Basic Facts about Expander Graphs

Oded Goldreich

Abstract. In this survey we review basic facts regarding expander graphs that are most relevant to the theory of computation.

Keywords: Expander Graphs, Random Walks on Graphs.

This text has been revised based on [8, Apdx. E.2].

1 Introduction

Expander graphs found numerous applications in the theory of computation, ranging from the design of sorting networks [1] to the proof that undirected connectivity is decidable in determinstic log-space [15]. In this survey we review basic facts regarding expander graphs that are most relevant to the theory of computation. For a wider perspective, the interested reader is referred to [10].

Loosely speaking, expander graphs are regular graphs of small degree that exhibit various properties of cliques.[1] In particular, we refer to properties such as the relative sizes of cuts in the graph (i.e., relative to the number of edges), and the rate at which a random walk converges to the uniform distribution (relative to the logarithm of the graph size to the base of its degree).

Some technicalities. Typical presentations of expander graphs refer to one of several variants. For example, in some sources, expanders are presented as bipartite graphs, whereas in others they are presented as ordinary graphs (and are in fact very far from being bipartite). We shall follow the latter convention. Furthermore, at times we implicitly consider an augmentation of these graphs where self-loops are added to each vertex. For simplicity, we also allow parallel edges.

We often talk of expander graphs while we actually mean an infinite collection of graphs such that each graph in this collection satisfies the same property (which is informally attributed to the collection). For example, when talking of a d-regular expander (graph) we actually refer to an infinite collection of graphs such that each of these graphs is d-regular. Typically, such a collection (or family) contains a single N-vertex graph for every $N \in \mathbb{S}$, where $\mathbb{S}$ is an infinite subset of $\mathbb{N}$. Throughout this section, we denote such a collection by $\{G_N\}_{N\in\mathbb{S}}$, with the understanding that G_N is a graph with N vertices and $\mathbb{S}$ is an infinite set of natural numbers.

[1] Another useful intuition is that expander graphs exhibit various properties of random regular graphs of the same degree.

2 Definitions and Properties

We consider two definitions of expander graphs, two different notions of explicit constructions, and two useful properties of expanders.

2.1 Two mathematical definitions

We start with two different definitions of expander graphs. These definitions are qualitatively equivalent and even quantitatively related. We start with an algebraic definition, which seems technical in nature but is actually the definition typically used in complexity theoretic applications, since it directly implies various "mixing properties" (see §2.3). We later present a very natural combinatorial definition (which is the source of the term "expander").

The algebraic definition (eigenvalue gap). Identifying graphs with their adjacency matrix, we consider the eigenvalues (and eigenvectors) of a graph (or rather of its adjacency matrix). Any d-regular graph $G = (V, E)$ has the uniform vector as an eigenvector corresponding to the eigenvalue d, and if G is connected and non-bipartite then the absolute values of all other eigenvalues are strictly smaller than d. The eigenvalue bound, denoted $\lambda(G) < d$, of such a graph G is defined as a tight upper-bound on the *absolute value* of all the other eigenvalues. (In fact, in this case it holds that $\lambda(G) < d - \Omega(1/d|V|^2)$.)[2] The algebraic definition of expanders refers to an infinite family of d-regular graphs and requires the existence of a *constant* eigenvalue bound that holds for all the graphs in the family.

Definition 1 (eigenvalue gap): *An infinite family of d-regular graphs, $\{G_N\}_{N\in\mathbb{S}}$, where $\mathbb{S} \subseteq \mathbb{N}$,* satisfies the eigenvalue bound β *if for every $N \in \mathbb{S}$ it holds that $\lambda(G_N) \leq \beta$. In such a case, we say that $\{G_N\}_{N\in\mathbb{S}}$ is a family of* (d, β)-expanders, *and call $d - \beta$ its* eigenvalue gap.

It will be often convenient to consider relative (or normalized) versions of the foregoing quantities, obtained by division by d.

The combinatorial definition (expansion). Loosely speaking, expansion requires that any (not too big) set of vertices of the graph has a relatively large set of neighbors. Specifically, a graph $G = (V, E)$ is c-expanding if, for every set $S \subset V$ of cardinality at most $|V|/2$, it holds that

$$\Gamma_G(S) \stackrel{\text{def}}{=} \{v : \exists u \in S \text{ s.t. } \{u, v\} \in E\} \tag{1}$$

has cardinality at least $(1 + c) \cdot |S|$. Assuming the existence of self-loops on all vertices, the foregoing requirement is equivalent to requiring that $|\Gamma_G(S) \setminus S| \geq$

[2] This follows from the connection to the combinatorial definition (see Theorem 3). Specifically, the square of this graph, denoted G^2, is $|V|^{-1}$-expanding and thus it holds that $\lambda(G)^2 = \lambda(G^2) < d^2 - \Omega(|V|^{-2})$.

$c \cdot |S|$. In this case, every connected graph $G = (V, E)$ is $(1/|V|)$-expanding.[3] The combinatorial definition of expanders refers to an infinite family of d-regular graphs and requires the existence of a *constant* expansion bound that holds for all the graphs in the family.

Definition 2 (expansion): *An infinite family of d-regular graphs,* $\{G_N\}_{N \in \mathbb{S}}$ is c-expanding *if for every* $N \in \mathbb{S}$ *it holds that* G_N *is c-expanding.*

The two definitions of expander graphs are related (see [6, Sec. 9.2] or [10, Sec. 4.5]). Specifically, the "expansion bound" and the "eigenvalue bound" are related as follows.

Theorem 3 (equivalence of the two definitions [3, 5]): *Let G be a d-regular graph having a self-loop on each vertex.*[4]

1. *The graph G is c-expanding for* $c \geq (d - \lambda(G))/2d$.
2. *If G is c-expanding then* $d - \lambda(G) \geq c^2/(4 + 2c^2)$.

Thus, any non-zero bound on the combinatorial expansion of a family of d-regular graphs yields a non-zero bound on its eigenvalue gap, and vice versa. Note, however, that the back-and-forth translation between these measures is not tight. We note that most applications in complexity theory refer to the algebraic definition, and that the loss incurred in Theorem 3 is immaterial for them.

Amplification. The "quality of expander graphs improves" by raising these graphs to any power $t > 1$ (i.e., raising their adjacency matrix to the t^{th} power), where this operation corresponds to replacing t-paths (in the original graphs) by edges (in the resulting graphs). Specifically, when considering the algebraic definition, it holds that $\lambda(G^t) = \lambda(G)^t$, but indeed the degree also gets raised to the power t. Still, the ratio $\lambda(G^t)/d^t$ deceases with t. An analogous phenomenon occurs also under the combinatorial definition, provided that some suitable modifications are applied. For example, if for every $S \subseteq V$ it holds that $|\Gamma_G(S)| \geq \min((1 + c) \cdot |S|, |V|/2)$, then for every $S \subseteq V$ it holds that $|\Gamma_{G^t}(S)| \geq \min((1 + c)^t \cdot |S|, |V|/2)$.

[3] In contrast, a bipartite graph $G = (V, E)$ is not expanding, because it always contains a set S of size at most $|V|/2$ such that $|\Gamma_G(S)| \leq |S|$ (although it may hold that $|\Gamma_G(S) \setminus S| \geq |S|$).

[4] Recall that in such a graph $G = (V, E)$ it holds that $\Gamma_G(S) \supseteq S$ for every $S \subseteq V$, and thus $|\Gamma_G(S)| = |\Gamma_G(S) \setminus S| + |S|$. Furthermore, in such a graph all eigenvalues are greater than or equal to $-d+1$, and thus if $d - \lambda(G) < 1$ then this is due to a positive eigenvalue of G. These facts are used for bridging the gap between Theorem 3 and the more standard versions (see, e.g., [6, Sec. 9.2]) that refer to variants of both definitions. Specifically, [6, Sec. 9.2] refers to $\Gamma_G^+(S) = \Gamma_G(S) \setminus S$ and $\lambda_2(G)$, where $\lambda_2(G)$ is the second largest eigenvalue of G, rather than referring to $\Gamma_G(S)$ and $\lambda(G)$. Note that, in general, $\Gamma_G(S)$ may be attained by the difference between the smallest eigenvalue of G (which may be negative) and $-d$.

The optimal eigenvalue bound. For every d-regular graph $G = (V, E)$, it holds that $\lambda(G) \geq 2\gamma_G \cdot \sqrt{d-1}$, where $\gamma_G = 1 - O(1/\log_d |V|)$. Thus, for any infinite family of (d, λ)-expanders, it must holds that $\lambda \geq 2\sqrt{d-1}$.

2.2 Two levels of explicitness

Towards discussing various notions of explicit constructions of graphs, we need to fix a representation of such graphs. Specifically, throughout this section, when referring to an infinite family of graphs $\{G_N\}_{N \in \mathbb{S}}$, we shall assume that the vertex set of G_N equals $[N]$. Indeed, at times, we shall consider vertex sets having a different structure (e.g., $[m] \times [m]$ for some $m \in \mathbb{N}$), but in all these cases there exists a simple isomorphism of these sets to the canonical representation (i.e., there exists an efficiently computable and invertible mapping of the vertex set of G_N to $[N]$).

Recall that a mild notion of explicit constructiveness refers to the *complexity of constructing the entire object* (i.e., the graph). Applying this notion to our setting, we say that an infinite family of graphs $\{G_N\}_{N \in \mathbb{S}}$ is explicitly constructible if there exists a *polynomial-time algorithm that, on input* 1^N (where $N \in \mathbb{S}$), *outputs the list of the edges in the N-vertex graph G_N*. That is, the entire graph is constructed in time that is polynomial in its size (i.e., in poly(N)-time).

The foregoing (mild) level of explicitness suffices when the application requires holding the entire graph and/or when the running-time of the application is lower-bounded by the size of the graph. In contrast, other applications refer to a huge virtual graph (which is much bigger than their running time), and only require the computation of the neighborhood relation in such a graph. In this case, the following stronger level of explicitness is relevant.

A strongly explicit construction of an infinite family of (d-regular) graphs $\{G_N\}_{N \in \mathbb{S}}$ is a *polynomial-time algorithm that on input* $N \in \mathbb{S}$ (in binary), *a vertex v in the N-vertex graph G_N* (i.e., $v \in [N]$), *and an index* $i \in [d]$, *returns the i^{th} neighbor of v*. That is, the "neighbor query" is answered in time that is polylogarithmic in the size of the graph. Needless to say, this strong level of explicitness implies the basic (mild) level.

An additional requirement, which is often forgotten but is very important, refers to the "tractability" of the set $\mathbb{S}$. Specifically, we require the existence of an *efficient algorithm that given any* $n \in \mathbb{N}$ *finds an* $s \in \mathbb{S}$ *such that* $n \leq s < 2n$. Corresponding to the two foregoing levels of explicitness, "efficient" may mean either running in time poly(n) or running in time poly($\log n$). The requirement that $n \leq s < 2n$ suffices in most applications, but in some cases a smaller interval (e.g., $n \leq s < n + \sqrt{n}$) is required, whereas in other cases a larger interval (e.g., $n \leq s <$ poly(n)) suffices.

Greater flexibility. In continuation to the foregoing paragraph, we comment that expanders can be combined in order to obtain expanders for a wider range of graph sizes. For example, given two d-regular c-expanding graphs, $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$ where $|V_1| \leq |V_2|$ and $c \leq 1$, we can obtain a $(d+1)$-regular $c/2$-expanding graph on $|V_1| + |V_2|$ vertices by connecting the two graphs using a perfect matching of V_1 and $|V_1|$ of the vertices of V_2 (and adding self-loops to the remaining vertices of V_2). More generally, combining the d-regular c-expanding graphs $G_1 = (V_1, E_1)$ through $G_t = (V_t, E_t)$, where $N' \stackrel{\rm def}{=} \sum_{i=1}^{t-1} |V_i| \leq |V_t|$, yields a $(d+1)$-regular $c/2$-expanding graph on $\sum_{i=1}^{t} |V_i|$ vertices (by using a perfect matching of $\cup_{i=1}^{t-1} V_i$ and N' of the vertices of V_t).

2.3 Two properties

The following two properties provide a quantitative interpretation to the statement that expanders approximate the complete graph (or behave approximately like a complete graph). When referring to (d, λ)-expanders, the deviation from the behavior of a complete graph is represented by an error term that is linear in λ/d.

The mixing lemma. Loosely speaking, the following (folklore) lemma asserts that in expander graphs (for which $\lambda \ll d$) the fraction of edges connecting two large sets of vertices approximately equals the product of the densities of these sets. This property is called *mixing.*

Lemma 4 (Expander Mixing Lemma): *For every d-regular graph $G = (V, E)$ and for every two subsets $A, B \subseteq V$ it holds that*

$$\left| \frac{|(A \times B) \cap \vec{E}|}{|\vec{E}|} - \frac{|A|}{|V|} \cdot \frac{|B|}{|V|} \right| \leq \frac{\lambda(G)\sqrt{|A| \cdot |B|}}{d \cdot |V|} \leq \frac{\lambda(G)}{d} \qquad (2)$$

where $\vec{E}$ denotes the set of directed edges (i.e., vertex pairs) *that correspond to the undirected edges of G* (i.e., $\vec{E} = \{(u, v) : \{u, v\} \in E\}$ and $|\vec{E}| = d|V|$).

In particular, $|(A \times A) \cap \vec{E}| = (\rho(A) \cdot d \pm \lambda(G)) \cdot |A|$, where $\rho(A) = |A|/|V|$. It follows that $|(A \times (V \setminus A)) \cap \vec{E}| = ((1 - \rho(A)) \cdot d \pm \lambda(G)) \cdot |A|$.

Proof: Let $N \stackrel{\rm def}{=} |V|$ and $\lambda \stackrel{\rm def}{=} \lambda(G)$. For any subset of the vertices $S \subseteq V$, we denote its density in V by $\rho(S) \stackrel{\rm def}{=} |S|/N$. Hence, Eq. (2) is restated as

$$\left| \frac{|(A \times B) \cap \vec{E}|}{d \cdot N} - \rho(A) \cdot \rho(B) \right| \leq \frac{\lambda\sqrt{\rho(A) \cdot \rho(B)}}{d} .$$

We proceed by providing bounds on the value of $|(A \times B) \cap \vec{E}|$. To this end we let $\overline{a}$ denote the N-dimensional Boolean vector having 1 in the $i^{\rm th}$ component if and only if $i \in A$. The vector $\overline{b}$ is defined similarly. Denoting the adjacency matrix of the graph G by $M = (m_{i,j})$, we note that $|(A \times B) \cap \vec{E}|$ equals $\overline{a}^\top M \overline{b}$ (because $(i, j) \in (A \times B) \cap \vec{E}$ if and only if it holds that $i \in A$, $j \in B$ and $m_{i,j} = 1$). We consider the *orthogonal eigenvector basis*, $\overline{e_1}, ..., \overline{e_N}$, where $\overline{e_1} = (1, ..., 1)^\top$ and $\overline{e_i}^\top \overline{e_i} = N$ for each i, and write each vector as a linear combination of the vectors in this basis. Specifically, we denote by a_i the coefficient of $\overline{a}$ in the direction of $\overline{e_i}$; that is, $a_i = (\overline{a}^\top \overline{e_i})/N$ and $\overline{a} = \sum_i a_i \overline{e_i}$. Note that $a_1 = (\overline{a}^\top \overline{e_1})/N = |A|/N = \rho(A)$ and $\sum_{i=1}^N a_i^2 = (\overline{a}^\top \overline{a})/N = |A|/N = \rho(A)$. Similarly for $\overline{b}$. It now follows that

$$\begin{aligned}|(A \times B) \cap \vec{E}| &= \overline{a}^\top M \sum_{i=1}^N b_i \overline{e_i} \\ &= \sum_{i=1}^N b_i \lambda_i \cdot \overline{a}^\top \overline{e_i}\end{aligned}$$

where λ_i denotes the i^{th} eigenvalue of M. Note that $\lambda_1 = d$ and for every $i \geq 2$ it holds that $|\lambda_i| \leq \lambda$. Thus,

$$\begin{aligned}\frac{|(A \times B) \cap \vec{E}|}{dN} &= \sum_{i=1}^N \frac{b_i \lambda_i \cdot a_i}{d} \\ &= \rho(A)\rho(B) + \sum_{i=2}^N \frac{\lambda_i a_i b_i}{d} \\ &\in \left[\rho(A)\rho(B) \pm \frac{\lambda}{d} \cdot \sum_{i=2}^N a_i b_i\right]\end{aligned}$$

Using $\sum_{i=1}^N a_i^2 = \rho(A)$ and $\sum_{i=1}^N b_i^2 = \rho(B)$, and applying Cauchy-Schwartz Inequality, we bound $\sum_{i=2}^N a_i b_i$ by $\sqrt{\rho(A)\rho(B)}$. The lemma follows. ■

The random walk lemma. Loosely speaking, the first part of the following lemma asserts that, as far as remaining "trapped" in some subset of the vertex set is concerned, a random walk on an expander approximates a random walk on the complete graph.

Lemma 5 (Expander Random Walk Lemma): *Let $G = ([N], E)$ be a d-regular graph, and consider walks on G that start from a uniformly chosen vertex and take $\ell - 1$ additional random steps, where in each such step we uniformly selects one out of the d edges incident at the current vertex and traverses it.*

- *Let W be a subset of $[N]$ and $\rho \stackrel{\text{def}}{=} |W|/N$. Then the probability that such a random walk stays in W is at most*

$$\rho \cdot \left(\rho + (1 - \rho) \cdot \frac{\lambda(G)}{d}\right)^{\ell - 1} \tag{3}$$

- *For any $W_0, ..., W_{\ell-1} \subseteq [N]$, the probability that a random walk of length ℓ intersects $W_0 \times W_1 \times \cdots \times W_{\ell-1}$ is at most*

$$\sqrt{\rho_0} \cdot \prod_{i=1}^{\ell-1} \sqrt{\rho_i + (\lambda/d)^2}, \tag{4}$$

where $\rho_i \stackrel{\text{def}}{=} |W_i|/N$.

The basic principle underlying Lemma 5 was discovered by Ajtai, Komlos, and Szemerédi [2], who proved a bound as in Eq. (4). The better analysis yielding the first part of Lemma 5 is due to [12, Cor. 6.1]. A more general bound that refer to the probability of visiting W for a number of times that approximates $|W|/N$ is given in [9], which actually considers an even more general problem (i.e., obtaining Chernoff-type bounds for random variables that are generated by a walk on an expander). An alternative approach to obtaining such Chernoff-type bounds has been recently presented in [11].

Proof of Equation (4). The basic idea is viewing events occuring during the random walk as an evolution of a corresponding probability vector under suitable transformations. The transformations correspond to taking a random step in G and to passing through a "sieve" that keeps only the entries that correspond to the current set W_i. The key observation is that the first transformation shrinks the component that is orthogonal to the uniform distribution, whereas the second transformation shrinks the component that is in the direction of the uniform distribution. Details follow.

Let A be a matrix representing the random walk on G (i.e., A is the adjacency matrix of G divided by d), and let $\hat{\lambda} \stackrel{\text{def}}{=} \lambda(G)/d$ (i.e., $\hat{\lambda}$ upper-bounds the absolute value of every eigenvalue of A except the first one). Note that the uniform distribution, represented by the vector $\overline{u} = (N^{-1}, ..., N^{-1})^\top$, is the eigenvector of A that is associated with the largest eigenvalue (which is 1). Let P_i be a 0-1 matrix that has 1-entries only on its diagonal such that entry (j, j) is set to 1 if and only if $j \in W_i$. Then, the probability that a random walk of length ℓ intersects $W_0 \times W_1 \times \cdots \times W_{\ell-1}$ is the sum of the entries of the vector

$$\overline{v} \stackrel{\text{def}}{=} P_{\ell-1}A \cdots P_2AP_1AP_0\overline{u}. \tag{5}$$

We are interested in upper-bounding $\|\overline{v}\|_1$, and use $\|\overline{v}\|_1 \leq \sqrt{N} \cdot \|\overline{v}\|$, where $\|\overline{z}\|_1$ and $\|\overline{z}\|$ denote the L_1-norm and L_2-norm of $\overline{z}$, respectively (e.g., $\|\overline{u}\|_1 = 1$ and $\|\overline{u}\| = N^{-1/2}$). The key observation is that the linear transformation P_iA shrinks every vector.

Main Claim. For every $\overline{z}$, it holds that $\|P_iA\overline{z}\| \leq (\rho_i + \hat{\lambda}^2)^{1/2} \cdot \|\overline{z}\|$.

Proof. Intuitively, A shrinks the component of $\overline{z}$ that is orthogonal to $\overline{u}$, whereas P_i shrinks the component of $\overline{z}$ that is in the direction of $\overline{u}$. Specifically, we decompose $\overline{z} = \overline{z_1} + \overline{z_2}$ such that $\overline{z_1}$ is the projection of $\overline{z}$ on $\overline{u}$ and $\overline{z_2}$ is the component orthogonal to $\overline{u}$. Then, using the triangle inequality and other obvious facts (which imply $\|P_iA\overline{z_1}\| = \|P_i\overline{z_1}\|$ and $\|P_iA\overline{z_2}\| \leq \|A\overline{z_2}\|$), we have

$$\begin{aligned}
\|P_iA\overline{z_1} + P_iA\overline{z_2}\| &\leq \|P_iA\overline{z_1}\| + \|P_iA\overline{z_2}\| \\
&\leq \|P_i\overline{z_1}\| + \|A\overline{z_2}\| \\
&\leq \sqrt{\rho_i} \cdot \|\overline{z_1}\| + \hat{\lambda} \cdot \|\overline{z_2}\|
\end{aligned}$$

where the last inequality uses the fact that P_i shrinks any uniform vector by eliminating $1-\rho_i$ of its elements, whereas A shrinks the length of any eigenvector except $\overline{u}$ by a factor of at least $\hat{\lambda}$. Using the Cauchy-Schwartz inequality[5], we get

$$\begin{aligned}\|P_i A \overline{z}\| &\leq \sqrt{\rho_i + \hat{\lambda}^2} \cdot \sqrt{\|\overline{z_1}\|^2 + \|\overline{z_2}\|^2} \\ &= \sqrt{\rho_i + \hat{\lambda}^2} \cdot \|\overline{z}\|\end{aligned}$$

where the equality is due to the fact that $\overline{z_1}$ is orthogonal to $\overline{z_2}$. □

Recalling Eq. (5) and using the Main Claim (and $\|\overline{v}\|_1 \leq \sqrt{N} \cdot \|\overline{v}\|$), we get

$$\begin{aligned}\|\overline{v}\|_1 &\leq \sqrt{N} \cdot \|P_{\ell-1} A \cdots P_2 A P_1 A P_0 \overline{u}\| \\ &\leq \sqrt{N} \cdot \left(\prod_{i=1}^{\ell-1} \sqrt{\rho_i + \hat{\lambda}^2} \right) \cdot \|P_0 \overline{u}\|.\end{aligned}$$

Finally, using $\|P_0\overline{u}\| = \sqrt{\rho_0 N \cdot (1/N)^2} = \sqrt{\rho_0/N}$, we establish Eq. (4). ■

Rapid mixing. A property related to Lemma 5 is that a random walk starting at any vertex converges to the uniform distribution on the expander vertices after a logarithmic number of steps. Specifically, we claim that *starting at any distribution* $\overline{s}$ (including a distribution that assigns all weight to a single vertex) *after* ℓ *steps on a* (d,λ)*-expander* $G = ([N], E)$ *we reach a distribution that is* $\sqrt{N} \cdot (\lambda/d)^\ell$*-close to the uniform distribution over* $[N]$. Using notation as in the proof of Eq. (4), the claim asserts that $\|A^\ell \overline{s} - \overline{u}\|_1 \leq \sqrt{N} \cdot \hat{\lambda}^\ell$, which is meaningful only for $\ell > 0.5 \cdot \log_{1/\hat{\lambda}} N$. The claim is proved by recalling that $\|A^\ell \overline{s} - \overline{u}\|_1 \leq \sqrt{N} \cdot \|A^\ell \overline{s} - \overline{u}\|$ and using the fact that $\overline{s} - \overline{u}$ is orthogonal to $\overline{u}$ (because the former is a zero-sum vector). Thus, $\|A^\ell \overline{s} - \overline{u}\| = \|A^\ell(\overline{s} - \overline{u})\| \leq \hat{\lambda}^\ell \|\overline{s} - \overline{u}\|$ and using $\|\overline{s} - \overline{u}\| < 1$ the claim follows.

3 Constructions

Many explicit constructions of (d,λ)-expanders are known. The first such construction was presented in [14] (where $\lambda < d$ was not explicitly bounded), and an optimal construction (i.e., an optimal eigenvalue bound of $\lambda = 2\sqrt{d-1}$) was first provided in [13]. Most of these constructions are quite simple (see, e.g., §3.1), but their analysis is based on non-elementary results from various branches of mathematics. In contrast, the construction of Reingold, Vadhan, and Wigderson [16], presented in §3.2, is based on an iterative process, and its analysis is based on a relatively simple algebraic fact regarding the eigenvalues of matrices.

[5] That is, we get $\sqrt{\rho_i}\|z_1\| + \hat{\lambda}\|z_2\| \leq \sqrt{\rho_i + \hat{\lambda}^2} \cdot \sqrt{\|z_1\|^2 + \|z_2\|^2}$, by using $\sum_{i=1}^n a_i \cdot b_i \leq \left(\sum_{i=1}^n {a_i}^2\right)^{1/2} \cdot \left(\sum_{i=1}^n {b_i}^2\right)^{1/2}$, with $n = 2$, $a_1 = \sqrt{\rho_i}$, $b_1 = \|z_1\|$, etc.

Before turning to these explicit constructions we note that it is relatively easy to prove the existence of 3-regular expanders, by using the Probabilistic Method (cf. [6]) and referring to the combinatorial definition of expansion.[6]

3.1 The Margulis–Gabber–Galil Expander

For every natural number m, consider the graph with vertex set $\mathbb{Z}_m \times \mathbb{Z}_m$ and the edge set in which every $\langle x,y\rangle \in \mathbb{Z}_m \times \mathbb{Z}_m$ is connected to the vertices $\langle x \pm y, y\rangle$, $\langle x \pm (y+1), y\rangle$, $\langle x, y \pm x\rangle$, and $\langle x, y \pm (x+1)\rangle$, where the arithmetic is modulo m. This yields an extremely simple 8-regular graph with an eigenvalue bound that is a constant $\lambda < 8$ (which is independent of m). Thus, we get:

Theorem 6 *There exists a strongly explicit construction of a family of* $(8, 7.9999)$*-expanders for graph sizes* $\{m^2 : m \in \mathbb{N}\}$*. Furthermore, the neighbors of a vertex in these expanders can be computed in logarithmic-space.*[7]

An appealing property of Theorem 6 is that, for every $n \in \mathbb{N}$, it directly yields expanders with vertex set $\{0,1\}^n$. This is obvious in case n is even, but can be easily achieved also for odd n (e.g., use two copies of the graph for $n-1$, and connect the two copies by the obvious perfect matching).

Theorem 6 is due to Gabber and Galil [7], building on the basic approach suggested by Margulis [14]. We mention again that the (strongly explicit) (d, λ)-expanders of [13] achieve the optimal eigenvalue bound (i.e., $\lambda = 2\sqrt{d-1}$), but there are annoying restrictions on the degree d (i.e., $d-1$ should be a prime congruent to 1 modulo 4) and on the graph sizes for which this construction works.[8]

[6] This can be done by considering a 3-regular graph obtained by combining an N-cycle with a random matching of the first $N/2$ vertices and the remaining $N/2$ vertices. It is actually easier to prove the related statement that refers to the alternative definition of combinatorial expansion that refers to the relative size of $\Gamma_G^+(S) = \Gamma_G(S) \setminus S$ (rather than to the relative size of $\Gamma_G(S)$). In this case, for a sufficiently small $\varepsilon > 0$ and all sufficiently large N, a random 3-regular N-vertex graph is "ε-expanding" with overwhelmingly high probability. The proof proceeds by considering a (not necessarily simple) graph G obtained by combining three uniformly chosen perfect matchings of the elements of $[N]$. For every $S \subseteq [N]$ of size at most $N/2$ and for every set T of size $\varepsilon|S|$, we consider the probability that for a random perfect matching M it holds that $\Gamma_M^+(S) \subseteq T$. The argument is concluded by applying a union bound.

[7] In fact, for m that is a power of two (and under a suitable encoding of the vertices), the neighbors can be computed by a on-line algorithm that uses a constant amount of space. The same holds also for a variant in which each vertex $\langle x,y\rangle$ is connected to the vertices $\langle x \pm 2y, y\rangle$, $\langle x \pm (2y+1), y\rangle$, $\langle x, y \pm 2x\rangle$, and $\langle x, y \pm (2x+1)\rangle$. This variant yields a better known bound on λ, i.e., $\lambda \leq 5\sqrt{2} \approx 7.071$.

[8] The construction in [13] allows graph sizes of the form $(p^3 - p)/2$, where $p \equiv 1 \pmod 4$ is a prime such that $d-1$ is a quadratic residue modulo p. As stated in [4, Sec. 2], the construction can be extended to graph sizes of the form $(p^{3k} - p^{3k-2})/2$, for any $k \in \mathbb{N}$ and p as in the foregoing.

3.2 The Iterated Zig-Zag Construction

The starting point of the following construction is a very good expander G of *constant size*, which may be found by an exhaustive search. The construction of a large expander graph proceeds in iterations, where in the i^{th} iteration the current graph G_i and the fixed graph G are combined, resulting in a larger graph G_{i+1}. The combination step guarantees that the expansion property of G_{i+1} is at least as good as the expansion of G_i, while G_{i+1} maintains the degree of G_i and is a constant times larger than G_i. The process is initiated with $G_1 = G^2$ and terminates when we obtain a graph G_t of approximately the desired size (which requires a logarithmic number of iterations).

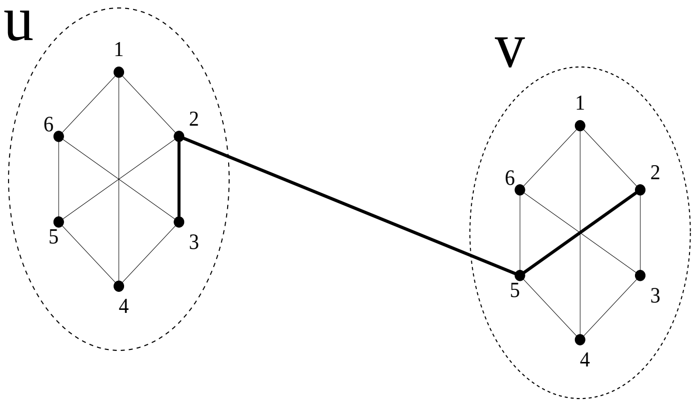

Fig. 1. Detail of the Zig-Zag product of G' and G. In this example G' is 6-regular and G is a 3-regular graph having six vertices. In the graph G' (not shown), the 2nd edge of vertex u is incident at v, as its 5th edge. The wide 3-segment line shows one of the corresponding edges of $G'ⓩG$, which connects the vertices $\langle u, 3\rangle$ and $\langle v, 2\rangle$.

The Zig-Zag product. The heart of the combination step is a new type of "graph product" called *Zig-Zag product*. This operation is applicable to any pair of graphs $G = ([D], E)$ and $G' = ([N], E')$, provided that G' (which is typically larger than G) is D-regular. For simplicity, we assume that G is d-regular (where typically $d \ll D$). The Zig-Zag product of G' and G, denoted $G'ⓩG$, is defined as a graph with vertex set $[N] \times [D]$ and an edge set that includes an edge between $\langle u, i\rangle \in [N] \times [D]$ and $\langle v, j\rangle$ if and only if $\{i, k\}, \{\ell, j\} \in E$ and the k^{th} edge incident at u equals the ℓ^{th} edge incident at v. That is, $\langle u, i\rangle$ and $\langle v, j\rangle$ are connected in $G'ⓩG$ if there exists a "three step sequence" consisting of a G-step from $\langle u, i\rangle$ to $\langle u, k\rangle$ (according to the edge $\{i, k\}$ of G), followed by a G'-step from $\langle u, k\rangle$ to $\langle v, \ell\rangle$ (according to the k^{th} edge of u in G' (which is the ℓ^{th} edge of v)), and a final G-step from $\langle v, \ell\rangle$ to $\langle v, j\rangle$ (according to the edge $\{\ell, j\}$ of G). See Figure 1 as well as further formalization (which follows).

It will be convenient to represent graphs like G' by their edge-rotation function, denoted $R' : [N] \times [D] \to [N] \times [D]$, such that $R'(u,i) = (v,j)$ if $\{u,v\}$ is the i^{th} edge incident at u as well as the j^{th} edge incident at v. That is, R' rotates the pair (u,i), which represents one "side" of the edge $\{u,v\}$ (i.e., the side incident at u as its i^{th} edge), resulting in the pair (v,j), which represents the other side of the same edge (which is the j^{th} edge incident at v). For simplicity, we assume that the (constant-size) d-regular graph $G = ([D], E)$ is edge-colorable with d colors, which in turn yields a natural edge-rotation function (i.e., $R(i,\alpha) = (j,\alpha)$ if the edge $\{i,j\}$ is colored α). We will denote by $E_\alpha(i)$ the vertex reached from $i \in [D]$ by following the edge colored α (i.e., $E_\alpha(i) = j$ iff $R(i,\alpha) = (j,\alpha)$). The Zig-Zag product of G' and G, denoted $G'\textcircled{z}G$, is then defined as a graph with the vertex set $[N] \times [D]$ and the edge-rotation function

$$(\langle u,i\rangle, \langle \alpha,\beta\rangle) \mapsto (\langle v,j\rangle, \langle \beta,\alpha\rangle) \quad \text{if } R'(u, E_\alpha(i)) = (v, E_\beta(j)). \tag{6}$$

That is, edges are labeled by pairs over $[d]$, and the $\langle \alpha,\beta\rangle^{\text{th}}$ edge out of vertex $\langle u,i\rangle \in [N] \times [D]$ is incident at the vertex $\langle v,j\rangle$ (as its $\langle \beta,\alpha\rangle^{\text{th}}$ edge) if $R(u, E_\alpha(i)) = (v, E_\beta(j))$, where indeed $E_\beta(E_\beta(j)) = j$. Intuitively, based on $\langle \alpha,\beta\rangle$, we first take a G-step from $\langle u,i\rangle$ to $\langle u, E_\alpha(i)\rangle$, then viewing $\langle u, E_\alpha(i)\rangle \equiv (u, E_\alpha(i))$ as a side of an edge of G' we rotate it (i.e., we effectively take a G'-step) reaching $(v,j') \stackrel{\text{def}}{=} R'(u, E_\alpha(i))$, and finally we take a G-step from $\langle v,j'\rangle$ to $\langle v, E_\beta(j')\rangle$.

Clearly, the graph $G'\textcircled{z}G$ is d^2-regular and has $D \cdot N$ vertices. The key fact, proved in [16] (using techniques as in §2.3), is that the relative eigenvalue-value of the zig-zag product is upper-bounded by the sum of the relative eigenvalue-values of the two graphs; that is, $\bar{\lambda}(G'\textcircled{z}G) \leq \bar{\lambda}(G') + \bar{\lambda}(G)$, where $\bar{\lambda}(\cdot)$ denotes the relative eigenvalue-bound of the relevant graph. The (qualitative) fact that $G'\textcircled{z}G$ is an expander if both G' and G are expanders is very intuitive (e.g., consider what happens if G' or G is a clique). Things are even more intuitive if one considers the (related) replacement product of G' and G, denoted $G'\textcircled{r}G$, *where there is an edge between $\langle u,i\rangle \in [N] \times [D]$ and $\langle v,j\rangle$ if and only if either $u = v$ and $\{i,j\} \in E$ or the i^{th} edge incident at u equals the j^{th} edge incident at v.*

The iterated construction. The iterated expander construction uses the aforementioned zig-zag product as well as graph squaring. Specifically, the construction starts[9] with the d^2-regular graph $G_1 = G^2 = ([D], E^2)$, where $D = d^4$ and $\bar{\lambda}(G) < 1/4$, and proceeds in iterations such that $G_{i+1} = G_i^2\textcircled{z}G$ for $i = 1, 2, ..., t-1$, where t is logarithmic in the desired graph size. That is, in each iteration, the current graph is first squared and then composed with the fixed (d-regular D-vertex) graph G via the zig-zag product. This process maintains the following two invariants:

1. The graph G_i is d^2-regular and has D^i vertices.

[9] Recall that, for a sufficiently large constant d, we first find a d-regular graph $G = ([d^4], E)$ satisfying $\bar{\lambda}(G) < 1/4$, by exhaustive search.

(The degree bound follows from the fact that a zig-zag product with a d-regular graph always yields a d^2-regular graph.)

2. The relative eigenvalue-bound of G_i is smaller than one half (i.e., $\bar{\lambda}(G_i) < 1/2$).
   (Here we use the fact that $\bar{\lambda}(G_{i-1}^2 \textcircled{z} G) \leq \bar{\lambda}(G_{i-1}^2) + \bar{\lambda}(G)$, which in turn equals $\bar{\lambda}(G_{i-1})^2 + \bar{\lambda}(G) < (1/2)^2 + (1/4)$. Note that graph squaring is used to reduce the relative eigenvalue of G_i before increasing it by zig-zag product with G.)

In order to show that we can actually construct G_i, we show that we can compute the edge-rotation function that correspond to its edge set. This boils down to showing that, given the edge-rotation function of G_{i-1}, we can compute the edge-rotation function of G_{i-1}^2 as well as of its zig-zag product with G. Note that this entire computation amounts to two recursive calls to computations regarding G_{i-1} (and two computations that correspond to the constant graph G). But since the recursion depth is logarithmic in the size of the final graph (i.e., $t = \log_D |\text{vertices}(G_t)|$), the total number of recursive calls is polynomial in the size of the final graph (and thus the entire computation is polynomial in the size of the final graph). This suffices for the minimal (i.e., "mild") notion of explicitness, but not for the strong one.

The strongly explicit version. To achieve a *strongly explicit construction*, we slightly modify the iterative construction. Rather than letting $G_{i+1} = G_i^2 \textcircled{z} G$, we let $G_{i+1} = (G_i \times G_i)^2 \textcircled{z} G$, where $G' \times G'$ denotes the *tensor product of* G' *with itself*; that is, if $G' = (V', E')$ then $G' \times G' = (V' \times V', E'')$, where

$$E'' = \{\{\langle u_1, u_2\rangle, \langle v_1, v_2\rangle\} : \{u_1, v_1\}, \{u_2, v_2\} \in E'\}$$

(i.e., $\langle u_1, u_2\rangle$ and $\langle v_1, v_2\rangle$ are connected in $G' \times G'$ if for $i = 1, 2$ it holds that u_i is connected to v_i in G'). The corresponding edge-rotation function is

$$R''(\langle u_1, u_2\rangle, \langle i_1, i_2\rangle) = (\langle v_1, v_2\rangle, \langle j_1, j_2\rangle),$$

where $R'(u_1, i_1) = (v_1, j_1)$ and $R'(u_2, i_2) = (v_2, j_2)$. We still use $G_1 = G^2$, where (as before) G is d-regular and $\bar{\lambda}(G) < 1/4$, but here G has $D = d^8$ vertices.[10] Using the fact that tensor product preserves the relative eigenvalue-bound while squaring the degree (and the number of vertices), we note that the modified iteration $G_{i+1} = (G_i \times G_i)^2 \textcircled{z} G$ yields a d^2-regular graph with $(D^{2^i - 1})^2 \cdot D = D^{2^{i+1}-1}$ vertices, and that $\bar{\lambda}(G_{i+1}) < 1/2$ (because $\bar{\lambda}((G_i \times G_i)^2 \textcircled{z} G) \leq \bar{\lambda}(G_i)^2 + \bar{\lambda}(G)$). Computing the neighbor of a vertex in G_{i+1} boils down to a constant number of such computations regarding G_i, but due to the tensor product operation the depth of the recursion is only double-logarithmic in the size of the final graph (and hence logarithmic in the length of the description of vertices in this graph).

[10] The reason for the change is that $(G_i \times G_i)^2$ will be d^8-regular, since G_i will be d^2-regular.

Digest. In the first construction, the zig-zag product was used both in order to increase the size of the graph and to reduce its degree. However, as indicated by the second construction (where the tensor product of graphs is the main vehicle for increasing the size of the graph), the primary effect of the zig-zag product is reducing the graph's degree, and the increase in the size of the graph is merely a side-effect.[11] In both cases, graph squaring is used in order to compensate for the modest increase in the relative eigenvalue-bound caused by the zig-zag product. In retrospect, the second construction is the "correct" one, because it decouples three different effects, and uses a natural operation to obtain each of them: Increasing the size of the graph is obtained by tensor product of graphs (which in turn increases the degree), the desired degree reduction is obtained by the zig-zag product (which in turn slightly increases the relative eigenvalue-bound), and graph squaring is used in order to reduce the relative eigenvalue-bound.

Stronger bound regarding the effect of the zig-zag product. In the foregoing description we relied on the fact, proved in [16], that the relative eigenvalue-bound of the zig-zag product is upper-bounded by the sum of the relative eigenvalue-bounds of the two graphs (i.e., $\bar{\lambda}(G' \textcircled{z} G) \leq \bar{\lambda}(G') + \bar{\lambda}(G))$). Actually, a stronger upper-bound is proved in [16]: It holds that $\bar{\lambda}(G' \textcircled{z} G) \leq f(\bar{\lambda}(G'), \bar{\lambda}(G)))$, where

$$f(x,y) \overset{\text{def}}{=} \frac{(1-y^2)\cdot x}{2} + \sqrt{\left(\frac{(1-y^2)\cdot x}{2}\right)^2 + y^2} \tag{7}$$

Indeed, $f(x,y) \leq (1-y^2)\cdot x + y \leq x+y$. On the other hand, for $x \leq 1$, we have $f(x,y) \leq \frac{(1-y^2)\cdot x}{2} + \frac{1+y^2}{2} = 1 - \frac{(1-y^2)\cdot(1-x)}{2}$, which implies

$$\bar{\lambda}(G' \textcircled{z} G) \;\leq\; 1 - \frac{(1-\bar{\lambda}(G)^2)\cdot(1-\bar{\lambda}(G'))}{2}\,. \tag{8}$$

Thus, $1 - \bar{\lambda}(G' \textcircled{z} G) \geq (1-\bar{\lambda}(G)^2)\cdot(1-\bar{\lambda}(G'))/2$, and it follows that the zig-zag product has a positive eigenvalue-gap if both graphs have positive eigenvalue-gaps (i.e., $\lambda(G' \textcircled{z} G) < 1$ if both $\lambda(G) < 1$ and $\lambda(G') < 1$). Furthermore, if $\bar{\lambda}(G) < 1/\sqrt{3}$ then $1 - \bar{\lambda}(G' \textcircled{z} G) > (1-\bar{\lambda}(G'))/3$. This fact plays an important role in the celebrated proof that undirected connectivity is decidable in determinstic log-space [15].

References

1. M. Ajtai, J. Komlos, E. Szemerédi. An $O(n \log n)$ Sorting Network, In *15th ACM Symposium on the Theory of Computing*, pages 1–9, 1983.
2. M. Ajtai, J. Komlos, E. Szemerédi. Deterministic Simulation in LogSpace. In *19th ACM Symposium on the Theory of Computing*, pages 132–140, 1987.

[11] We mention that this side-effect may actually be undesired in some applications. For example, in the proof of [15] we would rather not have the graph grow in size, but we can tolerate the constant size blow-up (caused by zig-zag product with a constant-size graph).

3. N. Alon. Eigenvalues and expanders. *Combinatorica*, Vol. 6, pages 83–96, 1986.
4. N. Alon, J. Bruck, J. Naor, M. Naor and R. Roth. Construction of Asymptotically Good, Low-Rate Error-Correcting Codes through Pseudo-Random Graphs. *IEEE Transactions on Information Theory*, Vol. 38, pages 509–516, 1992.
5. N. Alon and V.D. Milman. λ_1, Isoperimetric Inequalities for Graphs and Superconcentrators, *J. Combinatorial Theory, Ser. B*, Vol. 38, pages 73–88, 1985.
6. N. Alon and J.H. Spencer. *The Probabilistic Method.* John Wiley & Sons, Inc., 1992. Second edition, 2000.
7. O. Gabber and Z. Galil. Explicit Constructions of Linear Size Superconcentrators. *Journal of Computer and System Science*, Vol. 22, pages 407–420, 1981.
8. O. Goldreich. *Computational Complexity: A Conceptual Perspective.* Cambridge University Press, 2008.
9. A. Healy. Randomness-Efficient Sampling within NC1. *Computational Complexity*, to appear. Preliminary version in *10th RANDOM*, 2006.
10. S. Hoory, N. Linial, and A. Wigderson. *Expander Graphs and their Applications. Bull. AMS*, Vol. 43 (4), pages 439–561, 2006.
11. R. Impagliazzo and V. Kabanets. Constructive Proofs of Concentration Bounds. *ECCC*, TR10-072, 2010.
12. N. Kahale. Eigenvalues and Expansion of Regular Graphs. *Journal of the ACM*, Vol. 42 (5), pages 1091–1106, September 1995.
13. A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan Graphs. *Combinatorica*, Vol. 8, pages 261–277, 1988.
14. G.A. Margulis. Explicit Construction of Concentrators. *Prob. Per. Infor.*, Vol. 9 (4), pages 71–80, 1973 (in Russian). English translation in *Problems of Infor. Trans.*, pages 325–332, 1975.
15. O. Reingold. Undirected ST-Connectivity in Log-Space. In *37th ACM Symposium on the Theory of Computing*, pages 376–385, 2005.
16. O. Reingold, S. Vadhan, and A. Wigderson. Entropy Waves, the Zig-Zag Graph Product, and New Constant-Degree Expanders and Extractors. *Annals of Mathematics*, Vol. 155 (1), pages 157–187, 2001. Preliminary version in *41st FOCS*, pages 3–13, 2000.