

Decomposing a Graph Into Expanding Subgraphs

Guy Moshkovitz*

Asaf Shapira[†]

Abstract

A paradigm that was successfully applied in the study of both pure and algorithmic problems in graph theory can be colloquially summarized as stating that *any graph is close to being the disjoint union of expanders*. Our goal in this paper is to show that in several of the instantiations of the above approach, the quantitative bounds that were obtained are essentially best possible. Three examples of our results are the following:

- A classical result of Lipton, Rose and Tarjan from 1979 states that if \mathcal{F} is a hereditary family of graphs and every graph in \mathcal{F} has a vertex separator of size $n/(\log n)^{1+o(1)}$, then every graph in \mathcal{F} has $O(n)$ edges. We construct a hereditary family of graphs with vertex separators of size $n/(\log n)^{1-o(1)}$ such that not all graphs in the family have $O(n)$ edges.
- Trevisan and Arora-Barak-Steurer have recently shown that given a graph G , one can remove only 1% of its edges to obtain a graph in which each connected component has good expansion properties. We show that in both of these decomposition results, the expansion properties they guarantee are essentially best possible, even when one is allowed to remove 99% of G 's edges.
- Sudakov and the second author have recently shown that every graph with average degree d contains an n -vertex subgraph with average degree at least $(1-o(1))d$ and vertex expansion $1/(\log n)^{1+o(1)}$. We show that one cannot guarantee a better vertex expansion even if allowing the average degree to be $O(1)$.

The above results are obtained as corollaries of a new family of graphs which we construct in this paper. These graphs have a super-linear number of edges and nearly logarithmic girth, yet each of their subgraphs has (optimally) poor expansion properties.

1 Introduction

In recent years a certain paradigm has emerged which roughly says that any graph is close to being a vertex-disjoint union of expanders. Unlike the ϵ -regular partitions of Szemerédi [29] and the cut decompositions of Frieze-Kannan [8] which are relevant only for dense graphs (i.e., with $\Theta(n^2)$ edges), this paradigm is applicable for graphs of arbitrary density. Generally speaking, its usefulness stems from the fact that it allows for a divide-and-conquer approach, in essence reducing a problem on general graphs to the special case of expander graphs.

*School of Mathematics, Tel Aviv University, Tel Aviv, Israel 69978. Email: guymosko@tau.ac.il. Supported in part by ISF grant 224/11.

[†]School of Mathematics, Tel Aviv University, Tel Aviv, Israel 69978. Email: asafico@tau.ac.il. Supported in part by ISF Grant 1028/16 and ERC-Starting Grant 633509.

For algorithmic problems, this paradigm has been useful, for example, in designing approximation algorithms related to the Unique Games Conjecture [13], as in the seminal work of Arora-Barak-Steurer [4] and of Trevisan [30], as well as approximation algorithms for problems such as the Traveling Salesman Problem [3]. It has also seen applications in property testing algorithms [6, 9], as well as in data structure design [23]. In graph theory, instantiations of this paradigm include the theorem of Lipton-Rose-Tarjan [19] on vertex separators in hereditary families, the results of Linial-Saks [18] and Leighton-Rao [16] on low-diameter decompositions, as well as results in the field of graph minors [14, 26] and cycle packing [5]. See Section 2 for some further discussion.

We note that in all applications of the regularity lemmas [8, 29] one uses the same notion of expansion¹. On the other hand, when dealing with sparse graphs and applying the above-mentioned paradigm, each application calls for a different notion of expansion. Our main goal in this paper is to show that in several of the above-mentioned applications, the different notions of expansion that were used are quantitatively best possible.

As it turns out, all our results can be deduced from a single construction of graphs whose main property is that every subgraph has a small *edge separator*.

Definition 1.1. *An edge separator in a graph G is a set of edges whose removal leaves no connected component with more than $\frac{2}{3}|V(G)|$ vertices. The minimum cardinality of an edge separator in G is denoted $\text{sep}(G)$.*

We recall the definition of edge expanders. For a graph $G = (V, E)$ and a non-empty subset $A \subseteq V$, we denote by $\partial_G(A)$ the set of edges of G with precisely one endpoint in A , and by $\phi_G(A) = |\partial_G(A)|/|A|$ the *edge expansion* of A in G . We say that G has *edge expansion* α if $\phi_G(A) \geq \alpha$ for every $A \subseteq V$ with $0 < |A| \leq |V|/2$. In the paper, we will frequently rely (without an explicit reference) on the following trivial relation between separators and expansion.

Fact 1.2. *The edge expansion of every graph G is at most $3 \text{sep}(G)/|V(G)|$.*

Throughout, the *girth* of a graph G is the minimum length of a cycle in G , and the maximum degree of G is denoted $\Delta(G)$. We write $\log(\cdot)$ for $\log_2(\cdot)$. Our main technical result, stated next, implies that there exist graphs with a super-linear number of edges, girth (essentially) logarithmic and whose subgraphs are all either trees or have a small edge separator.

Theorem 1. *For any n, k with $2 \leq k \leq \frac{1}{648} \log \log n$ there is an n -vertex graph $G = G_{n,k}$ satisfying:*

- i. G has average degree at least k and maximum degree at most $6k$.*
- ii. G has girth at least $\log n / (6k)^2$.*
- iii. Every t -vertex subgraph H of G with $t \geq \log n / (6k)^2$ satisfies*

$$\text{sep}(H) \leq \frac{t}{\log t} \cdot (\log \log t)^2. \tag{1}$$

The quantitative estimates in Theorem 1 are, in fact, best possible up to $\log \log$ factors, and here we briefly explain why. First, as we discuss in Section 6, any graph satisfying item (iii) has

¹In the setting of the regularity lemma, expansion is referred to as being ϵ -regular.

average degree at most $(\log \log n)^{O(1)}$. Second, the girth of any graph of average degree at least 4 is at most $O(\log n)$, which matches item (ii) up to $\log \log n$ factors. Third, it follows from Theorem 4 below that every graph G has a subgraph H (of roughly the same average degree) with $\text{sep}(H)$ the same as in (1) up to some $\log \log t$ factors.

Crucially, $\text{sep}(H)$ is bounded in (1) only in terms of t , meaning that this small-separator property is *hereditary*. Notice that the fact that forest subgraphs are ignored in item (iii) is essential, as any graph G contains $K_{1,\Delta(G)}$ (the star with $\Delta(G)$ leaves) and $\text{sep}(K_{1,\Delta}) \geq \Delta/3$. Nevertheless, we still have the following fact (a proof of which can be deduced from Lemma 2.3 below).

Fact 1.3. *Every forest H has a vertex separator of size 1.² Consequently, $\text{sep}(H) \leq \Delta(H)$.*

Finally, we mention that the proof of Theorem 1 actually shows that in (1) we can replace the $(\log \log t)^2$ term by $(\log \log t)^{1+o(1)}$, but we opted to use the simpler/cleaner expression.

Let us make a few remarks about the proof of Theorem 1. Constructing graphs satisfying properties (i) and (ii) of the theorem is of course an easy application of the probabilistic deletion method. However, random graphs will also have the property that most of their subgraphs will have excellent expansion properties. Therefore, instead of considering random graphs, we pick a random subgraph of the boolean hypercube. As it turns out, showing that such a randomly chosen graph satisfies properties (i) and (ii) still follows from simple combinatorial arguments. Proving property (iii) is more interesting: to this end we borrow a well-known argument from the theory of metric embedding, that was first used by Linial, London and Rabinovich [17] in order to prove that the shortest path metric of a bounded degree n -vertex expander cannot be embedded into ℓ_1 with distortion $o(\log n)$.

1.1 Small set expansion

Next we describe a strengthening of Theorem 1 which will be important for some of our applications. Notice that Theorem 1 gives a graph G whose every subgraph H has a *large* subset (i.e., consisting of at least $1/3$ of its vertices) that does not expand well. One may instead ask for a graph G that does not contain even *small set expanders*; that is, a graph G whose every subgraph H has a small subset (i.e., consisting of $o(1)$ -fraction of the vertices) that does not expand well. The notion of small set expanders has recently received much attention in theoretical computer science (see, e.g., [4, 27, 28]).

We prove the following “small set” counterpart of Theorem 1. In fact, we prove the stronger property that one can nearly *partition* every subgraph H into small non-expanding subsets.

Theorem 2. *Suppose that, for n, k with $2 \leq k \leq \frac{1}{648} \log \log n$, the graph G satisfies properties (i),(ii) and (iii) of Theorem 1, and let $t \geq \log n / (6k)^2$ and $1/\sqrt{t} \leq \mu \leq 1/2$. For every t -vertex subgraph H of G there are at least $1/(8\mu)$ mutually disjoint subsets $A_i \subseteq V(H)$ of size $\mu t/3 \leq |A_i| \leq \mu t$ satisfying*

$$\phi_H(A_i) \leq \frac{\log(1/\mu)}{\log t} \cdot (14 \log \log t)^2. \quad (2)$$

²I.e., a single vertex whose removal leaves no connected component with more than $\frac{2}{3}|V|$ vertices.

In words, Theorem 2 gives a graph G whose every (large enough) subgraph H has a weak expansion profile, meaning that H has subsets with μt vertices and edge expansion at most roughly $\log(1/\mu)/\log t$, where $t = |V(H)|$. We note that by choosing $\mu = 1/\log t$ say, the resulting subsets have both size $o(1)$ and edge expansion $o(1)$ as t tends to infinity. As in Theorem 1, the quantitative estimates in Theorem 2 are essentially best possible. Indeed, it follows from Theorem 4 below that any graph G has a t -vertex subgraph H whose every subset of size μt has edge expansion at least $(\log(1/\mu)/\log t) \cdot (\log \log t)^{-2}$, which matches (2) up to $\log \log t$ factors.

The proof of Theorem 2 relies on the hereditary non-expansion properties of the graphs produced by Theorem 1 stated in item (iii). This allows us to devise a careful process which iteratively produces smaller and smaller sets within each subgraph H in such a way that the expansion of the sets does not increase too much.

1.2 Paper overview

In Section 2 we describe six applications of Theorems 1 and 2, showing that many decomposition-type results that were used in different areas of research are essentially best possible. The proof of Theorem 1 is given in Section 3 and the proof of Theorem 2 is given in Section 4. We give an overview of each of the proofs at the beginning of each section. Section 5 contains some deferred proofs from other sections and Section 6 contains some concluding remarks and open problems. We note that throughout the paper we use expressions such as $(n/\log n)(\log \log n)$, by which we of course mean $\frac{n}{\log n} \cdot \log \log n$.

2 Applications of Main Results

Henceforth, a *vertex separator* is a set of vertices whose removal leaves no connected component with more than $\frac{2}{3}|V(G)|$ vertices. The minimum cardinality of a vertex separator in G is denoted $\text{sep}_V(G)$. Note that $\text{sep}_V(G) \leq \text{sep}(G)$, by removing an arbitrary endpoint of each separator edge.

2.1 Edge density of hereditary families with small separators

A family of graphs is said to be *hereditary* if it is closed under taking induced subgraphs. The well-known Planar Separator theorem of Lipton-Tarjan [20] asserts that any n -vertex planar graph has a vertex separator of cardinality at most $O(\sqrt{n})$. This influential result led to many extensions for other hereditary families of graph (such as minor-free families, see [2]). The notion of vertex separator in graphs has found numerous applications to problems in graph theory, both pure and algorithmic (see [25] and the references therein). One of the first applications was given by Lipton, Rose and Tarjan [19] in their work on the nested dissection method. Intuitively, their result states that the reason planar graphs (and more generally minor-free graphs) have linearly many edges is that they have small separators.

Theorem 3 ([19], Theorem 10). *For every $\epsilon > 0$ there is $C > 0$ for which the following holds. Let \mathcal{F} be a hereditary family of graphs such that every n -vertex graph $G \in \mathcal{F}$ has $\text{sep}_V(G) \leq n/(\log n)^{1+\epsilon}$. Then every n -vertex graph in \mathcal{F} has at most Cn edges.*

We mention that Fox and Pach [7] strengthened Theorem 3 by proving that even separators of size $n/(\log n(\log \log n)^{1+\epsilon})$ guarantee³ that every graph in \mathcal{F} has $O(n)$ edges.

Using Theorem 1 we next show that the separation requirement in Theorem 3 cannot be improved much beyond $n/\log n$.

Corollary 2.1. *There is a hereditary family of graphs \mathcal{F} such that every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_V(G) \leq 1 + (n/\log n) \cdot (\log \log n)^2$, yet there is no $C > 0$ such that every n -vertex graph in \mathcal{F} has at most Cn edges.*

Proof. Put $n_0 = 2^{2^{648}}$ and note that the graphs $G_{n, \log \log n/648}$ in Theorem 1 exist for every $n \geq n_0$. Let \mathcal{F} be the family of graphs defined as follows;

$$\mathcal{F} = \{G : G \text{ is an induced subgraph of } G_{n, \log \log n/648} \text{ for some } n \geq n_0\}.$$

Note that \mathcal{F} is, by definition, a hereditary family of graphs. By items (ii), (iii) of Theorem 1, together with Fact 1.3, every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}_V(G) \leq 1 + (n/\log n)(\log \log n)^2$. On the other hand, item (i) of Theorem 1 implies that \mathcal{F} contains, for every $n \geq n_0$, an n -vertex graph with $\Omega(n \log \log n)$ edges. This completes the proof. \square

2.2 Finding a single vertex-expanding subgraph

For a graph G , a subset $S \subseteq V(G)$ is said to have *vertex expansion* α if $|N(S)| = \alpha |S|$, where the *vertex boundary* $N(S)$ of S is the set of vertices outside of S that have a neighbor in S . Motivated by extremal problems related to graph minors, it was shown in [26] that every graph G contains a subgraph with good vertex expansion properties and almost the same average degree as G .

Theorem 4 ([26], Lemma 1.2). *Let $0 < \epsilon \leq 2^{-8}$ and let G be a graph of average degree k . There is a t -vertex subgraph H of G , with average degree at least $(1 - \epsilon)k$, such that every subset of $V(H)$ of size μt with $1/t \leq \mu \leq 1/2$ has vertex expansion in H at least $\epsilon \cdot \log(1/\mu)/(4 \log t (\log \log t)^2)$.*

We now show that the expansion guarantee in Theorem 4 is tight up to a $(\log \log t)^4$ factor for all $1/\sqrt{t} \leq \mu \leq 1/2$.

Corollary 2.2. *Let G be any graph from Theorem 1. For every t -vertex subgraph H of G , and every $1/\sqrt{t} \leq \mu \leq 1/2$, there is a subset of $V(H)$ of size in $[\mu t/3, \mu t]$ that has vertex expansion in H at most $(\log(1/\mu)/\log t) \cdot (14 \log \log t)^2$.*

Proof. Let H be a t -vertex subgraph of G , and let $1/\sqrt{t} \leq \mu \leq 1/2$. Suppose H contains a cycle. By item (ii) of Theorem 1, $t \geq \log n/(6k)^2$. Theorem 2 thus implies that H has a subset of size in $[\mu t/3, \mu t]$ that has edge expansion—and therefore vertex expansion—at most $(\log(1/\mu)/\log t)(14 \log \log t)^2$. Otherwise, if H does not contain a cycle, Lemma 2.3 below guarantees the existence of a subset of size in $[\mu t/2, \mu t]$ that has vertex expansion at most $2/\mu t \leq 2/\sqrt{t} \leq (14 \log \log t)^2/\log t$, as desired. This completes the proof. \square

Lemma 2.3. *Let T be a tree on t vertices. For every $1 \leq a \leq t$ there is a subset $S \subseteq V(T)$ of size $a/2 \leq |S| \leq a$ that has vertex boundary of size at most 1.*

The proof of Lemma 2.3 is deferred to Section 5. We note that the special case of Lemma 2.3 with $a = 2t/3$ can be used to give a proof of Fact 1.3.

³The result of [7] actually shows that having separators of size $n/\log n \log \log n(\log \log \log n)^2$ suffices.

2.3 Decomposing a graph into edge expanders

Trevisan [30], in his work on approximation algorithms for constraint satisfaction problems known as *unique games* (see [30, 4] for background), devised a decomposition of any graph into disjoint expanders. Given this decomposition, Trevisan’s approximation algorithm proceeds to solve the problem on each connected component, exploiting their expansion properties. We note that we will revisit this paradigm for solving unique games in Subsection 2.4, where we discuss the well-known Arora-Barak-Steurer [4] algorithm.

Recall that a graph $G = (V, E)$ has *edge expansion* α if every subset $S \subseteq V$ with $|S| \leq |V|/2$ satisfies $|\partial(S)| \geq \alpha |S|$, where the *edge boundary* $\partial(S)$ of S is the set of edges with precisely one endpoint in S . Trevisan’s decomposition result asserts⁴ that, given any graph, one can remove few of its edges in order to partition it into connected components that each have good expansion⁵.

Theorem 5 ([30], Lemma 10). *Let $0 < \epsilon \leq 1/2$. From any n -vertex graph one can remove at most an ϵ -fraction of the edges to obtain a graph whose every connected component has edge expansion at least $\epsilon/12 \log n$.*

Using Theorem 1, we show in Corollary 2.5 below that, already for $\epsilon = 1/2$, the expansion guarantee in Theorem 5 cannot be improved much beyond $1/\log n$. In fact, we show that the same conclusion holds even when $\epsilon = 1 - o(1)$, that is, even if one is allowed to remove all but a $o(1)$ -fraction of the edges of the graph! To this end, we will first need the following observation regarding the graphs $G_{n,k}$ from Theorem 1.

Claim 2.4. *Any subgraph of $G_{n,k}$ with average degree at least 4 has at least $n^{1/72k^2}$ vertices.*

Proof. Let H be a t -vertex subgraph of $G_{n,k}$ with average degree at least 4. Recall the well-known fact that the girth of a t -vertex graph of average degree at least 4 is at most $2 \log t$ (in fact, average degree at least 3 suffices, see [1]). Item (ii) in Theorem 1 thus implies $\log n / (6k)^2 \leq 2 \log t$, completing the proof. \square

Corollary 2.5. *Any subgraph of $G_{n,k}$ with average degree at least 4 has a connected component with edge expansion at most $(1/\log n)(15k \log \log n)^2$.*

Proof. Consider a subgraph of $G_{n,k}$ with average degree at least 4, and note that it has a connected component H with average degree at least 4. Writing t for the number of vertices of H , Claim 2.4 implies that $\log t \geq \log n / 72k^2$. It follows from item (iii) in Theorem 1 that the edge expansion of H is at most $(t/\log t)(\log \log t)^2 \cdot (3/t) \leq (1/\log n)(15k \log \log n)^2$. \square

By Corollary 2.5, even after removing a fraction of $1 - 4/k$ of the edges of $G_{n,k}$, the remaining graph has a connected component with edge expansion at most $O((1/\log n)(k \log \log n)^2)$. So for example, taking $k = 8$ gives a bounded-degree graph such that even after removing up to half of its edges, the remaining graph has a connected component with edge expansion at most $O((1/\log n)(\log \log n)^2)$. At the other extreme, if $k = \log \log n / 648$ then even after removing a $(1 - o(1))$ -fraction of the edges, the remaining graph has a connected component with edge expansion at most $(1/\log n)(\log \log n)^4$.

⁴We note that prior to [30], Goldreich and Ron [9] have implicitly proved a result of the same spirit, whose exact quantitative properties are somewhat more complicated to state. See also [12]

⁵Trevisan originally used the notion of conductance; our lower bound applies even for edge expansion.

2.4 Threshold-rank decomposition

We now describe another example of a decomposition result used in devising approximation algorithms for unique games. While Trevisan’s decomposition, mentioned in the previous subsection, guarantees that every connected component has good expansion (meaning that its second eigenvalue is small compared to the first), a different decomposition was devised by Arora, Barak and Steurer [4], which guarantees that every connected component has relatively few eigenvalues larger than a specified threshold. As noted in [4], this decomposition was the main component in their breakthrough paper which gave the first subexponential time algorithm for unique games. We now turn to formally describe the properties of this decomposition.

For a d -regular graph G , denote by $\text{rank}_\tau(G)$ the number (with multiplicities) of eigenvalues λ of the adjacency matrix of G satisfying $|\lambda| > \tau d$. Let $R_n(\eta, \epsilon)$ denote the minimum integer such that for any n -vertex graph G , one can remove at most an ϵ -fraction of its edges so that each connected component H of the new graph satisfies $\text{rank}_{1-\eta}(H^*) \leq R_n(\eta, \epsilon)$, where H^* is obtained from H by adding self-loops to make it $\Delta(G)$ -regular. The decomposition result of [4] shows that one can remove few of the edges of any given graph, so that each connected component of the new graph has small threshold rank. This can be formally stated as follows.

Theorem 6 ([27], Theorem 5.6). *For every $0 < \eta, \epsilon \leq 1$ we have*

$$R_n(\eta, \epsilon) \leq n^{O((\eta/\epsilon^2)^{1/3})} .$$

So for example, for every d -regular n -vertex graph G , one can efficiently remove at most, say, 1% of the edges of G so that for each connected component in the resulting graph, its “regularized” adjacency matrix has at most $n^{O(\eta^{1/3})}$ eigenvalues larger than $(1 - \eta)d$.

The algorithm of [4] for approximating unique games runs in time exponential in $R_n(\eta, \epsilon)$. The bound in Theorem 6 therefore implies a subexponential running time. Note that to get a polynomial running time it thus suffices to prove the bound $R_n(\eta, \epsilon) \leq O(\log n)$ for constant η, ϵ . In fact, the bound $R_n(\eta, \epsilon) \leq n^{o(1)}$ already suffices to disprove the Unique Games Conjecture under the so-called Exponential Time Hypothesis.

Here we prove that the bound on $R_n(\eta, \epsilon)$ in Theorem 6 is in fact essentially tight. Namely, we show that one can derive from Theorem 2 a nearly polynomial lower bound on $R_n(\eta, \epsilon)$.

Corollary 2.6. *There is a positive integer n_0 such that for every $n \geq n_0$ and $0 < \eta \leq 1$ we have*

$$R_n(\eta, 0.99) \geq n^{\Omega(\eta/(\log \log n)^2)} .$$

We note that a proof similar to that of Corollary 2.6 can give roughly the same bound on $R_n(\eta, \epsilon)$ even when $\epsilon = 1 - o(1)$, that is, even when one is allowed to remove all but a $o(1)$ -fraction of the graph’s edges. We omit the details. We also note that we can improve the exponent in Corollary 2.6 to $\Omega(\eta/(\log \log n)^{1+o(1)})$, but we opted for the simpler/cleaner bound. See Section 6 for more details on this.

For the proof of Corollary 2.6 we will need the “easy” direction of the higher-order Cheeger inequality that first appeared in [15] (see Lemma 2.1 there). For completeness, we reproduce the proof in full details in Section 5.

Lemma 2.7. *Let $G = (V, E)$ be a d -regular (multi-)⁶graph, and let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_{|V|}$ be the eigenvalues of the adjacency matrix of G . Then for every k ,*

$$\frac{d - \lambda_k}{2} \leq \min_{S_1, \dots, S_k} \max_{1 \leq i \leq k} \phi_G(S_i),$$

where the minimum is over all collections of k mutually disjoint non-empty subsets of V .

Proof of Corollary 2.6. Let $G = G_{n,400}$ be a graph from Theorem 1. Put $r = \frac{1}{8}n^{\eta/(4000 \log \log n)^2}$. Let H be a subgraph of G with average degree at least 4. We will show that

$$\text{rank}_{1-\eta}(H^*) \geq r. \quad (3)$$

This would prove $R_n(\eta, 0.99) \geq r$ since removing at most a 0.99-fraction of the edges of G leaves a subgraph of average degree at least $0.01 \cdot 400 = 4$, which therefore must have a connected component H for which (3) applies.

Writing $t = |V(H)|$, put $\mu = 1/t^{\eta/(\log \log t)^2}$. Claim 2.4 implies that

$$t \geq n^{1/4000^2} \geq 16 \quad (4)$$

assuming $n \geq n_0$ for an appropriate n_0 . Note that (4) implies that $1/(8\mu) \geq r$, and that we may assume $\mu \leq 1/2$, as otherwise $r \leq 1/(8\mu) < 1$ so (3) trivially holds. Moreover, note that $\mu \geq 1/\sqrt{t}$, since $1/(\log \log t)^2 \leq 1/(\log \log(16))^2 \leq 1/2$ by (4). Apply Theorem 2 on H with μ , using that $1/\sqrt{t} \leq \mu \leq 1/2$, as well as the fact that $t \geq \log n / (6 \cdot 400)^2$ by (4) and the assumption $n \geq n_0$. We thus obtain r mutually disjoint non-empty subsets of H , each with edge expansion in H at most $(\log \frac{1}{\mu} / \log t)(14 \log \log t)^2 < 200\eta$. Let $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_t$ be the eigenvalues of the adjacency matrix of H^* (recall that H^* is obtained by adding to H sufficiently many self-loops so as to make it $\Delta(G)$ -regular). Note that adding self-loops does not alter the expansion of any subset. Therefore, Lemma 2.7 implies that $\Delta(G) - \lambda_r < 400\eta$. It follows that for every $1 \leq i \leq r$ we have $\lambda_i \geq \lambda_r > \Delta(G) - 400\eta \geq (1 - \eta)\Delta(G)$. This implies (3) and therefore completes the proof. \square

2.5 Hyperfinite families of graphs

A graph is said to be (ϵ, q) -hyperfinite if one can remove an ϵ -fraction of its edges and thus decompose it into connected components of size at most q each. A family of graphs is said to be hyperfinite if there is a function q such that for every $\epsilon > 0$, every graph in the family is $(\epsilon, q(\epsilon))$ -hyperfinite. Hyperfinite families of graphs have been extensively studied in recent years, mainly because of their role in the theory of graph limits of sparse graphs (see [21]). Motivated by certain questions related to the design of property-testing algorithms, it was shown in [6] that a hereditary family of graphs in which every graph has a small edge separator must be hyperfinite. More precisely, the following holds.

Theorem 7 ([6], Corollary 3.2). *Let \mathcal{F} be a hereditary family of graphs such that every n -vertex graph $G \in \mathcal{F}$ satisfies $\text{sep}(G) \leq n/(\log n (\log \log n)^2)$. Then, \mathcal{F} is hyperfinite.*

Using Theorem 1 we show that the edge-separation requirement in Theorem 7 cannot be improved much beyond $n/\log n$.

⁶By this we mean that we allow self-loops.

Corollary 2.8. *The graphs $G_{n,8}$ from Theorem 1 are not $(\frac{1}{2}, n^{1/4608})$ -hyperfinite. In particular,*

$$\mathcal{F}_8 := \{G : G \text{ is an induced subgraph of } G_{n,8} \text{ for some } n\}$$

is a hereditary family of graphs that is not hyperfinite, despite the fact that every n -vertex graph $G \in \mathcal{F}_8$ satisfies $\text{sep}(G) \leq (n/\log n)(\log \log n)^2 + 48$.

Proof. To see that the first assertion holds, note that after removing at most half of the edges of $G_{n,8}$, we obtain a graph with average degree at least 4. This graph has a connected component H of average degree at least 4. By Claim 2.4, the number of vertices in H is least $n^{1/4608}$. As to the second assertion of the corollary, note that the first assertion clearly means that \mathcal{F}_8 is not hyperfinite. Also, note that the fact that $\text{sep}(G) \leq \max\{(n/\log n)(\log \log n)^2, 6 \cdot 8\}$ holds for every n -vertex $G \in \mathcal{F}_8$ follows from items (i),(ii) and (iii) of Theorem 1 together with Fact 1.3. \square

2.6 Locally constructing spanning graphs

Motivated by the growing literature on local algorithms (see Rubinfeld et al. [24] and the references in [22]), let us consider the problem of constructing a spanning tree of a graph G in a local manner. By this we mean being able to decide if a given edge of G belongs to the tree in constant time (and in particular, without constructing the entire tree). It is easy to see that there actually cannot exist a local algorithm \mathcal{A} for constructing a spanning tree. Indeed, if G is a cycle then \mathcal{A} must answer negatively on one edge while if G is a path then \mathcal{A} must answer positively on all edges, yet the two graphs cannot be distinguished without making a linear number of queries. Thus, it is natural to require that the constructed subgraph is merely sparse.

Formally, an (ϵ, q) -local sparse spanning graph algorithm makes at most q queries to the incidence-lists representation⁷ of the (bounded-degree, connected) input graph G , and provides query access⁸ to a connected subgraph G' of G that has fewer than $(1 + \epsilon)n$ edges. The question of which graphs have a local spanning graph algorithm that uses only a constant number of queries was studied in [22], where it was shown that the answer is given by the same hereditary notion of expansion considered in the current paper. A graph is said to be f -non-expanding if every t -vertex subgraph H has edge expansion at most $f(t)$. It was shown in [22] that if the input graphs are f -non-expanding with $f(x) = \Omega(1/(\log x (\log \log x)^2))$ then a “localized” version of Kruskal’s algorithm constructs a sparse spanning graph in a constant number of queries. As for the more challenging lower bound question, the following theorem shows that the above expansion requirement is essentially sharp (even if the algorithm is allowed to use randomization).

Theorem 8 ([22], Theorem 3). *For infinitely many n , there is an f -non-expanding n -vertex graph G with $f(x) = (1/\log x) \cdot (70 \log \log x)^2$ such that even if the input graph is guaranteed to be isomorphic to G , every $(\frac{1}{2}, q)$ -local sparse spanning graph algorithm satisfies $q \geq \Omega(\log \log n)$.*

The proof of Theorem 8 relies heavily on Theorem 1. It uses $G_{n,k}$ in order to construct a 3-regular graph that is still non-expanding and of girth $\Omega(\log \log n)$. For that graph, it is shown that any local sparse spanning graph algorithm must query a number of edges proportional to the girth.

⁷Queries are of the form “which is the i -th neighbor of v ?”.

⁸I.e., on input $(u, v) \in E$ the algorithm returns whether $(u, v) \in E(G')$, and for any sequence of queries it answers consistently with the same G' .

3 Proof of Theorem 1

To construct the graphs $G_{n,k}$ of Theorem 1 we take an appropriate random subgraph of the Boolean hypercube. Using standard probabilistic and combinatorial arguments, we show that this random subgraph satisfies items (i) and (ii) of Theorem 1 (see Subsection 3.1). For the proof of the main property, item (iii), we use a technique similar to the one Linial, London and Rabinovich [17] used to prove that any embedding of an expander into ℓ_1 has logarithmic distortion. We prove that every low-degree t -vertex subgraph of the hypercube has edge expansion at most roughly $1/\log t$ (see Lemma 3.3). This gives a small cut in the subgraph, which we finally “boost” into a small balanced cut, or in other words, a small edge separator (see Subsection 3.3).

3.1 Construction

The d -cube, denoted Q_d , is the graph with vertex set $\{0, 1\}^d$ where two vertices are adjacent if their corresponding vectors differ in exactly one coordinate. Notice Q_d is a d -regular graph on 2^d vertices. The following gives the construction of $G_{n,k}$ and the proof of items (i) and (ii) of Theorem 1.

Lemma 3.1. *For any $d \in \mathbb{N}$ and $2 \leq k \leq d$ there is a subgraph $Q = Q_{d,k}$ of the d -cube, on all 2^d vertices, such that:*

- i. The average degree of Q is at least k .*
- ii. The maximum degree of Q is at most $3k$.*
- iii. The girth of Q is at least $d/(3k)^2$.*

We will need the following upper bound on the number of cycles of a given length in the cube.

Claim 3.2. *The number of cycles of length 2ℓ in the d -cube is at most $2^d(d\ell)^\ell$.*

Proof. We will show that the number of closed walks of length 2ℓ , starting and ending at a given vertex, is at most $(d\ell)^\ell$; this would immediately imply the stated bound. We claim that each such closed walk corresponds to a sequence $(x_1, \dots, x_{2\ell}) \in [d]^{2\ell}$ with the property that $|\{i \in [2\ell] : x_i = t\}|$ is even for every $t \in [d]$. To see this, recall that each vertex in the graph corresponds to an element of $\{0, 1\}^d$. The claim follows by considering, for each edge along a closed walk beginning (and ending) with a given vertex, the unique index in which the bit “flips”. So all that is left is to give an upper bound on the number of sequences as above. Observe that all these sequences can be generated by first partitioning the set of indices $\{1, \dots, 2\ell\}$ into ℓ pairs and then assigning to each pair a value from $\{1, \dots, d\}$ (of course, this process will generate some sequences several times). Since the number of ways one can pair the elements of $\{1, \dots, 2\ell\}$ is given by

$$(2\ell - 1)!! = (2\ell - 1)(2\ell - 3) \cdots 3 \cdot 1 \leq 2^\ell \ell! / 2 \leq 2^\ell (\ell/2)^\ell = \ell^\ell,$$

we get that the number of sequences is bounded from above by $d^\ell \cdot \ell^\ell$. □

To prove the existence of a graph as in Lemma 3.1 we apply standard arguments using the so-called probabilistic “deletion method”.

Proof of Lemma 3.1. Denote by G' the random subgraph of the d -cube where each edge is independently retained with probability $p = 3k/d$. Let the random variable X count the number of edges of G' . Then

$$\mathbb{E}[X] = p \cdot d2^{d-1} = 3k \cdot 2^{d-1}. \quad (5)$$

Set $L = d/9k^2$ and let the random variable Y count the number of cycles of length at most L in G' . By Claim 3.2,

$$\mathbb{E}[Y] \leq \sum_{\ell=2}^{L/2} p^{2\ell} 2^d (d\ell)^\ell = 2^d \sum_{\ell=2}^{L/2} (p^2 d\ell)^\ell = 2^d \sum_{\ell=2}^{L/2} (\ell/L)^\ell \leq 2^d \sum_{\ell=2}^{L/2} (1/2)^\ell \leq 2^{d-1}. \quad (6)$$

Let the random variable Z count the total number of “excess” edges in G' , that is,

$$Z = \sum_{v: \deg(v) > 3k} (\deg(v) - 3k).$$

We claim that $\mathbb{E}[Z] \leq \sqrt{3k} \cdot 2^{d-1}$. Indeed, for each vertex v the random variable $\deg(v)$ follows the binomial distribution $B(d, 3k/d)$, so $\mathbb{E}[\deg(v)] = 3k$. We have $\mathbb{E}[Z] = \frac{1}{2} \sum_v \mathbb{E} |\deg(v) - 3k|$, since $\sum_v |\deg(v) - 3k| = \sum_v (3k - \deg(v)) + 2Z$. By Jensen’s inequality, for each v we have

$$(\mathbb{E} |\deg(v) - 3k|)^2 \leq \mathbb{E}[(\deg(v) - 3k)^2] = \text{Var}[\deg(v)] \leq 3k,$$

implying that

$$\mathbb{E}[Z] = \frac{1}{2} \sum_v \mathbb{E} |\deg(v) - 3k| \leq \sqrt{3k} \cdot 2^{d-1}, \quad (7)$$

as claimed. Combining (5), (6) and (7) we get

$$\mathbb{E}[X - Y - Z] \geq (3k - 1 - \sqrt{3k}) \cdot 2^{d-1} \geq k \cdot 2^{d-1}, \quad (8)$$

where the last inequality can be easily checked to hold for any $k \geq 2$. Let Q be obtained from G' by removing an arbitrary edge from each cycle of length at most $d/9k^2$, as well as removing, for each vertex v with $\deg(v) > 3k$, arbitrary $\deg(v) - 3k$ adjacent edges. Clearly, Q satisfies the last two requirements in the statement. Moreover, we have from (8) that the expected average degree of Q is at least k . The existence of a subgraph of the d -cube as required immediately follows. \square

3.2 Expansion in the hypercube

In this subsection we prove that large subgraphs of the d -cube are not good edge expanders. Recall that a graph $G = (V, E)$ is said to have edge expansion α if every subset $S \subseteq V$ with $|S| \leq |V|/2$ satisfies $|\partial(S)| \geq \alpha |S|$ (the edge boundary $\partial(S)$ is the set of edges with exactly one vertex in S). Our goal in this subsection is to prove the following lemma.

Lemma 3.3. *Any (not necessarily induced) subgraph of the d -cube with t vertices and average degree r has edge expansion at most $2r \log d / \log(t/2)$.*

We first need the following easy claim, in which the *distance* between two vertices u, v in a graph G , denoted $\delta_G(u, v)$, is the length of a shortest path connecting u and v .

Claim 3.4. For every graph G of maximum degree $\Delta > 1$, and every t -vertex subset $S \subseteq V(G)$, the average distance $\sum_{\{u,v\} \in \binom{S}{2}} \delta_G(u,v) / \binom{t}{2}$ is at least $\log(t/2)/(2 \log \Delta)$.

Proof. Let $v \in S$. We claim that there are at least $t/2$ vertices in S of distance at least $\ell = \log(t/2)/\log \Delta$ from v . Indeed, the number of vertices of distance at most $\ell - 1$ from v is at most $\sum_{i=0}^{\ell-1} \Delta^i < \Delta^\ell = t/2$. It follows that the average distance is at least $\ell/2$. \square

We are now ready to prove the main result of this subsection.

Proof of Lemma 3.3. Let $H = (S, E)$ be a t -vertex subgraph of the d -cube Q_d of average degree r , and let α denote the edge expansion of H . We need to show that $\alpha \leq 2r \log d / \log(t/2)$. Note that we may assume $\alpha > 0$, as otherwise there is nothing to prove. For each vertex $v \in S$, write $(v_1, \dots, v_d) \in \{0, 1\}^d$ for the corresponding binary vector. Notice $\delta_{Q_d}(u, v) = \sum_{i=1}^d |u_i - v_i|$. Observe that

$$\frac{\sum_{\{u,v\} \in E} \delta_{Q_d}(u, v)}{\sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v)} = \frac{\sum_{i=1}^d \sum_{\{u,v\} \in E} |u_i - v_i|}{\sum_{i=1}^d \sum_{\{u,v\} \in \binom{S}{2}} |u_i - v_i|} \geq \min_i \frac{\sum_{\{u,v\} \in E} |u_i - v_i|}{\sum_{\{u,v\} \in \binom{S}{2}} |u_i - v_i|}, \quad (9)$$

with the minimum over all i for which the denominator is nonzero, where we used the elementary inequality $(\sum_{i=1}^d x_i) / (\sum_{i=1}^d y_i) \geq \min_{i: y_i \neq 0} x_i / y_i$ which holds for all non-negative reals $x_1, y_1, \dots, x_d, y_d$. Let $i \in [d]$ achieve the minimum in (9), and set $T = \{v \in S : v_i = 1\}$. Note that $0 < |T| < |S|$ and that we can assume $0 < |T| \leq |S|/2$, as otherwise we replace T with $S \setminus T$. Since in the right hand side of (9) the numerator is $|\partial T|$ and the denominator is $|T| |S \setminus T|$, we deduce that

$$\frac{\sum_{\{u,v\} \in E} \delta_{Q_d}(u, v)}{\sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v)} \geq \frac{|\partial T|}{|T| |S \setminus T|} \geq \min_{\substack{S' \subseteq S: \\ 0 < |S'| \leq |S|/2}} \frac{|\partial S'|}{|S'| |S \setminus S'|} \geq \frac{\alpha}{|S| - 1}.$$

Therefore,

$$\frac{1}{\binom{|S|}{2}} \sum_{\{u,v\} \in \binom{S}{2}} \delta_{Q_d}(u, v) \leq \frac{2}{\alpha |S|} \sum_{\{u,v\} \in E} \delta_{Q_d}(u, v) = \frac{2|E|}{\alpha |S|} = \frac{r}{\alpha}. \quad (10)$$

Applying Claim 3.4 with $G = Q_d$ and the set S , the left hand side of (10) is at least $\log(t/2)/2 \log d$. The desired bound on α follows. \square

3.3 Putting it all together

To prove item (iii) of Theorem 1 we will show that the graph G constructed in Lemma 3.1 also has, in addition to items (i) and (ii) of Theorem 1, the property that every (large enough) subgraph H has a small edge separator. Towards proving this, we first show that every such H has small expansion. We consider two cases: (1) H has at most $2d^{1/3}$ vertices and (2) H has at least $2d^{1/3}$ vertices. In the first case we use the fact that high expansion implies the existence of a short cycle (see Lemma 3.5 below), which contradicts the girth property of G . In the second case, as the maximum degree of G is small, we use Lemma 3.3 to bound the expansion of the t -vertex subgraph H by roughly $(\log d)^2 / \log t$, which in this case is at most roughly $(\log \log t)^2 / \log t$. Finally, seeing

as the above holds for every subgraph H , we show that it is possible to boost this “hereditary” non-expansion property of G in order to construct an edge separator in H that is as small as required in (1) (see Lemma 3.6 below).

To execute the above proof strategy we will need the following two lemmas, whose proofs appear in Section 5. Recall that we use $\text{sep}(G)$ to denote the size of the smallest edge separator in G .

Lemma 3.5. *Every connected n -vertex graph with edge expansion α (> 0) and maximum degree Δ that is not a tree has girth at most $20\Delta \log(n)/\alpha$.*

Lemma 3.6. *Let G be an n -vertex graph whose every t -vertex subgraph has edge expansion at most $f(t)$ for every $t \geq 2n/3$, where $f : [2n/3, n] \rightarrow \mathbb{R}$ is decreasing. Then $\text{sep}(G) \leq (2n/3)f(2n/3)$.*

Finally, note that the assumption $2 \leq k \leq \frac{1}{648} \log \log n$ in Theorem 1 trivially yields the relation $\log n \geq 2^{2 \cdot 648}$, which in particular implies

$$\log n \geq (\log \log n)^6. \quad (11)$$

We now give the proof of our main theorem.

Proof of Theorem 1. First, suppose that n is a power of 2, and write $n = 2^d$. Let $G = Q_{d,2k}$ be a subgraph of the d -cube as guaranteed by Lemma 3.1. The assertion of Lemma 3.1 implies that G has average degree at least $2k$, maximum degree at most $6k$, and girth at least $d/(6k)^2$. So to complete the proof (for $n = 2^d$) we only need to establish item (iii) of Theorem 1.

For the rest of the proof set $t_0 = \frac{2}{3}d/(6k)^2$ (≥ 256). We will use the following inequalities, which can be deduced from (11) and the theorem’s assumption that $k \leq \log d/648$;

$$d^{1/3}/(6k)^3 \geq 20, \quad (12)$$

$$6k \leq t_0/(3 \log t_0). \quad (13)$$

Put $f(x) = (1/\log x)(\log \log x)^2$, and note that $f(x)$ is increasing for $x \geq 256$.

We first show that every t -vertex subgraph H of G with $t \geq t_0$ has edge expansion at most $f(t)$. If H is cycle-free then $\text{sep}(H) \leq \Delta(G) \leq 6k \leq t_0/(3 \log t_0) \leq t/(3 \log t)$, where in the first inequality we used Fact 1.3, in the third inequality we used (13) and in the last inequality the fact that the function $x/\ln x$ is increasing. It follows that the edge expansion of H is at most $(\frac{1}{3}t/\log t)/(\frac{1}{3}t) = 1/\log t \leq f(t)$, as desired. Thus, we henceforth assume that H contains a cycle. We next consider two cases. Suppose first that $t < 2^{d^{1/3}}$. Assuming for contradiction that H has edge expansion at least $1/\log t$, Lemma 3.5 implies that H contains a cycle of length at most $20(6k) \log^2(t) \leq 20(6k)d^{2/3} \leq d/(6k)^2$, where the last inequality uses (12). This contradicts the fact that G does not contain a cycle this short, hence we deduce that the edge expansion of H is at most $1/\log t \leq f(t)$, as needed. Suppose now that $t \geq 2^{d^{1/3}}$. By Lemma 3.3 and the fact that the average degree of H is at most $6k$, the edge expansion of H is at most $24k \log d/\log t$. Thus, to prove our claim it suffices to show that $\sqrt{24k \log d} \leq \log \log t$. This indeed follows from the theorem’s assumption that $k \leq \log d/648$, as it implies that

$$\sqrt{24k \log d} \leq (\log d)/3 = \log(d^{1/3}) \leq \log \log t.$$

Having established that every subgraph of G on at least t_0 vertices has small edge expansion, we now wish to show that every t -vertex subgraph H of G with $t \geq d/(6k)^2 (= 3t_0/2)$ has a small edge separator. We apply Lemma 3.6 on the graph H with the function f , noting that, as required by the lemma, every t' -vertex subgraph of H with $t' \geq 2t/3 (\geq t_0)$ has edge expansion at most $f(t')$, and that $f : [t_0, t] \rightarrow \mathbb{R}$ is decreasing. We deduce that

$$\text{sep}(H) \leq (2t/3)f(2t/3) \leq ((2t/3)/\log(2t/3))(\log \log t)^2 \leq (t/\log t)(\log \log t)^2,$$

as required by item (iii) in Theorem 1. This completes the proof of the theorem for the case $n = 2^d$.

Finally, let us consider the case of an arbitrary n , that is, not necessarily a power of 2. In this case we set $d = \lceil \log n \rceil$. Since $n > 2^d/2$, a random n -vertex subgraph of $Q_{d,2k}$ will have, with positive probability, average degree at least k . Let G be such a graph. Then the maximum degree of G is still at most $6k$, its girth is at least $\log n/(6k)^2$, and it is easy to see that every t -vertex subgraph H of G with $t \geq \log n/(6k)^2$ satisfies $\text{sep}(H) \leq (t/\log t)(\log \log t)^2$. This completes the proof. \square

4 Proof of Theorem 2

In order to prove Theorem 2 we first show that the hereditary nature of item (iii) of Theorem 1 can be used to iteratively find smaller and smaller subsets while controlling their expansion. This enables us to find a single small subset having small expansion. We note that this special case of Theorem 2, stated as Claim 4.1 below, already suffices for the application in Corollary 2.2. To prove Theorem 2 we first iteratively pick sets A_1, A_2, \dots via Claim 4.1 in such a way that each A_i has small expansion in the graph obtained by removing A_1, \dots, A_{i-1} from H . We then complete the proof by showing that at least half of these sets have small expansion in H itself. For the rest of this section, G is the graph from Theorem 2 and $t_0 = \log n/(6k)^2$.

Claim 4.1. *For every t -vertex subgraph H of G with $t \geq \frac{1}{2}t_0$, and every $1/\sqrt{t} \leq \mu \leq 2/3$, there is a subset $A \subseteq V(H)$ of size $\mu t/3 \leq |A| \leq \mu t$ satisfying*

$$\phi_H(A) \leq \frac{\log(1/\mu)}{\log t} \cdot (4 \log \log t)^2. \quad (14)$$

Proof. Put $f(x) = (1/\log x) \cdot (\log \log x)^2$. One can check that $f(x)$ is decreasing for $x \geq 256$. Throughout the proof, unless otherwise mentioned, we write $\phi(\cdot) = \phi_H(\cdot)$ for the edge expansion in H , and $\partial(\cdot) = \partial_H(\cdot)$ for the edge boundary in H . In order to obtain a subset $A \subseteq V(H)$ satisfying (14) we will iteratively find smaller and smaller subsets, such that in each step the edge expansion (in H) of our subset does not increase by much. One can verify that, by the theorem's assumption on k and t , we have by (11) the inequality

$$\Delta(G) \leq \frac{\sqrt{t_0/2}}{\log \log n}. \quad (15)$$

Let $S \subseteq V(H)$ with $|S| \geq \sqrt{t}$ (≥ 256). We claim that there is a subset $S_1 \subseteq S$ of size $\frac{1}{3}|S| \leq |S_1| \leq \frac{2}{3}|S|$ such that $\phi(S_1) \leq \phi(S) + 2f(|S|)$. Before proving this claim, note that the induced subgraph $H[S]$ satisfies

$$\text{sep}(H[S]) \leq |S| f(|S|). \quad (16)$$

Indeed, if $|S| \geq t_0$ this is item (iii) of Theorem 1, while if $|S| < t_0$ we have

$$\text{sep}(H[S]) \leq \Delta(H) \leq \sqrt{t}/\log(t_0) \leq |S| f(|S|) ,$$

where the first inequality follows from item (ii) of Theorem 1 and Fact 1.3, and the second inequality from (15) and the theorem's assumption $t \geq \frac{1}{2}t_0$. To prove our claim, consider an edge separator in $H[S]$. Namely, let $S = S' \cup S''$ be a partition of S with $|S'|, |S''| \leq \frac{2}{3}|S|$ such that the number of edges between S' and S'' in $H[S]$, and thus in H , is $\text{sep}(H[S])$. Consider now the edge boundary of S' and S'' in H . We have

$$|\partial(S')| + |\partial(S'')| = |\partial(S)| + 2 \text{sep}(H[S]) .$$

Moving to edge expansion, observe that this means that

$$\min\{\phi(S'), \phi(S'')\} \leq \frac{|\partial(S')| + |\partial(S'')|}{|S'| + |S''|} = \frac{|\partial(S)| + 2 \text{sep}(H[S])}{|S|} \leq \phi(S) + 2f(|S|) ,$$

where in the left inequality we used the elementary inequality $(a+b)/(c+d) \geq \min\{a/c, b/d\}$, and in the right inequality we used (16). Thus, we can take either S' or S'' as the set S_1 above, proving our claim.

Applying our claim with $S = S_0 := V(H)$ yields a subset $S_1 \subseteq S_0$ of size $\frac{1}{3}|S_0| \leq |S_1| \leq \frac{2}{3}|S_0|$ satisfying $\phi(S_1) \leq \phi(S_0) + 2f(|S_0|) = 2f(|S_0|)$, where we used the fact that $\phi(S_0) = 0$. Having picked S_i , for some $i \geq 1$, we can then pick a new set $S_{i+1} \subseteq S_i$ satisfying $\frac{1}{3}|S_i| \leq |S_{i+1}| \leq \frac{2}{3}|S_i|$ and $\phi(S_{i+1}) \leq \phi(S_i) + 2f(|S_i|)$. Let j be the smallest integer such that $|S_j| \leq \mu t$. Note that $|S_j| \geq \mu t/3$ and that $j \leq \lceil \log_{3/2}(1/\mu) \rceil$. We get from the above observations that

$$\phi(S_j) \leq \sum_{i=0}^{j-1} 2f(|S_i|) \leq j \cdot 2f(\mu t) \leq 2 \log_{3/2}(1/\mu) \cdot 2f(\sqrt{t}) \leq \frac{\log(1/\mu)}{\log t} \cdot 16(\log \log t)^2 ,$$

where in the second inequality we used the fact that $f(x)$ is decreasing and in the third inequality we used both bounds in the assumption $1/\sqrt{t} \leq \mu \leq 2/3$ (the upper bound is used for the inequality $\lceil x \rceil \leq x + 1 \leq 2x$ for $x \geq 1$). Taking S_j to be the set A in the statement completes the proof. \square

Proof of Theorem 2. Let $H = (V, E)$ be a t -vertex subgraph with $t \geq t_0$, and recall $1/\sqrt{t} \leq \mu \leq 1/2$. We construct $r := \lceil 1/(4\mu) \rceil$ mutually disjoint sets $A_1, \dots, A_r \subseteq V$, each of size in $[\mu t/3, \mu t]$, as follows. Denoting $H_{i-1} = H[V \setminus \bigcup_{j=1}^{i-1} A_j]$ and $t_i = |V(H_{i-1})|$, we obtain A_i by applying Claim 4.1 on H_{i-1} with $\mu_i = \mu t/t_i$. To see why we may do so (i.e why we can apply Claim 4.1 this many times), note that $t_i \geq t - (i-1)\mu t \geq t - (r-1)\mu t \geq \frac{3}{4}t$, hence $t_i \geq \frac{1}{2}t_0$ and

$$1/\sqrt{t_i} \leq \sqrt{t}/t_i \leq \mu_i = \mu t/t_i \leq 4\mu/3 \leq 2/3 .$$

It follows that for every $1 \leq i \leq r$ we have

$$\phi_{H_{i-1}}(A_i) \leq \frac{\log(t_i/\mu t)}{\log t_i} (4 \log \log t_i)^2 \leq \frac{\log(1/\mu)}{\log t} (4 \log \log t)^2 , \quad (17)$$

where in the last inequality we used the fact that the function $\log(x/a)(4 \log \log x)^2/\log x$ is increasing for $x \geq a$.

To finish the proof we show that the expansion of at least half of the r subsets A_i is small in H as well. First, note that

$$\sum_{i=1}^r |\partial_H(A_i)| \leq 2 \sum_{i=1}^r |\partial_{H_{i-1}}(A_i)| = 2 \sum_{i=1}^r |A_i| \cdot \phi_{H_{i-1}}(A_i) \leq 2r \cdot \mu t \cdot \frac{\log(1/\mu)}{\log t} (4 \log \log t)^2, \quad (18)$$

where in the first inequality we used the fact that each edge is counted (at most) twice, and in the last inequality we used (17) and the fact that each A_i is of size at most μt . Dividing both sides of (18) by $\mu t/3$ (which is a lower bound on the size of the sets A_i), we get

$$\sum_{i=1}^r \phi_H(A_i) \leq \sum_{i=1}^r \frac{|\partial_H(A_i)|}{\mu t/3} \leq 6r \cdot \frac{\log(1/\mu)}{\log t} (4 \log \log t)^2.$$

Hence, by averaging, at least $\frac{1}{2}r \geq 1/(8\mu)$ of the sets A_i satisfy

$$\phi_H(A_i) \leq 12 \frac{\log(1/\mu)}{\log t} (4 \log \log t)^2 \leq \frac{\log(1/\mu)}{\log t} (14 \log \log t)^2,$$

thus completing the proof. \square

5 Missing Proofs

5.1 Proof of Lemma 2.3

Proof of Lemma 2.3. Root T at an arbitrary vertex, and start a walk from the root down the tree. At each step, move to an arbitrary child that has more than a vertices in the subtree rooted at it, if such a child exists. Clearly, this walk ends at some vertex v . Thus, v has at least a vertices in its subtree (more if v is not the root), and none of its children have more than a vertices. Let n_1, \dots, n_k be the number of vertices in the subtrees rooted at each of the k children of v , meaning $\sum_{i=1}^k n_i \geq a$ and $n_i \leq a$ for every i . Let $I \subseteq [k]$ maximize $s(I) := \sum_{i \in I} n_i$ under the restriction that $s(I) \leq a$. Then $s(I) \geq a/2$, for otherwise $a_j \geq a/2$ for some j , which implies $a/2 \leq s(\{j\}) \leq a$, a contradiction. Let $S \subseteq V(T)$ be the set of all vertices in the subtrees corresponding to I . Then $a/2 \leq |S| \leq a$ and, since the only vertex in the boundary of a subtree is its parent, only v lies in the boundary of S . This completes the proof. \square

5.2 Proof of Lemma 3.5

Proof of Lemma 3.5. Let $e = uv$ be an edge in our graph $G = (V, E)$ that lies on a cycle, and let G' be obtained from G by removing e . We claim that the edge expansion of G' is at least $\alpha/2$. Indeed, this follows from the fact that for every $S \subseteq V$ we have $|\partial_{G'}(S)| \geq |\partial_G(S)| - 1 \geq |\partial_G(S)|/2$, where the last inequality uses the fact that $|\partial_G(S)| \geq 2$ since e is not a cut-edge in G .

Put $\beta = \alpha/(2\Delta)$ and $n' = \lfloor n/2 \rfloor + 1$. For a vertex $x \in V$ and a non-negative integer r , write $B_x(r) = \{y : \delta_{G'}(x, y) \leq r\}$. We claim that $|B_x(r)| \geq n'$ for every $x \in V$ and $r' := \lceil \log(n')/\log(1 + \beta) \rceil$. More generally, we claim that $|B_x(r)| \geq \min\{(1 + \beta)^r, n'\}$. We prove this by

induction on r . For the induction basis we have $|\mathcal{B}_x(0)| = 1$, and for the induction step we have, provided $|B_x(r)| \leq n/2$, that

$$|B_x(r+1)| \geq |B_x(r)| + |\partial_{G'}(B_x(r))|/\Delta \geq (1+\beta)|B_x(r)| \geq (1+\beta)^{r+1}.$$

Having completed the inductive proof, we in particular deduce for the endpoints of the edge $e = uv$ above that $B_u(r') \cap B_v(r') \neq \emptyset$, implying that $\delta_{G'}(u, v) \leq 2r'$. Since u and v are adjacent in G but not in G' , we conclude that G has a cycle of length at most $2r'+1 \leq 5 \log(n)/\log(1+\beta) \leq 20\Delta \log(n)/\alpha$, where the last inequality uses the bound $\log(1+\beta) \geq \beta/2$ for $0 \leq \beta \leq 1$. \square

5.3 Proof of Lemma 3.6

Proof of Lemma 3.6. Iteratively construct subsets $S_1, \dots, S_k \subseteq V(G)$ as follows. To obtain S_i , consider the induced subgraph $G_i = G[V \setminus S_{(i-1)}]$, where $S_{(i-1)} := \bigcup_{j=1}^{i-1} S_j$, and let $S_i \subseteq V(G_i)$ satisfy $|S_i| \leq n_i/2$ and $\phi_{G_i}(S_i) \leq f(n_i)$, where $n_i = |V(G_i)|$. Stop once $|S_{(k)}| \geq n/3$. Notice $n_k = n - |S_{(k-1)}|$ and $|S_{(k-1)}| \leq n/3$, which implies $n_k \geq 2n/3$. Therefore,

$$|S_{(k)}| \leq |S_{(k-1)}| + n_k/2 = (n + |S_{(k-1)}|)/2 \leq 2n/3.$$

Put $S = S_{(k)}$. Note that $\text{sep}(G) \leq |\partial_G(S)|$, thus it suffices to show $\phi_G(S) \leq f(2n/3)$. Observe that every edge in the edge boundary $\partial_G(S)$ is a member of some edge boundary $\partial_{G_i}(S_i)$. Hence,

$$\frac{|\partial_G(S)|}{|S|} \leq \frac{\sum_{i=1}^k |\partial_{G_i}(S_i)|}{|S|} = \sum_{i=1}^k \frac{|S_i|}{|S|} \phi_{G_i}(S_i) \leq \max_{1 \leq i \leq k} \phi_{G_i}(S_i) \leq \max_{1 \leq i \leq k} f(n_i) \leq f(2n/3),$$

where in the last inequality we used the monotonicity of f . This completes the proof. \square

5.4 Proof of Lemma 2.7

Proof of Lemma 2.7. Put $L = dI - A$ where A is the adjacency matrix of $G = (V, E)$. Note that the k -th eigenvalue of L is $d - \lambda_k$. By the Courant-Fischer min-max theorem we have

$$d - \lambda_k = \min_{w^1, \dots, w^k} \max_{0 \neq x \in \text{span}\{w^1, \dots, w^k\}} R_L(x), \quad (19)$$

where the minimum is over all collections of k mutually orthogonal nonzero vectors in $\mathbb{R}^{|V|}$, and $R_L(x) = x^t L x / x^t x$ is the Rayleigh quotient. Let $S_1, \dots, S_k \subseteq V$ be mutually disjoint non-empty subsets, and denote by $\mathbf{1}_{S_i}$ the characteristic $\{0, 1\}$ -vector of S_i . As it is easy to see that $x^t L y = \sum_{\{u, v\} \in E} (x_u - x_v)(y_u - y_v)$, we deduce that $\mathbf{1}_{S_i}^t L \mathbf{1}_{S_i} = |\partial(S_i)|$, and that for $i \neq j$ we have $\mathbf{1}_{S_i}^t L \mathbf{1}_{S_j} = -e(S_i, S_j)$ where $e(S_i, S_j)$ is the number of edges between S_i and S_j . Consider now any vector $x \in \text{span}\{\mathbf{1}_{S_1}, \dots, \mathbf{1}_{S_k}\}$ and write $x = \sum_{i=1}^k c_i \mathbf{1}_{S_i}$. Then

$$\begin{aligned} x^t L x &= \sum_{i,j=1}^k c_i c_j \cdot \mathbf{1}_{S_i}^t L \mathbf{1}_{S_j} = \sum_{i=1}^k c_i^2 |\partial(S_i)| - \sum_{i \neq j} c_i c_j e(S_i, S_j) \leq \sum_{i=1}^k c_i^2 |\partial(S_i)| + \sum_{i \neq j} \frac{c_i^2 + c_j^2}{2} e(S_i, S_j) \\ &= \sum_{i=1}^k c_i^2 |\partial(S_i)| + \sum_{i=1}^k c_i^2 \sum_{j \neq i} e(S_i, S_j) \leq 2 \sum_{i=1}^k c_i^2 |\partial(S_i)|. \end{aligned}$$

The proof now follows since (19) implies

$$d - \lambda_k \leq \max_{0 \neq x \in \text{span}\{\mathbf{1}_{S_1}, \dots, \mathbf{1}_{S_k}\}} R_L(x) \leq \frac{2 \sum_{i=1}^k c_i^2 |\partial(S_i)|}{\sum_{i=1}^k c_i^2 |S_i|} \leq 2 \max_{1 \leq i \leq k} \phi_G(S_i).$$

□

6 Concluding Remarks and Open Problems

- As we mentioned in Section 1, we can replace the $(\log \log t)^2$ factors in (1) with $(\log \log t)^{1+o(1)}$. It would be interesting to know whether the $\log \log t$ factors can be removed. One way to obtain this is to improve Lemma 3.3. In this regard, we conjecture the following “higher-order” isoperimetric property of the hypercube: Among all subgraphs of the d -cube on 2^i vertices, for every $i \in \mathbb{N}$ at least $d^{1/3}$ say, the subcubes have the largest normalized⁹ edge expansion. (One can also come up with an analogous isoperimetric conjecture for vertex expanders in the hypercube.) Since a subcube on $t = 2^i$ vertices has normalized edge expansion $1/i = 1/\log t$, proving the above conjecture would mean that the additional $\log d$ factor in Lemma 3.3 is unnecessary for $t \geq 2^{d^{1/3}}$ (without the lower bound on t , Lemma 3.3 is tight as witnessed by $K_{1,d}$). By the proof of Theorem 1, this would mean that the graphs $G_{n,k}$ are such that the $\log \log t$ terms in (1) and (2) can in fact be replaced by $O(k)$. Note that this is best possible, as witnessed by the subgraph $G_{n,k}$ itself.

Proving the above conjecture would essentially close the gaps in the applications mentioned in Section 2, as k is chosen there to be either constant or grow arbitrarily small with n . For example, this would imply that the $n^{\Omega(\eta/(\log \log n)^2)}$ bound on the threshold rank in Corollary 2.6 can be improved to a truly polynomial bound $n^{\Omega(\eta)}$, thus showing that one cannot improve the decomposition of [4] in order to approximate unique games in truly subexponential time $\exp(n^{o(1)})$.

- It is not hard to see that if every t -vertex subgraph H of a graph G satisfies $\text{sep}(H) \leq (t/\log t)(\log \log t)^c$ then G can have at most $O(n(\log \log n)^{c+1})$ edges. (Heuristically speaking, the upper bound on the number of edges admits a recursion of the form $f(n) \leq 2f(n/2) + (n/\log n)(\log \log n)^c$). We note that one can construct a graph G satisfying this requirement (i.e., for any given c) using arguments similar to those used to prove Theorem 1, and in this case the resulting graph would have $\Omega(n(\log \log n)^{c-1})$ edges. This means that, up to $\log \log n$ factors, the graph we construct in Theorem 1 has the maximum possible number of edges.

The situation for vertex separators, as in Theorem 3, is not so clear, so it would be interesting to understand the maximum number of edges of n -vertex graphs in a hereditary family with vertex separators of size at most $(n/\log n)(\log \log n)^c$. Although we can construct such a family \mathcal{F} that contains n -vertex graphs with $\Omega(n(\log \log n)^{c-1})$ edges, we cannot rule out the possibility of improving this to (say) $n \log n$.

- Let us mention two problems that seem to be related to the types of problems studied here but that (unfortunately) we cannot resolve. The first is a nice problem of G. Kalai. Suppose

⁹I.e., edge expansion divided by the subgraph’s average degree.

\mathcal{F} is a hereditary family of graphs and that every graph in the family has a vertex separator of size $n/f(n)$. Then, how fast should $f(n)$ grow so that \mathcal{F} has only $2^{O(n)}$ non-isomorphic graphs on n vertices? It was shown in [10, 11] that $f(n) \geq \log^{2+\epsilon} n$ suffices while $f(n) \leq \log^{1-\epsilon} n$ does not. Closing the gap between these bounds is still open.

The second problem is related to a graph decomposition result from a work of Pătraşcu and Thorup [23]. They proved that the edge set of every graph can be decomposed into $b = O(\log n)$ subsets E_1, \dots, E_b so that for every $1 \leq i \leq b$ the graph spanned by E_1, \dots, E_i has “edge expansion” $1/\log n$, where the notion of edge expansion used here is slightly different from the (standard) one we used throughout this paper. It would be interesting to decide whether the parameters in the construction of [23] are optimal.

Acknowledgment: We are grateful to Alex Samorodnitsky for very helpful discussions related to this work.

References

- [1] N. Alon, S. Hoory and N. Linial, The Moore bound for irregular graphs, *Graph. Combinator.* **18** (2002), 53–57. [2.3](#)
- [2] N. Alon, P. Seymour and R. Thomas, Planar separators, *SIAM J. Discrete Math.* **7** (1994), 184–193. [2.1](#)
- [3] N. Anari and S.O. Gharan, Effective-Resistance-Reducing Flows, Spectrally Thin Trees, and Asymmetric TSP, *Proc. FOCS 2015*, 20–39. [1](#)
- [4] S. Arora, B. Barak and D. Steurer, Subexponential Algorithms for Unique Games and Related Problems, *J. ACM* **62** (2015), 1–25. [1](#), [1.1](#), [2.3](#), [2.4](#), [2.4](#), [6](#)
- [5] D. Conlon, J. Fox and B. Sudakov, Cycle packing, *Random Struct. Algor.* **45** (2014), 608–626. [1](#)
- [6] A. Czumaj, A. Shapira and C. Sohler, Testing hereditary properties of non-expanding bounded-degree graphs, *SIAM J. Comput.* **38** (2009), 2499–2510. [1](#), [2.5](#), [7](#)
- [7] J. Fox and J. Pach, Separator theorems and Turán-type results for planar intersection graphs, *Adv. Math.* **219** (2008), 1070–1080. [2.1](#), [3](#)
- [8] A. Frieze and R. Kannan, Quick approximation to matrices and applications, *Combinatorica* **19** (1999), 175–220. [1](#)
- [9] O. Goldreich and D. Ron, A sublinear bipartiteness tester for bound degree graphs, *Combinatorica* **19** (1999), 335–373. [1](#), [4](#)
- [10] M.. Litvin, Counting graphs, counting edges and separation properties, M.Sc Thesis, Jerusalem 1996. [6](#)

- [11] G. Kalai, Separator theorems and small hereditary classes (Reserach Problem), *Discrete Mathematics* 257 (2002), 611–612. [6](#)
- [12] R. Kannan, S. Vempala and A. Vetta, On clusterings: Good, bad and spectral, *J. ACM* 51 (2004), 497–515. [4](#)
- [13] S. Khot, On the power of unique 2-prover 1-round games, *Proc. STOC 2002* 34 (2002), 767–775. [1](#)
- [14] A. Kostochka and L. Pyber, Small topological complete subgraphs of “dense” graphs, *Combinatorica* 8 (1988), 83–86. [1](#)
- [15] J. Lee, S. Oveis Gharan and L. Trevisan, Multiway spectral partitioning and higher-order Cheeger inequalities, *J. ACM* 61 (2014), 37:1–37:30. [2.4](#)
- [16] F. Leighton and S. Rao, Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms, *J. ACM* 46 (1999), 787–832. [1](#)
- [17] N. Linial, E. London and E. Rabinovich, The geometry of graphs and some of its algorithmic applications, *Combinatorica* 15 (1995), 215–245. [1](#), [3](#)
- [18] N. Linial and M. Saks, Low diameter graph decompositions, *Combinatorica* 13 (1993), 441–454. [1](#)
- [19] R. Lipton, D. Rose and R. Tarjan, Generalized nested dissections, *SIAM J. Numer. Anal.* 16 (1979), 346–358. [1](#), [2.1](#), [3](#)
- [20] R. Lipton and R. Tarjan, A separator theorem for planar graphs, *SIAM J. Appl. Math* 36 (1979), 177–189. [2.1](#)
- [21] L. Lovász, **Large networks and graph limits**, AMS, 2012. [2.5](#)
- [22] R. Levi, G. Moshkovitz, D. Ron, R. Rubinfeld and A. Shapira, Constructing near spanning trees with few local inspections, *Random Struct. Algor.*, to appear. [2.6](#), [8](#)
- [23] M. Pătraşcu and M. Thorup, Planning for fast connectivity updates, *Proc. FOCS 2007*, 263–271. [1](#), [6](#)
- [24] R. Rubinfeld, G. Tamir, S. Vardi and N. Xie, Fast local computation algorithms, *Proc. ICS 2010*, 223–238. [2.6](#)
- [25] A. L. Rosenberg and L. S. Heath, **Graph Separators, with Applications**, Springer, 2000. [2.1](#)
- [26] A. Shapira and B. Sudakov, Small complete minors above the extremal edge density, *Combinatorica* 35 (2015), 75–94. [1](#), [2.2](#), [4](#)
- [27] D. Steurer, On the complexity of unique games and graph expansion, Dissertation, 2010. [1.1](#), [6](#)

- [28] D. Steurer and P. Raghavendra, Graph expansion and the unique games conjecture, *Proc. STOC 2010*, 755–764. [1.1](#)
- [29] E. Szemerédi, Regular partitions of graphs, In: *Proc. Colloque Inter. CNRS* (J. C. Bermond, J. C. Fournier, M. Las Vergnas and D. Sotteau, eds.), 1978, 399–401. [1](#)
- [30] L. Trevisan, Approximation algorithms for unique games, *Proc. FOCS 2005*, 197–205. [1](#), [2.3](#), [5](#), [4](#)