

Realization Problems on Reachability Sequences

Matthew Dippel, Ravi Sundaram, and Akshar Varma*

Northeastern University, Boston, USA

mdippel@ccs.neu.edu, r.sundaram@northeastern.edu, akshar@ccs.neu.edu

Abstract. The classical Erdős-Gallai theorem kicked off the study of graph realizability by characterizing degree sequences. We extend this line of research by investigating realizability of directed acyclic graphs (DAGs) given both a local constraint via degree sequences and a global constraint via a sequence of reachability values (number of nodes reachable from a given node). We show that, without degree constraints, DAG reachability realization is solvable in linear time, whereas it is strongly NP-complete given upper bounds on in-degree or out-degree. After defining a suitable notion of bicriteria approximation based on consistency, we give two approximation algorithms achieving $O(\log n)$ -reachability consistency and $O(\log n)$ -degree consistency; the first, randomized, uses LP (Linear Program) rounding, while the second, deterministic, employs a k -set packing heuristic. We end with two conjectures that we hope motivate further study of realizability with reachability constraints.

Keywords: Reachability sequences · Graph realization · Bicriteria approximation · Strong NP-completeness

1 Introduction

Given a property P , the Graph Realization problem asks whether there exists a graph that satisfies the property P . Starting with the Erdős-Gallai paper [10] on degree sequences [18,21] many other properties have been considered in the literature ranging from eccentricities [6,25] to connectivity and flow [11,12]. The best studied among these remain extensions of realization given degree sequences [1,4] and variants focusing on different subclasses of graphs [7,8,19,20,26]. In addition to their theoretical significance, realization questions occur naturally in numerous application contexts, including network design [12], social networks [5,28], DNA sequencing [29], enumerating chemical compounds [2], and phylogeny and evolutionary tree reconstruction [17,30].

We consider the realization problem on digraphs: we are given as input a sequence of tuples $(r_i, I(i), O(i))$, where r_i is the reachability value¹, $I(i)$ the in-degree and $O(i)$ the out-degree of each node $i, 1 \leq i \leq n$; and we wish to determine the existence of a digraph such that each node has the prescribed reachability value and the prescribed in-degree and out-degree. This formulation

* This work was supported by the NSF-CNS-1718286 and NSF-CCF-1535929 grants.

¹ The reachability value of a node is the number of nodes reachable from that node.

extends the local properties considered by degree sequences to global properties captured by the reachability sequence. Like all realization problems this has connections [1] to the graph isomorphism problem and graph canonization.

The study of reachability sequences has applications in several contexts. In the scientific context, reachability and degree constraints can reflect measurements obtained from naturally occurring networks with the aim being to generate a model that explains the measurements. Alternatively, from an engineering perspective, the goal may be to find an implementation satisfying the desired properties as specified by the reachability sequences. As an example scenario, consider the spread of a malicious virus in a network. Perhaps the first step to preventing the spread of this disease may be to understand the reach of infected nodes (using reachability values) before controlling the spread by disconnecting infected nodes from neighbors (using degree information). Alternatively, the goal may be to construct resilient networks that restrict the spread of the virus.

1.1 Our Contributions

We characterize the complexity of the realizability of acyclic digraphs as summarized in Table 1. Instead of referring to bounded vs. unbounded in-degree, we simply talk about trees vs. DAGs since trees are DAGs with in-degree bounded by one and they capture the essential behavior of bounded in-degree digraphs. This usage makes the exposition of the hardness results more natural.

Table 1: Reachability realization for unbounded out-degree DAGs is linear time. The other three cases are strongly NP-complete with bicriteria approximation algorithms achieving an approximation factor of $(O(\log n), O(\log n))$.

		Out-degree Bounded	Out-degree Unbounded
In-Degree	Bounded (Trees)	$(O(\log n), O(\log n))$	$(O(\log n), O(\log n))$
	Unbounded (DAGs)	$(O(\log n), O(\log n))$	Linear-time

- In Section 3 we give linear time verifiable, necessary and sufficient conditions, for realizing unbounded out-degree DAGs (Theorem 1).
- We define a notion of bicriteria approximation in Section 2 and give two algorithms in Section 4 to solve the reachability realization problem, both achieving an $(O(\log n), O(\log n))$ -approximation.
 - Theorem 3: Randomized LP rounding algorithm that runs in time $O(n^{\frac{37}{18}})$ if that the matrix multiplication exponent $\omega \approx 2$ and its dual $\alpha \approx 1$ [23].
 - Theorem 5: Deterministic algorithm using k -set packing heuristics [13,22] that runs in $n^{O(k^3)}$ time.
- In Section 5 we prove the strong NP-completeness of reachability realization when there are degree constraints. This includes one of our most technically involved results, a reduction from a generalized version of 3-PARTITION to show the strong NP-completeness of reachability realization when both in-degree and out-degree are bounded (Theorem 6). We also give simpler reductions from 3-PARTITION for reachability realization when only one of in-degree or out-degree is bounded (Theorems 7, 8).

Both approximation algorithms work in the presence of non-uniform degree bounds, that is, each degree might be a different value. On the other hand, our hardness results, except Theorem 8, prove that reachability realization problems are strongly NP-complete even when the degree bounds are uniform. In particular, we note that the Theorems 6 and 7 which have uniform degree bounds rely on having the in-degree bounded while in Theorem 8, where only the out-degree is bounded, we are only able to show hardness in the non-uniform case.

2 Preliminaries

Let n denote the length of the given reachability sequence and V the set of nodes in the corresponding graph, so that $|V| = n$. For node i in graph G we let $C(i)$ denote the set of children of i , i.e., $C(i) = \{j \mid (i, j) \in G\}$. The **out-degree** of i , $O_G(i)$ is the number of its children, i.e., $|C(i)|$; the **in-degree** of node i in graph G , $I_G(i)$ is the number of nodes with arcs directed into i . The reachability value of node u is the number of nodes it can reach: $r_u = |\{v : \exists \text{ path from } u \text{ to } v \in G\}|$. If the graph is a tree the reachability value r_i can be recursively defined as $1 + \sum_{j \in C(i)} r_j$. A rooted tree with out-degree upper bounded by k is called a **k -ary tree** otherwise they are **general trees**. **Full trees** are k -ary trees where every out-degree is either k or 0. A **complete k -ary tree** is a k -ary tree with every level except possibly the last filled and all nodes in the last level filled from the left. The unique reachability sequence of such trees is denoted by $T_k^c(n)$.

We now define the appropriate notions for the purpose of approximation. We say that a graph G is **δ -in-degree consistent** with graph H if they have the same set of nodes and if for all nodes i the following holds: $I_H(i) \leq I_G(i) \leq \delta \cdot I_H(i)$. Here in-degree can be replaced with out-degree to get **δ -out-degree consistency**. If a graph G is both in-degree and out-degree consistent with graph H , then we say it is **δ -degree-consistent**. For **ρ -reachability consistency** we generalize the idea of reachability to get a similar notion of approximation as degree consistency. Given a tree we say that it is ρ -reachability consistent if for all nodes i the following holds: $a_i \leq 1 + \sum_{j \in C(i)} a_j \leq \rho \cdot a_i$, where a_i is the reachability label on node i in the approximate solution. The above notion of approximation can be extended to DAGs by replacing the inequality constraint with $a_i \leq O_G(i) + \max_{j \in C(i)} a_j \leq \rho \cdot a_i$. Finally, we utilize the language of bicriteria optimization (see [9,27]) to say that G **(ρ, δ) -approximates** graph H if it is ρ -reachability consistent with the reachability sequence of H and it is δ -degree consistent with H . This captures the intuition that G approximately matches both the structure of G and its reachability sequence.

3 Linear time algorithm for DAGs

We show that there exist polynomial-time verifiable, necessary and sufficient conditions that characterize reachability sequences of unbounded out-degree DAGs. This is reminiscent of conditions for the reconstruction of graphs given degree sequences (see [10]). However it is in contrast to the hardness result of

degree realization for DAGs [8,20]. The inequalities in this section are deceptively simple considering the hardness results we prove in Section 5. Readers fond of puzzles are invited to prove the inequalities in Theorem 1 themselves before reading on.

Theorem 1 (DAG reachability). *Given a sequence of natural numbers $\{r_1, r_2, \dots, r_n\}$ in non-decreasing order there exists a DAG for which the given sequence is the sequence of the reachability sizes of the DAG iff $r_i \leq i$ for all i .*

Proof. In a DAG, nodes can only reach nodes with a strictly lower reachability otherwise its reachability would increase to a higher value causing a contradiction. Since there are at most $i - 1$ nodes of lower reachability than r_i , it can reach at most i nodes including itself. Hence $r_i \leq i$ is a necessary condition. Next, for all i , connect i to the first $r_i - 1$ nodes. Observe that, excluding itself, i cannot reach more than $r_i - 1$ nodes since every node j it connects to can only connect to a node k with $k < j$ but i is already connected to k . Hence it can reach exactly r_i nodes and the inequality is also a sufficient condition. \square

4 Approximation Algorithms

We present two approximation algorithms that are ρ -reachability consistent and δ -degree consistent with $\rho = \delta = O(\log n)$ given the reachability sequence along with the degree sequence. Thus the (ρ, δ) -approximation factor for our algorithms comes out to be $(O(\log n), O(\log n))$. The randomized algorithm runs in $O(n^{\omega + \max\{\frac{1}{18}, \omega - 2, \frac{1-\alpha}{2}\}})$ time as detailed in Section 4.1 while the deterministic algorithm runs in $n^{O(k^3)}$ time as we detail in Section 4.2. We compare further trade-offs between the two algorithms in Section 4.3 including the motivation for the more technically involved deterministic algorithm. While the exposition for both the algorithms addresses details using the full k -ary tree case, the results extend to all acyclic digraph cases in a straightforward manner by replacing reachability and degree consistency conditions for trees with that of DAGs.

4.1 LP based randomized rounding (LPRR) algorithm

The intuition behind the LPRR algorithm is to model the desired graph G as a collection of flows. Between every pair of nodes r_i and r_j with $r_i > r_j$ we assume a flow f_{ij} on each edge out of i and into j . We have three constraints for each node: the sum of flows into it is $l(i)$ (in-degree requirement), the sum of flows out of it is $O_G(i)$ (out-degree requirement), and that the reachability consistency conditions are satisfied. Further, there cannot be an edge (f_{ij} must be 0) from node i to node j if node i has a smaller reachability value than node j .

The existence of G guarantees that the LP is feasible. After solving for a feasible set of f_{ij} values we round each edge ij to 1 with probability f_{ij} independently $24 \ln n$ times. Each time an edge is rounded to 1 it is added to the solution (initialized to a graph with all nodes in V but no edges). We then argue

that the resulting structure satisfies the approximate reachability and degree consistency requirements with high probability using concentration bounds.

$$\begin{aligned}
 & \min && 1 \\
 \text{s. t.} && \sum_j f_{ji} = l(i) \quad \forall i, && \text{In-degree requirement} \\
 && \sum_j f_{ij} = O_G(i) \quad \forall i, && \text{Out-degree requirement} \\
 && r_i = 1 + \sum_j f_{ij} \cdot r_j \quad \forall i, && \text{Reachability consistency} \\
 && f_{ij} = 0 \quad \forall i, j \text{ s.t. } r_i \leq r_j && \text{Acyclicity}
 \end{aligned}$$

In particular, the following multiplicative form of the Chernoff bound is used.

Theorem 2. *Let $X = \sum_{i=1}^n X_i$, where X_i are independent Bernoulli trials with $\Pr[X_i = 1] = p_i \forall 1 \leq i \leq n$ and let $\mu = \mathbb{E}[X] = \sum_{i=1}^n p_i$. Then, for $0 < \epsilon < 1$*

$$\Pr[|X - \mu| \geq \epsilon] \leq 2e^{-\mu\epsilon^2/3} \quad (1)$$

Theorem 3 (LPRR). *Given a reachability sequence for a full k -ary tree, T , there exists a randomized $O(n^{\omega + \max\{\frac{1}{18}, \omega - 2, \frac{1-\alpha}{2}\}})$ -time algorithm that constructs a DAG that is an $(O(\log n), O(\log n))$ -approximation to T .*

Proof. Analysis of running time: Clearly the bottleneck here is the LP solver and the state of the art solver runs in $O(n^{\omega + \max\{\frac{1}{18}, \omega - 2, \frac{1-\alpha}{2}\}})$ time [23] where ω is the best known matrix multiplication exponent [24] and α is its dual. Further, under the common belief that $\omega \approx 2$ and $\alpha \approx 1$, our algorithm runs in $O(n^{\frac{37}{18}})$.

Proof of correctness: First, observe that vertex i has a total flow of $l(i)$ coming into it. So in one rounding the expected in-degree will be $l(i)$ and after $24 \ln n$ roundings the expected in-degree value will be $\mu_1 = 24 \ln n \cdot l(i)$. Invoking Chernoff bound with $\epsilon = 1/2$ we get that the probability that the node's in-degree lies outside the range $[\mu_1/2, 3\mu_1/2]$ is at most $2/n^2$. Similarly, we get the expected out-degree to be $\mu_2 = 24 \ln n \cdot O_G(i)$ and the probability that the out-degree of any node lies outside the range $[\mu_2/2, 3\mu_2/2]$ is at most $2/n^2$. Further for any node i , the reachability $1 + \sum_j f_{ij} \cdot r_j$ will have expected value $\mu_3 = 24 \ln n \cdot r_i$ and the probability that it lies outside the range $[\mu_3/2, 3\mu_3/2]$ is at most $2/n^2$. Applying the union bound over the $3n$ constraints in the LP, the probability that any of them lies outside their prescribed range is at most $3n \cdot \frac{2}{n^2} = o(1)$ as n goes to infinity. Thus with high probability after rounding, all of the quantities are within their prescribed ranges, i.e., the degree and reachability consistency are guaranteed to be within a logarithmic factor giving us the required (ρ, δ) -approximation. \square

4.2 k -set packing based deterministic algorithm

We give the intuition behind the algorithm, DSHS (Deterministic Sieving using Hurkens-Schrijver) before presenting the technical details. DSHS runs in two

(essentially independent) sieving phases, each phase taking $O(\log n)$ rounds: The MatchChildren phase matches each node (other than the leaves) with a (valid) set of children. The MatchParent phase matches each node (other than the root) with a parent.² Each phase starts with the entire set of candidate nodes and in each round sets up a $(k + 1)$ -set packing problem. The problem of k -set packing is to find the largest *disjoint* sub-collection of a given collection of sets each of cardinality k . The (approximate) solution to this problem sieves or reduces the candidate set by a constant factor, allowing each phase to finish in $O(\log n)$ rounds. Putting the results from the two phases together we get the desired (ρ, δ) -approximation factor. We use the following improvement of Hurkens-Schrijver’s algorithm [22]. Note that a smaller ϵ can improve the approximation but comes at the cost of a worse running time.

Theorem 4 (Theorem 5, Furer-Yu [13]; with $\epsilon = \frac{1}{3}$). *The $(k+1)$ -Set Packing problem can be approximated to a factor $\frac{3}{k+3}$ in deterministic time $n^{O(k^3)}$.*

Theorem 5 (DSHS). *Given a reachability sequence $\{r_1, r_2, \dots, r_n\}$ for a full k -ary tree, T , there exists a deterministic $n^{O(k^3)}$ -time algorithm that constructs a DAG that is an $(O(\log n), O(\log n))$ -approximation to T .*

Proof. **DSHS (Deterministic Sieving using Hurkens-Schrijver):** We initialize using an empty DAG with all the n nodes and no edges.

Phase MatchChildren: Initialize C_1 to be the set of all candidate nodes: nodes other than leaves (which have value 1). In round t the universe consists of C_t along with an entire set of V . Note that this has cardinality $|C_t| + |V|$ and is not the same as $C_t \cup V$. We create a collection of all possible $(k + 1)$ -sets with each set consisting of an element i from C_t and k elements, j_1, j_2, \dots, j_k from V such that $r_i = 1 + r_{j_1} + r_{j_2} + \dots + r_{j_k}$. Note that each $(k + 1)$ -set is a possible match for i to its children. The existence of T guarantees that the optimal solution to this $(k + 1)$ -Set Packing problem has size $|C_t|$ – namely the sub-collection consisting of each candidate node and its k children in T . Invoking the Hurkens-Schrijver approximation algorithm from the above theorem we are guaranteed to find a collection of sets that is at least $(3/(k + 2)) \cdot |C_t|$. We use this sub-collection of sets to augment our solution DAG with the corresponding arcs from the node i to each of its children (j_1, j_2, \dots, j_k) for each set. We also remove the corresponding candidate nodes i from C_t to get C_{t+1} . Phase MatchChildren ends when the candidate set C_t becomes empty.

Phase MatchParent: Initialize P_1 to be the set of all candidate nodes: nodes other than leaves. In round t the universe consists of P_t along with an entire set of V . Note that this has cardinality $|P_t| + |V|$ and is not the same as $P_t \cup V$. We create a collection of all possible $(k + 1)$ -sets with each set consisting of one element, i from P_t and k elements from V , of which one, j is the parent and the remaining $k - 1$ nodes j_1, j_2, \dots, j_{k-1} are siblings of i such that $r_j = 1 + r_i + r_{j_1} + r_{j_2} + \dots + r_{j_{k-1}}$. Note that each $(k + 1)$ -set is a possible match for i to its parent j . The existence of T guarantees that the optimal solution to this $(k + 1)$ -Set Packing problem has

² The root will have value n and leaves will have value 1.

size $|P_t|$ – namely the sub-collection consisting of each candidate node and its parent and siblings in T . Invoking the Hurkens-Schrijver approximation algorithm from the above theorem we are guaranteed to find a collection of sets that is at least $(3/(k+2)) \cdot |P_t|$. We use this sub-collection of sets to augment our solution DAG with the corresponding arcs from the node j to i and to each of its $k-1$ siblings for each set. We also remove the corresponding candidate nodes, i from P_t to get P_{t+1} . Phase MatchParent ends when the candidate set P_t becomes empty.

Analysis of running time: The bottleneck step is running the Hurkens-Schrijver approximation algorithm for set-packing which takes $n^{O(k^3)}$ time. Each of the two phases takes a logarithmic number of rounds, $\log_{\frac{k+2}{k-1}} n$ to be precise, which is absorbed into the total $n^{O(k^3)}$ -time since the big-O is in the exponent.

Proof of correctness: Note that after phase MatchChildren every eligible node is matched to exactly k children satisfying the reachability consistency condition exactly. However, some nodes may not have parents and some may have too many parents. Still every node is guaranteed to get at most one parent per round and so no node has more than $O(\log n)$ parents at the end of Phase MatchChildren. Similarly, after phase MatchParent every eligible node has at least one parent. However some parents may get too many children. Yet, in each round a parent gets at most k children and so no node gets more than $O(k \log n)$ children. Thus at the end of the two phases we are guaranteed $O(\log n)$ -degree consistency. Now observe that in each round of either phase, each node i either gets a valid set of k children, that is children j_1, j_2, \dots, j_k such that $r_i = 1 + r_{j_1} + r_{j_2} + \dots + r_{j_k}$, or no children at all; and we know that at the end of Phase MatchChildren every node other than leaves gets at least one valid set of children. Hence, we are guaranteed an $O(\log n)$ -reachability consistent solution. Thus the solution DAG at the end of both phases is an (ρ, δ) -approximation to T . \square

4.3 Trade-offs between the two Approximation Algorithms

The major trade-off between the DSHS and LPRR algorithms is the running time; while DSHS runs in $n^{O(k^3)}$, the LPRR algorithm is independent of k and runs in $O(n^\omega)$. Hence, unlike the deterministic algorithm, the randomized algorithm can be used even when k is a function of n . We also note that LPRR can be derandomized using the method of conditional probability [3]. While these might suggest that the more complex and technically involved DSHS algorithm is inferior, that is not the case. The LPRR algorithm results in more complex solutions, in particular, LPRR may return digraphs with multi-edges while DSHS is guaranteed to return simple digraphs. Also, the multi-edges provide a tighter concentration of reachability consistency, albeit away from the reachability values, which may be a desirable property in applications where certainty is more important than consistency. While that is an important application, it is more common to require simple digraphs which the randomized algorithm cannot guarantee. This motivates the more technically involved DSHS algorithm.

5 Strong NP-completeness Results

When the in-degrees and/or the out-degrees are constrained by the degree sequence we prove strong NP-completeness [15] using pseudo-polynomial transformations [16]. The reductions embed an instance of problems like 3-PARTITION between two consecutive levels of a tree. We first present the proof for the full k -ary tree realization problem (all in-degrees 1 or 0 and all out-degrees k or 0) in Section 5.1 which illustrates all the technicalities involved. We give a simpler reduction (in Section 5.2) to prove that realization of general trees (no out-degree bound) is also strongly NP-complete. In Section 5.3 we give a reduction to prove strong NP-completeness when there is only an out-degree bound.

5.1 Hardness of Realization for Full k -ary Trees

We prove hardness of the realization problem for full k -ary trees by reduction from the K-PWT problem, which we show to be strongly NP-complete via an involved series of reductions in the full version. Since K-PWT is strongly NP-complete we can reduce from a subclass, Π_p , such that the largest number in the instance is polynomially bounded, formally, $\text{MAX}[I] \leq p(\text{LENGTH}[I])$, $\forall I \in \Pi_p$.

Problem 1 (K-PWT). Given a set X with $|X| = Km$, $K \geq 2$, sizes $s : X \mapsto \mathbb{Z}^+$ and a target vector $B = (b_1, \dots, b_m) \in \mathbb{N}^m$, can X be partitioned into m disjoint sets A_1, A_2, \dots, A_m , such that, $|A_i| = k$ and $\sum_{a \in A_i} s(a) = b_i$, for $1 \leq i \leq m$?

Theorem 6 (Full k -ary tree). *It is strongly NP-complete to determine the existence of a full k -ary tree whose reachability sequence equals a given sequence.*

Proof. The problem is clearly in NP since a tree acts as a certificate. Set $k = K$ and define a number M which is a power of K , is much greater in magnitude than any of the other numbers in the problem, and is polynomially bounded by the maximum integer in the K-PWT instance (Eq. 2). We also define m' and m'' such that $m + m'$ and $m + m''$ are powers of K (Eq. 3).

$$M_1 = \max(\{s(x_i) | x_i \in X\} \cup \{b_i | b_i \in B\}); M_2 = KmM_1; M = K^{\lceil \log_k M_2 \rceil} \quad (2)$$

$$m' = K^d - Km, \quad m'' = K^{d-1} - m = m'/K, \quad \text{where } d = \lceil \log_K(Km) \rceil \quad (3)$$

We make the sequence $S = C \cup P \cup G \cup D$ using four ‘‘component’’ sequences: the ‘‘child component’’ $C = C' \cup C''$, the ‘‘parent component’’ $P = P' \cup P''$, the ‘‘ancestor component’’ G and the ‘‘descendant component’’ D .

The ‘‘child component’’ C is the union of the C' and C'' while the ‘‘parent component’’ P is the union of the P' and P'' . C' is in one-to-one correspondence with the set X , using the sizes of elements from X with M added to them while P' is in one-to-one-correspondence with B with changes to accommodate those made to sizes of elements of X while making C' . The sets C'' and P'' ensure that the cardinality of C and P respectively are a power of K .

We construct the ‘‘ancestor component’’ in ‘‘levels’’. The lowest level l_{d-2} is constructed from P , by arbitrarily taking blocks of K elements, adding them all up

and incrementing the result by one. Formally, order the elements in P arbitrarily as $P_1, P_2, \dots, P_{K^{d-1}}$ and let $l_{d-2} = \{l_{d-2,i} \mid l_{d-2,i} = 1 + \sum_{j=1}^K P_{(i-1)K+j}, 1 \leq i \leq K^{d-2}\}$. Other levels l_{d-i} are constructed in a similarly from levels l_{d-i+1} . This is continued until l_0 which has only one element since $|P|$ is a power of K and the size of each level above reduces by a factor of K . The element in l_0 would be the largest number in the final instance. The “descendant component” is constructed using reachability sequences of complete trees on the elements $c_i \in C$. For each such c_i , we make a complete k -ary tree on c_i nodes and use its reachability sequence $T_k^c(c_i)$. The descendant component is then $D = \bigcup_i T_k^c(c_i)$. Since each $c_i \in C$ has the form $Kx + 1$, $T_k^c(c_i)$ will also be full trees.

$$C' = \{K(s(x) + M) + 1 \mid x \in X\}, C'' = \{\overbrace{KM + 1, \dots, KM + 1}^{m' \text{ times}}\} \quad (4)$$

$$P' = \{K(b_i + KM + 1) \mid b_i \in B\}, P'' = \{\overbrace{K^2M + K, \dots, K^2M + K}^{m'' \text{ times}}\} \quad (5)$$

$$G = \bigcup_{i=0}^{d-2} l_i, \text{ where “levels” } l_i \text{ are defined in text; } D = \bigcup_{i=1}^{K^d} T_k^c(c_i), \forall c_i \in C \quad (6)$$

Constructing S takes polynomial time as elements in P and C are derived directly from the K-PWT instance, there are a logarithmic number of levels each computed in polynomial time, and the polynomial number of trees in D each be computed in linear time. All that remains is to prove that the reduction is valid.

By construction, elements in G and P will form a partial full k -ary tree with the elements in P as the “leaves”, elements in C and D will make a forest with elements from C as roots of the trees in the forest, and C'' can be arbitrarily partitioned to connect to P'' elements. If there is a partition of X , we can partition C' accordingly to complete the tree.

To prove that a full k -ary tree implies a partition it is sufficient to prove that in any tree the set of children of P' is equal to C' , that is, the nodes of P' and C' occur in consecutive levels in any tree. The node in l_0 is the largest and will necessarily have to be the root. This will be followed by the nodes from l_1 since no other nodes are large enough to reach those in l_0 (given the out-degree bound of k). Continuing the argument, $l_i \in G$ will always appear in consecutive levels in any tree and that P will follow below G . Since the in-degree is 1 no node from $p \in P$ will be a child of any $p' \in P$. Further, nodes in D will all be less than M/K in value and hence k of them will not be enough to reach nodes in P thus necessitating that all children of nodes of P come from C . We note that nodes from C' can not be children of nodes from P'' and so the set of children of P'' have to be C'' . Since a value of the order of KM has to be reached for nodes in P' and all nodes in C' are of the order of M , all the nodes from C' will be used. Thus any tree will have nodes from P' and C' in consecutive levels and therefore have a partition. This proves NP-completeness and as the maximum integer used is polynomially bounded it also proves strong NP-completeness. \square

5.2 Hardness of Realization for General Trees

Theorem 7 (General Trees). *It is strongly NP-complete to realize a general tree given a reachability sequence.*

Proof. This problem is in NP as a tree acts as a certificate and all that remains is give a reduction from 3-PARTITION.

Problem 2 (3-PARTITION). Given a set A , a target B and a size function $s : A \rightarrow \mathbb{Z}^+$ such that $|A| = 3m$ and $B/4 < s(a_i) < B/2 \forall a_i \in A$, can A be partitioned into m disjoint sets A_1, A_2, \dots, A_m , such that for $1 \leq i \leq m$, $\sum_{a \in A_i} s(a) = B$?

Since this problem is strongly NP-complete [14], for the reduction, we use an instance of 3-PARTITION wherein the numbers are polynomially bounded in the input length. The notation $[n]$ is used for the set consisting of the first n natural numbers. The constructed instance is as follows:

$$S = \{m(B+1) + 1\} \cup \{B+1, \dots, m \text{ times } \dots, B+1\} \cup \left(\bigcup_{i=1}^m [s(a_i)] \right)$$

This construction is clearly polynomial time, all that remains is to show that this is valid reduction. In an abuse of notation the reachability size is often used to refer to the node with that reachability size.

In any potential tree, then the $m(B+1) + 1$ node is forced to be the root as it cannot be the child of any node. Further, it will have m children, all the $B+1$ nodes as they cannot be the children of any other node. Considering the remaining nodes in a bottom up fashion, we see that a fixed structure is enforced on all these nodes due to the in-degree constraint in rooted trees. Any 2 node can only have a 1 as it's child, exhausting all 1's and leaving only 2's to be single children of 3's and only 3's as children for 4's and so on. By construction, there are exactly as many nodes of size $s-1$ as there are nodes of size s for all $s < B+1$ and hence they all get exhausted. This enforces that each node labelled $s(a_i)$ is a root of a path consisting of nodes with sub-tree sizes $[s(a_i)]$.

These restrictions are always present and a tree can be realized iff the paths rooted at $s(a_i)$ can be correctly made children of the $B+1$ nodes. This happens iff there is a partition of the 3-PARTITION instance; if there is no partition then the paths cannot be joined to the partial tree above it to form a single tree. \square

5.3 Hardness of Realization with out-degree constraints

Theorem 8 (Bounded out-degree). *It is strongly NP-complete to realize an acyclic graph given a reachability sequence and out-degree constraints.*

Proof. We reduce from 3-PARTITION, again using an instance with the maximum number polynomially bounded in the length of the problem. Let $s_i := s(a_i)$ and $M := mB^2$ and note that M is much bigger than every number in the 3-PARTITION instance and that the $\sum_i s_i = mB$. Let the reachability sequence $S := S_s \cup S_b \cup S_a$ where $S_s = \{Ms_i\}$, S_b be a multiset with m copies of $MB+1$,

and S_1 be a multiset of $mB - |A|$ copies of 1. Let the out-degree constraint for $MB + 1$ nodes be equal to three and for the s_i nodes be equal to s_i .

To achieve out-degree constraints, each of the s_i nodes will need to pick up $s_i - 1$ ones exhausting nodes in S_1 . For the $MB + 1$ nodes to achieve their degree requirement they'll have to pick up exactly three nodes from S_s set which will be possible iff there is a valid partition of the 3-PARTITION instance. \square

6 Conclusion

In this paper we initiate the study of the realization problem for DAGs and rooted directed trees given a reachability sequence. We provide a linear time algorithm for DAGs with unbounded out-degrees and show hardness results for variants when we are also given a degree sequence bounding the in-degree and/or out-degree. We define a notion of bicriteria approximation based on reachability and degree consistency and give two $(O(\log n), O(\log n))$ -approximation algorithms for all of these problems. We conclude with two intriguing conjectures:

- Given a uniform out-degree bound and a reachability sequence the DAG realizability problem is solvable in poly-time.
- The general digraph realizability problem given a reachability sequence (with or without degree sequences) is strongly NP-complete.

References

1. Aigner, M., Triesch, E.: Realizability and uniqueness in graphs. *Discrete Mathematics* **136**(1-3), 3–20 (1994)
2. Akutsu, T., Nagamochi, H.: Comparison and enumeration of chemical graphs. *Computational and structural biotechnology journal* **5**(6), e201302004 (2013)
3. Alon, N., Spencer, J.H.: *The Probabilistic Method*. Wiley Publishing, 4th edn. (2016)
4. Bar-Noy, A., Choudhary, K., Peleg, D., Rawitz, D.: Realizability of graph specifications: Characterizations and algorithms. In: *International Colloquium on Structural Information and Communication Complexity*. pp. 3–13. Springer (2018)
5. Bar-Noy, A., Choudhary, K., Peleg, D., Rawitz, D.: Graph profile realizations and applications to social networks. In: *International Workshop on Algorithms and Computation*. pp. 3–14. Springer (2019)
6. Behzad, M., Simpson, J.E.: Eccentric sequences and eccentric sets in graphs. *Discrete Mathematics* **16**(3), 187–193 (1976)
7. Berger, A.: A note on the characterization of digraphic sequences. *Discrete Mathematics* **314**, 38–41 (2014)
8. Berger, A., Müller-Hannemann, M.: How to attack the np-complete dag realization problem in practice. In: *International Symposium on Experimental Algorithms*. pp. 51–62. Springer (2012)
9. Ehrgott, M., Gandibleux, X.: A survey and annotated bibliography of multiobjective combinatorial optimization. *OR-Spektrum* **22**(4), 425–460 (2000)
10. Erdos, P., Gallai, T.: Graphs with points of prescribed degrees. *Mat. Lapok* **11**(264-274), 132 (1960)

11. Frank, A.: Augmenting graphs to meet edge-connectivity requirements. *SIAM Journal on Discrete Mathematics* **5**(1), 25–53 (1992)
12. Frank, H., Chou, W.: Connectivity considerations in the design of survivable networks. *IEEE Transactions on Circuit Theory* **17**(4), 486–490 (1970)
13. Fürer, M., Yu, H.: Approximating the k -set packing problem by local improvements. In: *International Symposium on Combinatorial Optimization*. pp. 408–420. Springer (2014)
14. Garey, M.R., Johnson, D.S.: Complexity results for multiprocessor scheduling under resource constraints. *SIAM Journal on Computing* **4**(4), 397–411 (1975)
15. Garey, M.R., Johnson, D.S.: “Strong”NP-Completeness Results: Motivation, Examples, and Implications. *Journal of the ACM (JACM)* **25**(3), 499–508 (1978)
16. Garey, M.R., Johnson, D.S.: *Computers and Intractability: A guide to the theory of NP-completeness*. WH Free. Co., San Fr (1979)
17. Guindon, S., Gascuel, O.: A simple, fast, and accurate algorithm to estimate large phylogenies by maximum likelihood. *Systematic biology* **52**(5), 696–704 (2003)
18. Hakimi, S.L.: On realizability of a set of integers as degrees of the vertices of a linear graph I. *Journal of the Society for Industrial and Applied Mathematics* **10**(3), 496–506 (1962)
19. Harary, F.: A survey of the reconstruction conjecture. In: *Graphs and Combinatorics*, pp. 18–28. Springer (1974)
20. Hartung, S., Nichterlein, A.: Np-hardness and fixed-parameter tractability of realizing degree sequences with directed acyclic graphs. In: *Conference on Computability in Europe*. pp. 283–292. Springer (2012)
21. Havel, V.: A remark on the existence of finite graphs. *Casopis Pest. Mat* **80**(477-480), 1253 (1955)
22. Hurkens, C.A.J., Schrijver, A.: On the size of systems of sets every t of which have an sdr, with an application to the worst-case ratio of heuristics for packing problems. *SIAM Journal on Discrete Mathematics* **2**(1), 68–72 (1989)
23. Jiang, S., Song, Z., Weinstein, O., Zhang, H.: Faster Dynamic Matrix Inverse for Faster LPs. arXiv preprint arXiv:2004.07470 (2020)
24. Le Gall, F.: Powers of tensors and fast matrix multiplication. In: *Proceedings of the 39th international symposium on symbolic and algebraic computation*. pp. 296–303. ACM (2014)
25. Lesniak, L.: Eccentric sequences in graphs. *Periodica Mathematica Hungarica* **6**(4), 287–293 (1975)
26. Lovász, L.: A note on the line reconstruction problem. In: *Classic Papers in Combinatorics*, pp. 451–452. Springer (2009)
27. Marathe, M.V., Ravi, R., Sundaram, R., Ravi, S., Rosenkrantz, D.J., Hunt III, H.B.: Bicriteria network design problems. *Journal of algorithms* **28**(1), 142–171 (1998)
28. Mihail, M., Vishnoi, N.K.: On generating graphs with prescribed vertex degrees for complex network modeling. *Position Paper, Approx. and Randomized Algorithms for Communication Networks (ARACNE)* **142** (2002)
29. Mossel, E., Ross, N.: Shotgun assembly of labeled graphs. *IEEE Transactions on Network Science and Engineering* (2017)
30. Ng, M.P., Wormald, N.C.: Reconstruction of rooted trees from subtrees. *Discrete Applied Mathematics* **69**(1-2), 19–31 (1996)