ldots, k\}$ in $\pi^{l o c a l}$ will be more or less the same as the vertices with rank $\{1, \ldots, k\}$ in $\pi^{*}$. Thus, we expect that only a few vertices will be misplaced during the divide step.

How costly can it be, if just one or two vertices are misplaced during the divide step? Unfortunately, it can be quite costly. For example, if a misplaced vertex $x$, which has rank $\ell>k$ in $\pi^{\text {local }}$, has rank 1 in $\pi^{*}$, then misplacing just a single vertex $x$ may conceivably incur a cost of $\Theta(n)$. This is where the single-vertex moves come to the rescue. By optimality, $\pi^{*}$ cannot improve its cost by moving $x$ to position $\ell$, and by local optimality, $\pi^{l o c a l}$ cannot improve its cost by moving $x$ to position 1. It must be that about half of the edges between $x$ and $V_{1, k}$ are directed towards $x$ and about half are directed away from $x$, and so, it does not really matter whether we place $x$ in position 1 or in position $\ell$. This explains why the algorithm makes sense intuitively.

### 3.4 Reduction to analyzing $\pi^{\text {local }}$

We first reduce the problem to analyzing $\pi^{\text {local }}$ instead of $\pi^{\text {out }}$. Since the two orderings only differ inside the leaves of the recursion tree, we have:

Lemma 8. For every sequence of random choices of the algorithm, we have

$$
C_{V}\left(\pi^{o u t}\right)-C_{V}\left(\pi^{l o c a l}\right)=\sum_{S} \text { leaf }\left(C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{l o c a l}\right)\right)
$$

To prove Theorem 7, we replace $C_{V}\left(\pi^{\text {out }}\right)-C_{V}\left(\pi^{*}\right)$ by:

$$
\left(C_{V}\left(\pi^{o u t}\right)-C_{V}\left(\pi^{\text {local }}\right)\right)+\left(C_{V}\left(\pi^{\text {local }}\right)-C_{V}\left(\pi^{*}\right)\right)
$$

We then apply Lemma 8 and replace, in each leaf $S$, the quantity $C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{\text {local }}\right)$ by

$$
\left(C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{*}\right)\right)-\left(C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)\right) .
$$

Now all we need to do is analyze, on the one hand, the expectation of

$$
\begin{equation*}
C_{V}\left(\pi^{l o c a l}\right)-C_{V}\left(\pi^{*}\right)-\sum_{S \text { leaf }}\left(C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)\right), \tag{1}
\end{equation*}
$$

(this only involves $\pi^{\text {local }}$, not $\pi^{\text {out }}$, and will be the bulk of the proof), and on the other hand,

$$
\begin{equation*}
\sum_{S \text { leaf }}\left(C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{*}\right)\right) \tag{2}
\end{equation*}
$$

which we now deal with.
The following lemma motivates the stopping condition of Algorithm 3.

Lemma 9 (Leaf nodes). Let $S$ be a leaf of the recursion tree. Then $C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{*}\right) \leq \epsilon F_{S}$.

Proof. If $|S|=1$ then the claim is trivial. Otherwise, using Theorem 2, the definition of tree leaves, and defining $\alpha_{\text {min }}=\beta / \epsilon^{3} \leq \alpha$ :

$$
\begin{aligned}
C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{*}\right) & \leq C_{S}\left(\pi^{o u t}\right)-\mathrm{OPT}_{S} \\
& \leq \beta|S|^{2}=\alpha_{m i n} \epsilon^{3}|S|^{2} \\
& \leq\left(\alpha_{m i n} \epsilon^{3}\right)\left(F_{S} /\left(\alpha \epsilon^{2}\right)\right) \leq \epsilon F_{S}
\end{aligned}
$$

Using Lemma 9 in Expression (2), along with the obvious fact $\sum_{S} F_{S}=F_{V}$, and applying Lemma 5 yields

$$
\sum_{S \text { leaf }}\left(C_{S}\left(\pi^{o u t}\right)-C_{S}\left(\pi^{*}\right)\right) \leq(2 \epsilon / b)\left(C_{V}\left(\pi^{*}\right)+C_{V}\left(\pi^{\text {local }}\right)\right),
$$

which is one of the contributions to the right-hand side of Theorem 7.

In order to complete the proof of Theorem 7, all that is left to do is to argue that in expectation, Expression (1) is also $O(\epsilon / b)\left(C_{V}\left(\pi^{*}\right)+C_{V}\left(\pi^{l o c a l}\right)\right)$.

### 3.5 Cost decomposition

To analyze Expression (1), we need to compare $\pi^{\text {local }}$ to $\pi^{*}$. Define the displacement graph with vertex set $[1, n]$ (corresponding to the ranks), and arc set $\left(\pi^{\text {local }}(x), \pi^{*}(x)\right)$ for every $x$. For example, here is the displacement graph if $\pi^{*}$ is in almost perfect agreement with $\pi^{\text {local }}$, except for switching the first and last vertices in the ordering:


12 $n$

The local optimality of $\pi^{l o c a l}$ with respect to single-vertex moves helps bound the difference in cost between $\pi^{*}$ and $\pi^{l o c a l}$.

Let $f(e)$ denotes the cost, for edge $e=\{x, y\}$ of modifying $\pi^{\text {local }}$ by switching the relative order of $x$ and $y$ :

$$
f(x, y)= \begin{cases}w_{x y}-w_{y x} & \pi^{l o c a l}(x)<\pi^{l o c a l}(y) \\ w_{y x}-w_{x y} & \text { o.w. }\end{cases}
$$

Let $T_{S}$ denote the (negation of the) sum of the costs of the single vertex moves that modify $\pi^{l o c a l}$ by taking $x \in S$ and moving it from position $\pi^{l o c a l}(x)$ to position $\pi^{*}(x)$. More precisely:

$$
T_{S}=T_{S}^{(1)}+T_{S}^{(2)}, \text { with }
$$

$$
\begin{aligned}
T_{S}^{(1)} & =\sum_{x \in S} \sum_{y \in S: \pi^{l o c a l}(x)<\pi^{l o c a l}(y) \leq \pi^{*}(x)}-f(x, y) \\
T_{S}^{(2)} & =\sum_{x \in S} \sum_{y \in S: \pi^{*}(x) \leq \pi^{l o c a l}(y) \leq \pi^{l o c a l}(x)}-f(x, y)
\end{aligned}
$$

The two terms $T_{S}^{(1)}$ and $T_{S}^{(2)}$ correspond to arcs from left to right and right to left respectively in the displacement graph. (The negation is there because we are trying to upperbound $\left.C\left(\pi^{\text {local }}\right)-C\left(\pi^{*}\right), \operatorname{not} C\left(\pi^{*}\right)-C\left(\pi^{\text {local }}\right).\right) T_{S}$ considers the displacement arcs one at a time: this is in a sense a first order approximation to the change in cost, since it ignores the interaction between several displacements.

How does $T_{S}$ compare to $C_{S}\left(\pi^{\text {local }}\right)-C_{S}\left(\pi^{*}\right)$ ? Some pairs $\{x, y\}$ contribute the same to $T_{S}$ as to $C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)$, such as:


For some other pairs, there is a discrepancy, such as the following pair which appears in the sum for $T_{S}$ but not in the sum for $C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)$ :


This is what we call a "crossing pair" (see Section 3.7), and we need a corrective term, which we will denote by $\Phi_{V}$, taking those into account. Thus we define:

$$
\Phi_{S}=C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)-T_{S}
$$

The following two Lemmas, proved in later subsections, analyze the increase of the first order $(T)$ and second order $(\Phi)$ effects from the leaves to the root.

Recall that $\alpha \in\left[\alpha_{\text {min }}, 1\right]$ with $\alpha_{\text {min }}=9^{-\log _{3 / 2}\left(1 / \epsilon^{2}\right) / \epsilon}$. The choice of $\alpha$ and of the random $k$ 's determines the random choices in the execution. Let $B$ denote the choices of $k$ for an execution of Algorithm 3 with $\alpha=1$. Smaller $\alpha$ leads to earlier termination, so $B$ also defines the random choices made for any choice of $\alpha$. Thus, for a given $\pi^{\text {local }}$, the random choices are determined by $(B, \alpha)$.

Lemma 10. For a random execution defined by a random choice of $(B, \alpha)$, we have

$$
T_{V}-\mathbf{E}_{B, \alpha}\left[\sum_{S} T_{S}\right] \leq 14 \epsilon F_{V} .
$$

Lemma 11. Let $\gamma=\frac{\sqrt{5} / 3}{1-\sqrt{5} / 3}$. For any $\alpha$, for a random execution defined by a random choice of ( $B$ ), we have

$$
\Phi_{V}-\mathbf{E}_{B}\left[\sum_{S} \Phi_{S}\right] \leq 96 \gamma \epsilon F_{V}
$$

Now we can bound Expression (1). Since Lemma 11 is true for every $\alpha$, it is also true in expectation over $\alpha$. Adding Lemma 10 and substituting the definition of $\Phi$ to both sides yields:

$$
\begin{aligned}
& C_{V}\left(\pi^{\text {local }}\right)-C_{V}\left(\pi^{*}\right)-\mathbf{E}[ \left.\sum_{S}\left(C_{S}\left(\pi^{\text {leaal }}\right)-C_{S}\left(\pi^{*}\right)\right)\right] \\
& \leq(14+96 \gamma) \epsilon F_{V}=O(\epsilon) F_{V}
\end{aligned}
$$

applying Lemma 5 concludes the proof of Theorem 7. It only remains to prove Lemmas 11 and 10 .

### 3.6 Proof of Lemma 10 (analyzing $T_{V}$ )

In this part, we prove Lemma 10. We split $T_{V}-\sum_{S}$ leaf $T_{S}$ into separate terms for each vertex $y$, and bound each term by the displacement $\left|\pi^{*}(y)-\pi^{\text {local }}(y)\right|$ of that vertex times $\epsilon$.

Given two vertices $x$ and $y$, let $E(x)$ be the event that $\pi^{l o c a l}(x)$ is between $\pi^{\text {local }}(y)$ and $\pi^{*}(y)$ but $x$ is not in the leaf containing $y$. First we see that by definition of $T_{V}$ and $\left\{T_{S}\right\}$,

$$
T_{V}-\mathbf{E}_{B, \alpha}\left[\sum_{S \text { leaf }} T_{S}\right]=\mathbf{E}_{B, \alpha}\left[\sum_{y} \sum_{x: E(x)} f(x, y)\right]
$$

Let $s(y)$ be the size of the leaf containing $y$ (number of vertices in that set). We can split the sum for each $y$ into several parts depending on the relative size of the leaf containing $y$ and $\left|\pi^{*}(y)-\pi^{l o c a l}(y)\right|$. The cases correspond to

- small leaves: $\left|\pi^{*}(y)-\pi^{l o c a l}(y)\right|>s(y) / \epsilon$, and
- big leaves: $\epsilon s(y)>\left|\pi^{*}(y)-\pi^{l o c a l}(y)\right|$,
- intermediate leaves: $\epsilon s(y) \leq\left|\pi^{*}(y)-\pi^{l o c a l}(y)\right| \leq s(y) / \epsilon$.

Small leaves are easy thanks to local optimality.
Lemma 12 (Small Leaves). For any execution (determined by $B$ and $\alpha$ ), we have

$$
\sum_{y \in \text { Small leaf }} \sum_{x: E(x)}-f(x, y) \leq \epsilon F_{V} .
$$

Proof Sketch. Moving $y$ cannot improve $\pi^{\text {local }}$.
We handle the big leaves by proving that the probability that $\pi^{l o c a l}(y)$ and $\pi^{*}(y)$ are in different leaves (and hence able to contribute to the sum) is $O(\epsilon)$.

Lemma 13 (Big Leaves). For any $\alpha$ and for a random $B$, we have:

$$
\mathbf{E}_{B}\left[\sum_{y \in \text { big leaf }} \sum_{x: E(x)}-f(x, y)\right] \leq 12 \epsilon F_{V} .
$$

Proof Sketch. Let $\alpha$ be fixed. If $\pi^{l o c a l}(x)$ is between $\pi^{\text {local }}(y)$ and $\pi^{*}(y)$, but $x$ is not in leaf $(y)$, then the vertex whose rank in $\pi^{\text {local }}$ is $\pi^{*}(y)$ is also not in leaf $(y)$. Thus we can bound the expression on the left hand side by the expectation (over the random tree construction) of
$\sum_{y}\left|\pi^{*}(y)-\pi^{l o c a l}(y)\right| \cdot \mathbb{1}\binom{y \in \operatorname{big}$ leaf }{ and $\pi^{l o c a l-1}\left(\pi^{*}(y)\right) \notin \operatorname{leaf}(y)}$.
One can that for any vertex $y$, the event $E_{1}$ that " $y \in \operatorname{big}$ leaf and $\pi^{\text {local }}{ }^{-1}\left(\pi^{*}(y)\right) \notin \operatorname{leaf}(y) "$ has probability at most $9 \epsilon$ over the sequence $B$ of random choices defining the decomposition. The intuition is that the random choice of $k$ is unlikely to land between $\pi^{\text {local }}(y)$ and $\pi^{*}(y)$.

We bound the contribution of the intermediate leaves by using the variation in $\alpha$ to force a variation in the leaf size, making it unlikely that a given vertex $y$ will be in an intermediate leaf.

Lemma 14 (Intermediate Leaves). For any fixed $B$ and for a random $\alpha$, we have

$$
\mathbf{E}_{\alpha}\left[\sum_{y \in \text { intermediate leaf }} \sum_{x: E(x)}-f(x, y)\right] \leq \epsilon F_{V} .
$$

Proof Sketch. Let $B=\left(k_{S}\right)$ be fixed.
As in the beginning of the proof of Lemma 13, we can bound the expression on the left hand side by the expectation (over the random choice of $\alpha$ ) of

$$
\sum_{y} \begin{cases}0 & \text { if } s(y)=n \\ \sum_{x: E(x)}-f(x, y) & \text { if } s(y)=1 \\ \left|\pi^{*}(y)-\pi^{l o c a l}(y)\right| \cdot & \\ \mathbb{1}(y \in \text { intermediate leaf and } & \\ \left.\pi^{\text {local }}{ }^{-1}\left(\pi^{*}(y)\right) \notin \operatorname{leaf}(y)\right) & \text { otherwise. }\end{cases}
$$

One can argue that for any $y$, the event that " $y \in$ intermediate leaf, $\pi^{\text {local }}{ }^{-1}\left(\pi^{*}(y)\right) \notin \operatorname{leaf}(y)$ and $1<s(y)<n "$ has low probability; but here for the first time, the probabilistic space is over the random definition of $\alpha$.

### 3.7 Proof of Lemma 11 (analyzing $\Phi_{V}$ )

First, we need an algebraic formula expressing $\Phi_{V}$ in terms of contributions of vertex pairs. We say that a pair of vertices $x, y$ such that $\pi^{\text {local }}(x)<\pi^{l o c a l}(y)$ is a crossing pair if exactly one of the following four cases occur:

$$
\left\{\begin{array}{llll}
\pi^{\text {local }}(x) & <\pi^{*}(y) & <\pi^{*}(x) & <\pi^{l o c a l}(y) \\
\pi^{*}(y) & \leq \pi^{l o c a l}(x) & \leq \pi^{\text {local }}(y) & \leq \pi^{*}(x) \\
\pi^{*}(x) & \leq \pi^{*}(y) & \leq \pi^{\text {local }}(x) & \leq \pi^{l o c a l}(y) \\
\pi^{\text {local }}(x) & \leq \pi^{l o c a l}(y) & \leq \pi^{*}(x) & \leq \pi^{*}(y)
\end{array}\right.
$$

We define $\delta_{x y}$ to be equal to 1 in the first case, -1 in the other three cases, and to be equal to 0 for non-crossing pairs. Intuitively, $\{x, y\}$ is a crossing pair if, when we move $x$ continuously from position $\pi^{\text {local }}(x)$ to position $\pi^{*}(x)$, we pass exactly one of $\left\{\pi^{l o c a l}(y), \pi^{*}(y)\right\}$.

Lemma 15 (Mistake Decomposition). We have:

$$
\Phi_{S}=\sum_{x, y \in S}-f(x, y) \cdot \delta_{x y}
$$

Proof. $\Delta C \equiv C_{S}\left(\pi^{l o c a l}\right)-C_{S}\left(\pi^{*}\right)=\sum_{-}-f(x, y)$, where the sum is over pairs $x, y$ in $S$ such that $\pi^{\sum_{\text {ocal }}}(x)<\pi^{\text {local }}(y)$ and $\pi^{*}(x)>\pi^{*}(y)$. To prove the Lemma, it suffices to show that each term $f(x, y)$ appears with the same coefficient in $\Delta C$ as it does in $T_{S}+\Phi_{S}$. First rewrite $T_{S}^{(2)}$ by swapping the names $x$ and $y$ and using the fact that $f$ is symmetric to yield $T_{S}^{(2 \prime)}=\sum-f(x, y)$, where the sum is now over pairs $x, y$ in $S$ such that $\pi^{*}(y) \leq \pi^{\text {local }}(x)<\pi^{\text {local }}(y)$. Note that every pair $(x, y)$ in the sums $\Delta C, T_{S}^{(1)}, T_{S}^{(2 \prime)}$ and $\Phi_{S}$ has $\pi^{\text {local }}(x)<\pi^{\text {local }}(y)$. We divide the pairs $(x, y)$ into seven sets based on the truth of the inequalities $\pi^{*}(x)>\pi^{*}(y)$, $\pi^{*}(x) \geq \pi^{\text {local }}(y)$, and $\pi^{*}(y) \leq \pi^{\text {local }}(x)$ (The eighth possibility never happens because it would imply a contradiction with $\pi^{\text {local }}(x)<\pi^{l o c a l}(y)$.) The seven cases are shown in the following table. T indicates that the inequality in the heading is true, integers indicate the coefficient of $f(x, y)$ in the sums, and dots indicate false or zero.

| $\pi^{*}(x)>$ | $\pi^{*}(x) \geq$ | $\pi^{\text {local }}(x) \geq$ | $\Delta C$ | $T_{1}$ | $T_{2}$ | $\Phi$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\pi^{*}(y)$ | $\pi^{\text {local }}(y)$ | $\pi^{*}(y)$ |  |  |  |  |
| $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ | $\cdot$ |
| $\cdot$ | $\cdot$ | $T$ | $\cdot$ | $\cdot$ | 1 | -1 |
| $\cdot$ | $T$ | $\cdot$ | $\cdot$ | 1 | $\cdot$ | -1 |
| $T$ | $\cdot$ | $\cdot$ | 1 | $\cdot$ | $\cdot$ | 1 |
| $T$ | $\cdot$ | $T$ | 1 | $\cdot$ | 1 | $\cdot$ |
| $T$ | $T$ | $\cdot$ | 1 | 1 | $\cdot$ | $\cdot$ |
| $T$ | $T$ | $T$ | 1 | 1 | 1 | -1 |

Our proof of Lemma 11 is by induction on the nodes of the recursion tree. To analyze crossing pairs, the number of vertices such that the displacement arc $\left(\pi^{\text {local }}(x), \pi^{*}(x)\right)$ crosses $k$ is relevant. Let

$$
\begin{gathered}
\psi_{S}^{L}=\left|\left\{x \in S \mid \pi^{*}(x) \leq k<\pi^{l o c a l}(x)\right\}\right|, \text { and } \\
\psi_{S}^{R}=\left|\left\{x \in S \mid \pi^{l o c a l}(x) \leq k<\pi^{*}(x)\right\}\right|
\end{gathered}
$$

Lemma 16 (Core). Let $S=V_{i, j}$, let $L$ be the set of $x \in S$ such that $\pi^{\text {local }}(x) \leq k$, and $R=S \backslash L$, where $V_{i, j}$ and $k$ are defined as in Algorithm 3. Then:

$$
\Phi_{S}-\mathbf{E}_{k}\left[\Phi_{L}+\Phi_{R}\right] \leq \frac{32}{|S|} \sum_{x \in S}\left|\pi^{*}(x)-\pi^{l o c a l}(x)\right| \psi_{S}^{*}
$$

where $\psi_{S}^{*}=\max _{k \in S} \max \left(\psi_{S}^{L}(k), \psi_{S}^{R}(k)\right)$
Proof. Apply Lemma 15 to $\Phi_{S}, \Phi_{L}$ and $\Phi_{R}$. Restrict attention to pairs $x, y$ such that $\pi^{l o c a l}(x)<\pi^{l o c a l}(y)$ (the other pairs have $\delta_{x y}=0$.) Since $\left|\delta_{x y}\right| \leq 1$ and $|f(x, y)| \leq 1$, we can write

$$
\Phi_{S}-\Phi_{L}-\Phi_{R} \leq \mid\left\{(x, y) \in L_{k} \times R_{k}, \text { crossing pair }\right\} \mid,
$$

where we write $L=L_{k}$ and $R=R_{k}$ to make the dependency in $k$ explicit. Now, it is easy to see that the set $K$ from which Algorithm 3 randomly chooses $k$ satisfies $|K| \geq\left|V_{i j}\right| / 4$ whenever $i \neq j$. We thus have:

$$
\begin{aligned}
& \mathbf{E}_{k}\left[\Phi_{S}-\Phi_{L}-\Phi_{R}\right] \leq \\
& \left.\left.\frac{4}{|S|} \right\rvert\,\left\{(x, y, k):(x, y) \in L_{k} \times R_{k}, \text { crossing pair }\right\} \right\rvert\,
\end{aligned}
$$

In such a triple $(x, y, k)$, we must have that $k$ is either between $\pi^{\text {local }}(x)$ and $\pi^{*}(x)$ or between $\pi^{\text {local }}(y)$ and $\pi^{*}(y)$. Moreover, in order for $\{x, y\}$ to be a crossing pair, it must be that $y$ cross over either $\pi^{\text {local }}(x)$ or over $\pi^{*}(x)$, so $y$ is counted in $\psi_{S}^{L}\left(\pi^{l o c a l}(x)\right)$ or in $\psi_{S}^{R}\left(\pi^{\text {local }}(x)\right)$ or in $\psi_{S}^{L}\left(\pi^{*}(x)\right)$ or in $\psi_{S}^{R}\left(\pi^{*}(x)\right)$. Taking the union of these cases yields the Lemma.

Lemma 17. $\psi_{S}^{*} \leq F_{S}^{1 / 2}$.
Proof Sketch. $i \mapsto \psi_{S}(i)$ is a 1-Lipschitz function (it changes by at most 1 when going from $i$ to $i+1$ ), and its integral is bounded by $F_{S}$. It is easy to see that these observations imply that the maximum value of this function cannot be more than $\sqrt{F_{S}}$.

Corollary 18.

$$
\Phi_{S}-\mathbf{E}_{k}\left[\Phi_{L}+\Phi_{R}\right] \leq \frac{32}{|S|} F_{S}^{3 / 2}
$$

At this point we are prepared to show that at any one level the mistakes made are bounded by $32 \epsilon$ times the optimum cost. To do this, note that by the stopping condition $\sqrt{F_{S}} /|S|<\epsilon$, so $F_{S}^{3 / 2} /|S|<\epsilon F_{S}$. This shows that any one level of the tree does not contribute excessive mistakes. The following lemma proves that the sum of mistakes over all of the levels is dominated by the nodes near the leaves of the tree.

Lemma 19 (Inductive). For every $\alpha \in\left[\alpha_{\text {min }}, 1\right]$,

$$
\Phi_{V} \leq \mathbf{E}_{B}\left[\sum_{S \text { leaf }}\left[96 \gamma \epsilon F_{S}+\Phi_{S}\right]\right]
$$

where $\gamma=\frac{\sqrt{5}}{3-\sqrt{5}}$.
Proof. Let $\alpha \in\left[\alpha_{\text {min }}, 1\right]$ be fixed. If $V$ is a leaf, then the statement is true. Else, given $\alpha$, we first prove by induction that for any $l \leq m$, with $U=\left\{x \mid \ell \leq \pi^{l o c a l}(x) \leq m\right\}$ and $\mathbf{E}[\cdot \mid U \in B]$ meaning expectation conditioned on $\bar{U}$ being a leaf or internal node in execution tree $B$ :

$$
\begin{align*}
\Phi_{U} \leq & \mathbf{E}_{B}\left[\sum_{\substack{S \text { internal } \\
\text { descendant of } U}} 32 \frac{F_{S}^{3 / 2}}{|S|}\right. \\
& \left.+\sum_{\substack{S \text { leaf } \\
\text { descendant of } U}} \Phi_{S} \mid U \in B\right] . \tag{3}
\end{align*}
$$

The base case of $U$ being a leaf is trivial.
Let $k$ be the random variable used at $U$ to decompose $U$ into $L$ and $R$. Let $X_{k}=\Phi_{U}-\Phi_{L}-\Phi_{R}$. Let $B_{1}$ be the random sequence used to construct the decomposition of the left child, and $B_{2}$ be the random sequence used to construct the decomposition of the right child, so that $B=$ $\left(k, B_{1}, B_{2}\right)$. Note that given an interval $U, \Phi_{U}=\mathbf{E}_{B}\left[\Phi_{U}\right]$ is deterministic. Therefore:

$$
\Phi_{U} \mathbf{E}_{B}\left[X_{k}+\Phi_{L}+\Phi_{R}\right]=\mathbf{E}_{k}\left[X_{k}\right]+\mathbf{E}_{k}\left[\Phi_{L}+\Phi_{R}\right]
$$

Using the induction hypothesis for fixed $L, R$ we have:

$$
\begin{aligned}
\Phi_{L}+\Phi_{R} \leq & \mathbf{E}_{B_{1}, B_{2}}\left[\sum_{S \text { internal desc. of } L \text { or } R} 32 \frac{F_{S}^{3 / 2}}{|S|}\right. \\
& \left.+\sum_{S \text { leaf desc. of } L \text { or } R} \Phi_{S} \mid U, L, R \in B\right]
\end{aligned}
$$

Take expectation over $k$ and rewriting:

$$
\begin{aligned}
\mathbf{E}_{k}\left[\Phi_{L}+\Phi_{R}\right] \leq & \mathbf{E}_{B}\left[\sum_{S \text { internal strict desc. of } U} 32 \frac{F_{S}^{3 / 2}}{|S|}\right. \\
& \left.+\sum_{S \text { leaf desc. of } U} \Phi_{S} \mid U \in B\right]
\end{aligned}
$$

Adding $\mathbf{E}_{k}\left[X_{k}\right] \leq 32 F_{U}^{3 / 2} /|U|$ (from Corollary 18) to this and combining $32 F_{U}^{3 / 2} /|U|$ with the sum over internal nodes completes the induction, proving Equation 3

The rest of the proof is deterministic. We prove that for every tree decomposition $B$, the argument of the expectation on the right hand side of Equation 3 with $\ell, m=1, n$ and $U=V$ is bounded by $O(1) \epsilon F_{V}$. Therefore the expectation must also be bounded. The proof uses the following algebraic claim which is easily checked. Recall that $\gamma=\frac{\sqrt{5} / 3}{1-\sqrt{5} / 3}$. If $0 \leq F_{1} \leq F, s>0$, and $s_{1}>0$ such that $s_{1},\left(s-s_{1}\right) \geq s / 3$, then:

$$
\frac{F^{3 / 2}}{s}<\sqrt{5} / 3\left(\frac{F_{1}^{3 / 2}}{s_{1}}+\frac{\left(F-F_{1}\right)^{3 / 2}}{s-s_{1}}\right)
$$

Thus, for every internal node $S$, we have:

$$
F_{S}^{3 / 2} /|S| \leq \sqrt{5} / 3\left(F_{S_{1}}^{3 / 2} /\left|S_{1}\right|+F_{S_{2}}^{3 / 2} /\left|S_{2}\right|\right)
$$

Summing over the recursion tree, we obtain:

$$
\sum_{S \text { internal }} \frac{F_{S}^{3 / 2}}{|S|} \leq \sum_{S \text { leaf }} \frac{F_{S}^{3 / 2}}{|S|}\left(\sqrt{5} / 3+(\sqrt{5} / 3)^{2}+\cdots\right)
$$

which is bounded by $\gamma \sum_{S}$ leaf $\frac{F_{S}^{3 / 2}}{|S|}$. If $S$ is a leaf, then let $P$ be the parent of $S .^{2}$ The parent is not a leaf, so by the stopping condition $\sqrt{F_{P}} /|P| \leq \alpha \epsilon \leq \epsilon$. Thus:

$$
\frac{F_{S}^{3 / 2}}{|S|}=F_{S} \frac{\sqrt{F_{S}}}{|S|} \leq F_{S} \frac{\sqrt{F_{P}}}{|P|} \frac{|P|}{|S|} \leq 3 F_{S} \frac{\sqrt{F_{P}}}{|P|} \leq 3 \epsilon F_{S} .
$$

Summing yields the lemma.
Combining Lemma 19 with the simple fact that $\sum_{S}$ leaf $F_{S}=$ $F_{V}$ proves Lemma 11.

## 4. RUNNING TIME

### 4.1 Improved running time: Randomization

The running time of Algorithms 2 and 1 can be improved somewhat by replacing the deterministic additive error algorithm with the randomized one. To use the randomized version of the AddApprox algorithm set $\eta=\delta / n$ for Algorithm 2 and $\eta=\delta \epsilon / n^{4}$ for Algorithm 1. Consider the event

[^3]$E_{0}$ stating that "During execution of the algorithm, every call to the randomized version of AddApprox yields a result within the stated bounds." Each call fails with probability at most $\eta$. As shown in Section 4.2, there are at most $n$ (resp. $n^{4} / \epsilon$ ) such calls, so event $E_{0}$ has probability at least $1-\delta$. Modify the analysis by assuming throughout that event $E_{0}$ holds and do the analysis conditioned on $E_{0}$.

Using the randomized version of the AddApprox algorithm introduces a complication because the proof of Lemma 3 implicitly assumes determinism. This is easily remedied by reusing the random numbers used the previous time randomized AddApprox was called on the same vertices.

### 4.2 Analysis: Running time

Both Algorithm 1 and 2 have a greedy local search phase. The preprocessing step ensures that the edge weights are integer multiples of $\epsilon / n^{2}$, so the cost decreases by at least $\epsilon / n^{2}$ at each iteration. Since the cost is always between $n^{2}$ and zero, the total number of iterations is bounded by $n^{4} / \epsilon$.

The single vertex move local search phase takes time $O\left(n^{2}\right.$. $n^{4} / \epsilon$ ), where the first term comes from the time to search for an improving single vertex move and the second is a bound on the number of iterations.

The runtime is dominated by the calls to the additive error algorithm. There are at most $n$ calls for the randomized version (Algorithm 2) and at most $n^{2} \cdot n^{4} / \epsilon$ for the deterministic one (Algorithm 1), where the $n^{2}$ comes from the number of intervals that need to be checked each iteration to find an improving move. Thus the runtime is

$$
O\left(n^{6} / \epsilon+n f\left(n,(1 / \epsilon)^{O(1 / \epsilon)}, \delta / n^{4}\right)\right)
$$

randomized and

$$
O\left(n^{6} / \epsilon \cdot f\left(n,(1 / \epsilon)^{O(1 / \epsilon)}, 0\right)\right)
$$

deterministic, where $f(n, \beta, \eta)$ is the time required for the additive approximation algorithm to run on the problem of size $n$ with error parameter $\beta$ and failure probability $\eta$. The best-known additive approximation run time $f(n, \beta, \eta)$ is $\left(n / \beta^{4}+2^{O\left(1 / \beta^{2}\right)}\right) \log 1 / \eta$ randomized in [22], or $n^{O\left(1 / \beta^{4}\right)}$ for deterministic $(\eta=0)$.

## 5. CONCLUSIONS AND FUTURE WORK

The reader may wonder whether our results can extend to the dense case, with $\Omega\left(n^{2}\right)$ edges but no constraint that every pair of vertices must have an edge between them. This extension is not feasible because the dense case is as hard as the general case, which is APX-hard. To see this, consider an arbitrary directed graph with $n$ vertices that we want the feedback arc set for. Union this graph with an arbitrary acyclic dense graph on $n$ vertices. The result is a dense graph with $2 n$ vertices that is essentially equivalent to the original problem.

Interesting open questions include:

1. Does weighted feedback arc set with the triangle inequality admit a PTAS? This would give a PTAS for partial rank aggregation [2].
2. Can the runtime be improved to be only singly exponential in $\epsilon$ ? The doubly exponential runtime comes from the intermediate leaves case in the $T$ analysis (Lemma 14).
3. Is single vertex move local optimality sufficient to achieve a constant factor approximation? We have an example showing that such a constant factor is at least three Also does sorting by indegree followed by single vertex moves have a provably better approximation factor than either one by itself?

## Acknowledgements

The authors wish to thank Nir Ailon, Eli Upfal and anonymous reviewers for useful suggestions.

## 6. REFERENCES

[1] Nir Ailon, Moses Charikar, Alantha Newman. Aggregating inconsistent information: ranking and clustering. ACM Symposium on Theory of Computing (STOC), pages 684-693, 2005.
[2] Nir Ailon. Aggregation of Partial Rankings, p-Ratings and Top-m Lists. ACM-SIAM Symposium on Discrete Algorithms (SODA), 2007, 415-424.
[3] N. Alon, Ranking tournaments, SIAM J. Discrete Math. 20 (2006), 137-142.
[4] Sanjeev Arora, Alan Frieze, and Haim Kaplan. A New Rounding Procedure for the Assignment Problem with Applications to Dense Graph Arrangement Problems. IEEE Symposium on Foundations of Computer Science (FOCS) 1996: 24-33.
[5] F.B.Baker and L.J.Hubert, Applications of combinatorial programming to data analysis: seriation using asymmetric proximity measures, British J. Math. Statis. Psych. 30 (1977) 154-164.
[6] J. Borda. Mémoire sur les élections au scrutin. Histoire de l'Académie Royale des Sciences, 1781.
[7] M.C. Cai, X. Deng, W. Zang. An approximation algorithm for feedback vertex sets in tournaments. SIAM J. Computing v. 30 No. 6 pp. 1993-2007. 2001.
[8] M. Charikar, V. Guruswami, and A. Wirth. Clustering with qualitative information. In $F O C S$, pages 524-533, 2003.
[9] William W. Cohen, Robert E. Schapire and Yoram Singer. Learning to order things. Proceedings of the 1997 conference on Advances in neural information processing systems (NIPS): 451-457.
[10] M. Condorcet. Essai sur l'application de l'analyse à la probabilité des décisions rendues à la pluralité des voix. 1785.
[11] Conitzer, V. Computing Slater Rankings Using Similarities Among Candidates. Proceedings of the 21st National Conference on Artificial Intelligence (AAAI-06), Boston, MA, USA, 613-619. 2006.
[12] D. Coppersmith, Lisa Fleischer, Atri Rudra: Ordering by weighted number of wins gives a good ranking for weighted tournaments. ACM-SIAM Symposium on Discrete Algorithms (SODA) 2006: 776-782.
[13] C. Demetrescu and I. Finocchi. Break the "right" cycles and get the "best" drawing. In B.E. Moret and A.V. Goldberg, editors, Proceedings of the 2nd International Workshop on Algorithmic Engineering and Experiments (ALENEX), 171-182, 2000.
[14] P. Diaconis and R. Graham. Spearman's footrule as a measure of disarray. J. of the Royal Statistical Society, 39(2), pages 262-268, 1977.
[15] I. Dinur and S. Safra. The importance of being biased. In Proceedings on 34th Annual ACM Symposium on Theory of Computing, pages 33-42, 2002.
[16] Cynthia Dwork, Ravi Kumar, Moni Naor, and D. Sivakumar. Rank aggregation revisited. Manuscript. 2001.
[17] Cynthia Dwork, Ravi Kumar, Moni Naor, D. Sivakumar. Rank Aggregation Methods for the Web. WWW10. 2001.
[18] P. Erdös and J.W.Moon, On sets on consistent arcs in tournaments, Canadian Mathematical Bulletin 8 (1965) 269-271.
[19] G. Even, J. Naor, S. Rao, B. Schieber. Divide-and-conquer approximation algorithms via spreading metrics. Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 62-71, 1995.
[20] Fernandez de la Vega, W. On the maximal cardinality of a consistent set of arcs in a random tournament, $J$. Combinatorial Theory, Ser. B 35:328-332 (1983).
[21] M.M.Flood. Exact and Heuristic Algorithms for the Weighted Feedback Arc Set Problem: a Special Case of the Skew-Symmetric Quadratic Assignment Problem. Networks, 20:1-23, 1990.
[22] Alan M. Frieze, Ravi Kannan: Quick Approximation to Matrices and Applications. Combinatorica 19(2): 175-220 (1999).
[23] M. Grötschel, M. Jünger, G. Reinelt. Facets of the linear ordering polytope. Math. Programming 33, 1985, 43-60.
[24] M. Grötschel, M. Jünger, G. Reinelt. On the acyclic subgraph polytope. Math. Programming 33, 1985, 28-42.
[25] R. Hassin and S. Rubinstein. Approximations for the maximum acyclic subgraph problem. Information processing letters 51 (1994) 133-140.
[26] H.A. Jung, On subgraphs without cycles in tournaments, Combinatorial Theory and its Applications II, North Holland (1970) 675-677.
[27] M. Jünger, Polyhedral combinatorics and the Acyclic Subdigraph Problem (Heldermann Verlag, Berlin, 1985).
[28] Karp, R., 1972. Reducibility among combinatorial problems. In: Miller, R. E., Thatcher, J. W. (Eds.), Complexity of Computer Computations. Plenum Press, NY, pp. 85-103.
[29] J. Kemeny. Mathematics without numbers. Daedalus, 88:571-591, 1959.
[30] J. Kemeny and J. Snell. Mathematical models in the social sciences. Blaisdell, New York, 1962. Reprinted by MIT press, Cambridge, 1972.
[31] Claire Kenyon-Mathieu and Warren Schudy. How to rank with few errors: A PTAS for Weighted Feedback Arc Set on Tournaments ECCC Report TR06-144, Nov 2006.
[32] B. Korte, Approximative algorithms for discrete optimization problems, Ann. Discrete Math 4 (1979), 85-120.
[33] T. Leighton and S. Rao. An approximate max-flow min-cut theorem for uniform multicommodity flow problems with applications to approximation algorithms. In Proceedings of the 29th Annual Symposium on Foundations of Computer Science, pages 422-431. IEEE, October 1988.
[34] A. Lempel and I. Cederbaum. Minimum feedback arc and vertex set of a directed graph. IEEE Trans. Circuit Theory, 13(4): 399-403, 1966.
[35] K.B.Reid, On sets of arcs containing no cycles in tournaments, Canad. Math Bulletin 12 (1969) 261-264.
[36] K.D.Reid and E.T.Parker, Disproof of a conjecture of Erdös and Moser on tournaments, J. Combin. Theory 9 (1970) 225-238.
[37] S. Seshu and M.B. Reed, Linear Graphs and Electrical Networks, Addison-Wesley, Reading, MA, 1961. Attribute the problem of the minimum feedback arc set to Runyon.
[38] P.D.Seymour. Packing directed circuit fractionally. Combinatorica 15:281-288, 1995.
[39] P. Slater, Inconsistencies in a schedule of paired comparisons, Biometrika 48 (1961) 303-312.
[40] J. Spencer, Optimal ranking of tournaments, Networks 1(1971) 135-138.
[41] J. Spencer, Optimal ranking of unrankable tournaments, Period. Math. Hungar. 11-2 (1980) 131-144.
[42] K. Sugiyama, S. Tagawa, and M. Toda. Methods for visual understanding of hierarchical system structures. IEEE Trans. Syst. Man Cybern., 11(2):109-125, 1981.
[43] A.W.Tucker. On directed graphs and integer programs. Presented at 1960 Symposium on Combinatorial Problems, Princeton University, cited in [44].
[44] D.H. Younger. Minimum feedback arc sets for a directed graph. IEEE Trans. Circuit Theory 10 (1963), 238-245.
[45] A. van Zuylen. Deterministic approximation algorithms for ranking and clusterings. Cornell ORIE Tech report No. 1431. 2005.
[46] Anke van Zuylen, Rajneesh Hegde, Kamal Jain and David P. Williamson. Deterministic pivoting algorithms for constrained ranking and clustering problems. ACM-SIAM Symposium on Discrete Algorithms (SODA), 2007, 405-414.


[^0]:    *An extended version is available [31]
    ${ }^{\dagger}$ Corresponding author. Email: ws@cs.brown.edu

[^1]:    Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.
    STOC'07, June 11-13, 2007, San Diego, California, USA.
    Copyright 2007 ACM 978-1-59593-631-8/07/0006 ...\$5.00.

[^2]:    ${ }^{1}$ This problem is identical to FAS except it deals with symmetric transitive relations instead of antisymmetric ones.

[^3]:    ${ }^{2}$ If $S$ is the root, no parent is available but in that case the lemma is trivially true anyway.

