CHAPTER 2

Graphs and expansion

2.1. Definitions

In these notes 'graph' will always mean an undirected, simple and (unless explicitly stated) finite graph. Here is the definition.

Definition 2.1. A *graph* is a pair $\Gamma = (V, E)$, where V is a non-empty set, and E a subset of $V^{[2]} = \{X \subseteq V \mid |X| = 2\}$. The elements of V are called *vertices* of the graph Γ, those of E *edges*. The graph is *finite* if V (and thus E) is finite.

Given two vertices x, y of a graph $\Gamma = (V, E)$, we write $x \sim y$ if x and y are *adjacent*, that is $x \neq y$ and $\{x, y\} \in E$. If $e = \{x, y\} \in E$, we say that x, y are the extremes of e, and that the vertex x and the edge e (as well as y and e) are *incident*.
A *subgraph* of a graph $\Gamma = (V, E)$ is just a pair (A', E') with $\emptyset \neq A' \subseteq A$ and $E' \subseteq E$. More important for our purposes is the concept of *induced subgraph*, where if X is a non-empty subset of the set V then the subgraph of $\Gamma = (V, E)$ induced by X is the subgraph of Γ with vertex set X and edges all those of Γ whose extremes belong to X, in formal terms $(X, E \cap X^{[2]})$.

For $x \in V$, the number of vertices adjacent to x is called the *degree* of x and denoted $d_\Gamma(x)$. Clearly, $d_\Gamma(x)$ coincides with the number of different edges that are incident to x. An immediate double counting argument yields,

$$\sum_{x \in V} d_\Gamma(x) = 2|E|. \tag{2.1}$$

Definition 2.2. Let $1 \leq k \in \mathbb{N}$. A graph Γ is *k-regular* if $d_\Gamma(x) = k$ for every vertex x of Γ. We say that a graph is regular if it is k-regular for some positve integer k.

If the graph $\Gamma = (V, E)$ is k-regular then, from (2.1), $2|E| = k|V|$.

Example 10 (Complete graphs). Given a not empty set V, the *complete graph* on V is $K_V = (V, V^{[2]})$. If $|V| = n$ we denote it by K_n. Then K_n is $(n-1)$-regular, and has $\binom{n}{2}$ edges.

Example 11 (Cycles). Let $n \geq 3$. The *n-cycle* is the graph $C_n = (V, E)$, where $V = \{1, \dots, n\}$ and $E = \{\{i, i+1\} \mid i = 1, \dots n-1\} \cup \{\{n, 1\}\}$. C_n is 2-regular, and has as many edges than vertices.

EXAMPLE 12 (Hypercubes). Let $n \geq 1$. The *n-hypercube* is the graph Q_n where the set of vertices is the set $\mathbb{Z}_2^n$ of all n-tuples in $\{0,1\}$ and two vertices are adjacent if and only if they differ for exactly one coordinate. Q_n is n-regular.

DEFINITION 2.3. Let $\Gamma = (V, E)$ be a graph.

– A *walk* in Γ is a sequence of vertices $x_0, x_1, \ldots, x_n$ such that $x_{i-1} \sim x_i$ for every $i = 1, \ldots n$; the integer n is the *length* of the walk.

– A *path* is a walk in which all edges $\{x_{i-1}, x_i\}$ are distinct; a path $x_0, x_1, \ldots, x_n$ is *simple* if all vertices are distinct, except possibly for $x_0 = x_n$.

– A *cycle* is a simple path of length $n \geq 3$, with $x_0 = x_n$ (n-cycle).

EXERCISE 13. Show that if there is a walk from two vertices x and y of a graph Γ, then there is in Γ a simple path from x to y.

DEFINITION 2.4 (Connected Graphs). A graph $\Gamma = (V, E)$ is *connected* if for every $x, y \in V$ there exists a (possibly of length 0) walk from $x = x_0$ to $y = x_n$.

More in general, an equivalence relation is clearly defined on the set V of vertices of a graph Γ by saying that two vertices are related if there is a walk from one to the other. The subgraphs induced by the equivalence classes are the *connected components* of Γ. Clearly, every connected component is a maximal connected subgraph of Γ (and viceversa).
In a connected graph Γ a distance function is defined in V a natural way; given $x, y \in V$, the distance $d_\Gamma(x, y)$ is the smallest length of a path in Γ from x to y (included the trivial path of length 0 if $x = y$). This is indeed a true distance (for instance the triangular inequality follows from exercise 13) Then, if $\Gamma = (V, E)$ is connected the *diameter* of Γ is defined as

$$\mathrm{diam}(\Gamma) = \sup\{d_\Gamma(x, y) \mid x, y \in V\}.$$

For example, $\mathrm{diam}(G) = 1$ if and only if Γ is complete, $\mathrm{diam}(C_n) = [n/2]$, while, for the hypercubes, $\mathrm{diam}(Q_n) = n$.

EXERCISE 14. Show that a graph is 2-regular if and only if all of its connected components is a cycle.

EXERCISE 15. Let $k \geq 2$ and Γ a k-regular graph on n vertices. Find a function that bounds n in terms of $d = \mathrm{diam}(\Gamma)$. In particular, show that if Γ si cubic and $\mathrm{diam}(\Gamma) = d$ then $n \leq 3 \cdot 2^d - 2$.

DEFINITION 2.5. An *isomorphism* of the graph $\Gamma = (V, E)$ to the graph $\Gamma' = (V', E')$ is a bijective map $\alpha : V \to V'$ that preserves edges, i.e.

$$\forall x, y \in V : \ \{x, y\} \in E \Leftrightarrow \{\alpha(x), \alpha(y)\} \in E'.$$

An isomorphism of a graph to itself is an *automorphism*. It is immediate to verify that, under composition, the set $Aut(\Gamma)$ of all automorphisms of a graph Γ is a group.

DEFINITION 2.6. A graph $\Gamma = (V, E)$ is said to be *vertex–transitive* if $Aut(\Gamma)$ acts transitively on V.

Clearly, a vertex–transitive graph is regular (but not vice-versa).

EXERCISE 16. Show that all graphs in examples 10, 11 and 12 are vertex–transitive.

Trees. A *tree* is a connected graph with no cycles. The following collects a couple of largely used equivalent conditions (the easy proof is left as an exercise)

LEMMA 2.7. *Let $\Gamma = (V, E)$ be a connected graph. Then the following conditions are equivalent.*

(a) Γ *is a tree;*
(b) $|E| = |V| - 1$*;*
(c) *for every pair of distinct vertices $x, y \in V$ there is (one and) only one path from x to y.*

Bipartite graphs. A graph $\Gamma = (V, E)$ is *bipartite* if the set V of vertices may be partitioned in two non-empty subsets V_1 and V_2, such that every edge $e \in E$ has one extreme in V_1 and the other in V_2.

Thus, for example, a cycle C_n is bipartite if and only if n is even. Also, every hypercube (example 12) Q_n is bipartite; in fact a partition of the vertex set $\mathbb{Z}_2^n$, with the required property of bipartite graphs, is realized by the set V_0 of all n-tuples in $\mathbb{Z}_2$ that sum up to 0, and by the set V_1 of those n-tuples that sum up to 1 (the figure below shows the case of the true cube Q_3: red dots for vertices in V_1, green dots for V_2).

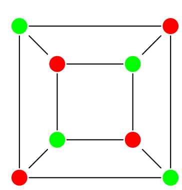

The following characterization of bipartite graphs is belongs to any first introduction to graphs.

PROPOSITION 2.8. *A graph is bipartite if and only if it does not have cycles of odd length.*

It turns out that bipartite graph are a rather distinguished class of graphs, and in many cases they have to be considered as somehow special items.

Cayley graphs. The concept of Cayley graph is the fundamental link between graphs and subset products in a group.

DEFINITION 2.9. Let G be a group and S a symmetric (that is $S^{-1} = S$) finite subset of G such that $1 \notin S$. The *Cayley graph* $\Gamma[G, S]$, is the graph whose vertex set is G, and whose edges are the 2-subsets

$$\{g, gx\} \quad \text{for} \quad g \in G, x \in S.$$

Observe that symmetry of S ensures that adjacency is a symmetric relation, while the condition $1 \notin S$ (this condition we will always tacitly assume when dealing with Cayley graphs) is introduced to avoid loops.

EXAMPLE 13. Let A_4 be the alternating group on 4 points. The Cayley graph $\Gamma[A_4, S]$ for $S = \{(123), (132), (12)(34)\}$

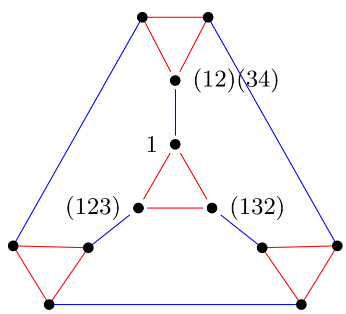

(123) and (132): red edges
(12)(34): blue edges.

EXAMPLE 14. Let $D = \langle x, y \mid x^2 = 1,\ y^x = y^{-1} \rangle$ (the infinite dihedral group); the Cayley graph $\Gamma[D, \{x, y\}]$ looks like

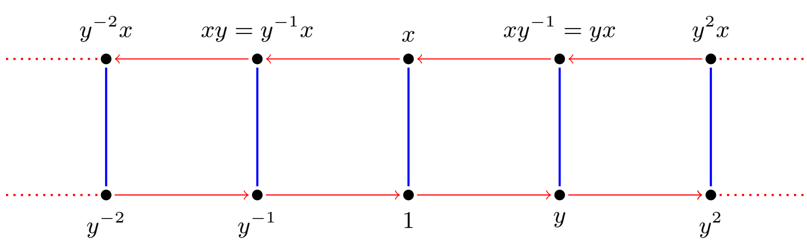

EXERCISE 17. Show that every hypercube (example 12) is a Cayley graph.

It is immediate to show that a Cayley graph $\Gamma[G, S]$ is connected if and only if S is a set of generators of G. Indeed, a walk of length n in $\Gamma[G, S]$, starting with $g \in G$, is just a sequence $g, gx_1, gx_1x_2, \ldots, gx_1x_2 \cdots x_n$, with $x_i \in S$; in other words, for $g, h \in G$ there is a walk of length n from g to h if and only if $h \in gS^n$. From this it follows that two elements of G are in the same connected component of $\Gamma[G, S]$

if and only if they belong to the same left coset modulo the subgroup $\langle S\rangle$ of G. Hence the number of connected components of $\Gamma[G,S]$ coincides with the index of $\langle S\rangle$ in G.

Even more important in our perspective is the fact that, if S is a symmetric set of generators of G, then

$$\mathrm{diam}(\Gamma) = \min\{d \geq 0 \mid S^d = G\}. \tag{2.2}$$

Clearly, because of the cancellation law, a Cayley graph $\Gamma[G,S]$ is k-regular, where $k = |S|$. Another relevant observation is that, for every g in G, left multiplication by g is an automorphism of every Cayley graph $\Gamma[G,S]$. In other words, the left-regular representation of G is also a faithful representation of G in the automorphims group of any Cayley graph on itself. Since this action is transitive on the vertex set G, we have

Proposition 2.10. *Every Cayley graph is vertex–transitive.*

Exercise 18. Let S be a symmetric set of generators of the group G; prove that the Cayley graph $\Gamma[G;S]$ is bipartite if and only if there exists a normal subgroup N of G such that $|G:N| = 2$ and $N \cap S = \emptyset$.

2.2. Adjacency matrix

Let $\Gamma = (V,E)$ be a graph with $|V| = n$. The *adjacency matrix* $A(\Gamma)$ is the $n \times n$ square matrix, whose entries are indexed on the vertices of the graph, and defined by the adjacency relation, i.e. for every $(x,y) \in V \times V$

$$A_{xy} = \begin{cases} 1 & if \quad x \sim y \\ 0 & if \quad x \not\sim y \end{cases}$$

(the actual matrix is obtained by fixing an enumeration $\{x_1,\ldots,x_n\}$ of V). Now, $A(\Gamma)$ is symmetric, real and has 0-entries on the diagonal. Hence $A(\Gamma)$ is diagonalizable, and its eigenvalues are all real. We agree in enumerating them in decreasing order, that is, counted with multiplicity,

$$\mu_0 \geq \mu_1 \geq \ldots \geq \mu_{n-1}$$

(occasionally, we may also specify the graph, $\mu_0(\Gamma), \mu_1(\Gamma)$, and so on). This the *spectrum* of $A(\Gamma)$, or by extension of language, the spectrum of Γ. Observe that, since $A(\Gamma)$ has trace 0,

$$\mu_0 + \mu_1 + \cdots + \mu_{n-1} = 0.$$

Example 15. The adjacency matrix of the complete graph K_n on n vertices, is $A(K_n) = J_n - I_n$, where J_n and I_n are, respectively, the $n \times n$ matrix with all entries 1 and the identity matrix of order n. Now, the eigenvalues of J_n are 0, with multiplicity $n-1$, and n, with multiplicity 1; hence the spectrum of K_n is $\mu_0 = n-1$, $\mu_1 = \ldots = \mu_{n-1} = -1$.

An effective way to look at the adjacency matrix $A(\Gamma)$ of a graph $\Gamma = (V, E)$ is to consider the *adjacency operator* associated to it, which we denote in the same way as $A(\Gamma)$. This is conveniently defined as a linear operator of the $\mathbb{C}$-space of functions

$$\ell^2(V) = \{f \mid f : V \to \mathbb{C}\}.$$

Namely (writing A for $A(\Gamma)$), for every $f \in \ell^2(V)$ and $x \in V$:

$$Af(x) = \sum_{y \in V} A_{xy} f(y) = \sum_{y \sim x} f(y). \tag{2.3}$$

Now, $\ell^2(V)$ is endowed with the standard hermitian product

$$\langle f, g \rangle = \sum_{x \in V} f(x)\overline{g(x)} \tag{2.4}$$

and the associated *norm*:

$$||f|| := ||f||_2 = \big(\sum_{x \in V} |f(x)|^2 \big)^{1/2}.$$

The operator $A = A(\Gamma)$ is thus hermitian (self-adjoint), i.e.

$$\langle Af, g \rangle = \langle f, Ag \rangle$$

for every $f, g \in \ell^2(V)$.

Write $A = A(\Gamma)$ the adjacency operator of a graph $\Gamma = (V, E)$, and let $\mathbf{1}_V$ be the constant map 1 on V; then for every $x \in V$,

$$A\mathbf{1}_V(x) = \sum_{y \sim x} 1 = d_\Gamma(x). \tag{2.5}$$

In particular, if Γ is k-regular, then $A\mathbf{1}_V = k\mathbf{1}_V$, that is k is an eigenvalue of Γ and $\mathbf{1}_V$ an eigenvector belonging to it. In fact, for regular graphs, we have the following.

Lemma 2.11. *For $k \geq 1$, let $\Gamma = (V, E)$ be a k-regular graph on n vertices, and let $\mu_0 \geq \mu_1 \geq \ldots \geq \mu_{n-1}$ be the spectrum of the adjacency matrix A of Γ. Then*

(1) *$\mu_0 = k$;*
(2) *$\mu_{n-1} \geq -k$;*
(3) *Γ is connected if and only if $k > \mu_1$;*
(4) *if Γ is connected, then $\mu_{n-1} = -k$ if and only if Γ is bipartite.*

Proof. We have already observed that k is an eigenvalue of A. Let μ be another eigenvalue, let $0 \neq f \in \ell^2(V)$ be an eigenvector relative to μ and $x \in V$ with $|f(x)|$ maximum. By replacing f with $\overline{f(x)}f$ we may suppose that $f(x)$ is a positive real number. Then

$$|\mu| f(x) = |Af(x)| = \big| \sum_{y \sim x} f(y) \big| \leq \sum_{y \sim x} |f(y)| \leq \sum_{y \sim x} f(x) = kf(x).$$

Hence $|\mu| \leq k$, proving points (1) and (2).
Let f be an eigenvector relative to k and $x \in V$ with $|f(x)|$ maximum. As before, we may suppose $\mathbb{R} \ni f(x) > 0$. Then

$$kf(x) = |Af(x)| = \Big|\sum_{y\sim x} f(y)\Big| \leq \sum_{y\sim x} |f(y)| \leq \sum_{y\sim x} f(x) = kf(x)$$

which forces $f(y) = f(x)$ for every y adjacent to x. This shows that f is constant on all vertices belonging to the same connected component containing x. If Γ is connected, $f = f(x)\mathbf{1}_V$ is a constant, and so $k = \mu_0$ has multiplicity 1, that is $\mu_1 < k$. Conversely, if V_1 and V_2 are the (non-empty) set of vertices of two distinct connected components of Γ, then $\mathbf{1}_{V_1}$ and $\mathbf{1}_{V_2}$ are linearly independent eigenvectors relative to k, and so $\mu_1 = k$. This proves (3).
To prove (4) assume Γ is bipartite and let $V = V_1 \cup V_2$ be a partition of V as in the definition of bipartire graph. Let $g = \mathbf{1}_{V_1} - \mathbf{1}_{V_2} \in \ell^2(V)$. If $x \in V_1$, all vertices adjacent to x are in V_2, and so

$$Ag(x) = \sum_{y\sim x} g(y) = -k = -kg(x);$$

similarly, if $x \in V_2$, $g(x) = -1$ and $Ag(x) = \sum_{y\sim x} g(y) = k = -kg(x)$. Therefore g is an eigenvector for the eigenvalue $-k$, and by point (2) we have $\mu_{n-1} = -k$.
Conversely, suppose that $\mu_{n-1} = -k$, let $0 \neq f \in \ell^2(V)$ be an eigenvector relative to it and, as before, $x \in V$ with $0 < f(x) \in \mathbb{R}$ the maximum among the modules of the images of f. Setting $V_1 = \{y \in V \mid f(y) = f(x)\}$ and $V_2 = \{y \in V \mid f(y) = -f(x)\}$, we prove that $V_1 \cup V_2$ is a graph partition of V . Clearly, $V_1 \cap V_2 = \emptyset$. Since Γ is connected, for every $y \in V$ we have a distance $d_\Gamma(x, y) = d_y$; by induction on d_y we show that $y \in V_1$ if d_y is even, while $y \in V_2$ if d_y is odd. In fact

$$|f(x)| = f(x) = -\frac{1}{k} Af(x) = -\frac{1}{k}\sum_{y\sim x} f(y) = \sum_{y\sim x} \frac{1}{k}(-f(y)),$$

hence $f(x)$ belongs to the inner of the convex hull of $\{-f(y) \mid y \sim x\}$; by the choice of x this is only possible if this hull reduces to a point: thus $f(y) = -f(x)$ and $y \in V_2$ for every $y \sim x$. This proves our claim for $d_y = 1$, and exactly the the same argument provides the inductive step. Therefore, Γ is bipartite. ■

Exercise 19. Let Γ be a k-regular graph.

(i) Prove that the multiplicity of k as eigenvalue of $A(\Gamma)$ equals the number of distinct connected components of Γ.
(ii) Prove that if Γ is bipartite then the spectrum of $A(\Gamma)$ is symmetric with respect to 0.

Exercise 20. Let C_n be the n-cycle ($n \geq 3$). For $t = 0, 1, \ldots, n-1$ let ζ_t be the n-th root of unity $\zeta_t = e^{\frac{2\pi i}{n}t}$. Prove that the eigenvalues of C_n are the real numbers

$$\zeta_t + \overline{\zeta_t} = 2\cos\frac{2\pi t}{n}$$

for $t = 0, \dots, n-1$. Thus, $\mu_0 = 2$ (with multiplicity 1), while all other eigenvalues have multiplicity 2, except for n even and $\mu_{n-1} = -2$, which has multiplicity 1.

EXERCISE 21. Let $A_n = A(Q_n)$ be the adjacency matrix of the hypercube Q_n. For $n \geq 2$, one may order the order of the vertices so that

$$A_{n+1} = \begin{pmatrix} A_n & I_n \\ I_n & A_n \end{pmatrix}.$$

Then, prove that the eigenvalues of $A(Q_n)$ are the integers $n - 2t$, with $t \in \mathbb{N}$ and $0 \leq t \leq n$, and that $n - 2t$ has multiplicity $\binom{n}{t}$.

Powers. The following useful fact is easily proved (e.g. by induction on m).

PROPOSITION 2.12. *Let $\Gamma = \Gamma = (V, E)$ be a graph, $m \geq 1$, and A^m the m-th power of its adjacency matrix (operator). Then, for every $x, y \in G$, A^m_{xy} coincides with the number of distinct walks of length m from x to y.*

PROOF. Exercise. ■

In particular, for each $x \in V$ we have $A^2_{xx} = d_\Gamma(x)$. Therefore, since the eigenvalues of A^m are the m-th powers of those of A,

$$\sum_{i=1}^{n-1} \mu_i(\Gamma)^2 = Tr(A^2) = \sum_{x \in V} d_\Gamma(x) = 2|E|$$

(where $n = |V|$). If Γ is k-regular we have $\sum_{i=1}^{n-1} \mu_i^2 = kn$.

The Laplace operator. If $\Gamma = (V, E)$ is a connected k-regular graph and $A = A(\Gamma)$, then the set of eigenvectors relative to $\mu_0 = k$ is the set of constant functions that we denote by $\mathcal{Z}$. Thus $\ell^2(V) = \mathcal{Z} \oplus \mathcal{Z}^\perp$, where

$$f \in \mathcal{Z}^\perp \;\Leftrightarrow\; \sum_{x \in V} f(x) = 0.$$

In particular we have that all eigenvectors relative to eigenvalues μ_i for $i > 0$ are 0-sum functions.

An *orientation* of the graph $\Gamma = (V, E)$ is just a total ordering on V. Any such orientation determines for every edge an intial and a final point, if $e \in E$ then $e = \{e_-, e_+\}$ with $e_- < e_+$.

Let $|E| = m$; given an orientation of Γ we define the *incidence matrix* M as the $n \times m$ matrix indexed on $V \times E$:

$$M_{xe} = \begin{cases} 1 & \text{if } x = e_+ \\ -1 & \text{if } x = e_- \\ 0 & \text{if } x \notin e \end{cases} \qquad \text{for } \quad (\mathrm{x}, \mathrm{e}) \in \mathrm{V} \times \mathrm{E}.$$

M is associated to an operator $\delta : \ell^2(E) \to \ell^2(V)$,

$$(\delta u)(x) = \sum_{x = e_+} u(e) - \sum_{x = e_-} u(e) \tag{2.6}$$

for $u \in \ell^2(E) := \{u \mid u : E \to \mathbb{C}\}$, and $x \in V$.
Its dual operator $\delta^* : \ell^2(V) \to \ell^2(E)$ is given by, for all $f \in \ell^2(V)$ and $e \in E$,

$$(\delta^* f)(e) = f(e_+) - f(e_-). \tag{2.7}$$

Then, for every $f \in \ell^2(V),\ u \in \ell^2(E)$, we have

$$\langle \delta^* f, u \rangle_E = \langle f, \delta u \rangle_V \tag{2.8}$$

(where the scalar products are those of the appropriate space). The *Laplace operator* associated to the graph Γ is

$$\delta\delta^* : \ell^2(V) \to \ell^2(V).$$

A simple computation shows that it is independent on the orientation, and that its matrix form is

$$L(\Gamma) = kI_n - A(\Gamma).$$

The Laplace operator $L = L(\Gamma)$ is real symmetric, has the same eigenspaces of A and its eigenvalues are

$$k - \mu_{n-1} \geq \cdots \geq k - \mu_1 \geq k - \mu_0 = 0.$$

If Γ is connected then $\ker(L)$, the eigenspace relative to $k - \mu_0 = 0$, is the set $\mathcal{Z}$ of all constant functions. While if $f \in \mathcal{Z}^\perp$ (the set of zero-sum functions), we have the standard Rayleigh bound:

$$\frac{\langle Lf, f \rangle}{\langle f, f \rangle} \in [k - \mu_1, k - \mu_{n-1}] \tag{2.9}$$

One way of using effectively the Laplace operator $L = L(\Gamma)$ is to apply the Rayleigh bound in connection to the nature of L, for in fact for every $f \in \ell^2(V)$,

$$\langle Lf, f \rangle = \langle \delta\delta^* f, f \rangle = \langle \delta^* f, \delta^* f \rangle_E$$

and then use the particularly simple description (2.7) of δ^*. This will be the case in the proof of Theorem 2.14; another application of this method is suggested in exercises 22 and 23 below.

2.3. Expansion in graphs

Definition 2.13. Let $\Gamma = (V, E)$ be a graph; for a subset $\emptyset \neq U \subseteq V$, the *boundary* ∂U of U is the set of all edges of Γ that have one extreme in U and the other in $V \setminus U$, that is

$$\partial(U) = \{e \in E \mid |e \cap U| = 1\}.$$

If Γ is finite, the *expanding* (or *isoperimetrical* or *Cheeger*) constant of Γ is

$$h(\Gamma) = \min\left\{ \frac{|\partial U|}{|U|} \middle| U \subseteq V,\ 0 < |U| < |V|/2 \right\}.$$

(There is also a definition that applies to infinite graphs, but we restrict to the finite case).

EXAMPLE 16. The first two items that follow are almost trivial; for (3) you may try to prove directly the claim, otherwise wait until after the next Theorem.

(1) For the complete graph K_n we have $h(K_n) = n - \left[\frac{n}{2}\right]$;
(2) if C_n is the n-cycle, then

$$h(C_n) = \frac{2}{[n/2]};$$

(3) if Q_n is the n-hypercube, then $h(Q_n) = 1$

THEOREM 2.14. *Let Γ be a k-regular, connected graph on n vertices, e let μ_1 be the second eigenvalue of the adjacency matrix of Γ. Then.*

$$\frac{k-\mu_1}{2} \leq h(\Gamma) \leq \sqrt{2k(k-\mu_1)}.$$

PROOF. We prove only the left inequality, which is the one that we need in the rest of these notes.
Let $U \subseteq V$, with $1 \leq |U| \leq |V|/2$. Fix an orientation on Γ such that $x < y$ for all $x \in U$, $y \in V \setminus U$, and consider $f \in \ell^2(V)$ defined by

$$f(x) = \begin{cases} |V| - |U| & \text{if } x \in U \\ -|U| & \text{if } x \in V \setminus U \end{cases}$$

Then $f \in \mathcal{Z}^{\perp}$, and so, by (2.9),

$$\frac{\langle Lf, f\rangle}{\langle f, f\rangle} \geq k - \mu_1. \tag{2.10}$$

Now, we have

$$\langle f, f\rangle = \sum_{x\in V} f(x)^2 = |U||V|(|V| - |U|),$$

while

$$\langle Lf, f\rangle = \langle \delta^* f, \delta^* f\rangle = \sum_{e\in E}(f(e_+) - f(e_-))^2 = \sum_{e\in\partial U} |V|^2 = |\partial U||V|^2$$

Therefore, substituting in (2.10),

$$k - \mu_1 \leq \frac{|\partial U||V|^2}{2|U||V|(|V| - |U|)} = \frac{|\partial U||V|}{|U|(|V| - |U|)}$$

and so

$$\frac{|\partial U|}{|U|} \geq \frac{(|V| - |U|)}{|V|}(k - \mu_1) \geq \frac{k-\mu_1}{2}.$$

The claim $h(\Gamma) \leq (k - \mu_1)/2$ follows by definition. ■

EXAMPLE 17. Consider the n-hypercube Q_n. Then $\mu_1 = n - 2$ by exercise 21, hence

$$h(Q_n) \geq \frac{n-(n-2)}{2} = 1$$

by Theorem 2.14. On the other hand, let U be the set of all n-tuples in $\mathbb{Z}_2^n = V$ whose first coordinate is 0; then $|U| = 2^{n-1} = |\partial U|$, and so $h(Q_n) \leq |\partial U|/|U| = 1$. Thus, $h(Q_n) = 1$.

EXERCISE 22. Let $c \geq 1$; a graph $\Gamma = (V, E)$ is c-colorable if there exists a function $\gamma : V \to \{1, 2, \ldots c\}$ such that for every $x, y \in V$,

$$x \sim y \;\Rightarrow\; \gamma(x) \neq \gamma(y).$$

The *chromatic number* $\chi(\Gamma)$ of Γ is the smallest positive integer c such that Γ is c-colorable. Also, we denote by $\alpha(\Gamma)$ the largest cardinality of a *stable* subset of V (a subset $U \subseteq V$ is said to be stable (or *independent*) if no pair of distinct elements of U are adjacent in Γ). Show that $\chi(\Gamma)\alpha(\Gamma) \geq |V|$.

EXERCISE 23. Let $\Gamma = (V, E)$ be k-regular with $|V| = n$. Show that

$$\chi(\Gamma) \geq 1 - k/\mu_{n-1}.$$

[Hint: for a stable set U of Γ with $|U| = \alpha = \alpha(\Gamma)$, consider $f \in \ell^2(V)$ given by

$$f(x) = \begin{cases} n-\alpha & \text{if } x \in U \\ -\alpha & \text{if } x \in V \setminus U \end{cases}$$

Show that $\langle Lf, f\rangle = n^2|\partial U| = n^2 k|U|$, then apply Rayleigh bound and exercise 22.]

EXERCISE 24. For every $n \geq 2$ the complete bipartite graph $K_{n,n}$ is the graph whose set of vertices $V = A \cup B$ is the disjoint union of two sets of order n, and every element of A is adjacent to all the elements of B, while there are no other edges. Determine the spectrum of the adjacency matrix of $K_{n,n}$, and the expansion constant $h(K_{n,n})$.

DEFINITION 2.15. Let $\varepsilon > 0$; a graph Γ such that $h(\Gamma) > \varepsilon$ is said to be an *ε-expander.*

Of course, every finite connected graph is a ε-expander for some $\varepsilon > 0$, so this definition seems to be rather useless. The point is that we look for large (in number of vertices) graphs which are ε-expanders for some constant value ε, and the problem arises when we are interested in graph that are sparse, in the generic sense that they have few edges. For instance, by bounding by a fixed $k \geq 1$ the ratio $|E|/|V|$ of the graph $\Gamma = (V, E)$, we may wonder whether $h(\Gamma)$ tends to 0 when $|V|$ goes to infinity, or not.

Families of expanders. The search for graphs that are at the same time sparse and highly connected lead to the following basic definition.

DEFINITION 2.16 (Bassalygo and Pinsker, 1972). An infinite family of finite graphs $\Gamma_n = (V_n, E_n)$, $1 \leq n \in \mathbb{N}$, is called a *family of expanders* if there exist an integer $k \geq 3$ and a $\varepsilon > 0$ such that

- Γ_n is k-regular for every $n \geq 1$;
- $\lim_{n\to\infty} |V_n| = \infty$;
- $h(\Gamma_n) \geq \varepsilon$ for every $n \geq 1$.

In view of Theorem 2.14, we have the following immediate algebraic equivalent formulation (which we adopt as definition)

DEFINITION 2.17. Let $k \geq 3$. An infinite family of finite connected k-regular graphs $\Gamma_n = (V_n, E_n)$, $1 \leq n \in \mathbb{N}$, with $|V_n| \to \infty$, is a family of expanders if and only there exists $\varepsilon > 0$ such that, for every $n \geq 1$,

$$k - \mu_1(\Gamma) \geq \varepsilon$$

where $\mu_1(\Gamma_n)$ is the largest eigenvalue of $A(\Gamma_n)$ different from k.

NOTE 1. Since all the graphs in the previous definitions are connected, for any finite subset of them there exists a non-zero lower bound for their expansion constants. Thus, the requirement that $h(\Gamma_n) \geq \varepsilon$ (or $k - \mu_1(\Gamma) \geq \varepsilon$) holds for every $n \geq 1$ may be replaced (as it will be often the case in actual proofs) by asking that there exists $n_0 \geq 1$ such that the bound applies for every $n \geq n_0$.

The definition we gave of ε-expander, and of families of expanders, is the more immediate and suggested by a neat graph-theoretical condition. However, from the point of view of algebraic manipulation it turns out that a slightly stronger condition seems to be more natural. Thus, given a k-regular (connected) graph Γ, we denote by $\boldsymbol{\mu} = \boldsymbol{\mu}(\Gamma)$ the largest *absolute value* of an eigenvalue $\neq k$ of $A(\Gamma)$, that is, if Γ has n vertices,

$$\boldsymbol{\mu}(\Gamma) = \max\{|\mu_i(\Gamma)| \mid i = 1, \dots n-1\}.$$

For $\varepsilon > 0$, we say that a k-regular graph Γ is an *algebraic* (or, following [**31**], a *two-sided*) ε-expander if

$$k - \boldsymbol{\mu} > \varepsilon. \tag{2.11}$$

and give a corresponding definition of a family of two-sided expanders. Observe that, by Lemma 2.11, bipartite graphs are never two-sided expanders.

Existence of families of expanders for every $k \geq 3$ is relatively easy to establish. Soon after the definition, by using a probabilistic approach, Pinsker proved the following result.

THEOREM 2.18 (Pinsker, 1973). *There exists $\delta > 0$ such that for every $n \geq 2$, and every $k \geq 3$ there exists a k-regular graph Γ on n vertices with $h(\Gamma) \geq \delta$.*

For a proof, see [**18**] or [**31**]. In fact, the following stronger result may be be proved along similar methods.

THEOREM 2.19. *Let $k \geq 3$. Then there exists $\delta > 0$ such that the probability that for a k-regular graph $\Gamma = (V, E)$ one has $h(\Gamma) \geq \delta$ tends to 1 as $|V| \to \infty$.*

As often in mathematics, more difficult is to provide explicit constructions of families of expanders. The first such example was due to Margulis (example 19) in 1973: it already exploited certain properties of the Cayley graph of groups. By now there are several methods available, but many rely on Cayley graphs and on group theoretical properties (and some on deep number theoretical results); one of these, developed by Bourgain and Gamburd, has many connections with set-products and is the one that we will treat in the next chapter and apply in the following one.

It deserves to be at least mentioned a notable recursive construction, the so called *Zig-Zag product* due to Reingold, Vadhan and Wigderson [**25**], that in principle is entirely graph-theoretical, and elementary, but we will not say anything else about it (recommending [**18**] to the interested reader).

Another issue that we will not touch is the question of optimal spectral gaps $k - \boldsymbol{\mu}$. In this respect, the basic result is the following.

PROPOSITION 2.20 (Alon and Boppana). *Let $k \geq 2$. If Γ is a k-regular connected graph with n vertices then*

$$\boldsymbol{\mu}(\Gamma) \geq 2\sqrt{k-1} - o_n(1),$$

where $o_n(1)$ is a positive quantity that, for a fixed k, tends to zero as $n \to \infty$.

A nice proof of this may be found in [**18**]. A graph Γ is called a *k-Ramanujan graph* if it is k-regular and

$$\boldsymbol{\mu}(\Gamma) \leq 2\sqrt{k-1}.$$

The existence of k-Ramanujan graphs of arbitrary large order has been established only for $k = q + 1$, when q a power of a prime. However, Marcus, Spielman and Srivastava [**21**] have recently proved that if, in the definition of Ramanujan graphs, one replaces $\boldsymbol{\mu}$ with the first non-trivial eigenvector μ_1 (i.e. if one looks for one-sided expansion), then for every $k \geq 3$ there exist bipartite graphs Γ of arbitrary large order such that $\mu_1(\Gamma) \leq 2\sqrt{k-1}$.

EXERCISE 25. Let Γ be a k-regular graph on n vertices. Prove that, for a fixed k,

$$\boldsymbol{\mu}(\Gamma) \geq \sqrt{k} - o_n(1)$$

where $o_n(1)$ is a positive quantity that goes to zero as $n \to \infty$. [hint: use the second power of the adjacency operator.]

EXERCISE 26. To any given graph $\Gamma = (V, E)$ we associate a bipartite graph $\mathcal{D}(\Gamma) = (V_1 \cup V_2, E^+)$ in the following way: $V_1 \cup V_2$ is the disjoint union of two copies of V, i.e. there are bijections $\alpha_1 : V_1 \to V$, $\alpha_2 : V_2 \to V$, and for $u_1 \in V_1,\ u_2 \in V_2$ let,

$$\{v_1, v_2\} \in E^+ \ \Leftrightarrow\ \{\alpha_1(u_1), \alpha_2(u_2)\} \in E.$$

Prove that $h(\mathcal{D}(\Gamma)) \geq h(\Gamma)/2$; than show that an infinite family $(\Gamma_n)_{n\geq 1}$ of graphs is a family of two-sided expanders if and only if $(\mathcal{D}(\Gamma_n))_{n\geq 1}$ is a family of one-sided expanders.

2.4. Expansion and random walks

We now give a brief and very basic outline of one of the important features of expanders graphs, which is related to probability, and is fundamental in many applications like those to derandomization and to the analysis of "Monte–Carlo" type algorithms. By doing this, we also introduce a perspective that will turn out to be almost indispensable to properly understand the ideas behind the method to establish expansion properties of Cayley graph treated in the next chapter.

We start with a simple but useful Lemma; if $\Gamma = (V, E)$ is a graph and X, Y are subset of V, we denote by $e(X, Y)$ the set of all adjacent pairs (x, y) with $x \in X$, $y \in Y$.

LEMMA 2.21 (Expander Mixing Lemma). *Let $\Gamma = (V, E)$ be a k-regular graph on n vertices. Then, for every $X, Y \subseteq V$,*

$$\left| |e(X,Y)| - \frac{k|X||Y|}{n} \right| \leq \boldsymbol{\mu}\sqrt{|X||Y|}. \tag{2.12}$$

PROOF. Let $A = A(\Gamma)$; then

$$\langle \mathbf{1}_X, A\mathbf{1}_Y \rangle = |e(X,Y)|. \tag{2.13}$$

Consider the decompositions $\mathbf{1}_X = c_X + f$, $\mathbf{1}_Y = c_Y + g$, con c_X, c_Y costanti (i.e. elements of $\mathcal{Z}$) and $f, g \in \mathcal{Z}^\perp$; then $c_X = |X|/n$ and $c_Y = |Y|/n$. Substituting in (2.13), because of orthogonaity, we have

$$|e(X,Y)| = \langle c_X, Ac_Y \rangle + \langle f, Ag \rangle = k\frac{|X||Y|}{n} + \langle f, Ag \rangle.$$

Hence,

$$\left| |e(X,Y)| - \frac{k|X||Y|}{n} \right| = |\langle f, Ag \rangle|. \tag{2.14}$$

Now, writing $f = \sum_{i=1}^{n-1} \alpha_i v_i$ and $g = \sum_{i=1}^{n-1} \beta_i v_i$ as linear combinations in an orthonormal basis $\{v_1, \ldots, v_{n-1}\}$ of $\mathcal{Z}^\perp$ (made of eigenvectors belonging to the eigenvalues $\mu_1, \ldots, \mu_{n-1}$), we have

$$|\langle f, Ag \rangle| = |\sum_{i=1}^{n-1} \mu_i \alpha_i \beta_i| \leq \sum_{i=1}^{n-1} |\mu_i||\alpha_i\beta_i| \leq \boldsymbol{\mu} \sum_{i=1}^{n-1} |\alpha_i\beta_i|.$$

Thus, by the Cauchy–Schwartz inequality,

$$|\langle f, Ag\rangle| \leq \boldsymbol{\mu} \cdot ||f|| \cdot ||g|| = \boldsymbol{\mu} \cdot ||\mathbf{1}_X|| \cdot ||\mathbf{1}_Y|| = \boldsymbol{\mu}\sqrt{|X||Y|}$$

and, by (2.14), the statement. ∎

Observe that the quantity $\frac{k|X||Y|}{n}$ represents the expected number of adjacent pairs between X and Y in a random graph of edge-density k; therefore, the left hand side in (2.12) measures the discrepancy between the number $|e(X,Y)|$ in the actual graph Γ and its expected value in the random case. The Mixing Lemma tells us that the smaller is $\boldsymbol{\mu}(\Gamma)$ (hence the larger the spectral gap $k - \boldsymbol{\mu}$) the closer the k-regular graph Γ approximates a random behavior.

This observation, admittedly generic, is reinforced by the analysis of *random walks* in a regular graph.

DEFINITION 2.22. Let $\Gamma = (V, E)$ be a graph. A *probability distribution* on Γ is a function $\pi : V \to [0, 1]$ such that $\sum_{x \in V} \pi(x) = 1$ (observe that $\pi \in \ell^2(V)$).

A *random walk* in a graph Γ is a walk $v_0, v_1, \ldots v_i, \ldots$ in Γ such that, independently for every $i \geq 0$, v_{i+1} is randomly selected with uniform probability $1/d_\Gamma(v_i)$ among the vertices that are adjacent to v_i.
The random walk process starts from a given *initial probability distribution* π_0 on Γ: for $x \in V$, $\pi_0(x) = P(x = v_0)$ is then the probability that a walk of length 0 starts (and ends) in x. Now, π_1 is defined as the probability distribution so that $\pi_1(x)$ is the probability that $x = v_1$ is in a random walk where the initial point v_0 is subject to the probability π_0; thus $\pi_1(x)$ is non-zero if and only if $x \sim y$ with $\pi_0(y) \neq 0$. Continuing this way, one produces a sequence of probability distributions π_i such that, for $x \in V$, $\pi_i(x) = P(x = v_i)$ in a random walk $v_0, v_1, \ldots v_i, \ldots$ of Γ.
It is a fundamental fact that if Γ is connected and not bipartite, the sequence $\{\pi_i\}_{i \in \mathbb{N}}$ converges in $\ell^2(V)$ to a stationary distribution, and it is not difficult to show that if moreover Γ is regular, then this limit is always the uniform distribution on V (that is $\frac{1}{|V|}\mathbf{1}_V$).
Let $\Gamma = (V, E)$ be a k-regular graph and write $\widehat{A} = \frac{1}{k}A(\Gamma)$. Fixed an initial probability distribution π_0 on V, for $x \in V$ we have

$$\pi_1(x) = \sum_{y \sim x} \frac{1}{k}\pi_0(y) = \widehat{A}\pi_0(x),$$

hence $\pi_1 = \widehat{A}\pi_0$. In the same way, for every $n \geq 0$,

$$\pi_{i+1} = \widehat{A}\pi_i = \widehat{A}^{i+1}\pi_0. \tag{2.15}$$

($\widehat{A}$ is called the Markov transition matrix of the process.)
Let $|V| = n$ and let $\mathbf{u} = \frac{1}{n}\mathbf{1}_V$ be the uniform probability distribution on V. Let $i \geq 1$; since π_i and $\mathbf{u}$ are probability measures, their difference $\pi_i - \mathbf{u}$ is a zero-sum

function on V, that is $\pi_i - \mathbf{u} \in \mathcal{Z}^\perp$. Observe that, being $\mathbf{u}$ a constant, $\widehat{A}\mathbf{u} = \mathbf{u}$. Therefore

$$\langle \pi_{i+1} - \mathbf{u}, \pi_{i+1} - \mathbf{u} \rangle = \langle \widehat{A}(\pi_i - \mathbf{u}), \widehat{A}(\pi_i - \mathbf{u}) \rangle = \langle \widehat{A}^2(\pi_i - \mathbf{u}), \pi_i - \mathbf{u} \rangle.$$

Now, $\boldsymbol{\mu}/k$ is the largest absolute value of an eigenvalue $\neq 1$ of $\widehat{A}$, and so

$$||\pi_{i+1} - \mathbf{u}||^2 = \langle \widehat{A}^2(\pi_i - \mathbf{u}), \pi_i - \mathbf{u} \rangle \leq \boldsymbol{\mu}^2/k^2 ||\pi_i - \mathbf{u}||^2$$

That is $||\pi_{i+1} - \mathbf{u}|| \leq (\boldsymbol{\mu}/k)||\pi_i - \mathbf{u}||$; and, by an immediate induction

$$||\pi_i - \mathbf{u}|| \leq \left(\frac{\boldsymbol{\mu}}{k}\right)^i \cdot ||\pi_0 - \mathbf{u}|| \leq \left(\frac{\boldsymbol{\mu}}{k}\right)^i, \tag{2.16}$$

for every $i \geq 0$.

Thus, the smaller is $\boldsymbol{\mu}$ with respect to k, the more rapidly $\{\pi_i\}_{i\in\mathbb{N}}$ converges to the uniform distribution on V. Let us be more precise. If, for $\varepsilon > 0$, $\Gamma = (V, E)$ is a two-sided ε-expander with $|V| = n$, then

$$\frac{\boldsymbol{\mu}}{k} \leq \frac{k - \varepsilon}{k} =: \alpha.$$

Let $C = -1/\log\alpha$ (observe that, since $\alpha < 1$, C is positive); then, from (2.16), for every integer $m \geq C\log n$ and every initial distribution π_0,

$$||\pi_m - \mathbf{u}|| \leq \alpha^m \leq \alpha^{C\log n} = n^{-1}.$$

When applied to a family of two sided expanders this is interesting enough to state as a Theorem.

Theorem 2.23. *Let $\Gamma_n = (V_n, E_n)$ be a family of two-sided expanders. Then there is a constant C such that for every $n \geq 1$*

$$||\pi_m - \frac{1}{|V_n|}\mathbf{1}_{V_n}|| \leq |V_n|^{-1},$$

for every initial distribution π_0 and any integer $m \geq C\log|V_n|$.

This is commonly referred to by saying that families of expanders have *logarithmic mixing time.*

Note 1. We have considered the convergence in the ℓ_2 norm, but similar considerations apply for other norms; the most relevant in this contest being the ℓ_1 and ℓ_∞ norms, defined, respectively, by

$$||f||_1 = \sum_{x\in V}|f(x)|, \quad ||f||_\infty = \max_{x\in V}|f(x)|.$$

Note 2. Theorem 2.23 is not true, as it stands, for families of *one-sided* expanders, because among any of them there might well be bipartite graphs (see Exercise 26), for which the random process does not even converge. In fact, let $\Gamma = (V_1 \cup V_2, E)$ be a bipartite graph, fix a vertex $v \in V_1$ and the initial distribution π_0 be the one concentrated in v; then, random walks starting at v alternatively jump from V_1 to

V_2 and back; thus, for every $i \geq 0$, $\pi_{2i}(x) = 0$ for all $x \in V_2$ while $\pi_{2i+1}(x) = 0$ for every $x \in V_1$, and the sequence $(\pi_i)_{i\geq 0}$ cannot possibly converge in $\ell^2(V)$ (in any of the three norms we have mentioned). Things may however be adjusted by adopting a different definition of random walks in bipartite graphs (see [**18**] or [**31**] for details).

Diameter of expanders. For $k \geq 3$, let $\Gamma_n = (V_n, E_n)$ be a family of connected k-regular graphs satisfying the conclusion of Proposition 2.23 for some constant C independent on n. Given $n \geq 1$, fix a vertex $v \in V_n$ and let π_0 the probability distribution on V_n concentrated in v, that is, for every $x \in V_n$,

$$\pi_0(x) = \begin{cases} 1 & \text{if } x = v \\ 0 & \text{if } x \neq v \end{cases}$$

Let $m = m(n)$ be the smallest integer greater or equal than $C \log |V_n|$, and let $\Delta = \{x \in V_n \mid \pi_m(x) = 0\}$. Then, writing $\mathbf{u} = \frac{1}{|V_n|}\mathbf{1}_{V_n}$,

$$|V_n|^{-2} \geq ||\pi_m - \mathbf{u}||^2 = \sum_{x\in V_n} (\pi_m(x) - |V_n|^{-1})^2 > \sum_{x\in\Delta} |V_n|^{-2} = |\Delta||V_n|^{-2}.$$

Thus $|\Delta| = 0$ and $\Delta = \emptyset$. This means that every element of V_n may be reached from v by a walk of length m. In particular, $d_{\Gamma_n}(v, x) \leq m$ for every $x \in V_n$, and this implies

$$diam(\Gamma_n) \leq 2m \leq D \log |V_n|$$

where $D = 2(C + 1)$ is a constant indepenent on n.

By this conclusion, one says that the family of graphs $\Gamma_n = (V_n, E_n)$ (or sometimes, for convenience, each of the graphs in the family) has *logarithmic diameter*. In conjunction with Theorem 2.23 we have,

THEOREM 2.24. *A family of expanders has logarithmic diameter.*

This has important algebraic consequences when applied to families of expander Cayley graphs.

EXAMPLE 18. For $n \geq 2$, define the square grid graph X_n as the graph that has $V_n = \mathbb{Z}/n\mathbb{Z} \times \mathbb{Z}/n\mathbb{Z}$ as set of vertices, and each vertex (a, b) is adjacent to all and only the vertices $(a \pm 1, b)$ and $(a, b \pm 1)$, where the sum is modulo n (notice that X_n is a Cayley graph). Thus, for every $n \geq 2$, X_n is a connected 4-regular graph. Let $F = \{(a, b) \in V_n \mid 0 \leq b \leq \frac{n-1}{2}\}$, then the edges in the boundary ∂F have one extreme in the two limit rows of F and for each vertex in these rows there is exactly one edge of ∂F incident to it, so $|\partial F| = 2n$. Hence, since $|F| = [n/2]n \leq n^2/2 = |V_n|/2$,

$$h(X_n) \leq \frac{|\partial F|}{|F|} = \frac{2n}{[n/2]^2} = \frac{4}{n-1}.$$

In particular, $\lim_{n\to\infty} |h(X_n)| = 0$, and so $(X_n)_{n\geq 3}$ is not a family of expanders.

We could have reached the same conclusion by invoking Theorem 2.24; in fact

$$diam(X_n) = n \neq O(\log n).$$

EXERCISE 27. Let $\Gamma[G_n, S_n]$ $(n \geq 1)$ be a family of k-regular Cayley graphs for some fixed $k \geq 3$, and suppose that all the groups G_n are abelian. Prove that $\Gamma[G_n, S_n]$ is not a family of expanders.

2.5. Cayley graphs as expanders

Possibly because the richness of tools that may be applied to study them, that of Cayley graphs soon appeared to be an elected area where to look for families of expanders. The graphs in first production by G. Margulis of such a family, described below, are not properly Cayley graphs, but belong to a generalization of them, called Schreier graphs (see example 20 for the definition); also, the expanding property of Margulis' example graphs is strictly related to the expanding property of a class of Cayley graphs.

EXAMPLE 19 (G. Margulis, 1973). For $m \geq 3$, M_m is the graph on the set of vertices

$$V_m = \mathbb{Z}/m\mathbb{Z} \times \mathbb{Z}/m\mathbb{Z}.$$

To define edges, let

$$e_1 = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad e_2 = \begin{pmatrix} 0 \\ 1 \end{pmatrix},$$

and consider the transformations of V_m given by

$$S = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}, \quad T = \begin{pmatrix} 1 & 0 \\ 2 & 1 \end{pmatrix}.$$

Edges are defined by setting that, for every $v \in V_m$, v is adjacent to the four vertices Sv, Tv, $Sv + e_1$, $Tv + e_2$ and to the other four vertices defined by the four inverse transformations. In this way an 8-regular graph M_m is obtained, of order m^2, for which it may be proved that

$$h(M_m) \geq \frac{8 - 5\sqrt{2}}{2} > \frac{4}{9}.$$

Hence, $(M_m)_{m \geq 3}$ is a family of expanders.

NOTE. M_m is not exactly a simple graph, as it has loops and multiple edges, but these are in number which is linear in m and do not affect in a significant way expansion.

The following Theorem, which is absolutely beyond the scope of this course, is the achievement of a research whose many steps where contributed by several authors; among those, G. Margulis, A. Lubotzky, M. Kassabov, N. Nikolov, P. Sarnak, E. Breuillard, B. Green, T. Tao, L. Pyber, E. Szab and others.

THEOREM 2.25. *There exist $k \in \mathbb{N}$ and $\varepsilon > 0$ such that, for every non-abelian finite simple group G, a symmetric set S of k generators may be selected such that the Cayley graph $\Gamma[G,S]$ is a two-sided ε-expander.*

For those who know something about the classification of finite simple groups, I will add that, as often, there is a certain difference between groups of Lie type and the alternating groups (the sporadics, being only a finite number, are not relevant in this contest). In fact, it has been proved that, for groups belonging to any given family of Lie type simple groups, most homogeneous choices (in a sense that for the moment I will not define precisely) of finite symmetric sets of generators produce Cayley graphs that form a family of expanders. The aim of the rest of this course is to provide a proof, as self-contained as feasible, for the case of the family of groups $SL_2(p)$, for p a prime number.

Expansion in a Cayley graph is immediately related to set-products, i.e. the theme of first Chapter. In fact, let G be a finite group, S a symmetric subset of G and $\Gamma = \Gamma[G,S]$ the associated Cayley graph; then, for every $F \subset G$, the set of vertices that are adjacent to some element of F is the product FS. Thus, the cardinality of the boundary of F in Γ is just $|FS - F|$, and so

$$h(\Gamma) = \min\{|FS - F|/|F| \mid F \subseteq G, 1 \leq |F| \leq |G|/2\}.$$

A straightforward argument yields that a family of Cayley graphs $(\Gamma[G_n, S_n]))_{n\geq 1}$, with $|S_n| = k$ for every $n \geq 1$, is a family of expanders if and only $|G_n| \to \infty$, as $n \to \infty$, and there exists $\varepsilon > 0$ such that

$$|F(S_n \cup \{1\})| \geq (1+\varepsilon)|F|,$$

for every $n \geq 1$ and every $\emptyset \neq F \subseteq G_n$ with $|F| \leq |G_n|/2$.

Convolutions. The fact that, in a Cayley graph $\Gamma[G,S]$, the set of vertices G is a group allows to endow the space $\ell^2(G)$ with the *convolution product* $*$. We have already introduced it in the case of abelian groups in Section 1.3; the definition for an arbitrary finite group G is the same, namely, for every $f, g \in \ell^2(G)$ and $x \in G$,

$$(f * g)(x) = \sum_{y \in G} f(xy^{-1})g(y) = \sum_{y \in G} f(y)g(y^{-1}x).$$

This product is associative but it is not (unless G is abelian) commutative. On the other hand it is bilinear, in the sense that, for every $f, g, h \in \ell^2(G)$, $a, b \in \mathbb{C}$,

$$(af + bg) * h = a(f * h) + b(g * h), \tag{2.17}$$

and the other way around.
For $f \in \ell^2(G)$ and $m \geq 1$ we write

$$f^{(m)} = f * f * \cdots * f \quad (m \text{ times}).$$

A function $f \in \ell^2(G)$ is *symmetrical* if $f(x^{-1}) = f(x)$ for all $x \in G$. Having said that, here is a number of other immediate properties of the convolution product: let $f, g \in \ell^2(G)$,

(1) if f is a constant then $f * g$ and $g * f$ are constants;
(2) if f is a sum-zero function (i.e. $f \in \mathcal{Z}^\perp$) then such are $f * g$ and $g * f$;
(3) if f is symmetrical then $f * f$ is symmetrical;
(4) if f, g are probability distributions then such is $f * g$.

We also collect some properties of the convolution with respect to norms in $\ell^2(G)$.

LEMMA 2.26 (Cauchy-Schwartz and Young's inequalities). *Let $f, g \in \ell^2(G)$; then*
(i) *(Cauchy-Schwarz)* $||f||_1 \leq |G|^{1/2}||f||$;
(ii) $||f * g|| \leq ||f||_1 \cdot ||g|| \leq |G|^{1/2}||f|| \cdot ||g||$;
(iii) $||f * g||_\infty \leq ||f||_\infty ||g||_1$.

PROOF. (i) This is an immediate application of the standard Cuchy-Schwarz inequality; in fact,

$$||f||_1^2 = \big(\sum_{x \in G} |f(x)|\big)^2 \leq |G| \sum_{x \in G} |f(x)|^2 = |G| ||f||^2.$$

(ii) Writing, for $x, y \in G$, $|f(xy^{-1})||g(y)| = |f(xy^{-1})|^{1/2}|f(xy^{-1})|^{1/2}|g(y)|$, we have, by Cauchy-Schwarz inequality,

$$|f * g(x)|^2 \leq \big(\sum_{y \in G} |f(xy^{-1})|\big)\big(\sum_{y \in G} |f(xy^{-1})||g(y)|^2\big) = ||f||_1 \sum_{y \in G} |f(xy^{-1})||g(y)|^2.$$

Then,

$$||f * g||^2 = \sum_{x \in G} |f * g(x)|^2 \leq ||f||_1 \sum_{x,y \in G} |f(xy^{-1})||g(y)|^2 \leq ||f||_1 \sum_{x,y \in G} |f(x)||g(y)|^2$$

and so

$$||f * g||^2 \leq ||f||_1 \sum_{x \in G} |f(x)| \sum_{y \in G} |g(y)|^2 = ||f||_1^2 ||g||^2.$$

(iii) By definition, we have

$$||f * g||_\infty = \max_{x \in G} \big|\sum_{y \in G} f(xy^{-1})g(y)\big| \leq (\max_{x \in G} |f(x)|) \sum_{y \in G} |g(y)| = ||f||_\infty \cdot ||g||_1.$$

■

Convolution product will play a fundamental role starting from the next section. For the moment, let us just observe a couple of immediate consequences of the definition in relation to Cayley graphs.

If A is the adjacency operator of the Cayley graph $\Gamma[G, S]$ then, keeping in mind that S is symmetrical, we have

$$Af(x) = \sum_{y \sim x} f(y) = \sum_{s \in S} f(xs) = \sum_{y \in G} f(xy^{-1})\mathbf{1}_S(y)$$

for every $f \in \ell^2(G)$ and $x \in G$. Hence, for every $f \in \ell^2(G)$

$$Af = f * \mathbf{1}_S \tag{2.18}$$

In particular, let S be a symmetrical finite subset of the group G; we denote by ν_S the probability distribution on G uniformly centered on S, that is

$$\nu_S = |S|^{-1}\mathbf{1}_S.$$

Suppose further that $1 \notin S$, with $k = |S|$, and let $\widehat{A} = \frac{1}{k}A$ be the normalized adjacency operator of the Cayley graph $\Gamma = \Gamma[G, S]$. We then take the probability distribution $\nu_0 = \delta_1$ concentrated in 1 as the initial distribution on Γ, and, for every $i \geq 1$, denote by ν_i its i-th Markov iterate as in section 2.4. Then $\nu_1 = \nu_S$, and, for every $i \geq 1$, by (2.18) and what said in section 2.4,

$$\nu_{i+1} = \widehat{A}\nu_i = \frac{1}{k}(\nu_i * \mathbf{1}_S) = \nu_i * \nu_S.$$

Therefore, for every $i \geq 1$,

$$\widehat{A}^i\delta_1 = \nu_i = \nu_S^{(i)}. \tag{2.19}$$

where δ_1 is the probability distribution concentrated at 1.
In other words, for every $x \in G$ and $i \geq 1$, $\nu_S^{(i)}(x)$ is the probability that x may be written as the product of exactly i factors from S.

Example 20 (Schreier graphs). Given a faithful action of the (finite) group G on the finite set Ω, that is, an injective group-homomorphism $\phi : G \to Sym(\Omega)$ (for $g \in G$ and $x \in \Omega$, we write, $x^g = \phi(g)(x)$ and $G \leq Sym(\Omega)$). Let $S \subseteq G$ be a finite symmetric subset of G which acts fixed-point-free on Ω (this means that $x^s \neq x$ for every $s \in S$ and $x \in \Omega$). The associated *Schreier graph* is the graph $\Sigma = (\Omega, E)$ where for $x, y \in \Omega$, $\{x, y\} \in E$ if and only if there is $s \in S$ such that $y = x^s$. Then Σ is $|S|$-regular and it is connected if and only if the action of $\langle S \rangle$ is transitive. A Cayley graph is a Schreier graph with respect to the regular action of G on itself by right multiplication.
If one drops the assumption that S acts fixed-point-free on Ω, one still speaks of a Schreier graph, which is in general a graph with loops and multiple edges.

Note. The first two sections of this Chapter are standard, and belong to any basic account of (algebraic) graph theory. A good introduction to graph theory, which includes random graphs, is the book by Diestel [**7**].
The best general survey on expander graphs, with several proofs and special emphasis on motivations and applications, remains the paper [**18**] by Hoory, Liniel and Wigderson. Along with Tao's book [**31**], and towards different directions, it is the place where to look for much more additional material, and find most of the proofs that are missing in this chapter.