

Greedy matroid base packings with applications to dynamic graph density and orientations

Pavel Arkhipov Vladimir Kolmogorov
 Institute of Science and Technology Austria
 {pavel.arkhipov,vnk}@ist.ac.at

Abstract

Greedy minimum weight spanning tree packings have proven to be useful in connectivity-related problems. We study the process of greedy minimum weight base packings in general matroids and explore its algorithmic applications.

When specialized to bicircular matroids, our results yield an algorithm for the approximate fully-dynamic densest subgraph density ρ . We maintain a $(1 + \varepsilon)$ -approximation of the density with a *worst-case* update time $O((\rho\varepsilon^{-2} + \varepsilon^{-4})\rho \log^3 m)$. It improves the dependency on ε from the current state-of-the-art worst-case update time complexity $O(\varepsilon^{-6} \log^3 n \log \rho)$ [Chekuri, Christiansen, Holm, van der Hoog, Quanrud, Rotenberg, Schwiegelshohn, SODA’24]. We also can maintain an implicit fractional out-orientation with a guarantee that all out-degrees are at most $(1 + \varepsilon)\rho$, with the *worst-case* update time $O(\varepsilon^{-4}\rho^2 \log^3 m)$. For any edge, the query time of its fractional orientation is $O(\varepsilon^{-2}\rho \log^2 m)$. In contrast, existing algorithms for (integer) dynamic out-orientation guarantee out-degree of at most $(1 + \varepsilon)\alpha$ where α is the fractional arboricity of the graph (and can be $\alpha = (1 + \Theta(1))\rho$ in the case of multigraphs).

Our algorithms above work by greedily packing pseudoforests, and require maintenance of a minimum-weight pseudoforest in a dynamically changing graph. We show that the latter problem can be solved in $O(\log n)$ worst-case time per edge insertion or deletion.

For general matroids, we observe two characterizations of the limit of the base packings (“the vector of ideal loads”). Specialized to graphic matroids, it implies the characterizations from [Cen, Fleischmann, Li, Li, Panigrahi, FOCS’25], namely, their entropy-minimization theorem and their bottom-up cut hierarchy.

Finally, we give combinatorial results on the greedy tree packings. We show that a tree packing of $O(\lambda^5 \log m)$ trees contains a tree crossing some min-cut once, which improves the bound $O(\lambda^7 \log^3 m)$ from [Thorup, Combinatorica’07]. We also strengthen the lower bound on the edge load convergence rate from [de Vos, Christiansen, SODA’25], showing that Thorup’s upper bound is tight up to a logarithmic factor.

Contents

1	Introduction	2
1.1	Greedy tree packings	3
1.2	Dynamic density and out-orientations	4
1.3	Summary of our results	4
2	General matroids	6
2.1	Background: Frank-Wolfe algorithm	6
2.2	Two characterizations of the limit of the greedy base packing	6
2.3	Convergence of $(\min_e x_e^k)^{-1}$ to ρ	8

3	Bicircular matroids: density and fractional orientations	9
3.1	Dynamic approximate density	9
3.2	Dynamic approximate fractional out-orientation	10
4	Graphic matroids: 1-respecting min-cuts and a lower bound on the convergence	11
4.1	A packing of $\Theta(\lambda^5 \log m)$ trees contains a tree crossing some min-cut once	11
4.2	A tight lower bound for $\ x^k - x^*\ _\infty$	15
5	Other observations	19
5.1	Theoretical observations	19
5.2	Experimental observations for tree packings	25
A	Proof of Theorem 2.4	25
B	Proof of Theorem 2.6	26
C	Dynamic minimum-weight pseudoforest	28

1 Introduction

We consider a greedy algorithm for packing bases in a matroid $M = (E, \mathcal{I})$ with $m = |E|$ elements that builds a sequence of bases $B_1, B_2, \dots$ such that B_k is a minimum-weight base with respect to the weights $\xi_{B_1} + \dots + \xi_{B_{k-1}}$. (Here $\xi_B \in \{0, 1\}^E$ is the indicator vector of a base B). For the graphic matroid this process is known as *greedy (spanning) tree packing*; it has many applications, and has been widely studied in the literature.

We investigate greedy packing in general matroids, and establish several new results. First, we observe that vectors $x^k := \frac{1}{k}(\xi_{B_1} + \dots + \xi_{B_k})$ converge to a certain vector x^* as $k \rightarrow +\infty$.¹ Following the terminology used in the context of spanning trees [26], we call the vector x^k the *relative loads* at step k , and x^* an *ideal loads vector*. An important property of this vector is that its minimum equals the inverse of the matroid density:

$$\frac{1}{\min_{e \in E} x_e^*} = \rho, \quad \rho = \rho(M) = \max_{\emptyset \neq H \subseteq E} \frac{|H|}{r(H)} \quad (1)$$

where $r(\cdot)$ is the rank function of the matroid. We specialize this result to two cases: the cyclic (or graphic) matroid and the bicircular matroid of an undirected graph $G = (V, E)$. In the former case (1) reduces to

$$\frac{1}{\min_{e \in E} x_e^*} = \alpha, \quad \alpha = \alpha(G) = \max_{S \subseteq V, |S| > 1} \frac{|E(S)|}{|S| - 1} \quad (2)$$

where α is known as the *fractional arboricity* of G . This relation has been recently proved in [10, Theorem 25] using graph-theoretic arguments specific to graphic matroids.

Let us now consider the bicircular matroid of the graph G . Recall that its independent sets are pseudoforests, i.e. subsets $H \subseteq E$ such that each connected component has at most one cycle. If G has at least one cycle then eq. (1) becomes

$$\frac{1}{\min_{e \in E} x_e^*} = \rho, \quad \rho = \rho(G) = \max_{\emptyset \neq S \subseteq V} \frac{|E(S)|}{|S|} \quad (3)$$

The quantity $\rho = \rho(G)$ is known as the *density* of G , and the set S maximizing the ratio in (3) is a *densest subgraph*. Using relation (3), we present a new algorithm for computing approximate

¹In literature, x^k and x^* are usually denoted by ℓ and ℓ^* respectively. We find the notation “ x ” to be more natural, since we will consider the loads vector in the context of Frank-Wolfe algorithm.

density in a dynamic graph, and also for a closely related problem of computing an approximate fractional orientation minimizing the maximum out-degree.

In Sections 1.1 and 1.2 we discuss these two special cases and related applications in more detail, and describe our contributions.

1.1 Greedy tree packings

Greedy minimum spanning tree (MST) packings turn out to be useful for cut algorithms. Karger gave an almost linear time algorithm for computing a minimum cut [19] using the greedy packing of a logarithmic number of MSTs. The key property is that a constant fraction of the trees cross some minimum cut at most twice. Given a tree, it is possible to quickly find the minimum cut that gets crossed at most twice, giving the desired algorithm.

Another use of the greedy MST packing are dynamic min-cut algorithms. Thorup described an algorithm for dynamically maintaining edge connectivity of a graph, up to a polylogarithmic value [26]. He used a greedy MST packing of $\Theta(\lambda^7 \log^3 m)$ trees, where λ is the graph's edge connectivity upper bound, updating the packing dynamically as the graph gets updated. The crucial property of the Thorup's MST packing is that it is guaranteed to contain a tree that crosses some minimum cut exactly once. The number of trees in the packing is somewhat high, but still polylogarithmic, if the value of the mincut is polylogarithmically bounded. Recently, Christiansen and de Vos showed that it is enough to maintain only $\Theta(\lambda^3 \log m)$ trees for the purposes of dynamic min-cut [10]. They show that they *either* have some tree that crosses some min-cut once *or* have some min-cut in the family of trivial cuts in some contracted graph.

We strengthen Thorup's bound as follows.

Theorem 1.1. *A greedy tree packing with $k = \Theta(\lambda^5 \log m)$ trees contains a tree crossing some min-cut only once.*

Given the result in [10], this does not lead to an improvement in the runtime for dynamic min-cut, but nonetheless shows an interesting structural property of greedy MST packings.

Greedy MST packings are also useful for the minimum k -way cut problem [27] and the network unreliability problem [6].

For the algorithms in [26], [10], [27], [6], it is crucial to bound the error between the ideal edge loads and the loads that one gets after greedily packing k MSTs. The following bound is due to Thorup ([26], Proposition 16, restated equivalently):

$$\|x^k - x^*\|_\infty \leq \sqrt{\frac{6 \log m}{k \lambda}}, \quad (4)$$

where x^* is the (m -dimensional) vector of ideal edge loads, x^k is the vector of edge loads after packing k MSTs, and λ is an upper bound on the edge connectivity of G .

One may ask how tight this bound is in the worst case. There exists an adversarial example ([10], Theorem 7, restated equivalently) of an MST packing that witnesses the error of the form

$$\|x^k - x^*\|_\infty \geq \Omega\left(\frac{1}{k^{2/3} \lambda^{1/3}}\right), \quad (5)$$

for the range of k such that $k/\lambda \leq O(\sqrt{n})$.

The gap between $\sim \frac{1}{\sqrt{k}}$ and $\sim \frac{1}{k^{2/3}}$ in the behavior of the ℓ_∞ -error was open. In this paper we close this gap by giving the following lower bound.

Theorem 1.2. *There exists a family of graphs in which $\|x^k - x^*\|_\infty \geq \Omega\left(\sqrt{\frac{1}{k \lambda}}\right)$ for $k = O(n^2)$.*

We also analyze the error bound for p -norms $\|x^k\|_p - \|x^*\|_p$ and show the upper bound of $\sim \frac{\log k}{k}$ for fixed p and for large enough k .

1.2 Dynamic density and out-orientations

Given a dynamic graph that can undergo edge insertions and deletions, one may ask for maintaining an approximation for the density of the densest subgraph ρ , defined as $\rho := \max_{S \subseteq V} \frac{|E[S]|}{|S|}$. This problem has been considered in [3, 12, 25, 7, 8]. In particular, the last three papers presented algorithms for maintaining $(1 + \varepsilon)$ -approximation of ρ with the following complexities for updates: $O(\varepsilon^{-6} \log^4 n)$ worst-case [25], $O(\varepsilon^{-6} \log^3 n \log \rho)$ worst-case or $O(\varepsilon^{-4} \log^2 n \log \rho)$ amortized [7], and $O(\varepsilon^{-6} \log^3 n \log \log n)$ worst-case or $O(\varepsilon^{-4} \log^2 n)$ amortized, with improved space requirements [8]. In this paper we present an algorithm with a better dependency on ε for the worst-case update time.

Theorem 1.3. *There exists a data structure that maintains an $(1 + \varepsilon)$ -approximation to the densest subgraph density ρ for a fully dynamic graph with worst-case update time $O((\rho \varepsilon^{-2} + \varepsilon^{-4}) \rho \log^3 m)$.²*

The density of the densest subgraph is closely related to out-orientations. Given an undirected graph, we wish to orient its edges to minimize the maximal out-degree of a vertex. It is a classical fact that the optimal value of the maximal out-degree equals $\lceil \rho \rceil$.³ If one allows orienting the edges fractionally, then the relation to the maximum density is even more direct: the optimal value of the maximal out-degree equals ρ . A fractional orientation is a collection of values $d_{u \rightarrow v} \in [0, 1]$ and $d_{v \rightarrow u} \in [0, 1]$ for any edge $(u, v) \in E$, satisfying $d_{u \rightarrow v} + d_{v \rightarrow u} = 1$.

Existing dynamic algorithms for out-orientations such as [17, 20, 2, 25, 9, 7] give approximation guarantees in terms of the fractional arboricity $\alpha := \max_{S \subseteq V, |S| > 1} \frac{|E[S]|}{|S| - 1} > \rho$. For example, [7] presents an algorithm for maintaining a dynamic orientation such that every vertex has out-degree at most $(1 + \varepsilon)\alpha + 2$, and the update time is (worst-case) $O(\varepsilon^{-6} \log^3 n \log \alpha)$. If graph G is simple then values α and ρ are close if they are large. However, in multigraphs α and ρ can differ by a factor of $1 + \Theta(1)$, even if they are large.

By greedily packing pseudoforests we will obtain

Theorem 1.4. *There exists a data structure for a fully dynamic graph with worst-case update time $O(\varepsilon^{-4} \rho^2 \log^3 m)$ that maintains an implicit approximate minimum fractional out-orientation. The query for a fractional orientation of an edge takes $O(\varepsilon^{-2} \rho \log^2 m)$ time. It is guaranteed that the out-degree of any vertex is at most $(1 + \varepsilon)\rho$.*

This result is not directly comparable to previous work on dynamic out-orientations: we obtain a better approximation guarantee (in terms of ρ , not α), but relax the set of solutions (since we compute a *fractional* out-orientation).

1.3 Summary of our results

The main contribution of this paper is a systematic study of greedy base packing in matroids. It turns out that the packing algorithm can be viewed as instance of the Frank-Wolfe algorithm applied to the problem of minimizing the norm over $\|x\|^2$ over the base polytope P_B of a given matroid; this was recently observed in [16] for graphic matroids. Using this connection, we give a simple proof of the following result that characterizes the ideal loads vector.

Theorem 1.5. *There exists a vector $x^* = \lim_{k \rightarrow +\infty} x^k$. This vector minimizes $\sum_e \varphi(x_e)$ over P_B for any convex function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$. In particular, (1) x^* is the projection of zero onto P_B , i.e. the point in P_B that minimizes $\|x\|_2$; (2) x^* minimizes the p -norm of x over P_B , for any $1 < p < \infty$.*

²When stating complexities, we assume that ρ is a known upper bound on the true value of the densest subgraph density. **Everywhere else ρ denotes the true density.**

³It can be viewed as a special case of the theorem by Edmonds [11] saying that the minimum number of independent sets that can cover a matroid equals the maximal density of the matroid rounded up.

Note, it was recently proved in [5, Theorem 6.3] that x^* maximizes the entropy $H(x) = \log(n+1) - \frac{1}{n+1} \sum_e x_e \log x_e$ over P_B (in the case of graphic matroids). This follows immediately from the result above, since the function $\varphi(t) = t \log t$ is convex.

We give convergence bounds for the sequence of edge loads $x^k \rightarrow x^*$, in terms of objective functions $\|x\|_2$ and $\|x\|_p$:

Theorem 1.6. *After greedily packing k trees, we have*

$$\|x^k\|_2 - \|x^*\|_2 \leq \sqrt{m} \cdot \frac{\log(k+1)}{k} \quad (6)$$

$$\|x^k\|_p - \|x^*\|_p \leq \frac{p}{2} \cdot m^{1/p} \frac{\log k}{k} (1 + o(1)) \quad \text{for any fixed integer } p \geq 2 \quad (7)$$

where the $o(1)$ function corresponds to $k \rightarrow \infty$.

Using a generic analysis of the Frank-Wolfe method, [16] recently showed that $\|x^k - x^*\|_2 \leq \varepsilon$ after $k = \Theta(\frac{m \log(m/\varepsilon)}{\varepsilon^2})$ steps. This can be seen as a corollary of the above theorem (see Corollary 5.2 and the remark after it).

We also establish how fast $(\min_e x_e^k)^{-1}$ converges to the density ρ of the given matroid.

Theorem 1.7. *If $\varepsilon \in (0, 1]$ and $k \geq \frac{20\rho \log m}{\varepsilon^2}$ then*

$$0 \leq \frac{1}{\min_e x_e^k} - \rho \leq \varepsilon \rho \quad (8)$$

In fact, we establish a similar bound for a more general version of the packing algorithm that removes elements with large values of x_e^k ; this version will be used for improving the complexity of a dynamic algorithm for approximate graph density.

Notice that, when specialized to graphic matroids, the above convergence bound does not explicitly depend on the edge-connectivity λ , unlike the known convergence bound of Thorup ([26], Proposition 16). This immediately strengthens Lemma 26 from [10], which states that one needs $\Theta\left(\frac{\alpha^2 \log m}{\lambda \varepsilon^2}\right)$ trees to $(1 + \varepsilon)$ -approximate the fractional arboricity α , shaving off a factor of $\alpha/\lambda \geq \Theta(1)$.

Next, we summarize our contributions related to the two special cases. For graphic matroids, we establish two new results:

- Improving Thorup's bound on the number of MSTs containing a 1-respecting tree from $\Theta(\lambda^7 \log^3 m)$ to $\Theta(\lambda^5 \log m)$ (Theorem 1.1).
- Giving a lower bound on the rate of convergence of the ∞ -norm that almost matches the upper bound in [26] (Theorem 1.2).

As for bicircular matroids, we present algorithms for two dynamic problems:

- Approximating the density of the densest subgraph in a dynamic graph (Theorem 1.3) with improved dependence on ε of worst-case update times.
- Minimum fractional out-orientation (Theorem 1.4) with better approximation guarantee compared to previous dynamic algorithms (in terms of ρ , not α).

Both algorithms rely on maintaining a minimum-weight pseudoforest in dynamically changing graphs (with edge insertions and deletions). In Appendix C we show that this can be done with $O(\log n)$ worst-case time per update, which we believe can be of independent interest. We contrast this with the problem of dynamically maintaining a minimum-weight spanning tree, for which only $O(\log^4 n)$ amortized [18] or Las-Vegas $O(n^{o(1)})$ worst-case (w.h.p.) [22] update times are known.

2 General matroids

Throughout this section, we fix matroid $M = (E, \mathcal{I})$ with the rank function $r(\cdot)$. $P_B \subset \mathbb{R}^E$ denotes the base polytope of this matroid, i.e. the convex hull of bases of M . It is known that

$$P_B = \{x \in \mathbb{R}_{\geq 0}^E \mid x(E) = r \text{ and } x(H) \leq r(H) \quad \forall H \subseteq E\}, \quad r = r(E). \quad (9)$$

$B_1, B_2, \dots$ will be the sequence of bases in greedy base packing, and $x^k = \frac{1}{k}(\xi_{B_1} + \dots + \xi_{B_{k-1}})$. Thus, B_k is a minimum-weight base in M with respect to weights x^{k-1} , where we take $x^0 = \mathbf{0}$.

As observed in [16] for graphic matroids, greedy base packing is an instance of the Frank-Wolfe algorithm. We review this method in the next subsection.

2.1 Background: Frank-Wolfe algorithm

This algorithm dates back to [13]. It is a simple iterative process for optimizing a convex differentiable function $f : P \rightarrow \mathbb{R}$ over a convex compact set $P \subseteq \mathbb{R}^d$, and works as follows.

1. Start from any point $x^0 \in P$.
2. Keep doing the following:
 - (a) Choose $v^k \in \arg \min_{v \in P} \langle \nabla f(x^k), v \rangle$
 - (b) $x^{k+1} = (1 - \eta_k)x^k + \eta_k v^k$

Here, η_k is the step size parameter. Popular choices are setting $\eta_k = \frac{2}{k+2}$ or doing a line search for η_t that minimizes $f(x^{k+1}) = f((1 - \eta_k)x^k + \eta_k v^k)$. As an output, we get a sequence of points $x^1, x^2, \dots$. If f is strictly convex, then the sequence x^k converges to the optimum $\arg \min_{x \in P} f(x)$.

There are lots of possible modifications of the Frank-Wolfe algorithms. For an introduction and surveys, we redirect the reader to [24], [21], [4].

The Frank-Wolfe method is attractive, because it does not require computing projections onto the feasible region, which could be tricky in general. We only need to be able to compute the gradient of f and solve the linear program when computing v^k . However, the vanilla Frank-Wolfe algorithm has a somewhat bad convergence rate of $O(1/k)$. More precisely, choosing $\eta_k = \frac{2}{k+2}$ for an L -smooth convex function f yields the following upper bound on the error:

$$f(x^k) - f(x^*) \leq \frac{2LD^2}{k+2}, \quad (10)$$

where x^* is the minimum point, D is the diameter of the feasible region P (see, for example, Theorem 4.4 in [24]). A function f is L -smooth if for all $x, y \in P$, we have $f(y) - f(x) \leq \langle \nabla f(x), y - x \rangle + \frac{L}{2} \|y - x\|^2$.

We will be interested in the Frank-Wolfe algorithm with the step size $\eta_k = \frac{1}{k+1}$. One can observe that this corresponds to simple averaging: $x^k = (v^1 + \dots + v^k)/k$. This step size rule is less popular than the choice $\eta_k = \frac{2}{k+2}$, but it has also been studied and known to yield convergence at rate $O\left(\frac{\log k}{k}\right)$ ([14], Section 2.3).

2.2 Two characterizations of the limit of the greedy base packing

We start with the following result.

Lemma 2.1. *There exists a limit $x^* \in \mathbb{R}^E$ of the sequence $x^1, x^2, \dots$ of relative loads in a greedy base packing. Moreover, this limit is the projection of zero onto the base polytope:*

$$\lim_k x^k = x^* = \operatorname{argmin}_{x \in P_B} \|x\|_2. \quad (11)$$

Proof. Let v^k be the indicator vector of B_{k+1} . It is a minimum weight base for the weights x^k , therefore

$$\langle v^k, x^k \rangle = \min_{v: \text{vertex of } P_B} \langle v, x^k \rangle = \min_{v \in P_B} \langle v, x^k \rangle. \quad (12)$$

Define a (strongly convex) function $f(x) = \|x\|_2^2$. Observe that $\nabla f(x) = 2x$. Then,

$$\langle v^k, \nabla f(x^k) \rangle = \min_{v \in P_B} \langle v, \nabla f(x^k) \rangle. \quad (13)$$

For the sequence of x^k , we have the update rule

$$x^{k+1} = \frac{k}{k+1}x^k + \frac{1}{k+1}v^k. \quad (14)$$

This is precisely the Frank-Wolfe process for the function f and the feasible set P_B with the step size $\eta_k = \frac{1}{k+1}$. It is known to converge to the minimum of f in P_B , which implies the lemma. $\square$

It is also not hard to prove the convergence by hand, without relying on the Frank-Wolfe results. For example, see the explicit convergence bound in Lemma 1.6 derived in an elementary way.

In fact, the limit x^* does not only minimize $\|x\|_2$ over P_B , but it also minimizes $\sum_e \varphi(x_e)$ over P_B for any convex function $\varphi : \mathbb{R} \rightarrow \mathbb{R}$.

Observation 2.2. *Let $f : \mathbb{R}^E \rightarrow \mathbb{R}$ be a function such that $f(x) = \sum_e \varphi(x_e)$, where $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly convex differentiable function. Then, for any i , B_i is also a minimum weight base with respect to the weights $\nabla f(x^k)$.*

Proof. The minimum weight base of a matroid is defined only by the ordering of the weights. If $x_e > x_f$, then $\varphi'(x_e) > \varphi'(x_f)$ for a strictly convex differentiable φ , so the ordering of the values of $\nabla f(x)$ is the same as the ordering of the values of x . $\square$

Corollary 2.3. *Let $f : \mathbb{R}^E \rightarrow \mathbb{R}$ be a function such that $f(x) = \sum_e \varphi(x_e)$, where $\varphi : \mathbb{R} \rightarrow \mathbb{R}$ is a strictly convex differentiable function. For a greedy base packing, the sequence $x^1, x^2, \dots$ converges to $\operatorname{argmin}_{x \in P_B} f(x)$.*

Proof. The observation above implies that the greedy base packing is a Frank-Wolfe process for the minimization of f over P_B . For a strictly convex function, the Frank-Wolfe iterations converge to the (unique) optimum. $\square$

It remains to cover the cases where φ is convex, but not *strictly* convex, and not necessarily differentiable. In this case, one can approximate φ with a strictly convex differentiable function arbitrarily well with respect to the functional C -norm, and use the convergence result for the approximation. Together, the observations above yield the characterization given in Theorem 1.5.

There are several well-known characterizations of the minimum-norm vector in the base polytope of a matroid, see e.g. [15, Chapter V]. For the proofs below we will need the following version. Recall that for a subset $H \subseteq E$, the *contraction* of H gives a new matroid M' on the ground set $E \setminus H$ and with a rank function $r'(A) = r(A \cup H) - r(H)$ for any set $A \subseteq E \setminus H$. If B_H is a base of the restriction matroid $M|_H$, then the bases of M/H are precisely such sets $B' \subseteq E \setminus H$ that $B' \cup B_H$ is a base of M (see, for example, the book [23], Chapter 3.1 “Contraction”). We will also use the identity $(M/H_1)/H_2 = M/(H_1 \cup H_2)$ for any disjoint $H_1, H_2 \subseteq E$. The following characterization of x^* as a minimum-norm vector in the base polytope appears in the book [15], Chapter V, Section 9 “Lexicographically Optimal Base”. For completeness, we give a self-contained proof in Appendix A.

Theorem 2.4 (Theorem 9.3 in [15], restated: characterization of $x^* = \operatorname{argmin}_{x \in P_B} \|x\|^2$ for general matroids). *While x^* is not defined for some elements, keep doing the following.*

1. Let $H = \operatorname{argmax}_{\emptyset \neq H \subseteq E} \frac{|H|}{r(H)}$. Moreover, choose H to be inclusion maximal.
2. Set $x_e^* = \frac{r(H)}{|H|}$ for every element $e \in H$.
3. Contract H .

Furthermore, the values $x_e^* = \frac{r(H)}{|H|}$ strictly increase during this process.

Recall that the maximal density⁴ of a matroid is defined as

$$\rho := \max_{\emptyset \neq H \subseteq E} \frac{|H|}{r(H)}. \quad (15)$$

From the monotonicity property in Theorem 2.4 we immediately get a connection between the maximal density and x^* .

Corollary 2.5. *The maximal density of a matroid is the inverse of the minimal x_e^* :*

$$\rho = \frac{1}{\min_e x_e^k}. \quad (16)$$

2.3 Convergence of $(\min_e x_e^k)^{-1}$ to ρ

Next, we will study the convergence of $(\min_e x_e^k)^{-1}$ in greedy base packing to ρ . In fact, we will consider a more general algorithm: *greedy base packing with pruning* that removes elements with large values of x_e^k . (This will be needed later on to improve the complexity of the algorithm for density estimation).

The algorithm is defined as follows. Given a matroid $M = (E, \mathcal{I})$ and input interval $[\rho^-, \rho^+]$, it produces sequences of sets $E = E_1 \supseteq E_2 \supseteq \dots$ and $B_1 \subseteq E_1, B_2 \subseteq E_2, \dots$ such that B_k is a base in the matroid $M_k = M|_{E_k}$ (the restriction of M to E_k). For $k \geq 1$ define vector $x^k \in [0, 1]^{E_k}$ via $x_e^k = \frac{1}{k} |\{B_j \mid e \in B_j, j \leq i\}|$ (and set $x^0 = \mathbf{0}$). The algorithm does the following for $k = 1, 2, \dots$:

- Let B_k be a minimum-weight base in M_k with respect to weights x^{k-1} .
- Let $E_{k+1} = E_k - \{e \in E_k : x_e^k > 2/\rho^- \text{ and if } k \geq 24\rho^+ \log m\}$.

Note, if $[\rho^-, \rho^+] = [0, +\infty)$ then this is equivalent to standard greedy packing without pruning.

To analyze this process, we will use an argument which is conceptually similar to Thorup's proof of [26, Claim 16.1]. It is also similar to the analysis of the Multiplicative Weights Update (MWU) algorithm (and yields the bound of $\Theta\left(\frac{\rho \log m}{\varepsilon^2}\right)$ iterations, which strongly resembles the bounds one gets in MWU type of algorithms). The following theorem is proved in Appendix B.

Theorem 2.6. *Suppose that $\rho = \rho(M) \in [\rho^-, \rho^+]$. After k iterations of greedy base packing with pruning we have $\rho(M_k) = \rho$. Furthermore, if $\varepsilon \in (0, 1]$ and $k \geq \frac{20\rho \log m}{\varepsilon^2}$ then*

$$0 \leq \frac{1}{\min_{e \in E_k} x_e^k} - \rho \leq \varepsilon \rho. \quad (17)$$

⁴If our matroid M has loops, then we can get zero in the denominator in the expression for density. In this case, the claims from this section work if one treats ρ as infinity. An alternative way is to forbid M to have loops.

3 Bicircular matroids: density and fractional orientations

Throughout this section, we will be dealing with the bicircular matroid of a graph $G = (V, E)$. The ground set of this matroid is E , and the independent sets are edge sets of all possible pseudoforests, i.e. such subgraphs that every connected component has at most one cycle. A base of the bicircular matroid is a maximal pseudoforest P . If G has no acyclic components then P can be characterized as follows: P contains $n = |V|$ edges, and every component of P contains exactly one cycle. The rank function of the bicircular matroid is $r(H) = |V(H)| - a(H)$, where $a(H)$ is the number of acyclic components of H , and $V(H)$ is the set of vertices adjacent to some edge in H .

Recall that the density of a graph is defined as $\rho(G) = \max_{\emptyset \neq S \subseteq V} \frac{|E(S)|}{|S|}$, and the subgraph $G[S]$ induced by a maximizer S of this ratio is a *densest subgraph* of G . This is related to the density of the bicircular matroid M as follows.

Observation 3.1. *If G is not a forest then $\rho(G) = \rho(M)$. Otherwise, if G is a forest, then a largest component of G is a densest subgraph.*

Proof. Let $S \subseteq V$ be the vertex set of the densest subgraph. If G has at least one cycle, then the densest subgraph has no acyclic components, and therefore, has rank $r(E[S]) = |S|$. It is straightforward to check that $E[S]$ witnesses the maximum of $\frac{|H|}{r(H)}$ over $H \subseteq E$. $\square$

It is easy to design an efficient dynamic algorithm that checks whether G is a forest, and if yes computes various quantities for it (using, for example, top trees data structure [1] discussed in Appendix C). Below we will often assume for simplicity that G is not a forest, and write just $\rho = \rho(G) = \rho(M)$.

3.1 Dynamic approximate density

With the above discussion, we can easily get a weaker version of Theorem 1.3.

Lemma 3.2 (Warm-up for dynamic density). *There exists a data structure that maintains an $(1 + \varepsilon)$ -approximation to the densest subgraph density ρ for a fully dynamic graph with worst-case update time $O(\varepsilon^{-4} \rho^2 \log^3 m)$.*

Proof. Maintain a greedy packing of $k = \Theta\left(\frac{\rho \log m}{\varepsilon^2}\right)$ maximal pseudoforests, so that $0 \leq \frac{1}{\min_e x_e^k} - \rho \leq \varepsilon \rho$ by Theorem 2.6. By a standard argument [28, Lemma 1], a single graph update can trigger $O(k^2)$ updates in the pseudoforests. Each update in a pseudoforest can be done in $O(\log n)$ worst-case time by Theorem C.1 (Appendix C). Thus, the update is done in time $O(k^2 \log n) = O(\varepsilon^{-4} \rho^2 \log^3 m)$.

For an estimator of ρ , output $\frac{1}{\min_e x_e^k}$, unless G is a forest. If G is a forest, output $\frac{N-1}{N}$, where N is the size of the largest component of G . $\square$

In the rest of this subsection, we shave off the factor ρ from the update time when $\varepsilon = O(1/\rho^2)$.

Theorem 3.3. *There exists a data structure that maintains an $(1 + \varepsilon)$ -approximation to the densest subgraph density ρ for a fully dynamic graph with worst-case update time $O((\rho \varepsilon^{-2} + \varepsilon^{-4}) \rho \log^3 m)$.*

To prove this result, we will use an approach which is similar in spirit to that in [10] (the latter improved the recourse for greedy packing of spanning trees). The main idea comes from the following observation: each edge e is present in kx_e^k pseudoforests, so adding / removing e may trigger at most $O(kx_e^k \cdot k)$ updates in the pseudoforests. If $x_e^k \sim 1/\rho$ then this improves the bound on the number of updates by a factor of $1/\rho$.

Recall that edges e of the densest subgraph have $x_e^* = 1/\rho$, and for other edges the value x_e^* can be very large, e.g. 1. We will use greedy packing with pruning that removes edges with large value of x_e^* .

Assume that we are given an upper bound $\rho_{\max}$ on the density; the estimation algorithm described below will be correct for a current graph if $\rho \leq \rho_{\max}$. Define $\rho_i = 2^{i-1}$ for $i = 1, \dots, r$ where $r = \lceil \log_2 \rho_{\max} \rceil + 2$. We will independently run r greedy packing algorithms with pruning, where for the i -th run we use interval $[\rho_i, \rho_{i+2}]$ and maintain $\Theta\left(\frac{\rho_i \log m}{\varepsilon^2}\right)$ maximal pseudoforests.

With this rule, at least one of the runs will give a correct estimate of ρ . To determine which one, we will also run standard greedy packing without pruning with $\Theta(\rho \log m)$ pseudoforests (using accuracy $\varepsilon = \Theta(1)$) so that it finds an interval $[\hat{\rho}, 2\hat{\rho}]$ containing ρ . We can then find i with $[\hat{\rho}, 2\hat{\rho}] \subseteq [\rho_i, \rho_{i+2}]$, and output the estimate of the i -th run.

Let us analyze the complexity of updating the i -th run during graph updates. We claim that each edge e can appear in at most $O((2^i + \frac{1}{\varepsilon^2}) \log m)$ pseudoforests. Indeed, $O(2^i \log m)$ is the maximum number of pseudoforests constructed at iterations with $k < 24\rho_{i+2} \log m$. By construction, if e appears in a pseudoforest at iteration $k \geq 24\rho_{i+2} \log m$ then it satisfies $x_k^e \leq 2/\rho_i$. Thus, the total number of such appearances is bounded by $\Theta\left(\frac{\rho_i \log m}{\varepsilon^2}\right) \cdot (2/\rho_i) = O\left(\frac{\log m}{\varepsilon^2}\right)$. The claim follows.

This leads to at most $O((2^i + \frac{1}{\varepsilon^2}) \log m) \cdot \Theta\left(\frac{\rho_i \log m}{\varepsilon^2}\right)$ updates in pseudoforests, where each update is done in $O(\log n)$ worst-case time. The total update time for the i -run is thus $O((2^{2i} \varepsilon^{-2} + 2^i \varepsilon^{-4}) \log^3 m)$. Summing over $i = 1, \dots, r$ gives the time $O((\rho_{\max} \varepsilon^{-2} + \varepsilon^{-4}) \rho_{\max} \log^3 m)$. Finally, the update time of the run without pruning is $O(\rho_{\max}^2 \log^3 m)$, since its accuracy is $\varepsilon = \Theta(1)$.

3.2 Dynamic approximate fractional out-orientation

For a given pseudoforest P , there is a way to orient its edges such that the out-degree of any vertex is at most 1: in each component, orient the cycle in some direction, and orient other edges towards the cycle. In acyclic components, choose an arbitrary root vertex and orient the edges towards the root. We will call this orientation of the edges of P just “the orientation of P ”. Given a pseudoforest packing that covers each edge at least once, we can define a fractional orientation of the graph.

Definition 3.4 (Orientation induced by a pseudoforest packing). *Let $P_1, \dots, P_k$ be maximal pseudoforests such that every edge of the graph is present in at least one P_i . The collection $P_1, \dots, P_k$ induces a fractional orientation defined as follows. For an edge $(u, v) \in E$, set*

$$d_{u \rightarrow v} = \frac{\#\{P_i \mid (u, v) \in P_i, \text{ oriented } u \rightarrow v \text{ in the orientation of } P_i\}}{\#\{P_i \mid (u, v) \in P_i\}}. \quad (18)$$

In other words, the fractional orientation of an edge e is the average orientation of e in those pseudoforests that contain e . Notice that this indeed defines a fractional orientation, since $d_{u \rightarrow v} + d_{v \rightarrow u} = 1$, and we never get a zero in the denominator since every edge is covered by some P_i . Also, there is freedom in defining the orientation for cycles in the pseudoforests. However, it does not affect the out-degrees of any vertices.

Denote the out-degree of a vertex v in a fractional orientation as $\text{outdeg}(v) = \sum_{(u, v) \in E} d_{u \rightarrow v}$.

We are going to see that the fractional orientation induced by the greedy pseudoforest packing is an approximate minimum-out-degree orientation.

Theorem 3.5. *A fractional orientation induced by the greedy minimum weight maximal pseudoforest packing of $k = \Theta\left(\frac{\rho \log m}{\varepsilon^2}\right)$ pseudoforests guarantees $\text{outdeg}(u) \leq (1 + \varepsilon)\rho$ for any vertex u , if G is not a forest.*

Proof. Consider a packing of $\Theta\left(\frac{\rho \log m}{\varepsilon^2}\right)$ pseudoforests, so that $0 \leq \frac{1}{\min_e x_e^k} - \rho \leq \varepsilon \rho$ by Theorem 2.6.

For any vertex u , consider its out-degree.

$$\begin{aligned} \text{outdeg}(u) &= \sum_{(u,v) \in E} d_{u \rightarrow v} = \sum_{(u,v) \in E} \frac{\#\{P_i \mid (u,v) \in P_i, \text{ oriented } u \rightarrow v \text{ in the orientation of } P_i\}}{\#\{P_i \mid (u,v) \in P_i\}} \leq \\ &\leq \left(\max_{(u,v) \in E} \frac{1}{\#\{P_i \mid (u,v) \in P_i\}} \right) \sum_{(u,v) \in E} \#\{P_i \mid (u,v) \in P_i, \text{ oriented } u \rightarrow v \text{ in the orientation of } P_i\}. \end{aligned} \quad (19)$$

Since $\#\{P_i \mid (u,v) \in P_i\} = kx_{(u,v)}^k$, we have $\max_{(u,v) \in E} \frac{1}{\#\{P_i \mid (u,v) \in P_i\}} \leq \frac{1}{k \min_e x_e^k} \leq \frac{(1+\varepsilon)\rho}{k}$. Also, since the out-degree of u in each P_i is at most one, we have $\sum_{(u,v) \in E} \#\{P_i \mid (u,v) \in P_i, \text{ oriented } u \rightarrow v \text{ in the orientation of } P_i\} \leq k$. Putting it together into the previous equation yields

$$\text{outdeg}(u) \leq \frac{(1+\varepsilon)\rho}{k} \cdot k = (1+\varepsilon)\rho, \quad (20)$$

as claimed. $\square$

For our dynamic fractional orientation, we are going to use the orientation induced by a greedy pseudoforest packing. For a single dynamic pseudoforest P and an edge $e \in P$, one can query an orientation of e in P in $O(\log n)$ time. Indeed, if $e = (u, v)$, one can delete e from the pseudoforest and query the connectivity of u and v to the special edge in the cycle that the component of the pseudoforest remembers. If only one of the two vertices stays connected to the special edge on the cycle, orient e towards this vertex. If both u and v stay connected to the edge in the cycle, then e was in the cycle. In this case, choose an orientation for the special edge in the cycle, and orient e consistently.

Now we can prove Theorem 1.4.

Proof of Theorem 1.4. Maintain a packing of $k = \Theta\left(\frac{\rho \log m}{\varepsilon^2}\right)$ pseudoforests. The naive update time is $O(k^2 \log n) = O(\varepsilon^{-4} \rho^2 \log^3 m)$. The desired orientation is the orientation induced by the packing.

To answer the query for an edge e , go through the pseudoforests containing e and determine orientations of e in these pseudoforests in time $O(\log n)$ per pseudoforest. Output the average such orientation, as in the definition of the induced orientation. The query takes $O(k \log n) = O(\varepsilon^{-2} \rho \log^2 m)$ time. $\square$

4 Graphic matroids: 1-respecting min-cuts and a lower bound on the convergence

4.1 A packing of $\Theta(\lambda^5 \log m)$ trees contains a tree crossing some min-cut once

The dynamic min-cut algorithm by Thorup [26] uses the following key combinatorial result.

Theorem 4.1 ([26], Theorem 9). *A greedy tree packing with $\Theta(\lambda^7 \log^3 m)$ trees contains a tree crossing some min-cut only once.*

We give a stronger bound:

Theorem 4.2. *A greedy tree packing with $k = \Theta(\lambda^5 \log m)$ trees contains a tree crossing some min-cut only once.*

We will closely follow some parts of Thorup's proof of his Theorem 9. Just like him, we will assume that every tree in a greedy MST packing crosses all min-cuts at least twice, and show that this leads to a contradiction, if $\|x^k - x^*\|_\infty \leq \frac{\varepsilon}{\lambda}$, for a suitably small value of ε . Thorup gives his bound for $\varepsilon = O\left(\frac{1}{\lambda^3 \log m}\right)$. We will show that having $\varepsilon = O\left(\frac{1}{\lambda^2}\right)$ is sufficient.

Following Thorup's line of proof, let MC be the set of edges in min-cuts, let $\ell_0 := \min_{e \in MC} x_e^*$, and let ε be such that $\|x^k - x^*\|_\infty \leq \frac{\varepsilon}{\lambda}$. For the sake of contradiction, assume that every tree in the tree packing crosses every min-cut at least twice. In his Lemma 18, Thorup shows a bound

$$\ell_0 > (1 - \varepsilon) \frac{2}{\lambda}. \quad (21)$$

For an integer $i \geq 0$, define $\ell_i = \ell_0 - i\varepsilon/\lambda$. For some real number ℓ , let $E_{\geq \ell}^* := \{e \in E \mid x_e^* \geq \ell\}$ and $E_{> \ell}^k := \{e \in E \mid x_e^k \geq \ell\}$. Define $\mathcal{P}_i$ to be the partition whose sets are the components of $(V, E_{< \ell_i}^*)$. Thorup proves that for any $i \geq 2$, some trivial $\mathcal{P}_i$ cut is not a min-cut ([26], Lemma 17).

Then, Thorup finds a parameter a such that $|E_{\geq \ell_{a+2}}^* \setminus MC| \leq (1 + \alpha)|E_{\geq \ell_a}^* \setminus MC|$ for some small parameter α . He upper-bounds a by $O(\log_{1+\alpha} m) = O\left(\frac{\log m}{\alpha}\right)$. This is where our proof diverges from Thorup's.

Proof of Theorem 4.2. Let us assume that $\lambda \geq 2$. If $\lambda = 1$, then the first tree crosses all min-cuts once. Next, we need some definitions.

- Let G_i be the graph with the vertex set $V_i := \mathcal{P}_i$ and the edge set $E_i := E_{\geq \ell_i}^*$. Denote $n_i := |V_i|$, $m_i := |E_i|$.
- Let $p_i = |\{S \in \mathcal{P}_i \mid |\partial S| = \lambda\}|$ and $q_i = |\{S \in \mathcal{P}_i \mid |\partial S| > \lambda\}| = n_i - p_i$.
- Let $D_i = \sum_{S \in \mathcal{P}_i} (|\partial S| - \lambda)$.

We will need the following fact.

Lemma 4.3 (Lemma 14 in [26], reformulated and specialized). *The values of x_e^* are decreasing in the sense that for each i , for each $S \in \mathcal{P}_i$, we have $x_e^* < \ell_i$ for any edge $e \in E[S]$.*

Notice that $D_i = \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (|\partial S| - \lambda)$. Also, $D_i = 2m_i - \lambda n_i$.

Claim 4.4. *The sequence $D_0, D_1, \dots$ is non-decreasing.*

Proof. Notice that $D_i = (n_i - 1) \left(2 \frac{m_i}{n_i - 1} - \lambda\right) - \lambda$. Now we have that the sequence D_i is non-decreasing if we show the following:

1. n_i is positive
2. $2 \frac{m_i}{n_i - 1} - \lambda$ is positive
3. n_i is non-decreasing
4. $2 \frac{m_i}{n_i - 1} - \lambda$ is non-decreasing

Items 1 and 3 are obvious. Item 2 follows from the fact that the minimum degree in G_i is at least λ , so $\frac{2m_i}{n_i} \geq \lambda$. Item 4 follows from Lemma 4.3, since $\frac{m_i}{n_i - 1}$ is the reciprocal of the average x_e^* for $e \in E_i$, and the average x_e does not increase. $\square$

Claim 4.5. *If $\varepsilon(i + 2)(\lambda + 1) \leq 1$ then $|E[S] \cap E_i| \leq \lambda(|\partial S| - \lambda)$ for any $S \in \mathcal{P}_0$.*

Proof. We can assume that $E[S]$ is non-empty (otherwise the statement is trivial). Choose a value $\ell_S \leq \ell_0$ (to be defined later), and let H_S be a graph with the vertex set $S/E_{<\ell_S}^*$ and the edge set $E[S] \cap E_{\geq \ell_S}^*$. Let H_S have $n_S \geq 2$ vertices and $m_S \geq 1$ edges. By the definition of ℓ_0 , strict subsets of S cannot be min-cuts, and thus, they have the edge boundary of size at least $\lambda + 1$. Then, we can double-count $E(H_S)$ and have the following bound:

$$(\lambda + 1)n_S \leq 2m_S + |\partial S|. \quad (22)$$

Two cases are possible.

- $|\partial S| \leq \lambda + 1$. We let $\ell_S := \max_{e \in E[S]} x_e^*$. Using the bound (22), we can write

$$\ell_S = \frac{n_S - 1}{m_S} \leq \frac{2(n_S - 1)}{(\lambda + 1)n_S - |\partial S|} \leq \frac{2(n_S - 1)}{(\lambda + 1)n_S - (\lambda + 1)} = \frac{2}{\lambda + 1}.$$

The assumption $\varepsilon(i + 2)(\lambda + 1) \leq 1$ and the bound (21) give $\ell_i = \ell_0 - \frac{i\varepsilon}{\lambda} > \frac{2}{\lambda} - (i + 2)\frac{\varepsilon}{\lambda} > \frac{2}{\lambda} - \frac{2}{\lambda(\lambda + 1)} = \frac{2}{\lambda + 1}$. We showed that the edges $e \in E[S]$ satisfy $x_e^* \leq \ell_S \leq \frac{2}{\lambda + 1}$ while the edges $e \in E_i$ satisfy $x_e^* \geq \ell_i > \frac{2}{\lambda + 1}$. Therefore, $E[S] \cap E_i$ is empty.

- $|\partial S| > \lambda + 1$. We let $\ell_S := \ell_i$, then $E_S = E[S] \cap E_i$. Using the bounds (21) and (22), we get

$$(1 - \varepsilon)\frac{2}{\lambda} - \frac{i\varepsilon}{\lambda} < \ell_0 - \frac{i\varepsilon}{\lambda} = \ell_i \leq \min_{e \in E(H_S)} x_e^* \leq \max_{e \in E(H_S)} x_e^* = \frac{n_S - 1}{m_S} \leq \frac{2m_S + |\partial S| - \lambda - 1}{(\lambda + 1)m_S}.$$

Rearranging and using the assumption $\varepsilon(i + 2)(\lambda + 1) < 1$ gives

$$\frac{|\partial S| - \lambda - 1}{m_S} > \frac{\lambda + 1}{\lambda}(2 - (i + 2)\varepsilon) - 2 \geq \frac{\lambda + 1}{\lambda} \left(2 - \frac{1}{\lambda + 1}\right) - 2 = \frac{1}{\lambda}$$

Therefore, $m_S \leq \lambda(|\partial S| - \lambda - 1) < \lambda(|\partial S| - \lambda)$, as claimed. $\square$

Corollary 4.6. *If $\varepsilon(i + 2)(\lambda + 1) \leq 1$ then*

$$n_i \leq n_0 + 2D_0. \quad (23)$$

Proof. Consider any $S \in \mathcal{P}_0$. If $|\partial S| = \lambda$ then $E[S] \cap E_i$ is empty and hence $S \in \mathcal{P}_i$. Otherwise, if $|\partial S| > \lambda$, we have $m_S \leq \lambda(|\partial S| - \lambda)$ where we use the notation from the proof above. Using $\frac{n_S - 1}{m_S} < \frac{2}{\lambda}$ (which is true since x^* decreases when we recurse inside S , by Lemma 4.3), we have

$$n_S \leq 2(|\partial S| - \lambda). \quad (24)$$

Therefore,

$$n_i = p_0 + \sum_{S \in \mathcal{P}_0 : |\partial S| > \lambda} n_S \leq n_0 + 2D_0. \quad (25)$$

$\square$

Suppose, we have a tree packing $\mathcal{T}$ of k trees that satisfies $\|x^N - x^*\| \leq \frac{\varepsilon/2}{\lambda}$ for any $N \geq k$, and let $\mathcal{T}'$ be a tree packing of $k' = 4k$ trees that extends $\mathcal{T}$. Assume that all of the trees cross all min-cuts at least 2 times. The edge load in $\mathcal{T}' \setminus \mathcal{T}$ deviates from the ideal edge load by at most $\frac{\varepsilon}{\lambda}$:

$$\left\| \frac{k'x^{k'} - kx^k}{k' - k} - x^* \right\|_{\infty} \leq \frac{\varepsilon}{\lambda}.$$

Claim 4.7.

$$n_{i+1} - n_i > \frac{D_i}{\lambda} (1 - (i/2 + 1)(\lambda + 1)\varepsilon). \quad (26)$$

Proof. Let T be a uniformly random tree from $\mathcal{T}' \setminus \mathcal{T}$. On the one hand, $|T \cap E_i| \leq n_{i+1} - 1 < n_{i+1}$, because, if we compute T greedily by adding edges with the smallest weight, then the edges from E_i will appear later than the edges in $E[S]$ for any $S \in \mathcal{P}_{i+1}$. Therefore,

$$\mathbb{E}[|T \cap E_i|] < n_{i+1}. \quad (27)$$

On the other hand, double-counting $T \cap E_i$ leads us to

$$\begin{aligned} 2\mathbb{E}[|T \cap E_i|] &= \mathbb{E} \left[\sum_{S \in \mathcal{P}_i} |T \cap \partial S| \right] = \sum_{S \in \mathcal{P}_i : |\partial S| = \lambda} \mathbb{E}[|T \cap \partial S|] + \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} \mathbb{E}[|T \cap \partial S|] \\ &\geq 2p_i + \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (1 - (i/2 + 1)\varepsilon) \frac{2}{\lambda} |\partial S|. \end{aligned} \quad (28)$$

We used the fact that if S is a mincut, then $|T \cap \partial S| \geq 2$, and the assumption about the loads of $\mathcal{T}' \setminus \mathcal{T}$: that for any edge $e \in E_i$, $\mathbb{E}[|T \cap \{e\}|] = \frac{k'x_e^{k'} - kx_e^k}{k' - k} \geq x_e^* - \frac{\varepsilon}{\lambda} \geq \ell_i - \frac{\varepsilon}{\lambda} = \ell_0 - \frac{(i+1)\varepsilon}{\lambda} > (1 - (i/2 + 1)\varepsilon) \frac{2}{\lambda}$. Putting together (28) and (27), we get

$$\begin{aligned} n_{i+1} &> \mathbb{E}[|T \cap E_i|] \geq p_i + \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (1 - (i/2 + 1)\varepsilon) \frac{1}{\lambda} |\partial S| \\ &= p_i + q_i - q_i + \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} \frac{|\partial S|}{\lambda} - \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (i/2 + 1)\varepsilon \frac{|\partial S|}{\lambda} \\ &= n_i + \frac{1}{\lambda} \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (|\partial S| - \lambda) - \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (i/2 + 1)\varepsilon \frac{|\partial S|}{\lambda} \\ &= n_i + \frac{D_i}{\lambda} - \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} (i/2 + 1)\varepsilon \frac{|\partial S|}{\lambda}, \end{aligned} \quad (29)$$

which implies

$$n_{i+1} - n_i > \frac{D_i}{\lambda} - \frac{(i/2 + 1)\varepsilon}{\lambda} \sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} |\partial S|. \quad (30)$$

Finally, notice that $\sum_{S \in \mathcal{P}_i : |\partial S| > \lambda} |\partial S| \leq (\lambda + 1)D_i$, and thus,

$$n_{i+1} - n_i > \frac{D_i}{\lambda} (1 - (i/2 + 1)(\lambda + 1)\varepsilon). \quad (31)$$

□

Finally, we are ready to reach the contradiction. Recall that the sequence $D_0, D_1, \dots$, is non-decreasing by Claim 4.4. Set $i = 3\lambda$, and $\varepsilon = \frac{1}{10\lambda^2}$. Summing the bounds from the previous claim, we have

$$n_i > n_0 + \sum_{j=0}^{3\lambda} \frac{D_j}{\lambda} (1 - (j/2 + 1)(\lambda + 1)\varepsilon) > n_0 + 2D_0. \quad (32)$$

This contradicts Corollary 4.6.

To have the desired value for $\varepsilon = \Theta(\frac{1}{\lambda^2})$, we need $\Theta(\lambda \log m / \varepsilon^2) = \Theta(\lambda^5 \log m)$ trees. This concludes the proof. □

4.2 A tight lower bound for $\|x^k - x^*\|_\infty$

In this section, we first describe a family of graphs with constant λ and MST packings in these graphs with a property $\|x^k - x^*\|_\infty \geq \Omega\left(\sqrt{\frac{1}{k}}\right)$, for $k < O(n^2)$. We then extend the family of graphs to higher values of λ to get $\|x^k - x^*\|_\infty \geq \Omega\left(\sqrt{\frac{1}{k\lambda}}\right)$, matching Thorup's upper bound (4) up to an $O(\sqrt{\log m})$ factor. Before, the strongest lower bound was $\|x^k - x^*\|_\infty \geq \Omega\left(\frac{1}{k^{2/3}\lambda^{1/3}}\right)$ (also for polynomially bounded k) due to [10].

First, we will consider the case $\lambda = 2$. Let G_d be a square grid graph with the size $2 \times d$. Clearly, for this graph, $n = 2d$, $m = 3d - 2$, $\lambda = 2$. The ideal loads vector x^* is uniform, since one can verify that $V = \operatorname{argmax}_{S \subseteq V, |S| > 1} \frac{|E[S]|}{|S|-1}$. Thus, $x_e^* = \frac{n-1}{m} = \frac{2d-1}{3d-2} = \frac{2}{3} + \frac{2}{9n} + O\left(\frac{1}{n^2}\right)$.

The minimum spanning tree at some iteration of the MST packing might be not unique. For example, the first tree can be any spanning tree. We will break the ties by fixing an ordering on the edge set, see Figure 1. If at some iteration several MSTs exist, we will choose the one that also minimizes the total index of the edge ordering.




Figure 1: The graph G_d for $d = 8$. The edge ordering is shown with edge labels.

We will first give the intuition why $\|x^k - x^*\|_\infty = \Theta\left(\sqrt{\frac{1}{k}}\right)$, and then give a proof. The result of the first 3 iterations of the MST packing is shown in Figure 2. Notice that 3 spanning trees almost ideally double-cover the edges of the graph, but there is one “extra edge”. When k is divisible by 3, the loads will be almost uniform except for $k/3$ “extra” edges that will accumulate on the left end of the graph.

If one packs $k = 3i$ trees, the edges in the far-right end of the graph have loads $2i = \frac{2}{3}k$. In Figure 3, the edge labels show the difference between the edge loads and the level $\frac{2}{3}k$. In other words, it shows how overpacked the edge is relative to the average load $\frac{2}{3}k + O(1/n)$.

As a result, after packing $k = O(j^2)$ trees, k being divisible by 3, we have a triangle-shaped plot of how overpacked the edge depending on the position of the edge in the graph. See Figure 4 for an example for the graph G_d , $d = 100$, and $k = 3000$. The area under this triangle equals the total number of “extra” edges. Since an extra edge appears every 3 iterations, the area is $k/3$. The sides of this triangle are $\Theta(\sqrt{k})$, because the slope is a constant, and the area is $\Theta(k)$. Thus, $\|x^k - x^*\|_\infty = \frac{1}{k}\Theta(\sqrt{k}) + O(1/n) = \Theta\left(\sqrt{\frac{1}{k}}\right)$, when $k < O(n^2)$.

We are now ready for the proof. Let E_i be the set of edges covered by exactly i trees.

We will now describe the structure of the edge loads, and prove the description by induction. For the graph G_d , $d \geq 8$, and k such that $54 \leq k < O(d^2)$, the subgraphs E_i can be the following:

1. empty
2. $\begin{array}{|c|} \hline \square \\ \hline \end{array}$, $\begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array}$, $\begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array}$ – at the left end of the graph. We will call these “beginning tiles”.
3. $\begin{array}{|c|} \hline \square \\ \hline \end{array}$, $\begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array}$, $\begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array}$, $\begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array}$, $\begin{array}{|c|c|c|} \hline \square & \square & \square \\ \hline \end{array}$, $\begin{array}{|c|c|} \hline \square & \square \\ \hline \end{array}$, $\begin{array}{|c|} \hline \square \\ \hline \end{array}$ – not at either end of the graph. We call these “middle tiles”.




Figure 2: The first 3 iterations of the tree packing. Edge labels show edge loads.

4. - at the end of the graph. We call these “ending tiles”.

5. E_{j-1} for E_j being the ending tile. If nonempty, this edge set is located at the end of the graph, to the right from where the ending tile begins.

Moreover, the graph is covered by a sequence of tiles, starting from a beginning tile, then containing some number of middle tiles, and ending with an end tile, such that the edge load decreases by 1 in each tile as we go from left to right.

Moreover, the middle tiles , , , are never adjacent to each other. This concludes the description of the induction hypothesis.

For example, in Figure 3, for $k = 54$, the state is , , , , and for $k = 96$, the state is , , , . The case $k = 54$ also gave us the induction base.

In other words, the edge loads are described by a string, where the first symbol is one of the beginning tiles then we have some number of middle tiles, and the last symbol is one of the ending tiles. Then the edge loads decrease by one per tile. If these tiles do not cover the whole graph, and there are some edges left (these edges are at the right end, adjacent to the ending tile), then the load of these edges is one less than the load of the edges in the ending tile.

Now we need the induction step. Pack one more tree into the graph. First, we show that the beginning tile remains a beginning tile. Suppose, the beginning tile was E_i . At the next step, the beginning tile is either E_i or E_{i+1} , depending on the beginning tile and the following middle tile. We manually consider all cases, summarized in Table 1. In the table, the first column contains the current beginning tile, the first row contains the following middle tile, and the rest of the table gives the new beginning tile at the next step. “x” means that these two tiles do not fit, and such a sequence could not happen.




Figure 3: The iterations with $k = 6j^2$ for $j = 1 \dots 5$. The edge labels show the differences between kx_e^k and $2k/3$.

Next, we show that the middle tiles remain middle tiles. Consider a middle tile E_i , and the tile immediately to the left E_{i+1} (could be another middle tile or a beginning tile). At the next step, the new tile E_{i+1} depends only on the current E_i and E_{i+1} . We manually consider all cases, summarized in Table 2. In the table, the first column contains the current E_{i+1} , the first row contains the current E_i , and the rest of the table gives the new tile E_{i+1} at the next step. “x” means that these two tiles do not fit. One can verify that $\{\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}\}$ is the sink of this change rule, so the other tiles $\{\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}\}$ can only live for a constant time after they are born next to the ending tile. In fact, they can only appear either immediately to the left of the ending tile or one tile away from the ending tile. One may also verify that the tiles $\{\begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}, \begin{smallmatrix} \square & \square \\ \square & \square \end{smallmatrix}\}$ are never adjacent to each other due to their limited lifetime.




Figure 4: The edge loads of the edges in the graph G_d for $d = 100$ after packing 3000 trees.

	x	x		x	x		x	x
	x		x			x		x
		x	x	x	x	x	x	
	x	x		x	x		x	x
	x		x			x		x

Table 1: Beginning tiles change rules

Finally, we consider the ending tiles. The next ending tile depends on the current ending tile and the tile immediately to the left. The evolution rules are shown in Table 3. The first column contains the current tile immediately to the left from the ending tile, the first row contains the current ending tile, and the rest of the table gives the new ending tile, possibly also with the new tile immediately to the left. In the table, “unreachable” means that this state of ending tile and the adjacent tile can never be reached. Notice that the evolution of the ending tile is tripartite: if $k \bmod 3 = 0$, the ending tile is one of ..., ..., ..., ..., if $k \bmod 3 = 1$, the ending tile is one of ..., ..., ..., ..., and if $k \bmod 3 = 2$, the ending tile is one of ..., ..., ..., ...

Together, the above three tables show the rules how the string of tiles defining the edge loads at step k is transformed into a sequence of tiles at the next step, thus, completing the induction. Notice that the ending tile gets shorter at some iterations. For this reason, we also need to upper bound k by $O(n^2)$, because at this iteration we will “run out” of the ending tile.

Finally, if $k \bmod 3 = 0$, we have the ending tile one of ..., ..., ..., ..., with the load $2k/3$. The excessive load of the beginning and middle tiles is $\Theta(k)$, and the length of each tile is $O(1)$. This implies that the excessive load of the beginning tile is $\Theta(\sqrt{k})$, and thus, the claim holds.

	x		x			x		x
	x	x		x	x		x	x
	x		x			x		x
		x	x	x	x	x	x	
	x	x		x	x		x	x
	x		x			x		x
		x	x	x	x	x	x	
	x	x		x	x		x	x

Table 2: Middle tiles change rules

																	
	x		x	x		x	x		x		x	x				x	
	x	x		x	x	x		x	x		x	x		x		x	x
	x		x	x		x	x		x	x		x				x	x
		x	x	x	x		x	x	x	x				x	x	x	x
	x		x	x		x	x		x	x	x		x			x	x
		x	x	x	x		x	x	x	x				x	x	x	
	x	x	x		x	x	x	x		x	x	x	x	x	x		
		x	x	x	x		x	x	x	x	unreachable	unreachable	x	x	x	x	x

Table 3: Ending tiles change rules

For higher values of λ , take graphs G_d^w defined as G_d with every edge duplicated w times. The connectivity of G_d^w is $\lambda = 2w$. Define an ordering of the edges like in G_d . Then, if one packs wj trees in G_d^w , the packing equals to the packing of j trees in G_d duplicated w times. Then, the difference between the maximal load and the average load is $\Theta(\sqrt{j}) = \Theta(\sqrt{k/\lambda})$. Then, the error in the *relative* loads is $\|x^k - x^*\|_\infty = \Theta(\sqrt{k/\lambda})/k = \Theta\left(\sqrt{\frac{1}{k\lambda}}\right)$, as claimed.

5 Other observations

5.1 Theoretical observations

Lemma 5.1. *After k iterations of the greedy base packing, we have*

$$\|x^k\|_2^2 - \|x^*\|_2^2 \leq 2r \frac{\log(k+1)}{k}, \quad (33)$$

where $r := rk(E)$ is the rank of the matroid M . Therefore,

$$\|x^k\|_2 - \|x^*\|_2 \leq \sqrt{m} \cdot \frac{\log(k+1)}{k}, \quad (34)$$

where $m := |E|$.

Proof. Define a function $f(x) = \|x\|_2^2$, and let $f^k := f(x^k)$. Also, let $f^* = f(x^*)$.

We will bound f^{k+1} in terms of f^k . Let v^k be the indicator vector of the the base packed at the step k . Recall that $x^{k+1} = \frac{k}{k+1}x^k + \frac{1}{k+1}v^k$.

$$f^{k+1} = \left\| \frac{k}{k+1}x^k + \frac{1}{k+1}v^k \right\|_2^2 = \frac{k^2}{(k+1)^2} \|x^k\|_2^2 + \frac{1}{(k+1)^2} \|v^k\|_2^2 + \frac{2k}{(k+1)^2} \langle x^k, v^k \rangle. \quad (35)$$

We will use the fact that $\|v_k\|_2^2 = r := rk(E)$ (in case of the MST packing, $\|v_k\|_2^2 = n - 1$). Also, since v^k minimizes the value of $\langle x^k, v^k \rangle$, we have $\langle x^k, v^k \rangle \leq \langle x^k, x^* \rangle \leq \|x^k\|_2 \|x^*\|_2 = \sqrt{f^k f^*}$. Therefore,

$$f^{k+1} \leq \frac{k^2}{(k+1)^2} f^k + \frac{r}{(k+1)^2} + \frac{2k}{(k+1)^2} \sqrt{f^k f^*}. \quad (36)$$

Define $\Delta f^k := f^k - f^*$. Then we have

$$\begin{aligned} \Delta f^{k+1} &\leq \frac{k^2}{(k+1)^2} \Delta f^k + \frac{r - f^*}{(k+1)^2} + \frac{2k}{(k+1)^2} \left(\sqrt{(f^* + \Delta f^k) f^*} - f^* \right) \\ &\leq \frac{k^2}{(k+1)^2} \Delta f^k + \frac{r - f^*}{(k+1)^2} + \frac{k}{(k+1)^2} \Delta f^k \\ &\leq \frac{k}{k+1} \Delta f^k + \frac{r}{(k+1)^2}. \end{aligned} \quad (37)$$

Now we will prove (33) by induction. Base case: for $k = 1$, x^1 is an indicator vector of some base, so $\|x^1\|_2^2 - \|x^*\|_2^2 = r - f^* \leq 2r \frac{\log(2)}{1}$. For the step case, using (37), we get

$$\Delta f^{k+1} \leq 2r \frac{\log(k+1)}{k+1} + \frac{r}{(k+1)^2}. \quad (38)$$

On the other hand,

$$2r \frac{\log(k+2)}{k+1} \geq \frac{2r}{k+1} \left(\log(k+1) + \frac{1}{k+1} - \frac{1}{(k+1)^2} \right) \geq 2r \frac{\log(k+1)}{k+1} + \frac{r}{(k+1)^2} \geq \Delta f^{k+1}, \quad (39)$$

concluding the proof of (33).

For the proof of (34), notice that $\|x^*\|_2^2 \geq \frac{r^2}{m}$, since $\sum_e x_e^* = r$. Therefore, $\|x^*\|_2 \geq \frac{r}{\sqrt{m}}$. Also, $\|x^k\|_2^2 - \|x^*\|_2^2 = (\|x^k\|_2 - \|x^*\|_2)(\|x^k\|_2 + \|x^*\|_2) \geq 2\|x^*\|_2(\|x^k\|_2 - \|x^*\|_2) \geq \frac{2r}{\sqrt{m}}(\|x^k\|_2 - \|x^*\|_2)$. Then we get

$$\|x^k\|_2 - \|x^*\|_2 \leq \frac{\sqrt{m}}{2r} \cdot 2r \frac{\log(k+1)}{k} = \sqrt{m} \cdot \frac{\log(k+1)}{k}. \quad (40)$$

□

Corollary 5.2. *After k iterations of the greedy base packing, we have*

$$\|x^k - x^*\|_2 \leq \sqrt{2r \frac{\log(k+1)}{k}}, \quad (41)$$

where $r := rk(E)$ is the rank of the matroid M .

Proof. Since x^* minimizes $\|x\|_2^2$ over the base polytope, we have the inequality $\langle x^k, x^* \rangle \geq \langle x^*, x^* \rangle = \|x^*\|_2^2$. Using this and the above lemma, we can write

$$\|x^k - x^*\|_2^2 = \|x^k\|_2^2 + \|x^*\|_2^2 - 2\langle x^k, x^* \rangle \leq \|x^k\|_2^2 - \|x^*\|_2^2 \leq 2r \frac{\log(k+1)}{k}. \quad (42)$$

Therefore, $\|x^k - x^*\|_2 \leq \sqrt{2r \frac{\log(k+1)}{k}}$, as claimed. □

In other words, to guarantee $\|x^k - x^*\|_2 \leq \varepsilon$, one needs $k = \Theta\left(\frac{r \log(r/\varepsilon)}{\varepsilon^2}\right)$ iterations.

Lemma 5.3. *For a fixed matroid M and a fixed integer $p \geq 2$, after k iterations of the greedy base packing, we have*

$$\|x^k\|_p^p - \|x^*\|_p^p \leq \frac{p^2}{2} \|x^*\|_{p-1}^{p-1} \frac{\log k}{k} (1 + o(1)), \quad (43)$$

where $o(1)$ is $o(1)$ as $k \rightarrow \infty$. Consequently,

$$\|x^k\|_p - \|x^*\|_p \leq \frac{p}{2} \cdot m^{1/p} \frac{\log k}{k} (1 + o(1)). \quad (44)$$

Proof. Define functions $f_p(x) = \|x\|_p^p$, and let $f_p^k := f_p(x^k)$. Also, let $f_p^* = f_p(x^*)$. We will bound f_p^{k+1} in terms of f_q^k for $q \in \{1, \dots, p\}$. Let v^k be the indicator vector of the base packed at the step k . Recall that $x^{k+1} = \frac{k}{k+1}x^k + \frac{1}{k+1}v^k$.

$$\begin{aligned} f_p^{k+1} &= \left\| \frac{k}{k+1}x^k + \frac{1}{k+1}v^k \right\|_p^p = \sum_e \left(\frac{k}{k+1}x_e^k + \frac{1}{k+1}v_e^k \right)^p \\ &= \sum_e \left(\frac{k}{k+1}x_e^k \right)^p + \sum_{i=0}^{p-1} \sum_e \frac{k^i}{(k+1)^p} \binom{p}{i} (x_e^k)^i (v_e^k)^{p-i}. \end{aligned} \quad (45)$$

Notice that since v^k is a 0/1 vector, we have $(v_e^k)^{p-i} = v_e^k$ whenever $p-i > 0$. Using this, we continue

$$f_p^{k+1} = \left(\frac{k}{k+1} \right)^p f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-1} k^i \binom{p}{i} \sum_e (x_e^k)^i v_e^k. \quad (46)$$

The sum $\sum_e (x_e^k)^i v_e^k$ is just the dot product $\langle (x^k)^i, v^k \rangle$. Since v^k minimizes the value of this dot product, we have $\langle (x^k)^i, v^k \rangle \leq \langle (x^k)^i, x^* \rangle$. Using Hölder's inequality, we also have $\langle (x^k)^i, x^* \rangle \leq \|x^k\|_{i+1}^i \|x^*\|_{i+1}$. Then,

$$f_p^{k+1} \leq \left(\frac{k}{k+1} \right)^p f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-1} k^i \binom{p}{i} \|x^k\|_{i+1}^i \|x^*\|_{i+1}. \quad (47)$$

Now, let us pass to $\Delta f_p^k := f_p^k - f_p^*$.

$$\begin{aligned} \Delta f_p^{k+1} &\leq \left(\frac{k}{k+1} \right)^p f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-1} k^i \binom{p}{i} \|x^k\|_{i+1}^i \|x^*\|_{i+1} - f_p^* \\ &= \left(\frac{k}{k+1} \right)^p \Delta f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-1} k^i \binom{p}{i} (\|x^k\|_{i+1}^i \|x^*\|_{i+1} - f_p^*). \end{aligned} \quad (48)$$

We will bound $\|x^k\|_{i+1}^i \|x^*\|_{i+1}$ from above. $\|x^k\|_{i+1}^i \|x^*\|_{i+1} = (f_{i+1}^* + \Delta f_{i+1}^k)^{\frac{i}{i+1}} (f_{i+1}^*)^{\frac{1}{i+1}} = f_{i+1}^* \left(1 + \frac{\Delta f_{i+1}^k}{f_{i+1}^*} \right)^{\frac{i}{i+1}} \leq f_{i+1}^* + \frac{i}{i+1} \Delta f_{i+1}^k$. Here we used the fact that the function $t^{\frac{i}{i+1}}$ is concave as a function of t . Thus,

$$\begin{aligned} \Delta f_p^{k+1} &\leq \left(\frac{k}{k+1} \right)^p \Delta f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-1} k^i \binom{p}{i} \left(f_{i+1}^* - f_p^* + \frac{i}{i+1} \Delta f_{i+1}^k \right) \\ &\leq \frac{k^{p-1}(k+p-1)}{(k+1)^p} \Delta f_p^k + \frac{1}{(k+1)^p} \sum_{i=0}^{p-2} k^i \binom{p}{i} \left(f_{i+1}^* + \frac{i}{i+1} \Delta f_{i+1}^k \right). \end{aligned} \quad (49)$$

From now on, we will pass to studying the asymptotics of the bounds. We will imagine p fixed, and $k \rightarrow \infty$. The remainder of the proof is a technical study of the above equation, treated as a recurrence for Δf_p^k . First, if k tends to infinity, then $f_{i+1}^* + \frac{i}{i+1} \Delta f_{i+1}^k = f_{i+1}^*(1 + o(1))$. Second, the sum that we have is dominated by the last term (the term that contains k^{p-2}). So, for large k , we have the second term in right-hand side approximated as follows:

$$\frac{1}{(k+1)^p} \sum_{i=0}^{p-2} k^i \binom{p}{i} \left(f_{i+1}^* + \frac{i}{i+1} \Delta f_{i+1}^k \right) = \frac{p(p-1)}{2} \cdot \frac{k^{p-2}}{(k+1)^p} f_{p-1}^* (1 + o(1)), \quad (50)$$

where $o(1)$ stands for $o(1)$ as $k \rightarrow \infty$.

Then, for large enough k , we have

$$\Delta f_p^{k+1} \leq \frac{k^{p-1}(k+p-1)}{(k+1)^p} \Delta f_p^k + \frac{p^2}{2} \cdot \frac{k^{p-2}}{(k+1)^p} f_{p-1}^*, \quad (51)$$

for $k \geq K$, where K is some constant.

Denote $C_p = \frac{p^2}{2} \cdot f_{p-1}^*$. Using this notation, we have

$$\Delta f_p^{k+1} \leq \frac{k^{p-1}(k+p-1)}{(k+1)^p} \Delta f_p^k + C_p \frac{k^{p-2}}{(k+1)^p} \quad (52)$$

for $k \geq K$.

Introduce a variable $y^k = k \Delta f_p^k$. Now we have

$$y^{k+1} \leq \left(\frac{k}{k+1} \right)^{p-2} \frac{(k+p-1)}{k+1} y^k + C_p \frac{k^{p-2}}{(k+1)^{p-1}}. \quad (53)$$

Denote $\alpha_k := \left(\frac{k}{k+1} \right)^{p-2} \frac{(k+p-1)}{k+1}$, $\beta_k := C_p \frac{k^{p-2}}{(k+1)^{p-1}}$. Now the previous inequality becomes

$$y^{k+1} \leq \alpha_k y^k + \beta_k \quad (54)$$

for $k \geq K$.

We will show that $0 \leq \alpha_k \leq 1$. Nonnegativity is obvious. For proving $\alpha_k \leq 1$, we will have two steps. First, for fixed p , we have α_k non-decreasing as a function of k , due to $\frac{d}{dk} \alpha_k = (p-1)(p-2) \frac{k^{p-3}}{(k+1)^p} \geq 0$. Second, α_k is asymptotically at most 1 for large k , because we can have a Taylor expansion: $\alpha_k = 1 - \frac{p^2-3p+2}{2} \frac{1}{k^2} + O\left(\frac{1}{k^3}\right)$. Together these two facts imply $\alpha_k \leq 1$. Then, at a cost of making our inequality a bit weaker, it becomes simpler: for $k \geq K$,

$$y^{k+1} \leq y^k + \beta_k. \quad (55)$$

Summing up these bounds, we get

$$y^k \leq y^K + \sum_{i=K}^{k-1} \beta_i. \quad (56)$$

Let us bound β_k : $\beta_k = C_p \frac{k^{p-2}}{(k+1)^{p-1}} = \frac{C_p}{k+1} \left(\frac{k}{k+1} \right)^{p-2} \leq \frac{C_p}{k+1}$. Now we have a bound

$$\begin{aligned} y^k &\leq y^K + \sum_{i=K}^{k-1} \beta_i \leq y^K + \sum_{i=K}^{k-1} \frac{C_p}{i+1} = y^K + C_p \log k (1 + o(1)) = C_p \log k (1 + o(1)) \\ &= \frac{p^2}{2} \cdot f_{p-1}^* \log k (1 + o(1)) = \frac{p^2}{2} \cdot \|x^*\|_{p-1}^{p-1} \log k (1 + o(1)). \end{aligned} \quad (57)$$

This concludes the first claim of the lemma.

To show (44), notice that $\|x^k\|_p^p = (\|x^*\|_p^p + p(\|x^k\|_p - \|x^*\|_p)\|x^*\|_p^{p-1})(1 + o(1))$. Then,

$$\|x^k\|_p - \|x^*\|_p = \frac{\|x^k\|_p^p - \|x^*\|_p^p}{p\|x^*\|_p^{p-1}}(1 + o(1)) \leq p \left(\frac{\|x^*\|_{p-1}}{\|x^*\|_p} \right)^{p-1} \frac{\log k}{k} (1 + o(1)). \quad (58)$$

Finally, one can bound $\left(\frac{\|x^*\|_{p-1}}{\|x^*\|_p} \right)^{p-1} \leq \frac{m}{m^{1-1/p}} = m^{1/p}$. Plugging it in the previous equation yields (44). $\square$

For both above lemmas, we used inequalities of the form $\langle x^k, v^k \rangle \leq \langle x^k, x^* \rangle$ and $\langle (x^k)^i, v^k \rangle \leq \langle (x^k)^i, x^* \rangle$, which we had due to the fact that v^k minimized functions $\langle x^k, v \rangle$ and $\langle (x^k)^i, v \rangle$ over the base polytope P_B . We would like to point out that in case x^* is in the relative interior of P_B , and (after passing to the affine subspace spanned by P_B) there is a ball of radius R around x^* that is contained in P_B , then we can use stronger inequalities instead:

$$\langle x^k, v^k \rangle \leq \left\langle x^k, x^* - R \frac{x^k}{\|x^k\|_2} \right\rangle = \langle x^k, x^* \rangle - R\|x^k\|_2, \quad (59)$$

and similarly,

$$\langle (x^k)^i, v^k \rangle \leq \langle (x^k)^i, x^* \rangle - R\|(x^k)^i\|_2. \quad (60)$$

This would lead to stronger convergence bounds of $\|x^k\|_2^2 - \|x^*\|_2^2 = O\left(\frac{1}{k^2}\right)$ and $\|x^k\|_p^p - \|x^*\|_p^p = O\left(\frac{1}{k^2}\right)$.

Fix a graph G . The corresponding graphic matroid $\mathcal{G}$ has the collection of independent sets that are acyclic edge sets, i.e. forests. The k -fold union of $\mathcal{G}$ with itself, $\mathcal{G}^k$, is also a matroid. Its independent sets are the edge sets partitionable into k forests. If G contains k edge-disjoint spanning trees, then the bases of $\mathcal{G}^k$ are precisely edge sets that are unions of k spanning trees. Next, we relate the limit of the base packing in $\mathcal{G}$ and $\mathcal{G}^k$.

Let P_{kST} be the convex hull of the indicator vectors of all possible edge-disjoint packings of k spanning trees in the graph G . If k edge-disjoint spanning trees do not exist, let P_{kST} be empty.

Lemma 5.4. *Suppose, P_{kST} is not empty. Let $\tilde{x}^* = \arg \min_{x \in P_{kST}} \|x\|_2^2$. Then $\tilde{x}^* = kx^*$.*

Proof. We will show that $kx^* \in P_{kST}$. Notice that

$$P_{kST} = kP_{ST} \cap \{x \leq 1\}. \quad (61)$$

It is enough to show that $kx_e^* \leq 1$ for any edge e . Recall that $\max_{e \in E} x_e^* = \phi$ is the optimal packing value of G , $\phi = \max_{\mathcal{P}: \text{partition of } V} \frac{|\mathcal{P}| - 1}{|E(\mathcal{P})|}$. Since we assume that P_{kST} is not empty, there exist k edge-disjoint spanning trees in G . By the classical Nash-Williams theorem, k edge-disjoint spanning trees in G exist if and only if for any partition of the vertex set $\mathcal{P}$, $|E(\mathcal{P})| \geq |\mathcal{P}| - 1$. Then, $\phi \leq \frac{1}{k}$. Then, $kx_e^* \leq k\phi \leq 1$. So, $kx^* \leq 1$, and $kx^* \in P_{kST}$. Then, $\|\tilde{x}^*/k\|_2 \leq \|x^*\|_2$, and kx^* minimizes the second norm over P_{kST} . Then $\tilde{x}^* = kx^*$, as claimed. $\square$

Corollary 5.5. *Suppose, G contains k edge-disjoint spanning trees. Let $\tilde{x}^*$ be the limit of the edge loads in the greedy packing of bases of $\mathcal{G}^k$. Then $\tilde{x}^* = kx^*$, where x^* is the usual vector of ideal edge loads.*

For a set of vertices $S \subseteq V$, let $d(S) := E[S]/\binom{|S|}{2}$. This is yet another notion of density, not to be confused with the density we were dealing with earlier. Also, let $d := d(V)$.

A graph G is called ε -quasirandom if for any subset of vertices $S \subseteq V$ such that $|S| \geq \varepsilon n$, we have $|d(S) - d| \leq \varepsilon$.

Lemma 5.6. *If G is an ε -quasirandom graph, and $\varepsilon < d$. Then almost all edges have the same value of x_e^* , in the sense that the fraction of these edges is at least $1 - \frac{3\varepsilon}{d} + O(n^{-1})$.*

Proof. Consider the minimum value of x^* , $x_{\min}^* := \min_{e \in E} x_e^*$. This value is attained at the edges inside some vertex set $S^* \subseteq V$, such that $E[S^*]/(|S^*| - 1)$ is maximized over all vertex subsets S^* . Moreover,

$$x_{\min}^* = \frac{|S^*| - 1}{E[S^*]}. \quad (62)$$

We claim that $|S^*| \geq \varepsilon n$. If we assume the contrary, then

$$\frac{|S^*| - 1}{E[S^*]} \geq \frac{|S^*| - 1}{(|S^*| - 1)|S^*|/2} = \frac{2}{|S^*|} > \frac{2}{\varepsilon n}. \quad (63)$$

However, consider the value $(|S| - 1)/E[S]$ for $S = V$:

$$\frac{|V| - 1}{E[V]} = \frac{n - 1}{m} = \frac{2}{dn} < \frac{2}{\varepsilon n} < \frac{|S^*| - 1}{E[S^*]}, \quad (64)$$

which contradicts the optimality of S^* .

Next, we will use the bound on $d(S^*)$:

$$x_{\min}^* = \frac{|S^*| - 1}{E[S^*]} = \frac{2}{d(S^*)|S^*|} \geq \frac{2}{(d + \varepsilon)|S^*|}. \quad (65)$$

Then we can bound the sum of x_e^* over all edges:

$$n - 1 = \sum_{e \in E} x_e^* \geq m \cdot x_{\min}^* \geq \frac{2m}{(d + \varepsilon)|S^*|}. \quad (66)$$

$$|S^*| \geq \frac{2m}{(d + \varepsilon)(n - 1)} = n \frac{d}{d + \varepsilon} \geq n \left(1 - \frac{\varepsilon}{d}\right). \quad (67)$$

So, S^* includes almost all vertices if ε is small. Every edge inside S^* has the same value of $x_e^* = x_{\max}^*$. Finally, we can bound the fraction of edges in S^* , which will conclude the proof.

$$\begin{aligned} E[S^*]/m &= \frac{d(S^*)|S^*|(|S^*| - 1)}{2m} \geq \frac{(d - \varepsilon)|S^*|(|S^*| - 1)}{2m} \\ &\geq \frac{n(n - 1)}{2m} \cdot (d - \varepsilon) \left(1 - \frac{\varepsilon}{d}\right)^2 + O(n^{-1}) \\ &= \left(1 - \frac{\varepsilon}{d}\right)^3 + O(n^{-1}) \geq 1 - \frac{3\varepsilon}{d} + O(n^{-1}). \end{aligned} \quad (68)$$

□

As shown above, in quasirandom graphs, for almost all the edges, the values x_e^* coincide. We believe that the following might be true.

Conjecture 5.7. *For any $\varepsilon > 0$, there exists an integer $M(\varepsilon)$ such that for any graph G on n vertices, there exists a set $X \subset \mathbb{R}$, $|X| \leq M(\varepsilon)$, and for all but εn^2 edges, we have $x_e^* \in X$.*

We believe so because of the Szemerédi's regularity lemma. After considering the regularity partition of the graph, in each regular pair, x^* might be the same for most of the edges, since an ε -regular pair is, in some sense, a bipartite quasirandom graph. This conjecture deviates significantly from this paper's topic, so we did not try to prove it.

5.2 Experimental observations for tree packings

In this section, we will show some experimental observations concerning the greedy MST packings. We simulate the greedy MST packings for different graphs and give comments on the observed dependencies of the errors $\|x^k - x^*\|_\infty$, $\|x^k\|_2 - \|x^*\|_2$ and $\|x^k\|_{10} - \|x^*\|_{10}$ on k . The drop rate of these errors can show a variety of qualitatively different behaviors for different graphs, so we believe that it is nice to show these observations.

Figure 5 contains data for a greedy MST packing for an example graph. On the top-right you can see the scatter plot of $\|x^k - x^*\|_\infty$ versus iteration number k . Thorup's upper bound is also plotted. It is visible that there is a gap in the behavior of the Thorup's bound $O\left(\frac{1}{\sqrt{k}}\right)$ and the real error $\|x^k - x^*\|_\infty$, that drops like $O\left(\frac{1}{k}\right)$. The errors in the second norm and the p -norm for $p = 10$ seem to drop like $\sim \frac{1}{k}$, which matches our bounds (34) and (44) up to a $(\log k)$ -factor.

The behavior of the errors can be more complex. In Figure 6, we show the simulation results for a random graph $G(n, p)$, $n = 50$, $p = 0.2$. One can see that both $\|x^k - x^*\|_2$ and $\|x^k - x^*\|_{10}$ start decreasing like $\sim \frac{1}{k^2}$, but then experience a break and transfer to a regime $\sim \frac{1}{k}$. The break seems to happen when the number of iteration k is much higher than the dimension of P_{ST} . Meanwhile, the error in the infinity-norm is still close to $\frac{1}{k}$.

While in the first two examples the errors decreased more or less monotonically, it can happen that the error plots oscillate (while still having a downward trend of $\sim \frac{1}{k^2}$, at least initially), like in Figure 7. Empirically, the period of oscillations in these cases is approximately λ . The error $\|x^k - x^*\|_\infty$, while oscillating, still does not deviate far from $\frac{1}{k}$.

In our final example, we consider a square grid graph (Figure 8). While k is much less than n , the $\|x^k - x^*\|_\infty$ decreases like $\sim \frac{1}{\sqrt{k}}$ instead of $\sim \frac{1}{k}$, like in the previous examples. However, the trend experiences a break, and if $k \gg n$, then $\|x^k - x^*\|_\infty$ starts decreasing like $\sim \frac{1}{k}$ again. The plots of $\|x^k - x^*\|_2$ and $\|x^k - x^*\|_{10}$ also experience breaks and transition from $\sim \frac{1}{k}$ behavior to $\sim \frac{1}{k^2}$. Notice that in Figure 6, we have also seen breaks in $\|x^k - x^*\|_2$ and $\|x^k - x^*\|_{10}$, but then they were the other way around: from $\sim \frac{1}{k^2}$ to $\sim \frac{1}{k}$.

A Proof of Theorem 2.4

The proof has three parts. First, we will show that the described x^* is feasible. Second, we will show that the value of x^* assigned to elements through the iterations, increases. This will be used in the third part, where we will recall that x^* is the (unique) minimizer of $\|x\|_2^2$, and show that $\langle \nabla \|x\|_2^2|_{x=x^*}, y - x^* \rangle \geq 0$ for all $y \in P_B$, thus proving the optimality of x^* .

Part 1: feasibility. We will show that x^* is a convex combination of indicator vectors of some bases. This would prove feasibility. The proof is by induction. After some number of iterations, let $C \subseteq E$ be the set of all elements that were contracted (so, x^* is defined only on C so far). We are going to show the existence of a collection of bases of the restriction matroid $M|_C$, such that $x^*|_C$ is a convex combination of the indicator vectors of these bases.

Base case: consider the first H -set found in step 1. For the restriction matroid $M|_H$, the uniform $x_e^* = \frac{r(H)}{|H|}$ is feasible. Indeed, it does not violate any constraints $x^*(H') \leq r_{M|H}(H') = r_M(H')$ for any $H' \subseteq H$, otherwise, H' would be discovered instead of H in step 1. Therefore, $x^*|_H$ is a convex combination of some bases of $M|_H$.

Step case: suppose, C is the set of edges that have been contracted so far. Let H be the next set found in step 1, and let $x_H := \frac{r(H)}{|H|}$ be the value of x_e^* for any element $e \in H$. Notice that here $r(\cdot)$ is the rank function for the contraction matroid M/C . By the same argument as above, the newly defined $x^*|_H$ is a convex combination of indicator vectors of some bases $B_1, \dots, B_k$ of $(M/C)|_H = (M|(C \cup H))/C$. Suppose, $x^*|_H = \lambda_1 I_{B_1} + \dots + \lambda_k I_{B_k}$ is this convex combination, where I_{B_i} is the indicator vector of B_i .

By the induction assumption, there exist bases $B'_1, \dots, B'_{k'}$ of $M|C$, such that $x^*|_C$ is their convex combination: $x^*|_C = \lambda'_1 I_{B'_1} + \dots + \lambda'_{k'} I_{B'_{k'}}$. We are going to “stitch” the two families of bases together to get a new family of kk' bases of $M|(C \cup H)$. Define the sets $\tilde{B}_{ij} := B_i \cup B'_j$ for all indices $i \in \{1, \dots, k\}$, $j \in \{1, \dots, k'\}$. First, $\tilde{B}_{ij}$ are bases of $M|(C \cup H)$. Second, $x^*|_{C \cup H}$ is their convex combination with coefficients $\lambda_i \lambda'_j$. Indeed, for any element $e \in C$,

$$\sum_{i=1}^k \sum_{j=1}^{k'} \lambda_i \lambda'_j I_{\tilde{B}_{ij}}(e) = \sum_{i=1}^k \lambda_i \sum_{j=1}^{k'} \lambda'_j I_{B'_j}(e) = x_e^*, \quad (69)$$

and similarly for $e \in H$. The induction step is done.

Part 2: monotonicity. We will show that, throughout the iterations, the value x_e^* assigned to the elements increases. This will help us in the optimality analysis.

Suppose, at some iteration, we found the set H_1 in step 1, and H_2 is the set found at the next iteration. Denote $|H_1| = a$, $r(H_1) = b$, $|H_2| = c$, $r(H_2) = d$, where the rank function $r(\cdot)$ is *at the moment of discovery of* H_1 . In this notation, the value of x^* in H_1 is $\frac{b}{a}$, and the value of x^* in H_2 is $\frac{r(H_1 \cup H_2) - b}{c}$.

Since the set $H_1 \cup H_2$ was not discovered instead of H_1 , we have the inequality $\frac{a}{b} > \frac{a+c}{r(H_1 \cup H_2)}$. This inequality implies $\frac{a}{b} > \frac{a+c-a}{r(H_1 \cup H_2) - c}$, and thus, $\frac{b}{a} < \frac{r(H_1 \cup H_2) - b}{c}$, which is exactly the desired monotonicity of x^* .

Part 3: optimality. Since $\nabla \|x\|_2^2 = 2x$, it is enough to show that $\langle x^*, y - x^* \rangle \geq 0$ for all $y \in P_B$. It is enough to consider y being a vertex of P_B . So, we need to prove that the weight of the minimum weight base under weights x^* equals $\langle x^*, x^* \rangle = \|x^*\|_2^2$.

The minimum weight base B can be found by greedily adding elements ordered by the weight to P , while P keeps being an independent set. In the first H -set, P contains $r(H)$ elements with total weight $r(H) \cdot \frac{r(H)}{|H|} = \sum_{e \in H} (x_e^*)^2$. After some number of steps, let C be the set of edges that have been contracted so far. Suppose, H is the next set discovered at step 1. B contains $r(H)$ elements in H , thus, the weight of $B|_H$ is $r(H) \cdot \frac{r(H)}{|H|} = \sum_{e \in H} (x_e^*)^2$. Summing all these equalities, we get that the total weight of B is $\sum_e (x_e^*)^2 = \|x^*\|_2^2$. This concludes the proof.

B Proof of Theorem 2.6

Define $H = \{e \in E : x_e^* = \frac{1}{\rho}\}$ where x^* is the optimal loads vector for M . By Theorem 2.4, H is the maximal set in $H \in \operatorname{argmax}_{\emptyset \neq H \subseteq E} \frac{|H|}{r(H)}$, and $|H| = \rho \cdot r(H)$. For a value $\alpha \in \mathbb{R}$ and integer $k \geq 1$ define

$$A_k^\alpha = \sum_{e \in E_k} (1 + \alpha)^{k \left(\frac{1}{\rho} - x_e^* \right)} \quad (70a)$$

$$B_k^\alpha = \sum_{e \in H \cap E_k} (1 + \alpha)^{k \left(\frac{1}{\rho} - x_e^* \right)} \quad (70b)$$

We will prove the following by induction on $k = 0, 1, 2, \dots$

$$A_k^\alpha \leq m(1 + \alpha)^{k/\rho} \left(1 - \frac{\alpha}{\rho(1 + \alpha)} \right)^k \quad \forall \alpha > -1 \quad (71a)$$

$$B_k^\alpha \leq m(1 + \alpha)^{k/\rho} \left(1 - \frac{\alpha}{\rho(1 + \alpha)} \right)^k \quad \forall \alpha \in (-1, 0) \quad (71b)$$

$$H \subseteq E_k \quad (71c)$$

Before proving this, we make the following observations for a fixed $k \geq 1$.

- By Theorem 2.4, (71c) implies that $\rho(M_k) = \rho$. It also implies that $0 \leq \frac{1}{\min_{e \in E_k} x_e^k} - \rho$. Indeed, let $v^k \in \{0, 1\}^{E_k}$ be the indicator vector of the base B_k at step k . Since v^k lies in the base polytope, we must have $v^k(H) \leq |r(H)|$. $x^k|_H$ is a convex combination of vectors $v^i|_H$ for $i \in [k]$, hence $x^k(H) \leq r(H) = \frac{1}{\rho}|H|$ and so $\min_{e \in H} x_e^k \leq \frac{1}{\rho}$.
- Eq. (71a) and (71b) imply that

$$\frac{1}{x_e^k} \leq \rho(1 + \varepsilon) \quad \text{if } e \in E_k, \varepsilon \in (0, 1] \text{ and } k \geq \frac{20\rho \log m}{\varepsilon^2} \quad (72a)$$

$$x_e^k \leq \frac{1 + \varepsilon}{\rho} \quad \text{if } e \in E_k \cap H, \varepsilon \in (0, 1] \text{ and } k \geq \frac{6\rho \log m}{\varepsilon^2} \quad (72b)$$

Indeed, consider an edge $e \in E_k$ for the first claim or an edge $e \in E_k \cap H$ for the second claim. The definitions of A_k^α, B_k^α and (71a),(71b) imply that

$$(1 + \alpha)^{k\left(\frac{1}{\rho} - x_e^k\right)} \leq m(1 + \alpha)^{k/\rho} \left(1 - \frac{\alpha}{\rho(1 + \alpha)}\right)^k$$

or equivalently

$$-x_e^k \cdot \log(1 + \alpha) \leq \frac{\log m}{k} + \log \left(1 - \frac{\alpha}{\rho(1 + \alpha)}\right) \quad (73)$$

Suppose that (72a) is false, i.e. $x_e^k < \frac{1}{\rho(1 + \varepsilon)}$. Plugging this into (73) with $\alpha \in (0, \varepsilon)$ gives

$$\begin{aligned} \frac{1}{\rho(1 + \varepsilon)} > x_e^k &\geq -\frac{\log m}{k \log(1 + \alpha)} - \frac{\log \left(1 - \frac{\alpha}{\rho(1 + \alpha)}\right)}{\log(1 + \alpha)} \\ &\geq -\frac{\log m}{k(\alpha - \alpha^2)} + \frac{\alpha/(\rho(1 + \alpha))}{\alpha} = \frac{1}{\rho(1 + \alpha)} - \frac{\log m}{k\alpha(1 - \alpha)} \end{aligned}$$

where we used inequalities $\log(1 + t) \leq t$ for $t > -1$ and $\log(1 + t) \geq t - \frac{t^2}{2}$ for $t \geq 0$. We conclude that $k < \frac{(1 + \alpha)(1 + \varepsilon)}{\alpha(1 - \alpha)(\varepsilon - \alpha)} \cdot \rho \log m$. Setting $\alpha = \varepsilon/4$ and recalling that $\varepsilon \in (0, 1]$ gives $k < \frac{20\rho \log m}{\varepsilon^2}$, contradicting the choice of k .

Now suppose that (72b) is false, i.e. $x_e^k > \frac{1}{\rho(1 - \varepsilon)}$. Plugging this into (73) with $\alpha = -\beta$, $\beta > 0$ gives

$$\begin{aligned} \frac{1 + \varepsilon}{\rho} < x_e^k &\leq \frac{\log m}{-k \log(1 - \beta)} + \frac{\log \left(1 + \frac{\beta}{\rho(1 - \beta)}\right)}{-\log(1 - \beta)} \\ &\leq \frac{\log m}{k\beta} + \frac{\beta/(\rho(1 - \beta))}{\beta} = \frac{1}{\rho(1 - \beta)} + \frac{\log m}{k\beta} \end{aligned}$$

We conclude that $k < \frac{1 - \beta}{\beta(\varepsilon - \beta - \varepsilon\beta)} \rho \log m$. Setting $\beta = \varepsilon/4$ and recalling that $\varepsilon \in (0, 1]$ gives $k < \frac{6\rho \log m}{\varepsilon^2}$, contradicting the choice of k .

We now proceed with the induction argument. Checking the base case $k = 0$ is straightforward. Suppose that (71) holds for $k \geq 0$; let us prove it for $k + 1$. Let v be the indicator vector of the next base B_{k+1} in the packing. Then, $x^{k+1} = \frac{kx^k + v}{k+1}$. Let x^* be the optimal loads vector in M_k . Eq. (71c) of the induction hypothesis implies that $\rho(M_k) = \rho$ and hence $x_e^* \geq 1/\rho$ for any $e \in E_k$.

We can write

$$\begin{aligned}
A_{k+1} &= \sum_{e \in E_{k+1}} (1+\alpha)^{(k+1)\left(\frac{1}{\rho} - \frac{kx_e^k + v_e}{k+1}\right)} \leq (1+\alpha)^{(k+1)/\rho} \sum_{e \in E_k} (1+\alpha)^{-kx_e^k} (1+\alpha)^{-v_e} \\
&\stackrel{(a)}{=} (1+\alpha)^{(k+1)/\rho} \left(\sum_{e \in E_k} (1+\alpha)^{-kx_e^k} - \sum_{e \in E_k} v_e \frac{\alpha}{1+\alpha} (1+\alpha)^{-kx_e^k} \right) \\
&\stackrel{(b)}{\leq} (1+\alpha)^{(k+1)/\rho} \left(\sum_{e \in E_k} (1+\alpha)^{-kx_e^k} - \sum_{e \in E_k} x_e^* \frac{\alpha}{1+\alpha} (1+\alpha)^{-kx_e^k} \right) \\
&\stackrel{(c)}{\leq} (1+\alpha)^{(k+1)/\rho} \left(\sum_{e \in E_k} (1+\alpha)^{-kx_e^k} - \sum_{e \in E_k} \frac{1}{\rho} \frac{\alpha}{1+\alpha} (1+\alpha)^{-kx_e^k} \right) \\
&= A_k (1+\alpha)^{1/\rho} \left(1 - \frac{\alpha}{\rho(1+\alpha)} \right) \stackrel{(d)}{\leq} m(1+\alpha)^{k+1/\rho} \left(1 - \frac{\alpha}{\rho(1+\alpha)} \right)^{k+1}
\end{aligned}$$

Here in (a) we used the fact that v is a 0/1 vector, and hence $(1+\alpha)^{-v_e} = 1 - v_e \frac{\alpha}{1+\alpha}$. The inequality in (b) follows from the following observation: the ordering of the edges with respect to the weights $w_e := -\frac{\alpha}{1+\alpha} (1+\alpha)^{-kx_e^k}$ coincides with the ordering of the edges with respect to the weights x_e^k , and hence v minimizes the value of $\sum_e w_e v_e$ over the base polytope. Finally, in (d) we used eq. 71a of the induction hypothesis. We proved eq. (71a) of the induction step.

To prove eq. (71b), we use the same derivation as above but replace A_i^α with B_i^α , and summations $\sum_{e \in E_k}$ and $\sum_{e \in E_{k+1}}$ with respectively $\sum_{e \in H}$ and $\sum_{e \in H \cap E_{k+1}}$. (Note that $H \cap E_k = H$ by the induction hypothesis). In (c) we now have an equality, since $x_e^* = 1/\rho$ for all $e \in H$. Let us justify the inequality in (b). Recall that x^* is a convex combination of some collection of bases. Consider any base B from this collection. It suffices to show that we have $\sum_{e \in H} w_e v_e = \sum_{e \in B^{k+1} \cap H} w_e \leq \sum_{e \in B \cap H} w_e$. The structure of x^* established in Theorem 2.4 implies that $B \cap H$ is a base in $M|H$. Set $B^{k+1} \cap H$ can be extended to a minimum-weight base $B' \supseteq B^{k+1} \cap H$ of $M|H$, with $\sum_{e \in B'} w_e \leq \sum_{e \in B \cap H} w_e$. Recall that for eq. (71b) we are assuming that $\alpha < 0$, and hence all edge weights w_e are nonnegative. The claim follows.

It remains to prove that $H \subseteq E_{k+1}$. Since $\rho \in [\rho^-, \rho^+]$ by assumption, each removed edge $e \in E_k$ satisfies $x_e^k > 2/\rho^- \geq 2/\rho$ and $k \geq 24\rho^+ \log m \geq 24\rho \log m$. Property (72b) with $\varepsilon = 1$ implies that $x_e^* > \frac{1}{\rho}$, and thus such edge is not in H .

C Dynamic minimum-weight pseudoforest

In this section we prove the following result.

Theorem C.1. *There exists an algorithm for maintaining a minimum-weight maximal pseudoforest in a dynamically changing graph that uses $O(\log n)$ time per update (edge insertion or deletion).*

We will rely on the top trees data structure for dynamic forests of Alstrup-Holm-De Lichtenberg-Thorup [1] that supports the following operations.

- **tree**(v): returns a pointer to the tree containing v .
- **link**(u, v, w): connects nodes u, v of distinct trees with an edge (u, v) of weight w .
- **cut**(e): removes existing edge e .

- **findmax**(u, v): returns an edge of maximum weight on the path between nodes u, v in the same tree.
- **findmin**(v): returns an edge of minimum weight in the tree containing node v .

Each one of them can be implemented in $O(\log n)$ worst-case time. (For operation **findmax**(u, v), see [1, Theorem 2.2]. For operation **findmin**(v), see [1, Theorem 2.4]. Operation **tree**(v) follows easily from the implementation, and can also be simulated using **findmin**(v): we can treat the minimum-weight edge in a given tree as the unique identifier of this tree after making all edge weights unique.) Note that one can check whether nodes u, v belong to the same tree in $O(\log n)$ time by comparing pointers **tree**(u) and **tree**(v).

We now proceed with the proof of Theorem C.1. Below we describe how updating graph $G = (V, E, w)$ affects a minimum-weight pseudoforest P . For a subset of edges $F \subseteq E$ we denote $V(F)$ to be the set of nodes incident to at least one edge in F . We also let P' be the pseudoforest after the update.

- Inserting edge $e = (u, v)$ with weight $w(e)$. There are at most two paths in P between u and v . Let Q be of them (or let $Q = \emptyset$ if u, v are in different components). For $w \in \{u, v\}$ let C_w be the cycle of the component to which w belongs (or $C_w = \emptyset$ if this component is acyclic). Define $F = Q \cup C_u \cup C_v$. Clearly, $P + e - f$ is a pseudoforest for an edge $f \in E - e$ if and only if $f \in F$. Therefore, we can take $f \in \arg \max_{f \in F} w(f)$ and then set $P' = P + e - f$ if $w(e) < w(f)$, or $P' = P$ otherwise.
- Removing edge e from G . If $e \notin P$ then $P' = P$. Otherwise let C be the component of P containing e . Define set $A \subseteq C - e$ as follows.
 - (1) If e belongs to a cycle in C then let $A = C$.
 - (2) Otherwise removing e from C splits C into two components. At least one of them must be acyclic; let A be such component. (Note, if both components are acyclic then $V(C)$ does not have incident edges in $E - P$ due to maximality of P).
 Let F be the set of edges that are incident to a node in $V(A)$ and are not present in P . Clearly, $P - e + f$ is a pseudoforest for an edge $f \in E - e$ if and only if $f \in F$. Therefore, if F is empty then $P' = P - e$, otherwise take $f \in \arg \min_{f \in F} w(f)$ and then set $P' = P - e + f$.

To implement these operations efficiently, we will store the following information. For each component C of P we will maintain some edge e_C on the cycle in C ; if C is acyclic then we let $e_C = \emptyset$. The tree $T_C = C - e_C$ will be stored in a top tree data structure $\mathcal{D}$, with edge weights as in the original graph. The edge e_C will be stored at the pointer corresponding to this tree. In addition, we will maintain a separate top tree data structure $\mathcal{D}'$ that will store trees T_C as above augmented with edges $(v, \bar{v})$ for all nodes v , where $\bar{v}$ is a unique copy of v . At each edge $(v, \bar{v})$ we will store a priority queue of edges incident to v in G and not present in P , with edge weights as the key. The minimum of these weights will be the weight of $(v, \bar{v})$ in $\mathcal{D}'$, while the weight of other edges in $\mathcal{D}'$ will be $+\infty$.

We claim that after each graph modification described above data structures $\mathcal{D}$ and $\mathcal{D}'$ can be updated via a constant number of calls to operations on dynamic trees. Consider, for example, the case when edge (u, v) is added to G but not to P . We then remove edge $(u, \bar{u})$ from $\mathcal{D}'$, insert (u, v) into the priority queue at $(u, \bar{u})$, and then insert $(u, \bar{u})$ back into $\mathcal{D}'$ with the new weight. The same is done for v . All other operations are similarly straightforward.

Finding $\arg \max_{f \in F} w(f)$ during edge insertion can be done by querying $\mathcal{D}$, while finding $\arg \min_{f \in F} w(f)$ during edge deletion can be done by querying $\mathcal{D}'$. This concludes the proof of Theorem C.1.

Acknowledgements

Pavel Arkhipov is grateful to Monika Henzinger, who originally encouraged him to think about tree packings.

References

- [1] Stephen Alstrup, Jacob Holm, Kristian De Lichtenberg, and Mikkell Thorup. Maintaining information in fully dynamic trees with top trees. *ACM Transactions on Algorithms*, 1(2):243–264, 2005.
- [2] Edvin Berglin and Gerth Stølting Brodal. A simple greedy algorithm for dynamic graph orientation. In *28th International Symposium on Algorithms and Computation (ISAAC)*, pages 12:1–12:12, 2017.
- [3] Sayan Bhattacharya, Monika Henzinger, Danupon Nanongkai, and Charalampos Tsourakakis. Space- and time-efficient algorithm for maintaining dense subgraphs on one-pass dynamic streams. In *Proceedings of the Forty-seventh Annual ACM Symposium on Theory of Computing (STOC)*, pages 173–182, 2015.
- [4] Immanuel M Bomze, Francesco Rinaldi, and Damiano Zeffiro. Frank–Wolfe and friends: a journey into projection-free first-order optimization methods. *4OR*, 19(3):313–345, 2021.
- [5] Ruoxu Cen, Henry Fleischmann, George Z. Li, Jason Li, and Debmalya Panigrahi. Fast algorithms for graph arboricity and related problems. In *Proceedings of the 66th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, 2025.
- [6] Ruoxu Cen, Jason Li, and Debmalya Panigrahi. Network unreliability in almost-linear time. In *Proceedings of the 57th Annual ACM Symposium on Theory of Computing*, STOC ’25, page 120–131, New York, NY, USA, 2025. Association for Computing Machinery.
- [7] Chandra Chekuri, Aleksander Bjørn Christiansen, Jacob Holm, Ivor van der Hoog, Kent Quanrud, Eva Rotenberg, and Chris Schwiegelshohn. Adaptive out-orientations with applications. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3062–3088. SIAM, 2024.
- [8] Chandra Chekuri and Kent Quanrud. $(1-\varepsilon)$ -approximate fully dynamic densest subgraph: linear space and faster update time. *arXiv preprint arXiv:2210.02611*, 2022.
- [9] Aleksander B. G. Christiansen and Eva Rotenberg. Fully-dynamic $\alpha + 2$ arboricity decompositions and implicit colouring. In *49th International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 42:1–42:20, 2022.
- [10] Tijn de Vos and Aleksander B. G. Christiansen. Tree-packing revisited: Faster fully dynamic min-cut and arboricity. In *Proceedings of the 2025 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2025.
- [11] Jack Edmonds. Minimum partition of a matroid into independent subsets. *Journal of Research of the National Bureau of Standards Section B Mathematics and Mathematical Physics*, page 67, 1965.
- [12] Alessandro Epasto, Silvio Lattanzi, and Mauro Sozio. Efficient densest subgraph computation in evolving graphs. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*, pages 300–310, 2015.

- [13] Marguerite Frank and Philip Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3(1-2):95–110, 1956.
- [14] Robert M Freund and Paul Grigas. New analysis and results for the Frank–Wolfe method. *Mathematical Programming*, 155(1):199–230, 2016.
- [15] Satoru Fujishige. *Submodular Functions and Optimization*, volume 58 of *Annals of Discrete Mathematics*. Elsevier, Amsterdam, 2nd edition, 2005.
- [16] Elfarouk Harb, Kent Quanrud, and Chandra Chekuri. Convergence to lexicographically optimal base in a (contra)polymatroid and applications to densest subgraph and tree packing. In *European Symposium on Algorithms (ESA)*, pages 181–193, 2023.
- [17] Meng He, Ganggui Tang, and Norbert Zeh. Orienting dynamic graphs, with applications to maximal matchings and adjacency queries. In *25th International Symposium on Algorithms and Computation (ISAAC)*, pages 128–140, 2014.
- [18] Jacob Holm, Kristian De Lichtenberg, and Mikkel Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge and biconnectivity. *J. ACM*, 47(4):723–760, 2001.
- [19] David R. Karger. Minimum cuts in near-linear time. *J. ACM*, 47(1):46–76, January 2000.
- [20] Tsvi Kopelowitz, Robert Krauthgamer, Ely Porat, and Shay Solomon. Orienting fully dynamic graphs with worst-case time bounds. In *41st International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 532–543, 2014.
- [21] Simon Lacoste-Julien and Martin Jaggi. On the global linear convergence of Frank-Wolfe optimization variants. *Advances in neural information processing systems*, 28, 2015.
- [22] Danupon Nanongkai, Thatchaphol Saranurak, and Christian Wulff-Nilsen. Dynamic minimum spanning forest with subpolynomial worst-case update time. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 950–961, 2017.
- [23] James G. Oxley. *Matroid Theory (Oxford Graduate Texts in Mathematics)*. Oxford University Press, Inc., USA, 2006.
- [24] Sebastian Pokutta. The Frank-Wolfe algorithm: a short introduction. *Jahresbericht der Deutschen Mathematiker-Vereinigung*, 126(1):3–35, 2024.
- [25] Saurabh Sawlani and Junxing Wang. Near-optimal fully dynamic densest subgraph. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing (STOC)*, pages 181–193, 2020.
- [26] Mikkel Thorup. Fully-dynamic min-cut*. *Combinatorica*, 27:91–127, 02 2007.
- [27] Mikkel Thorup. Minimum k-way cuts via deterministic greedy tree packing. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing, STOC '08*, page 159–166, New York, NY, USA, 2008. Association for Computing Machinery.
- [28] Mikkel Thorup and David R. Karger. Dynamic graph algorithms with applications. In *7th Scandinavian Workshop on Algorithm Theory (SWAT)*, pages 1–9, 2000.




Figure 5: Convergence plots of the MST packing for an example graph.




Figure 6: Convergence plots of the MST packing for a random graph $G(n, p)$, $n = 50$, $p = 0.2$.




Figure 7: Convergence plots of the MST packing for a random graph $G(n, p)$, $n = 50$, $p = 0.9$.




Figure 8: Convergence plots of the MST packing for a square grid graph.