Expanders – how to find them, and what to find in them

Michael Krivelevich[1]

Abstract

A graph $G = (V, E)$ is called an expander if every vertex subset U of size up to $|V|/2$ has an external neighborhood whose size is comparable to $|U|$. Expanders have been a subject of intensive research for more than three decades and have become one of the central notions of modern graph theory.

We first discuss the above definition of an expander and its alternatives. Then we present examples of families of expanding graphs and state basic properties of expanders. Next, we introduce a way to argue that a given graph contains a large expanding subgraph. Finally we research properties of expanding graphs, such as existence of small separators, of cycles (including cycle lengths), and embedding of large minors.

1 Introduction

Putting it somewhat informally, a graph G is an expander if every vertex subset U expands outside substantially, i.e., has an external neighborhood whose size is comparable to $|U|$. Since their introduction in the seventies, expanders have become one of the most central, and also one of the most applicable, notions of modern combinatorics. Their uses are numerous and wide-ranging, from extremal problems in graph theory and explicit constructions of graphs with desired properties to design and implementation of lean yet reliable communication networks. The reader is encouraged to consult the comprehensive and very readable survey of Hoory, Linial and Wigderson [25], devoted entirely to expanding graphs and covering many aspects of this subject; the very concise presentation of Sarnak [50] would be a (somewhat...) shorter introduction.

Ever since the seminal paper of Pinsker [47], randomness has been used to argue about expanders and their existence. In many cases, random graphs, sometimes after minor alterations, serve as typical examples of expanders with desired parameters. In fact, expanders are frequently viewed as synonymous to pseudo-random graphs, the latter being graphs whose edge distribution, and other properties, resemble closely those of truly random graphs with the same density. Pseudo-random graphs are covered in survey [38]. The notions of expanders and of pseudo-random graphs, while being close, are quite distinct — for example, a bipartite graph can be an excellent expander, but cannot quite model a random graph. Ideology and tools can be rather similar; in particular, spectral methods are used frequently to argue about both expansion and pseudo-randomness.

In this survey we undertake a (more) qualitative study of expansion and expanding graphs. Instead of looking for best possible quantitative results on expansion, strongest expanders possible, largest eigenvalue ratio etc., we aim here to study inherent properties of expanders, their derivation and consequences. We will rely less on the strength of expansion to derive our results; in a sense our assumption about expansion is fairly weak (we will sometimes even call our expanders weak explicitly). As it turns out, even the rather relaxed definition of expanders we embrace is deep enough to allow for meaningful study and quite interesting consequences.

The text is organized as follows. In the next section, we introduce a formal definition of expanders we use in this survey, and compare it with alternative but similar definitions. In Section 3 we discuss basic properties of expanding graphs. Examples of families of expanding graphs are presented in Section 4; we use the examples presented also to adjust our expectations, and to see which graph structures can and cannot be hoped to be found in weakly expanding graphs. The connection between expanders and separators is presented in Section 5. In Section 6 we introduce a general tool for proving the existence of a large expanding subgraph in a given graph, this tool is based on the notion of local sparseness; we discuss this notion in the context of sparse random graphs. Amongst the structures one can hope to find in expanders are long paths and cycles, and in Section 7 we show that indeed weak expanders are rich in cycles, in several well defined quantitative senses. Finally, in Section 8 we present results about embedding large minors in expanding graphs.

2 Definition(s) of an expander

Let us start with introducing (the rather standard) notation used throughout this survey. For a graph $G = (V, E)$ and a vertex set $U \subset V$ we denote by $N_G(U)$ the external neighborhood of U in G, i.e.,

$$N_G(U) = \{v \in V \setminus U : v \text{ has a neighbor in } U\}.$$

We also write, for two disjoint subsets $U, W \subset V$,

$$N_G(U, W) = \{v \in W : v \text{ has a neighbor in } U\},$$

thus $N_G(U) = N_G(U, V \setminus U)$. Also, $e_G(U, W)$ stands for the number of edges of G between U and W. Sometimes when the relevant graph G is clear from the context, we will omit the subscript in the above notation.

We can now give the most basic definition of this survey, that of an α-expander.

Definition 2.1. Let $G = (V, E)$ be a graph on n vertices, and let $\alpha > 0$. The graph G is an *α-expander* if $|N_G(U)| \geq \alpha|U|$ for every vertex subset $U \subset V$ with $|U| \leq n/2$.

It is important to note that although the above definition appears to cover only expansion of sets W up to half the order of G, it extends further to sets whose size exceeds $n/2$. Indeed, if $|U| > n/2$, then the set $W = V - (U \cup N_G(U))$ has size $|W| < n/2$, and satisfies $N_G(W) \subseteq N_G(U)$. This allows to bound $|N_G(U)|$ from below as a function of $|U|$, when assuming that G is an α-expander.

Our primary concern in this survey is neither getting strongest or best possible expanders, nor pushing the quantitative notion of expansion to the limit. Our goal is rather different — we mostly assume that we have a decent (one could even say weak) expander G, with α in the definition of an α-expander being a small constant, say, $\alpha = 0.01$, and aim to derive some nice properties of G. In principle, the expansion factor α in the above definition can even be allowed to be a vanishing function of the order n of G, however here we mostly stick to the assumption $\alpha = \Theta(1)$. We should add here that throughout this survey we allow ourselves to use expressions such as "weak expander", "strong expander", etc. rather informally, aiming to indicate the extent of expansion of the corresponding graph; of course formal definitions such as α-expansion should be used to measure the expansion formally and quantitatively.

The choice to stop the expansion assumption at $n/2$ in Definition 2.1 may appear somewhat arbitrary. As we will see shortly, this is certainly a convenient compromise, as it allows us to derive easily some basic properties of G, for example its connectedness. It is quite natural though to give also more general definitions, like the ones below.

Definition 2.2. Let $G = (V, E)$ be a graph, let I be a set of positive integers, and let $\alpha > 0$. The graph G is an (I, α)-*expander* if $|N_G(U)| \geq \alpha|U|$ for every vertex subset $U \subset V$ satisfying $|U| \in I$.

Definition 2.3. Let $G = (V, E)$ be a graph, let k be a positive integer, and let $\alpha > 0$. The graph G is a (k, α)-*expander* if $|N_G(U)| \geq \alpha|U|$ for every vertex subset $U \subset V$ satisfying $|U| \leq k$.

Notice, mostly to exercise the above defined notions, that an α-expander G on n vertices is (exactly) an (I, α)-expander for $I = \{1, \ldots, n/2\}$, and a (k, α)-expander for $k = n/2$.[2]

Are the above three notions of expansion radically different? Does an (I, α)-expander on n vertices with I being just a single value, say, $I = \{n/100\}$ have properties much different from an $(n/300, \alpha)$-expander, or from an α-expander? Well, yes in some senses (for example, the former two are easily seen not to guarantee connectedness, while the latter one does). However, since in this survey we will mostly research qualitative properties and will frequently settle for finding substructures in G of size linear in n, without trying too much to optimize the constants involved, in many other senses the above three definitions are morally equivalent. We will provide a formal statement supporting this paradigm shortly.

Definition 2.1 is a fairly commonly used notion of an expander, see, e.g., [2], or [6, Chapter 9]. For the case of d-regular graphs with constant d, or more generally, for graphs with bounded maximum degree, being an α-expander is essentially equivalent (qualitatively) to having edge expansion bounded away from 0. The latter notion can be quantified through the so-called Cheeger constant of a graph G. For a graph $G = (V, E)$ on n vertices and a subset $U \subseteq V$ set $\mathrm{vol}_G(U) = \sum_{v \in U} \deg_G(v)$. Define now

$$h(G) = \min_{\emptyset \neq U \subsetneq V} \frac{e_G(U, V \setminus U)}{\min(\mathrm{vol}_G(U), \mathrm{vol}_G(V \setminus U))},$$

the quantity $h(G)$ is called the *Cheeger constant* of G. (See, e.g., [13] for a general discussion.) Having $h(G)$ bounded away from 0 means the graph is a decent edge expander, and — assuming its maximum degree is bounded — is a decent vertex expander as well.

Let us now address the quantitative comparison between Definitions 2.1, 2.2 and 2.3. As announced, we will see that they are essentially equivalent.

Lemma 2.4. *Let G be an (I, α)-expander on n vertices with $I = \left[\frac{n}{3}, \frac{2n}{3}\right]$. Then there is a vertex subset $Z \subset V(G)$ of size $|Z| < n/3$ such that the graph $G' = G[V \setminus Z]$ is an α-expander.*

Proof. Start with $G' = G$, and as long as there is a vertex subset $A \subset V(G')$ of size $|A| \leq |V(G')|/2$ satisfying $|N_{G'}(A)| < \alpha|A|$, delete A and update $G' := G'[V(G') \setminus A]$. Observe that the disjoint union of the deleted sets Z obviously satisfies $|N_G(Z)| < \alpha|Z|$. Hence by our assumption on the expansion of G we derive that the size of Z is never in I during the execution of the procedure, so either $|Z| < n/3$ or $|Z| > 2n/3$. Suppose that after some iteration the size of Z exceeds $2n/3$ for the first time. Let A be the set deleted

---

[2] Here and later we allow ourselves to treat rounding issues somewhat casually.

at this iteration, and let $Z' = Z \setminus A$. Then $|Z'| < n/3$, and $|A| \leq (n - |Z'|)/2$, implying $|Z| = |Z'| + |A| \leq |Z'| + (n - |Z'|)/2 = n/2 + |Z'|/2 < n/2 + n/6 = 2n/3$ – a contradiction. This indicates that the deletion procedure halts with $|Z| < n/3$, and the subgraph $G[V \setminus Z]$ has the required expansion property. $\square$

Lemma 2.5. *Let G be a (k, α)-expander. Then G is a $\left(\frac{3k}{2}, \frac{\alpha}{6}\right)$-expander, or there is a subset $V_0 \subset V(G)$ of size $|V_0| \geq \frac{2k}{3}$ such that the induced subgraph $G[V_0]$ is an $(\alpha/2)$-expander.*

Proof. We can assume that G has more than $2k$ vertices, as otherwise the second alternative of the lemma is satisfied with $V_0 = V(G)$. If G does not satisfy the first alternative, then there exists a set X, $k < |X| \leq \frac{3k}{2}$ with $|N_G(X)| < \frac{\alpha|X|}{6}$. But then every subset $Y \subset X$ with $\frac{|X|}{3} \leq |Y| \leq \frac{2|X|}{3}$, satisfies:

$$|N_G(Y, X \setminus Y)| \geq |N_G(Y)| - |N_G(X)| \geq \alpha|Y| - \frac{\alpha|X|}{6} \geq \alpha|Y| - \frac{\alpha|Y|}{2} = \frac{\alpha|Y|}{2}.$$

Feeding the induced subgraph $G[X]$ to Lemma 2.4 produces the required $(\alpha/2)$-expander on at least $2|X|/3 \geq 2k/3$ vertices. $\square$

Lemma 2.6. *Let $G = (V, E)$ be a $(\{k\}, \alpha)$-expander with $0 < \alpha \leq 1$. Then there is a vertex subset $Z \subset V$ of size $|Z| < k$ such that the graph $G' = G[V \setminus Z]$ is an $\left(\frac{\alpha k}{3}, \frac{\alpha}{3}\right)$-expander.*

Proof. Start with $G' = G$, and as long as there is a vertex subset $A \subset V(G')$ of size $|A| \leq \frac{\alpha k}{3}$ with $|N_{G'}(A)| < \alpha|A|/3$, delete A and update $G' := G'[V(G') \setminus A]$. Observe that the union Z of deleted sets obviously satisfies $|N_G(Z)| < \alpha|Z|/3$. Assume $|Z|$ reaches k after some iteration. Let A be the set deleted at this iteration, and let $Z' = Z - A$. Then $|Z'| < k$, and $|A| \leq \alpha k/3$ with $|N_G(A, V - Z')| < \alpha|A|/3$. Choose an arbitrary subset $A' \subset A$ so that $|Z' \cup A'| = k$. Then

$$\alpha k \leq |N_G(Z' \cup A')| \leq |N_G(Z')| + |N_G(A', V \setminus Z')| \leq |N_G(Z')| + |N_G(A, V \setminus Z')| + |A \setminus A'|$$
$$< \frac{\alpha k}{3} + \frac{\alpha|A|}{3} + |A| < \frac{\alpha k}{3} + \frac{\alpha^2 k}{3} + \frac{\alpha k}{3} \leq \alpha k$$

– a contradiction. $\square$

We arrive at the following conclusion:

Theorem 2.7. *For every $0 < c, \alpha \leq 1$ there are $c', \alpha' > 0$ such that the following holds. Let G be a $(\{k\}, \alpha)$-expander on n vertices for $k = \lfloor cn \rfloor$. Then G contains an induced subgraph G' on at least $c'n$ vertices which is an α'-expander.*

The above theorem indicates clearly that the alternative definitions of expanders, as given by Definitions 2.1–2.3, are basically equivalent on the qualitative level: assuming a one-point expansion of Definition 2.2 with $I = \{k\}$ for $k = \Theta(n)$ delivers an α-expander of Definition 2.1 on linearly many vertices. Hence in the rest of this survey we will mostly work with Definition 2.1.

Observing Theorem 2.7, one may wonder whether assuming expansion of sublinear sets in a graph G, or even more ambitiously of sets of size up to k, where k is much smaller than $n = |V(G)|$, can guarantee the existence of a substantially sized expanding subgraph G' of G. For one, a (k, α)-expander on n vertices can be just a collection of disjoint cliques of size about $(1 + \alpha)k$ each, and so if $k = o(n)$, such G does not contain a linearly sized

expanding subgraph. Thus, the assumption $k = \Theta(n)$ in Theorem 2.7 is essential. A much more striking example is provided by Moshkovitz and Shapira [45]. They prove (see Theorem 1 of [45]) that for every n and $k = O(\log \log n)$, there is an n-vertex graph G with all degrees of order k and $\text{girth}(G) = \Omega(\log n/k^2)$, in which every induced subgraph H on $t \geq \text{girth}(G)$ vertices has a set of edges $E_0 \subseteq E(H)$ of size $|E_0| = o(t)$, whose removal from H creates a graph H' with all connected components having at most $2t/3$ vertices. (Rephrasing it, every induced subgraph H of G with less than $\text{girth}(G)$ vertices is a tree, while every induced H with at least $\text{girth}(G)$ vertices has a small edge-separator and thus cannot be an expander.) Such a graph G can easily be argued to be a good local expander; for example, every set $U \subset V(G)$ of size $|U| \leq (ck)^{(\text{girth}(G)-1)/2}$ satisfies $|N_G(U)| \geq 2|U|$, for an appropriately chosen constant $c > 0$, see Lemma 2.1 of [52]. Hence local expansion in a graph cannot always be traded for global expansion in a subgraph.

3 Basic properties of α-expanders

This short section discusses the most basic properties of α-expanders. It can serve as a vindication of our choice to stick to Definition 2.1 as the definition of choice for expanders.

For a graph $G = (V, E)$, a vertex $v \in V$ and a natural number t, denote by $B_t(v)$ the ball of radius t around v in G, i.e., $B_t(v)$ is the set of vertices of G at distance at most t from v.

Proposition 3.1. *Let $G = (V, E)$ be an α-expander on n vertices. Then for every $v \in V$ and every natural t, we have: $|B_t(v)| \geq \min\left\{\frac{n}{2}, (1+\alpha)^t\right\}$.*

Proof. By induction on t. Obvious for $t = 0$. For the induction step, observe that $B_t(v) \supseteq B_{t-1}(v)$. If $|B_{t-1}(v)| \geq \frac{n}{2}$, we are done. Otherwise, by the induction hypothesis we have $|B_{t-1}| \geq (1+\alpha)^{t-1}$. Denoting $U = B_{t-1}(v)$, we have: $B_t(v) = U \cup N(U)$, and $|N(U)| \geq \alpha|U|$. It follows that $|B_t(v)| \geq (1+\alpha)|U| \geq (1+\alpha)^t$, as required. $\square$

Corollary 3.2. *Let G be an α-expander on n vertices. Then the diameter $\text{diam}(G)$ satisfies:* $\text{diam}(G) \leq \left\lceil \frac{2(\log n - 1)}{\log(1+\alpha)} \right\rceil + 1 = O_\alpha(\log n)$.

Proof. Let $u \neq v \in V(G)$. Growing balls around u, v, we obtain by Proposition 3.1 that for $t = \left\lceil \frac{\log(n/2)}{\log(1+\alpha)} \right\rceil$ the balls $B_t(u), B_t(v)$ have at least $n/2$ vertices each. Then they intersect each other, or are connected by an edge. $\square$

In particular, an α-expander G is connected.

Turning now to the Cheeger constant, observe that an α-expander G of maximum degree $\Delta = \Delta(G)$ satisfies trivially: $h(G) \geq \frac{\alpha}{\Delta}$. Thus, assuming that both $\alpha, \Delta = \Theta(1)$, we derive that in this case the Cheeger constant of G is bounded away from zero.

Going from the Cheeger constant to random walks, we derive, using the classical result of Jerrum and Sinclair [27], that the mixing time T_{mix} of a (lazy) random walk on an α-expander G on n vertices with $\alpha = \Theta(1)$ and bounded maximum degree satisfies $T_{mix} = O(\log n)$. The assumption that the maximum degree of G is bounded is essential here, as for example taking G to be two disjoint cliques of size $n/2$ connected by a nice bounded degree expander, we observe that G is an α-expander with $\alpha = \Theta(1)$, and yet it takes a random walk a linear time in n on average to cross from one clique to the other. Since random walks are essentially outside the scope of this survey, we will not dwell on their aspects anymore, instead referring the interested reader to standard sources on random walks (say, [41]) for background, definitions and discussion.

We can already conclude that the definition of α-expanders we have adopted is capable of guaranteeing easily several basic and desirable properties of the graph. We aim however to set our goals much higher than that, and to derive further properties of α-expanders, such as the existence of linearly long paths and cycles, non-existence of sublinear separators, and embedding of large minors. These properties are to be discussed later in this survey.

4 Examples of α-expanders

In this section we present some archetypal (families of) examples of weak expanders. The aim here is not just to showcase concrete instances of α-expanders, but also to adjust our level of expectations of what can be found in α-expanders upon observing some concrete examples.

1. Bipartite graphs. The complete bipartite graph G with parts of size αn and $(1 - \alpha)n$ for $\alpha \leq 1/2$ is obviously an α-expander. This example can be sparsified by taking G to be a good expander with linearly many edges between parts of size $2\alpha n$ and $(1 - 2\alpha)n$ (for $\alpha \leq 1/4$).

This example shows us that we cannot always expect to find an odd cycle in an expander.

2. Spectral approach. This is a very powerful and appealing approach, allowing to connect between expansion properties of a graph and its eigenvalues. (See, e.g., [13, 12] for a general background on spectral graph theory.) For a graph G with vertex set $V(G) = [n]$, the adjacency matrix $A(G)$ is an n-by-n matrix whose entry a_{ij} is 1 if $(i, j) \in E(G)$ and $a_{ij} = 0$ otherwise. This is a symmetric matrix with n real eigenvalues, usually sorted in the non-increasing order $\lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$; they are usually called the eigenvalues of the graph G itself. For the simplest case of d-regular graphs, the first (or trivial) eigenvalue is $\lambda_1 = d$, and it is the second eigenvalue that governs the graph expansion. Specifically, the following is true (see, e.g., Corollary 9.2.2 of [6]):

Proposition 4.1. *Let G be a d-regular graph on n vertices with the second largest eigenvalue λ. Then G is an α-expander, with $\alpha = \frac{d-\lambda}{2d}$.*

Hence a d-regular graph with eigenvalue gap $d - \lambda_2 = \Theta(1)$ and $d = O(1)$ is a weak expander. This implication is reversible – a regular expander has its two first eigenvalues well separated, see [2].

The spectral approach extends to the general (non-regular) case. (See [13] for an extensive discussion of the subject and all missing definitions.) Let G be a graph on n vertices, and let $0 = \mu_0 \leq \mu_1 \leq \ldots \leq \mu_{n-1}$ be the eigenvalues of the normalized Laplacian $\mathcal{L}(G)$ of G, defined by $\mathcal{L} = I_n - D^{-1/2}AD^{1/2}$, where $D = D(G) = \mathrm{diag}(\deg(v_1), \ldots, \deg(v_n))$ is the degree matrix of G (we use the convention $D_{ii}^{-1} = 0$ when $\deg(v_i) = 0$). We set $\mu(G) = \mu_1$. Let $h(G)$ be the Cheeger constant of G. The famous *Cheeger inequality for graphs* [15, 4, 2] states that

$$\frac{h^2(G)}{2} \leq \mu(G) \leq 2h(G). \tag{4.1}$$

In particular, assuming that G has its degrees bounded, and recalling the prior discussion about expansion and the Cheeger constant, we conclude that having the first non-trivial eigenvalue of the normalized Laplacian $\mu(G)$ bounded away from 0 guarantees expansion. (Here we use the second part of the Cheeger inequality (4.1), the first (and more involved) part will be utilized later, when discussing algorithmic issues.)

3. Stretching edges of a bounded degree expander. This is a systematic way to obtain a (weaker) expander from a given expander G by stretching, or subdividing, the edges of G.

Proposition 4.2. *Let G be an α-expander of maximum degree $\Delta = O(1)$, and let ℓ be a positive integer. Subdividing each edge of G ℓ times produces an $\Omega_\Delta(\alpha/\ell)$-expander G'.*

While being a nice way of producing new expanders, the above proposition should turn all warning lights on. Indeed, even starting with a very strong bounded degree expander G, the new graph G' has no adjacent vertices of degree more than 2. Hence, one *cannot* hope to embed in G' any graph H in which some two vertices of degree at least three are adjacent. Even if H is a little tree with two vertices of degree three adjacent – H is not there in G'. This seems to indicate very gloomy prospects for this line of business, and for this survey in particular – there are seemingly not many structures one can seek to embed in an α-expander... Not all is lost however – weak expanders do have a rich enough structure, and we can argue about the existence of long paths and cycles, and also about embedding minors. This is indeed what we plan to do in the sequel.

4. Random regular graphs. Let $G_{n,d}$ be the probability space of d-regular graphs on n vertices. (See, e.g., [53], or [22, Chapter 10] for background.) For every fixed $d \geq 3$ and growing n a graph G drawn from $G_{n,d}$ is typically an expander. This can be shown directly using the so-called configuration model, as was done by Bollobás in [10]; another possible approach is to invoke Proposition 4.1 and known results about eigenvalues of random regular graphs [20].

5 Expanders and separators

Separators are one of the central concepts in graph theory in general, and in structural graph theory in particular.

Definition 5.1. Let $G = (V, E)$ be a graph on n vertices. A vertex set $S \subset V$ is a *separator* in G if there is a partition $V = A \cup B \cup S$ of the vertex set of G such that $|A|, |B| \leq 2n/3$, and G has no edges between A and B.

Separators serve to measure quantitatively the connectivity of large vertex sets in graphs; the fact that all separators in G are large indicates that it is costly to break G into large pieces not connected by any edge. Observe that if G is a bounded degree graph, then disconnecting G can be done cheaply by removing $O(1)$ neighbors of any vertex $v \in V(G)$; however, if G is well connected, then finding a small sized separator might be impossible.

Separators came into prominence with the celebrated result of Lipton and Tarjan [42], asserting that every planar graph on n vertices has a separator of size $O(\sqrt{n})$. This line of research, advancing the alternative "a small separator or a large minor" to be addressed later in this survey, has been quite fruitful over the years, see, e.g., [5, 48, 29]. We will return to it later in this survey.

It is easy to argue that expanders do not have small separators.

Proposition 5.2. *Let G be an α-expander on n vertices, and let S be a separator in G. Then $|S| \geq \frac{\alpha n}{3(1+\alpha)}$.*

Proof. Let S be a separator in G of size $|S| = s$, separating A and B, with $|A| = a, |B| = b$; we assume $a \leq b \leq 2n/3$. Then $a + s \geq n/3$. Clearly, $N_G(A) \subseteq S$. Since $a \leq n/2$, by

the definition of an α-expander we get $s - \alpha a \geq 0$. Multiplying this inequality by $1/\alpha$ and summing with $a + s \geq n/3$, we obtain: $s(1 + 1/\alpha) \geq n/3$, or $s \geq \frac{\alpha n}{3(1+\alpha)}$. $\qquad\square$

Perhaps more surprisingly, it turns out that the opposite implication is true in a well defined quantitative sense – graphs without small separators contain large induced expanders.

Proposition 5.3. *Let $\beta > 0$ be a constant. If all separators in a graph G on n vertices are of size at least βn, then G contains an induced $(3\beta/2)$-expander on at least $2n/3$ vertices.*

The proof is somewhat similar to Lemmas 2.4 and 2.6.

Proof. Start with $G' = G$, and delete repeatedly subsets A of size at most $|V(G')|/2$ not expanding themselves into $V(G')$ by at least the factor of $3\beta/2$, each time updating $G' := G' - A$. Let Z be the union of the deleted sets, clearly $|N_G(Z)| < 3\beta|Z|/2$. If the size of Z ever reaches $n/3$, focus on the moment it happens for the first time, let A be the last set deleted, and let $Z' = Z \setminus A$. Then $|A| \leq (n - |Z'|)/2$ and $|Z'| \leq n/3$. Combining these two inequalities, we get: $|Z| = |Z'| + |A| \leq 2n/3$. The set $N_G(Z)$ forms a separator in G (separating between Z and $V(G) \setminus (Z \cup N_G(Z))$), hence by the proposition's assumption its size is at least $\beta n \geq 3\beta|Z|/2$ – a contradiction. We conclude that the deletion process stops with $|Z| < n/3$, and the final graph of the process is a $(3\beta/2)$-expander on at least $2n/3$ vertices, as required. $\qquad\square$

The above arguments can easily be extended to sublinear separators and subconstant expansion.

We can learn a very important qualitative lesson here. It turns out that weak expanders and graphs without small separators are essentially the same thing (at least if we do not care that much about multiplicative constants involved). Thus, when aiming to prove results about graphs without sublinear separators, we can choose instead to operate in the realm of weak expanders. This simple yet powerful connection between two central graph theoretic notions (expanders and separators), usually perceived as belonging to quite different worlds (extremal graph theory and structural graph theory, respectively) can be quite fruitful and illuminating.

6 Finding large expanding subgraphs

Given the prominence of expanders and their usability, it is tempting to argue that every graph is an expander, or at least contains a large expander inside. This is obviously too much wishful thinking. Not every graph is an expander, moreover essentially every standard notion of expansion is rather fragile – adding an isolated vertex to a strong expander G produces a non-expanding graph G'. Not every graph, even of average degree bounded away from zero, contains a substantially sized expanding subgraph. Indeed, planar graphs, or more generally graphs of bounded genus have sublinear separators [42, 23]. Hence, applying Proposition 5.2, we conclude that such graphs do not contain expanding subgraphs of super-constant size. Recall also the example of Moshkovitz and Shapira [45], discussed in Section 2.

There is a partial remedy to the problem with this general approach. Shapira and Sudakov [51] and later Montgomery [44] argued that *every* graph G contains a very weak expander G^* of nearly the same average degree. Their notion of expansion is different – the expansion required is gradual in terms of subset sizes; for an m-vertex graph to be a weak expander, vertex sets of size m^c, $0 < c < 1$, should expand by a constant factor,

whereas linearly sized vertex sets are required to expand only by about $1/\log m$ factor. (We are rather informal here in our descriptions, see the actual papers [51, 44] for accurate definitions.) Neither of these results guarantees a (weakly) expanding subgraph on linearly many vertices. Much earlier, Komlós and Szemerédi, in their work on topological cliques in graphs [31, 32], presented a fairly general scheme for arguing about the existence of weak expanders in any given graph; their scheme does not provide – naturally – for finding linearly sized expanders, or expanders with constant expansion of subsets.

How natural or common is this condition? As the (very easy) proposition below shows, we now present a fairly simple sufficient condition guaranteeing the existence of a large expanding subgraph in a given graph. This condition is based on the quantitative notion of local sparseness.

Definition 6.1. Let $c_1 > c_2 > 1$, $0 < \beta < 1$. A graph $G = (V, E)$ on n vertices is called a (c_1, c_2, β)-*graph* if

1. $\frac{|E|}{|V|} \geq c_1$;

2. every vertex subset $W \subset V$ of size $|W| \leq \beta n$ spans fewer than $c_2|W|$ edges.

In words, the above condition says that relatively small sets are sizably sparser than the whole graph. It has been used in recent papers [35, 34] of the author.

How natural or common is this condition? As the (very easy) proposition below shows, most sparse graphs are locally sparse.

Proposition 6.2. *Let $c_1 > c_2 > 1$ be reals. Define $\beta = \left(\frac{c_2}{5c_1}\right)^{\frac{c_2}{c_2-1}}$. Let G be a random graph drawn from the probability distribution $G\left(n, \frac{c_1}{n}\right)$. Then with high probability every set of $k \leq \beta n$ vertices of G spans fewer than $c_2 k$ edges.*

Proof. The probability in $G(n, p = c_1/n)$ that there exists a vertex subset violating the required property is at most

$$\sum_{i \leq \beta n} \binom{n}{i}\binom{\binom{i}{2}}{c_2 i} \cdot p^{c_2 i} \leq \sum_{i \leq \beta n} \left(\frac{en}{i}\right)^i \cdot \left(\frac{eip}{2c_2}\right)^{c_2 i} = \sum_{i \leq \beta n} \left[\frac{en}{i} \cdot \left(\frac{ec_1 i}{2c_2 n}\right)^{c_2}\right]^i$$

$$= \sum_{i \leq \beta n} \left[\frac{e^{c_2+1}c_1^{c_2}}{(2c_2)^{c_2}} \cdot \left(\frac{i}{n}\right)^{c_2-1}\right]^i .$$

Denote the i-th summand of the last sum by a_i. Then, if $i \leq n^{1/2}$ we get: $a_i \leq \left(O(1)n^{-\frac{c_2-1}{2}}\right)^i$, implying $\sum_{1 \leq i \leq n^{1/2}} a_i = o(1)$. For $n^{1/2} \leq i \leq \beta n$, we have, recalling the expression for β:

$$a_i \leq \left[\frac{e^{c_2+1}c_1^{c_2}}{(2c_2)^{c_2}} \cdot \left(\frac{c_2}{5c_1}\right)^{c_2}\right]^i = \left(e \cdot \left(\frac{e}{10}\right)^{c_2}\right)^i = o(n^{-1}) .$$

It follows that $\sum_{i \leq \beta n} a_i = o(1)$, and the desired property of the random graphs holds with high probability. $\square$

One can also cap easily the typical maximum degree in (a nearly spanning subgraph of) a random graph.

Proposition 6.3. *For every $C > 0$ and all sufficiently small $\delta > 0$ the following holds. Let G be a random graph drawn from the probability distribution $G\left(n, \frac{C}{n}\right)$. Then with high probability every set of $\frac{\delta}{\ln \frac{1}{\delta}} n$ vertices of G touches fewer than δn edges.*

We omit the straightforward proof.

Observe that if $G = (V, E)$ satisfies the conclusion of the above proposition, then by deleting $\frac{\delta}{\ln \frac{1}{\delta}} n$ vertices of highest degrees in G, one obtains a spanning subgraph $G' = (V', E')$ on $|V'| = \left(1 - \frac{\delta}{\ln \frac{1}{\delta}}\right) n$ vertices and with $|E'| \geq |E| - \delta n$ edges, in which all degrees are at most $2\ln(1/\delta)$. (Otherwise, all deleted vertices are of degree at least $2\ln(1/\delta)$, forming a subset touching at least δn edges – a contradiction.)

We now formulate the main result of this section.

Theorem 6.4 ([34]). *Let $c_1 > c_2 > 1$, $0 < \beta < 1$, $\Delta > 0$. Let $G = (V, E)$ be a graph on n vertices, satisfying:*

1. $\frac{|E|}{|V|} \geq c_1$;

2. *every vertex subset $U \subset V$ of size $|U| \leq \beta n$ spans fewer than $c_2 |U|$ edges;*

3. $\Delta(G) \leq \Delta$.

Then G contains an induced subgraph $G^* = (V^*, E^*)$ on at least βn vertices which is an α-expander, for $\alpha = \frac{c_1 - c_2}{\Delta \cdot \left\lceil \log \frac{1}{\beta} \right\rceil}$.

Putting it informally, every locally sparse graph G of bounded maximum degree contains a linearly sized expander G^*. In our terminology, the first two conditions above say precisely that G is a (c_1, c_2, β)-graph. They are spelled out in full in the statement above so as to make it self-contained.

Proof of Theorem 6.4 (sketch). The proof proceeds in rounds. We describe the first round here, the rest is fairly similar.

Choose a constant $0 < \delta \ll c_1 - c_2$. Let $H = (W, F)$ be a minimal by inclusion non-empty induced subgraph of G, satisfying $|F|/|W| \geq c_1$. (Such a subgraph exists, as G itself meets this condition.) Due to the local sparseness assumption, we have $|W| > \beta n$. Also, every subset $U \subseteq W$ touches at least $c_1 |U|$ edges of H. Otherwise, deleting U is easily seen to produce a smaller induced subgraph H', still meeting the requirement stated in the definition of H – a contradiction. (We could add that H is connected due to its minimality, but we do not need this property here.)

Let us analyze now the expansion properties of H. If a subset $U \subset W$ has at most βn vertices, then it spans at most $c_2 |U|$ edges in G, and thus in H, and yet touches at least $c_1 |U|$ edges. This shows that at least $(c_1 - c_2)|U|$ edges of H leave U. Recalling the maximum degree assumption, we conclude that U has at least $(c_1 - c_2)|U|/\Delta$ neighbors outside of it in H.

At this point of the proof we have assured that H is a $(\beta n, \frac{c_1 - c_2}{\Delta})$-expander. We can complete the proof by appealing to the general (and somewhat inexplicit) statement of Theorem 2.7. We prefer however, just like in the original paper [34], to provide a self-contained argument, delivering also a better (and explicit) estimate for the order of the resulting subgraph and for its expansion strength.

Consider now medium-sized subsets U in H. Let $\beta n < |U| \leq |W|/2$. Such U touches at least $c_1|U|$ edges of U. If every such U spans at most $(c_1 - \delta)|U|$ edges, we get it expanding to $\delta|U|/\Delta$ vertices outside, thus indicating that H is a desired expander. Otherwise there is a subset U, $\beta n < |U| \leq |W|/2$, spanning at least $(c_1 - \delta)|U|$ edges. We can now switch our attention to the induced graph $H[U] = G[U]$, whose order is at most half that of the original graph, yet we did not give much in terms of its density, compared to the assumption on G – it is now at least $c_1 - \delta$.

Iterating this argument at most $O(\log(1/\beta))$ times, we eventually arrive at the desired linearly sized α-expander inside G. □

The assumption about bounded maximum degree $\Delta(G) = O(1)$ is used in the proof only to trade edge expansion for vertex expansion.

The proof of Theorem 6.4 is inherently non-algorithmic, as it uses the existence of a minimal by inclusion subgraph H of prescribed density, and such a subgraph is hard to find efficiently. It turns out that with some loss in constants, one can provide an algorithmic proof of the theorem. The argument for such a proof is based on the fact that the proof of (the first part of) the Cheeger inequality (4.1) can be made constructive in the sense that, given a graph $G = (V, E)$ on n vertices, one can find in time polynomial in n a subset $U \subset V$, $\mathrm{vol}_G(U) \leq \mathrm{vol}_G(V)/2$, satisfying $e_G(U, V \setminus U) \leq \sqrt{2\mu(G) \cdot \mathrm{vol}_G(U)}$. (See [2], or [13, Ch. 2], or [14], or [25, Sect. 4.5].) This yields the following algorithmic result.

Theorem 6.5 ([34]). *Let $c_1 > c_2 > 1$, $0 < \beta < 1$, $\Delta > 0$. There exist a constant $\alpha = \alpha(c_1, c_2, \beta, \Delta) > 0$ and an algorithm, that, given an n-vertex graph $G = (V, E)$ with $\Delta(G) \leq \Delta$ and $|E|/|V| \geq c_1$, finds in time polynomial in n a subset $U \subset V$ of size $|U| \leq \beta n$, spanning at least $c_2|U|$ edges in G, or an induced α-expander $G^* = (V^*, E^*) \subseteq G$ on at least βn vertices.*

See Section 2 of [34] for the proof.

Theorem 6.4 can be used to argue that a supercritical random graph $G(n, c/n)$, $c > 1$, contains **whp** an induced expander of linear size. This can be quite handy as it allows to extend immediately known properties of expanders to sparse random graphs.

Corollary 6.6. *For every $\epsilon > 0$ there exist $\alpha, \beta > 0$ such that a random graph $G \sim G\left(n, \frac{1+\epsilon}{n}\right)$ contains with high probability an induced bounded degree α-expander on at least βn vertices.*

Proof. Due to the standard monotonicity arguments we can assume that ϵ is small enough where necessary.

We will utilize several (very standard) facts about supercritical random graphs. It is known (see, e.g., [26, Ch. 5]) that **whp** $G \sim G\left(n, \frac{1+\epsilon}{n}\right)$ contains a connected component $C_1 = (V_1, E_1)$ (the so called giant component) satisfying:

$$
\begin{aligned}
|V_1| &= 2\epsilon(1 + o_\epsilon(1))n \,, \\
\frac{|E_1|}{|V_1|} &= 1 + (1 + o_\epsilon(1))\frac{\epsilon^2}{3} \,.
\end{aligned}
$$

Also, by Proposition 6.3 **whp** every $\frac{\epsilon^3}{10\ln\frac{1}{\epsilon}}n$ vertices of C_1 touch at most $\frac{\epsilon^3}{3}n$ edges. Deleting $\frac{\epsilon^3}{10\ln\frac{1}{\epsilon}}n$ vertices of highest degrees from C_1, one gets a graph $G_0 = (V_0, E_0)$ of

maximum degree $\Delta(G_0) \leq 7 \ln \frac{1}{\epsilon}$. In addition,

$$
\begin{aligned}
|V_0| &\geq \left(2\epsilon(1 + o_\epsilon(1)) - \frac{\epsilon^3}{10 \ln \frac{1}{\epsilon}} \right) n = 2\epsilon(1 + o_\epsilon(1))n \,, \\
|E_0| &\geq |E_1| - \frac{\epsilon^3}{3} n \geq |V_1| \left(1 + \frac{\epsilon^2}{3} + o(\epsilon^2) \right) - \frac{\epsilon^3}{3} n \\
&\geq |V_0| \left(1 + \frac{\epsilon^2}{3} + o(\epsilon^2) \right) - \frac{\epsilon^3}{3} n \geq |V_0| \left(1 + \frac{\epsilon^2}{7} \right).
\end{aligned}
$$

Finally, applying Proposition 6.2 with $c_1 = 1 + \epsilon$, $c_2 = 1 + \frac{\epsilon^2}{10}$, we get that **whp** every $k \leq \beta n$ vertices of G_0 (with $\beta = \beta(\epsilon)$ from Proposition 6.2) span fewer than $(1 + \epsilon^2/10)k$ edges. The conditions are set to call Theorem 6.4 and to apply it to G_0; we conclude that, given the above likely events, G_0 contains a linearly sized α-expander. $\qquad\square$

We remark here that in order to get this qualitative result, one does not really need to apply the heavy machinery of random graphs — it is enough actually to argue from the "first principles". Indeed, given the likely existence of the giant component C_1 in the supercritical regime, one can prove easily (for example, through sprinkling) that its density is typically above 1.

Let us give yet another illustrative example of how Theorem 6.4 can be utilized. This example comes from the realm of positional games. (The reader can consult monograph [24] for a general background on the subject.) In a *Maker–Breaker* game two players, called Maker and Breaker, claim alternately free edges of the complete graph K_n, with Maker moving first. Maker claims one edge at a time, while Breaker claims $b \geq 1$ edges (or all remaining fewer than b edges if this is the last round of the game). The integer parameter b is the so-called *game bias*. Maker wins the game if the graph of his edges in the end possesses a given monotone graph theoretic property, Breaker wins otherwise, with draw being impossible. In the *long cycle game* Maker's goal is to create as long as possible cycle. For $b \geq n/2$, Bednarska and Pikhurko proved [9] that Breaker has a strategy to force Maker's graph being acyclic. On the other hand, if $b = b(n)$ is such that Maker ends up with at least n edges, his graph contains a cycle (of some length) in the end. We can argue, using expanders, that already for $b = (1 - \epsilon)n/2$, Maker can force a linearly long cycle in his graph.

Theorem 6.7 ([35]). *For every $\epsilon > 0$ there exists $c > 0$ such that for all sufficiently large n, when playing the b-biased Maker–Breaker game on $E(K_n)$ with $b \leq (1 - \epsilon)\frac{n}{2}$, Maker has a strategy to create a cycle of length at least cn.*

Proof (sketch). Maker plays randomly (i.e., by taking a random free edge) during the first $(1 + \epsilon/2)n$ rounds of the game. Observe that at any moment during these rounds, the board still has $\Theta(n^2)$ free edges, and thus the probability of any edge to be chosen by Maker at a given round is $O(n^{-2})$. This allows us to apply standard union-bound type calculations to claim that with positive probability Maker, playing against any strategy of Breaker, can create a graph M_0 on n vertices with the following properties:

- M_0 has $\left(1 + \frac{\epsilon}{2} \right) n$ edges;

- every $k \leq \delta n$ vertices span at most $\left(1 + \frac{\epsilon}{8} \right) k$ edges;

- every δn vertices touch at most $\frac{\epsilon n}{4}$ edges,

for some $\delta = \delta(\epsilon) > 0$. Since the game analyzed is a perfect information game with no chance moves, it follows that in fact Maker has a (deterministic) strategy to create a graph M_0 with the above stated properties. Take such M_0 and delete δn vertices of highest degrees. The obtained graph M_1 has $(1 - \delta)n$ vertices, at least $\left(1 + \frac{\epsilon}{4}\right)n$ edges, maximum degree $\Delta(M_1) \leq \frac{\epsilon}{2\delta}$, and every $k \leq \delta n$ vertices span at most $\left(1 + \frac{\epsilon}{8}\right)k$ edges. Applying Theorem 6.4 shows that such M_1, being a part of Maker's graph by the end of the game, contains a linearly sized α-expander, for some $\alpha = \alpha(\epsilon) > 0$. As we will argue in the next section (see Theorem 7.4), Maker's graph contains then a linearly long cycle. $\qquad\square$

7   Long paths and cycles

Now we start investigating extremal properties of expanders directly. The first objects to explore are paths and cycles. As we explained in Section 4, it is unrealistic to expect appearance of many substructures in every expanding graphs. Luckily paths and cycles (the latter ones possibly with some restrictions on their lengths) are not in this excluded category, and indeed, as we will argue in this section, every α-expander on n vertices contains cycles whose length is linear in n.

We first describe the most basic – and very handy – tool for arguing about long paths and cycles in expanding graphs, which is the Depth First Search algorithm (DFS).

7.1   DFS algorithm

The *Depth First Search* is a well known graph exploration algorithm, usually applied to discover connected components of an input graph. As it turns out, due to its nature (of always pushing deeper, as its name suggests), this algorithm is particularly suitable for finding long paths and cycles in graphs. Some illustrative examples of its applications in the context of extremal graphs and random graphs include [1, 3, 11, 8, 40, 49, 36].

Recall that the DFS (Depth First Search) is a graph search algorithm that visits all vertices of a graph and eventually discovers the structure of the connected components of G. The algorithm receives as an input a graph $G = (V, E)$; it is also assumed that an order σ on the vertices of G is given, and the algorithm prioritizes vertices according to σ. The algorithm's output is a spanning forest F of G, whose connected components are identical to those of G. Each connected component C of F is a tree rooted at vertex r, the first vertex of C according to σ.

The algorithm maintains three sets of vertices, letting S be the set of vertices whose exploration is complete, T be the set of unvisited vertices, and $U = V \setminus (S \cup T)$, where the vertices of U are kept in a stack (the last in, first out data structure). It initializes with $S = U = \emptyset$ and $T = V$, and runs till $U \cup T = \emptyset$. At each round of the algorithm, if the set U is non-empty, the algorithm queries T for neighbors of the last vertex v that has been added to U, scanning T according to σ. If v has a neighbor u in T, the algorithm deletes u from T and inserts it into U. If v does not have a neighbor in T, then v is popped out of U and is moved to S. If U is empty, the algorithm chooses the first vertex of T according to σ, deletes it from T and pushes it into U.

The following are basic properties of the DFS algorithm:

(P1) at each round of the algorithm one vertex moves, either from T to U, or from U to S;

(P2) at any stage of the algorithm, it has been revealed already that the graph G has no edges between the current set S and the current set T;

(P3) the set U always spans a path;

(P4) let F be a forest produced by the DFS algorithm applied to G. If $(u, v) \in E(G) \setminus E(F)$, then one of $\{u, v\}$ is a predecessor of the other in F (along the unique path to the root).

Properties **(P1)**, **(P2)** are immediate upon a brief reflection on the description of the DFS algorithm. For Property **(P3)**, observe that when a vertex u is added to U, it happens because u is a neighbor of the last vertex v in U; thus u augments the path spanned by U, of which v is the last vertex; moving the last vertex v of U over to S deletes the last vertex of the path spanned by U, and hence still leaves the updated U spanning a path. Property **(P4)**, applied for example in [3], is perhaps less obvious, and therefore we supply its short proof here. Since u, v are connected by an edge, they obviously belong to the same connected component C of G (which is also a connected component of F, vertex-wise). Assume that v was discovered by the algorithm (and moved to U) earlier than u. If u is not under v in the corresponding rooted tree component T of F, then the algorithm has completed exploring v and moved it over to S before getting to explore u (and to move it to U) — obviously a contradiction.

7.2 Long paths

We now exploit the features of the DFS algorithm to derive the existence of long paths in expanding graphs. The following simple to prove statement is stated explicitly in [33] (see Proposition 2.1 there).

Proposition 7.1. *Let k, ℓ be positive integers. Assume that $G = (V, E)$ is a graph on more than k vertices, in which every vertex subset S of size $|S| = k$ satisfies: $|N_G(S)| \geq \ell$. Then G contains a path of length ℓ.*

Proof. Run the DFS algorithm on G, with σ being an arbitrary ordering of V. Look at the moment during the algorithm execution when the size of the set S of already processed vertices becomes exactly equal to k (there is such a moment due to Property **(P1)** above, as the vertices of G move into S one by one, till eventually all of them land there). By Property **(P2)**, the current set S has no neighbors in the current set T, and thus $N_G(S) \subseteq U$, implying $|U| \geq \ell$. The last move of the algorithm was to shift a vertex from U to S, so before this move U was one vertex larger. The set U always spans a path in G, by Property **(P3)**. Hence G contains a path of length ℓ. $\qquad\square$

It thus follows that if G is α-expander on n vertices, then G contains a path of length at least $\alpha n/2$ (apply Proposition 7.1 with $k = n/2$, $\ell = \alpha n/2$). As we have seen, with the right technology (DFS in this case) the proof of this result becomes essentially a one-liner. The estimate delivered by this argument is fairly tight, also in terms of the dependence on α – if G is taken to be the complete bipartite graph with sides of sizes αn and $(1 - \alpha)n$, then a longest path in G has $2\alpha n$ vertices.

We now describe an application of Proposition 7.1 to size Ramsey numbers. (See, e.g., [35] or [17] for the relevant background.) For graphs G, H and a positive integer r, we write $G \to (H)_r$ if every r-coloring of the edges of G produces a monochromatic copy of H. The *r-color size Ramsey number* $\hat{R}(H, r)$ is defined as the minimal possible number of edges in a graph G, for which $G \to (H)_r$. Let us discuss briefly the multi-color size Ramsey number $\hat{R}(P_n, r)$ of the path P_n on n vertices. Beck in his groundbreaking paper [7] proved

that $\hat{R}(P_n, 2) = O(n)$; his argument can easily be extended to show that $\hat{R}(P_n, r) = O_r(n)$ for any fixed $r \geq 2$. If so, the relevant question now is what is the hidden dependence on r in the above bound. Here we argue that $\hat{R}(P_n, r) = O(r^2 \log r)n$.

The following proposition is tailored to handle size Ramsey numbers of long paths.

Proposition 7.2. *Let $d, r > 0$. Let $G = (V, E)$ be a graph on $|V| = N$ vertices with average degree at least d. Assume that for an integer $0 < n < N$, every vertex subset $W \subset V$ of cardinality at most $2n$ spans at most $(d/8r)|W|$ edges. Then every edge subset $E_0 \subset E$ with $|E_0| \geq |E|/r$ contains a path of length n.*

Proof. Denote $G_0 = (V, E_0)$. Clearly the average degree of G_0 is at least d/r. Find a subgraph $G_1 = (V_1, E_1) \subseteq G_0$ of minimum degree at least $d/2r$. Since $|E_1| \geq (d/4r)|V_1|$, we have $|V_1| > 2n$ by the proposition's assumption. Let now $U \subset V_1$ be an arbitrary set of n vertices in G_1. Then the set $U \cup N_{G_1}(U)$ spans all edges of G_1 touching U, and there are at least $|U|(d/2r)/2 = nd/4r$ of them. Hence $|U \cup N_{G_1}(U)| \geq 2n$. We derive: $|N_{G_1}(U)| \geq n$. Applying Proposition 7.1 to G_1 with $k = \ell = n$, the proposition follows. $\square$

Note that the assumption in the proposition is fairly similar to the definition of locally sparse graphs as in Section 6.

We can now prove

Theorem 7.3. *There exists an absolute constant $C_0 > 0$ such that for any integer $r \geq 2$ and sufficiently large n, $\hat{R}(P_n, r) \leq C_0 r^2 \log r \cdot n$.*

This result can be read out from [35] (Theorem 4 there). An alternative proof has recently been offered by Dudek and Prałat [17], it also relies on a DFS-based argument. The estimate of Theorem 7.3 is close to being tight – Dudek and Prałat proved in a prior paper [16] that $\hat{R}(P_n, r) = \Omega(r^2)n$; a somewhat better absolute constant for the lower bound has been obtained by the author [35].

Proof (sketch). Let $N = 25rn$, and consider the binomial random graph $G(N, p)$ with $p = \frac{Cr \log r}{N}$, where $C > 0$ is sufficiently large constant. Then with high probability $G \sim G(N, p)$ has average degree $d = (1 + o(1))Np = (1 + o(1))Cr \log r$. In order to bound the typical local density of G, observe that for a subset $W \subset [N]$ of cardinality $|W| = i \leq 2n$, the number $e_G(W)$ of edges spanned by W in G is a binomial random variable with parameters $\binom{i}{2}$ and p. Recalling the easy estimate $\Pr[\text{Bin}(n, p) \geq k] \leq \binom{n}{k}p^k \leq (enp/k)^k$, we can write:

$$\Pr\left[e_G(W) \geq \frac{Ci \log r}{9}\right] \leq \left[\frac{\frac{ei^2}{2} \cdot \frac{Cr \log r}{N}}{\frac{Ci \log r}{9}}\right]^{\frac{Ci \log r}{9}} = \left(\frac{9er}{2} \cdot \frac{i}{N}\right)^{\frac{Ci \log r}{9}} = \left(\frac{9e}{50} \cdot \frac{i}{n}\right)^{\frac{Ci \log r}{9}}.$$

Summing over all $i \leq 2n$ and over all choices of $W \subset [N]$ with $|W| = i$, performing standard asymptotic manipulations, and recalling that $C > 0$ can be chosen to be sufficiently large, we derive that with high probability every such W spans at most $(d/8r)|W|$ edges.

Let now $f: E(G) \to [r]$ be an r-coloring of $E(G)$. Take E_0 to be the majority color, clearly $|E_0| \geq |E|/r$. Applying Proposition 7.2 to G and E_0, we conclude that E_0 contains a path of length n. It follows that every r-coloring of $E(G)$ contains a monochromatic path of length n. Hence $\hat{R}(P_n, r) \leq |E(G)| = O(r^2 \log r)n$. $\square$

7.3 Long cycles

With more effort (still based on the DFS algorithm and its properties) one can prove that every α-expander on n vertices contains a cycle of length linear in n. Notice that here, unlike in the case of paths, we cannot expect to get a cycle of any prescribed length, but rather a cycle at least as long as the lower bound given by the next theorem.

Theorem 7.4 ([35])**.** *Let $k > 0, \ell \geq 2$ be integers. Let $G = (V, E)$ be a graph on more than k vertices, satisfying:*

$$|N_G(W)| \geq \ell, \quad \text{for every } W \subset V, \; \frac{k}{2} \leq |W| \leq k\,.$$

Then G contains a cycle of length at least $\ell + 1$.

Proof (sketch). Observe first that G has a connected component C on more than k vertices. Let T be the tree obtained by applying the DFS algorithm to C, under an arbitrary order σ of its vertices, and let r be its root (which is the first vertex of C under σ). Then $|V(T)| = |V(C)| > k$. Now we find a vertex v and a subset X of its children in T such that the subtrees of T rooted at the vertices of X have between $k/2$ and k vertices in total. (This can be done by going from r down the tree till we find a vertex v whose subtree T_v has more than k vertices, but the subtrees T_x of its children in T are all less than k in their sizes; the required set X can then be recruited from the children of v in T.) Let $W = \bigcup_{x \in X} V(T_x)$. Denote by P the path in T from the root r to v. Then $k/2 \leq |W| \leq k$, and by Property **(P4)** of the DFS algorithm we have: $N_G(W) \subseteq V(P)$. Let $v^* \in V(P)$ be the farthest from v vertex in $N_G(W) \cap V(P) = N_G(W)$. Clearly its distance from v is at least $|N_G(W)| - 1 \geq \ell - 1$, by the theorem's assumption. Let w be a neighbor of v^* in W. Then the cycle, formed by the path in T from w to v^* (which includes the part of P from v to v^*) and the edge (w, v^*), has length at least $\ell + 1$. $\qquad\square$

Assuming that G is an α-expander on $n = \Omega(1/\alpha)$ vertices and applying Theorem 7.4 to G, we derive the existence of a cycle of length more than $\alpha n/4$ in G. The order of magnitude in n obtained is obviously optimal; moreover, the order of dependence on α in this bound is optimal as well, due to the same example of $K_{\alpha n, (1-\alpha)n}$.

7.4 Cycle lengths

Theorem 7.4 argues that if G is an α-expander on n vertices, then G has a cycle of length $\Omega(\alpha)n$. Also, due to Proposition 3.1, G has cycles as short as $O_\alpha(\log n)$. What can be said then about cycle lengths in G between the two extremes? Does G have linearly many in n cycle lengths? How well can one approximate an integer ℓ within the set of cycle lengths in G? The answer is given by the following theorem:

Theorem 7.5 ([21])**.** *For every $\alpha > 0$ there exist A, C, n_0 such that for every α-expander $G = (V, E)$ on $n \geq n_0$ vertices and every $\ell \in [C \ln n, n/C]$, the graph G has a cycle whose length is between ℓ and $\ell + A$.*

Hence the set $L(G)$ of cycle lengths in an α-expander G on n vertices is well spread and approximates every number ℓ in the range $[C \ln n, n/C]$ within an additive constant. In particular, G has linearly many cycle lengths.

The proof of Theorem 7.5 is somewhat involved. We outline instead the proof of a weaker statement, assuring that the set $L(G)$ of cycle lengths in an α-expander G on n

vertices has cardinality linear in n. The proof borrows some ideas from prior papers on cycle lengths, such as [52].

We start with the following lemma.

Lemma 7.6. *Let $\epsilon > 0$. There are constants $C_1 = C_1(\alpha, \epsilon)$ and $C_2 = C_2(\alpha, \epsilon) > 0$ such that for every $v \in V$ the graph G contains a spanning tree T rooted at v with levels $L_0 = \{v\}, L_1, \dots, L_{k_0}, \dots, L_{k_1}$ such that $k_0 = O(\log n), k_1 \le k_0 + C_1,$*

$$\left| \bigcup_{i=k_0}^{k_1} L_i \right| \ge (1 - \epsilon)n \,,$$

and the degrees of all $u \in \bigcup_{i=k_0}^{k_1} L_i$ in T are at most C_2.

The lemma guarantees that one can find a (spanning) tree T in G with few consecutive "thick" layers L_i, $i = O(\log n)$, containing nearly all vertices of G and such that the degrees in T of vertices in these thick layers are all bounded.

Let now T be the tree guaranteed by Lemma 7.6 with $\epsilon = \alpha/4$ and arbitrary v. Denote $W = \bigcup_{i=k_0}^{k_1} L_i$. Observe that every $U \subset W$, $|U| = n/2$, has at least $\alpha|U| - (n - |W|) \ge \alpha n/4$ neighbors in W. Hence W spans a path P of length at least $\alpha n/4$ by Proposition 7.1.

For each $u \in L_{k_0}$ let T_u be the full subtree of T rooted at u. Due to the bounded degree assumption on T we have $|T_u| \le C_2^{C_1}$.

For each $u \in L_{k_0}$ with $T_u \cap V(P) \ne \emptyset$ (there are at least $\frac{|V(P)|}{C_2^{C_1}} = \Theta(n)$ of those) choose a highest (closest to v) vertex $x \in T_u \cap V(P)$. Let X be the set of chosen vertices. Now, choose a maximum subset $X_0 \subseteq X$ such that the distance along P between any two vertices of X_0 is at least $C_1 + 1$. We have still $|X_0| = \Theta(n)$.

Let T' be the minimal subtree of T whose set of leaves is X_0, let r be its root. Due to the minimality T' branches at r. Let $A \subset X_0$ be the set of vertices of X_0 in one of the branches, and let $B = X_0 - A$. We can assume that $|B| \ge |A|$, implying $|B| \ge |X_0|/2$. Let a be an arbitrary vertex of A. At least $|B|/2 \ge |X_0|/4$ vertices of B are on the same side of a in P, denote this set by B_0.

For every $b \in B_0$ we get a cycle C_b in G as follows: first, the path from r to a in T', then we move from a to b along P, and finally take the path from b to r in T'. (These three pieces do not intersect internally, as a and b were chosen as highest vertices of P in their corresponding subtrees T_u, $u \in L_{k_0}$.) It is easy to check that the lengths of the cycles C_b are all distinct. Altogether we get $|B_0| = \Theta(n)$ different cycle lengths as promised.

8   Minors in expanding graphs

Minors is one of the most central notions in modern graph theory. It is thus only natural to expect to see meaningful research connecting expanding graphs and minors. And indeed, there have been several papers, addressing this subject directly or indirectly. We will mention some of them in this section, along with a description of new results due to Rajko Nenadov and the author.

First we recall the definition of a minor. Let $G = (V, E)$, $H = (U, F)$ be graphs with $U = \{u_1, \dots, u_k\}$. We say that G contains H as a *minor* if there is a collection $(U_1, \dots, U_k)$ of pairwise disjoint vertex subsets (supernodes) in V such that each U_i spans a connected subgraph in G, and whenever $(u_i, u_j) \in F$, the graph G has an edge between U_i and U_j. (Then contracting each U_i to a single vertex produces a copy of (a supergraph of) H.)

Notice that if G contains a minor of H then $|V(G)| \geq |V(H)|$ and $|E(G)| \geq |E(H)|$; these trivial bounds provide an obvious but meaningful benchmark for minor embedding statements.

As we discussed in Section 5, there is a very close connection between expanders and graphs without small separators. This connection enables to extend known results about embedding minors in graphs without small separators to claims about embedding minors in expanding graphs. In particular, it follows from the result of Plotkin, Rao and Smith [48] that an α-expanding graph on n vertices, $\alpha > 0$ a constant, contains a minor of the complete graph $K_{c\sqrt{n/\log n}}$. (In fact, the complete minor delivered by the result of [48] is shallow, i.e., each supernode U_i has diameter $O(\log n)$ in G.) An even stronger result has been announced by Kawarabayashi and Reed, who claimed in [29] that (translating it to the language of expanders) an α-expanding graph on n vertices contains a minor of $K_{c\sqrt{n}}$. Also, Kleinberg and Rubinfeld proved in [30] that an α-expander on n vertices of maximum degree Δ contains all graphs with $O(n/\log^\kappa n)$ vertices and edges as minors, for $\kappa = \kappa(\alpha, \Delta) > 1$. Phrasing it differently, a bounded degree α-expander is *minor universal* for all graphs with $O(n/\log^\kappa n)$ vertices and edges. This is optimal up to polylogarithmic factors due to the above stated trivial bound, as there exist n-vertex expanders with $\Theta(n)$ edges. (Formally, the number of vertices n is another bottleneck here.) It is worth mentioning here that the bounded maximum degree assumption $\Delta(G) = \Delta = O(1)$ is essential for the argument of Kleinberg and Rubinfeld, as the latter relies heavily on mixing properties of random walks; this issue has been briefly touched upon in Section 3.

We now present recent results, along with outlines of their proofs.

8.1 Large minors in expanding graphs

Our aim here is to present a universality-type result, guaranteeing the existence of a minor of *every* graph with bounded (as a function of n) number of vertices and edges in every α-expanding graph. This theorem is obtained jointly with Rajko Nenadov.

Theorem 8.1. *For every $\alpha > 0$ there exist c, n_0 such that the following holds. Let G be a graph on $n \geq n_0$ vertices, and let H be a graph with at most $cn/\log n$ vertices and edges. Then G has a separator of size at most αn, or G contains H as a minor.*

(We bound the number of vertices in H as well so as to avoid the pathological case of H having $O(n/\log n)$ edges, but many more vertices, say, more than n, thus making the embedding of H as a minor in G impossible.)

The guarantee $|V(H)| + |E(H)| = O(n/\log n)$ of the above theorem is asymptotically optimal up to multiplicative constants due to the following argument. Let G be an expanding graph of n vertices with logarithmic girth girth$(G) = \Theta(\log n)$, say, the celebrated Lubotzky-Phillips-Sarnak graph [43], and let H be a collection of k vertex-disjoint triangles. If H is a minor of G, then the image of every triangle in a minor embedding of H in G contains a cycle, and these cycles are disjoint for distinct triangles in H. It follows that $k \cdot \text{girth}(G) \leq n$, implying $k = O(n/\log n)$. Hence the largest collection of disjoint triangles that can be embedded in G as a minor counts $O(n/\log n)$ vertices.

Proof. The proof borrows substantially from the ideas of Plotkin, Rao and Smith [48]. Just as in many arguments of this sort, we start by assuming for convenience (of the proof) that the target graph $H = (U, F)$ has maximum degree (at most) three. This is possible as by splitting the vertices of H and edges emanating from them, we can construct a graph

$H' = (U', F')$ with $\Delta(H') \le 3$ and $|U'| = O(|U| + |F|)$ so that H is a minor of H'. Since minor containment is transitive, we can embed instead H' in G as a minor.

Rephrasing Corollary 3.2, we get the following handy lemma:

Lemma 8.2. *For every $\alpha > 0$ there exists $C_0 > 0$ such that every graph G on n vertices has diameter at most $C_0 \log n$, or contains a subset $U \subset V(G)$ satisfying: $|U| \le n/2$ and $|N_G(U)| \le \alpha n$.*

The proof of Theorem 8.1 is described as an algorithmic procedure that, given G and H with $|V(H)| = k \le \frac{\alpha n}{6 C_0 \log n}$ and $\Delta(H) \le 3$ (where C_0 is the constant from Lemma 8.2), outputs a small separator in G or a minor embedding of H in G. In fact, this procedure can be easily turned into a polynomial time algorithm performing this task. Let us assume $V(H) = [k]$. The algorithm maintains and updates a partition $V = A \cup B \cup C$ of the vertex set of G, where the set A will eventually contain a large and non-expanding set, witnessing the existence of a small separator in G, or the set B will contain the desired minor embedding of H; the set C will serve as a vertex reservoir. We will maintain $|N_G(A, C)| \le \alpha|A|$. We will also maintain and update a subset $I_0 \subset [k]$, which will describe the current induced subgraph of H, minor embedded in $G[B]$. Accordingly, there will be disjoint subsets $B_i \subseteq B$, $i \in I_0$, each spanning a connected subgraph in G; we commit ourselves to having always: $|B_i| \le 2 C_0 \log n$.

We initialize $A = B = \emptyset$, $C = V$, $I_0 = \emptyset$.

The procedure will repeat the following loop. Choose an arbitrary $i \in [k] \setminus I_0$. Let X be the set of neighbors of i in I_0, $|X| \le 3$. For $j \in X$, choose an arbitrary $v_j \in N_G(B_j, C)$. If for some $j \in X$ there is no such neighbor, we dump B_j into A and update accordingly: $A := A \cup B_j$, $B := B \setminus B_j$, $I_0 := I_0 \setminus \{j\}$. Otherwise, we apply Lemma 8.2 to the induced subgraph $G[C]$. In case it outputs a set U, $|U| \le |C|/2$, $|N_G(U, C \setminus U)| \le \alpha|U|$, we move U over to A and update: $A := A \cup U$, $C := C \setminus U$. In the complementary case, where $G[C]$ has logarithmic diameter, we find a subset $Y \subset C$ of size $|Y| \le 2 \cdot \operatorname{diam}(G[C]) \le 2 C_0 \log n$, such that $G[Y]$ is connected and contains vertices $v_j, j \in X$ (choose one of v_j's as the pivot vertex and connect it by paths of length at most $C_0 \log n$ to other v_j's). We then update: $B := B \cup Y$, $C := C \setminus Y$, $B_i := Y$, $I_0 = I_0 \cup \{i\}$. (In case $X = \emptyset$, we can simply take Y to be an arbitrary vertex of C and perform the same update.) The set B_i is connected and has an edge connecting it to B_j, $j \in X$.

Since we always add to the current A a piece U satisfying $|N_G(U, C)| \le \alpha|U|$, and the set C only shrinks as the algorithm proceeds, we observe that indeed at any point of the execution the current set A satisfies: $|N_G(A, C)| \le \alpha|A|$. Also, the set B is not too large – it is always composed of sets B_i, $i \in I_0$, and thus has at most $k \cdot 2 C_0 \log n = \alpha n/3$ vertices.

The above procedure eventually reaches the situation when A exceeds $n/3$ or $I_0 = [k]$. In the former case, it is easy to see that the first moment the size of A goes above $n/3$, it satisfies: $|A| \le 2n/3$. Also, $|N_G(A)| \le |B| + |N_G(A, C)| \le \alpha n/3 + \alpha|A| \le \alpha n/3 + 2\alpha n/3 = \alpha n$. In this case, the set $N_G(A)$ witnesses the existence of a small separator in G, separating between A and $V(G) \setminus (A \cup N_G(A))$. In case where the stopping criterion has been reached due to $I_0 = [k]$, the set B spans a minor of H in G, where the image of $u_i \in V(H)$ is the (connected) set B_i. $\square$

Corollary 8.3. *Let G be an α-expander on n vertices. Then G is minor universal for the class of all graphs H with at most $cn/\log n$ vertices and edges, for some $c = c(\alpha) > 0$.*

Let us add few remarks here. First, as we have indicated, the proof presented is algorithmic and allows to find in time polynomial in n a minor embedding of a given graph H

with $|V(H)|+|E(H)| \leq cn/\log n$. Secondly, the minor embedding produced is shallow, i.e., the image of every vertex $u \in V(H)$ in the embedding is a connected set U in G, spanning a graph of diameter logarithmic in n. Working out the details of the proof, we see that the constant c in Theorem 8.1, and thus in Corollary 8.3 depends quadratically on α. Finally, notice that the argument above can be trivially adapted to the case where $\alpha = \alpha(n)$ is vanishing.

8.2   Complete minors in expanding graphs

Theorem 8.1 guarantees that every expanding graph G on n vertices contains every graph H with $O(n/\log n)$ vertices and edges as a minor. As we argued in the previous subsection, this estimate on the size of H is optimal. For some particular types of H however we can do better. For example, due to Proposition 7.1 an expanding n-vertex graph G contains a path of length linear in n. One can also prove easily that such G contains a linearly sized star as a minor. (First find a linearly long path P in G, then argue that due to expansion one can find in G a matching M of linear size with every edge intersecting $V(P)$ in exactly one vertex; contracting $V(P)$ to a single vertex produces the required star.)

A particularly prominent class of graphs to try and find as minors is cliques. How large a clique minor can one hope to find in a weak expander? Denote by $ccl(G)$ the clique contraction number of G, which is the order of a largest clique minor to be found in G. Recall that if H is a minor of G, then $|E(H)| \leq |E(G)|$. Since weak expanders on n vertices can have linearly few in n edges, the best bound one can hope to get is $ccl(G) = \Theta\left(\sqrt{n}\right)$. Theorem 8.1 comes quite close and guarantees that a graph G on n vertices without separators of small linear size has $ccl(G) = \Omega\left(\sqrt{n/\log n}\right)$.

The above mentioned result of Kawarabayashi and Reed [29] improves this estimate, proving that a graph G on n vertices without sublinear separators satisfies: $ccl(G) = \Omega\left(\sqrt{n}\right)$. In fact, their result is much more general as it proves the following alternative: for any $h(n)$, an n-vertex graph G has a minor of the complete graph K_h, or a separator of order $O(h\sqrt{n})$.

Here we present a proof of the following theorem, a joint work with Rajko Nenadov.

Theorem 8.4. *For every $\alpha > 0$ there exist c, n_0 such that the following holds. Let G be a graph on $n \geq n_0$ vertices and maximum degree d. Assume that the edge isoperimetric number $i(G)$ satisfies: $i(G) \geq \alpha d$. Then $ccl(G) \geq c\sqrt{n}$.*

Here the edge isoperimetric number $i(G)$ of a graph $G = (V, E)$ is defined as:

$$i(G) = \min_{\substack{U \subset V \\ |U| \leq |V|/2}} \frac{e_G(U, V \setminus U)}{|U|}.$$

This quantity is fairly similar to the Cheeger constant $h(G)$ (to the extent that it is sometimes called the Cheeger constant of G); they are in fact exactly equal up to the factor d for the case of G being a d-regular graph.

The above theorem is superseded by the much more general result of Kawarabayashi and Reed. However, the proof of the latter is quite involved, and the full version of the conference paper [29] is still waiting to be published. We thus believe that presenting a reasonably short proof of Theorem 8.4 – whose outline will come shortly – certainly has merit.

Theorem 8.4 implies in particular that an α-expanding graph on n vertices with bounded degrees has a complete minor of the order of magnitude $\sqrt{n}$. The assumption about edge expansion being comparable to the maximum degree is an artifact of the proof and its techniques (specifically, the use of random walks), and can perhaps be lifted.

Proof. The proof is somewhat similar to the proof of Theorem 8.1, and uses some ideas from the arguments of Plotkin, Rao and Smith [48], and of Krivelevich and Sudakov [39]. Basically, we construct a complete minor of K_k in G, for an appropriately chosen k, a supernode by a supernode. Each time, when embedding the next supernode B_i in G, we need to make sure that B_i both spans a connected graph in G, and contains a neighbor of every presently embedded supernode B_j; due to expansion we can assume that the set of neighbors U_j of B_j is reasonably large. We will aim to do it economically, i.e., to get a relatively small set B_i, so as to allow enough room for embedding of k supernodes. A fairly natural idea for finding such a set, already suggested in [39], is to use random walks, and to take B_i to be the trace of a long enough random walk W, hoping it will hit each of the neighborhoods U_j. Our hopes are supported by the expansion properties of G, enabling to argue that W behaves similarly to a random set of the same size in terms of its hitting properties. In reality, it turns out to be more beneficial to take B_i to be the trace of W along with shortest paths from W to each U_j, arguing that typically the addition of these shortest paths does not add much to the size of W.

Now we get to the actual proof. We start it by introducing a basic tool based on random walks. Recall that a *lazy random walk* on a graph $G = (V, E)$ with the vertex set $V = \{1, \ldots, n\}$ is a Markov chain whose matrix of transition probabilities $P = P(G) = (p_{i,j})$ is defined by

$$p_{i,j} = \begin{cases} \frac{1}{2\deg_G(i)}, & \text{if } \{i, j\} \in E(G) \\ \frac{1}{2}, & \text{if } i = j, \\ 0, & \text{otherwise.} \end{cases}$$

This Markov chain has the stationary distribution π defined as $\pi(i) = \deg_G(i)/2e(G)$. As usual, for a subset $A \subseteq [n]$, we write $\pi(A) = \sum_{i \in A} \pi(i)$.

The following lemma gives an upper bound on the probability that a lazy random walk avoids a given subset $U \subseteq V(G)$.

Lemma 8.5. *Let G be a graph with n vertices and maximum degree $d = d(n)$. Then for any $U \subseteq V(G)$ the probability that a lazy random walk on G, which starts from the stationary distribution π and makes ℓ steps, does not visit U is at most*

$$\exp\left(-\frac{i(G)^3}{8d^3} \cdot \frac{|U|\ell}{n}\right).$$

We now describe the proof of Lemma 8.5.

Let $1 = \lambda_1 \geq \lambda_2 \geq \ldots \geq \lambda_n$ be the eigenvalues of the transition matrix P. The *spectral gap* of P (or of G) is defined as $\eta(G) = 1 - \lambda_2$. The following result of Mossel et al. [46] (more precisely, the first case of [46, Theorem 5.4]) relates the spectral gap to the probability that a random walk does not leave a specific subset. We state its version tailored for our needs.

Theorem 8.6 ([46])**.** *Let G be a connected graph with n vertices and let $\eta(G)$ be the spectral gap of the transition matrix $P = P(G)$. Then the probability that a lazy random walk of length ℓ, starting from a vertex chosen according to the stationary distribution π, does not leave a non-empty subset $A \subseteq V(G)$ is at most*

$$\pi(A)(1 - \eta(G)(1 - \pi(A)))^\ell.$$

In order to bound the spectral gap of G, we use the celebrated result of Jerrum and Sinclair [28, Lemma 3.3], relating the spectral gap of $P(G)$ to its *conductance* $\Phi(G)$, defined as

$$\Phi(G) = \min_{\substack{S \subseteq V \\ 0 < \pi(S) \leq 1/2}} \frac{\sum_{i \in S, j \notin S} \pi(i) p_{i,j}}{\pi(S)}.$$

Lemma 8.7 ([28])**.** *Let $G = (V, E)$ be a connected graph. Then the second eigenvalue λ_2 of the transition matrix P satisfies: $1 - \lambda_2 \geq \Phi(G)^2/2$.*

Finally, it is easy to verify that $\Phi(G) \geq i(G)/2d$, where as before $i(G)$ is the edge isoperimetric number of G, and d is its maximum degree.

Proof of Lemma 8.5. Consider a non-empty subset $U \subset V(G)$. Theorem 8.6 states that a lazy random walk never leaves the set $A = V(G) \setminus U$ (i.e., never visits U) with probability at most

$$\pi(A)(1 - \eta(G)(1 - \pi(A)))^\ell = (1 - \pi(U))(1 - \eta(G)\pi(U))^\ell$$

$$\leq \exp(-\eta(G)\ell\pi(U)) \leq \exp\left(-\frac{i(G)^2}{8d^2} \cdot \ell\pi(U)\right), \qquad (8.1)$$

where in the last inequality we used Lemma 8.7 and $\Phi(G) \geq i(G)/2d$. From the trivial bound $2e(G) \leq dn$ and $\delta(G) \geq i(G)$ we get

$$\pi(U) \geq \frac{|U|i(G)}{2e(G)} \geq |U|\frac{i(G)}{dn},$$

which after plugging into (8.1) gives the desired probability that a random walk misses U. $\qquad \square$

The next lemma implements our promise to utilize (extensions of) random walks as connected hitting sets.

Lemma 8.8. *For every $\beta > 0$ there exist positive $C = \Theta(1/\beta^3)$ and n_0 such that the following holds. Let G be a graph with $n \geq n_0$ vertices, maximum degree d, and $i(G) \geq \beta d$. Given k and s such that $ks \geq 2n$, and subsets $U_1, \ldots, U_k \subseteq V(G)$, where each U_j is of size $|U_j| \geq s$, there exists a connected set $Y \subseteq V(G)$ of size at most*

$$|Y| \leq C \cdot \frac{n}{s} \ln\left(\frac{ks}{n}\right)$$

which intersects every U_j.

We mention in passing that for many pairs $(k(n), s(n))$, by choosing subsets $U_i \subset [n]$, $|U_i| = s$, at random one can see that there is a familiy $\{U_i\}_{i=1}^k$, whose covering number has order of magnitude $(n/s)\log(ks/n)$. Thus Lemma 8.8 delivers a nearly optimal promise of the size of a hitting set, with an additional – and potentially important – benefit of this set spanning a connected subgraph in G.

Proof. Let

$$\ell = \frac{8}{\beta^3} \cdot \frac{n}{s} \ln\left(\frac{ks}{n}\right),$$

and consider a lazy random walk W in G which starts from the stationary distribution and makes ℓ steps. The desired connected subset Y is now constructed by taking the union of W with a shortest path between U_j and W, for each $j \in [k]$. We argue that with positive probability Y is of required size.

For $1 \leq j \leq k$, let X_j be the random variable measuring the distance from U_j to W in G. Then the set Y has expected size at most $\ell + \sum_{j=1}^{k} E[X_j]$.

In order to estimate the expectation of X_j, for a positive integer z write $p_{j,z} = Pr[X_j = z]$ and $p_{j,\geq z} = Pr[X_j \geq z]$. Trivially $p_{j,z} = p_{j,\geq z} - p_{j,\geq z+1}$. Hence

$$E[X_j] = \sum_{z \geq 1} z p_{j,z} = \sum_{z \geq 1} z \left(p_{j,\geq z} - p_{j,\geq z+1}\right) \leq \sum_{z \geq 1} p_{j,\geq z}.$$

Let $U_{j,z}$ be the set of vertices of G at distance at most z from U_j. The graph G is easily seen to be a β-expander, and similarly to Proposition 3.1 we have: $|U_{j,z}| \geq \min\left\{s(1+\beta)^z, \frac{n}{2}\right\}$. Applying Lemma 8.5 to $U_{j,z}$, we obtain for $z \leq \log(n/2)/\log(1+\beta)$:

$$p_{j,\geq z} \leq \exp\left\{-\frac{\beta^3(1+\beta)^z s\ell}{8n}\right\}.$$

Plugging in this estimate, recalling the value of ℓ, and doing fairly straightforward arithmetic, we can derive:

$$E[X_j] = \sum_{z \geq 1} p_{j,\geq z} = O\left(\left(\frac{n}{ks}\right)^{1+\beta}\right).$$

Hence

$$E[|Y|] = \ell + \sum_{j=1}^{k} E[X_j] = \ell + O\left(k\left(\frac{n}{ks}\right)^{1+\beta}\right) = O(\ell).$$

$\square$

We are now ready to complete the proof of Theorem 8.4. Let C be a constant given by Lemma 8.8 with $\beta = \alpha/4$, and set

$$b = \sqrt{\frac{8C}{\alpha} \cdot n} \qquad \text{and} \qquad k = \frac{\alpha n}{6b} = \Theta(\alpha^3)\sqrt{n}.$$

The argument here is somewhat similar to the proof of Theorem 8.1. We maintain and update a partition $V(G) = A \cup B \cup C$ so that always $e_G(A, C) < \beta d|A|$. There is also a set of indices $I_0 \subseteq [k]$ (corresponding to the set of supernodes $B_i, i \in I_0$), such that the subsets $B_i \subset B$ are pairwise disjoint, span connected subgraphs in G, and $|B_i| = b$. Moreover, if $i_1 \neq i_2 \in I_0$, then G has an edge between B_{i_1} and B_{i_2}.

Initially $A = B = \emptyset$, $C = V$, $I_0 = \emptyset$. As long as $|A| \leq n/3$ or $I_0 \subsetneq [k]$, we repeat the following loop. If there exists $i \in I_0$ such that $e_G(B_i, C) < \beta d|B_i|$, we move B_i to A and update: $A := A \cup B_i$, $B := B \setminus B_i$, $I_0 := I_0 \setminus \{i\}$. Otherwise, recalling that $\Delta(G) \leq d$, we derive that $|N_G(B_i, C)| \geq \beta b$ for all $i \in I_0$.

Now look into the induced subgraph $G[C]$. If there is a subset $U \subset C$, $|U| \leq |C|/2$ with $e_G(U, C \setminus U) < \beta d|U|$, we move U to A and update $A := A \cup U$, $C := C \setminus U$. Otherwise, the subgraph $G[C]$ satisfies $i(G[C]) \geq \beta d$. We now apply to it Lemma 8.8 with $s = \beta b$ and the sets $U_j = N_G(B_j, C)$, $j \in I_0$ (adding dummy sets U_j to get k' sets U_j altogether, all of size at least s, so that $2n < k's < 3n$) to get a connected set Y of size $|Y| \leq C\frac{n}{s}\ln 3 < b$.

Finally we extend Y to a connected set Y' of exactly b vertices in C, choose an arbitrary $i \in [k] \setminus I_0$ and update: $B := B \cup Y'$, $C := C \setminus Y'$, $B_i := Y$, $I_0 := I_0 \cup \{i\}$.

If the procedure terminated due to $I_0 = [k]$, then we have found the required K_k-minor. Suppose that this is not the case. We have $|B| < bk = \alpha n/6$. In the last step the set A has become larger than $n/3$, but still $|A| \leq 2n/3$. This implies $\min(|A|, |V \setminus A|) \geq n/3$, and hence $e_G(A, V \setminus A) \geq i(G)n/3 \geq \alpha nd/3$. On the other hand, $e_G(A, B) \leq d|B| \leq \alpha nd/6$, and $e_G(A, C) < \alpha d/4 \cdot |A| \leq \alpha dn/6$. Altogether $e_G(A, V \setminus A) < e_G(A, B) + e_G(A, C) < \alpha nd/6 + \alpha nd/6 = \alpha nd/3$. The obtained contradiction shows that the algorithm always runs to a successful end, outputting a minor of K_k in G. $\qquad\qquad\square$

Let us mention that [37] presents a result about large complete minors in expanding graph with growing degrees. Since the result's statement is somewhat involved and assumes stronger edge expansion of small sets, we decided not to state it here, referring the interested reader to [37] for details.

8.3 Large minors in random graphs

We conclude this section by presenting consequences of the above results for random graphs.

As we argued in Section 6 (see Corollary 6.6 there), a supercritical random graph $G \sim G(n, c/n)$, $c > 1$, typically contains an expanding subgraph G^* of bounded degree on linearly many vertices. Applying Corollary 8.3 and Theorem 8.4 to G^*, we derive the following corollaries.

Corollary 8.9. *For every $c > 1$ there exists $\delta > 0$ such that a random graph $G \sim G(n, c/n)$ is with high probability minor universal for the class of all graphs H with at most $\delta n/\log n$ vertices and edges.*

Corollary 8.10. *For every $c > 1$ there exists $\delta > 0$ such that a random graph $G \sim G(n, c/n)$ contains with high probability a minor of the complete graph on $\lfloor \delta\sqrt{n} \rfloor$ vertices.*

Both corollaries provide an asymptotically optimal order of magnitude. For Corollary 8.9, notice that $G \sim G(n, c/n)$ contains typically $o(n/\log n)$ cycles of length at most $c_0 \log n$, for some $c_0 = c_0(c) > 0$ small enough (straightforward first moment argument), and thus we can repeat the argument from Section 8.1. For Corollary 8.10, observe that $G \sim G(n, c/n)$ has typically $O(n)$ edges.

Corollary 8.10 reproves a result of Fountoulakis, Kühn and Osthus [18], obtained by applying direct ad hoc methods (and by working quite hard one may add).

Finally, recall that for $d \geq 3$ a random d-regular graph $G_{n,d}$ is an α-expander for $\alpha = \alpha(d) > 0$ (see Section 4 for discussion). We thus obtain:

Corollary 8.11. *For every integer $d \geq 3$ there exists $\delta > 0$ such that a random graph $G \sim G_{n,d}$ contains with high probability a minor of the complete graph on $\lfloor \delta\sqrt{n} \rfloor$ vertices.*

This has been proven in another paper of Fountoulakis, Kühn and Osthus [19], through the use of the so-called configuration model, and a quite substantial contiguity result about random graphs. The argument presented here provides an alternative, and perhaps more conceptual, way to derive this result.

References

[1] M. Ajtai, J. Komlós and E. Szemerédi, *The longest path in a random graph*, Combinatorica 1 (1981), 1–12.

[2] N. Alon, *Eigenvalues and expanders*, Combinatorica 6 (1986), 83–96.

[3] N. Alon, M. Krivelevich and P. Seymour, *Long cycles in critical graphs*, J. Graph Th. 35 (2000), 193–196.

[4] N. Alon and V. D. Milman, λ_1, *isoperimetric inequalities for graphs, and superconcentrators*, J. Combin. Th. Ser. B 38 (1985), 73–88.

[5] N. Alon, P. Seymour and R. Thomas, *A separator theorem for nonplanar graphs*, J. Amer. Math. Soc. 3 (1990), 801–808.

[6] N. Alon and J. H. Spencer, **The probabilistic method**, 4rd Ed., Wiley, New York, 2015.

[7] J. Beck, *On size Ramsey number of paths, trees, and circuits. I*, J. Graph Th. 7 (1983), 115–129.

[8] I. Ben-Eliezer, M. Krivelevich and B. Sudakov, *The size Ramsey number of a directed path*, J. Combin. Th. Ser. B 102 (2012), 743–755.

[9] M. Bednarska and O. Pikhurko, *Biased positional games on matroids*, Europ. J. Combin. 26 (2005), 271–285.

[10] B. Bollobás, *The isoperimetric number of random regular graphs*, Europ. J. Combin. 9 (1988), 241–244.

[11] S. Brandt, H. Broersma, R. Diestel and M. Kriesell, *Global connectivity and expansion: long cycles in f-connected graphs*, Combinatorica 26 (2006), 17–36.

[12] A. E. Brouwer and W. H. Haemers, **Spectra of graphs**, Springer, New York, 2012.

[13] F. R. K. Chung, **Spectral graph theory**, CBMS Regional Conf. Ser. Math. 92, Amer. Math. Soc., Providence, 1997.

[14] F. Chung, *Four proofs for the Cheeger inequality and graph partition algorithms*, Stud. Adv. Math. 48 (2010), 331–349.

[15] J. Dodziuk, *Difference equations, isoperimetric inequality and transience of certain random walks*, Trans. Amer. Math. Soc. 284 (1984), 787–794.

[16] A. Dudek and P. Prałat, *On some multicolor Ramsey properties of random graphs*, SIAM J. Discr. Math. 31 (2017), 2079–2092.

[17] A. Dudek and P. Prałat, *Note on the multicolour size-Ramsey number for paths*, Electr. J. Comb. 25 (2018), P 3.35.

[18] N. Fountoulakis, D. Kühn and D. Osthus, *The order of the largest complete minor in a random graph*, Random Struct. Alg. 33 (2008), 127–141.

[19] N. Fountoulakis, D. Kühn and D. Osthus, *Minors in random regular graphs*, Random Struct. Alg. 35 (2009), 444–463.

[20] J. Friedman, *A proof of Alon's second eigenvalue conjecture and related problems*, Memoirs Amer. Math. Soc., 2008.

[21] L. Friedman, M. Krivelevich and R. Nenadov, *Cycle lengths in expanding graphs*, manuscript.

[22] A. Frieze and M. Karoński, **Introduction to Random Graphs**, Cambridge Univ. Press, 2015.

[23] J. R. Gilbert, J. P. Hutchinson and R. E. Tarjan, *A separator theorem for graphs of bounded genus.* J. Algor. 5 (1984), 391–407.

[24] D. Hefetz, M. Krivelevich, M. Stojaković and T. Szabó, **Positional Games**, Birkhäuser, Basel, 2014.

[25] S. Hoory, N. Linial and A. Wigderson, *Expander graphs and their applications*, Bull. Amer. Math. Soc. (N.S.) 43 (2006), 439–561.

[26] S. Janson, T. Łuczak and A. Ruciński, **Random Graphs**, Wiley, New York, 2000.

[27] M. Jerrum and A. Sinclair, Conductance and the rapid mixing property for Markov chains: the approximation of the permanent resolved, Proc. 20^{th} Symp. Theory Comp. (STOC 88), 1988, pp. 235–244.

[28] M. Jerrum and A. Sinclair, *Approximate counting, uniform generation and rapidly mixing Markov chains*, Inform. Comp. 82 (1989), 93–133.

[29] K. Kawarabayashi and B. Reed, *A separator theorem in minor-closed classes*, Proc. 51st Symp. Found. Comp. Sci. (FOCS'10), 2010, 153–162.

[30] J. Kleinberg and R. Rubinfeld, *Short paths in expander graphs*, Proc. 37th Symp. Found. Comp. Sci. (FOCS'96), 1996, 86–95.

[31] J. Komlós and E. Szemerédi, *Topological cliques in graphs*, Combin. Probab. Comput. 3 (1994), 247–256.

[32] J. Komlós and E. Szemerédi, *Topological cliques in graphs II*, Combin. Probab. Comput. 5 (1996), 79–90.

[33] M. Krivelevich. Long paths and Hamiltonicity in random graphs. In: *Random Graphs, Geometry and Asymptotic Structure*, N. Fountoulakis and D. Hefetz, Eds., London Math. Soc. Student Texts Vol. 84, 2016, Cambridge Univ. Press, pp. 4–27.

[34] M. Krivelevich, *Finding and using expanders in locally sparse graphs*, SIAM J. Discr. Math. 32 (2018), 611–623.

[35] M. Krivelevich, *Long cycles in locally expanding graphs, with applications*, Combinatorica, to appear.

[36] M. Krivelevich, C. Lee and B. Sudakov, *Long paths and cycles in random subgraphs of graphs with large minimum degree*, Random Struct. Alg. 46 (2015), 320–345.

[37] M. Krivelevich and R. Nenadov, *Complete minors in graphs without sparse cuts*, submitted.

[38] M. Krivelevich and B. Sudakov, Pseudo-random graphs. In: *More sets, graphs and numbers*, E. Győri, G. O. H. Katona and L. Lovász, Eds., Bolyai Soc. Math. Stud. Vol. 15, 2006, pp. 199–262.

[39] M. Krivelevich and B. Sudakov, *Minors in expanding graphs*, Geom. Funct. Anal. 19 (2009), 294–331.

[40] M. Krivelevich and B. Sudakov, *The phase transition in random graphs — a simple proof*, Random Struct. Alg. 43 (2013), 131–138.

[41] D. A. Levin, Y. Peres and E. L. Wilmer, **Markov chains and mixing times**, American Math. Soc., Providence, 2009.

[42] R. J. Lipton and R. E. Tarjan, *A separator theorem for planar graphs*, SIAM J. Appl. Math. 36 (1979), 177–189.

[43] A. Lubotzky, R. Phillips and P. Sarnak, *Ramanujan graphs*, Combinatorica 8 (1988), 261–277.

[44] R. Montgomery, *Logarithmically small minors and topological minors*, J. Lond. Math. Soc. 91 (2015), 71–88.

[45] G. Moshkovitz and A. Shapira, *Decomposing a graph into expanding subgraphs*, Random Struct. Alg. 52 (2018), 158–178.

[46] E. Mossel, R. O'Donnell, O. Regev, J. Steif, and B. Sudakov, *Non-interactive correlation distillation, inhomogeneous Markov chains and the reverse Bonami-Beckner inequality*, Israel J. Math. 154 (2006), 299–336.

[47] M. S. Pinsker, On the complexity of a concentrator. In: *7th Int. Teletraffic Conf.*, pp. 318/1–318/4, 1973.

[48] S. Plotkin, S. Rao and W. Smith, *Shallow excluded minors and improved graph decompositions*, Proc. 5th Symp. Discr. Alg. (SODA'94), 1994, 462–470.

[49] O. Riordan, *Long cycles in random subgraphs of graphs with large minimum degree*, Random Struct. Alg. 45 (2014), 764–767.

[50] P. Sarnak, *What is ... an expander?*, Notices Amer. Math. Soc. 51 (2004), 762–763.

[51] A. Shapira and B. Sudakov, *Small complete minors above the extremal edge density*, Combinatorica 35 (2015), 75–94.

[52] B. Sudakov and J. Verstraëte, *Cycle lengths in sparse graphs*, Combinatorica 28 (2008), 357–372.

[53] N. C. Wormald. Models of random regular graphs. In: *Surveys in Combinatorics*, J. Lamb and D. Preece, Eds., London Math. Soc. Lect. Note Ser. Vol. 267, 1999, Cambridge Univ. Press, pp. 239–298.

School of Mathematical Sciences,
Raymond and Beverly Sackler Faculty of Exact Sciences,
Tel Aviv University,
Tel Aviv, 6997801,
Israel
krivelev@tauex.tau.ac.il