# Near-linear Time Algorithms for Approximate Minimum Degree Spanning Trees 

Ran Duan * Haoqing He ${ }^{\dagger}$ Tianyi Zhang ${ }^{\ddagger}$


#### Abstract

Given a graph $G=(V, E), n=|V|, m=|E|$, we wish to compute a spanning tree whose maximum vertex degree is as small as possible. Computing the exact optimal solution is known to be NP-hard, since it generalizes the Hamiltonian path problem. For the approximation version of this problem, a $\tilde{O}(m n)$ time algorithm that computes a spanning tree of degree at most $\Delta^{*}+1$ is previously known [Fürer, Raghavachari 1994]; here $\Delta^{*}$ denotes the optimal tree degree. In this paper we give the first near-linear time algorithm for this problem. Specifically speaking, we first propose a simple $\tilde{O}(m)$ time algorithm that achieves an $O\left(\Delta^{*} \log n\right)$ approximation; then we further improve this algorithm to obtain a $(1+\delta) \Delta^{*}+O\left(\frac{1}{\delta^{2}} \log n\right)$ approximation in $\tilde{O}\left(\frac{1}{\delta^{6}} m\right)$ time.


[^0]
## 1 Introduction

Computing minimum degree spanning trees is a fundamental problem that has inspired a long line of research. Let $G=(V, E)$ be an undirected graph, and we wish to compute a spanning tree of $G$ whose tree degree, or maximum vertex degree in the tree, is the smallest. Clearly this problem is NP-hard as the Hamiltonian path problem can be reduced to it, and so we could only hope for a good approximation in polynomial time. The optimal approximation of this problem was achieved in [6] where the authors proposed an $\tilde{1} \tilde{O}(m n)$ time algorithm that computes a spanning tree of tree degree $\leq \Delta^{*}+1$; conventionally $n=|V|, m=|E|$ and $\Delta^{*}$ denotes the optimal tree degree. However, polynomial time algorithms does not always mean efficient on large data sets, so finding approximation algorithms of almost linear time is a very popular and important topic nowadays.

### 1.1 Our results

The major results of this paper are two near-linear time algorithms for minimum degree spanning trees in undirected graphs. These are the first near-linear time algorithms for this problem. Formally we propose the following two theorems.

Theorem 1. There is an 2 2 $O\left(m \alpha(n) \log ^{2} n+n \log ^{3} n\right)$ time algorithm that computes a spanning tree with tree degree $O\left(\Delta^{*} \log n\right)$.

As in many algorithms of this problem such as [6, this algorithm iteratively improves the spanning tree T by finding replacement edge connecting two low-degree vertices. To achieve almost linear time, we fix a degree threshold $k \geq 3$, and repeatedly search for edges connecting two vertices of tree degree $\leq k-2$ such that the tree path between its two endpoints contains a vertex of tree degree $\geq k$. We can efficiently maintain the spanning tree by the link-cut tree structure [16. When there are not many vertices of tree degree $k-1$, we can argue a lower bound on $\Delta^{*}$ in terms of $k$. However, the algorithm may generate a large number of $(k-1)$-degree vertices which undermines the lower bound on $\Delta^{*}$. To circumvent such difficulties, we iteratively perform this procedure on larger and larger $k$ 's. If the number of vertices of degree $\geq k$ becomes smaller and smaller, we can finally bound the number of $(k-1)$-degree vertices. The crucial observation is that if a vertex which starts out as a low-degree vertex for previous $k$ now becomes $(k-1)$-degree, lots of high-degree vertices must have lost some tree neighbours. By carefully selecting a series of threshold $k$ 's, finally we can argue a lower bound on $\Delta^{*}$ or decrease the degree of T by a constant factor.
Theorem 2. For any constant $\delta \in\left(0, \frac{1}{6}\right)$, there is an algorithm that runs in time $O\left(\frac{1}{\delta^{6}} m \log ^{7} n\right)$ which computes a spanning tree with tree degree at most $(1+\delta) \Delta^{*}+\frac{5}{32 \delta^{2}} \log n$.

Theorem 2 refines Theorems approach by an augmenting path approach. In each iteration, the algorithm conducts a series of tree modifications to remove all augmenting paths of the shortest length, and so in the next iteration the shortest length of augmenting paths would increase. To facilitate our search for shortest augmenting paths, we divide the graph into $O\left(\frac{1}{\delta} \log n\right)$ layers and then look for edges that connect two different tree components on the bottom layer. If such an edge is successfully detected, then we add this edge to the tree and propagate a sequence of tree edge insertions and deletions upwards to higher layers. When no such edges can be found, we argue

[^1]that every layer can yield a lower bound on $\Delta^{*}$ which jointly proves a lower bound with a constant multiplicative error.

### 1.2 Related work

There is a line of works that are concerned with low-degree trees in weighted undirected graphs. In this scenario, the target low-degree that we wish to compute is constrained by two parameters: an upper bound $B$ on tree degree, an upper bound $C$ on the total weight summed over all tree edges. The problem was originally formulated in [3]. Two subsequent papers [10, 11] proposed polynomial time algorithms that compute a tree with cost $\leq w C$ and degree $\leq \frac{w}{w-1} b B+\log _{b} n, \forall b, w>1$. The cost was improved from $w C$ to $C$ in [2] while degree upper bound becomes $b B+2(b+1) \log _{b} n$; the authors also proposed an quasi-polynomial algorithm that finds a tree with cost $\leq C$ and degree $B+O(\log n / \log \log n)$. [2]'s result was improved by [8] where for all $k$, a spanning tree of degree $\leq k+2$ and of cost at most the cost of the optimum spanning tree of maximum degree at most $k$ can be computed in polynomial time. The degree bound was later further improved from $k+2$ to the optimal $k+1$ in [15].

Another variant is minimum degree Steiner trees which is related to network broadcasting [13, 14, 4]. For undirected graphs, authors of [6] showed that the same approximation guarantee and running time can be achieved as with minimum degree spanning trees in undirected graphs, i.e., a solution of tree degree $\Delta^{*}+1$ and a running time of $\tilde{O}(m n)$. For the directed case, [4] showed that directed minimum degree Steiner trees problem cannot be approximated within $(1-\epsilon) \log |D|, \forall \epsilon>0$ unless NP $\subseteq \operatorname{DTIME}\left(n^{\log \log n}\right)$, where $D$ is the set of terminals.

The minimum degree tree problem can also be formulated in directed graphs. This problem was first studied in [5] where the authors proposed a polynomial time algorithm that finds a directed spanning tree of degree at most $O\left(\Delta^{*} \log n\right)$. The approximation guarantee was improved to roughly $O\left(\Delta^{*}+\log n\right)$ in [12, 9] while the time complexity became $n^{O(\log n)}$. The problem becomes much easier when $G$ is acyclic, as shown in [17], where a directed spanning tree of degree $\leq \Delta^{*}+1$ is computable in polynomial time. The approximation was greatly advanced to $\Delta^{*}+2$ in [1] by an LP-based polynomial time algorithm, and this problem has become more-or-less closed since then.

## 2 Definitions

Logarithms are taken at base 2. Assume $G=(V, E)$ is a connected graph, $n=|V|, m=|E|$, and we assume $n \geq 4$. During the execution of our algorithm, a spanning tree T will be maintained and our algorithm will repeatedly modify T to reduce its degree $\Delta=\operatorname{deg}(\mathrm{T})$. For every $u \in V$, let $\operatorname{deg}(u)$ be the tree degree of $u$. For each pair $u, v \in V$, let $\rho_{u, v}$ be the unique tree path on T that connects $u$ and $v$. For each $1 \leq k \leq n$, define $S_{k}=\{u \mid \operatorname{deg}(u) \geq k\}, N_{k}=\{u \mid \operatorname{deg}(u)=k\}$, and let $d_{k}=\sum_{u \in S_{k}} \operatorname{deg}(u)$, that is, the total degree of all vertices of degree at least $k$.

## 3 An $O\left(\Delta^{*} \log n\right)$ Approximation

### 3.1 Main algorithm

Starting from an arbitrary spanning tree $T$ with degree $\Delta$, the core of the main algorithm is a near-linear time subroutine that, as long as $\Delta \geq 20 \log n$, either reduces $\Delta$ to $(1-\Omega(1)) \cdot \Delta$ or
terminates with the guarantee that $\Delta=O\left(\Delta^{*} \log n\right)$; the main algorithm simply repeatedly apply this subroutine until $\Delta<20 \log n$ or $\Delta=O\left(\Delta^{*} \log n\right)$. This subroutine consists of two parts: (1) a low-level fast degree reduction algorithm that, given any degree threshold $k$, modifies T to reduce the total number of high-degree vertices; (2) a high-level scheduling algorithm that selects a sequence of degree thresholds and feed them to the low-level degree reduction algorithm as inputs. For every $0 \leq i \leq \log n+1$, let us define a sequence of degree thresholds:

$$
k_{i}=\left\lfloor\frac{3}{4} \Delta+\frac{1}{4} \Delta \cdot\left(1-\left(1-\frac{1}{\log n}\right)^{i}\right)\right\rfloor
$$

Clearly $k_{i}$ 's are increasing as $\forall 0 \leq i \leq \log n$,

$$
k_{i+1}-k_{i} \geq \frac{1}{4} \Delta \cdot\left(1-\frac{1}{\log n}\right)^{i}-\frac{1}{4} \Delta \cdot\left(1-\frac{1}{\log n}\right)^{i+1}-1 \geq \frac{1}{16 \cdot \log n} \cdot \Delta-1>0
$$

The last two inequality holds as $\left(1-\frac{1}{\log n}\right)^{i} \geq 1 / 4$ and $\Delta \geq 20 \log n$.
The high-level scheduling algorithm (2) is described in Scheduling shown in Algorithm If it returns false, an upper bound $\Delta=O\left(\Delta^{*} \log n\right)$ would be established; otherwise when it returns true, it means $\Delta$ would be reduced to $(1-\Omega(1)) \cdot \Delta$. The low-level degree reduction algorithm (1) is described in FastDegreeReduction in Algorithm 2. The rough idea is that we repeatedly looks for edges that connect two vertices of tree degree $\leq k-2$ from different components of $\mathrm{T} \backslash S_{k}$ and add these edges to T, while at the same time we delete some edges incident on $S_{k}$ so T stays a tree. In order to implement this idea in near-linear time, we have to neglect those $(k-2)$-degree vertices that have once become ( $k-1$ )-degree. A key operation of our algorithm is marking. During one execution of FastDegreeReduction, a vertex gets marked whenever its tree degree becomes $k-1$ and it stays so even if its degree goes smaller, and instead of searching for edges between two ( $k-2$ )-degree vertices, we only care about edges between two unmarked vertices.

```
Algorithm 1: Scheduling
    input params: T with maximum tree degree \(\Delta \geq 20 \log n\);
    for \(i=1,2, \cdots, 1+\log n\) do
        count \(\leftarrow\left|S_{k_{i-1}}\right|\);
        if \(\left|S_{k_{i}}\right|>\frac{1}{2} \cdot\) count then
            invoke FastDegreeReduction with input \(k=k_{i}\);
            if \(\left|S_{k_{i}}\right|>\frac{1}{2}\) - count then
                return false;
    return true;
```


### 3.2 Implementation and running time

We specify some implementation details of FastDegreeReduction.
(1) To efficiently implement line-4, we enumerate all edges $(u, v) \in E$ one by one. Using the unionfind data structure [7, we check if one of $u, v$ is marked or both of $u, v$ belong to the same component, we move on to the next edge; otherwise we execute line- 5 through line-13. The total running time of this part would be $O(m \alpha(n))$.

```
Algorithm 2: FastDegreeReduction
    input params: \(k\);
    for every \(u \notin S_{k}\), let \(C_{u}\) be the connected component containing \(u\) of \(\mathrm{T} \backslash S_{k}\);
    mark all vertices in \(S_{k-1}\) and unmark all other vertices;
    while \(\exists(u, v) \in E\), both \(u, v\) are unmarked and belong to different components in \(T \backslash S_{k}\) do
        find a vertex \(w \in S_{k} \cap \rho_{u, v}\), and \((w, z)\) be an edge \(\in \rho_{u, v}\) which is incident on \(w\);
        \(\mathrm{T} \leftarrow \mathrm{T} \cup\{(u, v)\} \backslash\{(w, z)\} ;\)
        merge components \(C_{u}\) and \(C_{v}\);
        for \(x \in\{u, v\}\) do
            if \(\operatorname{deg}(x)=k-1\) then
                mark \(x\);
        for \(x \in\{w, z\} \cap S_{k}\) do
            if \(\operatorname{deg}(x)=k-1\) then
                remove \(x\) from \(S_{k}\) and merge tree components of \(\mathrm{T} \backslash S_{k}\) accordingly;
    return;
```

(2) For line-5, to efficiently retrieve a vertex $w \in S_{k} \cap \rho_{u, v}$ given $u, v$, we maintain T using the link-cut tree data structure [16. We set the weight of each $u \in S_{k}$ to be 1 , and weight of each $u \notin S_{k}$ equal to 0 . Then $w \in S_{k} \cap \rho_{u, v}$ can be found in $O(\log n)$ amortized time by querying the maximum weight vertex on the tree path $\rho_{u, v}$ using the link-cut tree data structure. Note that such $S_{k} \cap \rho_{u, v}$ is always non-empty because $u, v$ belong to different connected components of $\mathrm{T} \backslash S_{k}$. For line-6, edge updates to T can be handled using the link-cut tree as well. Since there are less than $n$ components in $\mathrm{T} \backslash B$, the total time would be $O(n \log n)$.
(3) On line-7, merging components $C_{u}$ and $C_{v}$ can be done in $O(\alpha(n))$ time using the union-find data structure.
(4) On line-13, when a vertex $x$ is removed from $S_{k}$, we need to add $x$ to $\mathrm{T} \backslash S_{k}$ and possibly merge some connected components. This can be done by enumerating $x$ 's incident tree edges and using the union-find data structure. The total cost of such operations would be $O(n \alpha(n))$.

To conclude, the overall running time complexity of FastDegreeReduction is by $O(m \alpha(n)+$ $n \log n)$. The running time of Scheduling then becomes $O\left(m \alpha(n) \log n+n \log ^{2} n\right)$ since $i$ can increase to at most $O(\log n)$.

To upper bound the running time of the main algorithm, we need the following lemma that characterize the performance of Scheduling.

Lemma 3. If Scheduling returns true, then the degree $\Delta$ of $T$ has at least dropped by a constant factor of $\frac{31}{32}$.

Proof. When Scheduling returns true, we claim that count declines by a factor of $\frac{1}{2}$ after each iteration of the while-loop. In fact, on the one hand, if the condition of line- 5 does not hold, i.e. $S_{k_{i-1}} \geq 2\left|S_{k_{i}}\right|$, then as count was previous set to $\left|S_{k_{i-1}}\right|$ and now equal to $\left|S_{k_{i}}\right|$, count declines by a factor of $\frac{1}{2}$. On the other hand, if the condition of line- 5 holds, then because Scheduling
returns true, the condition on line-7 always fails, i.e. $\left|S_{k_{i}}\right| \leq \frac{1}{2} \cdot$ count, and therefore when we set count $\leftarrow\left|S_{k_{i}}\right|$ the value of count would decrease by a factor of $\frac{1}{2}$. Hence, the while-loop of Scheduling can iterate for at most $\log n+1$ times before count becomes 0 ; that is to say, $\left|S_{k_{i}}\right|=0$ for some $i \leq \log n+1$. By definition,

$$
\begin{aligned}
k_{i} & \leq \frac{3}{4} \Delta+\frac{1}{4} \Delta \cdot\left(1-\left(1-\frac{1}{\log n}\right)^{i}\right) \\
& \leq \frac{3}{4} \Delta+\frac{1}{4} \Delta \cdot\left(1-\left(1-\frac{1}{\log n}\right)^{\log n+1}\right) \\
& \leq \frac{3}{4} \Delta+\frac{1}{4} \Delta \cdot\left(1-\frac{1}{8}\right)=\frac{31}{32} \cdot \Delta
\end{aligned}
$$

Here we use the fact that $\left(1-\frac{1}{\log n}\right)^{\log n+1} \geq \frac{1}{8}$. As $\left|S_{k_{i}}\right|=0$, $\operatorname{deg}(\mathbf{T})$ must now be smaller than $\frac{31}{32} \cdot \Delta$.

Now we can upper bound the running time of the main algorithm as stated in the following lemma.

Lemma 4. The running time of the main algorithm is $O\left(m \alpha(n) \log ^{2} n+n \log ^{3} n\right)$.
Proof. By Lemma 3, every invocation of Scheduling that returns true decreases $\Delta$ by a factor of $\frac{31}{32}$. Therefore, such kind of invocations can be at most $O(\log n)$ many. Also, there can be at most 1 instance of Scheduling that returns false because the main algorithm terminates immediately after that. Overall, the total running time of Scheduling would be $O\left(\log n \cdot\left(m \alpha(n) \log n+n \log ^{2} n\right)\right)=$ $O\left(m \alpha(n) \log ^{2} n+n \log ^{3} n\right)$.

### 3.3 Approximation guarantee

To prove approximation guarantee, we will utilize the following lemmas.
Lemma 5. Let $V_{1}, V_{2}, \cdots, V_{l} \subseteq V$ be disjoint vertex subsets. A set $W$ is called "boundary" (with respect to $V_{1}, V_{2}, \cdots, V_{l}$ ), if any edge incident on $\bigcup_{i=1}^{l} V_{i}$ whose both endpoints are not simultaneously contained in any single $V_{i}$, is incident on at least one vertex from $W$. Then, $\Delta^{*} \geq \frac{l-1}{|W|}$.
Proof. For any spanning tree, there are at least $l-1$ edges incident on $\bigcup_{i=1}^{l} V_{i}$ whose both endpoints are not simultaneously contained in any $V_{i}, 1 \leq i \leq l$. Then by definition of $W$, any one of these $l-1$ edges is incident on at least on vertex of $W$, and thus by the pigeon-hole principle, there exists a $u \in W$ whose tree degree is $\geq \frac{l-1}{|W|}$.

Lemma 6. For any vertex subset $B$, the number of connected components in $T \backslash B$ is at least $\sum_{u \in B} \operatorname{deg}(u)-2|B|+2$.

Proof. Note that there are at least $\sum_{u \in B} \operatorname{deg}(u)-|B|+1$ tree edges incident on $B$, and so removing all of these edges would break T into $\geq \sum_{u \in B} \operatorname{deg}(u)-|B|+2$ components. Therefore, excluding singleton components formed by $B$, there are $\geq \sum_{u \in B} \operatorname{deg}(u)-2|B|+2$ components are from $\mathrm{T} \backslash B$.

Now we prove when the main algorithm terminates, $\Delta=O\left(\Delta^{*} \log n\right)$. If the main algorithm terminates with $\Delta<20 \log n$, then automatically we have $\Delta<20 \log n=O(\log n)=O\left(\Delta^{*} \log n\right)$. Next we focus on the case when the main algorithm terminates on a false returned by Scheduling. In this case, there was an execution of Scheduling that returned on line-7. By the branching condition of line-6, we know that $\left|S_{k_{i}}\right|>\frac{1}{2} \cdot$ count by the end of this execution of Scheduling.

Consider the most recent invocation of FastDegreeReduction. By the end of this invocation, let $V_{1}, V_{2}, \cdots, V_{l}$ be the sequence of all different connected components spanned by $\mathrm{T} \backslash S_{k}$. Let $M$ be the set of all marked vertices. To apply Lemma 5, we claim $M$ is a boundary set with respect to $V_{1}, V_{2}, \cdots, V_{l}$; this is because, for any edge $(u, v)$ such that $u, v$ belong to different connected components of T $\backslash S_{k}$, one of $u, v$ must be marked since otherwise FastDegreeReduction would continue to merge $C_{u}$ and $C_{v}$ instead of terminating. Therefore, Lemma 5 immediately yields $\Delta^{*} \geq \frac{l-1}{|M|}$.

One last thing is to lower-bound $l$ and upper-bound $|M|$.
(1) Lower bounding $l$.

Let $S_{k_{i-1}}^{\prime}$ and $S_{k_{i}}^{\prime}$ be the snapshots of $S_{k_{i-1}}$ and $S_{k_{i}}$ before this instance of FastDegreeReduction began. So by the algorithm we have count $=\left|S_{k_{i-1}}^{\prime}\right|$ and thus $\left|S_{k_{i}}\right| \geq \frac{1}{2} \cdot$ count $=\frac{1}{2}\left|S_{k_{i-1}}^{\prime}\right| \geq$ $\frac{1}{2}\left|S_{k_{i}}^{\prime}\right|$. Then clearly, the number of connected components of $\top \backslash S_{k}$ is at least

$$
\begin{aligned}
l & \geq \sum_{u \in S_{k}} \operatorname{deg}(u)-2\left|S_{k}\right|+2 \\
& \geq \sum_{u \in S_{k_{i}}} \operatorname{deg}(u)-2\left|S_{k_{i}}\right|+2 \geq k_{i}\left|S_{k_{i}}\right|-2\left|S_{k_{i}}\right|+2 \\
& >k_{i}\left|S_{k_{i}}\right|-2\left|S_{k_{i}}\right| \geq\left(\frac{k_{i}}{2}-1\right) \cdot\left|S_{k_{i}}^{\prime}\right| \geq\left(\frac{3}{8} \Delta+\frac{\Delta}{8 \log n}-\frac{3}{2}\right) \cdot\left|S_{k_{i}}^{\prime}\right|>\frac{3}{8} \Delta\left|S_{k_{i}}^{\prime}\right|
\end{aligned}
$$

The first inequality holds by Lemma 6; the last two inequalities holds by $k_{i}=\left\lfloor\frac{3}{4} \Delta+\frac{1}{4} \Delta(1-\right.$ $\left.\left.\left(1-\frac{1}{\log n}\right)^{i}\right)\right\rfloor>\frac{3}{4} \Delta+\frac{\Delta}{4 \log n}-1$ and $\Delta \geq 20 \log n$.
(2) Upper bounding $|M|$.

There are two kinds of marked vertices.
(a) Either $u \in S_{k_{i}-1}^{\prime}$ was marked at the beginning, or $u \in S_{k_{i-1}}^{\prime} \backslash S_{k_{i}-1}^{\prime}$ whose tree degree later got increased to $k_{i}-1$ at some point while the algorithm kept modifying T . The total number of such vertices is at most $\left|S_{k_{i-1}}^{\prime}\right|<2\left|S_{k_{i}}^{\prime}\right|$.
(b) $u$ is a marked vertex and $u \notin S_{k_{i-1}}^{\prime}$. In this case, $\operatorname{deg}(u)<k_{i-1}$ before this instance of FastDegreeReduction began. Since $u$ is marked, $\operatorname{deg}(u)$ increases to $k_{i}-1$ at some point. Every time we modify T , at least one vertex in $S_{k_{i}}$ loses one degree and at most two unmarked vertices get one degree separately. So for a vertex $u \notin S_{k_{i-1}}^{\prime}$ to be marked, the vertices in $S_{k_{i}}$ loses at least $\frac{1}{2}\left(k_{i}-k_{i-1}\right)$ degrees on average. As each vertex will be removed from $S_{k_{i}}$ after it loses at most $\Delta-k_{i}+1$ tree degree, the total number of such vertex $u$ can be at most:

$$
\frac{2\left(\Delta-k_{i}+1\right)\left|S_{k_{i}}^{\prime}\right|}{k_{i}-k_{i-1}}<\frac{2\left(\frac{1}{4} \Delta\left(1-\frac{1}{\log n}\right)^{i}+2\right)\left|S_{k_{i}}^{\prime}\right|}{\frac{1}{4} \Delta\left(\left(1-\frac{1}{\log n}\right)^{i-1}-\left(1-\frac{1}{\log n}\right)^{i}\right)-1}<3\left|S_{k_{i}}^{\prime}\right| \log n
$$

Overall, $|M|<2\left|S_{k_{i}}^{\prime}\right|+3\left|S_{k_{i}}^{\prime}\right| \log n=(2+3 \log n)\left|S_{k_{i}}^{\prime}\right|$.
Summing up (1) and (2), we have

$$
\Delta^{*} \geq \frac{l-1}{|M|} \geq \frac{\frac{3}{8} \Delta\left|S_{k_{i}}^{\prime}\right|}{(2+3 \log n)\left|S_{k_{i}}^{\prime}\right|}=\Omega(\Delta / \log n)
$$

or equivalently, $\Delta=O\left(\Delta^{*} \log n\right)$.

## $4 \quad \mathbf{A}(1+\delta) \Delta^{*}+O\left(\frac{1}{\delta^{2}} \log n\right)$ Approximation

In this section we prove Theorem 2. To obtain an improved approximation of $(1+\delta) \Delta^{*}+O\left(\frac{\log n}{\delta^{2}}\right)$, the rough idea is that we refine the fast degree reduction algorithm in the previous section using an augmenting path technique.

Let $\epsilon \in\left(0, \frac{1}{48}\right)$ be a fixed parameter. The basic framework stays the same as in the previous section. One difference is that the new main algorithm consists of two phases. In the large-step phase, as long as $\Delta \geq \frac{10 \log ^{2} n}{\epsilon^{3}}$, we repeatedly apply a near-linear time algorithm LargeStepScheduling that either reduces $\Delta$ to $\leq(1-\epsilon) \cdot \Delta$ or terminates with $\Delta \leq(1+O(\epsilon)) \Delta^{*}$. In the small-step phase we need to deal with the situation where $\frac{20 \log n}{\epsilon^{2}} \leq \Delta<\frac{10 \log ^{2} n}{\epsilon^{3}}$; in this case we repeatedly run a weaker near-linear time algorithm SmallStepScheduling that either reduces $\Delta$ by 1 or provides evidence that $\Delta \leq(1+O(\epsilon)) \Delta^{*}+O\left(\frac{\log n}{\epsilon^{2}}\right)$.

Both algorithms LargeStepScheduling and SmallStepScheduling rely on a building block algorithm AugPathDegRed; similar to Scheduling, both scheduling algorithms run a while-loop and repeatedly feed inputs to AugPathDegRed. Algorithm AugPathDegRed efficiently reduces the total number of vertices of high tree degree using an augmenting path technique, which is a significant improvement over FastDegreeReduction.

For the rest of this section, we first propose and analyse the building block algorithm AugPathDegRed which underlies the core of our main algorithm. After that we specify how the two phases large-step phase and small-step phase work. Finally, we prove Theorem 2.

### 4.1 Degree reduction via augmenting paths

### 4.1.1 Algorithm description

Let $k \leq \Delta$ be a fixed threshold. This algorithm is, in some way, an extension of the previous algorithm FastDegreeReduction. As before, due to concerns of efficiency, a vertex gets marked if its tree degree is $\geq k-1$, and it stays marked (throughout one execution of AugPathDegRed) even if its tree degree decreases afterwards. Previously, we only look for a non-tree edge whose inclusion could directly reduce some tree degrees of vertices in $S_{k}$, and when such edges no longer exist the procedure terminates. In this case, AugPathDegRed would continue to explore possibilities of improving the tree structure using the idea of augmenting paths. Intuitively, an augmenting path consists of a sequence of non-tree edges that can jointly reduce tree degrees of $S_{k}$. Formally we give its definition below.

Definition 7 (augmenting paths). An h-length augmenting path consists of a sequence of distinct non-tree edges $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h}, z_{h}\right) \in E$ with the following properties.
(i) $\exists w_{0} \in \rho_{w_{1}, z_{1}} \cap S_{k}, w_{i} \in \rho_{w_{i+1}, z_{i+1}} \backslash\left(\bigcup_{j=i+2}^{h} \rho_{w_{j}, z_{j}}\right), \forall 0 \leq i<h$.
(ii) All $z_{i}$ 's are unmarked, $\forall 1 \leq i \leq h$; $w_{i}$ 's are marked for $1 \leq i<h$ and $w_{h}$ is unmarked.

Lemma 8 (tree modification). Given an augmenting path $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h}, z_{h}\right) \in E$, one can modify $T$ such that $d_{k}$ decreases and no vertices are added to $S_{k}$.
Proof. We modify T in an inductive way. For $i=h-1, h-2, \cdots, 0$, as $w_{i} \in \rho_{w_{i+1}, z_{i+1}}$, we can take an arbitrary tree edge $\left(w_{i}, x\right) \in \rho_{w_{i+1}, z_{i+1}}$, and then perform an update $\mathrm{T} \leftarrow \mathrm{T} \cup\left\{\left(w_{i+1}, z_{i+1}\right)\right\} \backslash\left\{\left(w_{i}, x\right)\right\}$ which guarantees that T is still a spanning tree. Note that this update also preserves the property that $w_{j} \in \rho_{w_{j+1}, z_{j+1}} \backslash\left(\bigcup_{l=j+2}^{h} \rho_{w_{l}, z_{l}}\right), \forall 0 \leq j<i$; this is because, when $w_{j} \notin \rho_{w_{i+1}, z_{i+1}}$, tree update $\mathrm{T} \leftarrow \mathrm{T} \cup\left\{\left(w_{i+1}, z_{i+1}\right)\right\} \backslash\left\{\left(w_{i}, x\right)\right\}$ does not change the connected components of $\mathrm{T} \backslash\left\{w_{j}\right\}$, and thus the condition $w_{j} \in \rho_{w_{j+1}, z_{j+1}} \backslash\left(\bigcup_{l=j+2}^{h} \rho_{w_{l}, z_{l}}\right)$ stays intact.

During the process, if any $\operatorname{deg}\left(z_{i}\right), 1 \leq i \leq h$ becomes $k-1$ during the process, mark $z_{i}$. By definition, $d_{k}$ decreases as $w_{0}$ loses a tree neighbour; plus, because all $\operatorname{deg}\left(w_{i}\right), 1 \leq i<h$ are unchanged, and no vertices are newly added to $S_{k}$ because $\operatorname{deg}\left(z_{i}\right) \leq k-2, \forall 1 \leq i \leq h$.

It is easy to notice that what FastDegreeRedection does is repeatedly looking for augmenting paths of length 1 and then apply Lemma 8. To extend this algorithm, when we can no longer find any augmenting paths of length 1 , we turn to search for augmenting paths of length 2 , and so on. Generally speaking, when the currently shortest augmenting paths have length $h$, we apply Lemma 8 to decrease the total number of shortest augmenting paths, and when no further progress of such kind can be made we argue the shortest length of augmenting paths must now increase. Finally our algorithm terminates when $h$ grows to $\geq 1+\log _{1+\epsilon} n$, and then we prove a lower bound on $\Delta^{*}$ based on the structure of T .

The algorithm for finding the shortest length of augmenting paths works as follows; actually the algorithm computes an auxiliary layering of the graph that will also help tree modification later. Initially we set $B_{0} \leftarrow S_{k}$. Inductively, suppose we have already computed $B_{0}, B_{1}, \cdots, B_{h}, h \geq 0$, then we compute the forest spanned by $\mathrm{T} \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$. Here is an extra notation: $\forall 0 \leq i \leq h$, for each $u \in V \backslash\left(\bigcup_{j=0}^{i} B_{j}\right)$, let $C_{u}^{i}$ be the connected component of $\mathrm{T} \backslash\left(\bigcup_{j=0}^{i} B_{j}\right)$ that contains $u$. If there exists an edge $(u, v) \in E$ such that both $u, v$ are unmarked vertices, and that $C_{u}^{h} \neq C_{v}^{h}$, then the algorithm terminates and reports that the shortest length of augmenting paths is equal to $h+1$; otherwise, we compute $B_{h+1}$ to be the set of all marked vertices $u \in V \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ such that there exists an unmarked adjacent vertex $v$ with $C_{u}^{h} \neq C_{v}^{h}$, and then continue. The above procedure is summarised as the following pseudo-code Layering shown in Algorithm 3. Note that once $B_{h+1}=\emptyset$, the algorithm would continue to compute $B_{h+2}=B_{h+3}=\cdots=B_{1+\log _{1+\epsilon} n}=\emptyset$.

After we have invoked Layering and computed a sequence of vertex subsets $B_{0}, B_{1}, \cdots, B_{h}$ which naturally divides the graph into $h+2$ layers, we should start to apply tree modifications of Lemma 8 to decrease the total number of shortest augmenting paths. The difficulty in searching for shortest augmenting paths is that, for a search that starts from a pair of adjacent and unmarked vertices $u, v$ satisfying $C_{u}^{h} \neq C_{v}^{h}$ and goes up the layers $B_{h}, B_{h-1}, \cdots, B_{1}, B_{0}$, not every route can reach the top layer $B_{0}$ because the augmentations of some previous $(h+1)$-length augmenting paths might have already blocked the road. Therefore, a depth-first search needs to performed. To save running time, some tricks are needed: if a certain vertex has been searched and failed to lead a way upwards to $B_{0}$, then we tag this vertex so that future depth-first searches may avoid this tagged vertex; if a certain edge has been searched before, then we tag this edge whatsoever. The following pseudocode AugDFS shown in Algorithm 4 may be a better illustration of this algorithm; the recursive

```
Algorithm 3: Layering
    \(B_{0} \leftarrow S_{k} ;\)
    \(h \leftarrow 0 ;\)
    while \(h \leq 1+\log _{1+\epsilon} n\) do
        compute the forest \(\left\{C_{u}^{h}\right\}\) spanned by \(\mathrm{T} \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)\);
        if exists unmarked \(u, v\) such that \((u, v) \in E, C_{u}^{h} \neq C_{v}^{h}\) then
            break;
        else
            compute \(B_{h+1}\) to be the set of all marked vertices in \(u \in V \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)\) such that
            there exists an unmarked adjacent vertex \(v\) with \(C_{u}^{h} \neq C_{v}^{h}\);
            \(h \leftarrow h+1 ;\)
    return \(B_{0}, B_{1}, \cdots, B_{h}\);
```

algorithm AugDFS searches for an $(h+1)$-length augmenting path $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h+1}, z_{h+1}\right)$ given input $\left(w_{h+1}, z_{h+1}\right) \in E$. Later we will prove, if AugDFS returns true, then the sequence $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h+1}, z_{h+1}\right)$ is indeed an augmenting path.

Now we come to describe the upper-level AugPathDegRed: basically, it repeatedly apply Layering followed by several rounds of AugDFS until $h \geq 1+\log _{1+\epsilon} n$. Here is the pseudo-code AugPathDegRed as shown in Algorithm [5

Before proving termination of AugPathDegRed, we first need to argue some properties of Layering. The following lemma will serve as the basis for our future lower bounds on $\Delta^{*}$.

Lemma 9 (the blocking property). Throughout each iteration of the repeat-loop in AugPathDegRed, for any $1 \leq i<h$ and any two adjacent vertices $u, v \in V \backslash\left(\bigcup_{j=0}^{i} B_{j}\right)$ such that $u$ is unmarked and $C_{u}^{i} \neq C_{v}^{i}$, then $v \in B_{i+1}$.

Proof. By rules of Layering, this blocking property holds right after Layering outputs them. This claim continuous to hold afterwards because tree modifications only merge components $C_{u}^{i}$ 's and never splits any $C_{u}^{i}$ 's.

Here is an important corollary of this Lemma 9 .
Corollary 10. Throughout each iteration of the repeat-loop, for any $w \in B_{i}, 1 \leq i \leq h$, suppose $w$ is adjacent to an unmarked $z$ such that $C_{w}^{i-1} \neq C_{z}^{i-1}$. Then $\rho_{w, z}$ only contains vertices from $V \backslash\left(\bigcup_{j=0}^{i-2} B_{j}\right)$.
Proof. Suppose otherwise, then there would be a vertex $x \in \rho_{w, z} \cap B_{j}, j \leq i-2$, then in this case $C_{w}^{j} \neq C_{z}^{j}$, and thus by Lemma $9 w \in B_{j+1}$ which is a contradiction as $j+1<i$.

Now we can argue correctness of AugDFS.
Lemma 11. If AugDFS returns true, $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h+1}, z_{h+1}\right)$ is an augmenting path.
Proof. Property (ii) in Definition 7 holds by rules of this algorithm. Now let us focus on property (i). Take an arbitrary $w_{0} \in \rho_{w_{1}, z_{1}} \cap B_{0}$. By the algorithm it must be $w_{i} \in B_{i}, \forall 0 \leq i \leq h$, then using Corollary 10 we know $\rho_{w_{i}, z_{i}}$ does not contain any $w_{j}, 0 \leq j \leq i-2$, so property (i) holds.

```
Algorithm 4: AugDFS
    1 global variables: \(\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h+1}, z_{h+1}\right)\);
    input params: \(i, 1 \leq i \leq h+1\);
    if \(i=1\) then
        return true;
    for untagged \(w \in \rho_{w_{i}, z_{i}} \cap B_{i-1}\) do
        \(w_{i-1} \leftarrow w ;\)
        for unmarked \(z\) such that \((w, z)\) is untagged and \(C_{z}^{i-2} \neq C_{w}^{i-2}\) do
            \(z_{i-1} \leftarrow z ;\)
            run AugDFS with input \(i-1\);
            \(\operatorname{tag}(w, z)\);
            if AugDFS has returned false then
                    continue;
            else
                if \(i=h+1\) then
                    modify T by augmenting path \(\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{h+1}, z_{h+1}\right)\) via Lemma 8
                        and mark new degree \(k-1\) vertices;
                return true;
        tag \(w\);
    return false;
```

```
Algorithm 5: AugPathDegRed
    input params: threshold \(k\);
    mark all vertices from \(S_{k-1}\) and unmark the rest;
    repeat
        run Layering which computes \(B_{0}, B_{1}, \cdots, B_{h}\);
        untag all vertices and edges;
        for \((u, v) \in E\) such that \(u, v\) are unmarked and adjacent, and that \(C_{u}^{h} \neq C_{v}^{h}\) do
            run AugDFS with input \(h+1\) and global params \(\left(w_{h+1}, z_{h+1}\right) \leftarrow(u, v)\);
    until \(h>1+\log _{1+\epsilon} n\);
    return T ;
```

Finally we conclude this subsection with the lemma below, from which termination of AugPathDegRed immediately follows.

Lemma 12. Every iteration of the repeat-loop, if not the last, increases $h$ by at least one.
Proof. By the rules of Layering, it is easy to see that at the beginning when Layering outputs $B_{0}, B_{1}, \cdots, B_{h}$, the shortest length of augmenting path is equal to $h+1$. So it suffices to prove that by the end of this iteration the shortest augmenting path has length $>h+1$.

First we need to characterize all augmenting paths using $B_{0}, B_{1}, \cdots, B_{h}$. Let the sequence $\left(w_{1}, z_{1}\right),\left(w_{2}, z_{2}\right), \cdots,\left(w_{l}, z_{l}\right)$ be an arbitrary augmenting path. We argue $l \geq h+1$, and more importantly, if $l=h+1$, it must be $w_{i} \in B_{i}, \forall 1 \leq i \leq h$. We inductively prove that $w_{i} \in \bigcup_{j=0}^{i} B_{j}$ for $i=0,1, \cdots, l-1$. The basis is obvious as is required by property (i) in Definition 7 Now assume $w_{i} \in B_{r}$ for some $r \leq i$. Then, by Corollary 10, it would not be hard to see $w_{i+1} \in$ $\bigcup_{j=0}^{r+1} B_{j} \subseteq \bigcup_{j=0}^{i+1} B_{j}$. Now, on the one hand by Corollary $10 \rho_{w_{l}, z_{l}} \cap \bigcup_{j=0}^{h-1} B_{j}=\emptyset$, and on the other hand $w_{l-1} \in \rho_{w_{l}, z_{l}} \cap \bigcup_{j=0}^{l-1} B_{j}$, so $l \geq h+1$. Plus, we can see from the induction that, when $l=h+1$ it must be $w_{i} \in B_{i}, \forall 0 \leq i \leq h$.

For any unmarked and adjacent vertices $u, v$ such that $C_{u}^{h} \neq C_{v}^{h}$, consider the instance of AugDFS with input $\left(w_{h+1}, z_{h+1}\right) \leftarrow(u, v)$. We make two claims.
(1) If there is an $(h+1)$-length augmenting path ending with $\left(w_{h+1}, z_{h+1}\right)=(u, v)$, AugDFS would succeed in finding one.
(2) If it has returned false, then there would be no $(h+1)$-augmenting path ending with $\left(w_{h+1}, z_{h+1}\right)=$ $(u, v)$ throughout the entire repeat-loop iteration.

If (1)(2) can be proved, then by the end of this repeat-loop iteration, there would be no $(h+1)$ length augmenting paths because at such augmenting path should end with a pair of adjacent unmarked vertices. Next we come to prove (1)(2).
(1) The depth-first search of AugDFS exactly coincides with the conditions that $w_{i} \in B_{i}$, except that it skips all tagged vertices and edges. Now we prove that omitting tagged vertices and edges does not miss any $(h+1)$-length augmenting paths. For a vertex $w$ to be tagged, we must have enumerated all of its untagged edges $(w, z)$ but failed to find any augmenting paths, and therefore any future depth-first searches on $w$ would still end up in vain. For an edge ( $w, z$ ) to be tagged, either a further recursion AugDFS on line-9 has succeeded or failed in finding an augmenting paths; in the former case, $C_{w}^{i-2}$ and $C_{z}^{i-2}$ has been merged, and so the condition $C_{w}^{i-2} \neq C_{z}^{i-2}$ would be violated afterwards; in the latter case, we would not need to recur on $(w, z)$ anyway.
(2) If AugDFS has once failed to find any augmenting paths starting with $(u, v)$, then all vertices $w \in \rho_{u, v} \cap B_{h}$ visited by this instance of AugDFS should be tagged and they would be omitted by all succeeding instances of AugDFS. Therefore $\rho_{u, v} \cap B_{h}$ would stay unchanged since then (although $\rho_{u, v}$ itself might change). Hence, image if we re-run AugDFS with $\left(w_{h+1}, z_{h+1}\right) \leftarrow$ $(u, v)$, it may return false without any recursion because all vertices in $\rho_{w_{h+1}, z_{h+1}} \cap B_{h}$ are tagged.

### 4.1.2 Lower bound on $\Delta^{*}$

Suppose AugPathDegRed has terminated with $B_{0}, B_{1}, \cdots, B_{\log _{1+\epsilon} n+1}$. Let us see it yields lower bounds on $\Delta^{*}$. To apply Lemma 5, we first need to specify a sequence of disjoint vertex subsets, which is what the following definition is about.

Definition 13. After an instance of AugPathDegRed has been executed, for an arbitrary component $C_{u}^{h}, 0 \leq h \leq \log _{1+\epsilon} n, u \in V \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$, it is called clean if all vertices in $C_{u}^{h}$ are unmarked.

Lemma 14. For any $0 \leq h \leq \log _{1+\epsilon} n$, suppose $T \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ has $l$ clean components, then a lower bound holds that $\Delta^{*} \geq \frac{-1-1}{\sum_{i=0}^{h+1}\left|B_{i}\right|}$.

Proof. By Lemma 9, any edge that connects a clean components of $\mathrm{T} \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ outwards must be incident on a vertex in $\bigcup_{i=0}^{h+1} B_{i}$. Therefore by Lemma 5 we have $\Delta^{*} \geq \frac{l-1}{\left|\bigcup_{i=0}^{h+1} B_{i}\right|}=\frac{l-1}{\sum_{i=0}^{h+1}\left|B_{i}\right|}$

From Lemma 14, it suffices to lower bound the total number of clean components. The next lemma describes a scenario in which $l$ must be large.

Lemma 15. Suppose an instance of AugPathDegRed has been executed. Let $d_{k}^{\prime}, d_{k-1}^{\prime}, S_{k-1}^{\prime}$ and $S_{k}^{\prime}$ be snapshots of $d_{k}, d_{k-1}, S_{k-1}$ and $S_{k}$ right before this instance of AugPathDegRed started; recall that $d_{k}, d_{k-1}, S_{k-1}$ and $S_{k}$ always refer to statistics of the current $T$ after this instance of AugPathDegRed has finished.

Assume the following three conditions:
(i) $k \geq \frac{4}{\epsilon}$
(ii) $d_{k}^{\prime} \geq \frac{1}{\epsilon(k-1)} \cdot d_{k-1}^{\prime}$
(iii) $d_{k} \geq\left(1-\frac{\epsilon^{2}}{2 \log n}\right) d_{k}^{\prime}$

Then, for each $0 \leq h \leq \log _{1+\epsilon} n$, the number of clean components in $T \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ is more than $k \cdot(1-4 \epsilon) \sum_{i=0}^{h}\left|B_{i}\right|$.

Proof. By Lemma 6, the number of tree components in $\mathrm{T} \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ is at least

$$
\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)-2\left|\bigcup_{i=0}^{h} B_{i}\right|+2
$$

Let $M$ be the set of all marked vertices $\notin S_{k-1}^{\prime}$ (i.e., vertices that are initially unmarked) by the end of AugPathDegRed. Then, the number of clean components in $\mathrm{T} \backslash\left(\bigcup_{i=0}^{h} B_{i}\right)$ is at least

$$
\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)-2 \sum_{i=0}^{h}\left|B_{i}\right|+2-\left|M \cup S_{k-1}^{\prime}\right|
$$

The argument consists of a lower bound on $\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)$ and an upper bound on $\left|M \cup S_{k-1}^{\prime}\right|$.
(1) Lower bound on $\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)$.

By the algorithm $B_{0}=S_{k}$, then we have $\sum_{u \in B_{0}} \operatorname{deg}(u)=d_{k}$.
For any vertex $u \in \bigcup_{i=1}^{h} B_{i}, \operatorname{deg}(u)=k-1$ by the time $u$ was first added to some $B_{i}$. After that, $\operatorname{deg}(u)$ could only decrease when we modify T by an augmenting path $\left(w_{1}, z_{1}\right), \cdots,\left(w_{t}, z_{t}\right)$ where $u=w_{j}$ for some $1 \leq j \leq t$. Since $t \leq 1+\log _{1+\epsilon} n$, during a tree modification, at least one vertex in $S_{k}$ loses one degree and at most $1+\log _{1+\epsilon} n$ vertices in $\bigcup_{i=1}^{h} B_{i}$ lose one degree separately. As the total number of the degree loss in $S_{k}$ is $\left(d_{k}^{\prime}-d_{k}\right)$, we have

$$
\sum_{u \in \bigcup_{i=1}^{h} B_{i}} \operatorname{deg}(u) \geq(k-1) \sum_{i=1}^{h}\left|B_{i}\right|-\left(d_{k}^{\prime}-d_{k}\right)\left(1+\log _{1+\epsilon} n\right)
$$

From above, we get a lower bound on $\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)$,

$$
\begin{aligned}
\sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u) & \geq d_{k}+(k-1) \sum_{i=1}^{h}\left|B_{i}\right|-\left(d_{k}^{\prime}-d_{k}\right)\left(1+\log _{1+\epsilon} n\right) \\
& \geq(k-1) \sum_{i=1}^{h}\left|B_{i}\right|+\left(1-\frac{\epsilon^{2}}{2 \log n}\right) d_{k}^{\prime}-\frac{\epsilon^{2}}{2 \log n} d_{k}^{\prime} \cdot\left(1+\log _{1+\epsilon} n\right) \\
& \geq(k-1) \sum_{i=1}^{h}\left|B_{i}\right|+\left(1-\frac{\epsilon^{2}}{2 \log n}-\epsilon\right) d_{k}^{\prime}
\end{aligned}
$$

(2) Upper bound on $|M|$.

The argument is similar to (1). An unmarked vertex $u$ is marked only when we modify T by an augmenting path $\left(w_{1}, z_{1}\right), \cdots,\left(w_{t}, z_{t}\right)$ where $u=z_{j}$ for some $1 \leq j \leq t$ or $u=w_{t}$. Since $t \leq 1+\log _{1+\epsilon} n$, during a tree modification, at least one vertex in $S_{k}$ loses one degree and at most $2+\log _{1+\epsilon} n$ unmarked vertices are marked. Then we get a upper bound on $|M|$.

$$
|M| \leq\left(d_{k}^{\prime}-d_{k}\right)\left(2+\log _{1+\epsilon} n\right) \leq \epsilon \cdot d_{k}^{\prime}
$$

(3) Upper bound on $\left|S_{k-1}^{\prime}\right|$.

By easy calculations, $d_{k}^{\prime} \geq \frac{1}{\epsilon(k-1)} d_{k-1}^{\prime} \geq \frac{1}{\epsilon}\left|S_{k-1}^{\prime}\right|$, and so $\left|S_{k-1}^{\prime}\right| \leq \epsilon \cdot d_{k}^{\prime}$.

Summing up (1)(2)(3), we can conclude the proof:

$$
\begin{aligned}
& \sum_{u \in \bigcup_{i=0}^{h} B_{i}} \operatorname{deg}(u)-2 \sum_{i=0}^{h}\left|B_{i}\right|+2-\left|M \cup S_{k-1}^{\prime}\right|>d_{k}+(k-1) \sum_{i=1}^{h}\left|B_{i}\right|-2 \sum_{i=0}^{h}\left|B_{i}\right|-3 \epsilon \cdot d_{k}^{\prime} \\
& \geq\left(1-\frac{\epsilon^{2}}{2 \log n}\right) \cdot d_{k}^{\prime}-2\left|B_{0}\right|+(k-3) \sum_{i=1}^{h}\left|B_{i}\right|-3 \epsilon \cdot d_{k}^{\prime} \\
& \geq\left(1-\frac{\epsilon^{2}}{2 \log n}-3 \epsilon\right) \cdot d_{k}^{\prime}-2\left|B_{0}\right|+(k-3) \sum_{i=1}^{h}\left|B_{i}\right| \\
& >k(1-3.5 \epsilon) \cdot\left|B_{0}\right|-2\left|B_{0}\right|+(k-3) \sum_{i=1}^{h}\left|B_{i}\right| \geq k(1-4 \epsilon) \cdot \sum_{i=0}^{h}\left|B_{i}\right|
\end{aligned}
$$

The last inequality holds by $k \geq \frac{4}{\epsilon}$.

### 4.1.3 Implementation and running time

We present an implementation of AugPathDegRed that runs in $O\left(\frac{1}{\epsilon^{2}} m \log ^{2} n\right)$ time. We discuss some implementation details of Layering, AugDFS and AugPathDegRed, and analyse their contributions to the total running time in a single run of AugPathDegRed.
(1) Layering.

For every instance of Layering, computing the forest $\left\{C_{u}^{h}\right\}_{u \in V \backslash\left(\cup_{i=0}^{h} B_{i}\right)}$ can be done in a single pass of breath-first search which takes $O(m)$ time. Computing $B_{h+1}$, if necessary, is easily done by scanning the edge set $E$ which also takes $O(m)$ time. As the while-loop iterates for at most $1+\log _{1+\epsilon} n$ times, and due to Lemma 12 Layering is invoked for at most $1+\log _{1+\epsilon} n$ times, the overall contribution of Layering is $O\left(\frac{1}{\epsilon^{2}} m \log ^{2} n\right)$.
(2) AugPathDegRed.

Excluding the contributions of AugDFS and Layering, all AugPathDegRed does is simply untagging all vertices and edges, as well as scanning the edge set $(u, v) \in E$ and deciding if $C_{u}^{h} \neq C_{v}^{h}$. As tree components only get merged and never split, we can use the union-find data structure to support querying whether $C_{u}^{h} \neq C_{v}^{h}$ in $O(\alpha(n))$ time. Hence, AugPathDegRed's exclusive contributions to the total running time would be $O\left(\frac{1}{\epsilon} m \alpha(n) \log n\right)$.
(3) AugDFS.

Now we analyse the overall time complexity induced by AugDFS invoked on line-7 of AugPathDegRed. There are three technical issues to be resolved.
(a) How to enumerate untagged vertices $\in \rho_{w_{i}, z_{i}} \cap B_{i-1}$ (line-5)?

For each $u \in B_{i}, \forall 0 \leq i \leq h$, assign $u$ a weight of $i$; vertices that do not belong to any $B_{i}$ have weight $h+1$. By Corollary 10, to enumerate vertices $\in \rho_{w_{i}, z_{i}} \cap B_{i-1}$, it suffices to enumerate the lightest vertices on $\rho_{w_{i}, z_{i}}$, which can be done using a link-cut tree data structure [16] built on T , each enumeration taking $O(\log n)$ amortized time. When a vertex gets tagged, we change its weight to $h+1$, and so future enumerations on $\rho_{w_{i}, z_{i}} \cap B_{i-1}$ may skip this tagged vertex.
(b) How to enumerate unmarked $z$ connected by an untagged edge $(w, z)$ such that $C_{z}^{i-2} \neq C_{w}^{i-2}$ (line-7)?
Each $w$ decrementally maintains a list of all its neighbours. While we scan the list, if the next edge $(w, z)$ satisfies both conditions that $C_{z}^{i-2} \neq C_{w}^{i-2}$ and $z$ is unmarked, then the algorithm starts a new iteration and recur; either way we cross the edge $(w, z)$ off the list. In this way, every edge appears on line- 7 for at most once. Thus the total time of this part is $O(m \alpha(n))$; the additional $\alpha(n)$ factor comes from the union-find data structure that helps deciding if $C_{z}^{i-2} \neq C_{w}^{i-2}$.
(c) How to implement the tree modification from Lemma 8?

Every tree modification involves insertions and deletions of $O\left(\frac{1}{\epsilon} \log n\right)$ edges, as well as merging $O\left(\frac{1}{\epsilon} \log n\right)$ pairs of some tree components $C_{u}^{i}$. Using the link-cut tree, every edge insertion and deletion takes update time $O(\log n)$, and every component-merging takes time $O(\alpha(n))$. Since every tree modification merges two components in $\mathrm{T} \backslash S_{k}$ (i.e., $\mathrm{T} \backslash B_{0}$ ), there can be at most $O(n)$ tree modifications throughout AugPathDegRed. Therefore, the overall contribution of tree modifications is $O\left(\frac{1}{\epsilon} n \log ^{2} n\right)$.

The total time of $(\mathrm{a})(\mathrm{b})$ is $O(m \log n)$ because every time a vertex is enumerated on line- 5 , either it gets tagged or one of its edge gets tagged. Thus, the overall complexity of AugDFS is $O\left(m \log n+\frac{1}{\epsilon} n \log ^{2} n\right)$.

Summing up $(1)(2)(3)$, the total running time of is dominated by time complexity of Layering which is $O\left(\frac{1}{\epsilon^{2}} m \log ^{2} n\right)$.

### 4.2 Large-step phase

### 4.2.1 Algorithm description and running time

Algorithm largeStepScheduling deals with the case $\Delta \geq \frac{10 \log n}{\epsilon^{2}}$ : it iterates over $k=(1-2 \epsilon) \Delta+$ $1,(1-2 \epsilon) \Delta+2, \cdots,(1-\epsilon) \Delta$ and if $d_{k-1} \leq 2 d_{k}$ it invoke AugPathDegRed with input $k$. During the course, if one instance fails to reduce $d_{k}$ significantly, then the algorithm reports a lower bound on $\Delta^{*}$ and terminates. Here is its pseudo-code 6.

```
Algorithm 6: LargeStepScheduling
    input params: T with maximum tree degree \(\Delta\);
    for \(k=(1-2 \epsilon) \Delta+1,(1-2 \epsilon) \Delta+2, \cdots,(1-\epsilon) \Delta\) do
        if \(d_{k}=0\) then
            break;
        else if \(d_{k-1} \leq 2 d_{k}\) then
            \(d_{k}^{\prime} \leftarrow d_{k} ;\)
            run AugPathDegRed with input \(k\);
            if \(d_{k}>\left(1-\frac{\epsilon^{2}}{2 \log n}\right) \cdot d_{k}^{\prime}\) then
                return false;
    return true;
```

From the previous subsection we already know AugPathDegRed runs in near-linear time, so here we only need to upper bound the total number of times AugPathDegRed gets invoked before the algorithm returns either true or false. We consider in each iteration, by how much would $d_{k}$ decrease. There are two cases.
(1) If the branching condition on line-5 does not hold, then $d_{k-1}>2 d_{k}$, and so $d_{k}$ has shrunk by a factor of at most $\frac{1}{2}$.
(2) In the case where the branching condition holds, if the algorithm does not return false within this iteration, then $d_{k} \leq\left(1-\frac{\epsilon^{2}}{2 \log n}\right) \cdot d_{k}^{\prime}$; that is, $d_{k}$ has declined by a factor of at most $1-\frac{\epsilon^{2}}{2 \log n}$ during this iteration.

Summing up (1)(2), the value of $d_{k}$ would decrease by a factor of at most $1-\frac{\epsilon^{2}}{2 \log n}$ in each iteration. Then the for-loop would break when $k \geq(1-2 \epsilon) \Delta+\frac{2 \log ^{2} n}{\epsilon^{2}}$ because by the time $d_{k}$ would decrease to 0 . Hence, the running time of LargeStepScheduling is bounded by $O\left(\frac{1}{\epsilon^{3}} m \log ^{3} n\right)$.

A byproduct of the above observation is that, when $\Delta \geq \frac{10 \log ^{2} n}{\epsilon^{3}}$, we have $(1-\epsilon) \Delta \geq(1-$ $2 \epsilon) \Delta+\frac{10 \log ^{2} n}{\epsilon^{2}}$, and thus if LargeStepScheduling eventually returns true, it must be $d_{(1-\epsilon) \Delta}=0$; in other words, $\Delta$ has been reduced by a factor of at most $1-\epsilon$.

### 4.2.2 Approximation guarantee

We prove when LargeStepScheduling returns false, it must be $\Delta \leq(1+8 \epsilon) \cdot \Delta^{*}$. Consider the most recent execution of AugPathDegRed before returning. By the previous subsection, this instance of AugPathDegRed has created a sequence of disjoint vertex subsets $B_{0}, B_{1}, \cdots, B_{1+\log _{1+\epsilon} n}$ that satisfies the blocking property.

Let $d_{k-1}^{\prime}, S_{k-1}^{\prime}$ and $S_{k}^{\prime}$ be snapshots of $d_{k-1}, S_{k-1}$ and $S_{k}$ right before this instance of AugPathDegRed started. Then $d_{k}^{\prime} \geq \frac{1}{2} d_{k-1}^{\prime} \geq \frac{1}{\epsilon(k-1)} d_{k-1}^{\prime}$ by the branching condition on line- 5 . Using Lemma 15 and Lemma 14, we can derive a lower bound

$$
\Delta^{*} \geq k(1-4 \epsilon) \cdot \frac{\sum_{i=0}^{h}\left|B_{i}\right|}{\sum_{i=0}^{h+1}\left|B_{i}\right|}, \forall 0 \leq h \leq \log _{1+\epsilon} n
$$

By the pigeon-hole principle, there exists an $h$ such that $\frac{\sum_{i=0}^{h}\left|B_{i}\right|}{\sum_{i=0}^{h+1}\left|B_{i}\right|} \geq \frac{1}{1+\epsilon}$, and hence

$$
\Delta^{*} \geq k(1-4 \epsilon) \cdot \frac{1}{1+\epsilon}>\frac{1-6 \epsilon+8 \epsilon^{2}}{1+\epsilon} \Delta
$$

or equivalently, $\Delta \leq \frac{1+\epsilon}{1-6 \epsilon+8 \epsilon^{2}} \Delta^{*}<(1+8 \epsilon) \Delta^{*}$ when $\epsilon \in\left(0, \frac{1}{48}\right)$.

### 4.3 Small-step phase

### 4.3.1 Algorithm description and running time

Algorithm SmallStepScheduling only deals with $\frac{20 \log n}{\epsilon^{2}} \leq \Delta<\frac{10 \log ^{2} n}{\epsilon^{3}}$. Set $c=12+6 \log _{1+\epsilon} n$ and define a potential function:

$$
\phi(\mathrm{T})=\sum_{i=\Delta+1-\log n}^{\Delta} c^{i} \cdot\left|N_{i}\right|
$$

SmallStepScheduling works by repetitively selecting a degree $k$ that maximizes $c^{k} \cdot\left|N_{k}\right|$ and then feed input $k$ to AugPathDegRed until $\Delta$ decreases; clearly $k$ must be larger than $\Delta-\log n>\frac{19 \log n}{\epsilon^{2}}$. We formulate this procedure as pseudo-code SmallStepScheduling as shown in Algorithm 7 below.

```
Algorithm 7: SmallStepScheduling
    input params: T with maximum tree degree \(\Delta\);
    while \(\Delta\) has not changed do
        pick a \(k \in \arg \max _{i \in[\Delta+1-\log n, \Delta]}\left\{c^{i} \cdot\left|N_{i}\right|\right\}\);
        \(d_{k}^{\prime} \leftarrow d_{k}\);
        run AugPathDegRed with input \(k\);
        if \(d_{k}>\left(1-\frac{\epsilon^{2}}{2 \log n}\right) \cdot d_{k}^{\prime}\) then
            return false;
    return true;
```

Now let us bound the running time of SmallStepScheduling. We study how many rounds of AugPathDegRed could be invoked before branching conditions on line- 2 or line- 6 can be triggered. In fact, an instance of AugPathDegRed has not triggered the condition on line-6, then $d_{k}$ has decreased by a factor of at most $1-\frac{\epsilon^{2}}{2 \log n}$. Let us analyse how $\phi(\mathrm{T})$ has decreased.

Let $N_{k}^{\prime}, k \in[\Delta+1-\log n, \Delta]$ be snapshots of $N_{k}$ right before we execute AugPathDegRed, and denote the potential of T at that time by $\phi^{\prime}(\mathrm{T})=\sum_{i=\Delta+1-\log n}^{\Delta} c^{i} \cdot\left|N_{k}^{\prime}\right|$. Every time a tree modification to T was made on line-15 of AugDFS, at least one vertex in $S_{k}$ would lose a tree edge and at most $2+\log _{1+\epsilon} n$ vertices would gain a tree edge, and then the total loss of $\phi(T)$ would be at least

$$
\begin{aligned}
& \left(c^{k}-c^{k-1}\right)-\left(2+\log _{1+\epsilon} n\right) \cdot\left(c^{k-1}-c^{k-2}\right) \geq\left(c^{k-1}-c^{k-2}\right)\left(c-2-\log _{1+\epsilon} n\right) \\
& =c^{k} \cdot\left(1-\frac{1}{c}\right) \cdot\left(1-\frac{2+\log _{1+\epsilon} n}{c}\right) \geq c^{k} \cdot\left(1-\frac{1}{c}\right) \cdot \frac{5}{6}>0.8 \cdot c^{k}
\end{aligned}
$$

Note that $d_{k}$ has decreased by $d_{k}^{\prime}-d_{k} \geq \frac{\epsilon^{2}}{2 \log n} d_{k}^{\prime} \geq \frac{\epsilon^{2}}{2 \log n} \cdot k\left|N_{k}^{\prime}\right|$, and so there are at least $\left(d_{k}^{\prime}-d_{k}\right) /(2 k) \geq \frac{\epsilon^{2}}{4 \log n} \cdot\left|N_{k}^{\prime}\right|$ such modifications to T , this is because a tree modification can at most move two vertices from $S_{k}$ to $S_{k-1}$ thus decreasing $d_{k}$ by $2 k$. Therefore,

$$
\phi(\mathrm{T}) \leq \phi^{\prime}(\mathrm{T})-\left(0.8 \cdot c^{k}\right) \cdot\left(\frac{\epsilon^{2}}{4 \log n} \cdot\left|N_{k}^{\prime}\right|\right) \leq\left(1-\frac{0.2 \epsilon^{2}}{\log ^{2} n}\right) \phi^{\prime}(\mathrm{T})
$$

The second inequality holds by maximality of $c^{k} \cdot\left|N_{k}^{\prime}\right|$ which implies $c^{k} \cdot\left|N_{k}^{\prime}\right| \geq \frac{1}{\log n} \cdot \phi^{\prime}(\mathbf{T})$
In a nutshell, whenever AugPathDegRed returns true, $\phi(\mathrm{T})$ has decreases by a factor of at most $1-\frac{0.2 \epsilon^{2}}{\log ^{2} n}$. As long as $\Delta$ has not changed, $\phi(\mathrm{T})$ belongs to the interval $\left(c^{\Delta-1}, n \cdot c^{\Delta}\right)$, and consequently, $\phi(\mathrm{T})$ could suffer at most $-\log _{1-\frac{0.2 \epsilon^{2}}{\log ^{2} n}} n \cdot c=O\left(\frac{\log ^{3} n}{\epsilon^{2}}\right)$ rounds of AugPathDegRed that returns true. Therefore, the overall running time of SmallStepScheduling before $\Delta$ decreases by 1 would be upper bounded by $O\left(\frac{1}{\epsilon^{3}} m \log ^{5} n\right)$.

### 4.3.2 Approximation guarantee

We prove when SmallStepScheduling returns false, it must be $\Delta \leq(1+8 \epsilon) \Delta^{*}+\log n$. Consider the most recent execution of AugPathDegRed before returning. By the previous subsection, this instance of AugPathDegRed has created a sequence of disjoint vertex subsets $B_{0}, B_{1}, \cdots, B_{1+\log _{1+\epsilon} n}$ that satisfies the blocking property.

Let $d_{k-1}^{\prime}, S_{k-1}^{\prime}$ and $S_{k}^{\prime}$ be snapshots of $d_{k-1}, S_{k-1}$ and $S_{k}$ right before this instance of AugPathDegRed started. By maximality of $c^{k} \cdot\left|N_{k}^{\prime}\right|$, we have $\left|N_{k}^{\prime}\right| \geq \frac{1}{c} \cdot\left|N_{k-1}^{\prime}\right|$. Then,

$$
\frac{d_{k}^{\prime}}{d_{k-1}^{\prime}}=\frac{\sum_{i=k}^{\Delta} i\left|N_{i}^{\prime}\right|}{\sum_{i=k-1}^{\Delta} i\left|N_{i}^{\prime}\right|}>\frac{k\left|N_{k}^{\prime}\right|}{k\left|N_{k}^{\prime}\right|+(k-1)\left|N_{k-1}^{\prime}\right|} \geq \frac{1}{1+\frac{c(k-1)}{k}}>\frac{1}{c+1} \geq \frac{1}{\epsilon(k-1)}
$$

The last inequality holds by $k \geq \frac{19 \log n}{\epsilon^{2}}$ and $c=6 \log _{1+\epsilon} n+12 \leq 18 \log _{1+\epsilon} n$.
As $d_{k} \geq\left(1-\frac{\epsilon^{2}}{2 \log n}\right) d_{k}^{\prime}$ by the branching condition on line- 6 , using Lemma 15 and Lemma 14, we can derive a lower bound

$$
\Delta^{*} \geq k(1-4 \epsilon) \cdot \frac{\sum_{i=0}^{h}\left|B_{i}\right|}{\sum_{i=0}^{h+1}\left|B_{i}\right|}, \forall 0 \leq h \leq \log _{1+\epsilon} n
$$

By the pigeon-hole principle, there exists an $h$ such that $\frac{\sum_{i=0}^{h}\left|B_{i}\right|}{\sum_{i=0}^{h+1}\left|B_{i}\right|} \geq \frac{1}{1+\epsilon}$, and hence

$$
\Delta^{*} \geq k(1-4 \epsilon) \cdot \frac{1}{1+\epsilon}>\frac{1-4 \epsilon}{1+\epsilon}(\Delta-\log n)
$$

or equivalently, $\Delta \leq \frac{1+\epsilon}{1-4 \epsilon} \Delta^{*}+\log n<(1+8 \epsilon) \Delta^{*}+\log n$ for $\epsilon \in\left(0, \frac{1}{48}\right)$.

### 4.4 Correctness of the main algorithm

As a conclusion we prove Theorem 2. If the main algorithm terminates in the large-step phase, then in previous subsections, we have proved $\Delta \leq(1+8 \epsilon) \cdot \Delta^{*}$; if it returns in the small-step phase , then either $\Delta \leq(1+8 \epsilon) \Delta^{*}+\log n$ or $\Delta<\frac{10 \log n}{\epsilon^{2}}$. So, in sum $\Delta \leq(1+8 \epsilon) \Delta^{*}+\frac{10 \log n}{\epsilon^{2}}$. Reassigning $\delta \leftarrow 8 \epsilon$ we can obtain the approximation guarantee promised in Theorem 2,

For the running time, the large-step phase can invoke at most $O\left(\frac{1}{\epsilon} \log n\right)$ rounds of LargeStepScheduling because each invocation that returns true reduces $\Delta$ by a factor of at most $1-\epsilon$; thus the total running time of the first while-loop never exceeds $O\left(\frac{1}{\epsilon^{4}} m \log ^{5} n\right)$. The small-step phase can invoke at most $O\left(\frac{\log ^{2} n}{\epsilon^{3}}\right)$ rounds of SmallStepScheduling because each invocation that returns true reduces $\Delta$ by at least 1 ; thus the total running time of the small-step phase is $O\left(\frac{1}{\epsilon^{6}} m \log ^{7} n\right)$ which is dominant. So the running time of the main algorithm is $O\left(\frac{1}{\epsilon^{6}} m \log ^{7} n\right)$.

## References

[1] Nikhil Bansal, Rohit Khandekar, and Viswanath Nagarajan. Additive guarantees for degreebounded directed network design. SIAM Journal on Computing, 39(4):1413-1431, 2009.
[2] Kamalika Chaudhuri, Satish Rao, Samantha Riesenfeld, and Kunal Talwar. What would edmonds do? augmenting paths and witnesses for degree-bounded msts. Lecture notes in computer science, 3624:26, 2005.
[3] Ted Fischer. Optimizing the degree of minimum weight spanning trees. Technical report, Cornell University, 1993.
[4] Pierre Fraigniaud. Approximation algorithms for minimum-time broadcast under the vertexdisjoint paths mode. AlgorithmsESA 2001, pages 440-451, 2001.
[5] Martin Fürer and Balaji Raghavachari. An nc approximation algorithm for the minimum degree spanning tree problem. In Proc. of the 28th Annual Allerton Conf. on Communication, Control and Computing, pages 274-281, 1990.
[6] Martin Furer and Balaji Raghavachari. Approximating the minimum-degree steiner tree to within one of optimal. Journal of Algorithms, 17(3):409-423, 1994.
[7] Harold N Gabow and Robert Endre Tarjan. A linear-time algorithm for a special case of disjoint set union. In Proceedings of the fifteenth annual ACM symposium on Theory of computing, pages 246-251. ACM, 1983.
[8] Michel X Goemans. Minimum bounded degree spanning trees. In Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on, pages 273-282. IEEE, 2006.
[9] Philip N Klein, Radha Krishnan, Balaji Raghavachari, and R Ravi. Approximation algorithms for finding low-degree subgraphs. Networks, 44(3):203-215, 2004.
[10] Jochen Könemann and R Ravi. A matter of degree: Improved approximation algorithms for degree-bounded minimum spanning trees. In Proceedings of the thirty-second annual ACM symposium on Theory of computing, pages 537-546. ACM, 2000.
[11] Jochen Könemann and R Ravi. Primal-dual meets local search: approximating mst's with nonuniform degree bounds. In Proceedings of the thirty-fifth annual ACM symposium on Theory of computing, pages 389-395. ACM, 2003.
[12] Radha Krishnan and Balaji Raghavachari. The directed minimum-degree spanning tree problem. In FSTTCS, volume 2245, pages 232-243. Springer, 2001.
[13] R Ravi. Rapid rumor ramification: Approximating the minimum broadcast time. In Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on, pages 202-213. IEEE, 1994.
[14] R Ravi, Madhav V Marathe, SS Ravi, Daniel J Rosenkrantz, and Harry B Hunt III. Many birds with one stone: Multi-objective approximation algorithms. In Proceedings of the twenty-fifth annual ACM symposium on Theory of computing, pages 438-447. ACM, 1993.
[15] Mohit Singh and Lap Chi Lau. Approximating minimum bounded degree spanning trees to within one of optimal. In Proceedings of the thirty-ninth annual ACM symposium on Theory of computing, pages 661-670. ACM, 2007.
[16] Daniel D Sleator and Robert Endre Tarjan. A data structure for dynamic trees. In Proceedings of the thirteenth annual ACM symposium on Theory of computing, pages 114-122. ACM, 1981.
[17] Guohui Yao, Daming Zhu, Hengwu Li, and Shaohan Ma. A polynomial algorithm to compute the minimum degree spanning trees of directed acyclic graphs with applications to the broadcast problem. Discrete Mathematics, 308(17):3951-3959, 2008.


[^0]:    *Institute for Interdisciplinary Information Sciences, Tsinghua University, duanran@mail.tsinghua.edu.cn
    ${ }^{\dagger}$ Institute for Interdisciplinary Information Sciences, Tsinghua University, hehq13@mails.tsinghua.edu.cn
    ${ }^{\ddagger}$ Institute for Interdisciplinary Information Sciences, Tsinghua University, tianyi-z16@mails.tsinghua.edu.cn

[^1]:    ${ }^{1} \tilde{O}(\cdot)$ hides poly-logarithmic factors.
    ${ }^{2} \alpha(\cdot)$ refers to the inverse Ackermann function.

