Chapter 1

Expander graphs: Basic theory

The objective of this text is to present a number of recent constructions of *expander graphs*, which are a type of sparse but "pseudorandom" graph of importance in computer science, the theory of random walks, geometric group theory, and in number theory. The subject of expander graphs and their applications is an immense one, and we will not possibly be able to cover it in full here. For instance, we will say almost nothing about the important applications of expander graphs to computer science, for instance in constructing good pseudorandom number generators, derandomising a probabilistic algorithm, constructing error correcting codes, or in building probabilistically checkable proofs. For such topics, see [**HoLiWi2006**].

Instead of focusing on applications, this text will concern itself much more with the task of *constructing* expander graphs. This is a surprisingly nontrivial problem. On one hand, we shall see that an easy application of the *probabilistic method* shows that a randomly chosen (large, regular, bounded-degree) graph will be an expander graph with very high probability, so expander graphs are extremely abundant. On the other hand, in many applications, one wants an expander graph that is more deterministic in nature (requiring either no or very few random choices to build), and of a more specialised form. For the applications to number theory or geometric group theory, it is of particular interest to determine the expansion properties of a very symmetric type of graph, namely a *Cayley graph*; we will also occasionally work with the more general concept of a *Schreier graph.* It turns out that such questions are related to deep properties of various groups G of Lie type (such as $\mathrm{SL}_2(\mathbf{R})$ or $\mathrm{SL}_2(\mathbf{Z})$), such as *Kazhdan's property (T)*, the first nontrivial eigenvalue of a Laplacian on a symmetric space G/Γ associated to G, the quasirandomness of G (as measured by the size of irreducible representations), and the product theory of subsets of G. These properties are of intrinsic interest to many other fields of mathematics (e.g., ergodic theory, operator algebras, additive combinatorics, representation theory, finite group theory, number theory, etc.), and it is quite remarkable[1] that a single problem — namely the construction of expander graphs — is so deeply connected with such a rich and diverse array of mathematical topics.

There are also other important constructions of expander graphs that are not related to Cayley or Schreier graphs, such as those graphs constructed by the *zigzag product construction*, but we will not discuss those types of graphs here; again, the reader is referred to [**HoLiWi2006**].

1.1. Expander graphs

We begin by defining formally the concept of an expander graph. As with many fundamentally important concepts in mathematics, there are a number of equivalent definitions of this concept. We will adopt a "spectral" perspective towards expander graphs, defining them in terms of a certain spectral gap, but will relate this formulation of expansion to the more classical notion of edge expansion later in this chapter.

We begin by recalling the notion of a graph. To avoid some very minor technical issues, we will work with undirected, loop-free, multiplicity-free graphs (though later, when we discuss Cayley graphs, we will allow loops and repetition).

Definition 1.1.1. A *graph* is a pair $G = (V, E)$, where V is a set (called the *vertex set* of G), and $E \subset \binom{V}{2}$ is a collection of unordered pairs $\{v, w\}$ of distinct elements v, w of V, known as the *edge set* of E. Elements of V or E are called *vertices* and *edges* of E. A graph is *finite* if the vertex set (and hence the edge set) is finite. If $k \geq 0$ is a natural number, we say that a graph $G = (V, E)$ is *k-regular* if each vertex of V is contained in exactly k edges in E; we refer to k as the *degree* of the regular graph G.

Example 1.1.2. The *complete graph* $(V, \binom{V}{2})$ on a vertex set V has edge set $\binom{V}{2} := \{\{v, w\} : v, w \in V, v \neq w\}$. If V has n elements, the complete graph is $n-1$-regular.

[1]Perhaps this is because so many of these fields are all grappling with aspects of a single general problem in mathematics, namely when to determine whether a given mathematical object or process of interest "behaves pseudorandomly" or not, and how this is connected with the symmetry group of that object or process.

In this course, we will mostly be interested in *constant-degree* large finite regular graphs, in which k is fixed (e.g., $k = 4$), and the number $n = |V|$ of vertices is going off to infinity.

Given a finite graph $G = (V, E)$, we let $\ell^2(V)$ be the finite-dimensional complex Hilbert space of functions $f \colon V \to \mathbf{C}$ with norm

$$\|f\|_{\ell^2(V)} := (\sum_{v \in V} |f(v)|^2)^{1/2}$$

and inner product

$$\langle f, g \rangle_{\ell^2(V)} := \sum_{v \in V} f(v)\overline{g(v)}.$$

We then can define the *adjacency operator* $A \colon \ell^2(V) \to \ell^2(V)$ on functions $f \in \ell^2(V)$ by the formula

$$Af(v) := \sum_{w \in V : \{v,w\} \in E} f(w),$$

thus $Af(v)$ is the sum of f over all of the neighbours of v; this is of course a linear operator. If one enumerates the vertices V as $v_1, \ldots, v_n$ in some fashion, then one can associate A with an $n \times n$ matrix, known as the *adjacency matrix* of G (with this choice of vertex enumeration).

As our graphs are undirected, the adjacency operator A is clearly self-adjoint (and the adjacency matrix is real symmetric). By the spectral theorem, A thus has (counting multiplicity) n real eigenvalues

$$\lambda_1 \geq \ldots \geq \lambda_n.$$

We will write λ_i as $\lambda_i(G)$ whenever we need to emphasise the dependence of the eigenvalues on the graph G.

The largest eigenvalue λ_1 is easily understood for k-regular graphs:

Lemma 1.1.3. *If G is a k-regular graph, then*

$$k = \lambda_1 \geq \lambda_n \geq -k.$$

Proof. Clearly $A1 = k1$ (where we write $1 \in \ell^2(V)$ for the constant function $v \mapsto 1$), and so k is an eigenvalue of A with eigenvector 1. On the other hand, for any $f, g \in \ell^2(V)$ with norm one, one has

$$\begin{aligned} |\langle Af, g\rangle_{\ell^2(V)}| &= \left| \sum_{v,w \in V : \{v,w\} \in E} f(w)\overline{g(v)} \right| \\ &\leq \frac{1}{2} \sum_{v,w \in V : \{v,w\} \in E} |f(w)|^2 + |g(v)|^2 \\ &\leq \frac{1}{2} k \sum_{w \in V} |f(w)|^2 + \frac{1}{2} k \sum_{v \in V} |g(v)|^2 \\ &= k, \end{aligned}$$

and so A has operator norm k (or equivalently, all eigenvalues of A lie between $-k$ and k). The claim follows. □

Now we turn to the next eigenvalue after λ_1.

Definition 1.1.4 (Expander family)**.** Let $\varepsilon > 0$ and $k \geq 1$. A finite k-regular graph is said to be a (one-sided) *ε-expander* if one has

$$\lambda_2 \leq (1 - \varepsilon)k$$

and a *two-sided ε-expander* if one also has

$$\lambda_n \geq -(1 - \varepsilon)k.$$

A sequence $G_i = (V_i, E_i)$ of finite k-regular graphs is said to be a one-sided (resp. two-sided) *expander family* if there is an $\varepsilon > 0$ such that G_i is a one-sided (resp. two-sided) ε-expander for all sufficiently large i.

Remark 1.1.5. The operator $\Delta := 1 - \frac{1}{k}A$ is sometimes known as[2] the *graph Laplacian*. This is a positive semidefinite operator with at least one zero eigenvalue (corresponding to the eigenvector 1). A graph is an ε-expander if and only if there is a *spectral gap* of size ε in Δ, in the sense that the first eigenvalue of Δ exceeds the second by at least ε. The graph Laplacian is analogous to the classical Laplacian in Euclidean space (or the Laplace-Beltrami operator on Riemannian manifolds); see Section 3.3 for a formalisation of this analogy.

The original definition of expander graphs focused on one-sided expanders, but it will be slightly more natural in this text to focus on the two-sided expanders. (But for Cayley graphs the two notions are almost equivalent; see Exercise 5.0.5.)

Strictly speaking, we have not defined the notion of a (one- or two-sided) expander graph in the above definition; we have defined an ε-expander graph

[2] It is also common to use the normalisation $k - A$ instead of $1 - \frac{1}{k}A$ in many texts, particularly if one wishes to generalise to graphs that are not perfectly regular.

for any given parameter $\varepsilon > 0$, and have defined the notion of an expander family, which is a *sequence* of graphs rather than for an individual graph. One could propose defining an expander graph to be a graph that is a one or two-sided ε-expander for some $\varepsilon > 0$ (or equivalently, a graph G such that the constant sequence $G, G, \ldots$ is an expander family), but this definition collapses to existing concepts in graph theory:

Exercise 1.1.1 (Qualitative expansion)**.** Let $k \geq 1$, and let $G = (V, E)$ be a finite k-regular graph.

(i) Show that $\lambda_2 = k$ if and only if G is not *connected*.

(ii) Show that $\lambda_n = -k$ if and only if G contains a nonempty *bipartite graph* as a connected component.

Thus, a graph is a one-sided expander for some $\varepsilon > 0$ if and only if it is connected, and a two-sided expander for some $\varepsilon > 0$ if and only if it is connected and not bipartite.

To obtain a more interesting theory, it is therefore necessary to either keep a more quantitative track of the ε parameter, or[3] work with expander families (typically involving vertex sets whose cardinality goes to infinity) rather than with individual graphs. Nevertheless, we will often informally drop the ε parameter (or the use of families) and informally refer simply to "expander graphs" in our discussion.

By taking the trace of the adjacency matrix or its square, one obtains some basic identities concerning the eigenvalues of a k-regular graph:

Exercise 1.1.2 (Trace formulae)**.** Let G be a k-regular graph on n vertices for some $n > k \geq 1$.

(i) Show that $\sum_{i=1}^n \lambda_i = 0$.

(ii) Show that $\sum_{i=1} \lambda_i^2 = nk$.

(iii) Show that $\max(|\lambda_2|, |\lambda_n|) \geq \sqrt{k} - o_{n\to\infty;k}(1)$, where $o_{n\to\infty;k}(1)$ denotes a quantity that goes to zero as $n \to \infty$ for fixed k.

Remark 1.1.6. The above exercise places an upper bound on how strong of a two-sided expansion one can obtain for a large k-regular graph. It is not quite sharp; it turns out that one can obtain the improvement

$$\max(|\lambda_2|, |\lambda_n|) \geq 2\sqrt{k-1} - o_{n\to\infty;k}(1),$$

a result of Alon and Boppana; see [**HoLiWi2006**] for a proof. Graphs with $\max(|\lambda_2|, |\lambda_n|) \leq 2\sqrt{k-1}$ are known as *Ramanujan graphs*, and (as the

[3] Alternatively, one could adopt a *nonstandard analysis* viewpoint and work with *ultra expander graphs* — i.e., ultraproducts of expander families — but we will postpone using this sort of viewpoint until Section 5.5.

name suggests) have connections to number theory, but we will not discuss this topic here; see for instance [**DaSaVa2003**] for more discussion.

We will give a probabilistic construction of an expander family later, but let us first give an example of a family of regular graphs that is *not* an expander family.

Exercise 1.1.3. For each $n \geq 3$, let G_n be the 2-regular graph whose vertex set is the cyclic group $\mathbf{Z}/n\mathbf{Z}$, and whose edge set is the set of pairs $\{x, x+1\}$ for $x \in \mathbf{Z}/n\mathbf{Z}$. (This is a basic example of a *Cayley graph*; such graphs will be discussed in more depth in Chapter 2.)

(i) Show that the eigenvalues of the adjacency operator A_n associated to G_n are $2\cos(2\pi j/n)$ for $j = 0, \ldots, n-1$. (*Hint:* You may find the *discrete Fourier transform* to be helpful.)

(ii) Show that G_n is not a one-sided expander family (and is thus not a two-sided expander family either). This is despite G_n always being connected (and nonbipartite for n odd).

The next exercise shows that the complete graph (Exercise 1.1.2) is an excellent expander; the whole point, though, of expander graph constructions is to come up with much sparser graphs that still have many of the connectivity and expansion properties of the complete graph.

Exercise 1.1.4. Let G be the complete graph on n vertices, which is of course an $n-1$-regular graph. Show that

$$\lambda_2 = \ldots = \lambda_n = -1$$

and so G is a one-sided $1+\frac{1}{n-1}$-expander and a two-sided $1-\frac{1}{n-1}$-expander. (This is not a counterexample to Exercise 1.1.2(iii), because the error term $o_{n\to\infty;k}(1)$ is only negligible in the regime when k is either fixed or is a very slowly growing function of n, and this is definitely not the case for the complete graph.)

Exercise 1.1.5. Let $G = (V, E)$ be a k-regular graph on n vertices. Let $G^c = (V, \binom{V}{2}\backslash E)$ be the *complement graph*, consisting of all the edges connecting two vertices in V that are not in E, thus G^c is an $n-k-1$-regular graph. Show that

$$\lambda_i(G^c) = -1 - \lambda_{n+2-i}(G)$$

for all $2 \leq i \leq n$.

Exercise 1.1.6. Let $n \geq 2$ be an even number, and let $G = K_{n/2,n/2}$ be the complete bipartite graph between two sets of $n/2$ vertices each, thus G is an $n/2$-regular graph. Show that $\lambda_n = -n/2$ and $\lambda_2 = \ldots = \lambda_{n-1} = 0$. Thus, G is a one-sided 1-expander, but is not a two-sided expander at all.

Exercise 1.1.7 (Expansion and Poincaré inequality)**.** If $G = (V, E)$ is a k-regular graph and $f\colon V \to \mathbf{C}$ is a function, define the gradient magnitude $|\nabla f|\colon V \to \mathbf{C}$ by the formula

$$|\nabla f(v)| := \left(\sum_{w \in V : \{v,w\} \in E} |f(w) - f(v)|^2 \right)^{1/2}.$$

Show that G is a one-sided ε-expander if and only if one has the *Poincaré inequality*

$$\|\nabla f\|^2_{\ell^2(V)} \geq 2k\varepsilon \|f\|^2_{\ell^2(V)}$$

whenever f has mean zero.

Exercise 1.1.8 (Connection between one-sided and two-sided expansion)**.** Let $G = (V, E)$ be a k-regular graph, let $\varepsilon > 0$, and let $G' = (V', E')$ be the bipartite version of G in which $V' := V \times \{1, 2\}$ and $E' := \{\{(v, 1), (w, 2)\} : \{v, w\} \in E\}$. Show that G is a two-sided ε-expander if and only if G' is a one-sided ε-expander. Based on this observation, introduce a notion of two-sided expansion for *directed* graphs, related to the *singular values* of the adjacency matrix, and connect it to one-sided expansion of the (undirected) bipartite version of this graph.

1.2. Connection with edge expansion

The intuition to explain Exercise 1.1.3 should be that while G_n is, strictly speaking, connected, it is not very *strongly* connected; the paths connecting a typical pair of points are quite long (comparable to n, the number of vertices) and it is easy to disconnect the graph into two large pieces simply by removing a handful of edges.

We now make this intuition more precise. Given two subsets F_1, F_2 of the vertex set V in a graph $G = (V, E)$, define $E(F_1, F_2) \subset F_1 \times F_2$ to be the set of all pairs $(v_1, v_2) \in F_1 \times F_2$ such that $\{v_1, v_2\} \in E$. Note that the cardinality of this set can be expressed in terms of the adjacency operator as

$$|E(F_1, F_2)| = \langle A1_{F_1}, 1_{F_2} \rangle_{\ell^2(V)}.$$

Define the *boundary* ∂F of a subset F of V to be the set $\partial F := E(F, V \backslash F)$, thus ∂F is essentially the set of all edges that connect an element of F to an element outside of F. We define the *edge expansion ratio* $h(G)$ of the graph G to be given by the formula

$$h(G) := \min_{F \subset V : |F| \leq |V|/2} \frac{|\partial F|}{|F|},$$

where F ranges over all subsets of V of cardinality at most[4] $|F| \leq |V|/2$. The quantity $h(G)$ can be interpreted as a type of isoperimetric constant for G, analogous to the *Cheeger constant* [**Ch1970**] of a compact Riemannian manifold, and so $h(G)$ is sometimes known as the *Cheeger constant* of the graph G.

Note that $h(G)$ is nonzero precisely when G is connected. (If G is disconnected, at least one of the components F will have cardinality less than $|V|/2$.) We have an analogous statement for one-sided expansion:

Proposition 1.2.1 (Weak discrete Cheeger inequality)**.** *Let $k \geq 1$, and let G_n be a family of finite k-regular graphs. Then the following are equivalent:*

(i) *G_n form a one-sided expander family.*

(ii) *There exists $c > 0$ such that $h(G_n) \geq c$ for all sufficiently large n.*

Proof. Let n be a large number. We abbreviate G_n as $G = (V, E)$.

We first establish the easy direction of this proposition, namely that (i) implies (ii). If n is large enough, then from the hypothesis (i) we have $\lambda_2 \leq (1-\varepsilon)k$ for some $\varepsilon > 0$ independent of n.

Let F be a subset of V with $|F| \leq |V|/2$. We consider the quantity

$$\langle A 1_F, 1_F \rangle_{\ell^2(V)}. \tag{1.1}$$

We split 1_F into a multiple $\frac{|F|}{|V|}1$ of the first eigenvector 1, plus the remainder $1_F - \frac{|F|}{|V|}1$, Using the spectral decomposition of A, we can upper bound (1.1) by

$$k \left\| \frac{|F|}{|V|} 1 \right\|_{\ell^2(V)}^2 + (1-\varepsilon)k \left\| 1_F - \frac{|F|}{|V|} 1 \right\|_{\ell^2(V)}^2$$

which after a brief calculation evaluates to

$$(1-\varepsilon)k|F| + \varepsilon k \frac{|F|^2}{|V|} \leq (1-\varepsilon/2)k|F|.$$

On the other hand, (1.1) is also equal to the number of (ordered) pairs of adjacent vertices $v, w \in F$. Since each $v \in F$ is adjacent to exactly k vertices, we conclude that there are at least $\varepsilon k|F|/2$ pairs v, w such that $v \in F$ and $w \notin F$. Thus

$$|\partial F| \geq \varepsilon |F|/2,$$

and so $h(G_n) \geq \varepsilon/2$, and the claim (ii) follows.

Now we establish the harder direction, in which we assume (ii) and prove (i). Thus we may assume that $h(G) \geq c$ for some $c > 0$ independent of n.

[4] Some upper bound on F is needed to avoid this quantity from degenerating, since ∂F becomes empty when $F = V$.

The difficulty here is basically that the hypothesis (ii) only controls the action $A1_F$ of A on indicator functions 1_F, whereas the conclusion (ii) basically requires us to understand Af for arbitrary functions $f \in \ell^2(V)$. Indeed, from the spectral decomposition one has

$$\lambda_2 = \sup_{f:\|f\|_{\ell^2(V)}=1;\langle f,1\rangle_{\ell^2(V)}=0} \langle Af, f\rangle_{\ell^2(V)},$$

so it suffices to show that

$$\langle Af, f\rangle_{\ell^2(V)} \leq (1-\varepsilon)k, \tag{1.2}$$

whenever $f \in \ell^2(V)$ has norm one and mean zero, and $\varepsilon > 0$ is independent of n. Since A has real matrix coefficients, we may assume without loss of generality that f is real.

The mean zero hypothesis is needed to keep the function f away from 1, but it forces f to change sign. It will be more convenient to first establish a variant of (1.2), namely that

$$\langle Af, f\rangle_{\ell^2(V)} \leq (1-c)k\|f\|^2_{\ell^2(V)} \tag{1.3}$$

whenever f is nonnegative and supported on a set of cardinality at most $|V|/2$.

Let us assume (1.3) for now and see why it implies (1.2). Let $f \in \ell^2(V)$ have norm one and mean zero. We split $f = f_+ - f_-$ into positive and negative parts, where f_+, f_- are nonnegative with disjoint supports. Observe that $\langle Af_+, f_-\rangle_{\ell^2(V)} = \langle Af_-, f_+\rangle_{\ell^2(V)}$ is positive, and so

$$\langle Af, f\rangle_{\ell^2(V)} \leq \langle Af_+, f_+\rangle_{\ell^2(V)} + \langle Af_-, f_-\rangle_{\ell^2(V)}.$$

Also we have

$$1 = \|f_+\|^2_{\ell^2(V)} + \|f_-\|^2_{\ell^2(V)}.$$

At least one of f_+ and f_- is supported on a set of size at most $|V|/2$. By symmetry we may assume that f_- has this small support. Let $\sigma > 0$ be a small quantity (depending on c) to be chosen later. If f_- has $\ell^2(V)$ norm at least σ, then applying (1.3) to f_- (and the trivial bound $\langle Af, f\rangle_{\ell^2(V)} \leq k\|f\|^2_{\ell^2(V)}$ for f_+) we have

$$\langle Af, f\rangle_{\ell^2(V)} \leq (1 - c\sigma^2)k,$$

which would suffice. So we may assume that f_- has norm at most σ. By Cauchy-Schwarz, this implies that $\sum_{x\in V} f_-(x) \leq \sigma|V|^{1/2}$, and thus (as f has mean zero) $\sum_{x\in V} f_+(x) \leq \sigma|V|^{1/2}$. Ordering the values $f_+(x)$ and applying Markov's inequality, we see that we can split f_+ as the sum of a function f'_+ supported on a set of size at most $|V|/2$, plus an error of ℓ^2

norm $O(\sigma)$. Applying (1.3) to f'_+ and f_- and using the triangle inequality (and Cauchy-Schwarz) to deal with the error term, we see that

$$\langle Af, f\rangle_{\ell^2(V)} \leq (1 - c + O(\sigma))k,$$

which also suffices (if σ is sufficiently small).

It remains to prove (1.3). We use the "wedding cake" decomposition, writing f as an integral,

$$f = \int_0^\infty 1_{F_t}\, dt,$$

where $F_t := \{x \in V : |f(x)| > t\}$. By construction, all of the F_t have cardinality at most $|V|/2$ and are nonincreasing in t. Also, a computation of the ℓ^2 norm shows that

$$\|f\|^2_{\ell^2(V)} = \int_0^\infty 2t|F_t|\, dt. \tag{1.4}$$

Expanding $\langle Af, f\rangle_{\ell^2(V)}$ and using symmetry, we obtain

$$2\int_0^\infty \int_0^t \langle A1_{F_s}, 1_{F_t}\rangle\, dsdt.$$

We can bound the integrand in two ways. First, since $A1_{F_s}$ is bounded by k, one has

$$\langle A1_{F_s}, 1_{F_t}\rangle \leq k|F_t|.$$

Second, we may bound $\langle A1_{F_s}, 1_{F_t}\rangle$ by $\langle A1_{F_s}, 1_{F_s}\rangle$.

On the other hand, from the hypothesis $h(G) \geq c$ we see that $|\partial F_s| \geq c|F_s|$, and hence

$$\langle A1_{F_s}, 1_{F_t}\rangle \leq (k - c)|F_s|.$$

We insert the first bound for $s \leq (1-\varepsilon)t$ and the second bound for $(1-\varepsilon)t < s \leq t$, for some $\varepsilon > 0$ to be determined later, and conclude that

$$\langle Af, f\rangle_{\ell^2(V)} \leq 2\int_0^\infty k(1-\varepsilon)t|F_t|\, dt + 2\int_0^\infty \int_{(1-\varepsilon)t}^t (k-c)|F_s|\, dsdt.$$

Interchanging the integrals in the second integral, we conclude that

$$\langle Af, f\rangle_{\ell^2(V)} \leq 2\int_0^\infty k(1-\varepsilon)t|F_t|\, dt + 2\int_0^\infty (k-c)\frac{\varepsilon}{1-\varepsilon}s|F_s|\, ds.$$

For ε small enough, one can check that $k(1-\varepsilon) + (k-c)\frac{\varepsilon}{1-\varepsilon} < k(1-\varepsilon')$ for some $\varepsilon' > 0$ depending only on ε, k, c, and the claim (1.3) then follows from (1.4). □

Example 1.2.2. The graphs in Exercise 1.1.3 contain large sets with small boundary (e.g., $\{1, \dots, m\}$ mod n for $1 \leq m \leq n/2$), which gives a nonspectral way to establish that they do not form an expander family.

Exercise 1.2.1. Show that if $k \leq 2$, then the only expander families of k-regular graphs are those families of bounded size (i.e., the vertex sets V_n have cardinality bounded in n).

Remark 1.2.3. There is a more precise relationship between the edge expansion ratio $h(G)$ and the best constant ε that makes G a one-sided ε-expander, namely the *discrete Cheeger inequality*

$$\frac{\varepsilon}{2}k \leq h(G) \leq \sqrt{2\varepsilon}k, \tag{1.5}$$

first proven in [**Do1984**] and [**AlMi1985**] (see also [**Al1996**]), based on the continuous isoperimetric inequalities in [**Ch1970**], [**Bu1982**]. The first inequality in (1.5) is already implicit in the proof of the above lemma, but the second inequality is more difficult to establish. However, we will not use this more precise inequality here.

There is an analogous criterion for two-sided expansion, but it is more complicated to state. Here is one formulation that is quite useful:

Exercise 1.2.2 (Expander mixing lemma)**.** Let $G(V, E)$ be a k-regular graph on n vertices which is a two-sided ε-expander. Show that for any subsets F_1, F_2 of V, one has

$$\left| |E(F_1, F_2)| - \frac{k}{n}|F_1||F_2| \right| \leq (1-\varepsilon)k\sqrt{|F_1||F_2|}.$$

(Actually, the factors $|F_1|, |F_2|$ on the right-hand side can be refined slightly to $|F_1| - \frac{|F_1|^2}{n}$ and $|F_2| - \frac{|F_2|^2}{n}$ respectively.)

Thus, two-sided expanders behave analogously to the pseudorandom graphs that appear in the *Szemerédi regularity lemma* [**Sz1978**] (but with the caveat that expanders are usually sparse graphs, whereas pseudorandom graphs are usually dense).

Here is a variant of the above lemma that more closely resembles Proposition 1.2.1.

Exercise 1.2.3. Let $k \geq 1$, and let G_n be a family of finite k-regular graphs. Show that the following are equivalent:

(i) G_n form a two-sided expander family.

(ii) There exists $c > 0$ such that whenever n is sufficiently large and F_1, F_2 are subsets of V_n of cardinality at most $|V_n|/2$, then

$$|E(F_1, F_2)| \leq (k - c)\sqrt{|F_1||F_2|}.$$

The exercises below connect expansion to some other graph-theoretic properties. On a connected graph G, one can define the *graph metric* $d\colon G \times$

$G \to \mathbf{R}^+$ by defining $d(v, w)$ to be the length of the shortest path from v to w using only edges of G. This is easily seen to be a metric on G.

Exercise 1.2.4 (Expanders have low diameter)**.** Let G be a k-regular graph on n vertices that is a one-sided ε-expander for some $n > k \geq 1$ and $\varepsilon > 0$. Show that there is a constant $c > 0$ depending only on k and ε such that for every vertex $v \in V$ and any radius $r \geq 0$, the ball $B(v, r) := \{w \in V : d(v, w) \leq r\}$ has cardinality

$$|B(v, r)| \geq \min((1 + c)^r, n).$$

(*Hint:* First establish the weaker bound $|B(v, r)| \geq \min((1 + c)^r, n/2)$.) In particular, G has diameter $O(\log n)$, where the implied constant can depend on k and ε.

Exercise 1.2.5 (Expanders have high connectivity)**.** Let G be a k-regular graph on n vertices that is a one-sided ε-expander for some $n > k \geq 1$ and $\varepsilon > 0$. Show that if one removes m edges from G for some $m \geq 0$, then the resulting graph has a connected component of size at least $n - Cm$, where C depends only on k and ε.

Exercise 1.2.6 (Expanders have high chromatic number)**.** Let G be a k-regular graph on n vertices that is a two-sided ε-expander for some $n > k \geq 1$ and $\varepsilon > 0$.

(i) Show that any *independent set*[5] in G has cardinality at most $(1 - \varepsilon)n$.

(ii) Show that the *chromatic number*[6] of G is at least $\frac{1}{1-\varepsilon}$. (Of course, this bound only becomes nontrivial for ε close to 1; however, it is still useful for constructing bounded-degree graphs of high chromatic number and large *girth*[7].)

Exercise 1.2.7 (Expansion and concentration of measure)**.** Let G be a k-regular graph on n vertices that is a one-sided ε-expander for some $n > k \geq 1$ and $\varepsilon > 0$. Let $f\colon G \to \mathbf{R}$ be a function which is Lipschitz with some Lipschitz constant K, thus $|f(v) - f(w)| \leq Kd(v, w)$ for all $v, w \in V$. Let M be a median value of f (thus $f(v) \geq M$ for at least half of the vertices

[5] A set of vertices in a graph G is *independent* if there are no edges in G that connect two elements in this set.

[6] The *chromatic number* of a graph G is the fewest number of colours needed to colour the vertices of the graph in such a way that no two vertices of the same colour are connected by an edge.

[7] The *girth* of a graph is the length of the shortest cycle in the graph, or infinity if the graph does not contain any cycles.

v, and $f(v) \leq M$ for at least half the vertices v; note that the median may be nonunique in some cases). Show that

$$|\{v \in V : |f(v) - M| \geq \lambda K\}| \leq Cn \exp(-c\lambda)$$

for all $\lambda > 0$ and some constants $C, c > 0$ depending only on k, ε.

1.3. Random walks on expanders

We now discuss a connection between expanders and random walks, which will be of particular importance in this text as a tool for demonstrating expansion. Given a k-regular graph G for some $k \geq 1$, and an initial vertex $v_0 \in G$, we define the random walk on G starting at v_0 to be a random sequence $v_0, v_1, v_2, \ldots$ of vertices in G defined recursively by setting, once $v_0, \ldots, v_i$ have been chosen, v_{i+1} to be one of the k neighbours of v_i, chosen at random[8]. For each i, let $\mu^{(i)} \colon G \to \mathbf{R}^+$ be the probability distribution of v_i, thus

$$\mu^{(i)}(v) = \mathbf{P}(v_i = v).$$

Thus $\mu^{(0)}$ is the Dirac mass δ_{v_0} at v_0, and we have the recursion

$$\mu^{(i+1)} = \frac{1}{k} A \mu^{(i)}$$

and thus

$$\mu^{(i)} = k^{-i} A^i \delta_{v_0}.$$

Among other things, this shows that the quantity $\|\mu^{(i)} - \frac{1}{|V_n|}\|_{\ell^2(V_n)}$, which measures the extent to which v_i is uniformly distributed, is nonincreasing in i. The rate of this decrease is tied to the expansion properties of the graph:

Exercise 1.3.1. Let $k \geq 1$, and let $G_n = (V_n, E_n)$ be a family of finite k-regular graphs. Let $\alpha > 1/2$. Show that the following are equivalent:

(i) The G_n are a two-sided expander family.

(ii) There is a $C > 0$ independent of n, such that for all sufficiently large n one has $\|\mu^{(i)} - \frac{1}{|V_n|}\|_{\ell^2(V_n)} \leq |V_n|^{-\alpha}$ for all $i \geq C \log |V_n|$, and all choices of initial vertex v_0.

Informally, the above exercise asserts that a two-sided expander on n vertices is one for which random walks (from an arbitrary starting point) become very close to uniform in just $O(\log n)$ steps. (Compare with, say,

[8] The existence of such a random process can be easily justified by using the *Kolmogorov extension theorem*; see, e.g., [**Ta2011**, Theorem 2.4.3]. Alternatively, one can select a random real number from $[0, 1]$ and express it in base k, obtaining an infinite string of digits in $\{0, \ldots, k-1\}$ that can be used (after arbitrarily ordering the edges emanating from each vertex in G) to generate the random sequence $v_0, v_1, \ldots$.

the graphs in Exercise 1.1.3, in which the random walks do not come close to mixing until time well beyond n^2, as indicated by the central limit theorem.) This rapid mixing is useful for many applications; for instance, it can be used in computer science to generate almost perfectly uniformly distributed random elements of various interesting sets (e.g., elements of a finite group); see [**HoLiWi2006**] for more discussion. It is also useful in number theory to facilitate certain sieving estimates, as will be discussed in Chapter 7.

Remark 1.3.1. In the above exercise, we assumed that the initial vertex v_0 was deterministic rather than random. However, it is easy to see (from Minkowski's inequality) that the exercise also holds if we permit v_0 to be drawn from an arbitrary probability distribution on V_n, rather than being a single deterministic vertex. Thus one can view the uniform distribution in a two-sided expander to be a very strong attractor for all the other probability distributions on the vertex set.

Remark 1.3.2. The ℓ^2 norm is the most convenient norm to use when using the spectral theorem, but one can certainly replace this norm if desired by other norms (e.g., the ℓ^1 norm, which in this finite setting is the same thing as the *total variation norm*), after adjusting the lower bound of A slightly, since all norms on a finite-dimensional space are equivalent (though one typically has to concede some powers of $|V_n|$ to attain this equivalency). One can also use some other norm-like quantities here to measure distance to uniformity, such as *Shannon entropy*, although we will not do so here.

Note that for graphs that are only one-sided expanders instead of two-sided expanders, the random walk is only partially mixing, in that the probability distribution tends to flatten out rapidly, but not converge as rapidly (or at all) to the uniform distribution. For instance, in the case of the complete bipartite graph $K_{n/2,n/2}$, it is clear that the random walk simply alternates between the two vertex sets of size $n/2$ in the bipartite graph (although it is uniformly distributed in each set). But this lack of rapid mixing can be dealt with by replacing the random walk with the *lazy random walk* $w_0, w_1, w_2, \ldots$, which is defined similarly to the random walk $v_0, v_1, v_2, \ldots$ except that w_{i+1} is only set equal to a randomly chosen neighbour of w_i with probability $1/2$, and remains equal to w_i with probability $1/2$. (Here, one can also select probabilities other than $1/2$, as this will not significantly affect the exercise below.) Indeed, we have:

Exercise 1.3.2. Let $k \geq 1$, and let $G_n = (V_n, E_n)$ be a family of finite k-regular graphs. Let $\alpha > 1/2$. Show that the following are equivalent:

(i) The G_n are a one-sided expander family.

(ii) There is a $C > 0$ independent of n, such that for all sufficiently large n one has $\|\nu^{(i)} - \frac{1}{|V_n|}\|_{\ell^2(V_n)} \leq |V_n|^{-\alpha}$ for all $i \geq C \log n$, and all choices of initial vertex w_0, where $\nu^{(i)}$ is the law of the random walk w_i.

1.4. Random graphs as expanders

We now turn to the task of constructing expander families of k-regular graphs. The first result in this direction was by Pinsker [**Pi1973**] (with a closely related result also established by Barzdin and Kolmogorov [**KoBa1967**]), who showed that if one chose a k-regular graph on n vertices randomly, and then sent n to infinity along a sufficiently sparse sequence, the resulting sequence would be almost surely[9] be an expander family. We will not quite prove this result here (because it requires some understanding of the probability distribution of k-regular graphs, which has some subtleties as the parity obstruction already indicates), but establish a closely related result, in which k is restricted to be even (to avoid parity problems) and sufficiently large (for convenience) and the k-regular graph is drawn from a slightly nonuniform distribution.

Before we do this, though, let us perform a heuristic computation as to why, when k is fixed but large, and n goes to infinity, one expects a "random" k-regular graph $G = (V, E)$ on n vertices to be an expander. For simplicity, we work with the one-sided expansion condition. By Proposition 1.2.1, we would then like to say that with high probability, one has

$$|\partial F| \geq c|F|$$

for all $F \subset V$ with cardinality at most $n/2$, and some $c > 0$ independent of n. An equivalent formulation would be to say that the neighbourhood $N(F)$ of F has to have cardinality at least $(1 + c')|F|$ for all $F \subset V$ with cardinality at most $n/2$, and some $c' > 0$ independent of n. Thus, we wish to exclude the possibility that there are sets $F \subset F' \subset V$ with $|F| \leq n/2$ and $|F'| \leq (1 + c')|F|$, for which all the edges starting from F end up in F'.

To bound this failure probability, we use the union bound: we try each pair F, F' in turn, bound the probability that the claim follows for that particular value of F, F', and then sum in F, F'. The goal is to obtain a total probability bound of $o_{n\to\infty;k}(1)$, that goes to zero as $n \to \infty$ for fixed k.

Accordingly, pick $1 \leq r \leq n/2$, and then pick $F \subset V$ of cardinality r, and then $F \subset F' \subset V$ of cardinality $r + r'$, where $r' := \lfloor c'r \rfloor + 1$ for some small constant $c' > 0$ to be chosen later. For each fixed r, there are $\frac{n!}{r!r'!(n-r-r')!}$

[9]Here, of course, one needs to avoid the parity obstruction that one cannot have a k-regular graph on n vertices if k and n are both odd.

choices of F and F'. For each F, F', there are kr edges emanating from F. Intuitively, if we choose the graph randomly, each edge has a probability about $\frac{r+r'}{n}$ of landing back in F', so the probability that they all do is about $(\frac{r+r'}{n})^{kr}$. So the failure rate should be about

$$\sum_{1 \le r \le n/2} \frac{n!}{r!r'!(n-r-r')!} \left(\frac{r+r'}{n}\right)^{kr}.$$

We can bound $\frac{n!}{r!r'!(n-r-r')!}$ somewhat crudely as $\frac{n^{r+r'}}{(r+r')!} O(1)^r$. Applying Stirling's formula, we obtain a bound of

$$\sum_{1 \le r \le n/2} O(1)^r \left(\frac{r+r'}{n}\right)^{kr-r-r'}.$$

For c small enough, $(r+r')/n$ is less than 0.6 (say), and then (for k large enough) we see that this series is bounded by $o_{n\to\infty;k}(1)$ as required.

Now we turn to the rigorous construction of random expander graphs. We will assume that k is a large (but fixed) even integer, $k = 2l$. To build a $2l$-regular graph on n vertices $\{1, \dots, n\}$, what we will do is pick l permutations $\pi_1, \dots, \pi_l : \{1, \dots, n\} \to \{1, \dots, n\}$, and let G be the graph formed by connecting v to $\pi_i(v)$ for all $v \in \{1, \dots, n\}$ and $i = 1, \dots, l$. This is not always a $2l$-regular graph, but we will be able to show the following two claims (if k is large enough):

Proposition 1.4.1 (G can be k-regular). *The graph G is $2l$-regular with probability at least $c - o_{n\to\infty;k}(1)$, where $c > 0$ depends only on k.*

Proposition 1.4.2 (G usually expands). *There is an $\varepsilon > 0$ depending only on $k = 2l$, such that the probability that G is $2l$-regular but not a one-sided ε-expander is $o_{n\to\infty;k}(1)$.*

Putting these two propositions together, we conclude:

Corollary 1.4.3. *With probability at least $c - o_{n\to\infty;k}(1)$, G is a $2l$-regular one-sided ε-expander.*

In particular, this allows us to construct one-sided expander families of k-regular graphs for any fixed large even k.

Remark 1.4.4. If one allowed graphs to have multiple edges and loops, then it would be possible to dispense with the need for Proposition 1.4.1, and show that G (now viewed as a $2l$-regular graph with multiple edges and loops) is a one-sided ε-expander with probability $1 - o_{n\to\infty;k}(1)$. (This requires extending results such as the weak discrete Cheeger inequality to the case when there are multiple edges and loops, but this turns out to be straightforward.) However, we will not do so here as it requires one to introduce a slight amount of additional notation.

Let us prove Proposition 1.4.2 first, which will follow the informal sketch at the beginning of this section. By Proposition 1.2.1, it suffices to show that there is a $c > 0$ depending only on k such that the probability that G is $2l$-regular with $h(G) \leq c$ is $o_{n\to\infty;k}(1)$. As in the sketch, we first bound for each $1 \leq r \leq n/2$ and $F \subset F' \subset \{1,\dots,n\}$ with $|F| = r$ and $|F'| = r + r'$, where $r' := \lfloor cr \rfloor + 1$, the probability that all the edges from F end up in F'. A necessary condition for this to occur is that $\pi_i(F) \subset F'$ for each $i = 1,\dots,k$. For each i, and for fixed r, F, F', the probability that the random permutation π_i does this, is

$$\frac{\binom{r+r'}{r}}{\binom{n}{r}} \leq \left(\frac{r+r'}{n}\right)^r,$$

so the total failure probability can be bounded by

$$\sum_{1\leq r\leq n/2} \frac{n!}{r!r'!(n-r-r')!}\left(\frac{r+r'}{n}\right)^{kr/2},$$

which is acceptable as discussed previously.

Now we turn to Proposition 1.4.1. We observe that G will be $2l$-regular unless there are distinct i, j and a vertex $v \in \{1,\dots,n\}$ such that either $\pi_i(v) = \pi_j(v)$, $\pi_i(v) = \pi_j^{-1}(v)$, $\pi_i(v) = v$, or $\pi_i(v) = \pi_i^{-1}(v)$ (as such cases lead to repeated edges or loops). Unfortunately, each of these events can occur with a fairly sizeable probability (e.g., for each i, j, the probability that $\pi_i(v) = \pi_j(v)$ for some v is about $1 - 1/e$, by the classical theory of *derangements*), so the union bound is not enough here. Instead, we will proceed by an interpolant between the union bound and the *inclusion-exclusion formula*, known as the *Bonferroni inequalities*:

Exercise 1.4.1 (Bonferroni inequalities)**.**

(i) Show that if $n, k \geq 0$ are natural numbers, that

$$\sum_{j=0}^{k}(-1)^j\binom{n}{j} \geq 1_{n=0}$$

when k is even, and

$$\sum_{j=0}^{k}(-1)^j\binom{n}{j} \leq 1_{n=0}$$

when k is odd, where we adopt the convention that $\binom{n}{j} = 0$ when $j > n$.

(ii) Show that if $N, k \geq 0$ and $E_1, \ldots, E_N$ are events, then

$$\sum_{j=0}^{k}(-1)^j \sum_{1 \leq i_1 < \ldots < i_j \leq N} \mathbf{P}(E_{i_1} \cap \ldots \cap E_{i_j}) \geq \mathbf{P}\left(\overline{\bigcup_{i=1}^{N} E_i}\right)$$

when k is even, and

$$\sum_{j=0}^{k}(-1)^j \sum_{1 \leq i_1 < \ldots < i_j \leq N} \mathbf{P}(E_{i_1} \cap \ldots \cap E_{i_j}) \leq \mathbf{P}\left(\overline{\bigcup_{i=1}^{N} E_i}\right)$$

when k is odd. (Note that the $k = 1$ case of this inequality is essentially the union bound, and the $k = N$ case is the inclusion-exclusion formula.) Here we adopt the convention that the empty intersection occurs with probability 1.

We return to the proof of Proposition 1.4.1. By a conditioning argument, it suffices to show the following:

Proposition 1.4.5. *Let $1 \leq i \leq k$, and suppose that the permutations $\pi_1, \ldots, \pi_{i-1}$ have already been chosen (and are now viewed as fixed deterministic objects). Let $\pi_i \colon \{1, \ldots, n\} \to \{1, \ldots, n\}$ be a permutation chosen uniformly at random. Then with probability at least $c - o_{n \to \infty; k}(1)$ for some $c > 0$ depending only on k, one has $\pi_i(v) \neq \pi_j(v), v, \pi_i^{-1}(v)$ for all $v \in \{1, \ldots, n\}$ and $j = 1, \ldots, i-1$.*

It remains to establish Proposition 1.4.5. We modify an argument of Bollobas [**Bo2001**], based on some technical counting asymptotics which we defer to the exercises. Consider the set $\Omega \subset \{1, \ldots, n\} \times \{1, \ldots, n\}$ of pairs

$$\begin{aligned} \Omega := \{&(v, \pi_j(v)) : v = 1, \ldots, n; j = 1, \ldots, i-1\} \\ &\cup \{(v, \pi_j(v)^{-1}) : v = 1, \ldots, n; j = 1, \ldots, i-1\} \\ &\cup \{(v, v) : v = 1, \ldots, n\}. \end{aligned}$$

Each pair (v, w) in Ω gives rise to a bad event $E_{(v,w)} := (\pi_i(v) = w)$. Each pair $(v, w) \in \{1, \ldots, n\}^2 \backslash \Omega$ gives rise to another bad event

$$E_{(v,w)} := (\pi_i(v) = w) \wedge (\pi(w) = v).$$

Note that $E_{(v,w)} = E_{(w,v)}$ for $(v, w) \in \{1, \ldots, n\}^2 \backslash \Omega$. To avoid this collision problem, we work in the space

$$\Omega' := \Omega \cup \{(v, w) \in \{1, \ldots, n\}^2 \backslash \Omega : v < w\}.$$

Our task is now to show that

$$\mathbf{P}\left(\overline{\bigcup_{s \in \Omega'} E_s}\right) \geq c - o_{n \to \infty; k}(1).$$

It is easy to check (by direct counting arguments) that

$$\mathbf{P}(E_s) = \frac{1}{n} + O\left(\frac{1}{n^2}\right) \tag{1.6}$$

for $s \in \Omega$, and

$$\mathbf{P}(E_s) = \frac{1}{n^2} + O\left(\frac{1}{n^3}\right) \tag{1.7}$$

for $s \notin \Omega'$. In particular, if we set

$$\alpha := \sum_{s \in \Omega'} \mathbf{P}(E_s),$$

then we have

$$\alpha = 2i - \frac{1}{2} + O\left(\frac{1}{n}\right).$$

More generally, for any fixed j, the same sort of counting arguments (which we leave as an exercise) gives the approximate independence property

$$\mathbf{P}(E_{s_1} \cap \ldots \cap E_{s_j}) = (1 + O_{j,k}\left(\frac{1}{n}\right))\mathbf{P}(E_{s_1}) \ldots \mathbf{P}(E_{s_j}) \tag{1.8}$$

whenever $s_1 = (v_1, w_1), \ldots, s_j = (v_j, w_j)$ are distinct elements of Ω' with the disjointness property $\{v_i, w_i\} \cap \{v_{i'}, w_{i'}\} = \emptyset$ for all $1 \leq i < i' \leq j$. If the disjointness property fails, we can still obtain the weaker bound

$$\mathbf{P}(E_{s_1} \cap \ldots \cap E_{s_j}) = O_{j,k}(\mathbf{P}(E_{s_1}) \ldots \mathbf{P}(E_{s_j})). \tag{1.9}$$

(The subscripts in the $O()$-notation indicate that the implied constant in that notation can depend on the subscripted parameters.) Summing this, we conclude that

$$\sum_{s_1, \ldots, s_j \in \Omega', \text{ distinct}} \mathbf{P}(E_{s_1} \cap \ldots \cap E_{s_j}) = \alpha^j + O_{j,k}\left(\frac{1}{n}\right) \tag{1.10}$$

and thus by the Bonferroni inequalities

$$\mathbf{P}(\overline{\bigcup_{s=1}^{2in} E_s}) \geq \sum_{j=1}^{m} (-1)^{j-1} \frac{\alpha^j}{j!} + O_{m,k}\left(\frac{1}{n}\right)$$

for any odd m. But from Taylor series expansion, the sum on the right-hand side converges to the positive quantity $e^{-\alpha}$, and the claim follows by taking m to be a sufficiently large odd number depending on k.

Exercise 1.4.2. Verify the estimates (1.6), (1.8), (1.9), (1.10).

Remark 1.4.6. Another way to establish Proposition 1.4.1, via the "swapping method", was pointed out to me by Brendan McKay. The key observation is that if l permutations $\pi_1, \ldots, \pi_l$ have m problematic edges (i.e., edges that are either repeated or loops), and then one applies m random transpositions to the $\pi_1, \ldots, \pi_l$ (as selected at random), then with probability $\gg_m n^{-m}$ (if n is sufficiently large depending on m), all problematic edges are erased and one obtains a random regular graph. Conversely, if one starts with a random regular graph and applies m random transpositions, the probability of obtaining m problematic edges as a result is $O_m(n^{-m})$. Combining the two facts, we see that p_m is the probability of having m problematic edges, then $p_0 \gg p_m$ for each fixed m (and n sufficiently large depending on m). Since the expected number of problematic edges is bounded, the desired bound $p_0 \gg 1$ then follows from Markov's inequality and the pigeonhole principle.

Exercise 1.4.3. Show that "one-sided" can be replaced with "two-sided" in Proposition 1.4.2 and hence in Corollary 1.4.3.

Remark 1.4.7. It turns out that the random k-regular graphs formed by taking l permutations as indicated above, and conditioning on the event that there are no "collisions" (so that one genuinely gets a k-regular graph) does not quite give a uniform distribution on the k-regular graphs. However, it is close enough to one that any property which is true with probability $1 - o_{n\to\infty;k}(1)$ for this model of random k-regular graph, is also true with probability $1-o_{n\to\infty;k}(1)$ for uniform k-regular graphs, and conversely. This fact (known as *contiguity* of the two random models, and analogous to the concept of mutually absolutely continuous measures in measure theory) is established, for instance, in [**Wo1999**]. As a consequence of this fact (and a more refined version of the above analysis), one can show that $1-o_{n\to\infty;k}(1)$ of all k-regular graphs on n vertices are ε-expanders for some $\varepsilon = \varepsilon_k > 0$ if $k \geq 3$ (assuming of course the parity requirement that nk be even, otherwise there are no k-regular graphs at all).