



Additive Sparsification of CSPs

Eden Pelleg 

Mathematical Institute, University of Oxford, UK

Stanislav Živný  

Department of Computer Science, University of Oxford, UK

Abstract

Multiplicative cut sparsifiers, introduced by Benczúr and Karger [STOC'96], have proved extremely influential and found various applications. Precise characterisations were established for sparsifiability of graphs with other 2-variable predicates on Boolean domains by Filtser and Krauthgamer [SIDMA'17] and non-Boolean domains by Butti and Živný [SIDMA'20].

Bansal, Svensson and Trevisan [FOCS'19] introduced a weaker notion of sparsification termed “additive sparsification”, which does not require weights on the edges of the graph. In particular, Bansal et al. designed algorithms for additive sparsifiers for cuts in graphs and hypergraphs.

As our main result, we establish that *all* Boolean Constraint Satisfaction Problems (CSPs) admit an additive sparsifier; that is, for every Boolean predicate $P : \{0, 1\}^k \rightarrow \{0, 1\}$ of a fixed arity k , we show that $\text{CSP}(P)$ admits an additive sparsifier. Under our newly introduced notion of all-but-one sparsification for non-Boolean predicates, we show that $\text{CSP}(P)$ admits an additive sparsifier for *any* predicate $P : D^k \rightarrow \{0, 1\}$ of a fixed arity k on an arbitrary finite domain D .

2012 ACM Subject Classification Theory of computation \rightarrow Design and analysis of algorithms; Theory of computation \rightarrow Sparsification and spanners

Keywords and phrases additive sparsification, graphs, hypergraphs, minimum cuts

Digital Object Identifier 10.4230/LIPIcs.EISA.2021.75

Related Version *Full Version*: <https://arxiv.org/abs/2106.14757> [29]

Funding 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 714532). The paper reflects only the authors’ views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein. *Stanislav Živný*: Royal Society University Research Fellowship.

1 Introduction

Graph sparsification is the problem of, given a graph $G = (V, E)$ with quadratically many (in $|V|$) edges, finding a sparse subgraph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$ such that important properties of G are preserved in G_ε . Sparse in this context usually means with sub-quadratically many edges, though in this work we require (and can achieve) linearly many edges.

One of the most studied properties of preservation is the size of cuts. If $G = (V, E, w)$ is an undirected weighted graph with $w : E \rightarrow \mathbb{R}_{>0}$, given some $S \subseteq V$, the cut of S in G is

$$\text{Cut}_G(S) = \sum_{\substack{\{u,v\} \in E \\ |\{u,v\} \cap S| = 1}} w(\{u,v\}),$$

the sum of weights of all edges connecting S and $S^c = V \setminus S$. In an influential paper, Benczúr and Karger [11] introduced cut sparsification with a multiplicative error. In particular, [11] showed that for any graph $G = (V, E, w)$ and any error parameter $0 < \varepsilon < 1$, there exists a sparse subgraph $G_\varepsilon = (V, E_\varepsilon \subseteq E, w')$ with $O(n(\log n)\varepsilon^{-2})$ edges (and new weights w' on



© Eden Pelleg and Stanislav Živný;
licensed under Creative Commons License CC-BY 4.0
29th Annual European Symposium on Algorithms (ESA 2021).

Editors: Petra Mutzel, Rasmus Pagh, and Grzegorz Herman; Article No. 75; pp. 75:1–75:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

the edges in E_ε), such that for every $S \subseteq V$ we have

$$\text{Cut}_{G_\varepsilon}(S) \in (1 \pm \varepsilon)\text{Cut}_G(S).$$

This was later improved by Batson, Spielman and Srivastava [10] to a subgraph with $O(n\varepsilon^{-2})$ many edges. Andoni, Chen, Krauthgamer, Qin, Woodruff and Zhang showed that the dependency on ε is optimal [4].

The ideas from cut sparsification paved the way to various generalisations, including streaming [1], sketching [4], cuts in hypergraphs [23, 27], spectral sparsification [33, 32, 34, 21, 31] and the consideration of other predicates besides cuts [20]. In this work, we focus on the latter.

The cut sparsification result in [10] was explored for other Boolean binary predicates by Filtser and Krauthgamer [20], following a suggestion to do so by Kogan and Krauthgamer in [23]. Filtser and Krauthgamer found [20] a necessary and sufficient condition on the predicate for the graph to be sparsifiable (in the sense of [10]). In particular, [20] showed that not all Boolean binary predicates are sparsifiable. Later, Butti and Živný [14] generalised the result from [20] to arbitrary finite domain binary predicates.

We remark that [20, 14] use the terminology of *constraint satisfaction problems* (CSPs) with a fixed predicate P . This is equivalent to a (hyper)graph G with a fixed predicate. Indeed, the vertices of G correspond to the variables of the CSP and the (hyper)edges of G correspond to the constraints of the CSP. If the fixed predicate P is not symmetric, the (hyper)edges of G are directed. We will mostly talk about sparsification of (hyper)graphs with a fixed predicate but this is equivalent to the CSP view.

Recently, while trying to eliminate the requirement for the introduction of new weights for the sparse subgraph, Bansal, Svensson and Trevisan [7] have come up with a new sparsification notion with an additive error term. They have shown (cf. Theorem 3 in Section 2) that under their notion any undirected unweighted hypergraph has a sparse subhypergraph which preserves all cuts up to some additive term.

Motivation

The relatively recent notion of additive sparsification has not yet been explored to the same extent as the notion of multiplicative sparsification has been. We believe that this notion has a lot of potential for applications as the sparsifiers are not weighted, unlike multiplicative sparsifiers. Indeed, the main restriction of multiplicative sparsifiers in applications appears to be the number of distinct weights required in sparsifiers. For some graphs (such as the “barbell graph” – two disjoint cliques joined by a single edge), any nontrivial multiplicative sparsifier requires edges of different weights. In any case, the authors find the notion of additive sparsification interesting in its own right, independently of applications.

The goal of our work is to understand how the notion of additive sparsification developed in [7] for cuts behaves on (hyper)graphs with other predicates (beyond cuts), deriving inspiration from the generalisations of cuts to other predicates in the multiplicative setting established in [20, 14]. In particular, already Boolean binary predicates include interesting predicates such as the *uncut edges* (using the predicate $P(x, y) = 1$ iff $x = y$), *covered edges* (using the predicate $P(x, y) = 1$ iff $x = 1$ or $y = 1$), or *directed cut edges* (using the predicate $P(x, y) = 1$ iff $x = 0$ and $y = 1$). While such graph problems are well-known and extensively studied, it is not clear whether one should expect them to be sparsifiable or not. For instance, as mentioned before, not all (even Boolean binary) predicates are sparsifiable multiplicatively [20]. Are there some predicates that are not additively sparsifiable?

1.1 Contributions

Boolean predicates

Our main result, Theorem 9 in Section 3, shows that *all* hypergraphs with *constant* uniformity k , directed or undirected, admit additive sparsification with respect to all Boolean predicates $P : \{0, 1\}^k \rightarrow \{0, 1\}$; the number of hyperedges of the sparsifier with error $\varepsilon > 0$ is $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$, where the $O(\cdot)$ hides a factor that depends on k . This result has three ingredients. First, we observe that the result in [7] also holds true for directed hypergraphs. Second, we use a reduction via the k -partite k -fold covers of hypergraphs to the already solved case of Boolean Cut. Finally, we use linear algebra to prove the correctness of the reduction. While the reduction via the k -partite k -fold cover was used in previous works on multiplicative sparsification [20, 14], the subsequent non-trivial linear-algebraic analysis (Proposition 13) is novel and constitutes our main technical contribution, as well as our result that, unlike in the multiplicative setting, all Boolean predicates can be (additively) sparsified. We also show that our results immediately apply to the more general setting where different hyperedges are associated with different predicates (cf. Remark 10). This corresponds to CSPs with a fixed constraint language (of a finite size) rather than just a single predicate.

Non-Boolean predicates

We introduce a notion of sparsification that generalises the Boolean case to predicates on non-Boolean domains, i.e. a notion capturing predicates of the form $P : D^k \rightarrow \{0, 1\}$, where D is an arbitrary fixed finite set with $|D| \geq 2$. We call this type of sparsification “all-but-one” sparsification since the additive error term includes the maximum volume of $|D| - 1$ (out of $|D|$) parts, where the volume of a subset is the sum of the degrees in the subset. (The precise definition can be found in Section 4.) By building on the techniques used to establish our main result, we show that all hypergraphs (again, directed or undirected) admit additive all-but-one sparsification with respect to all predicates. This is stated as Theorem 20 in Section 4. We also show, in Section 5, that our notion of all-but-one sparsification is, in some sense, optimal.

Comparison to previous work

As mentioned above, our sparsifiability result is obtained by a reduction via the k -partite k -fold cover to the cut case established in [7]. A reduction via the k -partite k -fold cover was also used (for $k = 2$) in previous work on multiplicative sparsification [20, 14]. In particular, the correctness of the reduction for Boolean binary predicates in [20] is done via an ad hoc case analysis for 11 concrete predicates. In the generalisation to binary predicates on arbitrary finite domains in [14], the correctness is proved via a combinatorial property of bipartite graphs without a certain 4-vertex graph¹ as a subgraph and a reduction to cuts with more than two parts.

In our case, we use the same black-box reduction via the k -partite k -fold cover. Thus the reduction itself is pretty straightforward, although the analysis is not. In fact, we find it surprising and unexpected that the k -partite k -fold cover works in the additive setting. Our key contribution is the proof of its correctness. A few simple reductions get us to the most technically involved case, in which k is even and the k -ary predicate satisfies $P(1, \dots, 1) = 0$.

¹ A bipartite graph on four vertices with each part of size two and precisely one edge between the two parts.

Additive sparsifiability of such predicates is established in Proposition 13. Unlike in the multiplicative setting, it is not clear how to do this in a straightforward way similar to [20, 14]. Instead, we associate with a given predicate P a vector v_P in an appropriate vector space, identify special vectors that can be shown additively sparsifiable directly, show that linear combinations preserve sparsifiability, and argue that v_P can be generated by the special vectors. The latter is the most technical part of the proof. While there are several natural ideas how to achieve this in a seemingly simpler way (such as arguing that the special vectors form a basis), we have not managed to produce a simpler or shorter proof.

The result in [7] also works for non-constant k . We emphasise that we deal with constant k , which is standard in the CSP literature in that the predicate (or a set of predicates) is fixed and not part of the input. For constant k , the representation of predicates is irrelevant (cf. Remark 18). Thus we do not keep track of (and have not tried to optimise) the precise dependency of the reduction on the predicate arity k (or the domain size $q = |D|$).

Related work

The already mentioned spectral sparsification [33] is a stronger notion than cut sparsification as it requires that not only cuts but also the Laplacian spectrum of a given graph should be (approximately) preserved [32, 34, 21, 31, 7].

Our focus in this article is on *edge sparsifiers* (of cuts and generalisations via local predicates). There are also vertex sparsifiers, in which one reduces the number of vertices. Vertex sparsifiers have been studied for cut sparsification (between special vertices called terminals) [22, 26, 25, 15] as well as for spectral sparsification [24].

Sparsification in general is about finding a sparse sub(hyper)graph while preserving important properties of interest. In addition to cut sparsifiers, another well studied concept is that of *spanners*. A spanner of a graph is a (sparse) subgraph that approximately preserves distances of shortest paths. Spanners have been studied in great detail both in the multiplicative [5, 28, 3, 17, 6, 9, 30] and additive [2, 18, 12, 8, 35, 16] setting. Emulators are a generalisation of spanners in which the sparse graph is not required to be a subgraph of the original graph. We refer the reader to a nice recent survey of Elkin and Neimain for more details [19]. Some of the proofs are deferred to the full version of this paper [29].

2 Preliminaries

For an integer k , we denote by $[k]$ the set $\{0, 1, \dots, k-1\}$. All graphs and hypergraphs² in this paper are unweighted.

For an assignment $a : V \rightarrow S$ from the set of vertices of a (hyper)graph to some set S containing 0, we denote by $Z_a = \{v \in V : a(v) = 0\}$ the set of vertices mapped to 0.

If $0 \leq i \leq r^k - 1$ is an integer, we denote by $\text{rep}_{r,k}(i)$ the representation of i in base r as a vector in \mathbb{R}^k , where the first coordinate stands for the most significant digit, and the last coordinate for the least significant digit. For the special case $r = 2$, we use the notation $\text{bin}_k(i)$ for the binary representation of i .

We denote by $v[j]$ the j -th coordinate of the vector v , counting from 0.

For an integer $0 \leq i \leq 2^k - 1$, we use $\text{zeros}_k(i) = \{\ell \in [k] : \text{bin}_k(i)[\ell] = 0\}$; for example $\text{zeros}_6(52) = \{2, 4, 5\}$, since $\text{bin}_6(52) = (1, 1, 0, 1, 0, 0)$.

We now define the value of an assignment on a hypergraph with a fixed predicate.

² We use the standard definition of hypergraphs, in which every hyperedge is an ordered tuple of vertices.

► **Definition 1.** Let $G = (V, E)$ be a directed k -uniform hypergraph and let $P : D^k \rightarrow \{0, 1\}$ be a k -ary predicate on a finite set D . Given an assignment $a : V \rightarrow D$ of G , the value of a is defined by $\text{Val}_{G,P}(a) = \sum_{(v_1, \dots, v_k) \in E} P(a(v_1), \dots, a(v_k))$. If G is undirected and P is order invariant,³ we define $\text{Val}_{G,P}(a) = \sum_{\{v_1, \dots, v_k\} \in E} P(a(v_1), \dots, a(v_k))$.⁴

The notion of additive sparsification was first introduced in [7] for cuts in graphs and hypergraphs. In order to define it, we will need the $\text{Cut} : \{0, 1\}^k \rightarrow \{0, 1\}$ predicate defined by $\text{Cut}(b_1, \dots, b_k) = 1 \iff \exists i, j, b_i \neq b_j$. Given a hypergraph $G = (V, E)$ and a set $U \subseteq V$, we denote by $\text{vol}_G(U)$ the *volume* of U , defined as the sum of the degrees in G of all vertices in U .

► **Definition 2.** Let $G = (V, E)$ be an undirected k -uniform hypergraph, and denote $|V| = n$. We say that G admits additive cut sparsification with error ε using $O(f(n, \varepsilon))$ hyperedges if there exists a subhypergraph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$ with $|E_\varepsilon| = O(f(n, \varepsilon))$, called an additive sparsifier of G , such that for every assignment $a : V \rightarrow \{0, 1\}$ we have

$$\left| \frac{|E|}{|E_\varepsilon|} \text{Val}_{G_\varepsilon, \text{Cut}}(a) - \text{Val}_{G, \text{Cut}}(a) \right| \leq \varepsilon(d_G |Z_a| + \text{vol}_G(Z_a)), \quad (1)$$

where d_G is the average degree of G .

Note that (1) can also be written as

$$\frac{|E|}{|E_\varepsilon|} \text{Val}_{G_\varepsilon, \text{Cut}}(a) \in \text{Val}_{G, \text{Cut}}(a) \pm \varepsilon(d_G |Z_a| + \text{vol}_G(Z_a)),$$

which explains the use of the term “additive” for the error.

Bansal, Svensson and Trevisan [7] showed the following sparsification result:

► **Theorem 3** (Additive Cut Sparsification [7, Theorem 1.3]). Let $G = (V, E)$ be an undirected n -vertex k -uniform hypergraph, and $\varepsilon > 0$. Then G admits additive cut sparsification with error ε using $O\left(\frac{n}{k} \varepsilon^{-2} \log\left(\frac{k}{\varepsilon}\right)\right)$ hyperedges.

► **Remark 4.** We call a predicate P *symmetric* if it is order invariant (as in Definition 1). Since Theorem 3 deals with only *undirected* hypergraphs, it is not clear how to generalise it to non-symmetric predicates directly, since the value of such predicates on undirected hypergraphs is not defined. Therefore, our course of action will be first to prove it for the case of directed hypergraphs, and then generalise it to other predicates on directed hypergraphs. In fact, by doing this we also prove the result for undirected hypergraphs with symmetric predicates, since hyperedges can be given arbitrary directions without changing the average degree of G , or the volume in G , or the value of the predicate in any assignment.

► **Remark 5.** Throughout this paper we only discuss the existence of sparsifiers and do not mention the time complexity to find them. However, the (implicit) time complexity results from [7] apply in our more general setting as well since the sparsifiers we find are in fact the same sparsifiers for all predicates, including cuts (cf. Remark 17).

An important tool we use to prove our results is the k -partite k -fold cover of a hypergraph. This construction is a well known one, and has been used for multiplicative sparsification (for $k = 2$) in [20] and [14].

³ $P(b_1, \dots, b_k) = P(b_{\sigma(1)}, \dots, b_{\sigma(k)})$ for all $b_1, \dots, b_k \in D$ and every permutation σ on the set $\{1, \dots, k\}$.

⁴ The terms are well defined since P is order invariant.

► **Definition 6.** Let $G = (V, E)$ be a directed k -uniform hypergraph. The k -partite k -fold cover of G is the hypergraph $\gamma(G) = (V^\gamma, E^\gamma)$ where

$$V^\gamma = \{v^{(0)}, v^{(1)}, \dots, v^{(k-1)} : v \in V\},$$

$$E^\gamma = \{(v_1^{(0)}, v_2^{(1)}, \dots, v_k^{(k-1)}) : (v_1, \dots, v_k) \in E\}.$$

If G is undirected we define the cover in the same way except

$$E^\gamma = \{(v_1^{(0)}, v_2^{(1)}, \dots, v_k^{(k-1)}) : \{v_1, \dots, v_k\} \in E\},$$

so for each hyperedge in G we get $k!$ hyperedges in $\gamma(G)$ in this case.

If $k = 2$ then $\gamma(G)$ corresponds to the well-known *bipartite double cover* of G [13].

3 Sparsification of Boolean Predicates

As mentioned in Section 1, we begin by observing that Theorem 3 also works for directed hypergraphs. The simple proof of this fact can be found in the full version [29]. (We emphasise that we treat k as a constant, cf. Remark 18.)

► **Proposition 7.** Let $G = (V, E)$ be a directed n -vertex k -uniform hypergraph, and $\varepsilon > 0$. Then G admits additive cut sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.

From now on, whenever we say a “hypergraph”, we mean a “directed hypergraph” with n vertices. By Remark 4, the results also apply to undirected hypergraphs (whenever it makes sense, i.e. if the associated predicate is symmetric). We also omit the word additive when discussing sparsification. The following notion of sparsification is a natural generalisation of cut sparsification (Definition 2) to arbitrary predicates.

► **Definition 8.** Let P be a k -ary Boolean predicate and $G = (V, E)$ a k -uniform hypergraph. We say that G admits P -sparsification with error ε using $O(f(n, \varepsilon))$ hyperedges if there exists a subhypergraph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$ with $|E_\varepsilon| = O(f(n, \varepsilon))$, called a P -sparsifier of G , such that for every assignment $a : V \rightarrow \{0, 1\}$ we have

$$\left| \frac{|E|}{|E_\varepsilon|} \text{Val}_{G_\varepsilon, P}(a) - \text{Val}_{G, P}(a) \right| \leq \varepsilon (d_G |Z_a| + \text{vol}_G(Z_a)), \quad (2)$$

where d_G is the average degree of G .

The following theorem is our main result, extending Proposition 7 to all k -ary predicates with Boolean domains.

► **Theorem 9 (Main).** For every k -uniform hypergraph G (k is a constant), every k -ary Boolean predicate $P : \{0, 1\}^k \rightarrow \{0, 1\}$, and every $\varepsilon > 0$, G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.

Theorem 9 can be informally restated as “every k -uniform hypergraph is sparsifiable with respect to all k -ary Boolean predicates” or “for every Boolean predicate P of constant arity, $\text{CSP}(P)$ is sparsifiable”.

► **Remark 10.** It is possible to consider an even more general case where each hyperedge in G has its own predicate. In this case, we can apply Theorem 9 to each of the hypergraphs obtained by taking only hyperedges corresponding to a specific predicate, and so get a sparsifier for each such predicate. Taking the union of all their hyperedges, we get a new


hypergraph G_ε , which is a sparsifier of the original hypergraph. Indeed, it has $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges since it is the union of a constant number of hypergraphs. (The number of predicates $P : \{0, 1\}^k \rightarrow \{0, 1\}$ is constant, since k is constant.) It also satisfies (2) for any given assignment up to some constant factor, since all the sparsifiers it is composed of do. This constant factor can be eliminated by choosing $\varepsilon_0 = \frac{\varepsilon}{m}$ for an appropriate m that depends only on k .

The main work in the proof of Theorem 9 is for even values of k ; a simple reduction (Proposition 16) then reduces the case of k odd to the even case.

In order to prove Theorem 9 for even k , we use the k -partite k -fold cover of G and apply Proposition 7 to various assignments of it. For a k -ary Boolean predicate $P : \{0, 1\}^k \rightarrow \{0, 1\}$, we consider the vector $v_P \in \mathbb{R}^{2^k}$, defined by $v_P[i] = P(\text{bin}_k(i))$. For instance, for the Cut predicate on a 3-uniform hypergraph, we have $v_{\text{Cut}} = (0, 1, 1, 1, 1, 1, 0)$.

For a given hypergraph G and an assignment a , we consider the vector $v_{G,a} \in \mathbb{R}^{2^k}$ defined by $v_{G,a}[i] = |\{(v_1, \dots, v_k) \in E : (a(v_1), \dots, a(v_k)) = \text{bin}_k(i)\}|$. In other words, each coordinate of $v_{G,a}$ counts the hyperedges in G whose vertices are assigned some specific set of values by a .

► **Example 11.** Given the graph $G = (V, E)$ in Figure 1 (so $k = 2$) and the assignment $a : V \rightarrow \{0, 1\}$ defined as $a(v_1) = a(v_2) = a(v_3) = 0$ and $a(v_4) = a(v_5) = a(v_6) = a(v_7) = 1$, we have $v_{G,a} = (2, 3, 1, 5)$, since there are two edges with assignment $(0, 0)$, namely (v_1, v_2) and (v_2, v_3) , three edges with assignment $(0, 1)$, namely (v_1, v_4) , (v_2, v_6) , and (v_2, v_7) , etc.



■ **Figure 1** Graph from Example 11.

Under these notations, we get $\text{Val}_{G,P}(a) = \langle v_P, v_{G,a} \rangle$, where $\langle \cdot, \cdot \rangle$ is the standard inner product in \mathbb{R}^{2^k} . We begin by proving the following useful lemma.

► **Lemma 12.** Let $G = (V, E)$ be a k -uniform hypergraph, P_1, \dots, P_m be k -ary Boolean predicates (m is a constant). Suppose that for every $\varepsilon > 0$ and $1 \leq i \leq m$, G admits P_i -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges, and that the same subhypergraph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$ is a P_i -sparsifier for all P_i . Suppose that P is some k -ary Boolean predicate for which we have $v_P = \sum_{i=1}^m \lambda_i v_{P_i}$ for some constants $\lambda_1, \dots, \lambda_m \in \mathbb{R}$. Under these conditions, G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.

Proof. Let $\varepsilon > 0$ and denote $\varepsilon_i = \frac{\varepsilon}{m|\lambda_i|}$ (if $\lambda_i = 0$ take $\varepsilon_i = 1$ instead) and $\varepsilon_0 = \min\{\varepsilon_1, \dots, \varepsilon_m\}$. Let $G_{\varepsilon_0} = (V, E_{\varepsilon_0})$ be the common witness subhypergraph for ε_0 promised by the assumption. We know that every P_i satisfies

$$\left| \frac{|E|}{|E_{\varepsilon_0}|} \text{Val}_{G_{\varepsilon_0}, P_i}(a) - \text{Val}_{G, P_i}(a) \right| \leq \varepsilon_0 (d_G |Z_a| + \text{vol}_G(Z_a)) \tag{3}$$

for every assignment $a : V \rightarrow \{0, 1\}$. We also have

$$\text{Val}_{G,P}(a) = \langle v_P, v_{G,a} \rangle = \left\langle \sum_{i=1}^m \lambda_i v_{P_i}, v_{G,a} \right\rangle = \sum_{i=1}^m \lambda_i \langle v_{P_i}, v_{G,a} \rangle = \sum_{i=1}^m \lambda_i \text{Val}_{G, P_i}(a),$$

75:8 Additive Sparsification of CSPs

and similarly

$$\text{Val}_{G_{\varepsilon_0}, P}(a) = \sum_{i=1}^m \lambda_i \text{Val}_{G_{\varepsilon_0}, P_i}(a).$$

Therefore, for every assignment a we get

$$\begin{aligned} \left| \frac{|E|}{|E_{\varepsilon_0}|} \text{Val}_{G_{\varepsilon_0}, P}(a) - \text{Val}_{G, P}(a) \right| &= \left| \frac{|E|}{|E_{\varepsilon_0}|} \sum_{i=1}^m \lambda_i \text{Val}_{G_{\varepsilon_0}, P_i}(a) - \sum_{i=1}^m \lambda_i \text{Val}_{G, P_i}(a) \right| \\ &\leq \sum_{i=1}^m |\lambda_i| \left| \frac{|E|}{|E_{\varepsilon_0}|} \text{Val}_{G_{\varepsilon_0}, P_i}(a) - \text{Val}_{G, P_i}(a) \right| \\ &\leq \sum_{i=1}^m |\lambda_i| \varepsilon_0 (d_G |Z_a| + \text{vol}_G(Z_a)) \\ &\leq \varepsilon (d_G |Z_a| + \text{vol}_G(Z_a)), \end{aligned}$$

where the second line is due to the triangle inequality, the third is due to (3) and the fourth is by the definition of ε_0 .

Furthermore, since m and all λ_i are constants,

$$|E_{\varepsilon_0}| = O\left(n\varepsilon_0^{-2} \log \frac{1}{\varepsilon_0}\right) = O\left(n\varepsilon^{-2} \log \frac{1}{\varepsilon}\right),$$

and so G_{ε_0} is a witness for the P -sparsification of G . ◀

The core of the proof of Theorem 9 is in the next proposition, which establishes the result for Boolean predicates on even uniformity hypergraphs, with a small restriction.

► **Proposition 13.** *Let k be an even number and G be a k -uniform hypergraph. Let P be a k -ary Boolean predicate with $P(1, 1, \dots, 1) = 0$. Then for every $\varepsilon > 0$, G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.*

Proof. Let $\varepsilon > 0$. We consider $\gamma(G)$, the k -partite k -fold cover of G . Let $\gamma(G)_\varepsilon$ be a subhypergraph of $\gamma(G)$ promised by Proposition 7, and $G_\varepsilon = (V, E_\varepsilon)$ the corresponding subhypergraph of G , i.e. the subhypergraph which satisfies $\gamma(G_\varepsilon) = \gamma(G)_\varepsilon$ (by taking the hyperedges corresponding to the ones of $\gamma(G)_\varepsilon$).

Let $a : V \rightarrow \{0, 1\}$. For every subset $T \subseteq [k]$, we look at the assignment $a_T : V^\gamma \rightarrow \{0, 1\}$ defined by $a_T(v^{(i)}) = 0$ if $i \in T$ and $a(v) = 0$, and $a_T(v^{(i)}) = 1$ otherwise. We therefore have


$$\left| \frac{|E^\gamma|}{|E_\varepsilon^\gamma|} \text{Val}_{\gamma(G)_\varepsilon, \text{Cut}}(a_T) - \text{Val}_{\gamma(G), \text{Cut}}(a_T) \right| \leq \varepsilon (d_{\gamma(G)} |Z_{a_T}| + \text{vol}_{\gamma(G)}(Z_{a_T})). \quad (4)$$

Define the vector $u_T \in \mathbb{R}^{2^k}$ as follows:

$$u_T[j] = \begin{cases} 1 & T \cap \text{zeros}(j) \neq \emptyset, [k] \\ 0 & \text{otherwise} \end{cases}.$$

In other words, the vector u_T is 1 in index j if and only if there exists an index $i \in T$ in which the binary representation of j has a zero, with the exception of $u_{[k]}[0] = 0$. Denote by P_T the predicate corresponding to u_T , that is $P_T(\text{bin}_k(j)) = 1 \iff u_T[j] = 1$. Observe that

$$\text{Val}_{\gamma(G), \text{Cut}}(a_T) = \text{Val}_{G, P_T}(a),$$



■ **Figure 2** An example of a representation of an assignment on $\gamma(G)$. $Z_a^{(i)}$ consists of all vertices in $V^{(i)}$ which are a copy of a vertex $v \in V$ with $a(v) = 0$, and $\bar{Z}_a^{(i)}$ consists of the rest of $V^{(i)}$. Each hyperedge has a unique path from left to right (but a path might belong to multiple hyperedges), choosing one of $Z_a^{(i)}, \bar{Z}_a^{(i)}$ for each i . Each such path is also in 1-1 correspondence with a coordinate in u_T . In this example $T = \{0, 3, k - 1\}$ and the shaded sets represent $a_T^{-1}(0)$. By green dotted lines we indicated a path corresponding to a hyperedge counted in $\text{Val}_{\gamma(G), \text{Cut}}(a_T)$, and by red dashed lines we indicated a path which does not. The green dotted path corresponds to a value of 1 in the coordinate of u_T with binary representation $(1, 1, 0, 0, 1, \dots, 0, 1)$, and the red dashed path to a value 0 in the coordinate with binary representation $(1, 0, 1, 1, 0, \dots, 1, 1)$. Note that if $T = [k]$ then any hyperedge corresponding to a path only on $Z_a^{(i)}$ is not counted.

since they both count exactly hyperedges (v_1, \dots, v_k) which have some vertex v_i with $a(v_i) = 0$ with $i \in T$, but if $T = [k]$ then they do not count hyperedges which have $a(v_i) = 0$ for all $i = 1, \dots, k$ (see example in Figure 2). The same is true for any hypergraph, and in particular for G_ε , that is

$$\text{Val}_{\gamma(G_\varepsilon), \text{Cut}}(a_T) = \text{Val}_{G_\varepsilon, P_T}(a).$$

Putting these results in (4), we get

$$\begin{aligned} \left| \frac{|E|}{|E_\varepsilon|} \text{Val}_{G_\varepsilon, P_T}(a) - \text{Val}_{G, P_T}(a) \right| &\leq \varepsilon(d_{\gamma(G)}|Z_{a_T}| + \text{vol}_{\gamma(G)}(Z_{a_T})) \\ &\leq \varepsilon(d_G|Z_a| + \text{vol}_G(Z_a)), \end{aligned}$$

so G admits P_T sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges for every $T \subseteq [k]$, and for every ε the sparsification is witnessed by the same subhypergraph G_ε . (Notice that Proposition 7, when applied to $\gamma(G)$ which has kn vertices, gives us a subhypergraph with $O(kn\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges, and recall that k is a constant.)

Our next goal is to show that the vector v_P is a linear combination of the vectors u_T for all $T \in [k]$. To show that, we show that every vector e_r in the standard basis of \mathbb{R}^{2^k} , with $r \neq 2^k - 1$, is a linear combination of these vectors. This is sufficient since the last coordinate of v_P is 0 by the assumption. First we need to order the various sets T . We order them in the following decreasing lexicographic order $T_0, T_1, \dots, T_{2^k-1}$, where $T_j = \text{zeros}(j)$, so $T_0 = [k], T_1 = [k] \setminus \{k - 1\}, T_2 = [k] \setminus \{k - 2\}, T_3 = [k] \setminus \{k - 1, k - 2\}, T_4 = [k] \setminus \{k - 3\}$ and so on, until $T_{2^k-1} = \emptyset$.

Let e_r be a vector in the standard basis of \mathbb{R}^{2^k} . We introduce the following coefficients for $0 \leq m \leq 2^k - 1$:

$$\lambda_{r,m} = \frac{1}{2}(-1)^{\text{Ham}(r \oplus m) + (1 - \mathbb{1}_{r \& m})},$$

where $\oplus, \&$ are the Xor and And binary functions respectively,⁵ Ham is the Hamming weight

⁵ The Xor of two integers is defined as the bitwise Boolean Xor of their binary representations, where the

75:10 Additive Sparsification of CSPs

function, and $\mathbf{1}_d$ returns 1 if $d \neq 0$ and 0 if $d = 0$. Denote

$$f_1(m) = \text{Ham}(r \oplus m) \quad , \quad f_2(m) = (1 - \mathbf{1}_{r \& m}).$$

We shall prove that

$$e_r = \sum_{m=0}^{2^k-1} \lambda_{r,m} u_{T_m}. \tag{5}$$

▷ **Claim 14.** The sum of all coefficients is 0; i.e., $\sum_{m=0}^{2^k-1} \lambda_{r,m} = 0$.

Proof. Let $b_1 b_2 \dots b_k$ be the binary representation of r . Since $r < 2^k - 1$, there exists some $1 \leq i \leq k$ for which $b_i = 0$. We can partition the coefficients into pairs, such that $\lambda_{r,m_1}, \lambda_{r,m_2}$ is a pair if and only if m_1, m_2 differ in the i -th coordinate only. This is clearly a partition. For each pair, f_1 gives m_1, m_2 different parity values, and f_2 gives them the same value (since $b_i = 0$), so $\lambda_{r,m_1}, \lambda_{r,m_2}$ have opposite signs, so their sum is zero. This is true for every pair, so the overall sum is zero, and the claim is proved. ◁

We prove (5) coordinate-wise. First we look at the coordinate r . Consider the set W of all vectors u_{T_m} for which the coordinate r is 0. If we show that the sum of the corresponding coefficients of the vectors in W is -1 , using Claim 14 we will deduce the result in this case. We distinguish 2 cases:

Case (I): $r = 0$. By the definition of u_T , in this case the set W contains two vectors, $u_{[k]}$ and u_\emptyset . The corresponding coefficients are $\lambda_{r,0} = -\frac{1}{2}$ and $\lambda_{r,2^k-1} = -\frac{1}{2}$ (since k is even), which sum up to -1 .

Case (II): $r > 0$. As in the proof of Claim 14, let $b_1 b_2 \dots b_k$ be the binary representation of r , and choose a coordinate $1 \leq i \leq k$ for which $b_i = 1$. Partition the vectors in W into pairs where $u_{T_{m_1}}, u_{T_{m_2}}$ is a pair if and only if m_1, m_2 differ in the i -th coordinate only. This is clearly a partition of all vectors, and by the definition of u_T , each such pair is either contained in W or disjoint from W so this is indeed a partition of W . (Note that $u_{T_{m_1}}[r]$ is determined by $T_{m_1} \cap \text{zeros}(r)$ which is in fact $\text{zeros}(m_1) \cap \text{zeros}(r)$, and the same for m_2 . Since m_1, m_2 differ in the i -th coordinate only, and r is not zero in this coordinate, this coordinate contributes nothing to the intersections, and so both these intersections are empty or non-empty together. The intersection never equals $[k]$ since $r > 0$.) For every such pair in W , if it does not contain the negation of $\text{bin}(r)$, then there is some other index $j \neq i$ in which r, m_1, m_2 are all 1. (This is because in all other coordinates m_1, m_2 are equal, and since they are not the negation of r , there is some coordinate $j \neq i$ in which they are equal to the j -th coordinate of r . These coordinates cannot be all 0, since this would imply $u_{T_{m_1}}, u_{T_{m_2}} \notin W$.) This implies that f_2 gives m_1, m_2 the same value, and clearly f_1 gives them different parity values, so $\lambda_{r,m_1} + \lambda_{r,m_2} = 0$. However, for the pair which contains the negation of r (this pair is clearly in W), suppose without loss of generality the negation is m_1 . Then f_2 gives m_1, m_2 the values 1, 0 respectively, and f_1 gives m_1 an even value and m_2 an odd value (since k is even), and so $\lambda_{r,m_1} = \lambda_{r,m_2} = -\frac{1}{2}$, and the overall sum is -1 . This finishes the proof of (5) in the coordinate r .

Now let $r' \neq r$ be some other coordinate, and let $c_1 c_2 \dots c_k$ be its binary representation. First, if $r' = 2^k - 1$ then for all m we have $u_{T_m}[r'] = 0$ by definition, so the linear combination

Boolean Xor of two bits is their sum modulo 2. The And of two integers is defined the same way with the Boolean And function which is defined as $\text{And}(i, j) = 1 \iff i = j = 1$.

of this coordinate is 0. So suppose $r' < 2^k - 1$. As before let W be the set of all vectors u_{T_m} for which the coordinate r' is 0. We show that the sum of the corresponding coefficients is zero, and again deduce the result using Claim 14. Now, there exists some index i for which $b_i \neq c_i$. Again we have two cases:

Case (1): $b_i = 0, c_i = 1$. Partition the vectors in W into pairs where $u_{T_{m_1}}, u_{T_{m_2}}$ is a pair if and only if m_1, m_2 differ in the i -th coordinate only. This is clearly a partition of all the vectors, and by the definition of u_T , each such pair is either contained in W or disjoint from W , so this is indeed a partition of W . For every such pair in W , f_1 gives m_1, m_2 different parity values, and f_2 gives them the same value (since $b_i = 0$), so $\lambda_{r, m_1}, \lambda_{r, m_2}$ have opposite signs, so their sum is zero. This is true for every pair in W , so the overall sum is zero.

Case (2) $b_i = 1, c_i = 0$. Here we consider two sub-cases:

Case (2a): $r' = 0$. The only vectors in W in this case are $u_{[k]}$ and u_\emptyset . The corresponding coefficients are $\lambda_{r, 0} = \frac{1}{2}(-1)^{\text{Ham}(r)+1}$ and $\lambda_{r, 2^k-1} = \frac{1}{2}(-1)^{\text{Ham}(\neg r)}$, where $\neg r$ denotes the negation of the binary representation of r . Since k is even, we know that $r, \neg r$ have the same parity, and so the sum of the two coefficients is 0.

Case (2b): $r' \neq 0$. Choose some j for which $c_j = 1$. Partition the vectors in W into pairs where $u_{T_{m_1}}, u_{T_{m_2}}$ is a pair if and only if m_1, m_2 differ in the j -th coordinate only. The argument for this being a partition of W is similar to the argument in Case (1). For each pair in W , f_1 gives m_1, m_2 a different parity as always, and f_2 gives them the same value, since r, m_1, m_2 are all 1 in the index i (similar argument as before), so the sum of coefficients is 0 for each pair, and so for all coefficients corresponding to vectors in W .

This finishes the proof of (5), and so v_P is a linear combination of the vectors u_T . From the result above and Lemma 12 we deduce that G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges, as required. ◀

To complete the picture for even k , we reduce to Proposition 13 by a simple “complementarity trick”; the proof can be found in the full version [29].

► **Proposition 15.** *Let k be an even number, and G a k -uniform hypergraph. Let P be a k -ary Boolean predicate. Then for every $\varepsilon > 0$, G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.*

The final piece in the jigsaw, proved in in the full version [29], shows how to reduce sparsification of k -uniform hypergraphs with k odd to the case of $(k+1)$ -uniform hypergraphs by adding a universal vertex and extending the original predicate by one dimension.

► **Proposition 16.** *Let k be an odd number, and $G = (V, E)$ a k -uniform hypergraph. Let P be a k -ary Boolean predicate. Then for every $\varepsilon > 0$, G admits P -sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.*

Propositions 15 and 16 complete the proof of Theorem 9.

► **Remark 17.** In the proof of Proposition 13 the hypergraph G_ε was chosen independently of the predicate P . Since Propositions 15 and 16 reduce to that case, we have in fact shown that for every $\varepsilon > 0$, Theorem 9 is witnessed by the same subhypergraph G_ε for all different predicates P . This will be important in the proof of Theorem 20.

► **Remark 18.** We note that our main result, Theorem 9, extends Theorem 3 in the regime where k is a constant, which is the main focus of this paper. However, Theorem 3 also works for non-constant k [7]. If k is not a constant, it can be seen from the proof of Lemma 12 that the number of hyperedges of the sparse subhypergraph is multiplied by a factor of $O(m^2)$ (since $O(m)$ is the proportion between ε and ε_0 given that the coefficients λ_i are constant).

75:12 Additive Sparsification of CSPs

In Proposition 13 we have $m = 2^k$, and so for k not constant we get an additional factor of 4^k . Furthermore, in Propositions 7 and 13 we obtain extra factors of k , by considering the k -partite k -fold cover. While the regime with non-constant k is interesting for cuts, for arbitrary predicates one needs to be careful about representation as the natural (explicit) representation of (non-symmetric) predicates requires exponential space in the arity k .

4 Sparsification of Non-Boolean Predicates

We now focus on non-Boolean predicates; i.e., predicates of the form $P : D^k \rightarrow \{0, 1\}$ with $|D| > 2$. Without loss of generality, we assume $D = [q]$ for some $q \geq 2$. The most natural way of generalising Theorem 9 to larger domains appears to be to use the same bound with $Z_a = \{v \in V : a(v) = 0\}$. This, however, cannot give the desired sparsification result (cf. Section 5). Instead we use a different and somewhat weaker kind of generalisation of the Boolean case, and show that all hypergraphs are still sparsifiable with respect to all predicates using this definition.

► **Definition 19.** *Let $P : D^k \rightarrow \{0, 1\}$ be a k -ary predicate where $D = [q]$. We say that a k -uniform hypergraph $G = (V, E)$ admits all-but-one P -sparsification with error ε using $O(f(n, \varepsilon))$ hyperedges if there exists a subhypergraph $G_\varepsilon = (V, E_\varepsilon \subseteq E)$ with $|E_\varepsilon| = O(f(n, \varepsilon))$ such that for every assignment $a : V \rightarrow D$ we have*

$$\left| \frac{|E|}{|E_\varepsilon|} \text{Val}_{G_\varepsilon, P}(a) - \text{Val}_{G, P}(a) \right| \leq \varepsilon(d_G |M_a| + \text{vol}_G(N_a)), \quad (6)$$

where M_a is the largest set among the sets $\{v \in V : a(v) = i\}$, N_a is the set with the largest volume among the sets $\{v \in V : a(v) = i\}$ for $0 \leq i \leq q - 2$, and d_G is the average degree in G .

Observe that the maximum in Definition 19 is over $i = 0, \dots, q - 2$ without $i = q - 1$, hence the name “all-but-one”. We note that there is nothing special about $q - 1$ and any value from $[q]$ could be chosen in Definition 19.

Under Definition 19, Theorem 9 generalises (proof of can be found in the full version [29]).

► **Theorem 20.** *For every k -uniform hypergraph $G = (V, E)$, every k -ary predicate $P : D^k \rightarrow \{0, 1\}$ with $D = [q]$ (k, q are constants), and every $\varepsilon > 0$, G admits P all-but-one sparsification with error ε using $O(n\varepsilon^{-2} \log \frac{1}{\varepsilon})$ hyperedges.*

Note that in the case of $q = 2$ we have $P : \{0, 1\}^k \rightarrow \{0, 1\}$, and Definition 19 and Theorem 20 coincide with Definition 8 and Theorem 9. This is because when $q = 2$ the definitions of M_a, N_a coincide with the definition of Z_a in the Boolean case.

5 Optimality of All-But-One Sparsification

One might wonder if there is a different, perhaps stronger way to define sparsification for predicates on non-Boolean domains. The following example shows that all-but-one sparsification is optimal.

For a hypergraph $G = (V, E)$ and a fixed assignment $a : V \rightarrow [q]$ denote $S_i = \{v \in V : a(v) = i\}$ (so $S_0 = Z_a$). The definition of all-but-one sparsification lets us take a bound which depends on the sizes and volumes of all the sets S_i except for S_{q-1} . In fact, if we try to take a bound which depends on fewer of these sets, the definition fails to generalise even the most basic case of the Cut predicate. To see this, it is sufficient to consider the graph case,

i.e. $k = 2$. Let us suppose, without loss of generality, that our bound does not depend on S_{q-2}, S_{q-1} . Consider the predicate $\text{Cut} : [q]^2 \rightarrow \{0, 1\}$ defined by $\text{Cut}(x, y) = 1 \iff x \neq y$. A simple (but lengthy) argument in the full version [29] shows that cliques do not have a Cut-sparsifier using such a definition. In fact, the same argument works for any predicate P with $P(q-2, q-1) = P(q-1, q-2) = 1$ and $P(q-2, q-2) = P(q-1, q-1) = 0$. Thus if a definition does not depend on more than just S_{q-2}, S_{q-1} , it specifically does not depend on these two, so the same argument still works. Therefore, no definition with a bound which depends on “less” is possible, under the current assumptions.

References

- 1 Kook Jin Ahn and Sudipto Guha. Graph sparsification in the semi-streaming model. In *Proceedings of the 36th International Colloquium on Automata, Languages and Programming (ICALP'09), Part II*, volume 5556 of *Lecture Notes in Computer Science*, pages 328–338. Springer, 2009. doi:10.1007/978-3-642-02930-1_27.
- 2 Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999. doi:10.1137/S0097539796303421.
- 3 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993. doi:10.1007/BF02189308.
- 4 Alexandr Andoni, Jiecao Chen, Robert Krauthgamer, Bo Qin, David P. Woodruff, and Qin Zhang. On sketching quadratic forms. In *Proceedings of the 7th ACM Conference on Innovations in Theoretical Computer Science (ITCS'16)*, pages 311–319. ACM, 2016. doi:10.1145/2840728.2840753.
- 5 Baruch Awerbuch. Complexity of network synchronization. *J. ACM*, 32(4):804–823, 1985. doi:10.1145/4221.4227.
- 6 Baruch Awerbuch, Yossi Azar, Avrim Blum, and Santosh S. Vempala. New approximation guarantees for minimum-weight k-trees and prize-collecting salesmen. *SIAM J. Comput.*, 28(1):254–262, 1998. doi:10.1137/S009753979528826X.
- 7 Nikhil Bansal, Ola Svensson, and Luca Trevisan. New notions and constructions of sparsification for graphs and hypergraphs. *Proceedings of the 60th IEEE Annual Symposium on Foundations of Computer Science (FOCS'19)*, pages 910–928, 2019. doi:10.1109/FOCS.2019.00059.
- 8 Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. New constructions of (alpha, beta)-spanners and purely additive spanners. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*, pages 672–681. SIAM, 2005. URL: <http://dl.acm.org/citation.cfm?id=1070432.1070526>.
- 9 Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Struct. Algorithms*, 30(4):532–563, 2007. doi:10.1002/rsa.20130.
- 10 Joshua Batson, Daniel A. Spielman, and Nikhil Srivastava. Twice-Ramanujan sparsifiers. *SIAM Review*, 56(2):315–334, 2014. doi:10.1137/130949117.
- 11 András A. Benczúr and David R. Karger. Approximating s-t Minimum Cuts in $\tilde{O}(n^2)$ Time. In *Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC'96)*, pages 47–55. ACM, 1996. doi:10.1145/237814.237827.
- 12 Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. *SIAM J. Discret. Math.*, 19(4):1029–1055, 2005. doi:10.1137/S0895480103431046.
- 13 Richard A. Brualdi, Frank Harary, and Zevi Miller. Bigraphs versus digraphs via matrices. *Journal of Graph Theory*, 4(1):51–73, 1980. doi:10.1002/jgt.3190040107.
- 14 Silvia Butti and Stanislav Živný. Sparsification of binary CSPs. *SIAM J. Discret. Math.*, 34(1):825–842, 2020. doi:10.1137/19M1242446.

- 15 Parinya Chalermsook, Syamantak Das, Bundit Laekhanukit, Yunbum Kook, Yang P Liu, Richard Peng, Mark Sellke, and Daniel Vaz. Vertex sparsification for edge connectivity. In *Proceedings of the 32nd Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'21)*, 2021. URL: <https://arxiv.org/abs/2007.07862>.
- 16 Shiri Chechik. New additive spanners. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA'13)*, pages 498–512. SIAM, 2013. doi:10.1137/1.9781611973105.36.
- 17 Edith Cohen. Fast algorithms for constructing t-spanners and paths with stretch t. *SIAM J. Comput.*, 28(1):210–236, 1998. doi:10.1137/S0097539794261295.
- 18 Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. *SIAM J. Comput.*, 29(5):1740–1759, 2000. doi:10.1137/S0097539797327908.
- 19 Michael Elkin and Ofer Neiman. Near-additive spanners and near-exact hopsets, A unified view. *Bull. EATCS*, 130, 2020. URL: <http://bulletin.eatcs.org/index.php/beatcs/article/view/608/624>.
- 20 Arnold Filtser and Robert Krauthgamer. Sparsification of two-variable valued constraint satisfaction problems. *SIAM J. Discret. Math.*, 31(2):1263–1276, 2017. doi:10.1137/15M1046186.
- 21 Wai Shing Fung, Ramesh Hariharan, Nicholas J. A. Harvey, and Debmalya Panigrahi. A general framework for graph sparsification. *SIAM J. Comput.*, 48(4):1196–1223, 2019. doi:10.1137/16M1091666.
- 22 Torben Hagerup, Jyrki Katajainen, Naomi Nishimura, and Prabhakar Ragde. Characterizing multiterminal flow networks and computing flows in networks of small treewidth. *J. Comput. Syst. Sci.*, 57(3):366–375, 1998. doi:10.1006/jcss.1998.1592.
- 23 Dmitry Kogan and Robert Krauthgamer. Sketching cuts in graphs and hypergraphs. In *Proceedings of the 6th ACM Conference on Innovations in Theoretical Computer Science (ITCS'15)*, pages 367–376. ACM, 2015. doi:10.1145/2688073.2688093.
- 24 Rasmus Kyng, Yin Tat Lee, Richard Peng, Sushant Sachdeva, and Daniel A. Spielman. Sparsified cholesky and multigrid solvers for connection laplacians. In *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing (STOC'16)*, pages 842–850. ACM, 2016. doi:10.1145/2897518.2897640.
- 25 Frank Thomson Leighton and Ankur Moitra. Extensions and limits to vertex sparsification. In *Proceedings of the 42nd ACM Symposium on Theory of Computing (STOC'2010)*, pages 47–56. ACM, 2010. doi:10.1145/1806689.1806698.
- 26 Ankur Moitra. Approximation algorithms for multicommodity-type problems with guarantees independent of the graph size. In *Proceedings of the 50th Annual IEEE Symposium on Foundations of Computer Science (FOCS'09)*, pages 3–12. IEEE Computer Society, 2009. doi:10.1109/FOCS.2009.28.
- 27 Ilan Newman and Yuri Rabinovich. On multiplicative lambda-approximations and some geometric applications. *SIAM J. Comput.*, 42(3):855–883, 2013. doi:10.1137/100801809.
- 28 David Peleg and Alejandro A Schäffer. Graph spanners. *Journal of graph theory*, 13(1):99–116, 1989. doi:10.1002/jgt.3190130114.
- 29 Eden Pelleg and Stanislav Živný. Additive Sparsification of CSPs, 2021. arXiv:2106.14757.
- 30 Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP'05)*, volume 3580 of *Lecture Notes in Computer Science*, pages 261–272. Springer, 2005. doi:10.1007/11523468_22.
- 31 Tasuku Soma and Yuichi Yoshida. Spectral sparsification of hypergraphs. In *Proceedings of the 30th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA'19)*, pages 2570–2581, 2019. doi:10.1137/1.9781611975482.159.
- 32 Daniel A. Spielman and Nikhil Srivastava. Graph sparsification by effective resistances. *SIAM J. Comput.*, 40(6):1913–1926, 2011. doi:10.1137/080734029.
- 33 Daniel A. Spielman and Shang-Hua Teng. Nearly-linear time algorithms for graph partitioning, graph sparsification, and solving linear systems. In *Proceedings of the 36th Annual ACM*

- Symposium on Theory of Computing (STOC'04)*, pages 81–90. ACM, 2004. doi:10.1145/1007352.1007372.
- 34 Daniel A. Spielman and Shang-Hua Teng. Spectral sparsification of graphs. *SIAM J. Comput.*, 40(4):981–1025, 2011. doi:10.1137/08074489X.
- 35 David P. Woodruff. Additive spanners in nearly quadratic time. In *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP'10)*, volume 6198 of *Lecture Notes in Computer Science*, pages 463–474. Springer, 2010. doi:10.1007/978-3-642-14165-2_40.