Greedy Can Beat Pure Dynamic Programming

Stasys Jukna[1], Hannes Seiwert

Institute of Computer Science, Goethe-University Frankfurt, Robert-Mayer-Str. 11-15, Frankfurt am Main, Germany

Abstract

Many dynamic programming algorithms for discrete 0-1 optimization problems are "pure" in that their recursion equations only use min/max and addition operations, and do not depend on actual input weights. The well-known greedy algorithm of Kruskal solves the minimum weight spanning tree problem on n-vertex graphs using only $O(n^2 \log n)$ operations. We prove that any pure DP algorithm for this problem must perform $2^{\Omega(\sqrt{n})}$ operations. Since the greedy algorithm can also badly fail on some optimization problems, easily solvable by pure DP algorithms, our result shows that the computational powers of these two types of algorithms are incomparable.

Keywords: Spanning tree; arborescence; arithmetic circuit; tropical circuit; lower bound

1. Introduction and result

A dynamic programming (DP) algorithm is *pure* if it only uses min or max and addition operations in its recursion equations, and the equations do not depend on the actual values of the input weights. Notable examples of such DP algorithms include the Bellman–Ford–Moore algorithm for the shortest s-t path problem [7, 15, 1], the Floyd–Warshall algorithm for the all-pairs shortest paths problem [6, 18], the Held–Karp algorithm for the Travelling Salesman Problem [8], and the Dreyfus–Levin–Wagner algorithm for the weighted Steiner tree problem [3, 14].

It is well known and easy to show that, for some optimization problems, already pure DP algorithms can be much better than greedy algorithms. Namely, there are a lot of optimization problems which are easily solvable by pure DP algorithms (exactly), but the greedy algorithm cannot even achieve any finite approximation factor: maximum weight independent set in a path, or in a tree, the maximum weight simple s-t path in a transitive tournament problem, etc.

In this paper, we show that the converse direction also holds: for some optimization problems, greedy algorithms can also be much better than pure dynamic programming. So, the computational powers of greedy and pure DP algorithms are *incomparable*. We will show that the gap occurs on the (undirected) minimum weight spanning tree problem, by first

Email addresses: `stjukna@gmail.com` (Stasys Jukna), `seiwert@thi.cs.uni-frankfurt.de` (Hannes Seiwert)

deriving an exponential lower bound on the monotone arithmetic circuit complexity of the corresponding polynomial.

Let K_n be the complete undirected graph on $[n] = \{1, \ldots, n\}$. Assume that edges e have their associated nonnegative real weights x_e, considered as formal variables. Let $\mathcal{T}_n$ be the family of all $|\mathcal{T}_n| = n^{n-2}$ spanning trees T in K_n, each viewed as its *set* of edges.

It is well known that $\mathcal{T}_n$ is the family of bases of a matroid, known as *graphic matroid*; so, on this family of feasible solutions, both optimization problems (minimization and maximization) can be solved by standard greedy algorithms. On the other hand, the theorem below states that the polynomial corresponding to $\mathcal{T}_n$ has exponential monotone arithmetic circuit complexity; due to special properties of this polynomial, the same lower bound also holds on the number of operations used by pure DP algorithms solving minimization and maximization problems on $\mathcal{T}_n$ (see Lemma 1).

The *spanning tree polynomial* (known also as the *Kirchhoff polynomial* of K_n) is the following homogeneous, multilinear polynomial of degree $n - 1$:

$$f_n(x) = \sum_{T \in \mathcal{T}_n} \prod_{e \in T} x_e \,.$$

For a multivariate polynomial f with positive coefficients, let $L(f)$ denote the minimum size of a monotone arithmetic $(+, \times)$ circuit computing f. Our goal is to prove that $L(f_n)$ is exponential in n.

Theorem 1.
$$L(f_n) = 2^{\Omega(\sqrt{n})} \,.$$

A "directed version" of f_n is the *arborescence polynomial* $\vec{f}_n$. An *arborescence* (known also as a *branching* or a *directed spanning tree*) on the vertex-set $[n]$ is a directed tree with edges oriented away from vertex 1 such that every other vertex is reachable from vertex 1. Let $\vec{\mathcal{T}}_n$ be the family of all arborescences on $[n]$. Jerrum and Snir [11] have shown that $L(\vec{f}_n) = 2^{\Omega(n)}$ holds for the arborescence polynomial

$$\vec{f}_n(x) = \sum_{T \in \vec{\mathcal{T}}_n} \prod_{\vec{e} \in T} x_{\vec{e}} \,.$$

Note that here variables $x_{i,j}$ and $x_{j,i}$ are treated as *distinct*, and cannot both appear in the same monomial. This dependence on orientation was crucially utilized in the argument of [11, p. 892] to reduce a trivial upper bound $(n-1)^{n-1}$ on the number of monomials in a polynomial computed at a particular gate till a non-trivial upper bound $(3n/4)^{n-1}$. So, this argument does not apply to the *undirected* version f_n (where $x_{i,j}$ and $x_{j,i}$ stand for the *same* variable). To handle the undirected case, we will use an entirely different argument.

Relation to pure DP algorithms. Every pure DP algorithm is just a special (recursively constructed) *tropical* $(\min, +)$ or $(\max, +)$ circuit, that is, a circuit using only min (or max) and addition operations as gates; each input gate of such a circuit holds either one of the variables x_i or a nonnegative real number. So, lower bounds on the size of tropical circuits yield the same lower bounds on the number of operations used by pure DP algorithms. For optimization problems, whose feasible solutions all have the same cardinality, the task of proving lower bounds on their tropical circuit complexity can be solved by proving lower bounds on the size of monotone arithmetic circuits.

Recall that a multivariate polynomial is *monic* if all its nonzero coefficients are equal to 1, *multilinear* if no variable occurs with degree larger than 1, and *homogeneous* if all monomials have the same degree. Every monic and multilinear polynomial $f(x) = \sum_{S \in \mathcal{F}} \prod_{i \in S} x_i$ defines two optimization problems: compute the minimum or the maximum of $\sum_{i \in S} x_i$ over all $S \in \mathcal{F}$.

Lemma 1 ([11, 12]). *If a polynomial f is monic, multilinear and homogeneous, then every tropical circuit solving the corresponding optimization problem defined by f must have at least $L(f)$ gates.*

This fact was proved by Jerrum and Snir [11, Corollary 2.10]; see also [12, Theorem 9] for a simpler proof. The proof idea is fairly simple: having a tropical circuit, turn it into a monotone arithmetic $(+, \times)$ circuit, and use the homogeneity of f to show that, after removing some of the edges entering $+$-gates, the resulting circuit will compute our polynomial f.

Greedy can beat pure DP. The (weighted) *minimum spanning tree* problem $MST_n(x)$ is, given an assignment of nonnegative real weights to the edges of K_n, compute the minimum weight of a spanning tree of K_n, where the weight of a graph is the sum of weights of its edges. So, this is exactly the minimization problem defined by the spanning tree polynomial f_n:

$$MST_n(x) = \min_{T \in \mathcal{T}_n} \sum_{e \in T} x_e\,.$$

Since the family $\mathcal{T}_n$ of feasible solutions of this problem is the family of bases of the (graphic) matroid, the problem *can* be solved by the standard greedy algorithm. In particular, the well-known greedy algorithm of Kruskal [13] solves MST_n using only $O(n^2 \log n)$ operations.

On the other hand, since the spanning tree polynomial f_n is monic, multilinear and homogeneous, Theorem 1 together with Lemma 1 implies that any $(\min, +)$ circuit solving the problem MST_n must have at least $L(f_n) = 2^{\Omega(\sqrt{n})}$ gates and, hence, at least so many operations must be performed by any pure DP algorithm solving MST_n. This gap between pure DP and greedy algorithms is our main result.

Directed versus undirected spanning trees. The arborescence polynomial $\vec{f}_n$ is also monic, multilinear and homogeneous, so that Lemma 1, together with the above mentioned lower bound on $L(\vec{f}_n)$ due to Jerrum and Snir [11], also yields the same lower bound on the size of $(\min, +)$ circuits solving the minimization problem on the family $\vec{\mathcal{T}}_n$ of arborescences.

But this does not separate DP from greedy, because the downward closure of $\vec{\mathcal{T}}_n$ is *not* a matroid: it is only an intersection of two matroids (see Edmonds [4]). So, greedy algorithms are only able to *approximate* the minimization problem on $\vec{\mathcal{T}}_n$ within the factor 2. Polynomial time algorithms solving this problem *exactly* were found by several authors, starting from Edmonds [5]. The fastest algorithm for the problem is due to Tarjan [16], and solves it in time $O(n^2 \log n)$, that is, with the same time complexity as Kruskal's greedy algorithm for undirected graphs [13]. But these are not greedy algorithms. So, $\vec{\mathcal{T}}_n$ does not separate standard, matroid based greedy and pure DP algorithms.

2. Proof of Theorem 1

A *rectangle* is specified by giving two families $\mathcal{A}$ and $\mathcal{B}$ of forests in the complete graph K_n on $[n] = \{1, \ldots, n\}$ such that for all forests $A \in \mathcal{A}$ and $B \in \mathcal{B}$ (viewed as sets of their

edges), we have $A \cap B = \emptyset$ (the forests are edge-disjoint), and $A \cup B$ is a spanning tree of K_n. The rectangle itself is the family

$$\mathcal{R} = \mathcal{A} \vee \mathcal{B} := \{A \cup B \colon A \in \mathcal{A} \text{ and } B \in \mathcal{B}\}$$

of all resulting spanning trees. A rectangle $\mathcal{R} = \mathcal{A} \vee \mathcal{B}$ is *balanced* if $(n-1)/3 \leq |A|, |B| \leq 2(n-1)/3$ holds for all forests $A \in \mathcal{A}$ and $B \in \mathcal{B}$; recall that every spanning tree of a graph on n vertices has $n-1$ edges. Let $\tau(n)$ be the minimum number of balanced rectangles whose union gives the family of all spanning trees of K_n.

Lemma 2. *For the spanning tree polynomial f_n, we have $L(f_n) \geq \tau(n)$.*

Proof. Let $t = L(f_n)$. The spanning tree polynomial f_n is multilinear and homogeneous of degree $n-1$: every spanning tree T of K_n has $|T| = n-1$ edges. Since the polynomial f_n is homogeneous of degree $n-1$, and since f_n can be computed by a monotone arithmetic circuit of size t, the well-known decomposition result, proved by Hyafil [10, Theorem 1] and Valiant [17, Lemma 3], implies that f_n can be written as a sum $f_n = g_1 \cdot h_1 + \cdots + g_t \cdot h_t$ of products of nonnegative homogeneous polynomials, each of degree at most $2(n-1)/3$; a polynomial is *nonnegative* if it has no negative coefficients.

Every monomial of f_n is of the form $\prod_{e \in T} x_e$ for some spanning tree T. Since the polynomials g_i and h_i in the decomposition of f_n are nonnegative, there can be no cancellations. This implies that all the monomials of $g_i \cdot h_i$ must be also monomials of f_n, that is, must correspond to spanning trees. Moreover, since the polynomial f_n is multilinear, the forests of g_i must be edge-disjoint from the forests of h_i. So, if we let $\mathcal{A}_i$ be the family of forests corresponding to monomials of the polynomial g_i, and $\mathcal{B}_i$ be the family of forests corresponding to monomials of the polynomial h_i, then $\mathcal{A}_1 \vee \mathcal{B}_1, \ldots, \mathcal{A}_t \vee \mathcal{B}_t$ are balanced rectangles, and their union gives the family of all spanning trees of K_n. This shows $\tau(n) \leq t = L(f_n)$, as desired. $\qquad\square$

So, it is enough to prove an exponential lower bound on $\tau(n)$. When doing this, we will concentrate on spanning trees of K_n of a special form. Let m and d be positive integer parameters satisfying $(d+1)m = n$, $m = \Theta(\sqrt{n})$ and $m \leq d/32$; we will specify these parameters later.

A *star* is a tree with one vertex, the *center*, adjacent to all the others, which are *leaves*. A *d-star* is a star with d leaves. A *spanning star-tree* consists of m vertex-disjoint d-stars whose centers are joined by a path. A *star factor* is a spanning forest of K_n consisting of m vertex-disjoint d-stars. Note that each spanning star-tree contains a unique star factor (obtained by removing edges between star centers).

Let $\mathcal{F}$ be the family of all star factors of K_n. For a rectangle $\mathcal{R}$, let $\mathcal{F}_\mathcal{R}$ denote the family of all star factors F of K_n contained in at least one spanning tree of $\mathcal{R}$; in this case, we also say that the factor F is *covered* by the rectangle $\mathcal{R}$.

Lemma 3. *There is an absolute constant $c > 0$ such that for every balanced rectangle $\mathcal{R}$, we have $|\mathcal{F}_\mathcal{R}| \leq |\mathcal{F}| \cdot 2^{-c\sqrt{n}}$.*

Note that this lemma gives a lower bound $\tau(n) \geq 2^{c\sqrt{n}}$ on the minimum number of balanced rectangles containing all spanning trees of K_n. Indeed, let $\mathcal{R}_1, \ldots, \mathcal{R}_t$ be $t = \tau(n)$ balanced rectangles whose union is the family of all spanning trees of K_n. Every star factor $F \in \mathcal{F}$ is contained in at least one spanning tree (in fact, in many of them). So, every star

4

factor $F \in \mathcal{F}$ must be covered by at least one of these t rectangles. But Lemma 3 implies that none of these rectangles can cover more than $h := |\mathcal{F}| \cdot 2^{-c\sqrt{n}}$ star factors $F \in \mathcal{F}$. So, we need $\tau(n) = t \geq |\mathcal{F}|/h \geq 2^{c\sqrt{n}}$ rectangles. Together with Lemma 2, this yields the desired lower bound $L(f_n) \geq 2^{c\sqrt{n}}$ on the monotone arithmetic circuit complexity of the spanning tree polynomial f_n.

The rest of the paper is devoted to the proof of Lemma 3.

Proof of Lemma 3. We can construct every star factor $F \in \mathcal{F}$ using the following procedure.

1. Choose a subset of m centers in $[n]$; $\binom{n}{m}$ possibilities.

2. Divide the remaining $n - m$ vertices into m blocks of size d, and connect all vertices of the ith block to the ith largest of the chosen centers; there are $\binom{n-m}{d,\dots,d} = \frac{(n-m)!}{d!^m}$ possibilities to do this.

Since different realizations of this procedure lead to different star factors, we have

$$|\mathcal{F}| = \binom{n}{m}\frac{(n-m)!}{d!^m}. \tag{1}$$

Fix a balanced rectangle $\mathcal{R} = \mathcal{A} \vee \mathcal{B}$ containing at least one spanning star-tree $T_0 = A_0 \cup B_0$ with $A_0 \in \mathcal{A}$ and $B_0 \in \mathcal{B}$, and let $c_1, \dots, c_m$ be the centers of stars of T_0. Every vertex $v \in [n] \setminus \{c_1, \dots, c_m\}$ is connected in T_0 by a *unique* edge e_v to one of the centers $c_1, \dots, c_m$. This gives us a partition $U \cup V$ of the vertices in $[n] \setminus \{c_1, \dots, c_m\}$ into two sets determined by the forests A_0 and B_0:

$$U = \{v \colon e_v \in A_0\} \quad \text{and} \quad V = \{v \colon e_v \in B_0\}.$$

We will concentrate on the bipartite complete subgraph $U \times V$ of K_n, and call the edges of K_n lying in this subgraph *crossing edges*. Since our rectangle $\mathcal{R}$ is balanced, we know that both $|A_0|$ and $|B_0|$ lie between $(n-1)/3$ and $2(n-1)/3$. So, since $m = o(n)$, for n large enough, we have

$$|U|, |V| \geq \tfrac{1}{3}(n-1) - m \geq \tfrac{1}{4}n. \tag{2}$$

The property that every graph $A \cup B$ with $A \in \mathcal{A}$ and $B \in \mathcal{B}$ must be cycle-free (must be a spanning tree of K_n) gives the following restriction on the rectangle $\mathcal{R} = \mathcal{A} \vee \mathcal{B}$.

Claim 1. *For all forests $A \in \mathcal{A}$ and $B \in \mathcal{B}$, and vertices $u \in U$ and $v \in V$, we have $|A \cap (\{u\} \times V)| \leq m$ and $|B \cap (U \times \{v\})| \leq m$.*

That is, no forest $A \in \mathcal{A}$ can contain more than m crossing edges incident to one vertex in U, and no forest $B \in \mathcal{B}$ can contain more than m crossing edges incident to one vertex in V.

Proof. Assume contrariwise that some vertex $u \in U$ has $l \geq m+1$ crossing edges $\{u, v_1\}, \dots, \{u, v_l\}$ in the forest A. Since these edges are crossing and $u \in U$, all vertices $v_1, \dots, v_l$ belong to V. In the (fixed) spanning star-tree $T_0 = A_0 \cup B_0$ (determining the partition $U \cup V$ of vertices in $[n] \setminus \{c_1, \dots, c_m\}$) each of these l vertices is joined by an edge of the forest B_0 to one of the centers $c_1, \dots, c_m$ of stars of T_0.

Since $l > m$, some two of these vertices v_i and v_j must be joined in B_0 to the same center $c \in \{c_1, \dots, c_m\}$. Since $\mathcal{R}$ is a rectangle, the graph $A \cup B_0$ must be a (spanning) tree. But the edges $\{u, v_i\}, \{u, v_j\}$ of A together with edges $\{v_i, c\}, \{v_j, c\}$ of B_0 form a cycle $u \to v_i \to c \to v_j \to u$ in $A \cup B_0$, a contradiction.

5

The proof of the inequality $|B \cap (U \times \{v\})| \leq m$ is the same by using the forest A_0 instead of B_0. $\square$

So far, we only used one fixed spanning tree T_0 in the rectangle $\mathcal{R}$ to define the subgraph $U \times V$ of K_n. We now use the entire rectangle $\mathcal{R} = \mathcal{A} \vee \mathcal{B}$ to color the *edges* of K_n in red and blue. When doing this, we use the fact that the sets $E_\mathcal{A} := \bigcup_{A \in \mathcal{A}} A$ and $E_\mathcal{B} := \bigcup_{B \in \mathcal{B}} B$ of edges of K_n must be disjoint:

- Color an edge $e \in K_n$ *red* if $e \in E_\mathcal{A}$, and color e *blue* if $e \in E_\mathcal{B}$.

This way, the edges of every spanning tree $T \in \mathcal{R}$ will receive their colors. The remaining edges of K_n (if there are any) can be colored arbitrarily.

Recall that an edge e of K_n is *crossing* if $e \in U \times V$. Assume that at least half of the crossing edges is colored in *red*; otherwise, we can consider blue edges. This assumption implies that the set $E_{\mathrm{red}} \subseteq U \times V$ of red crossing edges has $|E_{\mathrm{red}}| \geq \frac{1}{2}|U \times V|$ edges. For a vertex $u \in U$, the set of its *good neighbors* is the set

$$V_u = \{v \in V : \{u, v\} \in E_{\mathrm{red}}\}$$

of vertices that are connected to u by red crossing edges. Claim 1 gives the following structural restriction on star factors covered by the rectangle $\mathcal{R}$.

Claim 2. *For any star factor $F \in \mathcal{F}_\mathcal{R}$, and for any center z of F, if $z \in U$, then $|F \cap (\{z\} \times V_z)| \leq m$.*

That is, if a star factor F is covered by the rectangle $\mathcal{R}$, then every star of F centered in some vertex $z \in U$ can only have m or fewer (out of all $|V_z|$ possible) red crossing edges.

Proof. Take a star factor $F \in \mathcal{F}_\mathcal{R}$ having some star whose center z belongs to U. Since F is covered by the rectangle $\mathcal{R}$, $F \subseteq A \cup B$ holds for some forests $A \in \mathcal{A}$ and $B \in \mathcal{B}$. By the definition of the edge-coloring, we have $B \cap (\{z\} \times V_z) = \emptyset$: all edges in $\{z\} \times V_z$ are red, while those in B are blue. So, all edges of $F \cap (\{z\} \times V_z)$ belong to the forest A, and Claim 1 yields $|F \cap (\{z\} \times V_z)| \leq |A \cap (\{z\} \times V_z)| \leq m$. $\square$

We call a vertex u of K_n *rich* if $u \in U$ and at least one quarter of the vertices in V are good neighbors of u, that is, if $|V_u| \geq \frac{1}{4}|V|$ holds. By (2), every rich vertex u has $|V_u| \geq n/16$ good neighbors. Split the family $\mathcal{F}_\mathcal{R}$ of star factors covered by the rectangle $\mathcal{R}$ into the family $\mathcal{F}_\mathcal{R}^1$ of star factors $F \in \mathcal{F}_\mathcal{R}$ with *no* rich center, and the family $\mathcal{F}_\mathcal{R}^2$ of all star factors $F \in \mathcal{F}_\mathcal{R}$ with *at least* one rich center. We will upper-bound the number of star factors in $\mathcal{F}_\mathcal{R}^1$ and in $\mathcal{F}_\mathcal{R}^2$ separately.

The intuition behind this splitting is that star factors $F \in \mathcal{F}_\mathcal{R}^1$ have the restriction (given by Claim 3 below) that only relatively "few" potential vertices of K_n can be used as centers of stars, while the restriction for the star factors $F \in \mathcal{F}_\mathcal{R}^2$ (given by Claim 2) is that at least one of its stars $S_z \subset F$ (centered in a rich center z) has relatively "few" potential vertices of K_n which can be taken as leaves.

To upper-bound $|\mathcal{F}_\mathcal{R}^1|$, let us first show that the set $U^* = \{u \in U : |V_u| \geq \frac{1}{4}|V|\}$ of all rich vertices is large enough.

Claim 3. *There are $|U^*| \geq \frac{1}{4}|U| \geq n/16$ rich vertices.*

Proof. The second inequality follows from (2). To prove the first inequality, assume contrariwise that there are only $|U^*| < \frac{1}{4}|U|$ rich vertices in U. Since $|V_u| < \frac{1}{4}|V|$ holds for every vertex $u \in U \setminus U^*$, we obtain

$$\tfrac{1}{2}|U \times V| \leq |E_{\mathrm{red}}| = \sum_{u \in U^*} |V_u| + \sum_{u \in U \setminus U^*} |V_u| < \tfrac{1}{4}|U| \cdot |V| + |U| \cdot \tfrac{1}{4}|V| = \tfrac{1}{2}|U \times V|\,,$$

a contradiction. $\square$

Each star factor in $\mathcal{F}_{\mathcal{R}}^1$ can be constructed in the same way as we constructed any star factor $F \in \mathcal{F}$ above (before (1)), with the difference that centers can only be chosen from $[n] \setminus U^*$, not from the entire set $[n]$. Thus,

$$\frac{|\mathcal{F}_{\mathcal{R}}^1|}{|\mathcal{F}|} \leq \binom{n - |U^*|}{m} \cdot \binom{n}{m}^{-1} \leq \mathrm{e}^{-|U^*| \cdot m/n} = 2^{-\Omega(m)}\,. \tag{3}$$

Here we used Claim 3 together with the second of the two simple inequalities holding for all $b \leq b + x < a$:

$$\left(\frac{a - b - x}{a - x}\right)^x \leq \binom{a - x}{b}\binom{a}{b}^{-1} \leq \left(\frac{a - b}{a}\right)^x\,. \tag{4}$$

To upper bound $|\mathcal{F}_{\mathcal{R}}^2|$, we will use the restriction given by Claim 1. Recall that every star factor $F \in \mathcal{F}_{\mathcal{R}}^2$ has at least one rich center. So, consider the following (nondeterministic) procedure of constructing a star factor F in $\mathcal{F}_{\mathcal{R}}^2$.

1. Choose a rich center $z \in U^*$; there are at most $|U^*| \leq |U| \leq n$ possibilities to do this.

2. For the center z, do the following:

   (a) choose a subset of $i \leq m$ vertices from the set V_z of all good neighbors of z, and connect these vertices to z by (crossing) edges; for each $i \leq m$ there are $\binom{|V_z|}{i}$ possibilities.

   (b) choose a subset of $d - i$ vertices in $[n] \setminus (V_z \cup \{z\})$ and connect them to z; here we have at most $\binom{n - |V_z| - 1}{d - i} \leq \binom{n - |V_z|}{d - i}$ possibilities.

3. Choose a subset of $m - 1$ distinct centers from the remaining $n - d - 1$ vertices. There are at most $\binom{n - d - 1}{m - 1} \leq \binom{n - 1}{m - 1} = \frac{m}{n}\binom{n}{m}$ possibilities to do this.

4. Choose a partition of the remaining $n - m - d$ vertices into $m - 1$ blocks of size d, and connect the ith largest of the $m - 1$ chosen centers to all vertices in the ith block. There are at most $\binom{n - m - d}{d, \ldots, d} = \frac{(n - m - d)!}{d!^{m-1}}$ possibilities to do this.

Claim 4. *Every star factor $F \in \mathcal{F}_{\mathcal{R}}^2$ can be produced by the above procedure.*

Proof. Take a star factor $F \in \mathcal{F}_{\mathcal{R}}$ containing a star $S_z \subset F$ centered in a rich vertex $z \in U^*$. The star z can be picked by Step 1 of the procedure. By Claim 2, the star S_z can only have $i := |F \cap (\{z\} \times V_z)| \leq m$ good neighbors of z (those in V_z) as leaves, and Step 2(a) of our procedure can pick all these i leaves of S_z. The remaining $d - i$ leaves of the star S_z must belong to the set $[n] \setminus (V_z \cup \{z\})$. So, Step 2(b) can pick these $d - i$ leaves of S_z. Since the remaining two steps 3 and 4 of the procedure can construct *any* star factor of $K_n \setminus S_z$, the rest of the star factor F can be constructed by these steps. $\square$

7

The number of possibilities in Step 2 of our procedure is related to the probability distribution

$$h(K, n, d, i) := \Pr\{X = i\} = \frac{\binom{K}{i}\binom{n-K}{d-i}}{\binom{n}{d}}$$

of a hypergeometric random variable X: the probability of having drawn exactly i white balls, when drawing uniformly at random without replacement d times, from a vase containing K white and $n - K$ black balls. The number of possibilities in Step 2 of the procedure (for a center z picked in Step 1) is then at most $H(|V_z|, n, d, m) \cdot \binom{n}{d}$, where

$$H(K, n, d, m) := \Pr\{X \leq m\} = \sum_{i=0}^{m} h(K, n, d, i),$$

is the probability of having drawn at most m white balls. For fixed n, d and m, the function $H(K, n, d, m)$ is non-increasing in K, implying that the maximum of $H(|V_z|, n, d, m)$ over all rich centers $z \in U^*$ is achieved for $K := \min\{|V_z| : z \in U^*\}$. Hence, for every rich center $z \in U^*$, the number of possibilities in Step 2 is at most

$$H(|V_z|, n, d, m) \cdot \binom{n}{d} \leq H \cdot \binom{n}{d},$$

where $H := H(K, n, d, m)$. From the first inequality of (4) (applied with $x := m$, $a := n$ and $b := d$) we have $\binom{n}{d} \leq C \cdot \binom{n-m}{d}$, where $C = \left(\frac{n-m}{n-d-m}\right)^m \leq \exp\left(\frac{md}{n-d-m}\right)$ is a constant since $md = O(n)$ and $m, d = o(n)$.

Thus the total number of possibilities in all steps 1–4 and, by Claim 4, also the number $|\mathcal{F}_{\mathcal{R}}^2|$ of star factors in $\mathcal{F}_{\mathcal{R}}^2$, is at most a constant times

$$\underbrace{n \cdot H \cdot \binom{n-m}{d}}_{\text{Steps 1 and 2}} \underbrace{\frac{m}{n}\binom{n}{m}}_{\text{Step 3}} \underbrace{\frac{(n-m-d)!}{d!^{m-1}}}_{\text{Step 4}} = m \cdot H \cdot \underbrace{\binom{n}{m}\frac{(n-m)!}{d!^m}}_{= |\mathcal{F}|}.$$

Known tail inequalities for the hypergeometric distribution (see Hoeffding [9], or Chvátal [2] for a direct proof) imply that, if $m \leq (K/n - \epsilon)d$ for $\epsilon > 0$, then

$$H(K, n, d, m) = \Pr\{X \leq m\} \leq e^{-2\epsilon^2 d}. \tag{5}$$

Remark. In both papers [9] and [2], this upper bound is only stated for the event $X \geq (K/n + \epsilon)d$, but using the duality $h(K, n, d, i) = h(n - K, n, d, d - i)$ (count black balls instead of white), the same upper bound holds also for the event $X \leq (K/n - \epsilon)d$.

In our case, $K = \min\{|V_z| : z \in U^*\} \geq \frac{1}{4}|V| \geq n/16$. Recall that, so far, we have only used the conditions $(d + 1)m = n$ and $m = \Theta(\sqrt{n})$ on the parameters m and d. We now use the last condition $m \leq d/32$. For $\epsilon = 1/32$, we then have $m \leq d/32 \leq (K/n - \epsilon)d$, and (5) yields

$$H = H(K, n, d, m) \leq \Pr\{X \leq d/32\} \leq e^{-d/512} = 2^{-\Omega(d)}.$$

By taking $d := 6\sqrt{n}$ and $m := n/(d + 1)$, all three conditions on the parameters m and d are fulfilled, and we obtain $|\mathcal{F}_{\mathcal{R}}^2| \leq m|\mathcal{F}| \cdot 2^{-\Omega(d)}$. Together with the upper bound (3), the desired upper bound on $|\mathcal{F}_{\mathcal{R}}|$ follows:

$$\frac{|\mathcal{F}_{\mathcal{R}}|}{|\mathcal{F}|} = \frac{|\mathcal{F}_{\mathcal{R}}^1| + |\mathcal{F}_{\mathcal{R}}^2|}{|\mathcal{F}|} \leq 2^{-\Omega(m)} + m2^{-\Omega(d)} = 2^{-\Omega(\sqrt{n})}. \qquad \square$$

References

[1] Bellman, R., 1958. On a routing problem. Quarterly of Appl. Math. 16, 87–90.

[2] Chvátal, V., 1979. The tail of the hypergeometric distribution. Discrete Math. 25, 285–287.

[3] Dreyfus, S., Wagner, R., 1971. The Steiner problem in graphs. Networks 1 (3), 195–207.

[4] Edmonds, J., 1967. Optimum branchings. J. of Res. of the Nat. Bureau of Standards 71B (4), 233–240.

[5] Edmonds, J., 1973. Edge-disjoint branchings. In: Rustin, B. (Ed.), Combinatorial Algorithms. Academic Press, pp. 91–96.

[6] Floyd, R., 1962. Algorithm 97, shortest path. Comm. ACM 5, 345.

[7] Ford, L., 1956. Network flow theory. Tech. Rep. P-923, The Rand Corp.

[8] Held, M., Karp, R., 1962. A dynamic programming approach to sequencing problems. SIAM J. on Appl. Math. 10, 196–210.

[9] Hoeffding, W., 1963. Probability inequalities for sums of bounded random variables. J. of the Amer. Statistical Association 58 (301), 13–30.

[10] Hyafil, L., 1979. On the parallel evaluation of multivariate polynomials. SIAM J. Comput. 8 (2), 120–123.

[11] Jerrum, M., Snir, M., 1982. Some exact complexity results for straight-line computations over semirings. J. ACM 29 (3), 874–897.

[12] Jukna, S., 2015. Lower bounds for tropical circuits and dynamic programs. Theory of Comput. Syst. 57 (1), 160–194.

[13] Kruskal, J., 1956. On the shortest spanning subtree of a graph and the traveling salesman problem. Proc. of AMS 7, 48–50.

[14] Levin, A., 1971. Algorithm for the shortest connection of a group of graph vertices. Sov. Math. Dokl. 12, 1477–1481.

[15] Moore, E., 1957. The shortest path through a maze. In: Proc. Internat. Sympos. Switching Theory. Vol. II. pp. 285–292.

[16] Tarjan, R., 1977. Finding optimum branchings. Networks 7 (1), 25–35.

[17] Valiant, L., 1980. Negation can be exponentially powerful. Theor. Comput. Sci. 12, 303–314.

[18] Warshall, S., 1962. A theorem on boolean matrices. J. ACM 9, 11–12.