

Single pass sparsification in the streaming model with edge deletions

Ashish Goel*

Michael Kapralov†

Ian Post‡

March 22, 2012

Abstract

In this paper we give a construction of cut sparsifiers of Benczúr and Karger in the *dynamic* streaming setting in a single pass over the data stream. Previous constructions either required multiple passes or were unable to handle edge deletions. We use $\tilde{O}(1/\epsilon^2)$ time for each stream update and $\tilde{O}(n/\epsilon^2)$ time to construct a sparsifier. Our ϵ -sparsifiers have $O(n \log^3 n/\epsilon^2)$ edges. The main tools behind our result are an application of sketching techniques of Ahn *et al.* [SODA'12] to estimate edge connectivity together with a novel application of sampling with limited independence and sparse recovery to produce the edges of the sparsifier.

1 Introduction

We study the problem of graph sparsification on dynamic graph streams. Graph sparsification was introduced by Benczúr and Karger [BK96], who gave a near linear time procedure that takes as input an undirected graph G on n vertices and constructs a weighted subgraph H of G with $O(n \log n/\epsilon^2)$ edges such that the value of every cut in H is within a $1 \pm \epsilon$ factor of the value of the corresponding cut in G . This algorithm has subsequently been used to speed up algorithms for a host of applications involving cuts and flows such as finding approximately minimum or sparsest cuts in graphs ([BK96, KRV09]) as well as other applications (e.g. [KL02]). Spielman and Teng introduced a stronger class of sparsifiers called spectral sparsifiers [ST08]. Subsequent work has developed a number of efficient algorithms for constructing cut and spectral sparsifiers [BK96, SS08, BSS09, KMST10, FHHP11, KP12].

The algorithms developed in [BK96, SS08, FHHP11, KP12] take near-linear time in the size of the graph and produce very high quality sparsifiers, but require random access to the edges of the input graph G , which is often prohibitively expensive in applications involving massive data sets. The streaming model of computation, which restricts algorithms to use a small number of passes over the input and space polylogarithmic in the size of the input, has been studied extensively in various application domains—see [Mut06] for an overview—but has proven too restrictive for even the simplest graph algorithms. Even testing $s - t$ connectivity requires $\Omega(n)$ space [HRR99]. The less restrictive semi-streaming model, in which the algorithm is restricted to use $\tilde{O}(n)$ space, is more suited for graph algorithms [FKM⁺05, McG09].

The problem of constructing graph sparsifiers in the semi-streaming model was first considered by Ahn and Guha [AG09], who gave a one-pass algorithm for finding Benczúr-Karger type sparsifiers with a slightly larger number of edges than the original Benczúr-Karger algorithm, $O(n \log n \log \frac{m}{n}/\epsilon^2)$ as opposed to $O(n \log n/\epsilon^2)$ using $\tilde{O}(n)$ space. Subsequently, [KL11] obtained an algorithm for constructing stronger *spectral* sparsifiers of

*Departments of Management Science and Engineering and (by courtesy) Computer Science, Stanford University. Email: ashishg@stanford.edu. Research supported in part by NSF grants 0915040 and 0904314.

†Institute for Computational and Mathematical Engineering, Stanford University. Email: kapralov@stanford.edu. Research supported by NSF grant 0904314 and a Stanford Graduate Fellowship.

‡Department of Computer Science, Stanford University, Email: itp@stanford.edu. Research supported by NSF grants 0915040 and 0904314.

size $O(n \log n / \epsilon^2)$ in a single pass in the streaming model. All of these algorithms work only in the *incremental model*, where edges can be added to the graph but not removed.

In a recent paper [AGM12] Ahn, Guha and McGregor introduced a beautiful graph sketching approach to streaming computations in dynamic streams, i.e. allowing both edge additions and deletions. They showed that connectivity can be determined in $\tilde{O}(n)$ space in this setting, and gave a *multi-pass* algorithm for obtaining cut sparsifiers in small space. Their techniques center around the use of linear sketches, which have been heavily studied in the field of compressed sensing/sparse recovery originating in [CRT06, Don06]. See [GI10] for a survey. The focus of this paper is to provide a *single-pass* implementation of cut sparsification on dynamic streams in the semi-streaming model.

Our results: Our main result is an algorithm for constructing cut sparsifiers in a single pass in dynamic streams with edge deletions. We prove

Theorem 1. *There exists a single-pass streaming algorithm for constructing an ϵ -cut sparsifier of an unweighted, undirected graph $G = (V, E)$ with n vertices and m edges in the dynamic model using $\tilde{O}(n/\epsilon^2)$ space. The size of the sparsifier is $O(n \log^3 n / \epsilon^2)$, and the runtime of the algorithm is $\tilde{O}(1/\epsilon^2)$ per update. At each point in the stream we can recover the edges of the sparsifier in time $\tilde{O}(n/\epsilon^2)$.*

Our sparsification algorithm works by sampling edges at a rate inversely proportional to their edge connectivity, which was shown to work in [FHHP11]. In order to do this we maintain two sets of data structures. The first estimates connectivities, and the second does the actual sampling. We estimate connectivities by sampling edges of the input graph at a geometric sequence of sampling rates and recovering connected components of these samples using a result of [AGM12]. The second set of data structures stores a linear sketch of the actual samples we use in our sparsifier, also sampling at a geometric sequence of rates. Using sparse recovery and the linearity of our sketch we are able to reconstruct the necessary samples when needed.

Organization: We start by giving preliminaries on graph sparsification in Section 2. We then describe the algorithm in Section 3 and Section 4. Maintaining our samples in the dynamic model requires knowing, for each edge in the graph, whether or not it was included in each sample. This can be easily achieved if we assume that our algorithm has access to $\tilde{\Theta}(n^2)$ independent random bits, which, however, is not feasible in $\tilde{O}(n)$ space. For simplicity of presentation, we first describe our algorithm assuming that it has access to $\tilde{\Theta}(n^2)$ random bits in Section 3. We show how to obtain sufficiently good estimates of edge connectivities in a single pass, as well as recover the edges of a sparsifier using sparse recovery techniques. In Section 4 we show how to remove the assumption that the algorithm has access to $\tilde{\Theta}(n^2)$ independent random bits using random variables with limited independence, obtaining a $\tilde{O}(n/\epsilon^2)$ space single-pass algorithm for sparsification in the dynamic model.

2 Sparsification preliminaries

We will denote by $G(V, E)$ the undirected input graph with vertex set V and edge set E with $|V| = n$ and $|E| = m$. For our purposes G will be unweighted, but the results in this section apply to weighted graphs. For any $\epsilon > 0$, we say that a weighted graph $G'(V, E')$ is an ϵ -sparsification of G if the weight of every cut in G' is within $(1 \pm \epsilon)$ of the corresponding cut in G .

Sparsification algorithms work by sampling edges with probabilities inversely proportional to some measure of connectivity. The simplest of these is edge-connectivity:

Definition 2. *A graph G is k -connected if the value of each cut in G is at least k , and an edge e has edge-connectivity c_e if c_e is the value of the minimum cut separating its endpoints.*

Graphs with high k -connectivity are particularly simple to sample:

Theorem 3 ([Kar99]). *Let $G = (V, E)$ be a k -connected graph on n nodes, and let G' be obtained from G by sampling edges independently with probability $p = \Theta(\log n / (\epsilon^2 k))$, and giving sampled edges weight $1/p$. Then G' is an ϵ -sparsifier of G with high probability.*

The Benczúr-Karger algorithm samples according to a more strict notion of connectivity, referred to as *strong connectivity*, defined as follows:

Definition 4 ([BK96]). *A k -strong component is a maximal k -connected vertex-induced subgraph. The strong connectivity of an edge e , denoted by s_e , is the largest k such that a k -strong component contains e , and we say e is k -strong if its strong connectivity is k or more, and k -weak otherwise.*

The following two lemmas will be useful in our analysis:

Lemma 5 ([BK96]). *The number of k -weak edges in a graph on n vertices is bounded by $k(n - 1)$.*

Lemma 6 ([BK96]). *Let $G = (V, E)$ denote an undirected graph on n nodes. For an edge $e \in E$ let s_e denote the strong connectivity of e . Then $\sum_{e \in E} 1/s_e \leq n - 1$.*

We also rely on Benczúr and Karger's main result, which is as follows:

Theorem 7 ([BK96]). *Let G' be obtained by sampling edges of G with probability $p_e = \min\{\rho / (\epsilon^2 s_e), 1\}$, where $\rho = 16(d + 2) \ln n$, and giving each sampled edge weight $1/p_e$. Then G' is an ϵ -sparsification of G with probability at least $1 - n^{-d}$. Moreover, the expected number of edges in G' is $O(n \log n)$.*

It follows easily from the proof of Theorem 7 in [BK96] that if we over-sample by using an *underestimate* of edge strengths, the resulting graph is still an ϵ -sparsification.

Corollary 8. *Let G' be obtained by sampling each edge of G with probability $\tilde{p}_e \geq p_e$ and give every sampled edge e weight $1/\tilde{p}_e$. Then G' is an ϵ -sparsification of G with probability at least $1 - n^{-d}$.*

Recently Fung *et al.* [FHHP11] proved that a more aggressive sampling method, namely sampling using edge connectivities as opposed to strong connectivities, also produces cut sparsifiers, and we will also require this result.

Theorem 9 ([FHHP11]). *Let G' be obtained from a weighted graph G by independently sampling edge e with probability $p_e = \rho / c_e$, where $\rho = \Theta(\log^2 n / \epsilon^2)$. Then, G' contains $O(n \log^2 n / \epsilon^2)$ edges in expectation and is an ϵ -sparsification whp.*

3 Sparsification with free randomness

In this section we present a dynamic sparsifier under the assumption that the algorithm has access to $\tilde{\Theta}(n^2)$ random words. We will remove this assumption in Section 4.

Our input graph G is undirected and unweighted. As in [AGM12], we will use the following representation of G :

Definition 10. *Given an unweighted graph $G = (V, E)$, let A_G be the $n \times \binom{n}{2}$ matrix with entry $(u, (v, w)) \in [n] \times \binom{[n]}{2}$ and $v < w$ given by*

$$a_{u, (v, w)} = \begin{cases} 1 & \text{if } u = v \text{ and } (v, w) \in E \\ -1 & \text{if } u = w \text{ and } (v, w) \in E \\ 0 & \text{otherwise} \end{cases}$$

Updates to the graph G in the form of the addition or deletion of an edge arrive one at a time in a streaming fashion. An update cannot delete an edge that does not exist or add one that already does, but other than these restrictions the order is adversarial, and the stream can be arbitrarily long. We need to maintain a data structure using only $\tilde{O}(n)$ space that allows us to efficiently construct an ϵ -sparsifier of the current graph G after any sequence of updates. We will accomplish this using a collection of linear sketches of the rows of A_G .

Our algorithm has two components: the first will maintain an estimate of the connectivity of each edge and therefore its sampling rate (as discussed in Section 2), and the second will store the actual samples. The former uses the tools developed by Ahn *et al.* [AGM12], and the latter is based on the technique of sparse recovery developed in the sketching and compressed sensing literature [GI10].

Before delving into the details, we elaborate on our use of randomness. In this section we assume that for each pair $(u, v) \in [n]^2$ the algorithm has access to a uniformly random number $h_{(u,v)} \in [0, 1]$. In fact, we will need $O(\log n)$ independent copies of these random numbers, which we will denote by $h_{(u,v)}^b$, $b = 1, \dots, O(\log n)$. These random variables will be used to estimate sampling rates for edges of G . We will also assume access to independent random numbers $g_{(u,v)}^b \in [0, 1]$, $b = 1, \dots, O(\log n)$, which we will use to determine a partition of the vertex set needed for sampling. Finally, we also assume access to independent random numbers $g_{(u,v)}^* \in [0, 1]$ which will be used to sample edges of the sparsifier. Once g , g^* and h are sampled, they are fixed for the duration of the algorithm. This is important for handling deletions, as it ensures that edges can be removed from exactly those sketches to which they have been added.

It will be convenient to think of all these numbers as independent in this section, even though this is not feasible in subquadratic space in the semi-streaming model. In Section 4 we will show that using numbers that are only $\tilde{O}(1)$ -wise independent for fixed u and independent for different v is sufficient, leading to a space-efficient solution. Some of our subroutines will also require their own internal entropy, but the total used will be only $\tilde{O}(n)$ words.

3.1 Estimating edge connectivity

The building block of our connectivity estimates is the following result from [AGM12] that finds connected components by sketching the rows of A_G :

Theorem 11 ([AGM12]). *There is a single-pass, linear sketch-based algorithm supporting edge additions and deletions that uses $O(n \log^3 n)$ space and returns a spanning forest of the graph with high probability.*

Our sketch is simple. We consider random samples of the input graph at geometric sampling rates and find connected components in each sample. Specifically, for each $a = 1, \dots, O(\log n)$ denote by $G_a^b = (V, E_a^b)$ a subsample of the edges of G obtained by setting

$$E_a^b = \left\{ (u, v) \in E : \min\{h_{(u,v)}^b, h_{(v,u)}^b\} < 2^{-a} \right\}. \quad (1)$$

We maintain connectivity data structures C_a^b for each of the subgraphs G_a^b for $a = 1, \dots, O(\log n)$, $b = 1, \dots, O(\log n)$ using Theorem 11. It is important to note that use of the functions $h_{(u,v)}$ for sampling rather than fresh randomness allows us to handle deletions properly by deleting a removed edge only from the sketches that we added it to by simply sampling and using the consistent samples as input to the sketch in Theorem 11.

Remark 12. *Note that an edge (u, v) is included in E_a^b if the minimum of $h_{(u,v)}^b$ and $h_{(v,u)}^b$ is smaller than 2^{-a} . This will be important for the proof of correctness for hash functions with limited independence in Section 4.*

The connectivity structure of the subgraphs G_a^b allows us to associate a sampling rate with each edge. For an edge $(u, v) \in E$ we use the smallest sampling rate 2^{-a} at which u and v are still in the same component

as an estimate of the sampling rate for (u, v) . Independent repetition $O(\log n)$ times reduces the variance sufficiently to get precise estimates.

Define \mathcal{V}_a as the partition of vertices in V induced by the intersection of all partitions C_a^b , $b = 1, \dots, O(\log n)$. That is, vertices u and v are in the same connected component in \mathcal{V}_a if and only if they are connected in C_a^b for all $b = 1, \dots, O(\log n)$. For an edge (u, v) let $L(u, v)$ —the *level of (u, v)* —denote the largest a such that u and v are in the same component in \mathcal{V}_a , and for a vertex v let the *level $L(v)$* denote the largest a such that v is not a singleton in \mathcal{V}_a .

The level $L(e)$ of an edge serves as a proxy for its connectivity:

Lemma 13. *For all edges $e \in E$ one has*

$$\Theta(s_e / \log n) \leq 2^{L(e)} \leq 2c_e$$

with high probability, where s_e denotes strong connectivity and c_e denotes edge connectivity.

Proof. The first inequality follows from the fact that $\Theta(\log n) \cdot 2^L$ -strongly connected components will stay connected with high probability by Theorem 3 when the functions $h_{(u,v)}^b$ used for sampling are truly random. Lemma 26 from Section 4 gives the result for hash functions with limited independence.

The second inequality follows by noting that if there is a cut separating u and v of size at most $2^L/2$, then it will be empty with probability at least $1/2$ when we sample at rate 2^{-L} by Markov's inequality, so u and v will get disconnected in one of the $O(\log n)$ independent repetitions with high probability. \square

This implies the levels can be used as sampling rates:

Lemma 14. *Sampling edges independently at rate $p_e = O(\log^2 n / (\epsilon^2 2^{L(e)}))$ and weighting sampled edges with $1/p_e$ produces a sparsifier with $O(n \log^3 n / \epsilon^2)$ edges high probability.*

Proof. By Theorem 9 sampling at rate $O(\log^2 n / (\epsilon^2 c_e))$ works. By Lemma 13 $1/2^{L(e)} \geq 1/(2c_e)$, and oversampling only improves concentration. This proves the statement assuming that the sampling of edges is independent.

The expected size of the sample is bounded by $\sum_{e \in E} p_e = O(\log^3 n / \epsilon^2) \cdot \sum_{e \in E} 1/s_e = O(n \log^3 n / \epsilon^2)$, where we used Lemma 13 to bound $p_e = O(\log n) / s_e$ and the fact that $\sum_{e \in E} 1/s_e \leq n - 1$ by Lemma 6. \square

3.2 Maintaining edge samples

We now show how to maintain small space sketches that will allow us to reconstruct the edges of the sparsifier. Our basic tool is the technique of sparse recovery from the field of compressed sensing [GI10]. A vector A of dimension N is k -sparse if it has at most k non-zero entries, and a k -sparse, or approximately k -sparse, signal can be recovered with high probability from $O(k \log(N/k))$ non-adaptive linear measurements. Here we use the following result of Cormode and Muthukrishnan [CM06] that allows recovery in $\tilde{O}(k)$ time at the cost of slightly sub-optimal sketch size and error guarantees:

Theorem 15 ([CM06]). *We can construct a randomized 0/1 matrix T of dimension $O(ck \log^3 n / \epsilon^2) \times N$ such that for any k -sparse signal A of dimension N , given the transformation TA we can reconstruct A exactly with probability at least $1 - n^{-c}$ in time $O(c^2 k \log^3 n / \epsilon^2)$. The matrix T is constructed using $O(1)$ -wise independent hash functions, and individual entries can be queried efficiently.*

We will also need the following result by Indyk [Ind00] on sketching ℓ_1 norms:

Theorem 16 ([Ind00]). *There is a linear sketch-based algorithm using $O(c \log^2 N / \epsilon^2)$ space and $O(c \log^2 N / \epsilon^2)$ random bits that can estimate the ℓ_1 norm of a vector of dimension N to within a factor of $(1 \pm \epsilon)$ with probability $1 - N^{-c}$. The sketch can be updated in $O(\log N)$ time.*

For a row v of the matrix A_G let $S_a^r(v)$ for $r = 1, \dots, O(\log n)$ denote linear sketches guaranteed by Theorem 15 for $k = O(\log^3 n/\epsilon^2)$ where a corresponds to the geometric sequence of sampling rates, and $r = 1, \dots, O(\log n)$ are independent copies that are useful for recovery. More precisely, $S_a^r(v)$ is a sketch of row v in the matrix $A_{G'_{a,r}}$ where $G'_{a,r}$ has edges

$$E'_{a,r} = \left\{ (u, v) \in E : \min\{g_{(u,v)}^r, g_{(v,u)}^r\} < 2^{-a} \right\}. \quad (2)$$

Some of these sketches may accumulate more than k edges and consequently cannot be decoded on their own, but we will prove this is not an issue. We will also need sketches $d_a^r(v)$, $r = 1, \dots, O(\log n)$ for the ℓ_1 -norm of the v -th row of the matrix $A_{G'_{a,r}}$. Note row v of $A_{G'_{a,r}}$ contains a ± 1 entry for each edge incident on v in $G'_{a,r}$, so the ℓ_1 norm corresponds exactly to the degree of v in the sampled graph $G'_{a,r}$.

Remark 17. *Note that we are using $O(\log n)$ independent samples $G'_{a,r}$ for each sampling rate a . This will be important in the proof of Lemma 21 below.*

Finally, we will also need another set of independent samples of G that will be used to obtain the edges of the sparsifier. Let G_a^* be the graph with edges

$$E_a^* = \left\{ (u, v) \in E : \min\{g_{(u,v)}^*, g_{(v,u)}^*\} < 2^{-a} \right\}. \quad (3)$$

For each node v and sampling rate a we maintain sketches $S_a^*(v)$ of row v in the matrix $A_{G_a^*}$ using Theorem 15 for $k = O(\log^3 n/\epsilon^2)$. Here we do not need independent repetitions for each sampling rate a .

By the choice of the matrix A_G and the linearity of the sketches, if $\mathcal{S} \subseteq V$ is a cut then $\sum_{v \in \mathcal{S}} d_a^r(v)$ is a sketch for the size of the cut and $\sum_{v \in \mathcal{S}} S_a^r(v)$ is sketch of a sample of its edges. If v is a supernode obtained by contracting a set of vertices \mathcal{S} , we write $d_a^r(v)$ to denote $\sum_{u \in \mathcal{S}} d_a^r(u)$ and similarly for $S_a^r(v)$ and $S_a^*(v)$. For any fixed cut and fixed G'_a the estimate given by $d_a^r(v)$ is close to expectation with high probability.

Before specifying the algorithm formally, we give the intuition behind it. Recall that for every vertex u at level $L(u) = a$, we need to sample edges going from u to vertices in u 's component in \mathcal{V}_a with probability $\gamma \log^2 n / (\epsilon^2 2^a)$ for an appropriate constant γ . In order to do that, we will sample *all edges incident on u* with probability $\gamma \log^2 n / (\epsilon^2 2^a)$ and then throw away the ones that do not go to u 's component. In order to obtain such a sample, we will use the sketches $S_a^r(u)$ that were made with sampling at rate $\gamma \log^2 n / (\epsilon^2 2^a)$.

The main observation here is that if we contract connected components in \mathcal{V}_{a+1} into supernodes, the resulting graph will have only $(\gamma \log n \cdot 2^{a+1})$ -weak edges for a constant γ , so the average degree will be no larger than $\gamma \log n \cdot 2^{a+1}$. By repeatedly removing vertices with degree at most twice the average, nodes of this subgraph can be partitioned into sets W_1, \dots, W_t such that for each $i = 1, \dots, t$ and $u \in W_i$ the degree of u in $W_i \cup \dots \cup W_t$ is at most $4\gamma \log n 2^a$ and $t = O(\log n)$. Formally,

Lemma 18. *Suppose all edges in G are k -weak. Then V can be partitioned into $t = \log n$ sets W_1, \dots, W_t such that for all $v \in W_i$ the degree of v when restricted to $W_i \cup \dots \cup W_t$ is at most $2k$.*

Proof. By Lemma 5 if G has n' nodes then it has at most $k(n' - 1)$ edges. Let W_1 be the set of all nodes with degree at most $2k$. By Markov's inequality W_1 includes at least half the nodes. After removing W_1 and all incident edges we can repeat this process to find W_2 , etc. At each iteration we remove at least half the nodes, so it terminates in $\log n$ iterations. \square

This partition cannot actually be computed because we cannot properly update the degree sketches after removing W_1 , but its existence allows us to prove that the same procedure works when using the lower degree sample $G'_{a,r}$.

Let $\Delta = \log(\gamma \log^2 n / \epsilon^2)$. We need to use the samples $S_{a-\Delta}^r(u)$ for edges at level a . We first bound the degrees in $G'_{a-\Delta,r}$:

Lemma 19. *Let $g_{(u,v)}^r$ be $\Theta(\log^3 n/\epsilon^2)$ -wise independent for fixed u , and independent for different u . Then the degree of all $u \in W_i$ in $G'_{a-\Delta,r}$ restricted to nodes $W_i \cup \dots \cup W_t$ is at most $O(\log^3 n/\epsilon^2)$ with high probability.*

Proof. Consider a vertex $u \in W_i$ and let $N(u)$ denote the neighbors of u in $W_i \cup \dots \cup W_t$ in the full graph G . The size of its sampled neighborhood is bounded by

$$\sum_{v \in N(u)} \mathbf{1}_{g_{(u,v)}^r < \gamma \log^2 n / (\epsilon^2 2^a)} + \sum_{v \in N(u)} \mathbf{1}_{g_{(v,u)}^r < \gamma \log^2 n / (\epsilon^2 2^a)}$$

The number of terms is $O(\log n 2^a)$. The second sum consists of independent random variables, so standard Chernoff bounds apply. Concentration bounds from Theorem 25 apply to the first sum since they are sufficiently independent for expectation $O(\log^3 n/\epsilon^2)$. \square

Observe that if a node u satisfies the bound in Lemma 19, then the sketch $S_{a-\Delta}^r(u)$ can be decoded using Theorem 15. Let U_1 be the set of decodable nodes. We would like to argue that we can simply output a decoded edge (u, v) as part of the sparsifier if and only if $g_{(u,v)}^r < \gamma \log^2 n / (\epsilon^2 2^a)$ and v belongs to the same connected component as u in \mathcal{V}_a . Then, using the linearity of the sketches, we could subtract decoded edges of the form (u, v) , $u \in U_1, v \in V \setminus U_1$ from the sketches $S(v)$ and $d(v)$, effectively removing U_1 from the graph, and then move on to U_2 .

However, for technical reasons to avoid dependencies and ensure the algorithm works in small space, we cannot reuse the variables $g_{(u,v)}^r$. We need to calculate U_2 using the independent sketches S^{r+1} and d^{r+1} as opposed to S^r and d^r to avoid dependencies and also use the variables $g_{(u,v)}^*$ for the actual samples. The size of r will remain bounded since we will prove the process terminates in $O(\log n)$ steps, but switching to S^{r+1} introduces additional complications because to remove a vertex $u \in U_1$ from the graph, we must be able to recover the edges from $S_{a-\Delta}^{r+1}(u), \dots, S_{a-\Delta}^{O(\log n)}(u)$. The following lemma accomplishes this:

Lemma 20. *Let $g_{(u,v)}^r$ be $\Theta(\log^3 n/\epsilon^2)$ -wise independent for fixed u , and independent for different u , and suppose the degree of u in $G'_{a-\Delta,r}$ is at most $\alpha \log^3 n/\epsilon^2$ where α is the constant from Lemma 19. Then the degree of u in $G'_{a-\Delta,r'}$ is $O(\log^3 n/\epsilon^2)$ for all $r' \geq r$ with high probability.*

Proof. If the degree of u in $G'_{a-\Delta,r}$ is at most $\alpha \log^3 n/\epsilon^2$, we expect the degree in G to be at most $(\alpha/\gamma) \log n 2^a$, and concentration inequalities for sampling with limited independence (Theorem 25) show that with high probability its degree is at most, say, $2(\alpha/\gamma) \log n 2^a$. Applying Theorem 25 again shows u 's degree in $G'_{a-\Delta,r'}$ is $O(\log^3 n/\epsilon^2)$ for any r' , and taking a union bound over all r' finishes the proof. \square

We now state the algorithm formally. For each $a = 1, \dots, O(\log n)$ we denote the graph obtained by contracting all connected components in \mathcal{V}_{a+1} into supernodes by H_a .

Algorithm 1: PARTITION(a)

- 1: Let $H_a^1 \leftarrow H_a$.
 - 2: **for** $r \leftarrow 1$ to $O(\log n)$ **do**
 - 3: Estimate the degree of each $v \in H_a^r$ from sketches $d_{a-\Delta}^r(v)$
 - 4: $U_a^r \leftarrow \{v \in H_a^r \mid d_{a-\Delta}^r(v) \leq 4\alpha \log^3 n/\epsilon^2\}$
 - 5: **for** $u \in U_a^r, j = r+1, \dots, O(\log n)$ **do**
 - 6: Run sparse recovery on $S_{a-\Delta}^j(u)$
 - 7: For all edges (u, v) recovered from $S_{a-\Delta}^j(u)$, subtract (u, v) from $S_{a-\Delta}^j(v)$ and $d_{a-\Delta}^j(v)$
 - 8: **end for**
 - 9: $H_a^{r+1} \leftarrow H_a^r \setminus U_a^r$.
 - 10: **end for**
 - 11: **return** $\{U_a^r\}_{r=1, \dots, O(\log n)}$
-

Here γ is a constant such that sampling the edges of a k -connected graph at rate $\gamma \log n/k$ produces a connected subgraph with probability at least $1 - n^{-10}$, and α is a constant bounding the degree in Lemma 19.

We first prove

Lemma 21. *For all $a = 1, \dots, O(\log n)$, Algorithm 1 recovers a partition of H_a such that for all $r = 1, \dots, O(\log n)$ for each $u \in U_a^r$ the degree of u in $U_a^{r+1} \cup \dots \cup U_a^{O(\log n)}$ in graph $G_{a-\Delta}^*$ is $O(\log^3 n/\epsilon^2)$ with high probability.*

Proof. We first prove that the constructed sets cover all of H_a . Consider the set U_1 . By Lemma 19, all $u \in W_1$ have degree $O(\log^3 n/\epsilon^2)$ in $G'_{a-\Delta,1}$ with high probability, so removing all nodes with degree at most $4\alpha \log^3 n/\epsilon^2$ for large enough α will include all $u \in W_1$. Lemma 20 implies that the sparse recovery in line 6 will succeed for all j with high probability, so we can completely remove U_a^r from the graph. Lemma 20 also bounds the degree of $u \in U_a^r$ in graph $G_{a-\Delta}^*$, by replacing $G'_{a-\Delta,r'}$ with $G_{a-\Delta}^*$ in the statement of the lemma.

We now note that the identity of the set U_a^r is independent of the randomness used for samples and sketches $S_{a-\Delta}^j, d_{a-\Delta}^j, j = r+1, \dots, O(\log n)$. Thus, the same bounds on node degrees follow by a recursive application of the argument to $H_a \setminus U_a^1$. Furthermore, it follows by induction on r that after removing U_a^1, \dots, U_a^r we have removed all of W_1, \dots, W_r with high probability, so the algorithm terminates in $O(\log n)$ iterations. \square

Given the partition $U_a^1, \dots, U_a^{O(\log n)}$, the algorithm for recovering the edges of the sparsifier is as follows:

Algorithm 2: RECOVER(a)

- 1: **for** $r = 1, \dots, O(\log n), u \in U_a^r$ **do**
 - 2: Run sparse recovery on $S_{a-\Delta}^*(u)$
 - 3: Output each recovered edge $(u, v), u \in U_r$ only if $g_{(u,v)}^* < \gamma \log^2 n/(\epsilon^2 2^a)$ and $L(u, v) = a$.
 - 4: Subtract recovered edges from $S_{a-\Delta}^*(v)$ for all $v \in U_{a,r+1} \cup \dots \cup U_{a,O(\log n)}$.
 - 5: **end for**
-

We can now prove

Theorem 22. *For each $v \in V(G)$ Algorithm 2 recovers a sample of edges incident on v , where edges are picked with probability $\gamma \log^2 n/(\epsilon^2 2^{L(v)})$.*

Proof. Note that sparse recovery succeeds whp by the degree bound in Lemma 21. Finally, note that the structure of the partition $U_1 \cup \dots \cup U_r$ maps to each edge a single random variable $g_{(u,v)}$, so the probability of an edge being sampled is correct. \square

We will need the following definition in Section 4:

Definition 23. *An edge $e = (u, v) \in E$ is controlled by a vertex $u \in V$ if e is sampled using $g_{(u,v)}^*$. We denote the set of edges controlled by u by E_u .*

We will also need

Lemma 24. *Let E^* be a set of edges. For each $u \in V$ one has $\mathbb{E}[|E^* \cap E_u|] = O(\log^4 n/\epsilon^2)$.*

Proof. Consider a vertex $u \in V$. By Lemma 21 u controls $O(\log^3 n 2^a/\epsilon^2)$ edges at level a . Hence, the expected number of edges sampled at each level a is $O(\log^3 n/\epsilon^2)$. Hence, the expected number of edges controlled by u across all levels is $O(\log^4 n/\epsilon^2)$. \square

3.3 Runtime

We now briefly summarize the time required to update the sketches and to construct a sparsifier. We do not optimize the log factors but only show updates require $\tilde{O}(1/\epsilon^2)$ time and building a sparsifier requires $\tilde{O}(n/\epsilon^2)$. Each addition or deletion of an edge requires updating C_a^b , S_a^r and d_a^r for $a, b, r \leq O(\log n)$. The sketches C_a^b are built using ℓ_0 -samplers (see [AGM12]) and can be updated in $\tilde{O}(1)$ time. For an edge (u, v) we update all $O(\log n)$ copies of $S(u)$, $S(v)$, $d(u)$ and $d(v)$. By Theorems 15 and 16, these can each be updated in $\tilde{O}(1/\epsilon^2)$ time. For $S(u)$ this is done by querying only the $\tilde{O}(1/\epsilon^2)$ entries of the matrix T we need.

Construction of the sparsifier is more expensive. If we query the sparsifier after each graph update it may need to be recomputed from scratch each time due to edge deletions, so we cannot amortize its cost across the updates. However, we will show it requires only $\tilde{O}(n/\epsilon^2)$ time. Building all $O(\log n)$ \mathcal{V}_a requires $\tilde{O}(n)$ operations each if ℓ_0 -sampling is done efficiently.

Running one iteration of Algorithm 1 requires $O(n)$ estimations of d , $O(n)$ decodings of S , $\tilde{O}(k)$ updates to S and d for each of the $O(n)$ sparse recoveries and $O(n)$ additional bookkeeping. Since S is $\tilde{O}(1/\epsilon^2)$ -sparse by Theorem 15 decoding takes $\tilde{O}(1/\epsilon^2)$ time. Summing over $\tilde{O}(1)$ values of r and a , we use a total of $\tilde{O}(n/\epsilon^2)$ time. Algorithm 2 also does $O(n)$ sparse recoveries and $\tilde{O}(n/\epsilon^2)$ updates to S^* per iteration, which also totals to $\tilde{O}(n/\epsilon^2)$ summing over all r and a .

3.4 Weighted graphs

We note that even though we stated the algorithm for unweighted graphs, the following simple reduction yields a single pass dynamic sparsifier for weighted graphs, as long as when an edge is removed or updated, its weight is given together with the identity of its endpoints. Suppose that edge weights are integers between 1 and W (the general case can be reduced to this one with appropriate scaling and rounding). Consider graphs $G_0, \dots, G_{\log_2 W}$, where an edge $e = (u, v)$ belongs to the edge set of G_b iff the binary expansion of w_e has 1 in position b , for $b = 0, \dots, \log_2 W$. Note that in order to preserve cuts in G to a multiplicative factor of $1 \pm \epsilon$, it is sufficient to preserve cuts in each of G_0, \dots, G_b to the same factor. To do that, it is sufficient to maintain $\log_2 W$ copies of our algorithm operating on the graphs G_b (this is feasible due to the assumption that edges are either added or completely removed, i.e. the weight of the removed edge is given at the time of removal). The space used and the number of edges in the sparsifier will both increase by a factor of $\log_2 W$.

4 Sparsification with limited independence

In this section we remove the assumption that the algorithm has access to $\tilde{\Theta}(n^2)$ bits of randomness by using sampling with limited independence. We prove the following two statements. First, we show in Lemma 26 that sampling edges of a k -connected graph at rate $\gamma \log n/k$ yields a connected graph with high probability, *even with limited independence*. In particular, it is sufficient to ensure that random variables used for sampling edges incident to any given vertex are only $\tilde{O}(1/\epsilon^2)$ -wise independent. This lemma is used in Section 3 to show that our estimation of sampling rates is accurate. We then show that our algorithm for constructing a sparsifier by sampling at rates proportional to edge connectivities yields a sparsifier with high probability even when the sampling is done using limited independence. We note that the second claim does not subsume the first due an extra $\log n$ factor that is needed for sampling with edge connectivities to go through.

We will use tail bounds for t -wise independent random variables proved in [SSS95], Theorem 5:

Theorem 25. *Let X_1, \dots, X_n be random variables each of which is confined to $[0, 1]$. Let $X = \sum_{i=1}^n X_i$, $\mu = E[X]$. Let $p = \mu/n$, and suppose that $p \leq 1/2$. Then if X_i are $\lceil \epsilon\mu \rceil$ -wise independent, then*

$$\Pr[|X - \mu| \geq \epsilon\mu] < e^{-\lfloor \epsilon^2\mu/3 \rfloor},$$

if $\epsilon < 1$, and

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

otherwise.

We now prove

Lemma 26. *Let $G = (V, E)$ be a k -connected graph on n nodes. For edges $e = (u, v) \in E$ let random numbers $h_{u,v} \in [0, 1]$ be such that*

1. $h_{u,v}$ is independent of $h_{u',v'}$ for all $u' \neq u$;
2. $h_{u,v}$ are $\lceil 4\gamma \log n \rceil$ -wise independent for fixed u , where $\gamma \geq 20$.

Also, let $X_{u,v}$ be 0/1 random variables such that $X_{u,v} = 1$ if $h_{u,v} \leq (\gamma \log n)/k$ and 0 otherwise. If G' is obtained by including each edge $(u, v) \in E$ such that $X_{u,v} = 1$ or $X_{v,u} = 1$, then G' is connected whp.

Proof. First, for each $e = (u, v) \in E$ let $\hat{X}_{u,v}$ denote 0/1 random variables such that $X_{u,v} = 1$ if $h_{u,v} < (\gamma \log n)/s_e$, where s_e is the strong connectivity of e . Define \hat{G}' as the graph obtained by including each edge $(u, v) \in E$ such that $\hat{X}_{u,v} = 1$ or $\hat{X}_{v,u} = 1$. Note that \hat{G}' is a subgraph of G' , so it is sufficient to show that \hat{G}' is connected whp.

Suppose that $F = (V_F, E_F)$ is a k -connected graph without $2k$ -strongly connected components for some k . Recall from Lemma 18 that V_F can be partitioned into $\log |V_F|$ sets $W_1, \dots, W_{\log |V_F|}$ such that the degree of any $u \in W_r$ in $W_r \cup \dots \cup W_{\log |V_F|}$ is at most $4k$. For each node $u \in W_r$ let E_u denote the edges incident on u that go to nodes in $W_r \cup \dots \cup W_{\log |V_F|}$ (if an edge $e = (u, v)$ goes between two nodes in W_r , include it either in E_u or E_v arbitrarily). We will say vertex $u \in W_r$ controls edges $e \in E_u$. Note that $|E_u| \leq 4k$ for all $u \in W_r$, and hence the expected number of edges sampled in E_u is at most $4\gamma \log n$. Note that this definition of control is slightly different from the one given in Definition 23. In particular, this is because Definition 23 pertains to the actual sampling procedure that our algorithm uses, while here we are concerned with the estimation step.

Let $j_{max} = \lfloor \log_2 n \rfloor$. We will show by induction on $j = j_{max}, \dots, 0$ that all 2^j -connected components are connected with probability at least $1 - (j_{max} - j + 1)n^{-3}$.

Base: $j = j_{max}$ We have $\kappa = 2^{j_{max}}$. Apply the decomposition above to the κ -strongly connected components of G , which does not have any 2κ -connected components since $2\kappa > n$. Let \hat{G}'' denote the subgraph of G obtained by including for each $u \in V$ edges $e = (u, v) \in E_u$ when $X_{u,v} = 1$. Denote the set of sampled edges by E' . Recall that for all $u \in V(G)$ one has $\mathbf{E}[|E' \cap E_u|] \leq 4\gamma \log n$.

Fix a cut $(C, V \setminus C)$. For each vertex $u \in C$ let $X_u = \sum_{(u,v) \in E_u, v \notin C} X_{u,v}$.

By setting $\epsilon = 1$ in Theorem 25 we get

$$\Pr[X_u = 0] < e^{-\mathbf{E}[X_u]/3}.$$

Since $X_u, X_{u'}$ are independent for $u \neq u'$, the probability that the cut is empty is at most

$$\prod_{u \in C} e^{-\mathbf{E}[X_u]/3} = e^{-\gamma |C| \log n / (3k)}.$$

By Karger's cut counting lemma, the number of cuts of value at most αk is at most $n^{2\alpha}$. A union bound over all cuts, we get failure probability at most

$$\sum_{\alpha \geq 1} n^{2\alpha} e^{-\gamma \alpha \log n / 3} \leq n^{-4}$$

since $\gamma \geq 20$. Taking a union bound over all κ -connected components yields failure probability at most n^{-3} .

Inductive step: $j + 1 \rightarrow j$ We have $\kappa = 2^j$. By the inductive hypothesis, all 2^{j+1} -connected components will be connected with probability at least $1 - (j_{max} - (j + 1) + 1)n^{-3}$. We condition on this event and contract the connected components into supernodes.

We now have a union of vertex-disjoint κ -strongly connected components that do not contain any 2κ -connected components. The same argument as in the base case shows that each such component will be connected with probability at least $1 - n^{-4}$. A union bound over at most n such components completes the inductive step. □

In order to show that the results of [FHHP11] carry over to our setting, it is sufficient to show that the following version of Chernoff bounds holds under our limited independence assumptions (Theorem 2.2 in [FHHP11]):

Theorem 27. *Let X_1, \dots, X_n be n random variables such that X_i takes value $1/p_i$ with probability p_i and 0 otherwise. Then, for any p such that $p \leq p_i$, for each i , any $\epsilon \in (0, 1)$ and any $N \geq n$ the following holds:*

$$\Pr \left[\left| \sum_{i=1}^n X_i - n \right| > \epsilon N \right] < 2e^{-0.38\epsilon^2 p N}.$$

Indeed, an inspection of the proofs of Lemma 4.1 and Lemma 5.5 in [FHHP11] shows that the authors (a) only rely on independence of their sampling process to obtain Theorem 27 and (b) only apply Theorem 27 to subsets of edges of G , where p_i are sampling probabilities. Thus, proving an equivalent of Theorem 27 allows us to extend their results to our limited independence sampling approach.

We now prove

Lemma 28. *Let $G = (V, E)$ denote an unweighted undirected graph. Let $\gamma > 0$ be a sufficiently large constant such that sampling at rate $\gamma \log^2 n / c_e$ independently produces a sparsifier with probability at least $1 - n^{-2}$, where c_e is the edge connectivity of e . Let $X_e, e \in E$ be random variables corresponding to including edges from a set E^* into the sample such that X_e takes value $1/p_e$ with probability p_e and 0 otherwise, where p_e is the sampling probability used by Algorithm 2. Assume that sampling is $c \log^4 n / \epsilon^2$ -wise independent for a sufficiently large constant $c > 0$ that may depend on γ .*

There exists an event \mathcal{E} with $\Pr[\mathcal{E}] > 1 - n^{-2}$ such that for any $E^ \subseteq E$, any $p \leq p_e, e \in E^*$, any $\epsilon \in (0, 1)$ and any $N \geq |E^*|$*

$$\Pr \left[\left| \sum_{e \in E^*} X_e - |E^*| \right| > \epsilon N | \mathcal{E} \right] < e^{-\epsilon^2 p N / 6}.$$

Proof. For simplicity of exposition, we now assume that G is unweighted. Let X_1, \dots, X_n be random variables corresponding to picking edges of the graph. Recall that our sampling algorithm samples an edge (u, v) either depending on the value of $g_{(u,v)}$ or the value of $g_{(v,u)}$ (the choice depends on the partition of the node set $U_1 \cup \dots \cup U_r$ constructed in Algorithm 2). Recall that by Definition 23 a node u controls edge (u, v) if Algorithm 2 samples (u, v) using the value of $g_{(u,v)}$. Note that each edge is controlled by exactly one node. For a node u , as before, let E_u denote the set of edges controlled by u . By Lemma 24, one has $\mathbf{E}[|E^* \cap E_u|] = O(\log^4 n / \epsilon^2)$. Let \mathcal{E} denote the event that at most $2\gamma \log^4 n / \epsilon^2$ edges controlled by u are sampled, for all $u \in V$, where we are assuming that γ is sufficiently large. Since our random variables are $c \log^4 n / \epsilon^2$ -wise independent for sufficiently large c , by Theorem 25 and a union bound over all u one has $\Pr[\mathcal{E}] \geq 1 - n^{-2}$.

For each $e \in E$ let X_e be a Bernoulli random variable that takes value p/p_e if edge e is sampled, and 0 otherwise, so that $X_e \in [0, 1]$.

Consider a set of edges E^* . Partition E^* as $E^* = \bigcup_{u \in V} E_u^*$, where $E_u^* = E^* \cap E_u$. Thus, random variables $\mathcal{X}_u := \sum_{e \in E_u^*} X_e$ are independent for different u . Let $\mathcal{X} = \sum_{u \in V} \mathcal{X}_u$, $\mu = \mathbf{E}[\mathcal{X}]$.

Then by Markov's inequality

$$\Pr[\mathcal{X} \geq (1 + \delta)\mu | \mathcal{E}] \leq \frac{\mathbf{E}[e^{t\mathcal{X}} | \mathcal{E}]}{e^{t(1+\delta)\mu}}. \quad (4)$$

Recall that

$$\mathbf{E}[e^{t\mathcal{X}} | \mathcal{E}] = \sum_{j=0}^{\infty} \mathbf{E}[(t\mathcal{X})^j | \mathcal{E}] / j! = \sum_{j=0}^{\infty} t^j / j! \sum_{S \subseteq E^*, |S| \leq j} \sum_{\alpha_e \geq 0, \sum_{e \in S} \alpha_e = j} \mathbf{E} \left[\prod_{e \in S} X_e^{\alpha_e} | \mathcal{E} \right] \quad (5)$$

For any non-negative random variable \mathcal{Y} one has

$$\mathbf{E}[\mathcal{Y} | \mathcal{E}] \leq \mathbf{E}[\mathcal{Y}] / \Pr[\mathcal{E}].$$

Conditional on \mathcal{E} , one has $\prod_{e \in S} X_e^{\alpha_e} = 0$ for all $S \subseteq E^*$ such that $|S \cap E_u| > 2\gamma \log^4 n / \epsilon^2$ for at least one $u \in V$. For other S , setting $\mathcal{Y} = \prod_{e \in S} X_e^{\alpha_e}$, one gets

$$\mathbf{E} \left[\prod_{e \in S} X_e^{\alpha_e} | \mathcal{E} \right] \leq \mathbf{E} \left[\prod_{e \in S} X_e^{\alpha_e} \right] / \Pr[\mathcal{E}]. \quad (6)$$

Combining (6) and (5) one gets

$$\mathbf{E}[e^{t\mathcal{X}} | \mathcal{E}] \leq \frac{1}{\Pr[\mathcal{E}]} \sum_{j=0}^{\infty} t^j / j! \sum_{S \subseteq E^*, |S| \leq j, |S \cap E_u| \leq 2\gamma \log^4 n / \epsilon^2} \sum_{\alpha_e \geq 0, \sum_{e \in S} \alpha_e = j} \mathbf{E} \left[\prod_{e \in S} X_e^{\alpha_e} \right] \quad (7)$$

On the other hand, for all $S \subseteq E^*$ such that $|S \cap E_u| \leq 2\gamma \log^4 n / \epsilon^2$ one has

$$\mathbf{E} \left[\prod_{e \in S} X_e^{\alpha_e} \right] = \prod_{e \in S} \mathbf{E}[X_e^{\alpha_e}]$$

by $\gamma \log^4 n / \epsilon^2$ -wise independence. Thus, we get

$$\Pr[\mathcal{X} \geq (1 + \delta)\mu | \mathcal{E}] \leq \frac{1}{\Pr[\mathcal{E}]} \frac{\prod_{e \in E^*} \mathbf{E}[e^{tX_e}]}{e^{t(1+\delta)\mu}}, \quad (8)$$

which is the same bound as in the full independence case, except for a factor of $1/\Pr[\mathcal{E}] = 1 + O(1/n)$ in front. Now the same derivation as in the full independence case shows that the probability of overestimating is appropriately small.

We now bound the probability of underestimating. Consider a set of edges E^* . Partition E^* as $E^* = \bigcup_{i=1}^s E_i$, where $E_i \cap E_j = \emptyset, i \neq j$, so that

1. $\mathbf{E}[\sum_{e \in E_i} X_e] \leq c \log^4 n / \epsilon^2$ for a sufficiently large constant $c > 0$;
2. random variables $\sum_{e \in E_i} X_e$ are independent for different i ;
3. $s \leq \frac{\epsilon^2}{6 \log(4/\epsilon^2)} \mathbf{E}[\sum_{e \in E^*} X_e]$.

Note that this is feasible since our graphs are unweighted, so ϵ can be assumed to be larger than $1/n^2$. For each $i = 1, \dots, s$ let $\mathcal{X}_i := \sum_{e \in E_i} X_e$. Note that \mathcal{X}_i are independent, and X_e are $c \log^4 n / \epsilon^2$ -wise independent. Now by Theorem 25 for all $i = 1, \dots, s$ one has for all $\epsilon \in (0, 1)$

$$\Pr[\mathcal{X}_i < \mathbf{E}[\mathcal{X}_i] - \epsilon \mathbf{E}[\mathcal{X}_i]] < e^{-\epsilon^2 \mathbf{E}[\mathcal{X}_i]/3} \quad (9)$$

Let $\mathcal{X} = \sum_{i=1}^s \mathcal{X}_i$. For constant $\epsilon > 0$ let

$$\mathcal{K}(\epsilon) = \left\{ \mathbf{z} = (z_1, z_2, \dots, z_s) \in \left\{ 0, \frac{1}{4}\epsilon^2, \frac{1}{2}\epsilon^2, \frac{3}{4}\epsilon^2, \dots, 1 - \epsilon^2/4, 1 \right\}^s : \sum_{i=1}^s z_i \mathbf{E}[\mathcal{X}_i] \geq \epsilon \mathbf{E}[\mathcal{X}] \right\}.$$

We now have

$$\begin{aligned} \Pr[\mathcal{X} < \mathbf{E}[\mathcal{X}] - \epsilon \mathbf{E}[\mathcal{X}]] &\leq \sum_{\mathbf{z} \in \mathcal{K}(\epsilon)} \prod_{i=1}^s \Pr[\mathcal{X}_i < \mathbf{E}[\mathcal{X}_i] - (z_i - \epsilon^2/4) \mathbf{E}[\mathcal{X}_i]] \\ &\leq \sum_{\mathbf{z} \in \mathcal{K}(\epsilon)} \prod_{i=1}^s \Pr[\mathcal{X}_i < \mathbf{E}[\mathcal{X}_i] - z_i^2 \mathbf{E}[\mathcal{X}_i] + (\epsilon^2/2) \mathbf{E}[\mathcal{X}_i]] \end{aligned} \quad (10)$$

since every set of values for $\mathcal{X}_i - \mathbf{E}[\mathcal{X}_i]$ such that $\sum_i (\mathcal{X}_i - \mathbf{E}[\mathcal{X}_i]) < -\epsilon \mathbf{E}[\mathcal{X}]$ can be rounded to a point in $\mathcal{K}(\epsilon)$ with a loss of at most $(\epsilon^2/2) \mathbf{E}[\mathcal{X}_i]$ in each term. We now note that for any $\mathbf{z} \in [0, 1]^s$ such that $\sum_{i=1}^s z_i \mathbf{E}[\mathcal{X}_i] = \epsilon' \mathbf{E}[\mathcal{X}] \geq \epsilon \mathbf{E}[\mathcal{X}]$ one has $\sum_{i=1}^s z_i^2 \mathbf{E}[\mathcal{X}_i] \geq (\epsilon')^2 \mathbf{E}[\mathcal{X}] \geq \epsilon^2 \mathbf{E}[\mathcal{X}]$.

Next, since $s \leq \frac{\epsilon^2}{6 \log(4/\epsilon^2)} (\sum_{i=1}^s \mathbf{E}[\mathcal{X}_i])$, we have that

$$\Pr[\mathcal{X} < \mathbf{E}[\mathcal{X}] - \epsilon \mathbf{E}[\mathcal{X}]] \leq (4/\epsilon^2)^s e^{-\epsilon^2 \mathbf{E}[\mathcal{X}]/3} \leq e^{-\epsilon^2 \mathbf{E}[\mathcal{X}]/6}. \quad (11)$$

□

Theorem 29. *The set of edges returned by Algorithm 2 is a sparsifier whp.*

Proof. Lemma 28 can be used instead of Theorem 27 in [FHHP11]. □

References

- [AG09] K. Ahn and S. Guha. Graph sparsification in the semi-streaming model. *Automata, Languages and Programming*, pages 328–338, 2009. arXiv:0902.0140v2.
- [AGM12] Kook Jin Ahn, Sudipto Guha, and Andrew McGregor. Analyzing graph structure via linear measurements. *SODA*, pages 459–467, 2012. Available from: <http://siam.omnibooksonline.com/2012SODA/data/papers/017.pdf>.
- [BK96] András A. Benczúr and David R. Karger. Approximating s - t minimum cuts in $\tilde{O}(n^2)$ time. *Proceedings of the 28th annual ACM symposium on Theory of computing*, pages 47–55, 1996. doi:10.1145/237814.237827.
- [BSS09] Joshua D. Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-ramanujan sparsifiers. *STOC*, pages 255–262, 2009. arXiv:0808.0163v3.
- [CM06] G. Cormode and S. Muthukrishnan. Combinatorial algorithms for compressed sensing. *Structural Information and Communication Complexity*, 4056:280–294, 2006. doi:10.1007/11780823_22.

- [CRT06] E.J. Candès, J. Romberg, and T. Tao. Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *Information Theory, IEEE Transactions on*, 52(2):489–509, 2006. arXiv:0409186v1.
- [Don06] D.L. Donoho. Compressed sensing. *Information Theory, IEEE Transactions on*, 52(4):1289–1306, 2006. doi:10.1109/TIT.2006.871582.
- [FHHP11] Wai Shing Fung, Ramesh Hariharan, Nicholas J. A. Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. *STOC*, pages 71–80, 2011. doi:10.1145/1993636.1993647.
- [FKM⁺05] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. On graph problems in a semi-streaming model. *Theoretical Computer Science*, 348:207–216, 2005. doi:10.1016/j.tcs.2005.09.013.
- [GI10] A. Gilbert and P. Indyk. Sparse recovery using sparse matrices. *Proceedings of the IEEE*, 98(6):937–947, 2010. doi:10.1109/JPROC.2010.2045092.
- [HRR99] M.R. Henzinger, P. Raghavan, and S. Rajagopalan. Computing on data streams. In *External memory algorithms: DIMACS Workshop External Memory and Visualization, May 20-22, 1998*, volume 50, pages 107–118. American Mathematical Society, 1999. Available from: <https://www.hpl.hp.com/techreports/Compaq-DEC/SRC-TN-1998-011.pdf>.
- [Ind00] P. Indyk. Stable distributions, pseudorandom generators, embeddings and data stream computation. In *Proceedings of 41st Annual Symposium on Foundations of Computer Science*, pages 189–197. IEEE, 2000. doi:10.1109/SFCS.2000.892082.
- [Kar99] D. Karger. Random sampling in cut, flow, and network design problems. *Mathematics of Operations Research (Preliminary version appeared in the Proceedings of the 26th annual ACM symposium on Theory of computing)*, 24(2):383–413, 1999. doi:10.1145/195058.195422.
- [KL02] D. Karger and M. Levine. Random sampling in residual graphs. *STOC*, 2002. doi:10.1145/509907.509918.
- [KL11] Jonathan A. Kelner and Alex Levin. Spectral sparsification in the semi-streaming setting. *STACS*, pages 440–451, 2011. doi:10.4230/LIPIcs.STACS.2011.440.
- [KMST10] Alexandra Kolla, Yury Makarychev, Amin Saberi, and Shang-Hua Teng. Subgraph sparsification and nearly optimal ultrasparsifiers. *STOC*, pages 57–66, 2010. arXiv:0912.1623v1.
- [KP12] Michael Kapralov and Rina Panigrahy. Spectral sparsification via random spanners. *ICTS*, pages 393–398, 2012. doi:10.1145/2090236.2090267.
- [KRV09] R. Khandekar, S. Rao, and U. Vazirani. Graph partitioning using single commodity flows. *Journal of the ACM (JACM)*, 56(4):19:1–19:15, 2009. doi:10.1145/1538902.1538903.
- [McG09] A. McGregor. Graph mining on streams. *Encyclopedia of Database Systems*, pages 1271–1275, 2009. Available from: <http://people.cs.umass.edu/~mcgregor/papers/08-graphmining.pdf>.
- [Mut06] S. Muthukrishnan. *Data streams: algorithms and applications*. Now publishers, 2006. doi:10.1561/0400000002.

- [SS08] D.A. Spielman and N. Srivastava. Graph sparsification by effective resistances. *STOC*, pages 563–568, 2008. [arXiv:0803.0929v4](https://arxiv.org/abs/0803.0929v4).
- [SSS95] Jeanette P. Schmidt, Alan Siegel, and Aravind Srinivasan. Chernoff-hoeffding bounds for applications with limited independence. *SIAM J. Discrete Math.*, 8(2):223–250, 1995. doi: 10.1137/S089548019223872X.
- [ST08] D.A. Spielman and S.H. Teng. Spectral sparsification of graphs. *Arxiv preprint*, 2008. [arXiv:0808.4134v3](https://arxiv.org/abs/0808.4134v3).