

Towards Tight Bounds for Spectral Sparsification of Hypergraphs

Michael Kapralov*
École Polytechnique Fédérale de Lausanne
michael.kapralov@epfl.ch

Robert Krauthgamer†
Weizmann Institute of Science
robert.krauthgamer@weizmann.ac.il

Jakab Tardos‡
École Polytechnique Fédérale de Lausanne
jakab.tardos@epfl.ch

Yuichi Yoshida
National Institute of Informatics
yyoshida@nii.ac.jp

June 9, 2021

Abstract

Cut and spectral sparsification of graphs have numerous applications, including e.g. speeding up algorithms for cuts and Laplacian solvers. These powerful notions have recently been extended to hypergraphs, which are much richer and may offer new applications. However, the current bounds on the size of hypergraph sparsifiers are not as tight as the corresponding bounds for graphs.

Our first result is a polynomial-time algorithm that, given a hypergraph on n vertices with maximum hyperedge size r , outputs an ϵ -spectral sparsifier with $O^*(nr)$ hyperedges, where O^* suppresses $(\epsilon^{-1} \log n)^{O(1)}$ factors. This size bound improves the two previous bounds: $O^*(n^3)$ [Soma and Yoshida, SODA'19] and $O^*(nr^3)$ [Bansal, Svensson and Trevisan, FOCS'19]. Our main technical tool is a new method for proving concentration of the nonlinear analogue of the quadratic form of the Laplacians for hypergraph expanders.

We complement this with lower bounds on the bit complexity of any compression scheme that $(1+\epsilon)$ -approximates all the cuts in a given hypergraph, and hence also on the bit complexity of every ϵ -cut/spectral sparsifier. These lower bounds are based on Ruzsa-Szemerédi graphs, and a particular instantiation yields an $\Omega(nr)$ lower bound on the bit complexity even for fixed constant ϵ . In the case of hypergraph cut sparsifiers, this is tight up to polylogarithmic factors in n , due to recent result of [Chen, Khanna and Nagda, FOCS'20]. For spectral sparsifiers it narrows the gap to $O^*(r)$.

Finally, for directed hypergraphs, we present an algorithm that computes an ϵ -spectral sparsifier with $O^*(n^2r^3)$ hyperarcs, where r is the maximum size of a hyperarc. For small r , this improves over $O^*(n^3)$ known from [Soma and Yoshida, SODA'19], and is getting close to the trivial lower bound of $\Omega(n^2)$ hyperarcs.

*Supported in part by ERC Starting Grant 759471.

†Work partially supported by ONR Award N00014-18-1-2364, the Israel Science Foundation grant #1086/18, and a Minerva Foundation grant.

‡Supported by ERC Starting Grant 759471.

Contents

1	Introduction	1
1.1	Results	2
1.2	Related Work	4
1.3	Discussion	4
2	Preliminaries	5
2.1	Hypergraph and Expansion	5
2.2	Spectral Hypergraph Theory	5
3	Technical Overview	6
3.1	Spectral Sparsification of Expanders	6
3.2	General Spectral Sparsification of Hypergraphs	8
3.3	Lower Bounds	10
3.4	Directed Spectral Sparsification of Hypergraphs	12
4	Spectral Sparsification of Expanders	13
4.1	Construction	14
4.2	Correctness	14
4.3	Proofs of Claims 4.6, 4.7, 4.8, and 4.9	16
5	Expander Decomposition	21
5.1	Approximating Sparsest Cut	21
5.2	Proof of Lemma 5.1	23
6	General Spectral Sparsification of Hypergraphs	24
6.1	Construction	24
6.2	Correctness	26
6.3	Proofs of Claims 6.9, 6.10, and 6.11	29
6.4	Weighted Hypergraphs and Proof of Theorem 1.1	32
7	Lower Bounds	32
7.1	String Compression	32
7.2	Construction	34
8	Spectral Sparsification of Directed Hypergraphs	38
8.1	Preliminaries	38
8.2	Nearly Equally Overlapping Directed Hypergraphs	40
8.3	Proofs of Claims 8.9, 8.10, 8.11, and 8.12	43
8.4	Proof of Theorem 1.3	45
A	Technical lemmas	50
A.1	Concentration Inequalities	50
A.2	Proof of Hypergraph Cheeger’s Inequality	50

1 Introduction

Sparsification is an algorithmic paradigm where a dense object is replaced by a sparse one with similar features, which often leads to significant improvements in efficiency of algorithms, including running time, space complexity, and communication. We study edge-sparsification of hypergraphs, which replaces a hypergraph $G = (V, E, w)$ with a sparse hypergraph \tilde{G} that has the same vertex set V and only a few hyperedges, often a reweighted subset of E . This is a natural extension of edge-sparsification of ordinary graphs, which includes key concepts such as cut sparsifiers, spectral sparsifiers, and flow sparsifiers. These were studied extensively from numerous angles, including various constructions, tight size bounds, related variants, and practical applications. As this literature is too vast to cover here, we quickly recap the basics for graphs before discussing hypergraphs, which are our focus here.

Graphs. Let $G = (V, E, w)$ be an edge-weighted graph, where $w \in \mathbb{R}_+^E$. The *energy* of a vector $x \in \mathbb{R}^V$ in G is defined as

$$Q_G(x) = \sum_{uv \in E} w_{uv} (x_u - x_v)^2,$$

and can also be written as $x^\top L_G x$, where L_G is the Laplacian matrix of G . Spielman and Teng [ST11a] introduced the notion of an ϵ -spectral sparsifier of G , which is a graph $\tilde{G} = (V, \tilde{E}, \tilde{w})$ that satisfies (for $0 \leq \epsilon \leq 1/2$)

$$\forall x \in \mathbb{R}^V, \quad Q_{\tilde{G}}(x) \in (1 \pm \epsilon)Q_G(x). \quad (1)$$

The *size* of a spectral sparsifier \tilde{G} is $|\tilde{E}|$.

We say that an edge $e \in E$ is *cut* by $S \subseteq V$ if one endpoint of e belongs to S and another one belongs to $V \setminus S$. The total weight of edges cut by S is clearly $Q_G(1_S)$, where $1_S \in \mathbb{R}^V$ denotes the characteristic vector of a set $S \subseteq V$.

A spectral sparsifier \tilde{G} of a graph G preserves many important properties of G : its cuts have approximately the same weight as those in G ; its Laplacian $L_{\tilde{G}}$ approximates every eigenvalue of L_G ; electrical flows in \tilde{G} approximate those in G . It is extremely useful to have a spectral sparsifier with a small number of edges because algorithms that involve these quantities can be applied on the sparsifier \tilde{G} instead of on G , with only a small loss in accuracy.

A spectral sparsifier of size $O(n/\epsilon^2)$ can be computed in almost linear time [LS18], where n is the number of vertices in G .

Hypergraphs. A hypergraph is a natural extension of a graph, which can represent relations between three or more entities, and has proved useful to solve problems in practical areas such as computer vision [HLM09, OB12], bioinformatics [KHT09], and information retrieval [GKR00]. Many of those problems, such as semi-supervised learning [HSJR13, YNY⁺19, ZHTC20] and link prediction [YNN⁺20], involve the notion of energy for hypergraphs, where the *energy* of a vector $x \in \mathbb{R}^V$ in an edge-weighted hypergraph $G = (V, E, w)$ is defined as

$$Q_G(x) = \sum_{e \in E} w_e \max_{u, v \in e} (x_u - x_v)^2. \quad (2)$$

This definition matches the one for graphs when every hyperedge in G is of size two. As before, $Q_G(1_S)$ gives the total weight of hyperedges cut by S , where we regard a hyperedge $e \in E$ as *cut* if $e \cap S \neq \emptyset$ and $e \cap (V \setminus S) \neq \emptyset$.

Table 1: Bounds on the size of hypergraph sparsifiers

cut sparsification	spectral sparsification	reference
$\tilde{O}(n^2/\epsilon^2)$		[NR13] implicitly
$\tilde{O}(nr/\epsilon^2)$		[KK15]
	$\tilde{O}(n^3/\epsilon^2)$	[SY19]
	$\tilde{O}(nr^3/\epsilon^2)$	[BST19]
$\tilde{O}(n/\epsilon^2)$		[CKN20]
	$\tilde{O}(nr/\epsilon^{O(1)})$	Theorem 1.1

Spectral sparsification of hypergraphs was first defined by Soma and Yoshida [SY19], as follows. Similarly to graphs, an ϵ -spectral sparsifier of G is a hypergraph $G = (V, \tilde{E}, \tilde{w})$ that satisfies (1). This is a strictly stronger notion than that of the hypergraph cut sparsifier which has been previously studied in [NR13] and [KK15].

Besides the applications mentioned above, spectral sparsifiers for hypergraphs were used to show agnostic learnability of a certain subclass of submodular functions [SY19].

Soma and Yoshida [SY19] showed that every hypergraph G admits an ϵ -spectral sparsifier with $\tilde{O}(n^3/\epsilon^2)$ hyperedges,¹ which is non-trivial because a general hypergraph can have $2^n - 1$ (non-empty) hyperedges. Moreover, they provide an algorithm recovering this sparsifier, that runs in close to linear time (in the input size). Later, Bansal, Svensson and Trevisan [BST19] showed that every hypergraph G admits a spectral sparsifier with $\tilde{O}(nr^3/\epsilon^2)$ hyperedges, where r is the maximum size of a hyperedge in G . Note that this bound is incomparable to [SY19] because r could be as large as n .

1.1 Results

Spectral sparsification of undirected hypergraphs. Our first contribution is an algorithm that constructs an ϵ -spectral sparsifier of a hypergraph with only $\tilde{O}(nr/\epsilon^{O(1)})$ hyperedges, which improves upon the previous constructions mentioned above. (See Table 1 for known bounds for hypergraph sparsification.)

Theorem 1.1. *Given an r -uniform hypergraph $G = (V, E, w)$ and $1/n \leq \epsilon \leq 1/2$, one can compute in polynomial time with probability $1 - o(1)$ an ϵ -spectral sparsifier of G with $nr(\epsilon^{-1} \log n)^{O(1)}$ hyperedges. The running time is $O(mr^2) + n^{O(1)}$, where $m = |E|$.*

To simplify notation, our entire technical analysis considers a hypergraph $G = (V, E)$ that is unweighted (i.e., unit weight hyperedges), reserving the letter w for the edge weights in the sparsifier.

This is actually without loss of generality, see Section 6.4.

We stress that Theorem 1.1 in fact applies to hypergraphs with maximum size of a hyperedge at most r . Indeed, in our analysis every hyperedge is a multiset of vertices, and therefore a hyperedge with less than r vertices can be trivially extended to a multiset of exactly r vertices by copying an arbitrary vertex, without changing the energy (but it might affect vertex degrees).

¹Throughout, we write $\tilde{O}(\cdot)$ to suppress a factor of $\log^{O(1)} n$.

Bit-complexity lower bound. To complement Theorem 1.1, we consider lower bounds on the bit complexity of sparsifiers. Here, we consider ϵ -cut sparsifiers, which require that (1) holds only for vectors of the form $x = 1_S$. This notion actually predates spectral sparsification and was first defined by Benczúr and Karger [BK15] for graphs, and by Kogan and Krauthgamer [KK15] for hypergraphs. Obviously, lower bounds for cut sparsifiers directly imply the same lower bounds also for spectral sparsifiers.

The second contribution of this work is a surprising connection between a *Ruzsa-Szemerédi (RS) graph* [RS78], which is a well-studied notion in extremal graph theory, and a lower bound on the bit complexity of a hypergraph cut sparsifier. Here, an (ordinary) graph is called a (t, a) -RS graph if its edge set is the union of t induced matchings of size a . Then, we show the following.

Theorem 1.2. *Suppose that there exists a (t, a) -Ruzsa-Szemerédi graph on n vertices with $a \geq 6000\sqrt{n \log n}$. Assume also one can compress unweighted $(t + 1)$ -uniform hypergraphs $G = (V, E)$ on $2n$ vertices into k bits, from which $Q_G(1_S)$ can be approximated for every $S \subseteq V$ within factor $1 \pm \epsilon$, where $\epsilon = O(a/n)$. Then, $k = \Omega(at)$.*

For example, by instantiating Theorem 1.2 with the $(n^{\Omega(1/\log \log n)}, n/3 - o(n))$ -Ruzsa-Szemerédi graphs known due to Fischer et al. [FLN⁺02], we deduce that $\Omega(nr)$ bits are necessary to encode all the cut values of an arbitrary r -uniform hypergraph with $r = n^{O(1/\log \log n)}$, even within a fixed constant ratio $1 + \epsilon$.

This lower bound is in fact near-tight. Indeed, Chen, Khanna, and Nagda [CKN20] showed very recently that every hypergraph G admits an ϵ -cut sparsifier with $O(n \log n / \epsilon^2)$ hyperedges, which are actually sampled from G . Applying this construction with fixed ϵ and $r = n^{O(1/\log \log n)}$ yields a sparsifier of G with $O(n \log n)$ hyperedges; encoding a hyperedge (including its weight, which is bounded by n^r) takes at most $O(r \log n)$ bits, and thus one can encode all the cuts of G using $O(nr \log^2 n)$ bits. It follows that our lower bound is optimal up to a lower order factor $O(\log^2 n)$. Instantiating our lower bound with the original construction of Ruzsa and Szemerédi [RS78], we can rule out the possibility of compressing the cut structure of a hypergraph with n vertices and maximum hyperedge size r with significantly less than nr space, and a polynomial scaling in the error (that is with $nr^{1-\Omega(1)}\epsilon^{-O(1)}$ space), for any r . See Corollaries 7.11, 7.12 and 7.13 in Section 7 for more details.

In fact, our space lower bound for hypergraphs far exceeds the $O(n \log n / \epsilon^2)$ bits that suffices to approximately represent all the cuts of an (ordinary) graph by simply storing a cut sparsifier. We thus obtain the first provable separation between the bit complexity of approximating all the cuts of a graph vs. of a hypergraph.

Spectral sparsification of directed hypergraphs. We also consider spectral sparsification of directed hypergraphs. Here, a hyperarc e consists of two disjoint sets, called the *head* $h(e) \subseteq V$ and the *tail* $t(e) \subseteq V$, and the *size* of the hyperarc is $|t(e)| + |h(e)|$. A directed hypergraph $G = (V, E)$ then consists of a vertex set V and a hyperarc set E . For an edge-weighted directed hypergraph $G = (V, E, w)$ and a vector $x \in \mathbb{R}^V$, the *energy* of x in G is defined as

$$Q_G(x) = \sum_{e \in E} w_e \max_{u \in t(e), v \in h(e)} (x_u - x_v)_+^2, \quad (3)$$

where $(a)_+ = \max\{a, 0\}$. Again, it is defined so that $Q_G(1_S)$ is the total weight of hyperarcs that are cut by S , where a hyperarc e is *cut* if $t(e) \cap S \neq \emptyset$ and $h(e) \cap (V \setminus S) \neq \emptyset$.

It is not difficult to see that a spectral sparsifier might require (in the worst-case) at least $\Omega(n^2)$ hyperarcs, even for an ordinary directed graph. Indeed, consider a balanced bipartite clique directed

from one side of the bipartition towards the other. Here, every arc is the unique arc crossing some particular directed cut, and hence a sparsifier must keep all the $\Omega(n^2)$ arcs (see also [IT18, CPS20]). However, Soma and Yoshida [SY19] showed that every directed hypergraph admits an ϵ -spectral sparsifier with $\tilde{O}(n^3/\epsilon^2)$ hyperarcs. We tighten this gap by showing that $\tilde{O}(n^2/\epsilon^2)$ hyperarcs are sufficient when every hyperarc is of constant size.

Theorem 1.3. *Given a directed hypergraph $G = (V, E)$ with maximum hyperarc size at most r such that $11r \leq \sqrt{en}$, and a value $\epsilon \leq 1/2$, one can compute in polynomial time with probability $1 - o(1)$ an ϵ -spectral sparsifier of G with $O(n^2 r^3 \log^2 n / \epsilon^2)$ hyperarcs.*

We note that Theorem 1.3 is stated under the assumption $11r \leq \sqrt{en}$, which is useful for our analysis for technical reasons. For larger values of r the result of [SY19] gives a better bound on the number of hyperedges in the sparsifier, and therefore this assumption is not restrictive.

1.2 Related Work

The first construction of cut sparsifiers for hypergraphs was given by Kogan and Krauthgamer [KK15] and uses $O(n(r + \log n)/\epsilon^2)$ hyperedges. They also mention that an upper bound of $O(n^2 \log n / \epsilon^2)$ hyperedges follows implicitly from the results of Newman and Rabinovich [NR13]. Very recently (and independent of our work), Chen, Khanna, and Nagda [CKN20] improved this bound to $O(n \log n / \epsilon^2)$ hyperedges, which is near-optimal because the current lower bound is $\Omega(n/\epsilon^2)$ edges, and actually holds for (ordinary) graphs [ACK⁺16, CKST19].

Louis [Lou15] (later merged with Chan et al. [CLTZ18]) initiated the spectral theory for hypergraphs, in which the Laplacian operator $L : \mathbb{R}^V \rightarrow \mathbb{R}^V$ of a hypergraph is defined so that its “quadratic form” $xL(x)$ coincides with the energy (2). As opposed to the graph case, here the Laplacian operator is merely piecewise linear, and hence computing its eigenvalues/vectors is hard. He showed that $O(\log r)$ -approximation is possible, and that obtaining a better approximation ratio is NP-hard assuming the Small-Set Expansion (SSE) hypothesis [RS10]. He further showed a Cheeger inequality for hypergraphs, which implies that, given a vector $x \in \mathbb{R}^V$ with a small energy, one can efficiently find a set $S \subseteq V$ of small expansion, which roughly means that the number of hyperedges cut by S is small relative to the number of hyperedges incident to vertices in S (see Section 2 for details). Since then, several other algorithms for finding sets of small expansion have been proposed [TMIY20, IMTY18].

Yoshida [Yos16] proposed another piecewise linear Laplacian for directed graphs and used it to study structures of real-world networks. Generalizing the Laplacians for hypergraphs and directed graphs, Laplacian L for directed hypergraphs was proposed [LM18, Yos19], whose quadratic form $x^\top L(x)$ coincides with (3).

1.3 Discussion

An obvious open question is the existence of a spectral sparsifier with $\tilde{O}(n)$ hyperedges. As we will see in Section 3, our overall strategy to construct a spectral sparsifier is decomposing the input hypergraph into good expanders (in a non-trivial way) and then sparsifying each expander. Here a *good expander* is a hypergraph with the maximum possible expansion up to a constant factor (see Section 2.1 for the details). However, we do not even know whether we can spectrally sparsify hypergraph expanders with $\tilde{O}(n)$ hyperedges. To see the difficulty, note that a graph expander has expansion $\Theta(1)$ whereas an r -uniform hypergraph expander has expansion $\Theta(1/r)$. Let $x \in \mathbb{R}^V$ be a vector with $\sum_{v \in V} x_v^2 d(v) = 1$, where $d(v)$ is the degree of a vertex $v \in V$. Then by the Cheeger inequality for hypergraphs (Theorem 2.3), the energy of x in a graph expander is $\Omega(1)$ whereas

that in an r -uniform hypergraph expander is merely $\Omega(1/r)$. Hence preserving the latter energy is seemingly a harder problem.

2 Preliminaries

In the paper, we will often need to deal with additive or multiplicative errors of various approximations. For simplicity of notation we use $\tilde{A} = A \pm \delta$ to denote $A - \delta \leq \tilde{A} \leq A + \delta$, and we use $\tilde{A} = (1 \pm \epsilon)A$ to denote $(1 - \epsilon)A \leq \tilde{A} \leq (1 + \epsilon)A$.

2.1 Hypergraph and Expansion

A hypergraph $G = (V, E)$ on a vertex set V is usually defined so that E is a set of hyperedges, each of which is an arbitrary (non-empty) subset of V (as opposed to ordinary graphs, where it is a subset of size two). In a slight departure from the norm, we allow the hyperedges in E to be multisets instead. That is, a hyperedge may contain certain vertices multiple times. This may be thought of as a generalization of the use of self-loops in ordinary graphs, which can be considered as multisets containing a single vertex with multiplicity two — and thus having size two. This slight change in the definition allows us to consider r -uniform hypergraphs throughout the paper without loss of generality, which makes the analysis in Section 6 much simpler. We call a hypergraph r -uniform if all of its hyperedges have size r .

Let us denote the multiplicity of a vertex $v \in V$ in $e \in E$ by $\mu_e(v)$. Then the *size* of e is $\sum_{v \in V} \mu_e(v)$ (as is normal for multisets). The *degree* of a vertex v is $d(v) = \sum_{e \in E} \mu_e(v)$.

Furthermore, we also allow hyperedges in E to appear with multiplicity, i.e., parallel edges. This means that E itself is a multiset. We call a hypergraph that has neither multiset edges nor multiple instances of the same edge a *simple hypergraph*.

For a hypergraph $G = (V, E)$ and a set $S \subseteq V$, let $E(S) \subseteq E$ be the multiset of hyperedges $e \in E$ such that every vertex in e belongs to S . Then, let $G[S] = (S, E(S))$ denote the subgraph of G induced by S .

Let $G = (V, E)$ be a hypergraph and $S \subseteq V$ be a vertex set. The *volume* of S , denoted by $\text{vol}(S)$, is $\sum_{v \in S} d(v)$. We say that a hyperedge $e \in E$ is *cut* by S if $e \cap S \neq \emptyset$ and $e \cap (V \setminus S) \neq \emptyset$. In this context, we often call a pair $(S, V \setminus S)$ a *cut*. Let $E(S, V \setminus S)$ denote the set of hyperedges cut by S . Then, the *expansion* of S (or a cut $(S, V \setminus S)$) is

$$\Phi(S) = \frac{|E(S, V \setminus S)|}{\min\{\text{vol}(S), \text{vol}(V \setminus S)\}}.$$

The expansion of a hypergraph $G = (V, E)$ is defined to be $\Phi(G) := \min_{S \subseteq V} \Phi(S)$. For $\Phi \geq 0$, we say that G is a Φ -*expander* if $\Phi(G) \geq \Phi$.

2.2 Spectral Hypergraph Theory

We briefly review spectral theory for hypergraphs. See, e.g., [CLTZ18, Yos19] for more details.

Definition 2.1. Let $G = (V, E)$ be a hypergraph and $x \in \mathbb{R}^V$ be a vector. The energy of a hyperedge $e \in E$ with respect to x is defined as $Q_x(e) = \max_{a, b \in e} (x_a - x_b)^2$, and the energy of a subset of hyperedges $E' \subseteq E$ is $Q_x(E') = \sum_{e \in E'} Q_x(e)$, respectively. Finally, the entire energy of x is defined as the energy of all hyperedges combined, that is, $Q(x) = Q_x(E)$. If the underlying hypergraph G is unclear from context, we specify by writing $Q_G(x)$.

Definition 2.2. Let $G = (V, E)$ be a hypergraph and $\epsilon > 0$. $\tilde{G} = (V, \tilde{E}, w)$ is a weighted subgraph of G if w is a vector in \mathbb{R}_+^E , mapping each hyperedge $e \in E$ to a non-negative value, and \tilde{E} denotes $\{e \in E \mid w_e > 0\}$. Such a weighted subgraph is called an ϵ -spectral sparsifier if for any vector $x \in \mathbb{R}^V$, $\tilde{Q}(x) = (1 \pm \epsilon) \cdot Q(x)$, where \tilde{Q} denotes energy with respect to the graph \tilde{G} , that is

$$\tilde{Q}(x) = \sum_{e \in \tilde{E}} w_e \cdot Q_x(e).$$

The size of such a sparsifier is $|\tilde{E}|$.

Given a hypergraph $G = (V, E)$ and a vector $x \in \mathbb{R}^V$, we can define an ordinary graph $G_x = (V, E_x)$ so that the energy of x on G and that on G_x are equal. Specifically, we define E_x as the multiset

$$E_x = \left\{ \left(\operatorname{argmax}_{a \in e} x_a, \operatorname{argmin}_{b \in e} x_b \right) \mid e \in E \right\},$$

where ties are broken arbitrarily.

The following Cheeger’s inequality is a cornerstone of spectral hypergraph theory. Although a similar theorem has been proven in [CLTZ18, Theorem 6.1], we include the proof in Appendix A for completeness because we do not know whether their proof goes through when we allow for multiset hyperedges.

Theorem 2.3 (Hypergraph Cheeger’s inequality). *Let $G = (V, E)$ be an r -uniform hypergraph with expansion at least $\Phi \leq 2/r$. Then for any vector $x \in \mathbb{R}^V$ with $\sum_{v \in V} x_v d(v) = 0$, we have*

$$Q(x) \geq \frac{r\Phi^2}{32} \sum_{v \in V} x_v^2 d(v).$$

Remark 2.4. *In fact, for simple hypergraphs the requirement $\Phi \leq 2/r$ is unnecessary and the statements holds in full generality. In our setting, this requirement is crucial, as non-simple r -uniform hypergraphs may have expansion $\omega(1/r)$, in which case the statement clearly does not hold.*

3 Technical Overview

In this section we briefly outline the techniques used in the proofs of our main results.

3.1 Spectral Sparsification of Expanders

We begin by constructing spectral sparsifiers for “good” hypergraph expanders, where we call a hypergraph a good expander if it has expansion at least $\tilde{\Omega}(1/r)$. Even in this restricted case, no result better than $\tilde{O}(nr^3/\epsilon^2)$ [BST19] was known previously. Our plan will then be to partition general input hypergraphs into a series of good expanders. The expansion $\tilde{\Omega}(1/r)$ is in some sense the best we can hope for. In fact, r -uniform simple hypergraphs cannot have an expansion better than $\Theta(1/r)$ and consequently no expander decomposition algorithm can guarantee expansion more than that.

To construct our spectral sparsifier for a good expander, we apply importance sampling to the input hypergraph. We sample each hyperedge e independently with some probability p_e and scale it up with weight $1/p_e$ if sampled. This guarantees that $\mathbb{E}(\tilde{G}) = G$ and so for any vector $x \in \mathbb{R}^V$ we have $\mathbb{E}\tilde{Q}(x) = Q(x)$, where \tilde{Q} denotes the energy with respect to the sparsifier. In our case, p_e is inversely proportional to $\min_{v \in e} d(v)$, and then the expected number of sampled hyperedges is

proportional to n — simply charge each hyperedge e to a vertex $v \in e$ of minimum degree, then each vertex is in charge of $O(1)$ sampled hyperedges in expectation. It remains to prove that the random quantity $Q(x)$ concentrates well around its expectation *for all vectors x simultaneously*.

So far this is a known technique: similar approaches to constructing spectral sparsifiers in ordinary graphs have appeared in many works, starting from [ST11b, SS11]. However, all of these rely on concentration inequalities for linear functions of independent random variables related to the matrix Bernstein inequality — see, e.g., [Tro15]². Unfortunately, the energy of a hypergraph is not a linear transformation and such tools cannot be applied to it. Two recent works on spectral sparsification of hypergraphs developed methods for circumventing this problem, namely [SY19] and [BST19]. The former uses a rather crude union bound plus Chernoff bound argument, and loses a factor of n in the size of the sparsifier, both for undirected and directed hypergraphs. The latter, namely the recent work of [BST19] uses Talagrand’s comparison inequality and generic chaining to compare the hypergraph sampling process to effective resistance sampling of [SS11], and loses a factor of r^3 in the size of the sparsifier. In this work we derive a simultaneous concentration inequality for $\tilde{Q}(x)$ for all $x \in \mathbb{R}^V$ from more basic principles, and obtain a sparsifier with $\approx nr$ hyperedges as a result — a bound that is seemingly best that can be obtained through the expander decomposition route.

Note that for a single, fixed vector $x \in \mathbb{R}^V$, the concentration inequality $\tilde{Q}(x) = (1 \pm \epsilon)Q(x)$ holds with high probability by the Chernoff bound (Theorem A.1). Our broad strategy will be to prove concentration over individual choices of x , and combine these results through a union bound. An obvious issue is that x is a continuous variable, making a direct union bound infeasible. We therefore have to discretize it, rounding each x to some \tilde{x} from a finite net. Our plan then becomes to prove the chain of approximations

$$Q(x) \cong Q(\tilde{x}) \cong \tilde{Q}(\tilde{x}) \cong \tilde{Q}(x),$$

where the second approximation ($Q(\tilde{x}) \cong \tilde{Q}(\tilde{x})$) utilizes the idea above of a Chernoff bound for each \tilde{x} plus a union bound over the net.

This turns out to be too simplistic, and the analysis requires a more technical discretization of x . Recall that the energy of the whole hypergraph can be written as a sum of the energies of the individual hyperedges:

$$Q(x) = \sum_{e \in E} Q_x(e).$$

We categorize hyperedges based on a carefully chosen metric $\max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v)$, which we will call the hyperedge’s *power*. If a hyperedge’s power is approximately 2^{-i} , then it resides in the i^{th} category E_i (see Section 4.2). We have in total a logarithmic number of categories. This categorization is important, because the power of a hyperedge turns out to be closely related to the strength of the Chernoff bound applicable to it, as well as to the required accuracy of the approximation \tilde{x} . That is, some cruder approximation \tilde{x} may be sufficient to guarantee $Q_x(E_1) \cong Q_{\tilde{x}}(E_1)$, but it might not be able to guarantee the same for a later category. Conversely, the Chernoff bound is stronger (i.e., the failure probability is smaller) at larger values of i . Thus, for each i we discretize x into a different vector $x^{(i)}$ (rather than the same \tilde{x}) and we prove individually for each i that

$$Q_x(E_i) \cong Q_{x^{(i)}}(E_i) \cong \tilde{Q}_{x^{(i)}}(E_i) \cong \tilde{Q}_x(E_i).$$

²More precisely, the proof of the necessary concentration properties in [ST11b] heavily relies on linearity of the graph Laplacian (specifically, the proof proceeded by bounding the trace of a high power of a corresponding matrix using combinatorial methods), and the analysis of [SS11] relies on a concentration inequality for linear functions of independent random variables due to Rudelson and Vershynin [RV07]. Both of these proofs can also be reproduced using the matrix Bernstein inequality.

Here, “ \cong ” necessarily covers both multiplicative *and additive* errors. Indeed, we have no guarantee on the sizes of these categories. Some E_i could contain only a single hyperedge, in which case a simple Chernoff bound would yield no concentration whatsoever. This is where we utilize the additive-multiplicative version (Theorem A.2). Since we have $\Theta(\log n)$ categories to sum over, we naturally allow additive error $\Theta(\epsilon Q(x)/\log n)$.

Note that $x^{(i)}$ is a discretization of x specialized to preserve the energies of hyperedges in E_i . Intuitively, the energy of such a hyperedge e is dictated by the largest value of x_v^2 within it. This value necessarily belongs to a vertex satisfying $x_v^2 d(v) \gtrsim 2^{-i}$. Thus, it should be enough for our rounding to preserve the x -values of vertices that satisfy this. To this end, we round the x -values of vertices with $x_v^2 d(v) \gtrsim 2^{-i}$ *carefully* — by an inverse polynomial amount in n . However, we round the x -values of all other vertices to 0 — which is obviously a crude (non-careful) rounding. Thus, if there are only k_i vertices we have to be careful about, the number of possible settings of $x^{(i)}$ becomes $\approx \exp(k_i)$.

Recall the formula of the additive-multiplicative Chernoff bound from Theorem A.2. In our case, the allowable multiplicative error is always $\approx 1 + \epsilon$, while the allowable additive error is always $\approx \epsilon Q(x)$. The only quantity that varies from level to level is the range of the random variables involved. If a specific hyperedge is sampled, it is scaled up by $1/p_e \approx \min_{v \in e} d(v)$, and the energy of this weighted hyperedge can be upper bounded by $\approx \max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v)$ — exactly the power of the hyperedge. Thus, at level i , the additive-multiplicative Chernoff bound guarantees a failure probability of $\approx \exp(-2^i Q(x))$. (Here we omit the ϵ terms, along with others, for simplicity.)

Finally, we want to equate the terms in the exponents of the Chernoff bound with the enumeration of $x^{(i)}$'s, so as to bound the total failure probability. We use hypergraph Cheeger (Theorem 2.3) to relate k_i to $Q(x)$. Suppose that x is normalized in the sense that $\sum_{v \in V} x_v^2 d(v) = 1$. This immediately gives that $k_i \leq 2^i$ by definition. On the other hand, we can finally use our assumption that the input hypergraph G was a good expander, since hypergraph Cheeger gives us that $Q(x) \gtrsim 1/r$. This makes the error probability for individual $x^{(i)}$'s $\approx \exp(-2^i/r)$ (from Chernoff bounds), while the enumeration of all $x^{(i)}$ becomes $\approx \exp(2^i)$. To bridge this gap, we must sacrifice a factor r in the sampling ratio p_e , and correspondingly in the size of the output sparsifier (see proof of Claim 4.7).

The formal proof is far more involved, and can be found in Section 4.

3.2 General Spectral Sparsification of Hypergraphs

Having constructed spectral sparsifiers for good expanders, we move our attention to arbitrary input hypergraphs. We decompose the vertex set of the input hypergraph $G = (V, E)$ into clusters of good expansion, while being careful not to cut too many hyperedges between the clusters. We adapt well-known techniques to the setting of hypergraphs, and is detailed for completeness in Section 5. As is common for expander decompositions, we partition V into clusters C_1, \dots, C_k such that the internal expansion of each cluster (along with its induced hyperedges) is at least $\tilde{\Omega}(1/r)$ while cutting only a constant fraction of the hyperedges between the clusters.

In ordinary graphs, this would immediately yield the desired result: We could simply decompose G into expanders and sparsify these, then repeat this process on the discarded hyperedges. Since the number of hyperedges decreases by a constant factor at each level, this process terminates after $O(\log n)$ levels of expander decomposition; each vertex only participates in $O(\log n)$ expanders, and thus the size bound of the overall sparsifier only suffers a logarithmic factor compared to the sparsifiers of expanders. For hypergraphs, this is not the case. Even simple, r -uniform hypergraphs may have up to nearly n^r hyperedges. This means that such a decomposition process could require $r \log n$ levels to terminate, introducing another factor r in the size of the sparsifier.

To combat this problem, we contract clusters into individual supernodes after sparsifying them

(see Algorithm 2). This allows us to simply bound the number of clusters a single vertex can participate in, and consequently the size of the output sparsifier. However, proving the correctness of this more complicated algorithm introduces new challenges.

We denote the contracted version of the input hypergraph G by G/\approx , where $u \approx v$ if the two vertices u and v have been contracted into the same supernode. We can equate between the hyperedges of G and those of G/\approx using the natural bijection between them (this means that a hyperedge e in G refers also to the corresponding hyperedge in G/\approx , and vice versa). Note that this operation can produce multiple parallel hyperedges, as well as *vertices appearing within the same hyperedge with multiplicity*, even if these phenomena were not allowed in the input hypergraph. It is important to note that our expander sparsification algorithm from Section 4 works equally well in this setting. Furthermore, by allowing hyperedges to contain vertices with multiplicity higher than 1, we may continue to work with r -uniform hypergraphs throughout this process of repeatedly contracting vertices. This technicality is crucial, since our expander decomposition algorithm is designed for this setting, and does not work when hyperedges have different sizes (by more than a constant factor).

The main technical contribution of Section 6 is to show that a sparsifier computed *after contraction* still sufficiently approximates the energy of the input hypergraph *before contraction*. Here we take a simplified example: Suppose we wish sparsify a cluster $C \subset V$ and subsequently contract it into a supernode v_C . At a later level we might wish to sparsify some other cluster C' that contains v_C as one of its vertices (see Figure 1). The result is a (weighted) subset of hyperedges that well-approximates the spectral structure of C' , but will this still be the case when we un-contrast v_C ?




Figure 1: Illustration of the contraction process. Vertices inside C are contracted into a single supernode v_C . This is then contained in a later cluster C' .

Denote the hyperedges of C' by E' , and let their sparsifier be \tilde{E}' (which is a weighted subset of E'). Being a sparsifier with respect to the contracted hypergraph can be viewed as being a sparsifier on the original hypergraph *only when $x \in \mathbb{R}^V$ is uniform, i.e., takes the same value, on all vertices of C* , as in this case we can simply assign that same value to v_C , and the energy of the original and contracted hypergraphs will be the same. Unfortunately, we have to deal with general vectors $x \in \mathbb{R}^V$, so we quantify how far x is from satisfying that uniformity requirement. We

consider the maximum discrepancy between the x -values of C , defined as $\delta = \max_{u,v \in C} |x_u - x_v|$. We show in Section 6.3 that the additive error introduced by taking \tilde{E}' as a sparsifier to E' in the original hypergraph – as opposed to the contracted hypergraph where it is guaranteed to be a good sparsifier – is proportional to δ^2 per hyperedge (see the proof of Claim 6.13).

We handle this additive error by arguing that it is dwarfed by energy of x with respect to C . On the one hand, we introduce δ^2 error per hyperedge of C' for a total of at most $\approx \delta^2 d' n$, where d' is the typical degree in C' . On the other hand, we know that the range of x within C is δ , so by hypergraph Cheeger (Theorem 2.3) the energy of C is at least $\approx \delta^2 d/r$, where d is the typical degree in C . (Here we assume that there are no outlier vertices with extremely low degree, which can be guaranteed by a slight adaptation of the expander decomposition subroutine, Lemma 5.1.) Recall that the number of hyperedges — and therefore the typical degree — decreases by a constant factor per level. If we can simply guarantee that the sparsification of C precedes the sparsification of C' by at least $\Omega(\log n)$ levels, then d will dwarf d' by an arbitrarily large $n^{\Theta(1)}$ factor. We accomplish this by simply waiting $\Omega(\log n)$ levels to contract a cluster after sparsifying it (see Algorithm 2).

The formal proof is far more involved, but relies on the same concept of charging additive errors to previous clusters, until we ultimately achieve the desired overall error of $\epsilon Q(x)$. The details appear in Section 6.

3.3 Lower Bounds

The most common method for approximating the Laplacian of a (hyper)graph is to take a weighted subset of the original (hyper)edges. While asymptotically optimal for graphs [ACK⁺16, CKST19], this method has obvious limitations as a data structure: it is not hard to come up with an example where $\Omega(n)$ hyperedges are required even for the sparsifier to be connected, and if the input hypergraph is r -uniform, this translates into $\Omega(nr \log n)$ bit complexity, a linear loss in the arity r of the hypergraph. It is therefore natural to ask whether there are more efficient ways of storing a spectral approximation to a hypergraph. As concrete example, we could permit the inclusion of hyperedges not in the original hypergraph – could this or another scheme lead to a data structure that can approximate the spectral structure of a hypergraph using $\tilde{O}(n)$ space, avoiding a dependence on r ?

In Section 7, we study this question in full generality:

Is it possible to compress a hypergraph into a $o(n \cdot r)$ size data structure that can approximate the energy $Q_G(x)$ (defined in (2)) simultaneously for all $x \in \mathbb{R}^V$?

In Section 7, we show a space lower bound of $\Omega(nr)$ for sparsifying a hypergraph on n vertices with maximum hyperedge-size r^3 . In fact, our lower bound applies even to the weaker notion of cut sparsification (where one only wants to approximate $Q_G(x)$ for all $x \in \{0, 1\}^V$), and is tight by the recent result of [CKN20], who gave a sampling-based cut sparsification algorithm that produces hypergraph sparsifiers with $O(n \log^{O(1)} n)$ hyperedge. In what follows we give an outline of our lower bound.

We start by formally defining the data structure for approximating the cut structure of a hypergraph that we prove a lower bound for. A *hypergraph cut sparsification scheme (HCSS)* is an algorithm for compressing the cut structure of a hypergraph such that queries on the size of cuts can be answered within a small multiplicative error:

³With some limits on the range of r . For more formal statements of our results see Section 7.2.

Definition 3.1. Let $\mathfrak{H}(n, r)$ be the set of hypergraphs on a vertex set $[n]$ with each hyperedge having size at most r . A pair of functions $\text{SPARSIFY} : \mathfrak{H}(n, r) \rightarrow \{0, 1\}^k$ and $\text{CUT} : \{0, 1\}^k \times 2^{[n]} \rightarrow \mathbb{N}$ is said to be an (n, r, k, ε) -HCSS if for all inputs $G = (V, E) \in \mathfrak{H}(n, r)$ the following holds.

- For every query $S \in 2^{[n]}$, $|\text{CUT}(\text{SPARSIFY}(G), S) - |E(S, \bar{S})|| \leq \varepsilon \cdot |E(S, \bar{S})|$.

To argue a lower bound on the space requirement (parameter k above), we use a reduction to string compression. It is known that $\{0, 1\}$ -strings of length ℓ cannot be significantly compressed to a small space data structure that allows even extremely crude additive approximations to subset sum queries — see, e.g., the LP decoding paper of [DN03] (here we only need a lower bound for computationally unbounded adversaries), or Section 7.1. We manage to encode a $\{0, 1\}$ -string of length ℓ into the cut structure of a hypergraph H with *fewer hyperedges than ℓ* — a testament to the higher complexity of hypergraph cut structures, as opposed to the cut structures of ordinary graphs.

Our string encoding construction utilizes Ruzsa-Szemerédi graphs. These are (ordinary) graphs whose edge-sets are the union of *induced* matchings. Our construction works generally on any Ruzsa-Szemerédi graphs and as a result we get several lower bounds in various parameter regimes (values of the hyperedge arity r and the precision parameter ε) based on the specific Ruzsa-Szemerédi graph constructions we choose to utilize. In particular, for the setting where $r = n^{O(1/\log \log n)}$ we are able to conclude that any hypergraph cut sparsification scheme requires $\Omega(rn)$ bits of space even for constant ε , matching the upper bound of [CKN20] to within logarithmic factors. For larger r we get a lower bound of $n^{1-o(1)}r$ bits of space for $\varepsilon = n^{-o(1)}$. The latter in particular rules out the possibility of an ε -sparsifier that can be described with asymptotically fewer than $(\varepsilon^{-1})^{O(1)}nr$ bits of space.

Here we briefly describe how we encode strings into hypergraphs generated from Ruzsa-Szemerédi graphs. Let G be a bipartite Ruzsa-Szemerédi-graph (with bipartition $P \cup Q$) composed of t induced matchings of size a each. We can then use the $a \cdot t$ edges of the graph to encode a string s of length $\ell = at$: simply order the edges of G and remove any edges corresponding to 0 coordinates in s , while keeping edges corresponding to 1's. This graph — which we call G_s — already encodes s when taken as a whole. However, its cut structure is not sufficient for decoding it. For that we need to turn G_s into a hypergraph H_s as follows: For each vertex u on one side of the bipartition, say P , we combine all edges adjacent on u into one hyperedge containing $\{u\} \cup \Gamma(u)$. This means that each hyperedge will have only a single vertex in P , but many vertices in Q (see Figure 2).

To decode the original string s from the cut structure of H , we must be able to answer subset sum queries $q \subseteq [at]$, that is return how many 1-coordinates s has, restricted to q . (For more details see the definition of string compression – Definition 7.1 in Section 7.1.) To do this, consider each induced matching one at a time and decode s restricted to the corresponding coordinates. We measure the size of a carefully chosen cut in H_s . Consider Figure 2: We restrict our view to a single matching M_j supported on P_j and Q_j in the two sides of the bipartition. Suppose for simplicity that q is entirely contained in this matching, and we are interested in the Hamming-weight of s restricted to a subset of coordinates q . To create our cut, in the top half of the hypergraph (P), we take the endpoints of edges corresponding to q – we call this set A . In the bottom half (Q), we take everything except for Q_j . The cut, which we call S , is depicted in red in Figure 2.

Informally, the crux of the decoding is the observation that the number of hyperedges crossing from A to Q_j is exactly the quantity we want to approximate. Indeed, consider a coordinate in q . If it has value 1 in s , the corresponding hyperedge crosses from A to Q_j , thus crossing the cut S . If however this coordinate is 0 in s , the corresponding hyperedge does not cross to Q_j , thus not crossing the cut. These types of hyperedges are denoted by 1 in Figure 2.




Figure 2: Illustration of the decoding process. One side of the cut S is depicted in orange.

Unfortunately, there are more hyperedges crossing S , adding noise to our measurement of s . One might hope to prove that the noise is small, i.e., can be attributed to measurement error, but this is not the case. Instead, we show that while this noise is not small, it is predictable enough to subtract accurately without knowing s . Hyperedges denoted 2 in Figure 2 cross from $P_j \setminus A$ to $Q \setminus Q_j$. Here we observe that nearly all hyperedges from $P_j \setminus A$ do in fact cross the cut, for almost all choices of s . Hyperedges denoted 3 in Figure 2 cross from $P \setminus P_j$ to $Q \setminus Q_j$. Here we cannot say much about the quantity of such hyperedges crossing the cut. However, we observe that this quantity does not depend on q , and therefore we can use Chernoff bounds (Theorem A.1) to prove that it concentrates around its expectation with high probability *over* s . This allows us to predict and subtract the noise caused by type 3 hyperedges, for whatever instance of Ruzsa-Szemerédi-graph we use (see the proof of Theorem 7.9).

Ultimately, we show that efficient cut sparsification for such hypergraphs would result in an equally efficient compression of $\{0, 1\}$ -strings, which implies our lower bounds. For more details see Section 7.

3.4 Directed Spectral Sparsification of Hypergraphs

In Section 8, we apply our discretization technique from Section 4 to the spectral sparsification of directed hypergraphs. As a testament to the versatility of this technique, we are able to produce an $O(n^2 r^3 \log^2 n / \epsilon^2)$ -sized ϵ -spectral sparsifier. This is a factor n better than the previous state of the art by [SY19], and nearly optimal in the setting where r is constant.

The broad arc of the proof is very similar to that of Section 4: We construct our sparsifier using importance sampling. We then divide the set of hyperarcs into a logarithmic number of categories, E_i . For each category separately, we show using discretization that the energy of the proposed sparsifier approximates the energy of the input hypergraph with respect to *all* $x \in \mathbb{R}^V$ simultaneously with high probability.

However, the details of each of these steps differ from their corresponding step in Section 4. Here we mention only a few key differences. Instead of looking at degrees or expansion, we define a novel quantity characterizing each hyperarc we call its *overlap*. Intuitively, this denotes the

highest density of an induced subgraph in which the particular hyperarc resides. We then sample each hyperarc with probability inverse proportional to its overlap. We show that this produces a sufficiently small sparsifier with high probability (see Lemma 8.2).

Perhaps the most crucial departure from Section 4 occurs during the discretization step when proving $Q_x(E_i) = \tilde{Q}_x(E_i)$. Instead of discretizing the vector $x \in \mathbb{R}^V$, we discretize the derived vector of energies on the hyperarcs, that is $Q_x \in \mathbb{R}^E$. So for each x and i we define a vector $Q_x^{(i)}$ — from a finite set of possibilities — such that, informally

$$Q_x(E_i) \cong Q_x^{(i)}(E_i) \cong \tilde{Q}_x^{(i)}(E_i) \cong \tilde{Q}_x(E_i).$$

For more details on the definition of $Q_x^{(i)}$, see the proof of Lemma 8.7. This additional trick is necessary; we do not know of a way to make the discretization argument work by rounding x itself.

For more details on the construction of directed hypergraph sparsifiers and their analysis see Section 8.

4 Spectral Sparsification of Expanders

In this section, we prove the following.

Theorem 4.1. *There is an algorithm that, given a parameter n , given $100/n \leq \epsilon \leq 1/2$ and an r -uniform hypergraph $G = (V, E)$ with $|V| \leq n$ and expansion at least $350\sqrt{(\log n)/(\epsilon rn)} \leq \Phi \leq 2/r$, outputs an ϵ -spectral sparsifier of G with $O(|V| \cdot (\frac{1}{\epsilon} \log n)^{O(1)}/(\Phi^2 r))$ hyperedges with probability $1 - O((\log n)/n^2)$ in $O(r|E|)$ time.*

Remark 4.2. *Note that n here does not denote the size of V but an arbitrary parameter larger than that. n serves only as an indirect error parameter, as the failure probability of the algorithm is allowed to be $1 - O((\log n)/n^2)$. The reason for this notation is that later on, in Section 6, we apply Theorem 4.1 to subgraphs of the input hypergraph. In this context, n will denote the size of the input hypergraph, whereas $|V|$ will denote the (potentially much smaller) size of the cluster within it, to be sparsified. Note that the size of the sparsifier scales linearly in the size of the cluster, but only logarithmically in the size of the input hypergraph. The latter is because the desired failure probability is always defined in terms of n .*

Remark 4.3. *The guarantee of Theorem 4.1 translates to $|\tilde{E}| = \tilde{O}(|V|r)$ when $\Phi(G) = \Omega(1/r)$, i.e., when G is a nearly-optimal expander.*

We show our construction of the sparsifier in Section 4.1 and discuss its correctness in Section 4.2, where some proofs are deferred to Section 4.3.

The following lemma is useful throughout this section.

Lemma 4.4. *For any hypergraph $G = (V, E)$, we have*

$$\sum_{e \in E} \frac{1}{\min_{v \in e} d(v)} \leq |V|.$$

Proof. Consider each hyperedge $e \in E$ to be directed towards its vertex with the lowest degree, i.e., $\operatorname{argmin}_{v \in e} d(v)$, breaking ties arbitrarily. Each vertex $v \in V$ has at most $d(v)$ incoming hyperedges, and each such hyperedge contributes to the above sum by $1/d(v)$. Hence the total contribution of all the incoming hyperedges to v is at most 1. It follows that the overall summation is at most $|V|$. \square

4.1 Construction

The construction of \tilde{G} is quite simple. Sample each hyperedge $e \in E$ with probability $p_e = \min\left(\frac{\lambda}{\min_{v \in e} d(v)}, 1\right)$ for

$$\lambda = (\epsilon^{-1} \log n)^{O(1)} / (\Phi^2 r). \quad (4)$$

Each sampled hyperedge e is given weight $w_e = 1/p_e$, and for every non-sampled hyperedge e define $w_e = 0$. Let \tilde{G} contain the sampled hyperedges, i.e., $\tilde{E} = \{e \in E \mid w_e > 0\}$. Notice that each random variable w_e has expectation $\mathbb{E}[w_e] = 1$, and thus informally $\mathbb{E}[\tilde{G}] = G$.

Clearly we can compute the output in time $O(r|E|)$. Also, we can bound the size of the sparsifier with high probability as follows.

Lemma 4.5. *We have*

$$\mathbb{P}[|\tilde{E}| \geq 2\lambda|V|] \leq O(1/n^2),$$

when $|\tilde{E}| = \Omega(\log n)$.

Proof. First, we have

$$\mathbb{E}[|\tilde{E}|] \leq \sum_{e \in E} p_e \leq \lambda \sum_{e \in E} \frac{1}{\min_{v \in e} d(v)} \leq \lambda|V|,$$

where the last inequality is due to Lemma 4.4. Noting that $|\tilde{E}|$ is a sum of independent indicator random variables, the claimed inequality is a direct consequence of the Chernoff bound (Theorem A.1). \square

4.2 Correctness

Let us now consider the spectral properties of \tilde{G} . We must prove that with high probability

$$\forall x \in \mathbb{R}^V, \quad \tilde{Q}(x) = (1 \pm \epsilon) \cdot Q(x). \quad (5)$$

We stress that this gives an error bound that holds for all x *simultaneously*. We may assume without loss of generality that $\sum_{v \in V} x_v d(v) = 0$ and $\sum_{v \in V} x_v^2 d(v) = 1$, because Equation (5) is invariant under translation and scaling of x . Let the set of such centered and normalized vectors be $\overline{\mathbb{R}^V}$. This guarantees that every non-isolated vertex v has $x_v^2 \leq 1/d(v) \leq 1$, and by Theorem 2.3 we get $Q(x) \geq \frac{r\Phi^2}{32}$.

Now fix one such vector $x \in \overline{\mathbb{R}^V}$, and use it to partition the hyperedge multiset E into $O(\log n)$ subsets as follows. For each $i = 1, \dots, i^*$, where $i^* = \lceil 2 \log n \rceil$, let

$$E_i = \left\{ e \in E \mid \max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v) \in (2^{-i}, 2^{-i+1}] \right\},$$

and let

$$E_* = E \setminus \bigcup_{i=1}^{i^*} E_i = \left\{ e \in E \mid \max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v) \leq 2^{-i^*} \right\}.$$

To justify the second equality in the equation above, note that $\sum_{v \in V} x_v^2 d(v) = 1$ implies $x(v)^2 \leq 1/d(v)$, and therefore for every $e \in E$

$$\max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v) \leq \max_{v \in e} 1/d(v) \cdot \min_{v \in e} d(v) = 1.$$

Informally, we would like to show that with high probability, for all x and all i we have $\tilde{Q}_x(E_i) \cong Q_x(E_i)$. Note that the multisets E_i and E_* are dependent on x , but we omit this from the notation for better readability. Our plan is to define another vector $x^{(i)} \in \mathbb{R}^V$ by rounding the coordinates of x , that preserves $Q(E_i)$ up to small multiplicative and additive error. Using this rounded vector, we will then show

$$\tilde{Q}_x(E_i) \cong \tilde{Q}_{x^{(i)}}(E_i) \cong Q_{x^{(i)}}(E_i) \cong Q_x(E_i),$$

and similarly also $\tilde{Q}_x(E_*) \cong Q_x(E_*)$.

Formally, for each $v \in V$ define $x_v^{(i)}$ as follows:

- If $x_v^2 d(v) \geq \epsilon^2 2^{-i} / 2500$, then round x_v to the nearest integer multiple of $1/(n^2 \sqrt{d(v)})$.
- If $x_v^2 d(v) < \epsilon^2 2^{-i} / 2500$, then round x_v to 0.

We implement the above plan using the following four claims.

First, we show that for every scale i the energy of E_i with respect to the rounded vector $x^{(i)}$ is quite close to the energy of E_i with respect to the original vector x :

Claim 4.6. *For all $x \in \overline{\mathbb{R}^V}$ and all $i = 1, \dots, i^*$,*

$$Q_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10}\right) Q_x(E_i) \pm \frac{20}{n}.$$

Next, we show that for every scale i our sampling process preserves energy of E_i on rounded version of all x simultaneously:

Claim 4.7. *For all $i = 1, \dots, i^*$,*

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10}\right) Q_{x^{(i)}}(E_i) \pm \frac{\epsilon Q(x)}{10 \log n} \right] \geq 1 - \frac{1}{n^2}.$$

We then relate the energy of the sampled E_i on rounded versions of x to the corresponding energy on original x :

Claim 4.8. *For all $i = 1, \dots, i^*$,*

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10}\right) \tilde{Q}_x(E_i) \pm \frac{60}{n} \right] \geq 1 - \frac{1}{n^2}.$$

Finally we bound the error introduced on the hyperedges of E_* .

Claim 4.9.

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_x(E_*) = Q_x(E_*) \pm \frac{12}{n} \right] \geq 1 - \frac{1}{n^2}.$$

Before proving these claims, which we do in the next section, let us show how to use them to show the correctness of the sparsifier.

Lemma 4.10. *The hypergraph \tilde{G} is an ϵ -spectral sparsifier of G with probability $1 - O((\log n)/n^2)$.*

Proof. Assume henceforth that the events in Claims 4.7, 4.8, and 4.9 all hold simultaneously for every i — we know that this happens with probability $1 - O(\log n/n^2)$ — and let us compare $Q_x(E_i)$ with $\tilde{Q}_x(E_i)$ for each i . If the above claims had no additive error, we could conclude that $\tilde{Q}_x(E_i) = (1 \pm 4\epsilon/10)Q_x(E_i)$. Similarly, if they had no multiplicative error, we could conclude that $|\tilde{Q}_x(E_i) - Q_x(E_i)| \leq \frac{80}{n} + \frac{\epsilon Q(x)}{10 \log n}$; we could then use the assumed lower bound on Φ to bound

$\frac{80}{n} \leq \frac{\epsilon Q(x)}{10 \log n}$, and sum up these additive errors over all $i = 1, \dots, i^*$ to a total that is bounded by $\frac{4}{10}\epsilon Q(x)$. These arguments extend easily also to E_* .

For the formal calculation, consider first one direction,

$$\begin{aligned}
\tilde{Q}_x(E_i) &\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{60}{n} + \tilde{Q}_{x^{(i)}}(E_i)\right] && \text{By Claim 4.8} \\
&\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{60}{n} + \frac{\epsilon Q(x)}{10 \log n} + \left(1 + \frac{\epsilon}{10}\right) Q_{x^{(i)}}(E_i)\right] && \text{By Claim 4.7} \\
&\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{60}{n} + \frac{\epsilon Q(x)}{10 \log n} + \left(1 + \frac{\epsilon}{10}\right) \left[\frac{20}{n} + \left(1 + \frac{\epsilon}{10}\right) Q_x(E_i)\right]\right] && \text{By Claim 4.6} \\
&\leq \frac{120}{n} + \frac{2\epsilon Q(x)}{10 \log n} + \left(1 + \frac{4\epsilon}{10}\right) Q_x(E_i),
\end{aligned}$$

since $\epsilon \leq 1/2$. Now sum this over all i and combine it with the bound from Claim 4.9 on the error introduced by E_* , to get

$$\begin{aligned}
\tilde{Q}(x) &= \tilde{Q}_x(E_*) + \sum_{i=1}^{i^*} \tilde{Q}_x(E_i) \\
&\leq \left[Q_x(E_*) + \frac{12}{n}\right] + i^* \left[\frac{120}{n} + \frac{2\epsilon Q(x)}{10 \log n}\right] + \left(1 + \frac{4\epsilon}{10}\right) \sum_{i=1}^{i^*} Q_x(E_i) \\
&\leq \frac{250 \log n}{n} + \frac{5\epsilon Q(x)}{10} + \left(1 + \frac{4\epsilon}{10}\right) Q(x) \\
&\leq (1 + \epsilon)Q(x).
\end{aligned}$$

The last inequality follows from $Q(x) \geq \Phi^2 r / 32$ (by Theorem 2.3), and the theorem's assumption that $\Phi \geq 350\sqrt{(\log n)/(\epsilon r n)}$.

The other direction, $\tilde{Q}(x) \geq (1 - \epsilon)Q(x)$, follows similarly. \square

Theorem 4.1 then follows by Lemmas 4.5 and 4.10 and a union bound.

4.3 Proofs of Claims 4.6, 4.7, 4.8, and 4.9

We begin by presenting a preliminary lemma about the effects of approximating x on a general quadratic form, which will be useful in proving the four claims, and will be useful later on in Section 6.

Lemma 4.11. *Let $G = (V, E)$ be a hypergraph and let x, \tilde{x} be two vectors in \mathbb{R}^V such that $|x_v - \tilde{x}_v| \leq \delta$ on every coordinate $v \in V$ for some $\delta \geq 0$. Then for any $e \in E$,*

$$|Q_x(e) - Q_{\tilde{x}}(e)| \leq 4\delta \left(\sqrt{Q_x(e)} + \delta\right).$$

Proof. Given $e \in E$, we begin by finding two vertices $u^*, v^* \in e$ such that

$$|Q_x(e) - Q_{\tilde{x}}(e)| \leq |(x_{u^*} - x_{v^*})^2 - (\tilde{x}_{u^*} - \tilde{x}_{v^*})^2|.$$

It is indeed possible to find such vertices. If $Q_x(e) \geq Q_{\tilde{x}}(e)$, set u^*, v^* such that $Q_x(e) = (x_{u^*} - x_{v^*})^2$, and we get

$$|Q_x(e) - Q_{\tilde{x}}(e)| = Q_x(e) - Q_{\tilde{x}}(e) \leq (x_{u^*} - x_{v^*})^2 - (\tilde{x}_{u^*} - \tilde{x}_{v^*})^2,$$

since $Q_{\tilde{x}}(e) \geq (\tilde{x}_u - \tilde{x}_v)^2$ for every $u, v \in e$, and in particular $Q_{\tilde{x}}(e) \geq (\tilde{x}_{u^*} - \tilde{x}_{v^*})^2$. Otherwise, i.e., $Q_{\tilde{x}} > Q_x(e)$, similarly set u^*, v^* such that $Q_{\tilde{x}}(e) = (\tilde{x}_{u^*} - \tilde{x}_{v^*})^2$.

Using these u^*, v^* , we have

$$\begin{aligned} |Q_x(e) - Q_{\tilde{x}}(e)| &\leq |(x_{u^*} - x_{v^*})^2 - (\tilde{x}_{u^*} - \tilde{x}_{v^*})^2| \\ &= |x_{u^*} + \tilde{x}_{u^*} - x_{v^*} - \tilde{x}_{v^*}| \cdot |x_{u^*} - \tilde{x}_{u^*} - x_{v^*} + \tilde{x}_{v^*}| \end{aligned}$$

Let us now bound each of these two factors. The second one is clearly bounded by 2δ by the lemma's assumption. To bound the first term, we use that $Q_x(e) = \max_{u,v \in e} (x_u - x_v)^2 \geq (x_{u^*} - x_{v^*})^2$, and therefore

$$|x_{u^*} + \tilde{x}_{u^*} - x_{v^*} - \tilde{x}_{v^*}| \leq 2|x_{u^*} - x_{v^*}| + |x_{u^*} - \tilde{x}_{u^*}| + |x_{v^*} - \tilde{x}_{v^*}| \leq 2\sqrt{Q_x(e)} + 2\delta.$$

Putting these two bounds together, we obtain the result of the lemma. \square

The following claim examines the effects of rounding from x to $x^{(i)}$ (from the previous section) on a single hyperedge of E_i . This is the main technical claim that allows us to then easily prove both Claims 4.6 and 4.8.

Claim 4.12. *For all $x \in \mathbb{R}^V$, all $i = 1, \dots, i^*$, and every hyperedge $e \in E_i$,*

$$Q_{x^{(i)}}(e) = \left(1 \pm \frac{\epsilon}{10}\right) Q_x(e) \pm \frac{20}{n^2 \min_{v \in e} d(v)}.$$

Proof. We examine the difference $|Q_x(e) - Q_{x^{(i)}}(e)|$. Recall that $x_v^{(i)}$ is either rounded to the nearest multiple of $1/(n^2 \sqrt{d(v)})$ or rounded to zero. We consider two cases:

1. No vertex in e is rounded to zero.
2. At least one vertex in e is rounded to zero.

For simplicity, denote $x_+ = \max_{v \in e} |x_v|$ and $d_- = \min_{v \in e} d(v)$. Recall that by definition of E_i ,

$$x_+^2 d_- \in (2^{-i}, 2^{-i+1}]. \quad (6)$$

In the first case, the value of x on every vertex $v \in e$ changes by at most $1/(n^2 \sqrt{d(v)}) \leq 1/(n^2 \sqrt{d_-})$. Thus, we can apply Lemma 4.11 with $\delta = 1/(n^2 \sqrt{d_-})$ to get

$$|Q_x(e) - Q_{x^{(i)}}(e)| \leq \frac{4}{n^2 \sqrt{d_-}} \cdot \left(\sqrt{Q_x(e)} + \frac{1}{n^2 \sqrt{d_-}} \right).$$

We can use (6) to bound $Q_x(e) \leq 4x_+^2 \leq 4 \cdot 2^{1-i}/d_- \leq 4/d_-$. Substituting this in, we get

$$|Q_x(e) - Q_{x^{(i)}}(e)| \leq \frac{4}{n^2 \sqrt{d_-}} \cdot \left(\frac{2}{\sqrt{d_-}} + \frac{1}{n^2 \sqrt{d_-}} \right) \leq \frac{20}{n^2 d_-}.$$

In the second case, the value of x on a vertex in e can still change by at most $1/(n^2 \sqrt{d_-})$ by rounding to a non-zero value. It can additionally be rounded to a zero, as long as $x_v^2 d(v) < \epsilon^2 2^{-i}/2500$, which amounts to additive error per coordinate of at most $|x_v| < \epsilon/\sqrt{2500 \cdot 2^i d(v)} \leq$

$\epsilon/\sqrt{2500 \cdot 2^i d_-}$. Therefore we can apply Lemma 4.11 with $\delta = \epsilon/\sqrt{2500 \cdot 2^i d_-} \geq 1/(n^2 \sqrt{d_-})$, which gives us that

$$|Q_x(e) - Q_{x^{(i)}}(e)| \leq \frac{4\epsilon}{\sqrt{2500 \cdot 2^i d_-}} \cdot \left(\sqrt{Q_x(e)} + \frac{\epsilon}{\sqrt{2500 \cdot 2^i d_-}} \right).$$

This time we use a lower bound on $Q_x(e)$. Recall that we assumed that at least one vertex in e is rounded to zero. Let one such vertex be v_0 . This means that $x_{v_0}^2 d(v_0) \leq \epsilon^2 2^{-i}/2500$, but at the same time $x_+^2 d(v_0) \geq x_+^2 d_- \geq 2^{-i}$. Using these two facts, we get our lower bound

$$\sqrt{Q_x(e)} \geq x_+ - |x_{v_0}| \geq x_+ - \epsilon x_+/50 \geq \frac{49}{50 \sqrt{2^i d_-}}.$$

Substituting this in, we get

$$|Q_x(e) - Q_{x^{(i)}}(e)| \leq \frac{4\epsilon \sqrt{Q_x(e)}}{49} \cdot \left(\sqrt{Q_x(e)} + \frac{\epsilon \sqrt{Q_x(e)}}{49} \right) \leq \frac{4\epsilon Q_x(e)}{49} \cdot \left(1 + \frac{\epsilon}{49} \right) \leq \frac{\epsilon}{10} Q_x(e).$$

In conclusion, in the first case we get the claimed additive error, while in the second case we get the claimed multiplicative error. \square

We are now ready to proceed to proving Claims 4.6 and 4.8.

Claim 4.6. For all $x \in \overline{\mathbb{R}^V}$ and all $i = 1, \dots, i^*$,

$$Q_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10} \right) Q_x(E_i) \pm \frac{20}{n}.$$

Proof. We can bound

$$\begin{aligned} |Q_x(E_i) - Q_{x^{(i)}}(E_i)| &\leq \sum_{e \in E_i} |Q_x(e) - Q_{x^{(i)}}(e)| \\ &\leq \sum_{e \in E_i} \left[\frac{\epsilon}{10} Q_x(e) + \frac{20}{n^2 \min_{v \in e} d(v)} \right] && \text{by Claim 4.12} \\ &\leq \frac{\epsilon}{10} Q_x(E_i) + \frac{20}{n^2} \sum_{e \in E} \frac{1}{\min_{v \in e} d(v)} \\ &\leq \frac{\epsilon}{10} Q_x(E_i) + \frac{20}{n} && \text{by Lemma 4.4 and } n \geq |V|, \end{aligned}$$

as claimed. \square

Claim 4.8. For all $i = 1, \dots, i^*$,

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10} \right) \tilde{Q}_x(E_i) \pm \frac{60}{n} \right] \geq 1 - \frac{1}{n^2}.$$

Proof. Similarly to the previous proof, we first bound

$$|\tilde{Q}_x(E_i) - \tilde{Q}_{x^{(i)}}(E_i)| \leq \sum_{e \in E_i} |\tilde{Q}_x(e) - \tilde{Q}_{x^{(i)}}(e)|.$$

Recall that $\tilde{Q}_x(e) = w_e Q_x(e)$ where w_e is a random variable (independent from all others) that takes value $1/p_e$ with probability p_e , and value 0 otherwise. Similarly, $\tilde{Q}_{x^{(i)}}(e) = w_e \tilde{Q}_{x^{(i)}}(e)$. Applying this along with Claim 4.12, we get

$$\begin{aligned} |\tilde{Q}_x(E_i) - \tilde{Q}_{x^{(i)}}(E_i)| &\leq \sum_{e \in E_i} w_e |Q_x(e) - Q_{x^{(i)}}(e)| \\ &\leq \sum_{e \in E_i} \left[\frac{\epsilon}{10} w_e Q_x(e) + w_e \cdot \frac{20}{n^2 \min_{c \in e} d(v)} \right] \\ &= \frac{\epsilon}{10} \tilde{Q}_x(e) + \sum_{e \in E_i} \frac{20 w_e}{n^2 \min_{v \in e} d(v)}. \end{aligned}$$

Note that in the sum here the term corresponding to e is zero unless e is sampled to \tilde{E} , in which case $w_e = 1/p_e \leq 1 + \min_{v \in e} d(v)/\lambda$. (Recall λ from equation 4.) Using also Lemmas 4.4 and 4.5, and the fact that $n \geq |V|$, we have that with high probability

$$\sum_{e \in E_i} \frac{20 w_e}{n^2 \min_{v \in e} d(v)} \leq \sum_{e \in \tilde{E}_i} \frac{20}{n^2 \min_{v \in e} d(v)} + \sum_{e \in \tilde{E}} \frac{20}{\lambda n^2} \leq \sum_{e \in E_i} \frac{20}{n^2 \min_{v \in e} d(v)} + |\tilde{E}| \cdot \frac{20}{\lambda n^2} \leq \frac{60}{n}. \quad \square$$

Claim 4.9.

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_x(E_*) = Q_x(E_*) \pm \frac{12}{n} \right] \geq 1 - \frac{1}{n^2}.$$

Proof. Note that

$$|\tilde{Q}_x(E_*) - Q_x(E_*)| \leq \tilde{Q}_x(E_*) + Q_x(E_*) = \sum_{e \in E_*} \max_{u, v \in e} (x_u - x_v)^2 + \sum_{e \in E_*} w_e \cdot \max_{u, v \in e} (x_u - x_v)^2.$$

Now, we bound each term using that $\max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v) \leq 1/n^2$ by definition of E_* . For the first term, we use Lemma 4.4,

$$\sum_{e \in E_*} \max_{u, v \in e} (x_u - x_v)^2 \leq 4 \sum_{e \in E_*} \max_{v \in e} x_v^2 \leq 4 \sum_{e \in E_*} \frac{1}{n^2 \min_{v \in e} d(v)} \leq \frac{4}{n}.$$

For the second term, we have by Lemma 4.5, and the fact that $n \geq |V|$, that with high probability,

$$\sum_{e \in E_*} w_e \cdot \max_{u, v \in e} (x_u - x_v)^2 \leq 4 \sum_{e \in E_*} w_e \cdot \max_{v \in e} x_v^2 \leq 4 \sum_{e \in E_*} \frac{w_e}{n^2 \min_{v \in e} d(v)} \leq 4 |\tilde{E}| \cdot \frac{1}{\lambda n^2} \leq \frac{8}{n}. \quad \square$$

Finally, we prove the technical crux of the theorem, Claim 4.7.

Claim 4.7. For all $i = 1, \dots, i^*$,

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10}\right) Q_{x^{(i)}}(E_i) \pm \frac{\epsilon Q(x)}{10 \log n} \right] \geq 1 - \frac{1}{n^2}.$$

Proof. We shall prove the stronger claim

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V}, \tilde{Q}_{x^{(i)}}(E_i) = \left(1 \pm \frac{\epsilon}{10}\right) Q_{x^{(i)}}(E_i) \pm \frac{\epsilon r \Phi^2 / 32}{10 \log n} \right] \geq 1 - \frac{1}{n^2}.$$

This is indeed stronger, since for all $x \in \overline{\mathbb{R}^V}$, we know that $Q(x) \geq r \Phi^2 / 32$ by the Hypergraph Cheeger inequality (Theorem 2.3). This allows us to argue that the probabilistic claim depends

on x only through $x^{(i)}$ and E_i . These are discrete which will allow for the use of union bound later on. We will first prove a deviation bound for a single instance of $(x^{(i)}, E_i)$ using an additive-multiplicative Chernoff bound, and then extend it to hold for all instances simultaneously using a union bound.

Fix i , $x^{(i)}$, and E_i . Notice that $\tilde{Q}_{x^{(i)}}(E_i) = \sum_{e \in E_i} w_e \cdot \max_{a,b \in e} (x_a^{(i)} - x_b^{(i)})^2$ is a sum of independent random variables whose expectation is $Q_{x^{(i)}}(E_i)$. Let us bound the maximum range of one summand, for some $e \in E_i$. If $p_e = 1$ the range is 0, and otherwise the range is bounded by

$$w_e \cdot \max_{a,b \in e} (x_a^{(i)} - x_b^{(i)})^2 \leq \max_{a,b \in e} \frac{2(x_a^2 + x_b^2)}{p_e} \leq \frac{4}{\lambda} \max_{v \in e} x_v^2 \cdot \min_{v \in e} d(v) \leq \frac{2^{-i+3}}{\lambda}.$$

We can thus apply Theorem A.2 and get

$$\begin{aligned} \mathbb{P} \left[|\tilde{Q}_{x^{(i)}}(E_i) - Q_{x^{(i)}}(E_i)| \geq \frac{\epsilon}{10} Q_{x^{(i)}}(E_i) + \frac{\epsilon r \Phi^2 / 32}{10 \log n} \right] &\leq 2 \exp \left(-\frac{\epsilon / 10 \cdot (\epsilon r \Phi^2 / 32) / (10 \log n)}{3 \cdot 2^{-i+3} / \lambda} \right) \\ &= 2 \exp \left(-\frac{\lambda 2^i \epsilon^2 r \Phi^2}{32 \cdot 2400 \log n} \right). \end{aligned}$$

Now we wish to extend this high-probability bound to hold simultaneously for all possible $x^{(i)}$ and E_i . How many possible settings of $(x^{(i)}, E_i)$ are there? Each non-zero coordinate v of $x^{(i)}$ has $x_v^2 d(v) \geq \epsilon^2 2^{-i} / 2500$, so there are at most $2500 \cdot 2^i / \epsilon^2$ such coordinates. Furthermore, each such coordinate $x_v^{(i)}$ is an integer multiple of $1 / (n^2 \sqrt{d(v)})$ in the range $[-1 / \sqrt{d(v)}, 1 / \sqrt{d(v)}]$, so there are only $2n^2$ possibilities per non-zero coordinate. Thus, the total number of vectors $x^{(i)}$ is at most

$$\binom{|V|}{2500 \cdot 2^i / \epsilon^2} \cdot (2n^2)^{2500 \cdot 2^i / \epsilon^2} \leq (2n^3)^{2500 \cdot 2^i / \epsilon^2}$$

We still need to enumerate the number of possible hyperedge multisets E_i given $x^{(i)}$. To know whether some hyperedge $e \in E$ is in E_i , we must know whether the value of $\max_{v \in e} x_v^2 \min_{v \in e} d(v)$ is in $(2^{-i}, 2^{-i+1}]$. Unfortunately, this depends on $\max_{v \in e} x_v^2$, which is not determined by $x^{(i)}$, due to the rounding error between x and $x^{(i)}$. Let $D = \{d(v) \mid v \in V\}$ be the set of all degrees in G . It suffices to know for each v corresponding to a non-zero coordinate of $x^{(i)}$ the two values

$$\min\{d \in D \mid x_v^2 d > 2^{-i}\} \quad \text{and} \quad \max\{d \in D \mid x_v^2 d \leq 2^{-i+1}\}.$$

Indeed, we need not worry about zero coordinates of $x^{(i)}$, i.e., vertices v with $x_v^2 d(v) < \epsilon^2 2^{-i} / 2500$, as these cannot attain $\max_{u \in e} x_u^2$ for a hyperedge $e \in E_i$. Thus, the total number of possible multisets E_i given $x^{(i)}$ is at most $(|V|^2)^{2500 \cdot 2^i / \epsilon^2} \leq (n^2)^{2500 \cdot 2^i / \epsilon^2}$.

We are now ready to apply a union bound,

$$\begin{aligned} \mathbb{P} \left[\forall x, |\tilde{Q}_{x^{(i)}}(E_i) - Q_{x^{(i)}}(E_i)| \leq \frac{\epsilon}{10} Q_{x^{(i)}}(E_i) + \frac{\epsilon(r\Phi^2/32)}{10 \log n} \right] &\leq (n^2 \cdot (2n^3))^{2500 \cdot 2^i / \epsilon^2} \cdot 2 \exp \left(-\frac{\lambda 2^i \epsilon^2 r \Phi^2}{32 \cdot 2400 \log n} \right) \\ &\leq 2 \exp \left(\frac{15000 \cdot 2^i \log n}{\epsilon^2} - \frac{\lambda 2^i \epsilon^2 r \Phi^2}{32 \cdot 2400 \log n} \right) \leq \frac{1}{n^2}, \end{aligned}$$

where the last inequality holds as long as $\lambda \geq 24 \cdot 10^8 \cdot \log^2 n / (\epsilon^4 \Phi^2 r)$, which is indeed how we set λ . \square

5 Expander Decomposition

This section provides a procedure to decompose an input hypergraph into expanders while cutting a small number of hyperedges. This is stated in the following lemma, which we fully prove for completeness, as we cannot find a useful reference for it. It is based on the standard technique of iteratively removing a sparse cut, with slight adaptations like a minimum-degree guarantee in each expander, and an approximation algorithm for sparsest cut (minimal expansion) in a hypergraph.

Lemma 5.1 (Expander Decomposition). *There exists a polynomial-time algorithm that, given an r -uniform hypergraph $G = (V, E)$ with n vertices and m hyperedges, outputs disjoint vertex subsets $C_1, \dots, C_k \subseteq V$ (not necessarily a partition) that satisfy*

1. $\Phi(G[C_j]) = \Omega(1/(r \log^2 n))$ for all $j = 1, \dots, k$;
2. $\min_{v \in C_j} d_{G[C_j]}(v) \geq m/(4n)$ for all $j = 1, \dots, k$; and
3. $|E \setminus \bigcup_{j=1}^k E(G[C_j])| \leq m/2$.

We first provide (in Section 5.1) an approximation algorithm for the sparsest cut problem on hypergraphs by slightly modifying a known algorithm from [FHL08]. We then use it (in Section 5.2) to prove Lemma 5.1, where we decompose an input hypergraph into expanders by iteratively deleting sparse cuts and low-degree vertices.

5.1 Approximating Sparsest Cut

Lemma 5.2. *There exists a polynomial-time algorithm that, given a hypergraph $G = (V, E)$, computes a cut $S \subseteq V$ with $\Phi(S) = O(\log n \cdot \Phi(G))$.*

We remark that a better approximation ratio $O(\sqrt{\log n})$ was shown by Louis and Makarychev [LM16]. Strictly speaking, their definition of expansion is slightly different, using $|S|$ rather than $\text{vol}(S)$, but their results probably extend also to our setting.

Our algorithm is an extension of known approximation algorithms for the sparsest cut problem on ordinary graphs. Specifically, we follow the rounding procedure techniques of Feige, Hajiaghayi and Lee [FHL08], which in turn build on the linear programming relaxation approach introduced by Leighton and Rao [LR99] and on the rounding procedures based on metric embeddings devised by Aumann and Rabani [AR98] and Linial, London and Rabinovich [LLR95].

For a hypergraph $G = (V, E)$ and a vertex set $S \subseteq V$, we define

$$\phi(S) = \frac{|E(S, V \setminus S)|}{\text{vol}(S) \cdot \text{vol}(V \setminus S)}$$

(notice the difference in the denominator from $\Phi(S)$), and let $\phi(G) = \min_{S \subseteq V} \phi(S)$. We note that

$$\frac{\text{vol}(S) \cdot \text{vol}(V \setminus S)}{\text{vol}(G)} \leq \min\{\text{vol}(S), \text{vol}(V \setminus S)\} \leq \frac{2 \cdot \text{vol}(S) \cdot \text{vol}(V \setminus S)}{\text{vol}(G)},$$

and hence a $\rho(n)$ -approximation algorithm for $\phi(G)$ immediately gives a $2\rho(n)$ -approximation for $\Phi(G)$.

For a vertex set $S \subseteq V$, let $1_S: V \rightarrow \{0, 1\}$ be the indicator function of S . Then, we have

$$\phi(S) = \frac{\sum_{e \in E} \max_{u, v \in e} |1_S(u) - 1_S(v)|}{\sum_{u, v \in V} d(u)d(v)|1_S(u) - 1_S(v)|}.$$

By relaxing this optimization over cut metrics (distances induced by indicator functions) to optimization over all pseudo-metrics, we obtain the following linear program (LP).

$$\begin{aligned}
& \text{minimize} && \sum_{e \in E} z(e), \\
& \text{subject to} && \sum_{u, v \in V} d(u)d(v)\ell(u, v) = 1, \\
& && z(e) \geq \ell(u, v) \quad \forall e \in E, u, v \in e, \\
& && \ell(u, w) \leq \ell(u, v) + \ell(v, w) \quad \forall u, v, w \in V, \\
& && \ell(u, v) = \ell(v, u) \geq 0 \quad \forall u, v \in V.
\end{aligned}$$

This LP has variables $\{\ell(u, v)\}_{u, v \in V}$ and $\{z(e)\}_{e \in E}$, and its size is $O(n^2 + mr^2)$. Our algorithm solves this LP and then rounds the solution as explained next.

Our rounding procedure is similar to [FHL08], who designed an approximation algorithm for the sparsest vertex-cut problem. They use the following embedding result due to Bourgain.

Theorem 5.3 (Bourgain's embedding [Bou85]). *Let $\ell : V \times V \rightarrow \mathbb{R}$ be a pseudo-metric on an n -point set V . Then there exists an embedding $f : V \rightarrow \mathbb{R}^k$ with distortion $D = O(\log n)$ in the following sense:*

$$\forall u, v \in V, \quad \max_{i \in [k]} |f_i(u) - f_i(v)| \leq \ell(u, v), \quad (7)$$

$$\forall u, v \in V, \quad \frac{1}{k} \sum_{i \in [k]} |f_i(u) - f_i(v)| \geq \frac{\ell(u, v)}{D}. \quad (8)$$

Linial, London and Rabinovich [LLR95] showed how to compute such an embedding with high probability in time $\tilde{O}(n^2)$.

Remark 5.4. *The above form of Bourgain's embedding is slightly stronger than the usual statement of embedding into ℓ_1 (see also [MR01]). Indeed, the usual statement follows easily by viewing f as an embedding into ℓ_1 , then*

$$\forall u, v \in V, \quad \frac{\ell(u, v)}{D} \leq \frac{1}{k} \|f(u) - f(v)\|_1 \leq \ell(u, v),$$

which means that scaling f by factor $1/k$ achieves distortion D .

Lemma 5.5. *Given an embedding $f : V \rightarrow \mathbb{R}^k$, one can find in polynomial time a cut $S^* \subseteq V$ such that*

$$\phi(S^*) \leq \min_{i \in [k]} \frac{\sum_{e \in E} \max_{u, v \in e} |f_i(u) - f_i(v)|}{\sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|}.$$

Proof. Let $i \in [k]$ be the index that minimizes the above ratio, and define the embedding $g : V \rightarrow \mathbb{R}$ by scaling and translating the corresponding f_i such that $\min_v g(v) = 0$ and $\max_v g(v) = 1$.

Pick uniformly at random a threshold $s \in [0, 1]$, and consider the (random) set $S = \{v \in V \mid g(v) > s\}$. Then by simple calculations

$$\begin{aligned}
\forall u, v \in V, & \quad \mathbb{E}|1_S(u) - 1_S(v)| = |g(u) - g(v)|, \\
\forall e \in E, & \quad \mathbb{E} \max_{u, v \in e} |1_S(u) - 1_S(v)| = \max_{u \in e} g(u) - \min_{v \in e} g(v).
\end{aligned}$$

It follows that there must exist $S^* = \{v \in V \mid g(v) > s^*\}$ that is non-trivial (i.e., $S^* \neq \emptyset, V$) for which

$$\begin{aligned} \phi(S^*) &= \frac{\sum_{e \in E} \max_{u, v \in e} |1_{S^*}(u) - 1_{S^*}(v)|}{\sum_{u, v \in V} d(u)d(v) |1_{S^*}(u) - 1_{S^*}(v)|} \leq \frac{\mathbb{E} \sum_{e \in E} \max_{u, v \in e} |1_S(u) - 1_S(v)|}{\mathbb{E} \sum_{u, v \in V} d(u)d(v) |1_S(u) - 1_S(v)|} \\ &= \frac{\sum_{e \in E} \max_{u, v \in e} |g(u) - g(v)|}{\sum_{u, v \in V} d(u)d(v) |g(u) - g(v)|} = \frac{\sum_{e \in E} \max_{u, v \in e} |f_i(u) - f_i(v)|}{\sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|}. \end{aligned}$$

A polynomial-time implementation can simply take the best sweep cut. \square

Proof of Lemma 5.2. The algorithm computes an optimal LP solution and then applies to it Bourgain's embedding (Theorem 5.3) and then Lemma 5.5 to find a cut $S^* \subset V$. This is a (randomized) polynomial-time algorithm, and with high probability its output $S^* \subseteq V$ satisfies

$$\begin{aligned} \phi(S^*) &\leq \min_{i \in [k]} \frac{\sum_{e \in E} \max_{u, v \in e} |f_i(u) - f_i(v)|}{\sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|} \leq \min_{i \in [k]} \frac{\sum_{e \in E} \max_{u, v \in e} \ell(u, v)}{\sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|} \\ &\leq \frac{\sum_{e \in E} z(e)}{\max_{i \in [k]} \sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|} \leq \frac{\sum_{e \in E} z(e)}{\mathbb{E}_{i \in [k]} \sum_{u, v \in V} d(u)d(v) |f_i(u) - f_i(v)|} \\ &\leq \frac{\sum_{e \in E} z(e)}{\sum_{u, v \in V} d(u)d(v) \ell(u, v) / D} = D \cdot \text{LP} \leq O(\log n) \cdot \phi(G). \end{aligned}$$

We thus conclude an $O(\log n)$ -approximation algorithm for $\phi(G)$. \square

5.2 Proof of Lemma 5.1

We now prove Lemma 5.1. We shall refer to the algorithm given in Lemma 5.2 as SPARSECUT.

Algorithm 1

```

1: procedure EXPANDERDECOMPOSITION( $G = (V, E)$ )
2:    $\mathcal{C} \leftarrow \{V\}$ 
3:   while true do
4:     if  $\exists C \in \mathcal{C}$  and  $v \in C$  such that  $d_{G[C]}(v) < m/(4n)$  then
5:       remove  $v$  from  $C$ .
6:     else if  $\exists C \in \mathcal{C}$  s.t. SPARSECUT( $G[C]$ ) finds a cut  $S$  with  $\Phi_{G[C]}(S) \leq 1/(4r \log n)$  then
7:        $\mathcal{C} \leftarrow (\mathcal{C} \setminus \{C\}) \cup \{S, C \setminus S\}$ 
8:     else
9:       return  $\mathcal{C}$ 

```

Proof of Lemma 5.1. Our algorithm is given in Algorithm 1. By Lemma 5.2, there exists a constant $K > 0$ such that, if there is a cut of expansion at most $1/(Kr \log^2 n)$, then SPARSECUT finds a cut of expansion at most $1/(4r \log n)$, and hence Guarantee 1 holds. Guarantee 2 also holds, since the algorithm only returns when when no vertices violate the min-degree condition, due to Line 4.

We now show Guarantee 3, which claims that at most half the hyperedges are omitted from this clustering. First, we bound the number of hyperedges cut by splitting C on Line 7. To this end, we charge some weight to vertices for each cut $(S, C \setminus S)$ found throughout the algorithm. Specifically, suppose $|S| \leq |C \setminus S|$ — then we charge cut hyperedges to vertices in S *proportionally to their degree*: Each vertex $v \in S$ is charged

$$E(S, C \setminus S) \cdot \frac{d_{G[C]}(v)}{\text{vol}_{G[C]}(S)} \leq d_{G[C]}(v) \cdot \frac{E(S, C \setminus S)}{\min(\text{vol}_{G[C]}(S), \text{vol}_{G[C]}(C \setminus S))} \leq \frac{d_{G[C]}(v)}{4r \log n} \leq \frac{d(v)}{4r \log n}.$$

If $|C \setminus S| \leq |S|$, vertices of $C \setminus S$ get charged similarly. Since each vertex is charged only when the size of its containing component is decreased by at least a factor of 2, each vertex can get charged a maximum of $\log n$ times. This means that in total v is charged at most $d(v)/(4r)$ cut hyperedges.

Hence, the total number of hyperedges cut by splitting components is bounded by

$$\sum_{v \in V} \frac{d(v)}{4r} = \frac{m}{4}.$$

The total number of hyperedges cut by discarding low-degree vertices in Line 5 is at most

$$\frac{m}{4n} \cdot n = \frac{m}{4}.$$

Altogether, the total number of hyperedges cut by the algorithm is at most $m/4 + m/4 = m/2$. \square

6 General Spectral Sparsification of Hypergraphs

In this section, we prove Theorem 1.1. We describe our construction in Section 6.1 and then prove its correctness in Section 6.2.

6.1 Construction

We shall call a hyperedge a *self-loop* if all of its vertices are identical (i.e., the number of distinct vertices in it is one). We explicitly prohibit self-loops in the next lemma for a technical reason inside the proof of Theorem 1.1.

Lemma 6.1. *There is an algorithm that, given a parameter n , given $100/n \leq \epsilon \leq 1/2$ and an r -uniform hypergraph $G = (V, E)$ with $|V| \leq n$ and expansion $\Phi(G) \geq \Omega(\frac{1}{r \log^2 n})$ and $r \leq \epsilon n / \log^6 n$, outputs an ϵ -spectral sparsifier of G with at most $V|r(\epsilon^{-1} \log n)^{O(1)}$ hyperedges and no self-loops with probability at least $1 - O((\log n)/n^2)$ in $O(r|E|)$ time. Let $\text{EXPANDERSPARSIFY}(G, \epsilon)$ be such an algorithm.*

Proof. This is a simple corollary of Theorem 4.1 when applied with $\Phi = \Omega(\frac{1}{r \log^2 n})$, the only difference being the exclusion of self-loops. One can simply remove all the self-loops from the sparsifier obtained from Theorem 4.1, as this does not change the spectral properties at all. \square

Definition 6.2. *Let \sim be a relation on a ground set V . The transitive closure \approx of \sim is a relation on V defined as follows*

$$a \approx b \iff \exists c_1, \dots, c_k \in V, a \sim c_1 \sim \dots \sim c_k \sim b.$$

Definition 6.3. *Let $G = (V, E)$ be a hypergraph and let \approx be an equivalence relation on V . Then the contraction of G with respect to \approx , denoted by G/\approx , is a hypergraph on the vertex set $V' = V/\approx$ and the multiset of hyperedges*

$$E' = \{ \{ [v] \mid v \in e \} \mid e \in E \},$$

where $[v]$ denotes the equivalence class of v with respect to \approx . We stress that each hyperedge in the contracted hypergraph is itself a multiset, i.e., for each $e \in E$ there is a corresponding hyperedge in E' of total multiplicity $|e|$. Moreover, E' is a multiset and thus may have copies of the same hyperedge (i.e., parallel hyperedges).

Remark 6.4. Since there is a bijection between E and E' , we shall slightly abuse notation and equate the hyperedges of G and of G/\approx (i.e., use one as a shorthand for the other). This occurs already in the next observation.

Observation 6.5. Let $G = (V, E)$ be a hypergraph and let \approx be an equivalence relation on V . Let \tilde{G} be an ϵ -spectral sparsifier of the contracted hypergraph G/\approx (not of G). Then it need not be true that the energy of G is always approximated by the energy of \tilde{G} in the sense that for all $x \in \mathbb{R}^V$,

$$\tilde{Q}(x) = (1 \pm \epsilon)Q(x).$$

Here \tilde{Q} is the energy of the sparsifier \tilde{G} when interpreted as a subgraph of G , not G/\approx . That is, we take the weighted hyperedges found in \tilde{G} and interpret them as hyperedges over the vertex set of G (see Remark 6.4).

However, the equation above does hold if x is constant on each equivalence class of \approx , that is whenever

$$u \approx v \Rightarrow x_u = x_v.$$

We are now ready to define our main algorithm for sparsifying arbitrary r -uniform input hypergraphs $G = (V, E)$. We use EXPANDERDECOMPOSITION to produce subsets of V that are good expanders and EXPANDERSPARSIFY to sparsify them. Since we get rid of at least half the edges each turn, this process, repeated until no hyperedges remain, would take potentially $\Omega(\log m) = \tilde{\Omega}(r \log n)$ rounds. To reduce the number of rounds, we contract clusters into supernodes shortly after sparsifying them. For more intuition on the workings of Algorithm 2 see Section 3.2.

Algorithm 2 Algorithm sparsifying an arbitrary hypergraph

```

1: procedure SPARSIFY( $G, \epsilon$ )
2:    $E^{(0)} \leftarrow E$ 
3:    $\tilde{E} \leftarrow \emptyset$ 
4:   for  $i = 0, \dots, \log m$  do
5:     if  $i \geq 10 \log n$  then
6:        $\sim_i \leftarrow$  the relation on  $V$  where  $u \sim_i v$  iff  $\exists i' \leq i - 10 \log n, \exists j, [u], [v] \in C_j^{(i')}$ 
7:        $\approx_i \leftarrow$  transitive closure of  $\sim_i$ 
8:        $G^{(i)} \leftarrow (V, E^{(i)}) / \approx_i$ 
9:     else
10:       $G^{(i)} \leftarrow (V, E^{(i)})$ 
11:       $(C_1^{(i)}, \dots, C_{k_i}^{(i)}) \leftarrow \text{EXPANDERDECOMPOSITION}(G^{(i)})$ 
12:      for  $j = 1, \dots, k_i$  do
13:         $G_j^{(i)} = (C_j^{(i)}, E_j^{(i)}) \leftarrow G^{(i)}[C_j^{(i)}]$ 
14:         $(C_j^{(i)}, \tilde{E}_j^{(i)}, w_j^{(i)}) \leftarrow \text{EXPANDERSPARSIFY}(G_j^{(i)}, \epsilon/10)$ 
15:         $\tilde{E} \leftarrow \tilde{E} \cup \tilde{E}_j^{(i)}$ 
16:         $w|_{E_j^{(i)}} = w_j^{(i)}$ 
17:       $E^{(i+1)} \leftarrow E^{(i)} \setminus \bigcup_{j=1}^{k_i} E_j^{(i)}$ 
18:   return  $\tilde{G} = (V, \tilde{E})$ 

```

G is assumed to be r -regular, and ϵ is assumed to be in $[1000/n, 1/2]$. In Line 6, $[u], [v] \in C_j^{(i')}$ means that the *supernodes containing u and v* were in the same expander of the decomposition at an

earlier level i' . Note also the abuse of notation in Lines 14 and 15 (in accordance with Remark 6.4): In Line 14, $\tilde{E}_j^{(i)}$ is defined as a multiset of hyperedges on the *contracted* vertex set $C_j^{(i)} \subseteq V / \approx_i$, but in Line 15 we treat it as containing hyperedges supported on V . This is justified because elements of $\tilde{E}_j^{(i)}$ have clearly corresponding counterparts in E (recall our definition of a sparsifier as a weighted subgraph and that a contraction maintains a bijection between the hyperedges), and it would only complicate the notation to make this distinction formal.

Also note that throughout the contractions, our graphs always remain r -regular, thanks to the use of multisets as hyperedges (see Definition 6.3). This — along with Guarantee 1 of Lemma 5.1 that the expansion of $G_j^{(i)}$ is $\Omega(\frac{1}{r \log^2 n})$ — allows the use of EXPANDERSPARSIFY in Line 14.

Line 16 simply means that we set the weights of hyperedges in $E_j^{(i)}$ as in the sparsifier computed in Line 14. This is consistent with the update of \tilde{E} in Line 15.

6.2 Correctness

First we bound the total size and number of clusters $C_j^{(i)}$:

Claim 6.6.

$$\sum_{i=0}^{\log m} \sum_{j=1}^{k_i} |C_j^{(i)}| \leq 21n \log n.$$

Proof. We first bound the number of *distinct* vertices (including supernodes) that appear throughout the execution of Algorithm 2, i.e.,

$$V^* = \bigcup_{i=0}^{\log m} \bigcup_{j=1}^{k_i} C_j^{(i)}.$$

This set includes vertices from V , as well as some contracted supernodes that are subsets of V . Note however, that these sets form a laminar family, since supernodes are only constructed by merging previous supernodes — no supernodes ever get broken apart. Therefore, $|V^*| \leq 2n - 1$.

Let us first bound the number of *non-singleton* clusters a single $v \in V^*$ can participate in. As soon as v participates in some cluster C of size at least 2 at level i , we know that at level $i + 10 \log n$ C will be contracted into a different vertex in V^* . Therefore, v can participate in at most $10 \log n$ clusters of size 2 or more.

On the other hand, any $v \in V^*$ can participate in at most one singleton cluster. Indeed as soon as v forms a cluster $\{v\}$ all of its self-loop hyperedges are removed, and it will not have self-loops until it is contracted again. Therefore, $\{v\}$ cannot appear again as a cluster as this would violate Guarantee 3 of Lemma 5.1.

Therefore, the total of the sum is at most $(2n - 1) \cdot (10 \log n + 1) \leq 21n \log n$, for large enough n . \square

As an immediate corollary to this, we can conclude that there are at most $21n \log n$ clusters considered throughout the entire algorithm, since the size of a cluster is always at least one. Therefore, EXPANDERSPARSIFY — the only non-deterministic step of our algorithm — gets called at most $21n \log n$ times, and succeeds *all* of these times with combined probability $1 - O((\log^2 n)/n)$. From this point on we consider only the event that all of these calls indeed succeed.

Let us consider the size of the sparsifier output by SPARSIFY(G, ϵ).

Lemma 6.7. *Let G be r -regular hypergraph and let $(r \log^6 n)/n \leq \epsilon \leq 1/2$. The hypergraph returned by SPARSIFY(G, ϵ) has $nr(\epsilon^{-1} \log n)^{O(1)}$ hyperedges with probability $1 - O((\log^2 n)/n)$.*

Proof. The hyperedge multiset of the hypergraph $\tilde{G} = (V, \tilde{E})$ is simply the union of $\tilde{E}_j^{(i)}$ for all i and j , and by Lemma 6.1 the size of each of them is bounded by $|\tilde{E}_j^{(i)}| \leq |C_j^{(i)}| r(\epsilon^{-1} \log n)^{O(1)}$ with probability $1 - O(\log n^2/n)$. So by Claim 6.6

$$|\tilde{E}| \leq \sum_{i=0}^{\log m} \sum_{j=1}^{k_i} |C_j^{(i)}| r(\epsilon^{-1} \log n)^{O(1)} \leq nr(\epsilon^{-1} \log n)^{O(1)},$$

as claimed. \square

Lemma 6.8. *Let G be r -regular hypergraph and let $(r \log^6 n)/n \leq \epsilon \leq 1/2$. Then hypergraph $\text{SPARSIFY}(G, \epsilon)$ an ϵ -spectral sparsifier to G in polynomial time with probability at least $1 - O((\log^2 n)/n)$.*

Proof. We will consider each cluster $C_j^{(i)}$ separately, it suffices to prove that for each i and j , the energy of $(V, \tilde{E}_j^{(i)})$ approximates the energy of $(V, E_j^{(i)})$ up to small additive and multiplicative errors. We must first verify that each hyperedge of E appears in exactly one of $E_j^{(i)}$. Indeed, any hyperedge can appear in at most one of them: $E_j^{(i)}$ gets removed from $E^{(i+1)}$ at the end of the main for loop in Line 17. Furthermore, all hyperedges eventually get removed this way: By Guarantee 3 of Lemma 5.1, after each round no more than half of the hyperedges remain, so after $\log m + 1$ rounds all hyperedges must be gone and $E^{(\log m + 1)} = \emptyset$.

Let us fix a single cluster $C_j^{(i)}$ and compare the spectral properties of $E_j^{(i)}$ and $\tilde{E}_j^{(i)}$. Let $\hat{C}_j^{(i)}$ be the set of vertices from V that make up $C_j^{(i)}$ after contraction by \approx_i , that is, formally

$$\hat{C}_j^{(i)} = \{v \in V \mid [v] \in C_j^{(i)}\},$$

where $[v]$ denotes the equivalence class of v with respect to \approx_i . At this point it is important to keep in mind the differences between the graphs $\hat{G}_j^{(i)} = (\hat{C}_j^{(i)}, E_j^{(i)})$, $G_j^{(i)} = (C_j^{(i)}, E_j^{(i)})$, $\tilde{G}_j^{(i)} = (C_j^{(i)}, \tilde{E}_j^{(i)})$, and $\check{G}_j^{(i)} = (\hat{C}_j^{(i)}, \tilde{E}_j^{(i)})$. Though $\hat{G}_j^{(i)}$ and $G_j^{(i)}$ share their hyperedge multiset, they are not the same, in fact $G_j^{(i)} = \hat{G}_j^{(i)} / \approx_i$. Similarly, $\tilde{G}_j^{(i)} = \check{G}_j^{(i)} / \approx_i$. Lemma 6.1 says that $\tilde{G}_j^{(i)}$ is an ϵ -spectral sparsifier to $G_j^{(i)}$, it makes no such guarantee about $\hat{G}_j^{(i)}$ and $\check{G}_j^{(i)}$.

We will show that for every vector $x \in \mathbb{R}^V$,

$$\check{Q}(x) = \left(1 \pm \frac{4\epsilon}{10}\right) \hat{Q}(x) \pm \frac{3\epsilon Q(x)}{n^2}, \quad (9)$$

where \hat{Q} and \check{Q} denote the energy with respect to $\hat{G}_j^{(i)}$ and $\check{G}_j^{(i)}$, respectively. Recall from Observation 6.5 that the equation

$$\check{Q}(x) = \left(1 \pm \frac{\epsilon}{10}\right) \hat{Q}(x) \quad (10)$$

holds when x is constant within all supernodes of $C_j^{(i)}$. Informally, our plan is to round x to \tilde{x} such that it satisfies this requirement and then show that

$$\hat{Q}(x) \cong \hat{Q}(\tilde{x}) \cong \check{Q}(\tilde{x}) \cong \check{Q}(x).$$

We may assume without loss of generality that $i \geq 10 \log n$, as otherwise no contraction takes place and Equation (10) holds trivially. Denote the maximum discrepancy between the values of x within supernodes of $C_j^{(i)}$ by

$$\delta = \max_{u, v \in \hat{C}_j^{(i)}: u \approx_i v} |x_u - x_v|.$$

Our rounding of x to \tilde{x} is mostly arbitrary, we enforce only that no coordinate changes by more than an additive δ .

We will show the following three claims in the next section.

Claim 6.9.

$$\widehat{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \widehat{Q}(x) \pm \frac{\epsilon Q(x)}{n^2}.$$

Claim 6.10.

$$\check{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \widehat{Q}(\tilde{x}).$$

Claim 6.11.

$$\check{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \check{Q}(x) \pm \frac{\epsilon Q(x)}{n^2}.$$

By combining Claims 6.9, 6.10, and 6.11 we get that the quadratic form \check{Q} is indeed close to the quadratic form \widehat{Q} , with small additive and multiplicative error, as claimed. As in the proof of Lemma 4.10, first note that the multiplicative error between $\widehat{Q}(x)$ and $\check{Q}(x)$ by itself would only amount to $(1 \pm 4\epsilon/10)$. Similarly, the additive error by itself would be exactly $2\epsilon Q(x)/n^2$, small enough even when summed over all possible clusters $C_j^{(i)}$. This is because the number of clusters throughout the whole algorithm is bounded by $21n \log n$ due to Claim 6.6.

Formally, we consider one direction of the bound first:

$$\begin{aligned} \check{Q}(x) &\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{\epsilon Q(x)}{n^2} + \check{Q}(\tilde{x}) \right] && \text{By Claim 6.11} \\ &\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{\epsilon Q(x)}{n^2} + \left(1 + \frac{\epsilon}{10}\right) \widehat{Q}(\tilde{x}) \right] && \text{By Claim 6.10} \\ &\leq \left(1 - \frac{\epsilon}{10}\right)^{-1} \left[\frac{\epsilon Q(x)}{n^2} + \left(1 + \frac{\epsilon}{10}\right) \left[\frac{\epsilon Q(x)}{n^2} + \left(1 + \frac{\epsilon}{10}\right) \widehat{Q}(x) \right] \right] && \text{By Claim 6.9} \\ &\leq \left(1 + \frac{4\epsilon}{10}\right) \widehat{Q}(x) + \frac{3\epsilon Q(x)}{n^2}, \end{aligned}$$

since $\epsilon \leq 1/2$. The other direction that

$$\check{Q}(x) \geq \left(1 - \frac{\epsilon}{10}\right) \widehat{Q}(x) - \frac{3\epsilon Q(x)}{n^2}$$

follows similarly, which concludes the proof of Equation (9).

Finally, we can sum over all clusters $C_j^{(i)}$, noting that their number cannot exceed $21n \log n$, to get that $Q_{\tilde{E}}(x) = (1 \pm \epsilon)Q(x)$. \square

Combining Lemmas 6.7 and 6.8, we get a polynomial-time algorithm that, given an unweighted r -uniform hypergraph, constructs an ϵ -spectral sparsifier with $nr(\epsilon^{-1} \log n)^{O(1)}$ hyperedges. We will discuss how to handle weighted hypergraphs and reduce the running time to $O(mr^2) + n^{O(1)}$ in Section 6.4.

6.3 Proofs of Claims 6.9, 6.10, and 6.11

Before proceeding on to Claims 6.9, 6.10, and 6.11 from the previous section, we prove a simple supporting lemma. This allows us to relate the total energy of G to the energies of the various $C_j^{(i)}$ clusters.

Lemma 6.12. *Let $G = (V, E)$ be an arbitrary hypergraph and let x be a vector in \mathbb{R}^V . Let \approx be an equivalence relation on V and define the contraction $G' = (V', E') = G/\approx$. Let $x' \in \mathbb{R}^{V'}$ be a specification of x on vertices of G' such that*

$$\forall v' \in V', \exists v \in v' : x'_{v'} = x_v.$$

That is each vertex in V' takes the value of one of the vertices in V from which it was contracted. Then

$$Q(x) \geq Q'(x),$$

where Q and Q' denote the energy with respect to G and G' respectively.

Proof. We examine each hyperedge of E separately. Let $e \in E$ and let the corresponding hyperedge in E' be e' . By definition $Q_x(e) = \max_{a,b \in e} (x_a - x_b)^2$ and $Q'_{x'}(e') = \max_{a,b \in e'} (x_a - x_b)^2$. By definition of x' , each value of x' in e' also appears as a value of x in e . Therefore, $Q_x(e)$ and $Q'_{x'}(e')$ are maximizations of the same formula, with the former having more choice in x_a and x_b , so $Q_x(e) \geq Q'_{x'}(e')$. Summing this over all hyperedges we get

$$Q(x) \geq Q'(x'). \quad \square$$

We are now ready to prove a claim bounding the effect of rounding from x to \tilde{x} on the energy of a single hyperedge. This is the main technical result of the section that allows us to do contraction in Algorithm 2.

Claim 6.13. *For all $e \in E_j^{(i)}$,*

$$Q_{\tilde{x}}(e) = \left(1 \pm \frac{\epsilon}{10}\right) Q_x(e) \pm \frac{\epsilon Q(x)}{2n^3 m_i},$$

where m_i is the size of $E^{(i)}$.

Proof. Recall that the additive error between x and \tilde{x} is at most δ , thus by Lemma 4.11

$$|Q_x(e) - Q_{\tilde{x}}(e)| \leq 4\delta \left(\sqrt{Q_x(e)} + \delta\right).$$

We distinguish between two cases based on the size of $\sqrt{Q_x(e)}$ relative to δ . First suppose that $\delta \leq \epsilon\sqrt{Q_x(e)}/50$. This is the simpler case, because we immediately get

$$|Q_x(e) - Q_{\tilde{x}}(e)| \leq \frac{4\epsilon}{50} \sqrt{Q_x(e)} \cdot \left(\sqrt{Q_x(e)} + \frac{\epsilon}{50} \sqrt{Q_x(e)}\right) \leq \frac{\epsilon}{10} Q_x(e).$$

Now consider the second case, $\delta \geq \epsilon\sqrt{Q_x(e)}/50$. This time we get

$$|Q_x(e) - Q_{\tilde{x}}(e)| \leq 4\delta \left(\frac{50\delta}{\epsilon} + \delta\right) \leq \frac{204\delta^2}{\epsilon}.$$

Based on the definition of δ , let $u, v \in \widehat{C}_j^{(i)}$ be such that $u \approx_i v$ and $|x_u - x_v| = \delta$. By Definition 6.2 there must exist a sequence of vertices, or *path*, $u = w_0, \dots, w_k = v$ such that $w_{\ell-1} \sim_i w_\ell$ for all $\ell \in [k]$ (\sim_i is defined in Line 6). Without loss of generality, assume the path length k is minimal and therefore at most n . Then by averaging, there exists $\ell' \in [k]$ such that $|x_{w_{\ell'-1}} - x_{w_{\ell'}}| \geq \delta/k \geq \delta/n$. By definition of \sim_i , there exist $i' \leq i - 10 \log n$ and j' such that $w_{\ell'-1}, w_{\ell'} \in C_{j'}^{(i')}$ but $w_{\ell'-1} \not\sim_{i'} w_{\ell'}$.

We now wish to relate $Q(x)$ to δ^2 . We will accomplish this by lower bounding $Q(x)$ by the energy of $G_{j'}^{(i')}$ with respect to some vector x' , as per Lemma 6.12. Let us define $x' \in \mathbb{R}^{G_{j'}^{(i')}}$ as in Lemma 6.12 such that the supernode of $w_{\ell'-1}$ retains the x -value of $w_{\ell'-1}$ and the supernode of $w_{\ell'}$ retains the x -value of $w_{\ell'}$. Formally,

$$\begin{aligned} x'_{[w_{\ell'-1}]} &= x_{w_{\ell'-1}}, \\ x'_{[w_{\ell'}]} &= x_{w_{\ell'}}, \end{aligned}$$

where $[w]$ denotes the equivalence class with respect to $\approx_{i'}$. All other coordinates of x' are defined arbitrarily. By applying Lemma 6.12 and discard unnecessary hyperedges we can conclude that

$$Q(x) \geq Q'(x'),$$

the energy of x on $G_{j'}^{(i')}$.

To lower bound the energy of x on $G_{j'}^{(i')}$, we can apply the hypergraph Cheeger inequality (Theorem 2.3). We know by Guarantee 1 of EXPANDERDECOMPOSITION that $\Phi\left(G_{j'}^{(i')}\right) \geq \Omega\left(\frac{1}{r \log^2 n}\right)$. Let $m_{i'}$ be the number of hyperedges in $E^{(i')}$. Then by Guarantee 2 of EXPANDERDECOMPOSITION, the minimum degree of $G_{j'}^{(i')}$ is at least $m_{i'}/4n$. In general, the hypergraph Cheeger inequality states that when x is centered, that is $\sum_{v \in V} x_v d(v) = 0$, we have

$$Q(x) \geq \frac{r\Phi^2}{32} \sum_v x_v^2 d(v).$$

Our vector x' is not centered with respect to the cluster $C_{j'}^{(i')}$. However, since the difference between $x_{[w_{\ell'}]}$ and $x_{[w_{\ell'+1}]}$ is at least δ/n , at least one of them will have absolute value at least $\delta/(2n)$ even when x is centered. Therefore, the terms corresponding to $[w_{\ell'-1}]$ and $[w_{\ell'}]$ already give

$$Q(x) \geq Q_{G_{j'}^{(i')}}(x) \geq \frac{r\Omega\left(\frac{1}{r \log^2 n}\right)^2}{32} \cdot \frac{m_{i'}}{4n} \cdot \left(\frac{\delta}{2n}\right)^2.$$

Putting these together, we have

$$|Q_x(e) - Q_{\bar{x}}(e)| \leq \frac{204\delta^2}{\epsilon} \leq \frac{204}{\epsilon} \cdot (2n)^2 \cdot \frac{32}{r \cdot \Omega\left(\frac{1}{r \log^2 n}\right)^2} \cdot \frac{4n}{m_{i'}} \cdot Q(x) \leq \frac{O(1) \cdot n^3 r \log^4 n}{\epsilon m_{i'}} \cdot Q(x)$$

To relate this to m_i , as stated in the claim, recall that by Guarantee 3 of EXPANDERDECOMPOSITION the total number of hyperedges decreases at least by a factor of 2 during every iteration of the outer for-loop. So $i' \leq i - 10 \log n$ implies that $m_i \leq m_{i'} n^{-10}$, and we get

$$|Q_x(e) - Q_{\bar{x}}(e)| \leq \frac{O(\log^4 n) \cdot n^3 r}{\epsilon m_i n^{10}} \cdot Q(x) \leq \frac{\epsilon Q(x)}{2n^3 m_i},$$

where the last inequality is because r , ϵ^{-1} , and the $O(\log^4 n)$ term are all smaller than n , by the assumptions of Lemma 6.8 and for large enough n .

Putting the two cases together gives us the additive and multiplicative error terms and completes the proof of Claim 6.13. \square

Claims 6.9 and 6.11 follow as a result of this, while Claim 6.10 is a simple consequence of Observation 6.5.

Claim 6.9.

$$\widehat{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \widehat{Q}(x) \pm \frac{\epsilon Q(x)}{n^2}.$$

Proof. This follows immediately from Claim 6.13, as $\widehat{G}_j^{(i)}$ has at most m_i hyperedges,

$$\left| \widehat{Q}(x) - \widehat{Q}(\tilde{x}) \right| \leq \sum_{e \in E_j^{(i)}} \left| Q_x(e) - Q_{\tilde{x}}(e) \right| \leq \sum_{e \in E_j^{(i)}} \left[\epsilon Q_x(e) + \frac{\epsilon Q_x(e)}{2n^3 m_i} \right] \leq \epsilon \widehat{Q}(x) + \frac{\epsilon Q(x)}{n^2}. \quad \square$$

Claim 6.10.

$$\check{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \widehat{Q}(\tilde{x}).$$

Proof. This follows immediately from Equation (10) because \tilde{x} is constant on all equivalence classes of \approx_i in $C_j^{(i)}$. \square

Claim 6.11.

$$\check{Q}(\tilde{x}) = \left(1 \pm \frac{\epsilon}{10}\right) \check{Q}(x) \pm \frac{\epsilon Q(x)}{n^2}.$$

Proof. The proof follows similarly to the proof of Claim 6.9. However,

$$\check{Q}(x) = \sum_{e \in \tilde{E}_j^{(i)}} w_e \cdot Q_x(e),$$

so we must prove that the total weight $\sum_{e \in \tilde{E}_j^{(i)}} w_e$ of the sparsifier is not too large. Note that this is not the same as the size of the sparsifier, which is guaranteed to be small by Lemma 5.1.

To do this, note that we can bound the total weight of hyperedges adjacent on a specific vertex, say v , by looking at the energy of the vector $\mathbb{1}_v$, which has value 1 on v and 0 everywhere else. (Here it is important that there are no self-loop hyperedges in $\tilde{E}_j^{(i)}$). So we have

$$\sum_{e \in \tilde{E}_j^{(i)}} w_e \leq \sum_{v \in C_j^{(i)}} \sum_{e \in \tilde{E}_j^{(i)}: v \in e} w_e = \sum_{v \in C_j^{(i)}} Q_{\tilde{G}_j^{(i)}}(\mathbb{1}_v) \leq \sum_{v \in C_j^{(i)}} (1 + \epsilon) Q_{G_j^{(i)}}(\mathbb{1}_v) \leq 2 \sum_{v \in C_j^{(i)}} |\tilde{E}_j^{(i)}| \leq 2nm_i.$$

Therefore,

$$\left| \check{Q}(x) - \check{Q}(\tilde{x}) \right| \leq \sum_{e \in \tilde{E}_j^{(i)}} w_e \cdot \left| Q_x(e) - Q_{\tilde{x}}(e) \right| \leq \sum_{e \in \tilde{E}_j^{(i)}} w_e \cdot \left[\epsilon Q_x(e) + \frac{\epsilon Q_x(e)}{2n^3 m_i} \right] = \epsilon \check{Q}(x) + \frac{\epsilon Q(x)}{n^2}. \quad \square$$

6.4 Weighted Hypergraphs and Proof of Theorem 1.1

We have so far dealt only with unweighted hypergraphs, so as not to further complicate our algorithms and notation. However, our techniques extend essentially unchanged to weighted ones as well.

One way to see this is to replace weighted hyperedges with parallel hyperedges. Our proofs throughout Sections 2–6 apply to hypergraphs that may contain hyperedges with multiplicity. Given a weighted graph, we can simply scale the weights up (or down if necessary) and approximate them with integer weights arbitrarily closely. A weighted hypergraph where the ratio between the largest and smallest weights is w_{\max}/w_{\min} can be approximated to within a multiplicative $1 \pm \eta$ error using $\log(\eta^{-1}) \cdot w_{\max}/w_{\min}$ parallel hyperedges to replace each weighted hyperedge. The parameter η can be set to $o(\epsilon)$ so as to still produce a good spectral approximation.

One might worry that this increases the running time since the number of hyperedges has technically increased. However, this turns out not to be the case: the running times of all of our algorithms scale polynomially with the number of *distinct* hyperedges. Indeed, parallel edges can be considered simultaneously at every step. In the expander sparsification algorithm of Section 4 the sampling probability of parallel hyperedges is the same, and at most $r(\epsilon^{-1} \log n)^{O(1)}$ of them are sampled. In Algorithm 1 of Section 6 each parallel instance of the same hyperedge gets cut by the same cuts and ends up in the same component. Consequently parallel edges end up on the same level in the same component in Algorithm 2 of Section 6 and are sampled at the same rate.

In fact, one can verify that our proofs extend even more directly to weighted hypergraphs, without the need for approximating hyperedge weights by integers.

Until now, all of our algorithms have claimed only polynomial running time. Surprisingly it is the above extension that allows us to accelerate the runtime to nearly linear — even in the case of unweighted input graphs.

Proof of Theorem 1.1. Given a hypergraph G , we can apply the previously known hypergraph sparsification algorithm of [SY19] to get a polynomial (in n) size sparsifier in nearly linear (that is $O(mr^2) + n^{O(1)}$) time. We can then further sparsify this using our own Algorithm 2 in time $n^{O(1)}$. Setting the error parameters of both algorithms to $\epsilon/3$ allows us to recover an ϵ -spectral sparsifier of G , as desired.

Note that we may drop the $(r \log^6 n)/n \leq \epsilon$ requirement of Lemma 6.8 without loss of generality. \square

7 Lower Bounds

In this section we prove our space lower bound for an arbitrary compression of the cut structure of a hypergraph. In Section 7.1 we introduce string compression, and reprove the corresponding lower bound result for completeness. In section 7.2 we construct our generic hard example in Theorem 7.9. We then state Corollaries 7.11, 7.12, and 7.13 which result from applying Theorem 7.9 to various specific Ruzsa-Szemerédi graph constructions.

7.1 String Compression

A *string compression scheme (SCS)* is an algorithm for compressing a long string into a short string, such that any subset sum query can be answered with small additive error. Formally, we define it as follows.

Definition 7.1. For positive integers ℓ, k and $\epsilon, g > 0$, a pair of functions $\text{ENCODE} : \{0, 1\}^\ell \rightarrow \{0, 1\}^k$ and $\text{DECODE} : \{0, 1\}^k \times 2^{[\ell]} \rightarrow \mathbb{N}$ is considered to be an (ℓ, k, ϵ, g) -SCS, if there exists a set of strings $\mathcal{G} \subseteq \{0, 1\}^\ell$, such that the following holds.

- $|\mathcal{G}| \geq g \cdot 2^\ell$.
- For every string $s \in \mathcal{G}$ and every query $q \in 2^{[\ell]}$, $|\text{DECODE}(\text{ENCODE}(s), q) - |s \cap q|| \leq \epsilon\ell/2$.

Remark 7.2. In general we use subsets of $[\ell]$ and elements of $\{0, 1\}^\ell$ interchangeably. For instance, in the above definition, in $|s \cap q|$, s is considered as a set.

Remark 7.3. It is important that although a compression scheme may only work on a subset of strings (\mathcal{G}), it must work on all queries. In fact, it is trivial to answer almost all queries on all inputs, by simply outputting $|q| \cdot |s|/\ell$.

The lower following lower bound on the space requirement of string compression schemes has been known, and appears, for example, in [DN03]. We reprove it here for completeness.

Theorem 7.4. Suppose $(\text{ENCODE}, \text{DECODE})$ is an (ℓ, k, ϵ, g) -SCS, where $\epsilon \leq 1/10$. Then

$$k \geq \frac{\log g + 3\ell/50}{\log 2} - 1.$$

Proof. We know that ENCODE maps \mathcal{G} into $\{0, 1\}^k$. Therefore, by pigeonhole principle, there must be some set of inputs \mathcal{G}_0 of size at least $|\mathcal{G}| \cdot 2^{-k} \geq g \cdot 2^{\ell-k}$ that maps to the same output, say c_0 . Let s_0 be an arbitrary string in \mathcal{G}_0 .

Define $B_H(s_0, 2\epsilon\ell)$ as the ball of radius $2\epsilon\ell$ in Hamming distance around s_0 , that is, the set of strings $s \in \{0, 1\}^\ell$ such that the number of coordinates where s and s_0 differ is at most $2\epsilon\ell$.

Claim 7.5. $\mathcal{G}_0 \subseteq B_H(s_0, 2\epsilon\ell)$.

Proof. Suppose there exists $s \in \mathcal{G}_0 \setminus B_H(s_0, 2\epsilon\ell)$, that is s and s_0 differ on more than $2\epsilon\ell$ coordinates. Without loss of generality, we may assume that there are more than $\epsilon\ell$ coordinates where s_0 is 0 but s_1 is 1; let the set of such coordinates be q . By the definition of a string compression scheme

$$\text{DECODE}(\text{ENCODE}(s_0), q) = \text{DECODE}(c_0, q) \leq |s_0 \cap q| + \epsilon\ell/2 = \epsilon\ell/2,$$

but

$$\text{DECODE}(\text{ENCODE}(s), q) = \text{DECODE}(c_0, q) \geq |s \cap q| - \epsilon\ell/2 = |q| - \epsilon\ell/2 > \epsilon\ell/2.$$

This is a contradiction. □

Claim 7.6. $|B_H(s_0, 2\epsilon\ell)| < 2^\ell \cdot 2 \exp\left(-\frac{\ell(1-4\epsilon)^2}{6}\right)$.

Proof. Indeed,

$$B_H(s_0, 2\epsilon\ell) = B_H(0^\ell, 2\epsilon\ell) = 2^\ell \cdot \mathbb{P}(w_H(x) \leq 2\epsilon\ell),$$

where x is a uniformly random vector in $\{0, 1\}^\ell$. By Chernoff's bound

$$\mathbb{P}[w_H(x) \leq 2\epsilon\ell] \leq \mathbb{P}\left[\left|w_H(x) - \frac{\ell}{2}\right| \geq \frac{\ell}{2}(1-4\epsilon)\right] \leq 2 \exp\left(-\frac{\ell(1-4\epsilon)^2}{6}\right),$$

since $\epsilon \leq 1/4$, and the claim holds. □

Combining Claims 7.5 and 7.6 we get that

$$\begin{aligned}
g \cdot 2^{\ell-k} &\leq 2^\ell \cdot 2 \exp\left(-\frac{\ell(1-4\epsilon)^2}{6}\right), \\
\Rightarrow \log g - k \log 2 &\leq \log 2 - \frac{\ell(1-4\epsilon)^2}{6}, \\
\Rightarrow k &\geq \frac{\log g - \log 2 + \frac{\ell(1-4\epsilon)^2}{6}}{\log 2} \geq \frac{\log g + 3\ell/50}{\log 2} - 1,
\end{aligned}$$

since $\epsilon \leq 1/10$. □

Corollary 7.7. *For $\ell \geq 200$, there does not exist an $(\ell, k, 1/10, 1/2)$ -SCS with $k < \ell/20$.*

7.2 Construction

We will derive a lower bound on k from the existence of a *Ruzsa-Szemerédi (RS) graph*, defined as follows.

Definition 7.8 (Ruzsa-Szemerédi graph). *We call an (ordinary) graph a (t, a) -RS graph if its edge set is the union of t induced matchings of size a .*

Recall the definition of hypergraph cut sparsification schemes from Section 3.3:

Definition 3.1. *Let $\mathfrak{H}(n, r)$ be the set of hypergraphs on a vertex set $[n]$ with each hyperedge having size at most r . A pair of functions $\text{SPARSIFY} : \mathfrak{H}(n, r) \rightarrow \{0, 1\}^k$ and $\text{CUT} : \{0, 1\}^k \times 2^{[n]} \rightarrow \mathbb{N}$ is said to be an (n, r, k, ϵ) -HCSS if for all inputs $G = (V, E) \in \mathfrak{H}(n, r)$ the following holds.*

- For every query $S \in 2^{[n]}$, $|\text{CUT}(\text{SPARSIFY}(G), S) - |E(S, \bar{S})|| \leq \epsilon \cdot |E(S, \bar{S})|$.

Theorem 7.9. *Suppose there exists a (t, a) -RS graph on n vertices where $a \geq 6000\sqrt{n \log n}$ and $at \geq 480n$. Then, any $(2n, t+1, k, \epsilon)$ -HCSS where $\epsilon \leq a/(60n)$ must have*

$$k = \Omega(at).$$

This is equivalent to Theorem 1.2.

Proof. Let us fix such a (t, a) -RS graph G on n vertices, and a $(2n, t+1, k, \epsilon)$ -HCSS $(\text{SPARSIFY}, \text{CUT})$. We will use this HCSS as a black box to construct a string compression scheme using k bits of space, then bound k by Corollary 7.7. First let us convert G into a bipartite graph G' . Let the vertex set of $G' = (V', E')$ be $V \times \{0, 1\}$ where $P = V \times \{0\}$ and $Q = V \times \{1\}$ are the two sides of the bipartition. For each edge $e = (u, v) \in E$, we add two edges to E' : $((u, 0), (v, 1))$ and $((v, 0), (u, 1))$, ensuring that G' is indeed bipartite. Note that E' is the union of t induced matchings of size $2a$. Let us call these matchings M_1, \dots, M_t and let each M_j be supported on P_j in P and Q_j in Q . The maximum degree in G' is t .

We will use G' to design a compression of strings of length $\ell = 2ta$. Note that there are exactly $2ta$ edges of G' . Let ϕ be an arbitrary bijection from E' to $[\ell]$. For a string $s \in \{0, 1\}^\ell$, Let E_s be the subset of E' defined as

$$E_s = \{e \in E' : s_{\phi(e)} = 1\}.$$

Thus the graph $G_s = (P \cup Q, E_s)$ encodes the string s . We then transform G_s into the hypergraph $H_s = (P \cup Q, E_s^H)$. Let E_s^H consist of one hyperedge corresponding to each vertex $u \in P$:

$$E_s^H = \{\{u\} \cup \Gamma_s(u) \mid u \in P\},$$

where Γ_s denotes the neighborhood in G_s .

Our compression function ENCODE is then simply to sparsify H_s using SPARSIFY. This can indeed be done, since H_s is a hypergraph with $2n$ vertices and each edge has cardinality at most $t + 1$. It remains to define the decoding function DECODE.

Given a query $q \subseteq [\ell]$, we must estimate the size of $s \cap q$, the number of coordinates of s within q having value 1. To do this, we partition q into segments q^1, \dots, q^t , and then estimate the size of each $s \cap q^j$. Specifically, let

$$q^j = \{i \in q \mid \phi^{-1}(i) \in M_j\}.$$

We can then define

$$\text{DECODE}(\text{SPARSIFY}(H_s), q) = \sum_{j=1}^t \text{DECODE}^j(\text{SPARSIFY}(H_s), q^j).$$

Here DECODE^j remains undefined for now. In what follows we will define it such that

$$\text{DECODE}^j(\text{SPARSIFY}(H_s), q^j) \cong |s \cap q^j|.$$

To estimate the size of $s \cap q^j$, we will observe the cut $E_s^H(S, \bar{S}) = E_s^H(S_s^j, \bar{S}_s^j)$ defined as follows:

- From P , S contains the subset of vertices in P_j corresponding to edges in q^j . Formally

$$S \cap P = \{P \cap e \mid e \in M_j \text{ s.t. } \phi(e) \in q^j\}.$$

- From Q , S contains all vertices except Q_j .

We will prove the the size of the cut (S, \bar{S}) is closely related to the size of $s \cap q$, as long as s satisfies some nice properties.

Note that each hyperedge in E_s^H corresponds to a vertex in P : for $u \in P$ we denote the hyperedge $\{u\} \cup \Gamma_s(u)$ as e_u . We will bound the contribution of e_u to the cut (S, \bar{S}) for all u in each of the following three categories:

1. $u \in P_j \cap S$:

Let the edge from M_j adjacent on u be f_u . For any such u , e_u crosses the cut if and only if $s_{\phi(f_u)} = 1$. Indeed, if $s_{\phi(f_u)} = 1$, then $f \in E_s$ and $f \cap Q \in Q_j \subseteq \bar{S}$. On the other hand, if $s_{\phi(f_u)} = 0$ then $f \notin E_s$ and all edges adjacent on u in E_s correspond to matchings different from M_j (that is M_k for $k \neq j$). Since M_j is induced by the property of RS-graphs, $\Gamma_s(u) \subseteq Q \setminus Q_j \subseteq S$. Therefore, the total amount of hyperedges crossing the cut from this category is exactly $|s \cap q^j|$.

2. $u \in P_j \setminus S$:

In this case e_u crosses the cut unless $d_s(u) < 2$. Indeed, if $d_s(u) \geq 2$ then at least one edge adjacent on u in G_s does *not* come from M_j . The other endpoint of this edge is in $Q \setminus Q_j \subseteq S$, whereas u itself is in \bar{S} by definition. In the the case where $d_s(u) < 2$ we cannot say whether e_u crosses the cut or not. Therefore, the number of hyperedges crossing the cut from this category is approximately $m - |q^j|$ (that is all of them), but with a possible error of

$$|\{u \in P_j \mid d_s(u) < 2\}|.$$

3. $u \in P \setminus P_j$:

In this case we cannot say anything about the number of edges crossing the cut, except that it is unlikely to deviate from its expectation when s is considered to be uniformly random on $\{0, 1\}^\ell$. Let

$$Z_j = |\{u \in P \setminus P_j \mid \Gamma_s(u) \not\subseteq Q_j\}|,$$

or the number of hyperedges in E_s^H from this category crossing the cut.

Overall, we can approximate the size of the cut (S, \bar{S}) in H_s by

$$|s \cap q^j| + (a - |q^j|) + \mathbb{E}_s Z_j, \quad (11)$$

with an maximum additive error of

$$|\{u \in P_j \mid d_s(u) < 2\}| + |Z_j - \mathbb{E}_s Z_j|. \quad (12)$$

Conversely, this allows us to approximate $|s \cap q^j|$ using the size of the same cut in our $(2n, t + 1, k, \varepsilon)$ -HCSS. Therefore, we define DECODE^j as follows:

$$\text{DECODE}^j(\text{SPARSIFY}(H_s), q^j) = \text{CUT}(\text{SPARSIFY}(H_s), S) - (a - |q^j|) - \mathbb{E}_s Z_j. \quad (13)$$

It remains to bound the total error introduced by the inaccuracies above.

We will define the set of good input strings, \mathcal{G} to be those where this additive error is small across all j 's, and we will prove that this contains a majority of possible input strings.

Claim 7.10. *Let \mathcal{G} be the set of strings $s \in \{0, 1\}^\ell$ such that*

$$\sum_{j=1}^t (|\{u \in P_j \mid d_s(u) < 2\}| + |Z_j - \mathbb{E}_s Z_j|) \leq 8n + 100t\sqrt{n \log n}.$$

Then $|\mathcal{G}| \geq 2^{\ell-1}$.

Proof. Consider s to be a random string, chosen uniformly on $\{0, 1\}^\ell$. We will prove that $\mathbb{P}[s \in \mathcal{G}] \geq 1/2$. We do this by considering the two bad events

$$\begin{aligned} \sum_{j=1}^t |\{u \in P_j \mid d_s(u) < 2\}| &> 8n, \\ \sum_{j=1}^t |Z_j - \mathbb{E} Z_j| &> 100t\sqrt{n \log n}, \end{aligned}$$

and prove that neither happens with probability more than $1/4$.

To bound the probability of the first event consider the expectation of the sum:

$$\begin{aligned} \mathbb{E} \sum_{j=1}^t |\{u \in P_j \mid d_s(u) < 2\}| &= \mathbb{E} \sum_{j=1}^t \sum_{u \in P_j} \mathbb{1}(d_s(u) < 2) = \sum_{j=1}^t \sum_{u \in P_j} \mathbb{P}[d_s(u) < 2] \\ &= \sum_{u \in P} \sum_{j: u \in P_j} \mathbb{P}[d_s(u) < 2] = \sum_{u \in P} |\{j \mid u \in P_j\}| \cdot \mathbb{P}[d_s(u) < 2] \\ &= \sum_{u \in P} d(u) \cdot (d(u) + 1) \cdot 2^{-d(u)} \leq 2n, \end{aligned}$$

as the function $d(d+1) \cdot 2^{-d}$ is bounded by 2 for all non-negative d .

This means, that by the Markov inequality

$$\mathbb{P} \left[\sum_{j=1}^t |\{u \in P_j \mid d_s(u) < 2\}| > 8n \right] \leq \frac{1}{4}.$$

Now, for the second bad event, we apply Chernoff bound (Theorem A.1). Note that

$$Z_j = |\{u \in P \setminus P_j \mid \Gamma_s(u) \not\subseteq Q_j\}|,$$

is the sum of $n - m$ independent random variables bounded by one. Therefore,

$$\mathbb{P} [|Z_j - \mathbb{E}Z_j| > \delta n] \leq 2 \exp \left(-\frac{\delta^2 n}{3} \right).$$

Setting δ to $100\sqrt{(\log n)/n}$ and taking union bound over $j = 1, \dots, t$ gives us that

$$\mathbb{P} \left[\sum_{j=1}^t |Z_j - \mathbb{E}Z_j| > 100t\sqrt{n \log n} \right] \leq \frac{1}{4}.$$

Putting the bounds on the first and second event together gives us the statement of the claim. \square

This \mathcal{G} will be our set of good inputs in our $(\ell, k, 1/10, 1/2)$ -SCS. Claim 7.10 essentially shows that the error in our estimate of $|s \cap q|$ would be at most $8n + 100t\sqrt{n \log n}$ without the inaccuracy introduced by our cut sparsifier. Since the size of the cut (S, \bar{S}) is at most n (the total number of hyperedges in the hypergraph H_s), this introduces an additional εn additive error.

Formally, when $s \in \mathcal{G}$

$$\begin{aligned} & \left| |s \cap q| - \text{DECODE}(\text{ENCODE}(s), q) \right| \\ &= \left| |s \cap q| - \sum_{j=1}^t \text{DECODE}^j(\text{SPARSIFY}(H_s), q^j) \right| \\ &\leq \sum_{j=1}^t \left| |s \cap q^j| - \text{DECODE}^j(\text{SPARSIFY}(H_s), q^j) \right| \\ &\leq \sum_{j=1}^t \left| |s \cap q^j| - \text{CUT}(\text{SPARSIFY}(H_s), S) + (a - |q^j|) + \mathbb{E}_s Z_j \right| && \text{by equation 13} \\ &\leq \sum_{j=1}^t \left(\left| |s \cap q^j| - |E_s^H(S_s^j, \bar{S}_s^j)| \right| + (a - |q^j|) + \mathbb{E}_s Z_j + \left| |E_s^H(S_s^j, \bar{S}_s^j)| - \text{CUT}(\text{SPARSIFY}(H_s), S_s^j) \right| \right) \\ &\leq \sum_{j=1}^t \left(|\{u \in P_j \mid d_s(u) < 2\}| + |Z_j - \mathbb{E}_s Z_j| + \varepsilon n \right) && \text{by equations 11 and 12} \\ &\leq \left(8n + 100t\sqrt{n \log n} \right) + \sum_{j=1}^t \varepsilon n && \text{since } s \in \mathcal{G} \\ &\leq 8n + 100t\sqrt{n \log n} + \varepsilon t n \end{aligned}$$

This is less than $\ell/20 = at/20$ due to the theorem's assumptions on the parameters. Therefore, (ENCODE, DECODE) is a $(\ell, k, 1/10, 1/2)$ -SCS with the set of good inputs being \mathcal{G} . By Corollary 7.7 k must be at least $\Omega(\ell) = \Omega(at)$. \square

We can now apply Theorem 7.9 to several RS-graph constructions known in the literature. Note that if there exists a (t, a) -RS graph, one can always reduce the parameters to get an (t', a') -RS graph for $t' \leq t$ and $a' \leq a$. We begin with Fischer et al. [FLN⁺02], which proves the existence of $(n^{\Omega(1/\log \log n)}, n/3 - o(1))$ -Ruzsa-Szemerédi graphs, resulting in the following corollary.

Corollary 7.11. *Any (n, r, k, ε) -HCSS with $r = n^{O(1/\log \log n)}$ and small constant ε requires $k = \Omega(nr)$ space.*

In other words, any data structure that can provide a $(1 + \varepsilon)$ -approximation to the size of all cuts in an r -uniform hypergraph with n vertices and $r = n^{O(1/\log \log n)}$ for small constant $\varepsilon \in (0, 1)$ requires $\Omega(nr)$ bits of space. This is tight due to the hypergraph cut sparsifier construction of [CKN20]. A different construction, also from [FLN⁺02], is able to achieve an $(n^c, n/O(\sqrt{\log \log n / \log n}))$ -RS graph for some small enough constant c . This results in the following:

Corollary 7.12. *For some constant c , any (n, r, k, ε) -HCSS with $r = O(n^c)$ and $\varepsilon = O(\sqrt{\log \log n / \log n})$ requires $k = \Omega(nr/\sqrt{\log n / \log \log n})$ space.*

Finally, the original construction of Ruzsa and Szemerédi [RS78] guarantees the existence of an $(n/3, n/2^{O(\sqrt{\log n})})$ -RS graphs, implying:

Corollary 7.13. *Any (n, r, k, ε) -HCSS with $\varepsilon = 2^{-\Omega(\sqrt{\log n})}$ requires $k = nr/2^{O(\sqrt{\log n})}$ space.*

These results imply that for any value of r , it is impossible to compress the cut structure of a hypergraph with n vertices and maximum hyperedge size r , with significantly less than nr space, and a polynomial scaling in the error (that is with $nr^{1-\Omega(1)}\varepsilon^{-O(1)}$ space).

8 Spectral Sparsification of Directed Hypergraphs

In this section, we discuss spectral sparsification of directed hypergraphs. First we introduce some notions and study basic properties of directed hypergraphs in Section 8.1. Then, we discuss spectrally sparsifying directed hypergraphs with hyperedges having nearly equal overlap (a concept to be defined in Section 8.1). Finally, we prove Theorem 1.3 in Section 8.4.

8.1 Preliminaries

A *directed hypergraph* $G = (V, E)$ is a pair of a vertex set V and a set E of hyperarcs, where a *hyperarc* $e \in E$ is an ordered pair of two disjoint vertex sets $h(e) \subseteq V$, the *head*, and $t(e) \subseteq V$, the *tail*. The *size* of a hyperarc $e \in E$ is $|h(e)| + |t(e)|$. We restrict ourselves to dealing with only *simple* directed hypergraphs, that is, in Section 8 E is always considered to be a set as opposed to a multiset.

We say that a vertex set $S \subseteq V$ *cuts* a hyperarc $e \in E$ if $S \cap t(e) \neq \emptyset$ and $(V \setminus S) \cap h(e) \neq \emptyset$. The *energy* of a hyperarc e with respect to a vector $x \in \mathbb{R}^V$ is defined as

$$\max_{a \in t(e), b \in h(e)} (x_a - x_b)_+^2,$$

where $(\alpha)_+ = \max\{\alpha, 0\}$. The energies of a set of arcs, or of an entire vector with respect to G , is defined identically to the undirected case. So in particular the energy of x with respect to G is

$$Q(x) = \sum_{e \in E} \max_{a \in t(e), b \in h(e)} (x_a - x_b)_+^2.$$

Note that $Q(1_S)$, where $1_S \in \mathbb{R}^V$ is the characteristic vector of S , is equal to the number of hyperarcs cut by S . Identically to Definition 2.2, for $\epsilon > 0$, a weighted subgraph \tilde{G} of G is said to be a ϵ -spectral sparsifier of G if

$$\tilde{Q}(x) = (1 \pm \epsilon)Q(x),$$

where $Q(x)$ and $\tilde{Q}(x)$ are energy of x with respect to G and \tilde{G} , respectively.

In constructing our sparsifier, a useful object to consider will be the *clique graph* of G , the directed (ordinary) multigraph we get by replacing each hyperarc in G with a directed bipartite clique. Formally, the clique of a hyperarc $e \in E$ is the set of arcs $C(e) = \{(a, b) \mid a \in t(e), b \in h(e)\}$. The clique graph of a set of hyperarcs $E' \subseteq E$ is the multi-union of the individual cliques $C(E') = \biguplus_{e \in E'} C(e)$. Finally, the clique graph of G itself is $C(G) = (V, C(E))$. In the following, we make some observation about the multiplicities of arcs in the clique graph.

Definition 8.1. *Given a hypergraph $G = (V, E)$, we say that a subset of hyperarcs $E' \subseteq E$ k -overlapping if every arc in $C(E')$ appears with multiplicity at least k . Furthermore, the overlap $k(e)$ of a single hyperarc $e \in E$ is defined as the largest k such that there exists a k -overlapping set of hyperarcs containing e .*

Informally, we will use the inverse overlap of each hyperarc as a sampling rate in constructing our sparsifier. Thus, the following lemma will be a useful bound on the sum of these rates:

Lemma 8.2. *Let $G = (V, E)$ be a directed hypergraph. Then, we have*

$$\sum_{e \in E} \frac{1}{k(e)} \leq n^2.$$

Proof. Consider the following simple algorithm:

Algorithm 3

```

1: procedure OVERLAPPEELING( $G = (V, E)$ )
2:    $E' \leftarrow E$ .
3:   for  $k = 1, \dots, 2^{n-2}$  do
4:      $E'_k \leftarrow E'$ .
5:     while there exists  $(u, v) \in C(E')$  with multiplicity at most  $k$  do
6:       for all hyperarcs  $e \in E'$  such that  $(u, v) \in C(e)$  do
7:          $f(e) \leftarrow (u, v)$ .
8:          $E' \leftarrow E' \setminus \{e\}$ .

```

This algorithm iterates through all possible overlaps (from 1 to 2^{n-2}) and peels off all hyperarcs with this overlap, until no hyperarcs remain. The algorithm maintains several variables (E'_k and $f(e)$) that are not used. However, these will be useful in proving the lemma.

Claim 8.3. *The set E'_k has overlap k for all k .*

Proof. Indeed, the variable k is augmented in the for-loop at Line 3 only after exiting the while-loop at Line 5. This means that there no longer existed any pairs (u, v) in $C(E')$ with multiplicity at most $k - 1$, and therefore E' was k -overlapping. (The exception to this argument is $k = 1$, however all sets are 1-overlapping by definition.) \square

Claim 8.4. *If a hyperarc e is removed at a time when $k = k^*$, then it has overlap exactly k^* .*

Proof. It is easy to see that e has overlap *at least* k^* , since it was an element of E'_{k^*} which is itself k^* -overlapping by Claim 8.1.

We prove that e has overlap *at most* k^* by induction. By induction, we can assume that all hyperarcs removed before e had overlap corresponding to the value of k at the time, that is, at most k^* . Let E^* be the current value of E' at the time just before e is removed. Suppose for contradiction that e is at least $(k^* + 1)$ -overlapping, or equivalently there exists a $(k^* + 1)$ -overlapping set containing e , say \tilde{E}_{k^*+1} . However, no hyperarc removed before e could be in this set, since we know they are at most k^* -overlapping. So $\tilde{E}_{k^*+1} \subseteq E^*$. But some arc in $C(e)$ has multiplicity only at most k^* in E^* , which is a contradiction. \square

Claim 8.5. *For any pair $(u, v) \in V^2$, we have*

$$\sum_{e: f(e)=(u,v)} \frac{1}{k(e)} \leq 1.$$

Proof. First note that all pairs (u, v) are only considered once in the while-loop of Line 5 throughout the whole algorithm. Indeed, once a pair is considered, all hyperarcs containing it are removed and (u, v) is no longer in $C(E')$.

Suppose (u, v) is removed in this way when $k = k^*$. Then all hyperarcs e such that $f(e) = (u, v)$ have overlap at most k^* . On the other hand, there are at most k^* such hyperarcs due to the condition in Line 5. This concludes the proof of the claim. \square

From here the lemma statement follows simply:

$$\sum_{e \in E} \frac{1}{k(e)} = \sum_{(u,v) \in V^2} \sum_{e: f(e)=(u,v)} \frac{1}{k(e)} \leq \sum_{(u,v) \in V^2} 1 = n^2. \quad \square$$

Remark 8.6. *Note that, by Claim 8.4, we can compute overlaps of hyperarcs by running Algorithm 3. Furthermore, we can make it run in polynomial time by, instead of incrementing k at Line 3, updating k to be the smallest multiplicity of an edge in $C(E')$.*

8.2 Nearly Equally Overlapping Directed Hypergraphs

In this section, we consider the simpler case where every hyperarc has a similar overlap.

Lemma 8.7. *There is an algorithm that, given $0 < \epsilon \leq 1/2$ and a directed hypergraph $G = (V, E)$ such that every hyperarc has overlap between k and $2k$ for some $k \geq 1$, and each hyperarc has size at most $r \leq \sqrt{\epsilon n}/11$, outputs in polynomial time a weighted subgraph $\tilde{G} = (V, \tilde{E}, w)$ of G satisfying the following with probability $1 - O(1/n)$:*

- \tilde{G} is an ϵ -spectral sparsifier of G ,
- $|\tilde{E}| = O(n^2 r^2 \log n / \epsilon^2)$.

Construction Let us construct $\tilde{G} = (V, \tilde{E})$ by sampling each hyperarc independently with the same probability $p = 1000r^2 \log n / (k\epsilon^2)$ and scaling them up by $1/p$. Let the weight of each hyperarc e in \tilde{G} be denoted as w_e . Then w_e is an independent random variable taking value $1/p$ with probability p and value 0 otherwise, for each e .

Clearly, we can compute the output in $O(m)$ time. Also, we can bound the size of \tilde{E} easily:

Lemma 8.8. *We have $\mathbb{E}[|\tilde{E}|] = 2000n^2r^2 \log n / \epsilon^2$ and*

$$\mathbb{P}\left[|\tilde{E}| > 4000n^2r^2 \log n / \epsilon^2\right] \leq 2 \exp\left(-\frac{2pkn^2}{3}\right).$$

Proof. Note that since the overlap of each hyperarc is at most $2k$, there are at most $2kn^2$ hyperarcs in total (in E) by Lemma 8.2. Each hyperarc is sampled with probability p to be in \tilde{E} , so $\mathbb{E}[|\tilde{E}|] = 2pkn^2 = 2000n^2r^2 \log n / \epsilon^2$, as claimed. By Chernoff bounds (Theorem A.1), the claimed concentration inequality holds. \square

Correctness We now examine the spectral properties of \tilde{G} . Recall that $C(G)$ is the clique graph of G . Let us denote by Q^C the energy with respect to the clique graph. We may assume without loss of generality that $Q^C(x) = 1$, since whether $Q(x) = (1 \pm \epsilon)\tilde{Q}(x)$ holds or not is unaffected by scaling x . Define \mathbb{R}^V to be the set of vectors x such that this is satisfied. Note that this means that $Q_x(E) \geq 1/r^2$. Indeed

$$\begin{aligned} Q_x(E) &= \sum_{e \in E} \max_{u \in t(e), v \in h(e)} (x_u - x_v)_+^2 = \sum_{e \in E} \max_{f \in C(e)} Q_x^C(f) \geq \frac{1}{r^2} \sum_{e \in E} \sum_{f \in C(E)} Q_x^C(f) \\ &= \frac{1}{r^2} \sum_{f \in C(E)} Q_x^C(f) = \frac{Q^C(x)}{r^2} = \frac{1}{r^2}. \end{aligned}$$

Let us categorize the arcs in $C(E)$ based on their contributions to the total energy $Q^C(x) = 1$ in $C(G)$. The categories are

$$C_i = \left\{ f \in C(E) \mid Q_x^C(f) \in \left(\frac{2^{-i}}{k}, \frac{2^{-i+1}}{k} \right] \right\},$$

for $i = 1, \dots, i^*$ where $i^* := \lceil 3 \log n \rceil$, as well as

$$C_* = \left\{ f \in C(E) \mid Q_x^C(f) \leq \frac{2^{-i^*}}{k} \right\}.$$

Recall that $C(E)$ is a multiset, and consequently so are C_i and C_* . Since each arc f appears with multiplicity at least k , any single arc can contribute at most $1/k$ to the energy. Therefore, all arcs of $C(G)$ are covered by these categories.

We then partition the hyperarcs into similar categories: A hyperarc e gets into category i (or E_i) if i is the smallest number for which $C(e)$ contains an arc in C_i . Formally

$$\begin{aligned} E_i &= \{e \in E \mid i = \max\{j \mid C(e) \cap C_j \neq \emptyset\}\} \quad (i = 1, \dots, i^*), \text{ and} \\ E_* &= \{e \in E \mid C(e) \subseteq C_*\}. \end{aligned}$$

To prove that \tilde{G} is an ϵ -spectral sparsifier, we will show that, for all i , $Q_x(E_i) \approx \tilde{Q}_x(E_i)$. Similarly to the proof of Theorem 4.1 we will introduce a discretization of $Q_x(E_i)$. However, unlike

in the proof of Theorem 4.1, instead of rounding the vertex potentials x_v , we will round the energies of hyperarcs, that is, $Q_x(e)$ for $e \in E$.

Let us first define $Q_x^{C,(i)}(f)$, the rounding of $Q_x^C(f)$. Firstly, if $Q_x^C(f) \leq 2^{-i}/k$, that is the arc f is not relevant to E_i , we define $Q_x^{C,(i)}(f)$ to be zero. Otherwise, let $Q_x^{C,(i)}(f)$ be the rounding of $Q_x^C(f)$ to the nearest integer multiple of $1/(kn^3)$. Analogously with the definition of Q_x , for $e \in E$ let

$$Q_x^{(i)}(e) = \max_{f \in C(e)} Q_x^{C,(i)}(f), \quad Q_x^{(i)}(E') = \sum_{e \in E'} Q_x^{(i)}(e),$$

$$\tilde{Q}_x^{(i)}(e) = w_e Q_x^{(i)}(e), \quad \tilde{Q}_x^{(i)}(E') = \sum_{e \in E'} \tilde{Q}_x^{(i)}(e).$$

Informally, we prove the following chain of approximations for each i :

$$Q_x(E_i) \cong Q_x^{(i)}(E_i) \cong \tilde{Q}_x^{(i)}(E_i) \cong \tilde{Q}_x(E_i),$$

as well as

$$Q_x(E_*) \cong \tilde{Q}_x(E_*).$$

We make this formal in the following claims:

Claim 8.9. *For all $x \in \mathbb{R}^V$ and all $i = 1, \dots, i^*$,*

$$Q_x^{(i)}(E_i) = Q_x(E_i) \pm \frac{2}{n}.$$

Claim 8.10. *For all $i = 1, \dots, i^*$,*

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x^{(i)}(E_i) = \left(1 \pm \frac{\epsilon}{2}\right) Q_x^{(i)}(E_i) \pm \frac{\epsilon Q(x)}{10 \log n} \right] \geq 1 - \frac{1}{n}.$$

Claim 8.11. *For all $i = 1, \dots, i^*$,*

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x^{(i)}(E_i) = \tilde{Q}_x(E_i) \pm \frac{4}{n} \right] \geq 1 - \frac{1}{n}.$$

Claim 8.12.

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x(E_*) = Q_x(E_*) \pm \frac{6}{n} \right] \geq 1 - \frac{1}{n}.$$

Before proving the above claims, which we do in the next section, we conclude the analysis of correctness of the sparsifier.

Lemma 8.13. *The directed hypergraph \tilde{G} is an ϵ -spectral sparsifier of G with probability $1 - O(1/n)$.*

Proof. The statements of Claims 8.10, 8.11, and 8.12 all hold with high probability. Let us consider the event that they all hold simultaneously, then by Claims 8.9, 8.10, and 8.11,

$$\begin{aligned} |Q_x(E_i) - \tilde{Q}_x(E_i)| &\leq |Q_x(E_i) - Q_x^{(i)}(E_i)| + |Q_x^{(i)}(E_i) - \tilde{Q}_x^{(i)}(E_i)| + |\tilde{Q}_x^{(i)}(E_i) - \tilde{Q}_x(E_i)| \\ &\leq \frac{2}{n} + \frac{\epsilon}{2} Q_x^{(i)}(E_i) + \frac{Q(x)}{10 \log n} + \frac{4}{n} \leq \frac{\epsilon}{2} Q_x(E_i) + \frac{Q(x)}{10 \log n} + \frac{6}{n}, \end{aligned}$$

using that $\epsilon \leq 1$. Summing this over $i = 1, \dots, i^* = \lceil 3 \log n \rceil$ and adding Claim 8.12 we get

$$\begin{aligned} |Q(x) - \tilde{Q}(x)| &\leq \sum_{i=1}^{i^*} |Q_x(E_i) - \tilde{Q}_x(E_i)| + |Q_x(E_*) - \tilde{Q}_x(E_*)| \\ &\leq \sum_{i=1}^{i^*} \left[\frac{\epsilon}{2} Q_x(E_i) + \frac{\epsilon Q(x)}{10 \log n} + \frac{6}{n} \right] + \frac{6}{n} \leq \frac{\epsilon}{2} Q(x) + \frac{4\epsilon}{10} Q(x) + \frac{12}{n} \leq \epsilon Q(x), \end{aligned}$$

since $\epsilon Q(x)/10 \geq \epsilon/(10r^2) \geq 12/n$ because $11r \leq \sqrt{\epsilon n}$. \square

Lemma 8.7 follows by Lemmas 8.8 and 8.13 and a union bound.

8.3 Proofs of Claims 8.9, 8.10, 8.11, and 8.12

We begin with a preliminary claim examining the difference between Q_x and $Q_x^{(i)}$ on a single hyperarc.

Claim 8.14. *For all $x \in \mathbb{R}^V$, all $i = 1, \dots, i^*$, and any hyperarc $e \in E_i$,*

$$Q_x^{(i)}(e) = Q_x(e) \pm \frac{1}{kn^3}.$$

Proof. Suppose first that $Q_x(e) \geq Q_x^{(i)}(e)$. Recall that $e \in E_i$ and by definition of E_i there exist arcs in $C(e) \cap C_i$. In this case let $f = \operatorname{argmax}_{f \in C(e)} Q_x^C(f)$, guaranteeing that $f \in C_i$. Therefore, by definition $Q_x^{C,(i)}(f)$ is not zero, but a rounding to the nearest integer multiple of $1/(kn^3)$. Therefore,

$$Q_x(e) - Q_x^{(i)}(e) \leq Q_x^C(f) - Q_x^{C,(i)}(f) \leq \frac{1}{kn^3}.$$

Now suppose that $Q_x(e) < Q_x^{(i)}(e)$. In this case we define f to be $\operatorname{argmax}_{f \in C(e)} Q_x^{C,(i)}(f)$. Now if $Q_x^{C,(i)}(f)$ is zero the claim holds trivially, so we may assume that $Q_x^{C,(i)}(f)$ is instead a rounding to the nearest integer multiple of $1/(kn^3)$:

$$Q_x^{(i)}(e) - Q_x(e) \leq Q_x^{C,(i)}(f) - Q_x^C(f) \leq \frac{1}{kn^3}. \quad \square$$

Claim 8.9. *For all $x \in \mathbb{R}^V$ and all $i = 1, \dots, i^*$,*

$$Q_x^{(i)}(E_i) = Q_x(E_i) \pm \frac{2}{n}.$$

Proof. We can simply sum over the hyperarcs in E_i . Since $|E_i| \leq |E| \leq 2kn^2$, we have that

$$|Q_x(E_i) - Q_x^{(i)}(E_i)| \leq \sum_{e \in E_i} |Q_x(e) - Q_x^{(i)}(e)| \leq \frac{|E_i|}{kn^3} \leq \frac{2}{n}. \quad \square$$

Claim 8.10. *For all $i = 1, \dots, i^*$,*

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x^{(i)}(E_i) = \left(1 \pm \frac{\epsilon}{2} \right) Q_x^{(i)}(E_i) \pm \frac{\epsilon Q(x)}{10 \log n} \right] \geq 1 - \frac{1}{n}.$$

Proof. We prove the stronger claim

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x^{(i)}(E_i) = \left(1 \pm \frac{\epsilon}{2}\right) Q_x^{(i)}(E_i) \pm \frac{\epsilon}{10r^2 \log n} \right] \geq 1 - \frac{1}{n},$$

replacing $Q(x)$ by $1/r^2$ in the allowable additive error, which depends on x only through $Q_x^{(i)}$ and E_i .

We first consider a single setting of x (and consequently E_i and $Q_x^{(i)}$). Since $\mathbb{E}[\tilde{Q}_x^{(i)}(e)] = Q_x^{(i)}(e)$, we can apply additive-multiplicative Chernoff (Theorem A.2) to get the desired bound. Each independent random variable ($\tilde{Q}_x^{(i)}(e)$ for $e \in E_i$) is in the range $[0, 2^{-i+1}/(pk)]$ by definition of E_i . Therefore we get

$$\begin{aligned} \mathbb{P} \left[\left| \tilde{Q}_x^{(i)}(E_i) - Q_x^{(i)}(E_i) \right| > \frac{\epsilon}{2} Q_x^{(i)}(E_i) + \frac{\epsilon}{10r^2 \log n} \right] &\leq 2 \exp \left(-\frac{\epsilon/2 \cdot \frac{\epsilon}{10r^2 \log n}}{3 \cdot 2^{-i+1}/(pk)} \right) \\ &= 2 \exp \left(-\frac{\epsilon^2 pk \cdot 2^i}{120r^2 \log n} \right). \end{aligned}$$

We will now use a union bound to prove that this holds simultaneously for all possible settings of E_i and $Q_x^{(i)}$. Recall that by definition $\bigcup_{j=1}^i C_j$ contains exactly arcs of $C(E)$ that contribute more than $2^{-i}/k$ energy to the total energy of $Q^C(x) = 1$. There are at most $k \cdot 2^i$ such arcs, but since each arc appears with multiplicity at least k , there are at most 2^i *distinct* arcs. The $Q_x^{C, (i)}$ -value of all arcs not in $\bigcup_{j=1}^i C_j$ is zero. To select this multiset of non-zero valued arcs there are

$$\binom{n^2}{2^i} \leq n^{2 \cdot 2^i} = \exp(2 \cdot 2^i \log n)$$

options. Furthermore, for each relevant arc, we must choose its $Q_x^{C, (i)}$ -value: This is an integer multiple of $1/kn^3$ in the range $[-1/k, 1/k]$ and so there are $2n^3$ options per arc—for a total of

$$(2n^3)^{2^i} \leq \exp(4 \cdot 2^i \log n)$$

options. Finally, for each relevant arc, we must choose which category among E_1, \dots, E_i it belongs to (as this may not be evident from the value of $Q_x^{C, (i)}$). This is an additional $i \leq 3 \log n + 1$ options per arc—for a total of

$$(3 \log n + 1)^{2^i} \leq \exp(2^i \log n),$$

options among all arcs.

Ultimately, there are

$$\exp(2 \cdot 2^i \log n + 4 \cdot 2^i \log n + 2^i \log n) = \exp(7 \cdot 2^i \log n)$$

possible settings of $(E_1, \dots, E_i, Q_x^{C, (i)})$.

Combining the above Chernoff bound for a single setting of x with this union bound we get the statement of the claim:

$$\begin{aligned} &\mathbb{P} \left[\forall x : \left| Q_x^{(i)}(E_i) - \tilde{Q}_x^{(i)}(E_i) \right| > \frac{\epsilon}{2} Q_x^{(i)}(E_i) + \frac{\epsilon}{10r^2 \log n} \right] \\ &\leq 2 \exp(7 \cdot 2^i \log n) \cdot \exp \left(-\frac{\epsilon^2 pk \cdot 2^{-i}}{120r^2} \right) \end{aligned}$$

$$\begin{aligned}
&= 2 \exp \left(2^i \cdot \left(7 \log n - \frac{\epsilon^2 pk}{120r^2} \right) \right) \\
&\leq \frac{1}{n},
\end{aligned}$$

since $pk = 1000r^2 \log n / \epsilon^2$. □

Claim 8.11. For all $i = 1, \dots, i^*$,

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x^{(i)}(E_i) = \tilde{Q}_x(E_i) \pm \frac{4}{n} \right] \geq 1 - \frac{1}{n}.$$

Proof. We consider the high probability event that $|\tilde{E}| \leq 4pkn^2$. (Lemma 8.8). Similarly to the proof of Claim 8.9 we simply sum over all edges of E_i . Note that if $e \in \tilde{E}$,

$$|\tilde{Q}_x(e) - \tilde{Q}_x^{(i)}(e)| \leq \sum_{e \in E_i} |\tilde{Q}_x(e) - \tilde{Q}_x^{(i)}(e)| = \sum_{e \in E_i \cap \tilde{E}} \frac{1}{p} |Q_x(e) - Q_x^{(i)}(e)| \leq \frac{|\tilde{E}|}{pkn^3} \leq \frac{4}{n}. \quad \square$$

Claim 8.12.

$$\mathbb{P} \left[\forall x \in \overline{\mathbb{R}^V} : \tilde{Q}_x(E_*) = Q_x(E_*) \pm \frac{6}{n} \right] \geq 1 - \frac{1}{n}.$$

Proof. Note that

$$|Q_x(E_*) - \tilde{Q}_x(E_*)| \leq Q_x(E_*) + \tilde{Q}_x(E_*).$$

We bound the two terms separately:

$$Q_x(E_*) \leq |E_*| \cdot \frac{1}{kn^3} \leq |E| \cdot \frac{1}{kn^3} \leq \frac{2}{n},$$

and

$$\tilde{Q}_x(E_*) \leq |\tilde{E}| \cdot \frac{1}{pkn^3} \leq \frac{4}{n},$$

with high probability by Lemma 8.8. □

8.4 Proof of Theorem 1.3

Proof of Theorem 1.3. Given the results of Lemma 8.7, we only need to decompose G into directed hypergraphs with their hyperedges having nearly the same overlap. We will repeatedly separate and sparsify hyperarcs of the highest overlap until no hyperarcs remain. This results in an ϵ -spectral sparsifier of G , since the quality of being an ϵ -spectral sparsifier is additive.

Consider the following simple algorithm, where UNIFORMSAMPLINGSPARSIFY denotes the sparsification algorithm given in Lemma 8.7:

Algorithm 4 Directed hypergraph sparsification

- 1: **procedure** SPARSIFY($G = (V, E)$)
 - 2: $\tilde{E} \leftarrow \emptyset$.
 - 3: **while** $E \neq \emptyset$ **do**
 - 4: $2k \leftarrow$ the largest overlap among hyperedges in E .
 - 5: $E' \leftarrow$ the maximal k -overlapping set in E .
 - 6: $E \leftarrow E \setminus E'$.
 - 7: $\tilde{E} \leftarrow \tilde{E} \cup \text{UNIFORMSAMPLINGSPARSIFY}(V, E')$.
 - 8: **return** (V, \tilde{E}) .
-

Note first that the maximal set of a certain multiplicity (in Line 5) is indeed unique. It follows from definition that the union of hyperarc sets of multiplicity k still has multiplicity k . Therefore, E' contains all hyperarcs of overlap at least k from (the current) E . Furthermore, removing E' from E reduces the maximum overlap of any hyperarc to below k , so by a factor of at least 2. Since the maximum overlap started out is at most n^{r-2} , Algorithm 8.4 terminates in at most $r \log n$ iterations. Since the size of \tilde{E} increased by at most $O(n^2 r^2 \log n / \epsilon^2)$ in each iteration, by Lemma 8.7, this results in $|\tilde{E}| = O(n^2 r^3 \log^2 n / \epsilon^2)$, as claimed.

The running time is polynomial because we can compute overlaps of hyperarcs in polynomial time by Remark 8.6, and hence can compute the largest overlap k and the maximal k -overlapping set in polynomial time. \square

Acknowledgement

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 759471).

References

- [ACK⁺16] Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In *Proceedings of the 2016 Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 311–319, 2016. doi:[10.1145/2840728.2840753](https://doi.org/10.1145/2840728.2840753).
- [Alo86] Noga Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. doi:[10.1007/bf02579166](https://doi.org/10.1007/bf02579166).
- [AM85] N Alon and V.D Milman. λ_1 , isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory, Series B*, 38(1):73–88, 1985. doi:[10.1016/0095-8956\(85\)90092-9](https://doi.org/10.1016/0095-8956(85)90092-9).
- [AR98] Y. Aumann and Y. Rabani. An $O(\log k)$ approximate min-cut max-flow theorem and approximation algorithm. *SIAM Journal on Computing*, 27(1):291–301, 1998. doi:[10.1137/S0097539794285983](https://doi.org/10.1137/S0097539794285983).
- [AS08] Noga Alon and Joel H. Spencer. *The Probabilistic Method, Third Edition*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 2008.
- [BK15] Andrs A. Benczr and David R. Karger. Randomized approximation schemes for cuts and flows in capacitated graphs. *SIAM Journal on Computing*, 44(2):290–319, 2015. doi:[10.1137/070705970](https://doi.org/10.1137/070705970).
- [Bou85] J. Bourgain. On Lipschitz embedding of finite metric spaces in Hilbert space. *Israel Journal of Mathematics*, 52(1-2):46–52, 1985. doi:[10.1007/bf02776078](https://doi.org/10.1007/bf02776078).
- [BST19] Nikhil Bansal, Ola Svensson, and Luca Trevisan. New notions and constructions of sparsification for graphs and hypergraphs. In *Proceedings of the IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 910–928, 2019. doi:[10.1109/focs.2019.00059](https://doi.org/10.1109/focs.2019.00059).

- [BV13] Ashwinkumar Badanidiyuru and Jan Vondrák. Fast algorithms for maximizing submodular functions. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1497–1514, 2013. doi:[10.1137/1.9781611973402.110](https://doi.org/10.1137/1.9781611973402.110).
- [CKN20] Yu Chen, Sanjeev Khanna, and Ansh Nagda. Near-linear size hypergraph cut sparsifiers. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020*, pages 61–72. IEEE, 2020. doi:[10.1109/FOCS46700.2020.00015](https://doi.org/10.1109/FOCS46700.2020.00015).
- [CKST19] Charles Carlson, Alexandra Kolla, Nikhil Srivastava, and Luca Trevisan. Optimal lower bounds for sketching graph cuts. In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2565–2569, 2019.
- [CLTZ18] T.-H. Hubert Chan, Anand Louis, Zhihao Gavin Tang, and Chenzi Zhang. Spectral properties of hypergraph Laplacian and approximation algorithms. *Journal of the ACM*, 65(3):1–48, 2018. doi:[10.1145/3178123](https://doi.org/10.1145/3178123).
- [CPS20] Yu Cheng, Debmalya Panigrahi, and Kevin Sun. Sparsification of balanced directed graphs. *CoRR*, abs/2006.01975, 2020. URL: <https://arxiv.org/abs/2006.01975>, arXiv:2006.01975.
- [DN03] Irit Dinur and Kobbi Nissim. Revealing information while preserving privacy. In *Proceedings of the 22nd ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS)*, pages 202–210, 2003. URL: <https://doi.org/10.1145/773153.773173>, doi:[10.1145/773153.773173](https://doi.org/10.1145/773153.773173).
- [FHL08] Uriel Feige, MohammadTaghi Hajiaghayi, and James R. Lee. Improved approximation algorithms for minimum weight vertex separators. *SIAM Journal on Computing*, 38(2):629–657, 2008. doi:[10.1137/05064299x](https://doi.org/10.1137/05064299x).
- [FLN⁺02] Eldar Fischer, Eric Lehman, Ilan Newman, Sofya Raskhodnikova, Ronitt Rubinfeld, and Alex Samorodnitsky. Monotonicity testing over general poset domains. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 474–483, 2002. doi:[10.1145/509907.509977](https://doi.org/10.1145/509907.509977).
- [GKR00] David Gibson, Jon Kleinberg, and Prabhakar Raghavan. Clustering categorical data: an approach based on dynamical systems. *The International Journal on Very Large Data Bases*, 8(3-4):222–236, 2000. doi:[10.1007/s007780050005](https://doi.org/10.1007/s007780050005).
- [HLM09] Yuchi Huang, Qingshan Liu, and Dimitris Metaxas. Video object segmentation by hypergraph cut. In *Proceedings of the 2009 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1738–1745, 2009. doi:[10.1109/cvpr.2009.5206795](https://doi.org/10.1109/cvpr.2009.5206795).
- [HSJR13] Matthias Hein, Simon Setzer, Leonardo Jost, and Syama Sundar Rangapuram. The total variation on hypergraphs - learning on hypergraphs revisited. In *Advances in Neural Information Processing Systems 26 (NIPS)*, pages 2427–2435, 2013.
- [IMTY18] Masahiro Ikeda, Atsushi Miyauchi, Yuuki Takai, and Yuichi Yoshida. Finding Cheeger cuts in hypergraphs via heat equation. *arXiv preprint arXiv:1809.04396*, 2018. URL: <https://arxiv.org/abs/1809.04396>.

- [IT18] Motoki Ikeda and Shin-ichi Tanigawa. Cut sparsifiers for balanced digraphs. In *Proceedings of the 16th International Workshop on Approximation and Online Algorithms (WAOA)*, pages 277–294, 2018. doi:[10.1007/978-3-030-04693-4_17](https://doi.org/10.1007/978-3-030-04693-4_17).
- [KHT09] Steffen Klamt, Utz-Uwe Haus, and Fabian Theis. Hypergraphs and cellular networks. *PLoS Computational Biology*, 5(5):e1000385, 2009. doi:[10.1371/journal.pcbi.1000385](https://doi.org/10.1371/journal.pcbi.1000385).
- [KK15] Dmitry Kogan and Robert Krauthgamer. Sketching cuts in graphs and hypergraphs. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 367–376, 2015. doi:[10.1145/2688073.2688093](https://doi.org/10.1145/2688073.2688093).
- [LLR95] Nathan Linial, Eran London, and Yuri Rabinovich. The geometry of graphs and some of its algorithmic applications. *Combinatorica*, 15(2):215–245, 1995. doi:[10.1007/bf01200757](https://doi.org/10.1007/bf01200757).
- [LM16] Anand Louis and Yury Makarychev. Approximation algorithms for hypergraph small-set expansion and small-set vertex expansion. *Theory of Computing*, 12(17):1–25, 2016. doi:[10.4086/toc.2016.v012a017](https://doi.org/10.4086/toc.2016.v012a017).
- [LM18] Pan Li and Olgica Milenkovic. Submodular hypergraphs: p -Laplacians, Cheeger inequalities and spectral clustering. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, volume 80, pages 3020–3029, 2018.
- [Lou15] Anand Louis. Hypergraph Markov operators, eigenvalues and approximation algorithms. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 713–722, 2015. doi:[10.1145/2746539.2746555](https://doi.org/10.1145/2746539.2746555).
- [LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM*, 46(6):787–832, 1999. doi:[10.1145/331524.331526](https://doi.org/10.1145/331524.331526).
- [LS18] Yin Tat Lee and He Sun. Constructing linear-sized spectral sparsification in almost-linear time. *SIAM Journal on Computing*, 47(6):2315–2336, 2018. doi:[10.1137/16m1061850](https://doi.org/10.1137/16m1061850).
- [MR01] J. Matoušek and Y. Rabinovich. On dominated l_1 metrics. *Israel Journal of Mathematics*, 123(1-2):285–301, 2001. doi:[10.1007/BF02784132](https://doi.org/10.1007/BF02784132).
- [NR13] Ilan Newman and Yuri Rabinovich. On multiplicative λ -approximations and some geometric applications. *SIAM Journal on Computing*, 42(3):855–883, 2013. doi:[10.1137/100801809](https://doi.org/10.1137/100801809).
- [OB12] P. Ochs and T. Brox. Higher order motion models and spectral clustering. In *Proceedings of the 2012 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 614–621, 2012. doi:[10.1109/cvpr.2012.6247728](https://doi.org/10.1109/cvpr.2012.6247728).
- [RS78] Imre Z Ruzsa and Endre Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Keszthely, 1976), Coll. Math. Soc. J. Bolyai*, 18:939–945, 1978.

- [RS10] Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *Proceedings of the 42nd Annual ACM Symposium on Theory of Computing (STOC)*, pages 755–764, 2010. doi:[10.1145/1806689.1806792](https://doi.org/10.1145/1806689.1806792).
- [RV07] Mark Rudelson and Roman Vershynin. Sampling from large matrices: An approach through geometric functional analysis. *J. ACM*, 54(4):21, 2007. URL: <https://doi.org/10.1145/1255443.1255449>, doi:[10.1145/1255443.1255449](https://doi.org/10.1145/1255443.1255449).
- [SS11] Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM Journal on Computing*, 40(6):1913–1926, 2011. doi:[10.1137/080734029](https://doi.org/10.1137/080734029).
- [ST11a] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM Journal on Computing*, 40(4):981–1025, 2011. doi:[10.1137/08074489x](https://doi.org/10.1137/08074489x).
- [ST11b] Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011. URL: <https://doi.org/10.1137/08074489X>, doi:[10.1137/08074489X](https://doi.org/10.1137/08074489X).
- [SY19] Tasuku Soma and Yuichi Yoshida. Spectral sparsification of hypergraphs. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2570–2581, 2019. doi:[10.1137/1.9781611975482.159](https://doi.org/10.1137/1.9781611975482.159).
- [TMIY20] Yuuki Takai, Atsushi Miyauchi, Masahiro Ikeda, and Yuichi Yoshida. Hypergraph clustering based on PageRank. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD)*, pages 1970–1978, 2020. doi:[10.1145/3394486.3403248](https://doi.org/10.1145/3394486.3403248).
- [Tro15] Joel A. Tropp. An introduction to matrix concentration inequalities. *Foundations and Trends® in Machine Learning*, 8(1-2):1–230, 2015. doi:[10.1561/22000000048](https://doi.org/10.1561/22000000048).
- [YNN⁺20] Naganand Yadati, Vikram Nitin, Madhav Nimishakavi, Prateek Yadav, Anand Louis, and Partha Talukdar. NHP: Neural hypergraph link prediction. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management (CIKM)*, pages 1705–1714, 2020. doi:[10.1145/3340531.3411870](https://doi.org/10.1145/3340531.3411870).
- [YNY⁺19] Naganand Yadati, Madhav Nimishakavi, Prateek Yadav, Vikram Nitin, Anand Louis, and Partha P. Talukdar. HyperGCN: A new method for training graph convolutional networks on hypergraphs. In *Advances in Neural Information Processing Systems 32*, pages 1509–1520, 2019.
- [Yos16] Yuichi Yoshida. Nonlinear Laplacian for digraphs and its applications to network analysis. In *Proceedings of the 9th ACM International Conference on Web Search and Data Mining (WSDM)*, pages 483–492, 2016. doi:[10.1145/2835776.2835785](https://doi.org/10.1145/2835776.2835785).
- [Yos19] Yuichi Yoshida. Cheeger inequalities for submodular transformations. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2582–2601, 2019. doi:[10.1137/1.9781611975482.160](https://doi.org/10.1137/1.9781611975482.160).
- [ZHTC20] Chenzi Zhang, Shuguang Hu, Zhihao Gavin Tang, and T-H. Hubert Chan. Re-revisiting learning on hypergraphs: Confidence interval, subgradient method, and extension to multiclass. *IEEE Transactions on Knowledge and Data Engineering*, 32(3):506–518, 2020. doi:[10.1109/tkde.2018.2880448](https://doi.org/10.1109/tkde.2018.2880448).

A Technical lemmas

A.1 Concentration Inequalities

The following concentration bound is standard.

Theorem A.1 (Chernoff bound, see e.g. [AS08]). *Let X_1, \dots, X_n be independent random variables in the range $[0, a]$. Let $\sum_{i=1}^n X_i = S$. Then for any $\delta \in [0, 1]$ and $\mu \geq \mathbb{E}S$,*

$$\mathbb{P}[|S - \mathbb{E}S| \geq \delta\mu] \leq 2 \exp\left(-\frac{\delta^2\mu}{3a}\right).$$

The following slight variation, allowing for both multiplicative and additive error, will be the most convenient for our purposes throughout the paper.

Theorem A.2 (Additive-multiplicative Chernoff bounds [BV13]). *Let X_1, \dots, X_n be independent random variables in the range $[0, a]$. Let $\sum_{i=1}^n X_i = S$. Then for all $\delta \in [0, 1]$ and $\alpha > 0$,*

$$\mathbb{P}[|S - \mathbb{E}S| \geq \delta\mathbb{E}S + \alpha] \leq 2 \exp\left(-\frac{\delta\alpha}{3a}\right).$$

A.2 Proof of Hypergraph Cheeger's Inequality

Proof of Theorem 2.3: Recall that $\sum_{v \in V} x_v d(v) = 0$. Suppose for contradiction that there exists a vector $x \in \mathbb{R}^V$ such that $Q(x) < \frac{r\Phi^2}{32} \sum_{v \in V} x_v^2 d(v)$. Let $\tilde{x} \in \mathbb{R}^V$ be such x shifted such that $\sum_{v \in V} x_v d_x(v) = 0$, where $d_x(v)$ denotes the degree of v in $G_x = (V, E_x)$. Then, we have

$$\begin{aligned} Q(x) &< \frac{r\Phi^2}{32} \sum_{v \in V} x_v^2 d(v) \leq \frac{r\Phi^2}{32} \sum_{v \in V} \tilde{x}_v^2 d(v) = \frac{r\Phi^2}{32} \sum_{e \in E} \sum_{v \in e} \tilde{x}_v^2 \\ &\leq \frac{r^2\Phi^2}{32} \sum_{e \in E} \max_{v \in e} \tilde{x}_v^2 \leq \frac{r^2\Phi^2}{32} \sum_{(a,b) \in E_x} (\tilde{x}_a^2 + \tilde{x}_b^2) = \frac{r^2\Phi^2}{32} \sum_{v \in V} \tilde{x}_v^2 d_x(v), \end{aligned}$$

The second inequality follows since x is centered, that is $\sum_{v \in V} x_v d(v) = 0$.

This means, by Cheeger's inequality for ordinary graphs [AM85, Alo86], that there exists a vertex set S of expansion $\frac{r\Phi}{4}$ in G_x . Moreover, S can be chosen to be a *sweep cut* with respect to x (regardless of the degree vector) in the sense that S is of the form $\{v \in V \mid x_v \leq \tau\}$ or $\{v \in V \mid x_v \geq \tau\}$ for some $\tau \in \mathbb{R}$. Let $S \subseteq V$ be the smaller side of the cut (in volume). Let $\eta := |E_x(S, V \setminus S)|$ and $\zeta := |E(S)|$. Then, we have

$$\eta \leq \frac{r\Phi}{4} \sum_{v \in S} d_x(v) = \frac{r\Phi}{4} (\eta + 2\zeta).$$

Since $\Phi \leq \frac{2}{r}$, $\frac{r\Phi}{4} \leq \frac{1}{2}$ and so $\zeta \geq \frac{\eta}{r\Phi}$. Since S is a sweep cut with respect to x , each edge of G_x crossing the cut $(S, V \setminus S)$ corresponds to a distinct hyperedge of G also crossing the cut, and each edge of G_x fully inside S translates to a hyperedge of G fully inside S . Therefore, the number of edges crossing the cut $(S, V \setminus S)$ in G is still η and $\sum_{v \in S} d(v) > r\zeta \geq \frac{\eta}{\Phi}$. Similarly, $\sum_{v \in V \setminus S} d(v) > \frac{\eta}{\Phi}$. Therefore, the expansion of the cut $(S, V \setminus S)$ in G is less than Φ , which is a contradiction. \square