Generating expanders from two permutations

Miklós Ajtai, János Komlós* and Endre Szemerédi

Abstract

Given $\alpha > 0$, we say that the sequence $\pi_1, \pi_2, \ldots, \pi_r$ of permutations of $\{1, 2, \ldots, n\}$ has the expanding property if, for all $X \subset \{1, 2, \ldots, n\}$ of size at least αn,

$$\left| \bigcup_{i=1}^{r} \pi_i X \right| > (1-\alpha)n.$$

We show that there exist pairs (σ, τ) of permutations such that the generated permutations $\pi_i = \sigma^i \tau^i$, $1 \leqslant i \leqslant r$, are expanding, where $r = O(1/\alpha^3)$. In fact, most pairs of permutations have this property.

No construction of such pairs of permutations is known.

0 Introduction

Expanders are of ever-increasing use in computer science. The random generation of expanders (Pinsker [8], Pippinger [9] and others) has been replaced by explicit constructions (Margulis [6], Gabber–Galil [4] and, recently, Lubotzky–Phillips–Sarnak [5]).

On the other hand, the recent characterization of expanders through eigenvalues by Alon [1] provides efficient testing for the expanding property, thus reviving interest in the random generation of expanders.

Here we offer a compromise approach: generating only a fraction of the bits at random and then extending them deterministically.

* Work supported by the NSF grant NSF-CCR 8505053.

Definitions

In the following, we will always assume that $0 < \alpha < 1$, $A > 1$ and $A\alpha < 1$ and, for simplicity of notation, we will also assume that αn is integral.

A *weak α-expander* on $2n$ vertices is a bipartite graph $G = (U, V, E)$ such that $|U| = |V| = n$ and, for any pair of subsets $X \subset U$ and $Y \subset V$, $|X| = |Y| = \alpha n$, there is at least one edge between X and Y.

A strong (A, α)*-expander* is a bipartite graph $G = (U, V, E)$ such that $|U| = |V| = n$ and, for any $X \subset U$, $|X| \leqslant \alpha n$, we have that $|N(X)| \geqslant A|X|$, where $N(X)$ is the neighbourhood-set of X:

$$N(X) = \{v \in V;\ (x, v) \in E \text{ for some } x \in X\}.$$

We will use the parameter value $A = (1-\alpha)/\alpha$, so that $\alpha = 1/(A+1)$, and write *strong α-expander* for strong (A, α)-expander.

Connections

We will only deal with regular bipartite graphs $G = (U, V, E)$, where we identify both vertex sets U and V by $[n] = \{1, 2, \ldots, n\}$. By König's theorem, an r-regular bipartite graph $G = (U, V, E)$ consists of r one-factors between U and V, i.e. permutations π_i, $1 \leqslant i \leqslant r$, of $\{1, 2, \ldots, n\}$, where, within the ith one-factor, vertex $u \in U$ is connected with vertex $v \in V$ such that $v = \pi_i u$.

Since very often the connections represent expensive modules (e.g. comparator switches), an alternative way to represent one-factors between U and V is to use two extra sets U' and V' of in-between vertices with fixed connections (the jth element of U' is connected to the jth element of V'). The set U is then connected to the set U' using a permutation connection σ, and similarly V is connected to V' by using some permutation τ.

This is, of course, equivalent to using the permutation $\pi = \tau^{-1}\sigma$ directly between U and V, but, while σ and τ may vary, the identity connection between U' and V' remains the same. The connections σ and τ are only used to communicate data from U to U' and from V to V', and the expensive part of the work is done along the fixed connection between U and V.

Thus, given permutations σ_i and τ_i, $1 \leqslant i \leqslant r$, we define the corresponding bipartite graph as

$$G = (U, V, E); \qquad |U| = |V| = n,$$
$$E = \{(\sigma_i^{-1}k, \tau_i^{-1}k);\ k \in [n],\ 1 \leqslant i \leqslant r\},$$

which is the same as the graph

$$G = (U, V, E); \qquad |U| = |V| = n,$$
$$E = \{(k, \pi_i k);\ k \in [n],\ 1 \leqslant i \leqslant r\},$$

where $\pi_i = \tau_i^{-1}\sigma_i$.

Notation

$c_0, c_1, c_2, \ldots$ are *absolute* constants, while c is a 'generic constant' with possibly different values at each occurrence.

log stands for the natural logarithm. $\log n$ and $\log\log n$ are truncated from below, so that their value is at least 1. $+$ will always mean addition modulo n. For simplicity of notation, we will always assume that αn is integral.

$[n] = \{1, 2, \ldots, n\}$. $[a, b]$ denotes the set of integers in the interval (a, b).

$\bar{A}$ denotes the complementary of the set A.

We will use the entropy function $h(\alpha) = -\alpha\log\alpha - (1-\alpha)\log(1-\alpha)$, together with the inequality $\binom{n}{\alpha n} \leqslant \exp h(\alpha) n$.

Random generation of expanders

It is known that, for most r-tuples of permutations $(\pi_1, \ldots, \pi_r)$, the corresponding bipartite graphs are strong α-expanders, where

$$r = r(\alpha) = c_1 \frac{1}{\alpha} \log \frac{1}{\alpha}.$$

(Note that Zarankiewicz's theorem [12] implies that *any* α-expander (even in the weak sense) has a degree of at least $c(1/\alpha)\log(1/\alpha)$.)

On the other hand, it is easy to see that, if we select only one permutation and generate the others from it as $\pi_i = \sigma^i$, $i = 1, 2, \ldots, r$, then, for large n, the corresponding graph is *never* an expander. (Here 'large' means large in terms of r, which, in turn, is determined by α.) The reason is that these permutations have the same cycle structures.

We are going to show, however, that **two permutations can already generate expanders**.

In what follows, $r = r(\alpha)$ will always be a fixed function, upper-bounded by a polynomial of $1/\alpha$. The proofs will show that

$$r(\alpha) = c_2 \left(\frac{1}{\alpha} \log \frac{1}{\alpha}\right)^2$$

suffices.

1 Formulation of results

We now define the crucial notion of *mixing* permutations. For the sake of better understanding, we define it in two equivalent ways.

Definition We say that the pair (σ, τ) of permutations of $[n]$ is *mixing* if

(first form): for all $X \subset [n]$, $|X| \geqslant n^{1-c_3}$, we have

$$\left| \bigcup_{1 \leqslant i \leqslant r(\alpha)} \pi_i X \right| > n - |X|, \tag{1}$$

where $\pi_i = \tau^{-i}\sigma^i$ and $\alpha = |X|/n$: in other words, for all $\alpha > n^{-c_3}$, the permutations $\pi_i = \tau^{-i}\sigma^i$, $1 \leqslant i \leqslant r(\alpha)$, form a weak α-expander;

(second form): for any two sets $X, Y \subset [n]$, $|X|, |Y| \geqslant n^{1-c_3}$, there exists an i, $1 \leqslant i \leqslant r(\alpha)$, such that

$$\sigma^i X \cap \tau^i Y \neq \varnothing, \tag{2}$$

where

$$\alpha = \frac{1}{n} \min\{|X|, |Y|\}:$$

in other words, for any $\alpha > n^{-c_3}$, the following graph G_α is a weak α-expander:

$$G_\alpha = (U, V, E); \qquad |U| = |V| = n,$$
$$E = \{(\sigma^{-i}k, \tau^{-i}k);\ k \in [n],\ 1 \leqslant i \leqslant r(\alpha)\}.$$

Probably the second form is the most illuminating.

Theorem 1 (*a*) *For every n, there exist mixing pairs (σ, τ) of permutations of $\{1, 2, \ldots, n\}$. In fact, most pairs (σ, τ) are mixing.*

Here 'most' means that the proportion of pairs that are not mixing approaches to zero as n tends to infinity. (It does so at an exponential rate.)

(*b*) *The same remains true if we restrict σ and τ to cyclic permutations (in which case, of course, proportions are calculated within the class of cyclic permutations of size n).*

(*c*) *The same remains true for most cyclic τ if we fix σ to be the right shift ($i \to i+1$, $n \to 1$).*

It is clear that (*c*) implies (*b*) and it is not hard to see that (*b*) implies (*a*). We will only prove (*c*).

Remark It is easy to see (by substituting $\frac{1}{2}\alpha$ for α) that (2) can be replaced by

$$|\sigma^i X \cap \tau^i Y| > \delta(\alpha) n,$$

where $\delta(\alpha) = \frac{1}{2}\alpha / r(\frac{1}{2}\alpha)$.

It is very likely that (2) can actually be replaced by

$$|\sigma^i X \cap \tau^i Y| > \tfrac{1}{2}\alpha^2 n.$$

Theorem 2 *Theorem 1 remains true if, in the definition of mixing, we replace the words 'weak α-expanders' by 'strong α-expanders'.*

In other words, most pairs (σ, τ) have the following property: for all $X \subset [n]$ and α, $\alpha \geq \max\{n^{-c_4}, |X|/n\}$, we have that

$$\left| \bigcup_{1 \leq i \leq r(\alpha)} \pi_i X \right| \geq \frac{1-\alpha}{\alpha} |X|, \tag{3}$$

where $\pi_i = \tau^{-i}\sigma^i$.

Remark Sarnak [10] has a simple construction (with a hard proof) for two permutations which generate expanders. Namely, if n is prime, then the graph corresponding to the permutations

$$\sigma k = k+1 \pmod n, \qquad \tau k = k^{-1} \pmod n$$

is expanding with a small factor. However, repeated application of the same permutations does not improve expansion, so the large amount of expansion required by our definitions cannot be achieved.

Using the above-described way of realizing connections through two extra layers, the theorems translate to generating expanding connections networks by using only five (fixed-wired) permutations, which give *3-regular fixed connection networks*: σ between U and U', τ between V and V', and the identity permutations between U and U', between U' and V' and between V and V'. U sends information to U' using σ and similarly V to V' using τ. Next, U' and V' communicate using the identity, then U' recycles information to U using the identity; same with V.

The process repeats a number of times, and we use the same network to get a better expansion by simply running it a few more times (increase r).

2 Proof of the theorems

For the proof of Theorem 1 we will use four lemmas, whose proofs we present in a separate section.

The following lemma has been conjectured by Minc [7] and proved by Brègman [3]. (For a simple proof, see Schrijver [11].)

Lemma 1 *The permanent of an $n \times n$ 0-1 matrix does not exceed*

$$\prod_{i=1}^{n} (r_i!)^{1/r_i}, \tag{4}$$

where r_i are the row-sums of the matrix.

Corollary 1 *The permanent of an $n \times n$ 0-1 matrix is at most $n! p^{c_5 n}$, where p is the proportion of ones in the whole matrix.*

We also use two lemmas that say roughly the following: There are two small families F_1 and F_2 of random-looking sets such that every set contains – as a subset – a member of F_1 and is contained in a member of F_2.

Definition A set R is called d-random if, for any set B of size at most d,

$$|R+B| \geqslant \tfrac{1}{2}|R||B|. \tag{5}$$

($R+B$ is defined as the set $\{r+b; r \in R, b \in B\}$; in particular, it is empty if either R or B is empty.)

Lemma 2 (Random-looking sets) *If $N > c_6 d \log d$, then any set A of size $c_7 dN$ contains a d-random subset R of size N. In fact, most subsets of A of size N are d-random.*

The lemma is true in any additive group; in particular, it is also valid modulo n.

Remark The lemma offers only a probabilistic construction. One may be tempted to exhibit such subsets R by constructing sets with small discrete Fourier transforms. This, however, presents unsurmountable technical difficulties.

For proving Theorem 2, the following counterpart of the previous lemma will be used.

Lemma 3 (Covering k-sets with random-looking sets) *Given n, k and $d > c_8$, let $l > c_9 k$ be such that*

$$\frac{l}{n} < \left(\frac{k}{n}\right)^{c_{10}/d}.$$

Then, there exists a family $\{S_1, S_2, \ldots, S_N\}$ of l-sets such that

(a) *the sets S_i, $1 \leqslant i \leqslant N$, cover all k-subsets of $[1, n]$;*

(b) *for any S_i and any d-set B,*

$$|\overline{\overline{S_i} + B}| < c_{11} k.$$

Furthermore,

$$N = 2\frac{\binom{n}{k}}{\binom{l}{k}}\log\binom{n}{k}$$

and most families of N l-sets satisfy the above two conditions.

Proof of (c) of Theorem 1 Fix a set X, $|X| = \alpha n$, and the permutation σ, and let Y, $|Y| = \alpha n$, vary together with the random permutation τ.

Define the sets $I(k) = \{i \leqslant r; \sigma^{-i}k \in X\} = (k-X)\cap[1,r]$, $1 \leqslant k \leqslant n$. If $\sigma^i \cap \tau^i = \varnothing$ for all $i \leqslant r$, then, for all k and $i \in I(k)$, we have $\tau^{-i} \notin Y$.

Let us represent Y by its location on τ, that is, define $\tilde{Y} = \{j : \tau^j 1 \in Y\}$. We now fix $\tilde{Y}$ (rather than Y) and compute the probability of the event

$$C = \{\tau : \text{for all } k \text{ and } i \in I(k),\ k \notin \tau^i Y\}.$$

(The probability of C is defined as $|C|/(n-1)!$.)

Now,

$$\begin{aligned} C &= \{\tau : \text{for all } k \text{ and } i \in I(k) \text{ and } j \in \tilde{Y},\ k \notin \tau^{i+j}1\} \\ &= \{\sigma : \text{for all } k,\ k \notin \tau^{\tilde{Y}+I(k)}1\}, \end{aligned}$$

where $+$ is meant modulo n.

Let us represent τ by the $(n-1)\times(n-1)$ permutation matrix

$$b_{ij} = \begin{cases} 1 & \text{if } i+1 = \tau^j 1, \\ 0 & \text{otherwise,} \end{cases}$$

where $1 \leqslant i, j \leqslant n-1$. Then any event C means that τ is chosen as a permutation from the 0-1 matrix A, whose $(k-1)$-th row, $2 \leqslant k \leqslant n$, has zero in the locations $\tilde{Y}+I(k) \pmod n$.

This, by Lemma 1, has a probability less than

$$\left(\frac{1}{n^2}\sum_{k\in[2,n]}|\overline{\tilde{Y}+I(k)}|\right)^{cn} < \exp\left\{-c\sum_{k\in[2,n]}\frac{|\tilde{Y}+I(k)|}{n}\right\} < e^{-c\alpha^2 n}, \quad (6)$$

since at least αn of the sets $I(k)$ are non-empty.

Unfortunately, the number of possible sets Y or, which is the same, the number of $\tilde{Y}$, is too large (about $\exp h(\alpha)n$). To cope with this, we select a small, random-looking subset Y' inside Y. This will drastically decrease the number of possible choices for Y', but, if r was chosen large enough, most of the sets $\tilde{Y}'+I(k)$ will still be as large as before.

More precisely, let us select a d-random subset $\tilde{Y}'$ of $\tilde{Y}$ of size

$$\frac{1}{d}|\tilde{Y}|, \qquad \text{where } d = c_{12}\frac{1}{\alpha}\log\frac{1}{\alpha}.$$

By Lemma 2, such a $\tilde{Y}'$ exists if $|\tilde{Y}| = |Y| = \alpha n \geqslant c_{13}d^2\log d$. Now, since $\tilde{Y}'$ is d-random, we have

$$|\tilde{Y}' + I(k)| \geqslant \tfrac{1}{2}|\tilde{Y}|\min\{d, |I(k)|\}. \tag{7}$$

We have to estimate from below the sum $S = \sum_k \min\{d, |I(k)|\}$, since equation (6) estimates the probability in question by $\exp\{-c|\tilde{Y}|S/n\} \leqslant \exp\{-c\alpha S\}$.

Another problem we have to deal with is the large number of choices for X. Just as for the set Y, we can represent X by its location on σ, that is, define $\tilde{X} = \{i : \sigma^i 1 \in X\}$ and, similarly, let $\tilde{k}$ be such that $k = \sigma^{\tilde{k}}1$. Thus, $I(k) = [1, r] \cap (\tilde{k} - \tilde{X})$.

To combine the two remaining tasks (shrinking X and estimating S from below), we use the following simple result.

Lemma 4 *Let $r = d^2$. For any Z, $|Z| = \alpha n$, there is a subset $Z' \subset Z$, $|Z'| = |Z|/d$, such that*

$$\frac{1}{n}\sum_{k=1}^{n} \min\{d, |Z' \cap [k+1, k+r]|\} \geqslant c_{14}\alpha d.$$

Using Lemma 4 and inequalities (6) and (7), we get the bound

$$\text{Prob}(C) \leqslant \exp\{-c_{15}\alpha'\alpha dn\} = \exp\{-c_{16}\alpha^2 n\}.$$

The number of choices for X' and Y' is at most $\exp 2h(\alpha')n$, where $\alpha' = c_{17}\alpha/d$. Hence, the probability that C occurs for some X and Y is less than $\exp\{-c_{18}\alpha^2 n\}$ as long as $h(\alpha') < c_{19}\alpha^2$.

This is satisfied by the choice

$$d = c\frac{1}{\alpha}\log\frac{1}{\alpha}$$

(with large c), which leads to

$$r = c_{20}\left(\frac{1}{\alpha}\log\frac{1}{\alpha}\right)^2. \qquad \square$$

Proof of Theorem 2 Since the proof is very similar to that of Theorem 1, we will only sketch it.

By changing α to $\frac{1}{2}\alpha$, we can reduce the problem to the following. Show that almost all pairs (σ, τ) of cyclic permutations satisfy the following condition.

For all $X \subset [n]$, $|X| \leqslant \alpha n$, we have

$$\left| \bigcup_{1 \leqslant i \leqslant r} \pi_i X \right| > A|X|, \tag{8}$$

where $A = 1/2\alpha$ and $\pi_i = \tau^{-i}\sigma^i$.

Or, equivalently, for any two sets $X, Y \subset [n]$, $|Y| = n - A|X| \geqslant \frac{1}{2}n$, there exists an i, $1 \leqslant i \leqslant r$, such that

$$\sigma^i X \cap \tau^i Y \neq \emptyset. \tag{9}$$

The proof is similar to that of Theorem 1. We select a small, random-looking set inside X using Lemma 2, and a random-looking set inside Y from the small family of such sets guaranteed by Lemma 3. For this latter selection, we choose an l such that

$$\left(\frac{k}{n}\right)^{c_{21}/A} < \frac{l}{n} < \left(\frac{k}{n}\right)^{c_{22}/d}.$$

The rest of the proof is standard calculus. □

Remark The proofs would be greatly simplified if we only set out to prove that the larger family of permutations $\tau^{-i}\sigma^j$, $1 \leqslant i, j \leqslant r$, is expanding.

3 Proof of the lemmas

Proof of Lemma 2 Write $n = c_7 dN$ and $p = c_{23}/d$, where $c_7 c_{23} > 1$. We choose a random subset of A by making independent randomizations for each point in A whether to include it in R. If each point has a probability p to succeed, then, with a very large probability, the obtained subset will be of size greater than N. Delete arbitrary points to make the size equal to N.

Given an arbitrary set $B = \{b_1, b_2, \ldots\}$ of size d, we write $R_i = R + b_i$. We have

$$|R+B| = \left| \bigcup_{i=1}^{d} R_i \right| \geqslant dN - \sum_{1 \leqslant i < j \leqslant d} |R_i R_j|. \tag{10}$$

Let us write $\|R\|$ for the maximum number of times a non-zero integer can appear as the difference of two elements in R:

$$\|R\| = \max_{\Delta \neq 0} |\{(r, r'); r, r' \in R, r - r' = \Delta\}|. \tag{11}$$

Clearly, $|R_i R_j| \leqslant \|R\|$, and thus

$$|R+B| \geqslant dN - \binom{d}{2}\|R\|. \tag{12}$$

We show now that, with a large probability, $\|R\| \leqslant N/d$, which implies that $|R+B| \geqslant \frac{1}{2}dN = \frac{1}{2}|R||B|$. (Since $\|R\| \leqslant N/d$ is the only property we use, it is clear that the proof also applies in the case $|B| < d$.)

Let us fix a non-zero integer Δ, and define a graph G on the points of A by connecting two elements if their difference is equal to Δ. It is clear that G is a union of (vertex-) disjoint paths and cycles (additive group!) and isolated vertices.

The subset R defines a (spanned) subgraph H of G, which is also a union of disjoint paths and cycles. We have to estimate the probability that H has more than $x = N/d$ edges.

For a specific choice of x edges from G, the probability that they all get into H is p^m, where m is the number of vertices these edges cover. (Actually, it is even less than that, since some of the vertices have been thrown away.)

If $N(m,x)$ is the number of ways to choose x edges from G which cover m vertices altogether, then we want to estimate the quantity

$$Q = \sum_m N(m,x)p^m.$$

Now, one can specify the obtained paths (and cycles) in H by specifying their 'left' endpoints and their lengths. If the x edges form k such intervals then $m = x+k$. Since the sum of the lengths is x, we get

$$N(m,x) \leqslant \sum_{k=1}^{x} \binom{n}{k}\binom{x-1}{k-1}$$

and obtain the estimate

$$Q \leqslant \sum_{k=1}^{x} \binom{n}{k}\binom{x-1}{k-1}p^{x+k} < \binom{n}{x}2^x p^{2x} < \left(\frac{2enp^2}{x}\right)^x$$
$$= (2ec_7c_{23}^2)^x < \frac{1}{n^3}$$

if $2ec_7c_{23}^2 < 1$ and c_6 was chosen large enough.

Since there are only $O(n^2)$ potential values for Δ, the lemma is proved. □

Proof of Lemma 3 Let us choose N sets of size l at random. For a fixed k-set, the probability that it is not contained in any of the l-sets is