



Faster Cut Sparsification of Weighted Graphs

Sebastian Forster¹ · Tijn de Vos¹

Received: 25 August 2022 / Accepted: 12 October 2022 / Published online: 1 November 2022
© The Author(s) 2022

Abstract

A cut sparsifier is a reweighted subgraph that maintains the weights of the cuts of the original graph up to a multiplicative factor of $(1 \pm \epsilon)$. This paper considers computing cut sparsifiers of weighted graphs of size $O(n \log(n)/\epsilon^2)$. Our algorithm computes such a sparsifier in time $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$, both for graphs with polynomially bounded and unbounded integer weights, where $\alpha(\cdot)$ is the functional inverse of Ackermann’s function. This improves upon the state of the art by Benczúr and Karger (SICOMP, 2015), which takes $O(m \log^2(n))$ time. For unbounded weights, this directly gives the best known result for cut sparsification. Together with preprocessing by an algorithm of Fung et al. (SICOMP, 2019), this also gives the best known result for polynomially-weighted graphs. Consequently, this implies the fastest approximate min-cut algorithm, both for graphs with polynomial and unbounded weights. In particular, we show that it is possible to adapt the state of the art algorithm of Fung et al. for unweighted graphs to weighted graphs, by letting the partial maximum spanning forest (MSF) packing take the place of the Nagamochi–Ibaraki forest packing. MSF packings have previously been used by Abraham et al. (FOCS, 2016) in the dynamic setting, and are defined as follows: an M -partial MSF packing of G is a set $\mathcal{F} = \{F_1, \dots, F_M\}$, where F_i is a maximum spanning forest in $G \setminus \bigcup_{j=1}^{i-1} F_j$. Our method for computing (a sufficient estimation of) the MSF packing is the bottleneck in the running time of our sparsification algorithm.

Keywords Graph algorithms · Cut sparsification · Maximum spanning forest · Minimum cut

1 Introduction

In many applications, graphs become increasingly large, hence storing and working with such graphs becomes a challenging problem. One strategy to deal with this

✉ Tijn de Vos
tdevos@cs.sbg.ac.at

¹ Department of Computer Science, University of Salzburg, Salzburg, Austria

issue is graph sparsification, where we model the graph by a sparse set of (reweighted) edges that preserve certain properties. Especially because the aim is to work with large input graphs, this process should be efficient with respect to the graph size. Among the different types of graph sparsifiers, there are spanners (preserving distances, see e.g. [1–4]), resistance sparsifiers (preserving effective resistances, see e.g. [5]), cut sparsifiers (preserving cuts, see e.g. [6–8]), and spectral sparsifiers (preserving Laplacian quadratic forms, see e.g. [9–12]). This paper focuses on cut sparsifiers, as first introduced by Benczúr and Karger in [6]. We say that a (reweighted) subgraph $H \subseteq G$ is a $(1 \pm \epsilon)$ -cut sparsifier for a weighted graph G if for every cut C , the total weight $w_H(C)$ of the edges of the cut in H is within a multiplicative factor of $1 \pm \epsilon$ of the total weight $w_G(C)$ of the edges of the cut in G .

The main approach to compute cut sparsifiers uses the process of *edge compression*: each edge $e \in E$ is part of the sparsifier with some probability p_e , and if selected obtains weight $w(e)/p(e)$. It is immediate that such a scheme gives a sparsifier in expectation, but it has to be shown that the result is also a sparsifier with high probability. The main line of research has been to select good *connectivity estimators* λ_e for each edge such that sampling with $p_e \sim 1/\lambda_e$ yields a good sparsifier. The simplest such result is by Karger [13], where we sample uniformly with each λ_e equal to the weight of the (global) minimum cut. Continuing along these lines are parameters as: edge connectivity [8], strong connectivity [6, 7], electrical conductance [10], and Nagamochi-Ibaraki (NI) indices [8, 14, 15]. The challenge within the approach of edge compression is to find a connectivity estimator that results in a sparse graph, but can be computed fast.

For weighted graphs, there are roughly three regimes for sparsification. The first regime consists of cut sparsifiers of size $O(n \log^2(n)/\epsilon^2)$. Fung, Hariharan, Harvey, and Panigrahi [8, 16] show that sparsifiers of this asymptotic size can be computed in linear time for polynomially-weighted graphs. For this they introduce a general framework of cut sparsification with a connectivity estimator, see Sect. 2.1. For unbounded weights, Hariharan and Panigrahi [17] give an algorithm to compute a sparsifier of size $O(n \log^2(n)/\epsilon^2)$ in time $O(m \log^2(n)/\epsilon^2)$.

The second regime consists of cut sparsifiers of size $O(n \log(n)/\epsilon^2)$. Benczúr and Karger [6, 7] show that these can be computed in time $O(m \log^2(n))$ for polynomially-weighted graphs, and in time $O(m \log^3(n))$ for graphs with unbounded weights. Note that these results can be optimized by preprocessing with the algorithms for the first regime.

A third regime, consists of sparsifiers of size $O(n/\epsilon^2)$. The known constructions in this regime yield *spectral* sparsifiers, which are more general than cut sparsifiers. Spectral sparsification was first introduced by Spielman and Teng in [9]. It considers subgraphs that preserve Laplacian quadratic forms. Lee and Sun [12] give an algorithm for finding $(1 \pm \epsilon)$ -spectral sparsifiers of size $O(n/\epsilon^2)$ in time $O(m \cdot \text{poly}(\log(n), 1/\epsilon))$. Analyzing their results, we believe that the poly-logarithmic factor contributes at least a factor of $\log^{10}(n)$. While this is optimal in size, both for spectral sparsifiers [18] and cut sparsifiers [19], it is not in time.

Algorithm	Size	Running time
<i>Unweighted</i>		
[8]	$O(n \log(n)/\epsilon^2)$	$O(m)$
<i>Polynomial weights</i>		
[7]	$O(n \log(n)/\epsilon^2)$	$O(m \log^2(n))$
[8]	$O(n \log^2(n)/\epsilon^2)$	$O(m)$
[8] + [7]	$O(n \log(n)/\epsilon^2)$	$O(m + n \log^4(n)/\epsilon^2)$
This paper	$O(n \log(n)/\epsilon^2)$	$O(m \log(n))$
This paper	$O(n \log(n)/\epsilon^2)$	$O(m\alpha(n) \log(m/n))$
[8] + this paper	$O(n \log(n)/\epsilon^2)$	$O(m + n (\log^2(n)/\epsilon^2) \cdot \alpha(n) \log(\log(n)/\epsilon))$
<i>Unbounded weights</i>		
[17]	$O(n \log^2(n)/\epsilon^2)$	$O(m \log^2(n)/\epsilon^2)$
[7]	$O(n \log(n)/\epsilon^2)$	$O(m \log^3(n))$
[17] + [7]	$O(n \log(n)/\epsilon^2)$	$O(m \log^2(n)/\epsilon^2 + n \log^5(n)/\epsilon^2)$
[12]	$O(n/\epsilon^2)$	$O(m \cdot \text{poly}(\log(n), 1/\epsilon))$
This paper	$O(n \log(n)/\epsilon^2)$	$O(m \log(n))$
This paper	$O(n \log(n)/\epsilon^2)$	$O(m\alpha(n) \log(m/n))$

Fig. 1 An overview of the state of the art algorithms for computing cut sparsifiers for undirected graphs with integer weights. Algorithm $A + B$ indicates that algorithm B is preprocessed with algorithm A

In this paper, we improve on the results in the second regime, both for graphs with polynomially bounded and unbounded weights.¹ For an overview of the previous best running times and our results, see Fig. 1. We present our sparsification algorithm in Sect. 4, with the special treatment of unbounded weights in Sect. 5. Our algorithm improves on the algorithm of Benczúr and Karger [6, 7] for bounded weights, which has been unchallenged for the last 25 years. It also improves on the algorithm of [17] for unbounded weights, which has been unchallenged for the last 10 years. We obtain the following theorem, where $\alpha(\cdot)$ refers to the functional inverse of Ackermann’s function, for a definition see e.g. [20]. For any realistic value x , we have $\alpha(x) \leq 4$.

Theorem 1.1 *There exists an algorithm that, given a weighted graph G and a freely chosen parameter $\epsilon \in (0, 1)$, computes a graph G_ϵ , which is a $(1 \pm \epsilon)$ -cut sparsifier for G with high probability. The running time of the algorithm is $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$ and the number of edges of G_ϵ is $O(n \log(n)/\epsilon^2)$.*

Using preprocessing with a result from [8] (see Theorem 2.5), we obtain the following corollary for polynomially-weighted graphs.

Corollary 1.2 *There exists an algorithm that, given a polynomially-weighted graph G and a freely chosen parameter $\epsilon \in (0, 1)$, computes a graph G_ϵ , which is a $(1 \pm \epsilon)$ -cut sparsifier for G with high probability. The running time of the algorithm is $O(m + n (\log^2(n)/\epsilon^2) \alpha(n) \log(\log(n)/\epsilon))$ and the number of edges of G_ϵ is $O(n \log(n)/\epsilon^2)$.*

¹ See Sect. 2.3 for our assumptions on the computational model in case of unbounded weights.

Following Benczúr and Karger [7], the computation of cut sparsifiers of graphs with fractional or even real weights can be reduced to integer weights. For the reduction see Appendix 2. Thus our algorithm also gives a speedup for such graphs. Since the integer case is the essential one, we follow prior works and only formulate our results for this particular case.

As a direct application of the cut sparsifier, we can use Theorem 1.1 and Corollary 1.2 to replace m by $n \log(n)/\epsilon^2$ in the time complexity of algorithms solving cut problems, at the cost of a $(1 \pm \epsilon)$ -approximation. We detail the effects for the minimum cut problem. Recently, Gawrychowski et al. [21] showed that one can compute the minimum cut of a weighted graph in $O(m \log^2(n))$ time. Using sparsification [7, 8] for preprocessing, the state of the art for $(1 + \epsilon)$ -approximate min-cut is $O(m + n \log^4(n)/\epsilon^2)$. When we use our new sparsification results, we obtain faster $(1 + \epsilon)$ -approximate min-cut algorithms when $m = \Omega(n \log(n)/\epsilon^2)$.

Corollary 1.3 *There exists an algorithm that, given a polynomially-weighted graph G and a freely chosen parameter $\epsilon \in (0, 1)$, with high probability computes an $(1 + \epsilon)$ -approximation of the minimum cut in time $O(m + n \log^3(n)/\epsilon^2)$.*

There exists an algorithm that, given a weighted graph G and a freely chosen parameter $\epsilon \in (0, 1)$, with high probability computes an $(1 + \epsilon)$ -approximation of the minimum cut in time $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)) + n \log^3(n)/\epsilon^2)$.

For unweighted graphs, even faster minimum cut algorithms exist: Ghaffari, Nowicki, and Thorup [22] show that we can find the minimum cut in $O(\min\{m + n \log^3(n), m \log(n)\})$ time. Combining this with the linear time cut sparsifier of Fung et al. [8], we get $(1 + \epsilon)$ -approximate minimum cut in unweighted graphs in $O(m + n \log(n) \min\{1/\epsilon + \log^2(n), \log(n)/\epsilon\})$ time.

The remainder of this article is organized as follows. The rest of the introduction consists of a technical overview of our algorithms. Section 2 contains a review of the general sparsification framework from [8] tailored to our needs, and can be skipped by readers that are already familiar with this work. We present our algorithm to compute the MSF indices in Sect. 3. This is used as a black box in our algorithm, which is presented and analyzed in Sect. 4. In Sect. 5, we show how the results of Sect. 4 generalize to graphs with unbounded weights.

1.1 Technical Overview

The high-level set-up of our sparsification algorithm is similar to the algorithm for unweighted graphs of Fung et al. [8]. Our main contribution consists of showing how to generalize this technique to weighted graphs, by using maximum spanning forest (MSF) indices instead of Nagamochi-Ibaraki (NI) indices. On a less significant note, we prove that by a tightening of the analysis one can show that the size and time bounds hold with high probability, and not only in expectation.

NI indices are defined by means of an NI forest packing: view graphs with integer weights as unweighted multigraphs, and repeatedly compute a spanning forest. The NI index is the (last) forest in which an edge appears (for details see Definition 2.4). The MSF index is also defined by a forest packing, but in this case the MSF packing:

we say $\mathcal{F} = \{F_1, \dots, F_M\}$ is an M -partial maximum spanning forest packing of G if for all $i = 1, \dots, M$, F_i is a maximum spanning forest in $G \setminus \bigcup_{j=1}^{i-1} F_j$. Now, we say that an edge e has MSF index i (w.r.t. to some (partial) MSF packing \mathcal{F}) if e appears in the i -th forest F_i of the (partial) MSF packing \mathcal{F} . The MSF index has been used previously in the context of dynamic graph sparsifiers (see Abraham et al. [23]). However, there it was only used because it rendered a faster running time, but using NI indices in the corresponding static construction would have been possible as well. In this paper, we use distinctive properties of the MSF index, and the NI index would not suffice. We show that using the MSF index, we can generalize the sparsification algorithm for unweighted graphs to an algorithm for weighted graphs, thereby demonstrating that the MSF index is a natural analogue for the NI index in the weighted setting. We provide an algorithm to compute an M -partial MSF packing in time $O(m \cdot \min(\alpha(n) \log(M), \log(n)))$ for polynomially-weighted graphs. We show that for unbounded weights we can compute a sufficient estimation, also in time $O(m \cdot \min(\alpha(n) \log(M), \log(n)))$.

An important distinction between the unweighted algorithm of Fung et al. and our weighted algorithm, is that the use of contractions to keep running times low throughout the algorithm is no longer possible: edges of different weights have to be treated differently, hence cannot be contracted. By using multiple iterations with an exponentially decreasing precision parameter we can overcome this problem.

In the case of a polynomially-weighted input graph, the algorithm consists of two main phases. In the first phase, we compute sets $F_0, F_1, \dots, F_\Gamma \subseteq E$, where edges satisfy some lower bound on the weight of any cut separating their endpoints. In the second phase, we sample edges from each set F_i with a corresponding probability.

We set a parameter $\rho = \Theta\left(\frac{\ln(n)}{\epsilon^2}\right)$ and start by computing a 2ρ -partial maximum spanning forest packing for G . We define F_0 to be the union of these 2ρ forests. We add the edges of F_0 to G_ϵ , which will become our sparsifier. We sample each of the remaining edges $E \setminus F_0$ with probability $1/2$ to construct X_1 . To counterbalance for the sampling, we will boost the weight of each sampled edge with a factor 2. Now we continue along these lines, but in each iteration we let F_i consist of an exponentially growing number of spanning forests: F_i is defined as the union of the forests in a $(2^{i+1} \cdot \rho)$ -partial MSF packing of X_i . Then, X_{i+1} is sampled from the remaining edges $X_i \setminus F_i$, where again each edge is included with probability $1/2$. We continue this process until there are sufficiently few edges left in X_{i+1} . We add these remaining edges to G_ϵ .

The second phase of the algorithm is to sample edges from the sets F_i and add these sampled edges to G_ϵ . Hereto, note that an edge e of F_i (for $i \geq 1$) was not part of F_{i-1} , meaning it was not part of any spanning forest in a $(2^i \cdot \rho)$ -partial MSF packing of X_{i-1} . This implies that for an edge $e \in F_i$ the weight of any cut C in X_{i-1} containing e is at least $2^i \cdot \rho \cdot w(e)$. Now we use the general framework for cut sparsification of Fung et al. [8], which boils down to the fact that this guarantee on the weights of cuts implies that we can sample edges from F_i with probability proportional to $1/(2^i w(e))$. We show that this results in a sufficiently sparse graph.

Intuitively, it might seem redundant to sample edges from $X_i \setminus F_i$ to form X_{i+1} . This is indeed not necessary to guarantee that the resulting graph is a sparsifier. However,

it ensures that the number of iterations is limited, which leads to better bounds on the size of the sparsifier and the running time. Since we sample edges with probability $1/2$ in each phase, we need to repeat the sampling $O(\log(m/(m_0)))$ times to get the size of X_i down to $O(m_0)$. As this number of steps depends on the initial number of edges m , we get better bounds for size and running time if m is already small. We will exploit this by preprocessing the graph with an algorithm from [8] that gives a cut sparsifier of size $O(n \log^2(n)/\epsilon^2)$ in linear time. Moreover, we can show that repeatedly calling our algorithm has no worse asymptotic time bound than calling it once, since the input graph becomes sparser very quickly. By doing so, we obtain a sparsifier of size $O(n \log(n)/\epsilon^2)$.

Since we only use that the MSF index gives a guaranteed lower bound on the connectivity of an edge, one might wonder why the NI index does not work here. After all, the NI indices of a graph can be computed in linear time, which would result in a significant speed-up. However, when computing the NI index, the weight of an edge influences the number of forests necessary, while computing the MSF index only requires the comparison of weights. Moreover, the number of trees in a MSF packing is always bounded by n . We can use this to bound the number of edges in the created sparsifier. The same technique with NI indices would make the size of the sparsifier depend on the maximum weight in the original graph.

To show that the algorithm outputs a cut sparsifier, it needs to be proven that both the sampling in the first and the second phase preserve cuts. We follow the lines of the analysis of [8], which makes use of cut projections and Chernoff bounds. We show that by partitioning the edge sets according to their weight this method extends to weighted graphs.

One part of the algorithm has remained unaddressed: the computation of the maximum spanning forests. The approach we use here is related to Kruskal's algorithm for computing minimum spanning trees [24]. We start by sketching the M -partial MSF packing algorithm for polynomial weights. We sort the edges according to their weights using radix sort in $O(m)$ time. We create M empty forests on n vertices. Starting with the heaviest edge, we add each edge e to the first forest in which it does not create a cycle. We can find this forest using a binary search in $\log(M)$ steps. By using a disjoint-forest representation for the union-find data structure necessary to carry out these steps, we achieve a total time of $O(m\alpha(n) \log(M))$.

When working with unbounded weights, the bottleneck is the initial sorting of the edges. Radix sort does not guarantee to be efficient for unbounded weights. Instead we could use a comparison-based algorithm, such as merge sort, which takes time $O(m \log(n))$. By employing a different data structure than before, we can guarantee total running time $O(m \log(n))$. However, we do not need the exact MSF indices for our sampling procedure; estimates suffice. We can apply a 'windowing' technique from [7] to split the graph into subgraphs, where we can rescale the weights to polynomial weights and apply our previously mentioned algorithm. We then achieve a total running time of $O(m\alpha(n) \log(M))$, as before. For more details on this, we refer to Sect. 3.1. So in total we have running time $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$.

2 Notation and Review

Throughout this paper, we consider $G = (V, E)$ to be an undirected, integer weighted graph on $|V| = n$ vertices with $|E| = m$ edges. We define a set of edges $C \subseteq E$ to be a *cut* if there exists a partition of the vertices V in two non-empty subsets A and B , such that C consists of all edges with one endpoint in A and the other endpoint in B . The weight of the cut is the sum of the weights of the edges of the cut: $w_G(C) = \sum_{e \in C} w_G(e)$. The *minimum cut* is defined as the cut with minimum weight. We say that a (reweighted) subgraph $H \subseteq G$ is a $(1 \pm \epsilon)$ -*cut sparsifier* for a weighted graph G if for every cut C in H , its weight $w_H(C)$ is within a multiplicative factor of $1 \pm \epsilon$ of its weight $w_G(C)$ in G . A key concept in the realm of cut sparsification is the connectivity of an edge.

Definition 2.1 Let $G = (V, E)$ be a graph, possibly weighted. We define the *connectivity* of an edge $e = (u, v) \in E$ to be the minimal weight of any cut separating u and v . We say that e is k -*heavy* if it has connectivity at least k . For a cut C , we define the k -*projection* of C to be the k -heavy edges of the cut C .

The following theorem from [8] bounds the number of distinct k -projections of a graph, it is a generalization of a preceding theorem by Karger, see [25, 26]. This result can be useful when showing that cuts are preserved by a sampling scheme. This is due to the fact that while there may be exponentially many different cuts, this theorem shows that there are only polynomially many cut projections. Hence if one can reduce a claim for cuts to their k -projections, a high probability bound can be obtained through the application of a Chernoff bound.

Theorem 2.2 For any $k \geq \lambda$ and any $\eta \geq 1$, the number of distinct k -projections in cuts of weight at most ηk in a graph G is at most $n^{2\eta}$, where λ is the weight of a minimum cut in G .

Throughout this paper, we say a statement holds *with high probability* (w.h.p.) if it holds with probability at least $1 - n^{-c}$, for some constant c . This constant can be modified by adjusting the constants hidden in asymptotic notation.

2.1 A General Framework for Cut Sparsification

We review the general framework for cut sparsification as presented in [8]. This section does not contain new results, and can be skipped by readers that are only interested in our contribution.

The framework shows that edges can be sampled using different notions of connectivity estimators. Although this scheme provides one proof for the validity of multiple parameters, it might be worth noting that an analysis tailored to the used connectivity estimator might provide a better result. For example, when the framework is applied with ‘edge strengths’, it produces a sparsifier of size $O(n \log^2(n)/\epsilon^2)$, a $\log(n)$ factor denser than the edge strength-based sparsifier of Benczúr and Karger [7].

Let $G = (V, E)$ be a graph with integer weights, and let $\epsilon \in (0, 1)$, $c \geq 1$ be parameters, corresponding to the approximation precision and failure probability

respectively. Given a parameter γ (possibly depending on n) and an integer-valued parameter λ_e for each $e \in E$. We obtain G_ϵ from G by independently *compressing* each edge e with parameter

$$p_e = \min \left(1, \frac{16(c + 7)\gamma \ln(n)}{0.38\lambda_e \epsilon^2} \right).$$

Compressing an edge e with weight $w(e)$ consists of sampling r_e from a binomial distribution with parameters $w(e)$ and p_e . If $r_e > 0$, we include the edge in G_ϵ with weight r_e/p_e .

In the following we describe a sufficient condition on the parameters γ and λ_e such that G_ϵ is a $(1 \pm \epsilon)$ -cut sparsifier for G with probability at least $1 - 4/n^c$. Hereto we partition the edges according to their value λ_e :

$$\Lambda := \left\lfloor \log \left(\max_{e \in E} \{\lambda_e\} \right) \right\rfloor;$$

$$R_i := \{e \in E : 2^i \leq \lambda_e \leq 2^{i+1} - 1\}.$$

Let $\mathcal{G} = \{G_i = (V, E_i) : 1 \leq i \leq \Lambda\}$ be a set of integer-weighted subgraphs such that $R_i \subseteq G_i$. Moreover suppose that $w_{G_i}(e) \geq w_G(e)$ for each $e \in R_i$. For a given set of parameters $\Pi = \{\pi_1, \dots, \pi_\Lambda\} \subseteq \mathbb{R}^\Lambda$, we define

- Π -connectivity: each edge $e \in R_i$ is π_i -heavy in G_i ;
- γ -overlap: for any cut C ,

$$\sum_{i=0}^{\Lambda} \frac{e_i^{(C)} 2^{i-1}}{\pi_i} \leq \gamma \cdot e^{(C)},$$

where $e^{(C)} = \sum_{e \in C} w_G(e)$ and $e_i^{(C)} = \sum_{e \in C \cap E_i} w_{G_i}(e)$.

The following theorem shows that compressing with parameters adhering to these conditions gives a cut sparsifier with high probability.

Theorem 2.3 (See [8, Theorem 1.14]) *Fix the parameters γ and λ_e for each edge e . If there exists \mathcal{G} satisfying Π -connectivity and γ -overlap for some Π , then G_ϵ is a $(1 \pm \epsilon)$ -cut sparsifier for G , with probability at least $1 - 4/n^c$, where G_ϵ is obtained by edge compression using parameters γ and λ_e 's.*

2.2 A First Application of the Framework

In this section, we review the application of the framework from the previous section with *Nagamochi–Ibaraki (NI) indices* as parameters, as presented in [8]. As the name suggests, NI indices were first introduced by Nagamochi and Ibaraki [14, 15]. The algorithm they provide gives a graph partitioning into forests, and subsequently a corresponding index for each edge, called the NI index.

Definition 2.4 Let $G = (V, E)$ be a graph, possibly weighted. We say an edge-disjoint sequence F_1, F_2, \dots of forests is a *Nagamochi-Ibaraki forest packing* for G if F_i is a spanning forest for $G \setminus \bigcup_{j=1}^{i-1} F_j$, where the weights of $\bigcup_{j=1}^{i-1} F_j$ are subtracted of G . If G is a weighted graph, each edge e must be contained in $w(e)$ contiguous forests. We define the *NI index*, denoted by l_e , to be the index of the (last if weighted) forest in which e appears.

Nagamochi and Ibaraki show that the NI indices can be computed in linear time for unweighted graphs and in $O(m + n \log(n))$ time for weighted graphs, see [14, 15]. As is shown in [8], we can use the NI index as the connectivity estimator in the sparsification framework to obtain the following result.

Theorem 2.5 *Let $G = (V, E)$ be a weighted graph, and let $\epsilon > 0$ be a constant. Let G_ϵ be obtained by independently compressing each edge with parameter $p_e = \min(1, \rho/l_e)$, where $\rho = \frac{224}{0.38} \ln(n)/\epsilon^2$ and l_e is the NI index. Then G_ϵ is a $(1 \pm \epsilon)$ -cut sparsifier for G with high probability.*

The sampling itself takes at most $O(m)$ time, as explained in Sect. 4.4. As the NI indices can be computed in $O(m + n \log(n))$ time, this implies that the total running time is $O(m + n \log(n))$. As a graph with $m \leq n \log(n)$ is already sparse, we can assume $m > n \log(n)$. Thus, for our purposes, the total running time is simply $O(m)$.

Next we provide a bound for the number of edges in the sparsifier G_ϵ . [8] proves this same bound in expectation, we provide a proof for this bound ‘with high probability’.

Lemma 2.6 *With high probability, the size of the graph G_ϵ in Theorem 2.5 is $O(n \log^2(n)/\epsilon^2)$.*

Proof Let $v \in V$ be a vertex with degree $d_v \geq O(\log^2(n)/\epsilon^2)$ in G . We denote the degree of v in G_ϵ by d'_v and we write $d' := \max_{v \in V} d'_v$. For each neighbor u of v in G , we compress the edge $e = (u, v)$ with parameter $p_e = \min\left(1, \frac{224 \ln(n)}{0.38 \epsilon^2 l_e}\right)$, where l_e is the NI index of e . For each edge, the probability that it remains after compression is $1 - (1 - p_e)^{w_e}$. From Bernoulli’s inequality we see $1 - (1 - p_e)^{w_e} \leq w_e p_e$. Let Y_e be the random variable that is 1 if e remains, and 0 else. We note that $\mathbb{E}\left[\sum_{e:v \in e} Y_e\right] \leq \frac{224}{0.38} \ln^2(n)/\epsilon^2$. Now we apply a Chernoff bound (Theorem A.2) to obtain

$$\mathbb{P}\left[d'_v \geq \delta \frac{224}{0.38} \ln^2(n)/\epsilon^2\right] \leq \exp\left(-0.38\delta \frac{224}{0.38} \ln^2(n)/\epsilon^2\right) = n^{-224\delta \ln(n)/\epsilon^2}.$$

Using a union bound we get the desired result

$$\mathbb{P}\left[d' \leq \delta \frac{224}{0.38} \ln^2(n)/\epsilon^2\right] \geq 1 - n^{1-224\delta \ln(n)/\epsilon^2}.$$

Consequently, we obtain that with high probability the number of edges of the sparsifier is at most $O(n \log^2(n)/\epsilon^2)$.

The state of the art for polynomially-weighted graphs is achieved by postprocessing this result with the algorithm by Benczúr and Karger [7]. Thus our improvement on [7] leads to an overall improved result.

2.3 The Computational Model

If we have an input graph $G = (V, E)$ with weights $w: E \rightarrow \{1, \dots, W\}$, we assume our computational model has word size $\Theta(\log(W) + \log(n))$. Note that for polynomial weights, this comes down to a word size of $\Theta(\log(n))$. Moreover, we assume that basic operations on such words have uniform cost, i.e., they can be performed in constant time. In particular, these basic operations are addition, multiplication, inversion, logarithm, and sampling a random bit string of word size precision. Such assumptions are in line with previous work [7, 8], where they are made implicitly. As detailed in Section 4.4, this can be done efficiently using the algorithm of Devroye [29], which uses addition, multiplication, inversion, logarithm, and sampling a random bit string of word size precision.

3 A Maximum Spanning Forest Packing

An important primitive in our algorithm is the use of the maximum spanning forest (MSF) index. The concept is similar to the Nagamochi-Ibaraki index, the important difference is that an edge e with weight $w(e)$ appears in $w(e)$ different NI forests. This means that the number of NI forests depends on the numerical values of the edge weights, and thus can grow far beyond $O(n)$. On the other hand, the number of maximum spanning forests in a MSF packing is bounded by the maximum degree in the graph, hence also by n . While this already has noteworthy implications for polynomially-weighted graphs, it is even more significant for superpolynomially-weighted graphs. We believe that this property might make them suitable for applications other than presented here.

Definition 3.1 Let $G = (V, E)$ be a weighted graph. We say $\mathcal{F} = \{F_1, \dots, F_M\}$ is an M -partial maximum spanning forest packing of G if for all $i = 1, \dots, M$, F_i is a maximum spanning forest in $G \setminus \bigcup_{j=1}^{i-1} F_j$. If we have that $\bigcup_{i=1}^M F_i = G$, then we call \mathcal{F} a (complete) maximum spanning forest packing of G . Moreover, for $e \in E$ we denote the MSF index of e (w.r.t. \mathcal{F}) by f_e , i.e., f_e is the unique index such that $e \in F_{f_e}$.

Note that we do not demand the $F_i \in \mathcal{F}$ to be non-empty, as this suits notation best in our applications. Also note that a (partial) MSF packing is fully determined by the MSF indices.

The following theorem states that computing the MSF indices up to M takes $O(m\alpha(n) \log(M))$ time for polynomially-weighted graphs.

Theorem 3.2 Let $G = (V, E)$ be a graph, where we allow parallel edges but no self-loops, and we suppose $m \leq n^2$. Suppose we have weights $w: E \rightarrow \{1, \dots, n^c\}$ for some $c \geq 0$. Then, for any $M > 0$, there exists an algorithm that computes an M -partial MSF packing in $O(m(\alpha(n) \log(M) + c))$ time.

Proof The outline of the algorithm is as follows.

1. Sort the edges by weight in descending order using radix sort in base n .²
2. Create empty forests F_1, \dots, F_M .
3. Iterate over the edges in descending order and for each edge $e = (u, v)$ do the following:
 - (a) Find the smallest index i such that u and v are not connected in F_i .
 - (b) Store i as the MSF index f_e of e . If u and v are connected in every F_i , store $f_e > M$.
 - (c) Add e to F_i .

We need at most M trees, since we only compute an M -partial MSF packing. By using radix sort, the initial sorting takes time $O(cm)$ time (for a time bound of radix sort, see e.g. [27]). We show that the remainder of the algorithm can be executed in $O(m\alpha(n) \log(M))$ time.

For every $1 \leq i \leq M$ we maintain the non-singular components of F_i with a union-find data structure (supporting the three operations MakeSet_i , Union_i , and FindSet_i). To be precise, we use the disjoint-set forest representation of Tarjan [20] (see e.g. [27, Chapter 21]). Additionally, for every node $v \in V$ we maintain $s(v)$, the smallest index i such that $\{v\}$ is a singleton component of F_i .

For the binary search in Step 3a it is sufficient to first search over indices $i < \min\{s(u), s(v)\}$. If this search is successful and we find such an index $i < \min\{s(u), s(v)\}$, then we perform $\text{Union}_i(u, v)$. Otherwise, we have learned that $\min\{s(u), s(v)\}$ is the smallest index i such that u and v are not connected in F_i . The algorithm then proceeds as follows:

- Let $j := \min\{s(u), s(v)\}$.
- If $j = s(u)$, then we perform $\text{MakeSet}_j(u)$ and increase $s(u)$ by one.
- If $j = s(v)$ (which could also be the case in addition to $j = s(u)$), we perform $\text{MakeSet}_j(v)$ and increase $s(j)$ by one.
- Finally, we perform $\text{Union}_j(u, v)$.

Now let φ_i , χ_i , and ψ_i denote the number of MakeSet_i -, Union_i -, and FindSet_i -operations in the i -th union-find data structure, respectively. Since we use the disjoint-set forest representation, we obtain a bound on the running time (see [27, Theorem 21.14]) of $O((\varphi_i + \chi_i + \psi_i)\alpha(\varphi_i + \chi_i + \psi_i))$. To obtain a bound on the total running time for all operations in the union-find data structures, we sum over all i :

$$\begin{aligned}
 & O\left(\sum_{i=1}^M (\varphi_i + \chi_i + \psi_i)\alpha(\varphi_i + \chi_i + \psi_i)\right) \\
 & \leq O\left(\sum_{i=1}^M (\varphi_i + \chi_i + \psi_i)\alpha\left(\max_{j=1, \dots, M} (\varphi_j + \chi_j + \psi_j)\right)\right).
 \end{aligned}$$

Now observe that in total we perform at most two MakeSet -operations per edge (one for each of its endpoints) and thus $\sum_{i=1}^M \varphi_i \leq 2m$. The number of UNION -

² Note that conversion to base n takes time $O(\log_n(w(e))) \leq O(\log_n(n^c)) = O(c)$ for each edge, so total time $O(mc)$.

operations is always bounded above by $\varphi_i - 1$, so $\sum_{i=1}^M \chi_i \leq 2m$. Furthermore, by using binary search, we perform $O(\log(M))$ FINDSET-operations per edge and thus $\sum_{i=1}^M \psi_i = O(m \log(M))$. Thus we get a total time of $O((m \log(M))\alpha(m + 2n)) \leq O(m\alpha(n) \log(M))$, which holds because $\alpha(m + 2n) \leq \alpha(n^4) \leq \alpha(n) + 2$.

We therefore arrive at a total running time of $O(cm + m\alpha(n) \log(M)) = O(m(\alpha(n) \log(M) + c))$.

If we want to compute the full maximum spanning forest packing, it suffices to set M to be the maximum degree in the graph. When M is large, managing the data structures slightly differently yields a better result.

Theorem 3.3 *There exists an algorithm that, given a weighted graph $G = (V, E)$, where we allow parallel edges but no self-loops, and parameter $M > 0$, computes an M -partial MSF packing in $O(m(\log(n) + \log(M)))$ time.*

Proof We use the same algorithm as in Theorem 3.2 with two simple changes. In step 1, we use an optimal comparison-based sorting algorithm, like merge sort, instead of radix sort. This takes time $O(m \log(n))$. In steps 3 and 3, we use a linked-list representation [27, Chapter 21] instead of the disjoint-set forest representation. To analyze the running time, recall the following notation. Let φ_i , χ_i , and ψ_i denote the number of MakeSet_{*i*}-, Union_{*i*}-, and FindSet_{*i*}-operations in the i -th union-find data structure, respectively. Here MAKESET_{*i*}, UNION_{*i*}, and FINDSET_{*i*} are the operations on the component F_i . By [27, Theorem 21.1], we obtain a bound on the running time of: $O(\chi_i + \psi_i + \varphi_i \log(\varphi_i))$. We sum over all i to obtain

$$O\left(\sum_{i=1}^M \chi_i + \psi_i + \varphi_i \log(\varphi_i)\right) = O\left(\sum_{i=1}^M \chi_i + \sum_{i=1}^M \psi_i + \sum_{i=1}^M \varphi_i \log(\varphi_i)\right).$$

As before, we have $\sum_{i=1}^M \chi_i \leq m$ and $\sum_{i=1}^M \varphi_i \leq 2m$. Also note $\varphi_i \leq n$, as F_i is a forest, so we have $\sum_{i=1}^M \varphi_i \log(\varphi_i) \leq 2m \log(n)$. Again, we perform at most $O(\log(M))$ FindSet-operations per edge, hence $\sum_{i=1}^M \psi_i = O(m \log(M))$. We conclude we have total time $O(m(\log(n) + \log(M)))$.

Note that if we do not have parallel edges, then $M \leq n$, so the running time simplifies to $O(m \log(n))$. Also note that the weights no longer need to be bounded for this result. In the next section, we consider an algorithm for *sparse* graphs with unbounded weights.

3.1 An Estimation for Unbounded Weights

For our purposes we do not need the exact MSF indices, but estimates suffice. The MSF index guarantees that if an edge $e = (u, v) \in E$ has MSF index f_e , then there are at least f_e paths from u to v , where every edge on such a path has weight at least $w(e)$. We relax this, to get the guarantee that if an edge $e = (u, v) \in E$ has estimated MSF index \tilde{f}_e , then there are at least \tilde{f}_e paths from u to v , where every edge on such a path has weight at least $(1 - 1/n)w(e)$. When we only compute estimates, we can

do this faster than when we compute exact indices. The following lemma is inspired by the windowing technique of Benczúr and Karger [7], which shows that strong connectivities can be computed efficiently for graphs with unbounded weights by ‘windowing’ these weights. This means we divide the graph into subgraphs according to an estimate and compute the sought connectivity estimators in these subgraphs. Hereto, we first compute a single maximum spanning forest F for G . Now we define $d(e)$ to be the minimum weight among the edges on the path from u to v in F , where $e = (u, v)$. This can be done in total time $O(m + n)$, see [28].

Lemma 3.4 *There exists an algorithm that, given a weighted graph $G = (V, E)$ and parameter $M > 0$, computes in time $O(m\alpha(n) \log(M))$ an MSF index estimator \tilde{f}_e for each edge $e \in E' := \{e \in E : w(e) > d(e)/n\}$ with $f_e \leq M$.*

Proof We will split the graph G into graphs $G^{(D)}$ for different values of D . In each $G^{(D)}$ we compute the estimator \tilde{f}_e for some subset of edges from E' . We iteratively define D to be the highest value among the $d(e)$ for which $e \in E'$ and \tilde{f}_e has not been computed yet. We look at the subgraph $G^{(D)} = (V_D, E_D)$ defined by contracting all edges with $w(e) > D$, and deleting self loops. Moreover, we delete all edges with $w(e) \leq D/n^2$.

We claim we that for each edge $e \in E'$ with $d(e) \in (D/n, D]$ the MSF index in G is equal to the MSF index in $G^{(D)}$. First we show that these edges actually appear in $G^{(D)}$. It is clear that $w(e) \leq d(e) \leq D$, so such an edge e is not contracted. Suppose e is deleted, then $w(e) \leq D/n^2$. If $d(e) \in (D/n, D]$, then $d(e) > D/n$, hence $D/n^2 < d(e)/n$. Consequently $w(e) \leq D/n < d(e)/n$, so then $e \notin E'$.

Now that we have established that edges $e \in E'$ with $d(e) \in (D/n, D]$ appear in $G^{(D)}$, it remains to show that we can compute \tilde{f}_E . First let us remark that if $e = (u, v)$, then no (relevant) path from u to v is eliminated, since each such path must have an edge e' with $w(e') \leq d(e) \leq D$, by definition of $d(e)$. Hence the only paths that are deleted, contain an edge e' with $w(e') \leq D/n^2 < d(e)/n \leq w(e)$, hence this path does not contribute to the MSF index.

Next we compute an estimator of the MSF index in $G^{(D)}$, by computing the MSF indices in a reweighted graph. We rescale the graph by multiplying all weights with n^3/D and rounding to the closest integer. This means that we have an error in the weight of at most D/n^3 . For an edge with $D/n^2 < w(e) \leq D$, this means that the error is at most $w(e)/n$. So using Theorem 3.2, we can compute the MSF indices in this reweighted sugraph with edge weights bounded by n^3 in time $O(m'\alpha(n) \log(M))$, where there is a multiplicative error in edge weights of at most $(1 \pm 1/n)$. Note that each edge appears in at most two subgraphs, hence we have a total time of $O(m\alpha(n) \log(M))$.

4 Cut Sparsification for Weighted Graphs

In this section, we present our algorithm for computing a $(1 \pm \epsilon)$ -cut sparsifier G_ϵ for a weighted graph G . This makes use of the framework as presented in Sect. 2.1 and the maximum spanning forest packing as treated in Sect. 3. This section works

towards proving the following theorem for polynomially-weighted graphs. In Sect. 5, we will generalize the techniques of this section to graphs with unbounded weights.

Theorem 4.1 *There exists an algorithm that, given a weighted graph $G = (V, E)$, and freely chosen parameter $\epsilon > 0$, computes a graph G_ϵ , which is a $(1 \pm \epsilon)$ -cut sparsifier for G with high probability. The algorithm runs in time $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$ and the number of edges of G_ϵ is $O(n (\log(n)/\epsilon^2) \log(m/(n \log(n)/\epsilon^2)))$.*

To be precise, we give an algorithm where the given bounds on both running time and size of the sparsifier hold with high probability. By simply halting when the running time exceeds the bound, and outputting an empty graph if we exceed the size bound, this gives the result above.

To achieve a better bound on the size of the sparsifier, we repeatedly apply this theorem to the input graph, with an exponentially decreasing precision parameter.

Theorem 4.2 (Restated) *There exists an algorithm that, given a weighted graph $G = (V, E)$, and freely chosen parameter $\epsilon \in (0, 1)$, computes a graph G_ϵ , which is a $(1 \pm \epsilon)$ -cut sparsifier for G with high probability. The algorithm runs in time $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$ and the number of edges of G_ϵ is $O(n \log(n)/\epsilon^2)$.*

Proof We obtain this result by repeatedly applying the algorithm from Theorem 4.1, for a total of $k := \log^* \left(\frac{m}{n \log(n)/\epsilon^2} \right)$ times. In iteration i , we set $\epsilon_i := \epsilon/2^{k-i+2}$ and denote the output of this iteration by G_i . This means that G_i is a $(1 \pm \epsilon/2^{k-i+2})$ -cut sparsifier for G_{i-1} . In total, we see that $G_\epsilon := G_k$ is a $(1 \pm \epsilon)$ -cut sparsifier for G since

$$\begin{aligned} \prod_{i=1}^k (1 + \epsilon/2^{k-i+2}) &\leq \exp \left(\sum_{i=1}^k \log(1 + \epsilon/2^{k-i+2}) \right) \leq \exp \left(\sum_{i=1}^k \epsilon/2^{k-i+2} \right) \\ &\leq \exp \left(\epsilon \sum_{j=2}^\infty 2^{-j} \right) = \exp(\epsilon/2) \leq 1 + \epsilon, \end{aligned}$$

as $\epsilon < 1$, and

$$\begin{aligned} \prod_{i=1}^k (1 - \epsilon/2^{k-i+2}) &\geq \prod_{j=0}^\infty 1 - \frac{\epsilon/4}{2^j} &&\geq (1 - \epsilon/8) \prod_{j=1}^\infty 1 - \frac{\epsilon/4}{j^2} \\ &= (1 - \epsilon/8) \frac{\sin(\pi \sqrt{\epsilon}/2)}{\sqrt{\epsilon}/2} &&\geq (1 - \epsilon/8)(1 - \pi^2/24\epsilon) \\ &\geq 1 - (1/8 + \pi^2/24)\epsilon + \frac{\pi^2}{192}\epsilon^2 &&\geq 1 - \epsilon. \end{aligned}$$

Since $k = \log^* \left(\frac{m}{n \log(n)/\epsilon^2} \right) = O(\log^*(n))$, all bounds hold with high probability simultaneously, and thus the end result holds with high probability.

Now for the size bound, we have that

$$m_i := |E(G_i)| \leq C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \log \left(\frac{m_{i-1}}{n \log(n)/\epsilon^2} \right) \right),$$

for some constant $C > 0$, where we denote $m_0 := m$. We will show by induction that

$$m_i \leq C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \cdot 2 \log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right),$$

which means in particular that $m_k = O(n \log(n)/\epsilon^2)$. The claim for m_1 is immediate. Suppose it holds for $i - 1$, then

$$\begin{aligned} m_i &\leq C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \log \left(\frac{m_{i-1}}{n \log(n)/\epsilon^2} \right) \right) \\ &\leq C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \log \left(C \cdot 4^{k-i+3} \cdot 2 \log^{(i-1)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right) \right) \\ &= C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \left((k-i) \log(4) + \log(C \cdot 2^7) + \log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right) \right) \\ &\leq C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \cdot 2 \log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right), \end{aligned}$$

since

$$\begin{aligned} (k-i) \log(4) + \log(C \cdot 2^7) &= \left(\log^* \left(\frac{m}{n \log(n)/\epsilon^2} \right) - i \right) \log(4) + \log(C \cdot 2^7) \\ &= \log^* \left(\log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right) \log(4) + \log(C \cdot 2^7) \\ &< \log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right), \end{aligned}$$

if $\frac{m}{n \log(n)/\epsilon^2} > D$, for some constant D . This can be assumed to hold, since if $\frac{m}{n \log(n)/\epsilon^2} \leq D$, then Theorem 4.1 immediately gives the desired result. The total running time becomes of the sum of the k iterations:

$$\begin{aligned} &\sum_{i=1}^k O(m_{i-1} \cdot \min(\alpha(n) \log(m_{i-1}/n), \log(n))) \\ &= O \left(\left(m + \sum_{i=1}^{k-1} m_i \right) \cdot \min(\alpha(n) \log(m_{i-1}/n), \log(n)) \right). \end{aligned}$$

Note that

$$\begin{aligned} \sum_{i=1}^{k-1} m_i &\leq \sum_{i=1}^{k-1} C \cdot \left(\frac{n \log(n)}{\epsilon^2} 4^{k-i+2} \cdot 2 \log^{(i)} \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right) \\ &= O \left(\frac{n \log(n)}{\epsilon^2} 4^k \log \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right). \end{aligned}$$

We have $\log^*(x) = O(\log \log(x))$, hence we obtain $4^{\log^*(x)} \log(x) = O(\log^2(x)) = O(x)$. Using this with $x = \frac{m}{n \log(n)/\epsilon^2}$ gives us total running time

$$\begin{aligned} &\sum_{i=1}^k O(m_{i-1} \cdot \min(\alpha(n) \log(m_{i-1}/n), \log(n))) \\ &= O \left(\left(m + \frac{n \log(n)}{\epsilon^2} 4^k \log \left(\frac{m}{n \log(n)/\epsilon^2} \right) \right) \cdot \min(\alpha(n) \log(m/n), \log(n)) \right) \\ &= O(m \cdot \min(\alpha(n) \log(m/n), \log(n))). \end{aligned}$$

4.1 The Algorithm

To sparsify the graph, two methods of sampling are used. One of which is the framework presented in Sect. 2.1. However, instead of applying the framework to the graph directly, there is another sampling process that precedes it.

To simplify equations, let us set $\rho := \frac{(7+c)1352 \ln(n)}{0.38\epsilon^2}$. If $|E| \leq 4\rho n \log(m/(n \log(n)/\epsilon^2))$, we do nothing. That is, we return $G_\epsilon = G$. If not, we start by an initialization step and continue with an iterative process, which ends when the remaining graph becomes sufficiently small.

In the initialization step, we define $X_0 := E$. We compute a $\lfloor 2\rho \rfloor$ -partial maximum spanning forest packing $T_1, \dots, T_{\lfloor 2\rho \rfloor}$ and we define $F_0 := \bigcup_{j=1}^{\lfloor 2\rho \rfloor} T_j$. The remaining edges $Y_0 := X_0 \setminus F_0$ move on to the next phase.

In iteration i , we create X_{i+1} from Y_i by sampling each edge with probability $1/2$. Next, we compute $k_i := \rho \cdot 2^{i+1}$ maximum spanning forests T_1, \dots, T_{k_i} . We define $F_i := \bigcup_{j=1}^{k_i} T_j$, and $Y_i := X_i \setminus F_i$.

We continue until Y_i has at most $2\rho n$ edges, and set Γ to be the number of iterations. We retain all edges in F_0 . In other words: add each edge $e \in F_0$ to G_ϵ with weight $w(e)$. The edges of Y_Γ are also retained, but they need to be scaled to counterbalance the $\Gamma - 1$ sampling steps: add each edge $e \in Y_\Gamma$ to G_ϵ with weight $2^{\Gamma-1} w(e)$.

Any other edge $e \in F_i$ is at least $k_i w(e)$ -heavy in X_{i-1} , as $e \notin F_{i-1}$. We exploit this heaviness to sample from these edges using the framework. For each $e \in F_i$ we:

- Define $n_e := 2^i w(e)$ and $p_e := \min \left(1, \frac{384}{169} \frac{1}{4^i w(e)} \right)$;
- Generate r_e from the binomial distribution with parameters n_e and p_e ;
- If r_e is positive, add e to G_ϵ with weight r_e/p_e .

Algorithm 1 SPARSIFY(V, E, w, ϵ, c)

Input: An undirected graph $G = (V, E)$, with integer weights $w: E \rightarrow \mathbb{N}^+$, and parameters $\epsilon \in (0, 1)$, $c \geq 1$.

Output: An undirected weighted graph $G_\epsilon = (V, E_\epsilon)$.

```

1: Set  $\rho \leftarrow \frac{(7+c)1352 \ln(n)}{0.38\epsilon^2}$ .
2: if  $|E| \leq 4\rho n \log(m/(n \log(n)/\epsilon^2))$  then
3:   return  $G_\epsilon = G$ .
4: end if
5: Compute an  $\lfloor 2\rho \rfloor$ -partial maximum spanning forest packing
    $T_1, T_2, \dots, T_{\lfloor 2\rho \rfloor}$  for  $G$ .
6: Set  $i \leftarrow 0$ .
7: Set  $X_0 \leftarrow E$ .
8: Set  $F_0 \leftarrow \bigcup_{j=1}^{\lfloor 2\rho \rfloor} T_j$ .
9: Set  $Y_0 \leftarrow X_0 \setminus F_0$ .
10: while  $|Y_i| > 2\rho n$  do
11:   Sample each edge in  $Y_i$  with probability  $1/2$  to construct  $X_{i+1}$ .
12:    $i \leftarrow i + 1$ .
13:   Set  $k_i \leftarrow \rho \cdot 2^{i+1}$ .
14:   Compute an  $k_i$ -partial maximum spanning forest packing
      $T_1, T_2, \dots, T_{k_i}$  for the graph  $G_i := (V, X_i)$ .
15:   Set  $F_i \leftarrow \bigcup_{j=1}^{k_i} T_j$ 
16:   Set  $Y_i \leftarrow X_i \setminus F_i$ .
17: end while
18: Set  $\Gamma \leftarrow i$ . //  $\Gamma$  is the number of elapsed iteration in the previous while-
   loop.
19: Add each edge  $e \in Y_\Gamma$  to  $G_\epsilon$  with weight  $2^{\Gamma-1}w(e)$ .
20: Add each edge  $e \in F_0$  to  $G_\epsilon$  with weight  $w(e)$ .
21: for  $j = 1, \dots, \Gamma$  do
22:   for  $e \in F_j$  do
23:     Set  $p_e \leftarrow \min\left(1, \frac{384}{169} \frac{1}{4^j w(e)}\right)$ .
24:     Generate  $r_e$  from  $\text{Binom}(2^j w(e), p_e)$ .
25:     if  $r_e > 0$  then
26:       Add  $e$  to  $G_\epsilon$  with weight  $r_e/p_e$ .
27:     end if
28:   end for
29: end for
30: return  $G_\epsilon = (V, E_\epsilon)$ .

```

The factor 2^i in calling upon the binomial distribution can be seen as boosting the weight of the edge by a factor 2^i , which is needed to counterbalance the i sampling steps in creating F_i .

Up to the computation method of the MSF packing, the presented algorithm is the same for polynomially and superpolynomially-weighted graphs. For the unbounded case, we use the MSF index estimator as presented in Sect. 3.1. In Sect. 5 we detail how

this influences the correctness of the algorithm, and the bounds on size and running time.

4.2 Correctness

We will prove that G_ϵ constructed in $\text{SPARSIFY}(V, E, w, \epsilon, c)$ is a $(1 \pm \epsilon)$ -cut sparsifier for G with probability at least $1 - 8/n^c$. Following the proof structure of [8], we first define

$$S := \left(\bigcup_{i=0}^{\Gamma} 2^i F_i \right) \cup 2^\Gamma Y_\Gamma,$$

where Γ is the maximum number such that $F_i \neq \emptyset$. We define $G_S := (V, S)$. And we prove the following two lemmas, that together yield the desired result.

Lemma 4.3 G_S is a $(1 \pm \epsilon/3)$ -cut sparsifier for G with probability at least $1 - 4/n^c$.

Lemma 4.4 G_ϵ is a $(1 \pm \epsilon/3)$ -cut sparsifier for G_S with probability at least $1 - 4/n^c$.

Let us start by proving Lemma 4.3. In creating the sets F_i , we repeatedly makes use of the MSF indices. The MSF index of an edge immediately ensures a certain connectivity of that edge. The following lemma makes this precise.

Lemma 4.5 Let $i \geq 0$ and $e \in Y_i$ be an edge, and set $k_i := \rho \cdot 2^{i+1}$. Then e is $w(e)k_i$ -heavy in $G'_{i,e} = (V, X'_{i,e})$, where $X'_{i,e} := \{e' \in X_i : w(e') \geq w(e)\}$. Consequently, e is also $w(e)k_i$ -heavy in $G_i = (V, X_i)$.

Proof Since $e \in Y_i = X_i \setminus F_i$, we know that e was not part of any maximum spanning forest in a k_i -partial MSF packing \mathcal{F}_i of G_i . Hence, by definition of the maximum spanning forests, each of the forests in \mathcal{F}_i has a path connecting the vertices of e , with all edges of weight at least $w(e)$. Thus any cut in G'_i picks up a contribution of at least $w(e)$ for each of the k_i paths. Hence the minimum cut in G'_i separating the vertices of e has value at least $w(e)k_i$, or equivalently e is $w(e)k_i$ -heavy in G'_i .

Next, we show in a general setting that certain ways of sampling preserve cuts. The following lemma is a generalization of Lemma 5.5 in [8].

Lemma 4.6 Let $R \subseteq Q$ be subsets of weighted edges on some set of vertices V , satisfying $0 < w(e) \leq 1$ for all $e \in Q$. Moreover, assume that each edge in R is π -heavy in (V, Q) . Suppose that each edge $e \in R$ is sampled with probability $p \in (0, 1]$, and if selected, given a weight of $w(e)/p$ to form a set of edges \hat{R} . We denote, for every cut C :

$$r^{(C)} := \sum_{e \in R \cap C} w(e), \quad q^{(C)} := \sum_{e \in Q \cap C} w(e), \quad \hat{r}^{(C)} := \sum_{e \in \hat{R} \cap C} w(e)/p.$$

Let $\zeta \in \mathbb{N}_{\geq 5}$, and $\delta \in (0, 1]$ such that $\delta^2 p\pi \geq \frac{\zeta \ln(n)}{0.38}$, then

$$\left| r^{(C)} - \widehat{r}^{(C)} \right| \leq \delta q^{(C)}$$

for all cuts C , with probability at least $1 - 4/n^{\zeta-4}$.

Proof For each $j \geq 0$, let \mathcal{C}_j be the set consisting of all cuts C with

$$2^j \cdot \pi \leq r^{(C)} \leq 2^{j+1} \cdot \pi - 1.$$

We will show that each j the statement of the lemma holds true with probability at least $1 - 2n^{(4-\zeta)2^j}$. Then the lemma follows from the union bound since

$$\sum_{j=0}^{\infty} 2n^{(4-\zeta)2^j} \leq 2n^{4-\zeta} \sum_{j=0}^{\infty} 2^{-(2^j-1)} \leq 2n^{4-\zeta} \sum_{k=0}^{\infty} 2^{-k} \leq 4n^{4-\zeta},$$

where we use that $n^{4-\zeta} \leq 1/2$.

Let $C \in \mathcal{C}_j$. For every $e \in R$, define the random variables Y_e that takes value $w(e)$ with probability p and 0 otherwise. We have $Y_e \in [0, 1]$, $\mathbb{E}[Y_e] = pw(e)$, and $\sum_{e \in R} Y_e = pr^{(C)}$. Now we apply Theorem A.1 with $\epsilon = \delta q^{(C)}/r^{(C)}$ and $\mu = pr^{(C)}$ to obtain

$$\begin{aligned} \mathbb{P} \left[\left| r^{(C)} - \widehat{r}^{(C)} \right| > \delta q^{(C)} \right] &= \mathbb{P} \left[\left| \sum_{e \in R} Y_e - \mu \right| > \delta \frac{q^{(C)}}{r^{(C)}} \cdot pr^{(C)} \right] \\ &\leq 2 \exp \left(-0.38\delta^2 \left(\frac{q^{(C)}}{r^{(C)}} \right)^2 pr^{(C)} \right) \\ &\leq 2 \exp \left(-0.38\delta^2 pq^{(C)} \right), \end{aligned}$$

where the last inequality holds as $r^{(C)} \leq q^{(C)}$ since $R \subseteq Q$. Now observe that $q^{(C)} \geq r^{(C)} \geq \pi \cdot 2^j$, hence

$$\begin{aligned} \mathbb{P} \left[\left| r^{(C)} - \widehat{r}^{(C)} \right| > \delta q^{(C)} \right] &\leq 2 \exp \left(-0.38\delta^2 p\pi 2^j \right) \\ &\leq 2 \exp \left(-\zeta \ln(n) 2^j \right) \\ &= n^{-\zeta 2^j}. \end{aligned}$$

As every edge in $R \cap C$ is π -heavy in (V, Q) , we can apply Theorem 2.2 to see that the number of distinct sets $R \cap C$ is at most:

$$n^{2 \frac{2^{j+1}\pi}{\pi}} = n^{4 \cdot 2^j}.$$

Thus the union bound gives us that the statement of the lemma holds true for all cuts $C \in \mathcal{C}_j$ with probability at least $1 - 2n^{(4-\zeta)2^j}$.

We want to apply this lemma to our sampling procedure. We do this by considering different weight classes separately. We define $X_{i,k} := \{e \in X_i : 2^k \leq w(e) \leq 2^{k+1} - 1\}$, and $x_{i,k}^{(C)} = \sum_{e \in X_{i,k} \cap C} w(e)$. We define $Y_{i,k}$ and $y_{i,k}^{(C)}$ analogously. Some rescaling is necessary to ensure that all weights lie in $(0, 1]$, as Lemma 4.6 requires. For $A \subseteq E$ and $\beta > 0$, we write βA to indicate we multiply the weight of the edges by a factor of β .

Corollary 4.7 *With probability at least $1 - 4/n^{4+c}$, for every cut C in G_i ,*

$$\left| 2^{-k} x_{i+1,k}^{(C)} - 2^{-k-1} y_{i,k}^{(C)} \right| \leq \frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=k}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)}.$$

Proof Any $e \in Y_{i,k}$ is $\rho \cdot 2^{i+1} w(e) \geq \rho \cdot 2^{i+k+1}$ -heavy in $\bigcup_{k'=k}^{\infty} X_{i,k'}$. A closer look shows us that we also have that any $e \in 2^{-k-1} Y_{i,k}$ is $\rho \cdot 2^i$ -heavy in $\bigcup_{k'=k}^{\infty} 2^{-k'-1} X_{i,k'}$. We set $R = 2^{-k-1} Y_{i,k}$, $Q = \bigcup_{k'=k}^{\infty} 2^{-k'-1} X_{i,k'}$, $\pi = \rho \cdot 2^i$, $p = 1/2$, and $\delta = \frac{\epsilon/13}{2^{i/2+1}}$, and we check that

$$\delta^2 p \pi = \frac{\epsilon^2/13^2}{2^{i+3}} \rho 2^i = \frac{\epsilon^2}{2^3 \cdot 13^2} \frac{(7+c)1352 \ln(n)}{0.38\epsilon^2} = \frac{(7+c) \ln(n)}{0.38}.$$

So we can apply Lemma 4.6 with these settings to obtain:

$$\left| 2^{-k} x_{i+1,k}^{(C)} - 2^{-k-1} y_{i,k}^{(C)} \right| \leq \frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=k}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)},$$

which holds for all cuts C with probability $1 - 4/n^{3+c}$.

Now we look at the general case, for which we sum all weight classes. Hereto, we define $x_i^{(C)} = \sum_{e \in X_i \cap C} w(e)$, $x_{i+1}^{(C)} = \sum_{e \in X_{i+1} \cap C} w(e)$, and $y_i^{(C)} = \sum_{e \in Y_i \cap C} w(e)$.


Corollary 4.8 *With probability at least $1 - 4/n^{1+c}$, for every cut C in G_i ,*

$$\left| 2x_{i+1}^{(C)} - y_i^{(C)} \right| \leq \frac{\epsilon/13}{2^{i/2}} \cdot x_i^{(C)}.$$

Proof We rescale and sum over k for each of the weight classes in Lemma 4.7 to get

$$\begin{aligned} \left| 2x_{i+1}^{(C)} - y_i^{(C)} \right| &= \left| \sum_{k=0}^{\infty} 2^{k+1} \left(2^{-k} x_{i+1,k}^{(C)} - 2^{-k-1} y_{i,k}^{(C)} \right) \right| \\ &\leq \sum_{k=0}^{\infty} 2^{k+1} \left| 2^{-k} x_{i+1,k}^{(C)} - 2^{-k-1} y_{i,k}^{(C)} \right| \end{aligned}$$

Fig. 2 A visualization of the area covered by $\sum_{k=0}^{\infty} \sum_{k'=k}^{\infty} 1 = \sum_{k'=0}^{\infty} \sum_{k=0}^{k'} 1$



$$\leq \sum_{k=0}^{\infty} 2^{k+1} \left(\frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=k}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)} \right)$$

Next, we want to interchange the sum over k with the sum over k' , a visual argument for the adjustment of the bounds can be found in Fig. 2.

$$\begin{aligned} \sum_{k=0}^{\infty} 2^{k+1} \left(\frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=k}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)} \right) &= \frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=0}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)} \sum_{k=0}^{k'} 2^{k+1} \\ &\leq \frac{\epsilon/13}{2^{i/2+1}} \sum_{k'=0}^{\infty} 2^{-k'-1} x_{i,k'}^{(C)} 2^{k'+2} \\ &\leq \frac{\epsilon/13}{2^{i/2}} x_i^{(C)}, \end{aligned}$$

which holds simultaneously for all cuts C with probability at least $1 - 4/n^{1+c}$. The reason is that at most $m \leq n^2$ of the $X_{i,k} \cap C$ are non-empty, hence a union bound gives the desired bound on the probability.

We will repeatedly apply this lemma. To show that the accumulated error does not grow beyond $\epsilon/3$, we use the following fact. For a proof we refer to [8].

Lemma 4.9 *Let $x \in (0, 1]$ be a parameter. Then for any $k \geq 0$,*

$$\begin{aligned} \prod_{i=0}^k \left(1 + \frac{x/13}{2^{i/2}} \right) &\leq 1 + x/3, \\ \prod_{i=0}^k \left(1 - \frac{x/13}{2^{i/2}} \right) &\geq 1 - x/3. \end{aligned}$$

As a final step towards proving Lemma 4.3, we prove a lemma that focusses on the sparsification occurring in the last $\Gamma - j + 1$ iterative steps of our algorithm.

Lemma 4.10 *Let*

$$S_j = \left(\bigcup_{i=j}^{\Gamma} 2^{i-j} F_i \right) \cup 2^{\Gamma-j} Y_{\Gamma}$$

for any $j \geq 0$. Then, S_j is a $(1 \pm (\epsilon/3)2^{-j/2})$ -cut sparsifier for $G_j = (V, X_j)$, with probability at least $1 - 4/n^c$.

Note that setting $j = 0$ gives us Lemma 4.3. Although this lemma is a generalization of the corresponding case for unweighted graphs in [8], the proof for the weighted case will be exactly the same: all the work that needed to be done is contained in the previous lemmas. We include the proof here for completeness.

Proof Let C be a cut. We define $s_j^{(C)} := \sum_{e \in S_j \cap C} w_{S_j}(e)$, and $f_i^{(C)}$, $x_i^{(C)}$, and $y_i^{(C)}$ analogously. We will show that the weight of C in S_j is at most $(1 + (\epsilon/3)2^{-j/2})$ times the weight of C in G_j .

$$\begin{aligned} s_j^{(C)} &= \sum_{i=j}^{\Gamma} 2^{i-j} f_i^{(C)} + 2^{\Gamma-j} y_{\Gamma}^{(C)} && \text{by definition of } S_j \\ &= \sum_{i=j}^{\Gamma-1} 2^{i-j} f_i^{(C)} + 2^{\Gamma-j} x_{\Gamma}^{(C)} && \text{by definition of } Y_{\Gamma} \\ &\leq \sum_{i=j}^{\Gamma-1} 2^{i-j} f_i^{(C)} + 2^{\Gamma-j-1} \left(y_{\Gamma-1}^{(C)} + \frac{\epsilon/13}{2^{(\Gamma-1)/2}} x_{\Gamma-1}^{(C)} \right) && \text{by Lemma 4.7} \\ &= \sum_{i=j}^{\Gamma-2} 2^{i-j} f_i^{(C)} + 2^{\Gamma-j-1} x_{\Gamma-1}^{(C)} \left(1 + \frac{\epsilon/13}{2^{(\Gamma-1)/2}} \right) && \text{by definition of } Y_{\Gamma-1} \end{aligned}$$

We repeat the last step $\Gamma - j - 1$ times to conclude

$$\begin{aligned} s_j^{(C)} &\leq x_j^{(C)} \prod_{i=j}^{\Gamma-1} \left(1 + \frac{\epsilon 2^{-j/2}/13}{2^{i/2}} \right) \\ &\leq x_j^{(C)} (1 + (\epsilon/3)2^{-j/2}) && \text{by Lemma 4.8} \end{aligned}$$

The proof of $s_j^{(C)} \geq x_j^{(C)} (1 - (\epsilon/3)2^{-j/2})$ is analogous. As we have that $\Gamma \leq n$, we can use a union bound to conclude that Lemma 4.8 holds for all simultaneously with probability at least $1 - 4/n^c$, which concludes the proof.

To prove Lemma 4.4, we will invoke the framework from [8], as given in Sect. 2.1. More specifically, we will apply Theorem 2.3. We set the parameter $\gamma := 64/3$, and for each $e \in F_i$ we set $\lambda_e := \rho \cdot 4^i w(e)$. This is in line with our choice for p_e :

$$\begin{aligned} \min \left(1, \frac{16(c + 7)\gamma \ln(n)}{0.38\lambda_e \epsilon^2} \right) &= \min \left(1, \frac{16(c + 7)\gamma \ln(n)}{0.38\rho \cdot 4^i w(e) \epsilon^2} \right) \\ &= \min \left(1, \frac{384}{169} \frac{1}{4^i w(e)} \right) = p_e. \end{aligned}$$

We have to provide a set of subgraphs \mathcal{G} and a set of parameters Π such that Π -connectivity and γ -overlap are satisfied.

To explore the connectivity of edges in $R_i := \{e \in E : 2^i \leq \lambda_e \leq 2^{i+1} - 1\}$ we partition these sets as follows:

$$R_{j,k} := \{e \in F_j : 2^k \leq \rho w(e) \leq 2^{k+1} - 1\}.$$

We will view these edges in the subgraph:

$$E_{j,k} := \bigcup_{j'=j-1}^{\Gamma} \bigcup_{k'=k}^{\infty} \rho \cdot 4^{\Gamma-j'+1} 2^{\Lambda-k'+j'} R_{j',k'}.$$

Lemma 4.11 *Each edge $e \in R_{j,k}$ is $\pi := \rho \cdot 4^{\Gamma} 2^{\Lambda}$ -heavy in $(V, E_{j,k})$.*

Proof Fix $e \in R_{j,k}$. This edge is $\rho \cdot 2^j w(e) \geq \rho \cdot 2^{j+k}$ -heavy in $\{e \in X_{j-1} : w(e') \geq w(e)\}$, see Lemma 4.5. Hence e is $\rho \cdot 2^{j+k}$ -heavy in $\{e' \in X_{j-1} : \rho w(e') \geq 2^k\}$. We can rescale this: e is $(\rho \cdot 2^{j+\Lambda})$ -heavy in $2^{\Lambda-k} \cdot \{e' \in X_{j-1} : \rho w(e') \geq 2^k\} = \bigcup_{k'=k}^{\infty} 2^{\Lambda-k} \cdot \{e' \in X_{j-1} : 2^{k'} \leq \rho w(e') \leq 2^{k'+1} - 1\}$. We rescale again to see e is $\rho \cdot 2^{2j+\Lambda-1}$ -heavy in $\bigcup_{k'=k}^{\infty} 2^{\Lambda-k+j-1} \cdot \{e' \in X_{j-1} : 2^{k'} \leq \rho w(e') \leq 2^{k'+1} - 1\}$. Next, we want to replace X_{j-1} with S_{j-1} . Hereto, we apply Lemma 4.7 with $\epsilon = 13 \cdot 2^{i/2+1}$, which shows that for each of the weight classes the cuts are preserved up to a factor 2. Hence we obtain e is $\rho \cdot 2^{2j+\Lambda-2}$ -heavy in $\tilde{E}_{j,k} := \bigcup_{j'=j-1}^{\Gamma} \bigcup_{k'=k}^{\infty} 2^{\Lambda-k+j'} \cdot \{e' \in F_{j'} : 2^{k'} \leq \rho w(e') \leq 2^{k'+1} - 1\}$.

Now let $e' \in R_{j,k}$ be any edge, and let C be a cut such that $e' \in C$. We need to show that the weight of this cut in $E_{j,k}$ is at least $\rho \cdot 4^{\Gamma} 2^{\Lambda}$. Let $e := \operatorname{argmin}_{e \in C} \{j_e : e \in R_{j_e, k_e} \text{ for some } k_e \geq k\}$ (in case e is not unique, pick any). By the above statement we have that e is $\rho \cdot 2^{2j_e+\Lambda-2}$ -heavy in $\tilde{E}_{j_e, k_e} \subseteq \tilde{E}_{j_e, k}$. Thus e is $\rho \cdot 4^{\Gamma} 2^{\Lambda}$ -heavy in $4^{\Gamma-j_e+1} \tilde{E}_{j_e, k}$. This is a subgraph of $E_{j_e, k}$, which in turn is a subgraph of $E_{j,k}$. Hence e is $\rho \cdot 4^{\Gamma} 2^{\Lambda}$ -heavy in $E_{j,k}$, and thus C has weight at least $\rho \cdot 4^{\Gamma} 2^{\Lambda}$.

Now we take all weight classes together to find the set of subgraphs \mathcal{G} for which Π -connectivity is satisfied.

Corollary 4.12 *Each edge in $e \in R_i$ is $\rho \cdot 4^{\Gamma} 2^{\Lambda}$ -heavy in $G_i = (V, E_i)$, with $E_i := \bigcup_{j=1}^{\min(i/2, \Gamma)} E_{j, i-2j}$.*

Proof Note that $e \in R_i$ satisfies $2^i \leq \rho \cdot 2^{2j} w(e) \leq 2^{i+1} - 1$ if $e \in F_j$. Hence $e \in R_{j,k}$ with $2j + k = i$. We are only considering edges in F_j with $1 \leq j \leq \Gamma$, thus we have $R_i = \bigcup_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} R_{j,i-2j}$, hence the claim follows directly from Lemma 4.11.

It remains to show that γ -overlap is satisfied.

Lemma 4.13 *For any cut C ,*

$$\sum_{i=0}^{\Lambda} \frac{e_i^{(C)} 2^{i-1}}{\rho \cdot 4^\Gamma 2^\lambda} \leq 64/3 \cdot e^{(C)},$$

where $e^{(C)} = \sum_{e \in C} w_{G_S}(e)$ and $e_i^{(C)} = \sum_{e \in C \cap E_i} w_{G_i}(e)$.

Proof We add F_0 and Y_Γ to G_ϵ , so we do not need to be concerned about the intersection of the cut C with these sets. This means we only intersect a cut C with F_j where $1 \leq j \leq \Gamma$. Hence we start our sum with $i = 2$. We consider the sum we need to bound:

$$\begin{aligned} \sum_{i=2}^{\Lambda} \frac{e_i^{(C)} 2^{i-1}}{\rho \cdot 4^\Gamma 2^\lambda} &= \sum_{i=2}^{\Lambda} \frac{\left(\sum_{e \in C \cap E_i} w_{G_i}(e)\right) 2^{i-1}}{\rho \cdot 4^\Gamma 2^\lambda} \\ &= \sum_{i=2}^{\Lambda} \sum_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} \frac{\left(\sum_{e \in C \cap E_{j,i-2j}} w_{G_i}(e)\right) 2^{i-1}}{\rho \cdot 4^\Gamma 2^\lambda} \\ &= \sum_{i=2}^{\Lambda} \sum_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} \sum_{j'=j-1}^{\Gamma} \sum_{k'=i-2j}^{\infty} \frac{\rho \cdot 4^{\Gamma-j'+1} 2^{\Lambda-k'+j'} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e)\right) 2^{i-1}}{\rho \cdot 4^\Gamma 2^\lambda} \\ &= \sum_{i=2}^{\Lambda} \sum_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} \sum_{j'=j-1}^{\Gamma} \sum_{k'=i-2j}^{\infty} 2^{-k'-j'+i+1} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right). \end{aligned}$$

Next, we want to interchange the sum over i and the sum over j and change the bounds accordingly. See Fig. 3a for a visual argument.

$$\begin{aligned} &\sum_{i=2}^{\Lambda} \sum_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} \sum_{j'=j-1}^{\Gamma} \sum_{k'=i-2j}^{\infty} 2^{-k'-j'+i+1} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \\ &= \sum_{j=1}^{\Gamma} \sum_{i=2j}^{\Lambda} \sum_{j'=j-1}^{\Gamma} \sum_{k'=i-2j}^{\infty} 2^{-k'-j'+i+1} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \end{aligned}$$

Interchanging the sum over i and j' does not change the bounds, as they are independent of each other. When interchanging the sum over i and the sum over k' we

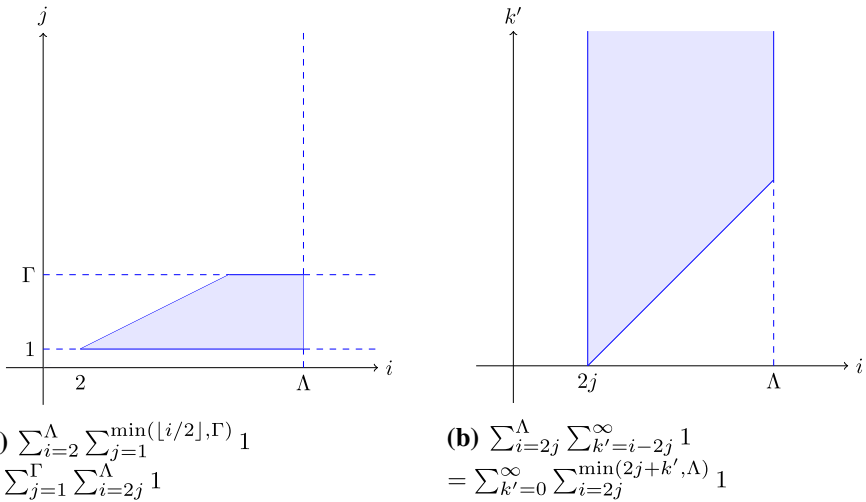


Fig. 3 Two visualizations of the area covered by a double sum


have to be more careful, see Fig. 3b for a visual argument.

$$\begin{aligned}
 & \sum_{j=1}^{\Gamma} \sum_{i=2j}^{\Lambda} \sum_{j'=j-1}^{\Gamma} \sum_{k'=i-2j}^{\infty} 2^{-k'-j'+i+1} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \\
 &= \sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} \sum_{k'=0}^{\infty} \sum_{i=2j}^{\min(2j+k', \Lambda)} 2^{-k'-j'+i+1} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \\
 &\leq \sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} \sum_{k'=0}^{\infty} 2^{-k'-j'+2j+k'+2} \left(\sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \\
 &= \sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} 2^{2j-j'+2} \left(\sum_{k'=0}^{\infty} \sum_{e \in C \cap E_{j',k'}} w_G(e) \right) \\
 &= \sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} 2^{2j-j'+2} \left(\sum_{e \in C \cap F_{j'}} w_G(e) \right).
 \end{aligned}$$

Next, we want to interchange the sum over j with the sum over j' , a visual argument can be found in Fig. 4.

$$\sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} 2^{2j-j'+2} \left(\sum_{e \in C \cap F_{j'}} w_G(e) \right) = \sum_{j'=0}^{\Gamma} 2^{-j'+2} \sum_{j=1}^{j'+1} 4^j \left(\sum_{e \in C \cap F_{j'}} w_G(e) \right)$$

Fig. 4 A visualization of the area covered by $\sum_{j=1}^{\Gamma} \sum_{j'=j-1}^{\Gamma} 1 = \sum_{j'=0}^{\Gamma} \sum_{j=1}^{j'+1} 1$



$$\begin{aligned}
 &\leq \sum_{j'=0}^{\Gamma} 2^{-j'+2} \frac{4^{j'+2}}{3} \left(\sum_{e \in C \cap F_{j'}} w_G(e) \right) \\
 &= \frac{64}{3} \sum_{j'=0}^{\Gamma} 2^{j'} \left(\sum_{e \in C \cap F_{j'}} w_G(e) \right) \\
 &= \frac{64}{3} \sum_{e \in C} w_{G_S}(e) \\
 &= \frac{64}{3} e^{(C)}.
 \end{aligned}$$

Together Corollary 4.12 and Lemma 4.13 show that the conditions of Theorem 2.3 are met with the given parameters. This proves Lemma 4.4, and then Theorem 4.1 follows.

4.3 Size of the Sparsifier

The sparsifier G_ϵ consists of F_0 , Y_Γ , and F' , where $F' = \cup_{i=1}^{\Gamma} F'_i$, with F'_i the sampled edges of F_i . First of all, note that $|F_0| = O(cn \ln(n)/\epsilon^2)$ and $|Y_\Gamma| = O(cn \ln(n)/\epsilon^2)$. Now take $e \in F_i$. This edge results to an edge in G_ϵ if the sample from the binomial distribution with parameters $n_e = 2^i w(e)$ and $p_e = \min\left(1, \frac{384}{169} \frac{1}{4^i w(e)}\right)$ is positive. The probability that this happens is

$$\mathbb{P}[\text{Binom}(n_e, p_e) > 0] = \sum_{k=1}^{n_e} \mathbb{P}[\text{Binom}(n_e, p_e) = k]$$

$$\begin{aligned}
 &\leq \sum_{k=1}^{n_e} k \mathbb{P}[\text{Binom}(n_e, p_e) = k] \\
 &= \sum_{k=0}^{n_e} k \mathbb{P}[\text{Binom}(n_e, p_e) = k] \\
 &= \mathbb{E}[\text{Binom}(n_e, p_e)] \\
 &= n_e p_e \\
 &\leq \frac{384}{169} 2^{-i}.
 \end{aligned}$$

Note that this probability is equal for all $e \in F_i$. Since F_i is the union of $k_i = \rho \cdot 2^{i+1}$ spanning forests, we know that $|F_i| \leq \rho 2^{i+1} n$. Hence the expected size of F'_i , the sampled edges in F_i , equals

$$\begin{aligned}
 \mathbb{E}[|F'_i|] &= \sum_{e \in F_i} \mathbb{P}[\text{Binom}(n_e, p_e) > 0] \\
 &\leq \sum_{e \in F_i} \frac{384}{169} 2^{-i} \\
 &= |F_i| \frac{384}{169} 2^{-i} \\
 &\leq \rho 2^{i+1} n \frac{384}{169} 2^{-i} \\
 &= \rho \frac{768}{169} n.
 \end{aligned}$$

We have that the total number of sampled edges equals

$$\mathbb{E}[|F'|] = \sum_{i=1}^{\Gamma} \mathbb{E}[|F'_i|] \leq \Gamma \rho \frac{768}{169} n,$$

so it remains to bound Γ , i.e., the number of F_i 's. Hereto, note that the while loop of lines 10–17 ends if $|Y_i| \leq 2\rho n$. We bound the number of edges in Y_i by bounding the number of edges of X_i , of which Y_i is a subset. Each edge in $Y_{i-1} \subseteq X_{i-1}$ is sampled with probability $1/2$ to form X_i . So $\mathbb{E}[|X_i|] \leq |X_{i-1}|/2$. Now by a Chernoff bound (see Theorem A.2) we obtain:

$$\mathbb{P}\left[|X_i| > \frac{2}{3} |X_{i-1}|\right] \leq \exp\left(-\frac{0.38}{36} |X_{i-1}|\right) > \exp\left(-\frac{cn \ln(n)}{36}\right) = n^{-cn/36},$$

since $|X_{i-1}| \geq |Y_{i-1}| \geq 2\rho n = 2 \cdot \frac{(7+c)1352 \ln(n)}{0.38\epsilon^2} n \geq \frac{cn \ln(n)}{0.38}$. We have at most n^2 sets X_i , so we can conclude that with high probability $|X_i| \leq \frac{2}{3} |X_{i-1}|$ in each step, and

by induction $|Y_i| < |X_i| \leq \left(\frac{2}{3}\right)^i m$. We see that

$$m \left(\frac{2}{3}\right)^\Gamma \leq 2\rho n = \frac{21632}{0.38\epsilon^2} cn \ln(n),$$

which is equivalent to

$$\left(\frac{2}{3}\right)^\Gamma \leq \frac{\frac{21632}{0.38\epsilon^2} cn \ln(n)}{m},$$

and that is equivalent to

$$\Gamma \geq \log \left(\frac{m}{\frac{21632}{0.38\epsilon^2} cn \ln(n)} \right) / \log(3/2).$$

So, we can conclude $\Gamma = O\left(\log\left(\frac{m}{cn \log(n)/\epsilon^2}\right)\right)$. This gives that the total number of sampled edges is, in expectation,

$$\mathbb{E}[|F'|] \leq \Gamma \rho \frac{768}{169} n = O(cn \log(n) \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2).$$

This compression process can also be seen as the sum of m independent random variables that take values in $\{1, 0\}$.³ We have just calculated that the expected value μ is at most $Bcn \ln(n) \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2$, for some $B > 0$. Using this, we apply a Chernoff bound (Theorem A.2) to get an upper limit for the number of sampled edges:

$$\begin{aligned} & \mathbb{P}\left[|F'| > 2Bcn \ln n \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2\right] \\ & \leq \exp\left(-0.38Bcn \ln(n) \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2\right) \\ & = n^{-0.38cnB \log(m/(cn \log(n)/\epsilon^2))/\epsilon^2}. \end{aligned}$$

We conclude that, with high probability, the number of sampled edges is

$$O(cn \log(n) \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2).$$

And finally, we conclude that with high probability the number of edges of G_ϵ is bounded by $|E(G_\epsilon)| = |F_0| + |Y_\Gamma| + |F'| = O(cn \log(n) \log\left(m/(cn \log(n)/\epsilon^2)\right) / \epsilon^2)$.

³ To be precise, we set the probability of an edge $e \notin \cup_i F_i$ to exist to 0.

4.4 Time Complexity

First off, if

$$\begin{aligned} m &\leq 4\rho n \log \left(m / (n \log(n) / \epsilon^2) \right) \\ &= O(cn \log(n) / \epsilon^2 \log \left(m / (n \log(n) / \epsilon^2) \right)), \end{aligned}$$

the algorithm does nothing and returns the original graph. So for this analysis we can assume $m > 4\rho n \log \left(m / (n \log(n) / \epsilon^2) \right)$. We analyze the time complexity of the algorithm in two phases. The first phase consists of computing the probabilities p_e for all $e \in E$. The second one is compressing edges, given these probabilities.

The first phase contains i iterations of the while loop (lines 10–17). In each iteration we sample edges from $Y_i \subseteq X_i$ with probability $1/2$ to form X_{i+1} . This takes time at most $O(|X_i|)$. Next, we compute a maximum spanning forest packing of the graph $G_{i+1} = (V, X_{i+1})$. We know that we can compute a M -partial maximum spanning forest packing of a polynomially-weighted graph with n vertices and m_0 edges in $O(m_0 \cdot \min(\alpha(n) \log(M), \log(n)))$ time (see Theorem 3.2 and Theorem 3.3). So this iteration takes at most $O(|X_{i+1}| \cdot (\min(\alpha(n) \log(k_{i+1}), \log(n))))$ time. As noted earlier, we have with high probability that $|X_i| \leq \left(\frac{2}{3}\right)^i m$. If $m\alpha(n) \log(m/n) \leq m \log(n)$, we conclude w.h.p. that the first phase takes total time at most

$$\begin{aligned} &\sum_{i=0}^{\Gamma} O(|X_i|) + O(|X_{i+1}| \alpha(n) \log(k_{i+1})) \\ &= \sum_{i=0}^{\Gamma} \left(\frac{2}{3}\right)^i O(m) + \left(\frac{2}{3}\right)^{i+1} O(m\alpha(n) \log(\rho 2^{i+2})) \\ &\leq 3 \cdot O(m) + 3 \cdot O(m\alpha(n) \log(\rho 2^{\Gamma})) \\ &= O(m\alpha(n) \log(m/n)). \end{aligned}$$

And if $m \log(n) < m\alpha(n) \log(m/n)$, we have that w.h.p. the first phase takes total time at most

$$\begin{aligned} \sum_{i=0}^{\Gamma} O(|X_i|) + O(|X_{i+1}| \log(n)) &= \sum_{i=0}^{\Gamma} \left(\frac{2}{3}\right)^i O(m) + \left(\frac{2}{3}\right)^{i+1} O(m \log(n)) \\ &\leq 3 \cdot O(m) + 3 \cdot O(m \log(n)) \\ &= O(m \log n). \end{aligned}$$

In the second phase, we sample each edge e from the binomial distribution with parameters n_e and p_e . We will show this can be done with a process that takes $T = O(m)$ time with high probability. Hereto, we use an algorithm from [29] for binomial sampling, for which the pseudocode is given in Algorithm 2.

Algorithm 2 BINOM(n, p)

Input: Two parameters n, p .

Output: A random sample from the binomial distribution with parameters n and p .

- 1: Set $k \leftarrow -1, S \leftarrow 0$.
 - 2: **while** $S < n$ **do**
 - 3: $k \leftarrow k + 1$.
 - 4: Generate $u \sim \mathcal{U}(0, 1)$.
 - 5: $S \leftarrow S + \lceil \log(u) / \log(1 - p) \rceil + 1$.
 - 6: **end while**
 - 7: **return** k
-

It is easy to see that this algorithm takes $O(1 + k)$ time, where k is the output. So if the sample from the binomial distribution is k , this takes time $O(1 + k)$. This means that the total time T equals m plus the total sum of all samples. Note that this is slightly different from what we did in Sect. 4.3 to bound the number of edges: there we needed to bound the number of positive samples.

For each edge $e \in F_i$ we need to draw from the binomial distribution with parameters n_e and p_e . We denote T_e for the time we need to sample e . By the above, we have $\mathbb{E}[T_e] = 1 + n_e p_e$. So, the expected number of successes is at most

$$\begin{aligned} \mathbb{E}[T] &= \sum_i \sum_{e \in F_i} \mathbb{E}[T_e] = \sum_i \sum_{e \in F_i} (1 + n_e p_e) \\ &= \sum_i |F_i| + O(cn \log(n) \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2), \end{aligned}$$

as shown in Sect. 4.3. Let $B > 0$ such that $\sum_i \sum_{e \in F_i} n_e p_e \leq Bcn \ln(n) \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2$. We can use a Chernoff bound (see Theorem A.2) on the sum of these $\sum_i \sum_{e \in F_i} n_e$ random variables to obtain:

$$\begin{aligned} &\mathbb{P}\left[T - \sum_i |F_i| > 2Bcn \ln(n) \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2\right] \\ &\leq \exp\left(-0.38Bcn \ln(n) \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2\right) \\ &= n^{-0.38Bcn \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2}. \end{aligned}$$

So we can say that with high probability we need

$$\begin{aligned} T &= \sum_i |F_i| + \left(T - \sum_i |F_i|\right) \\ &= O(m) + O(2Bcn \ln(n) \log(m / (n \log(n) / \epsilon^2)) / \epsilon^2) = O(m) \end{aligned}$$

time for the sampling.

Concluding, the algorithm takes

$$O(m \cdot \min(\alpha(n) \log(m/n), \log(n)) + O(m) = O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$$

time in total for polynomially-weighted graphs.

5 Adaptation to Unbounded Weights

In this section, we sketch how we can adapt the algorithm of the previous section to sparse graphs with unbounded weights. The key to this is Lemma 3.4, which shows that for unbounded weights we might not be able to compute the MSF indices exactly, but we can find an estimate for edges e with $w(e) > d(e)/n$. Recall the definition of $d(e)$: compute a single maximum spanning forest F for G and define $d(e)$ to be the minimum weight among the edges on the path from u to v in F , where $e = (u, v)$.

The only adaptation for unbounded weights is that the first time we compute maximum spanning forests in Algorithm 1, we set aside any edges $e \in E$ with $w(e) \leq d(e)/n$. We show that we can sample efficiently from these vertices, since they are well-connected by F_0 , the initial MSF that remains in our sparsifier. We will do this by sampling them with $\lambda_e = \rho \cdot d(e)$. Note that we only have to set aside vertices the first time we compute a MSF packing, after this the estimates $d(e)$ in a new graph can only decrease, so if a vertex satisfies $w(e) \leq d(e)/n$ in a certain subgraph, it also satisfied this in the initial graph.

For the remaining vertices, we apply the algorithm as presented in the previous section. The only difference is that we use Lemma 3.4 to compute a estimates of the MSF indices. This means that if an edge $e \in E$ obtains the estimate index \hat{f}_e w.r.t. some graph E' , we have that e is at least $f_e w_e (1 - 1/n)$ -heavy in E' . For simplicity, we use $1 - 1/n \geq 1/2$. We see that this impacts the analysis in two places where the heaviness is used: Lemmas 4.7 and 4.11.

When examining Lemma 4.7, we see that we apply Lemma 4.6 with $\delta^2 p \pi \geq \frac{\zeta \ln(n)}{0.38}$, for certain δ, p, π , and ζ . We want to apply this lemma but have $\tilde{\pi} = \pi/2$, hence we set $\tilde{\delta} = \sqrt{2}\delta$. If we want to end up with the original result of Lemma 4.7, we set the $\tilde{\epsilon} = \epsilon/\sqrt{2}$. This constant factor change gets absorbed in the asymptotic notation for size and running time of the algorithm.

The second lemma we investigate is Lemma 4.11, which is the Π -connectivity in the sampling. Here, there is an easy solution: we boost all edges in $E_{j,k}$ by a factor two, which ensures the Π -connectivity as desired. Consequently, all edges in E_i are boosted with a factor two, which propagates to a factor two in $e_i(C)$ as denoted in Lemma 4.13, resulting to a γ -overlap with $\gamma = \frac{128}{3}$, rather than $\frac{64}{3}$.

Summing this up, we can say that our original analysis holds when we call the algorithm with $\tilde{\epsilon} = \epsilon/\sqrt{2}$ and $\tilde{\rho} = \frac{(7+c)2704 \ln(n)}{0.38\epsilon^2}$, where the change in ρ is a direct consequence of the change in γ .

The last thing that remains, is to show that, when we sample, Π -connectivity is also satisfied for the edges $e \in E$ with $w(e) \leq d(e)/n$. This is an extension to Corollary 4.12.

Lemma 5.1 *Suppose $e \in R_i$ and $w(e) \leq d(e)/n$, then e is $\pi = \rho \cdot 4^\Gamma 2^\Lambda$ -heavy in $G_i = (V, E_i)$, with $E_i = \bigcup_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} E_{j, i-2j}$.*

Proof We know that e is $d(e)$ -heavy in F_0 , so we look for the occurrence of F_0 in E_i :

$$\begin{aligned}
 E_i &= \bigcup_{j=1}^{\min(\lfloor i/2 \rfloor, \Gamma)} \bigcup_{j'=j-1}^{\Gamma} \bigcup_{k'=i-2j}^{\infty} \rho \cdot 4^{\Gamma-j'+1} 2^{\Lambda-k'+j'} \{e' \in F_{j'} : 2^{k'} \leq \rho \cdot w(e') \leq 2^{k'+1} - 1\} \\
 &\supseteq \rho \cdot 4^{\Gamma+1} \bigcup_{k'=i-2}^{\infty} 2^{\Lambda-k'} \{e' \in F_0 : 2^{k'} \leq \rho \cdot w(e') \leq 2^{k'+1} - 1\}. \tag{1}
 \end{aligned}$$

We look more closely at the connectedness of e in this particular set. We note that $w(e') \geq d(e)$ for any edge on a path in F_0 from u to v for $e = (u, v)$, by definition of $d(e)$. So we only need to consider $e' \in F_0$ with $\rho \cdot w(e') \geq \rho \cdot d(e) = \lambda_e \geq 2^i$, as $e \in R_i$. This means that e is $d(e)$ -heavy in

$$\begin{aligned}
 &\bigcup_{k'=i}^{\infty} \{e' \in F_0 : 2^{k'} \leq \rho \cdot w(e') \leq 2^{k'+1} - 1\} \\
 &\subseteq \bigcup_{k'=i-2}^{\infty} \{e' \in F_0 : 2^{k'} \leq \rho \cdot w(e') \leq 2^{k'+1} - 1\}.
 \end{aligned}$$

We can rescale this to exploit the weights fully: e is 2^Λ -heavy in $\bigcup_{k'=i-2}^{\infty} 2^{\Lambda-k'} \{e' \in F_0 : 2^{k'} \leq \rho \cdot w(e') \leq 2^{k'+1} - 1\}$. Combining this with Equation 1 gives us that e is $\rho \cdot 4^{\Gamma+1} 2^\Lambda$ -heavy in E_i , which is a factor four more than we needed to show.

5.1 Size and Time Complexity

For the size of the resulting graph G_ϵ , the upper bound of the previous section still holds for the edges that are sampled according to their MSF index. It remains to show that the contribution of any edges with $w(e) \leq d(e)/n$ is small. For these edges we have $p_e = \frac{384}{169} \frac{1}{d(e)}$. We use $\mathbb{P}[\text{Binom}(n_e, p_e) > 0] \leq p_e n_e$, to see

$$\mathbb{P}[e \in G_\epsilon : \lambda_e = \rho \cdot d(e)] \leq \frac{384}{169} \frac{w(e)}{d(e)} \leq \frac{384}{169} \frac{w(e)}{n \cdot w(e)}.$$

As there can be at most n^2 edges with $w(e) \leq d(e)/n$, we obtain that the expected number of edges in G_ϵ originating from such edges is at most $O(n)$. By the same arguments as given in Sect. 4.3, this holds not only in expectation, but also with high probability.

Concerning the time complexity, we use Theorem 3.3 or Lemma 3.4 instead of Theorem 3.2. These run in time $O(m \log(n))$ and $O(m\alpha(n) \log(M))$ respectively.

Since the size of the sparsifier does not increase significantly, the time needed for sampling does not increase significantly either. Hence we obtain a total time of $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$. This makes the algorithm the fastest cut sparsification algorithm known for graphs with unbounded weights.

6 Conclusion

In this paper, we presented a faster $(1 \pm \epsilon)$ -cut sparsification algorithm for weighted graphs. We have shown how to compute sparsifiers of size $O(n \log(n)/\epsilon^2)$ in $O(m \cdot \min(\alpha(n) \log(m/n), \log(n)))$ time, for integer weighted graphs. Both algorithms apply a sampling technique where the MSF index is used as a connectivity estimator.

We have shown that we can compute an M -partial MSF packing in $O(m\alpha(m) \log(M))$ time for polynomially-weighted graphs. For graphs with unbounded integer weights, we have shown that we can compute a complete MSF packing in $O(m \log(n))$ time, and a sufficient estimation of an M -partial MSF packing can be computed in time $O(m\alpha(m) \log(M))$. An open question is whether a more efficient computation is possible. This would improve on our sparsification algorithm, but might also be advantageous in other applications. The NI index has shown to be useful in various applications. We believe to have shown that the MSF index is a natural analogue.

To develop an algorithm to compute an MSF packing, one might be inclined to build upon one of the algorithms that compute a minimum spanning tree faster than Kruskal's algorithm, such as the celebrated linear-time algorithm of Karger, Klein, and Tarjan [30]. However, this algorithm and many other fast minimum spanning tree algorithms make use of edge contractions. It is far from obvious how to generalize this to a packing: in that case, we need to work simultaneously on multiple trees, hence we cannot simply contract the input graph in favor of any single one. To make this work, a more meticulous use of data structures seems necessary.

Computation of the MSF indices in linear time would be an ultimate goal. However, for our application a slightly looser bound suffices. If we can reduce the running time to compute the MSF indices to $O(m + n \log(n))$, then we obtain a time bound of $O(m)$ for cut sparsification. Moreover, we do not need the exact MSF index, an estimate suffices. This can either be a constant-factor approximation of the MSF index for each edge, or an estimate in the weights used in the forests, as done for graphs with unbounded weights in Sect. 5.

Funding Open access funding provided by Austrian Science Fund (FWF). This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant Agreement No. 947702) and is supported by the Austrian Science Fund (FWF): P 32863-N.

Declarations

Conflict of interest The authors have no relevant financial or non-financial interests to disclose.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence,

and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

A Tail Bounds

To analyze the sampling methods used in Sect. 4, we make use of the well-known Chernoff bound to get a grasp on the tail of various distributions [31].

Theorem A.1 *Let Y_1, \dots, Y_n be n independent random variables such that each Y_i takes values in $[0, 1]$. Let $\mu = \sum_{i=1}^n \mathbb{E}[Y_i]$ and $\xi = 2 \ln(2) > 0.38$. Then for all $\epsilon > 0$*

$$\mathbb{P} \left[\left| \sum_{i=1}^n Y_i - \mu \right| > \epsilon \mu \right] \leq 2 \exp \left(-\xi \min(\epsilon, \epsilon^2) \mu \right).$$

At times, the expected value μ itself is not known. Fortunately an upper bound on the expected value also suffices.

Theorem A.2 *Let Y_1, \dots, Y_n be n independent random variables such that Y_i takes values in $[0, 1]$. Let $\mu = \sum_{i=1}^n \mathbb{E}[Y_i]$ and $\xi = 2 \ln(2) > 0.38$. Suppose $\mu' \geq \mu$. Then for all $\delta \geq 2$*

$$\mathbb{P} \left[\sum_{i=1}^n Y_i > \delta \mu' \right] \leq 2 \exp \left(-\xi (\delta - 1) \mu' \right).$$

Proof Let $\epsilon := (\delta - 1) \frac{\mu'}{\mu}$. We have $\epsilon \geq 1$, so $\min(\epsilon, \epsilon^2) = \epsilon$. The statement now follows directly from Theorem A.1.

B Reduction from Real to Integer Weights

In this section, we show how to reduce the computation of a cut sparsifier of a graph with non-negative real weights to integer weights, formalizing the procedure sketched by Benzúr and Karger [7]. Let $G = (V, E, w)$ be a weighted graph, where $w : E \rightarrow \mathbb{R}$.

Denote $W_{\max} := \max_{e \in E} w(e)$ and $W_{\min} := \min \left\{ 1, \min_{e \in E} w(e) \right\}$. Then the reduction consists of the following steps:

1. Compute W_{\min} and $r := -\lceil \log(\frac{\epsilon}{2} W_{\min}) \rceil$.
2. Create $w' : E \rightarrow \mathbb{R}$ by rounding the weights $w(e)$ to the closest multiple of 2^{-r} , and define $G' := (V, E, w')$.
3. Create $\hat{w} : E \rightarrow \mathbb{R}$ by $\hat{w}(e) := 2^r w'(e)$.
4. Compute a $(1 \pm \epsilon/3)$ -cut sparsifier $\hat{H} = (V, E_H, \hat{w}_H)$ of $\hat{G} = (V, E, \hat{w})$.

5. Output $H = (V, E_H, w_H)$ where $w_H(e) := 2^{-r} \hat{w}_H(e)$.

First, we show that the graph H is indeed a $(1 + \epsilon)$ -cut sparsifier of G . Hereto, we note that for any cut C we have

$$w_H(C) = 2^{-r} w_{\hat{H}}(C) \leq 2^{-r} (1 + \epsilon/3) w_{\hat{G}}(C) = (1 + \epsilon/3) w_{G'}(C) \leq (1 + \epsilon) w_G(C),$$

where the last inequality holds as each weight $w'(e)$ has at most an additive error of $2^{-r} \leq \frac{\epsilon}{2} W_{\min} \leq \frac{\epsilon}{2}$ with respect to $w(e)$, hence at most a multiplicative error of $\frac{\epsilon}{2}$. Analogously we obtain $w_H(C) \geq (1 - \epsilon) w_G(C)$.

By construction, \hat{G} has integer weights, which are bounded by $O(\frac{W_{\max}}{\epsilon W_{\min}})$. Steps 1, 2, 3, and 5 can be implemented in $O(m)$ time. So indeed we have reduced the problem to finding a cut sparsifier of a graph with integer weights. Moreover, note that if G has polynomially bounded real weights, in the sense that $W_{\max} = O(\text{poly}(n))$ and $W_{\min} = \Omega(1/\text{poly}(n))$, then the graph \hat{G} has polynomially bounded integer weights. We can state this independent of ϵ , since for $\epsilon \leq 1/m$ we can always output the entire input graph as a cut sparsifier of optimal size $O(n/\epsilon^2)$ [19].

References

1. Peleg, D., Schäffer, A.A.: Graph spanners. *J. Graph Theory* **13**(1), 99–116 (1989)
2. Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. *Discrete Comput. Geom.* **9**(1), 81–100 (1993)
3. Baswana, S., Sen, S.: A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms* **30**(4), 532–563 (2007). <https://doi.org/10.1002/rsa.20130>
4. Elkin, M., Neiman, O.: Efficient algorithms for constructing very sparse spanners and emulators. *ACM Trans. Algorithms* (2016). <https://doi.org/10.1145/3274651>
5. Diniz, M., Krauthgamer, R., Wagner, T.: Towards resistance sparsifiers. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM)*. LIPIcs, vol. 40, pp. 738–755 (2015)
6. Benczúr, A.A., Karger, D.R.: Approximating st minimum cuts in $\tilde{O}(n^2)$ time. In: *Proceedings of the Symposium on Theory of Computing (STOC)*, pp. 47–55 (1996)
7. Benczúr, A.A., Karger, D.R.: Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM J. Comput.* **44**(2), 290–319 (2015)
8. Fung, W.-S., Hariharan, R., Harvey, N.J.A., Panigrahi, D.: A general framework for graph sparsification. *SIAM J. Comput.* **48**(4), 1196–1223 (2019)
9. Spielman, D.A., Teng, S.-H.: Spectral sparsification of graphs. *SIAM J. Comput.* **40**(4), 981–1025 (2011)
10. Spielman, D.A., Srivastava, N.: Graph sparsification by effective resistances. *SIAM J. Comput.* **40**(6), 1913–1926 (2011)
11. Koutis, I., Xu, S.C.: Simple parallel and distributed algorithms for spectral graph sparsification. *ACM Trans. Parallel Comput.* **3**(2), 14–11414 (2016). <https://doi.org/10.1145/2948062>
12. Lee, Y.T., Sun, H.: An sdp-based algorithm for linear-sized spectral sparsification. In: *Proceedings of the Symposium on Theory of Computing (STOC)*, pp. 678–687 (2017)
13. Karger, D.R.: Random sampling in cut, flow, and network design problems. *Math. Oper. Res.* **24**(2), 383–413 (1999)
14. Nagamochi, H., Ibaraki, T.: Computing edge-connectivity in multigraphs and capacitated graphs. *SIAM J. Discret. Math.* **5**(1), 54–66 (1992)
15. Nagamochi, H., Ibaraki, T.: A linear-time algorithm for finding a sparse k -connected spanning subgraph of a k -connected graph. *Algorithmica* **7**(1–6), 583–596 (1992)

16. Fung, W.S., Hariharan, R., Harvey, N.J.A., Panigrahi, D.: A general framework for graph sparsification. In: Proceedings of the Symposium on Theory of Computing (STOC), New York, NY, USA, pp. 71–80 (2011). <https://doi.org/10.1145/1993636.1993647>
17. Hariharan, R., Panigrahi, D.: A general framework for graph sparsification. In: CoRR abs/1004.4080 (2010). <https://arxiv.org/abs/1004.4080>
18. Batson, J., Spielman, D.A., Srivastava, N.: Twice-Ramanujan sparsifiers. *SIAM J. Comput.* **41**(6), 1704–1721 (2012)
19. Andoni, A., Chen, J., Krauthgamer, R., Qin, B., Woodruff, D.P., Zhang, Q.: On sketching quadratic forms. In: Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science, pp. 311–319 (2016)
20. Tarjan, R.E.: Efficiency of a good but not linear set union algorithm. *J. ACM* **22**(2), 215–225 (1975). <https://doi.org/10.1145/321879.321884>
21. Gawrychowski, P., Mozes, S., Weimann, O.: Minimum Cut in $O(m \log^2 n)$ Time. In: Czumaj, A., Dawar, A., Merelli, E. (eds.) Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP). Leibniz International Proceedings in Informatics (LIPIcs), vol. 168, pp. 57–15715. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany (2020). <https://doi.org/10.4230/LIPIcs.ICALP.2020.57>
22. Ghaffari, M., Nowicki, K., Thorup, M.: Faster algorithms for edge connectivity via random 2-out contractions. In: Proceedings of the Symposium on Discrete Algorithms (SODA), pp. 1260–1279 (2020)
23. Abraham, I., Durfee, D., Koutis, I., Krinninger, S., Peng, R.: On fully dynamic graph sparsifiers. In: Proceedings of the Symposium on Foundations of Computer Science (FOCS), pp. 335–344 (2016)
24. Kruskal, J.B.: On the shortest spanning subtree of a graph and the traveling salesman problem. *Proc. Am. Math. Soc.* **7**(1), 48–50 (1956)
25. Karger, D.R.: Global min-cuts in $\mathcal{RN}\mathcal{C}$, and other ramifications of a simple min-cut algorithm. In: Proceedings of the Symposium on Discrete Algorithms (SODA), vol. 93, pp. 21–30 (1993)
26. Karger, D.R., Stein, C.: A new approach to the minimum cut problem. *J. ACM* **43**(4), 601–640 (1996)
27. Cormen, T.H., Leiserson, C.E., Rivest, R.L., Stein, C.: Introduction to Algorithms. MIT Press, Cambridge, Massachusetts (2009)
28. Thorup, M.: Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM* **46**(3), 362–394 (1999)
29. Devroye, L.: Generating the maximum of independent identically distributed random variables. *Comput. Math. Appl.* **6**(3), 305–315 (1980). [https://doi.org/10.1016/0898-1221\(80\)90039-5](https://doi.org/10.1016/0898-1221(80)90039-5)
30. Karger, D.R., Klein, P.N., Tarjan, R.E.: A randomized linear-time algorithm to find minimum spanning trees. *J. ACM* **42**(2), 321–328 (1995)
31. Chernoff, H.: A measure of asymptotic efficiency for tests of a hypothesis based on the sum of observations. *Ann. Math. Stat.* **23**(4), 493–507 (1952)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.