

Deconstructing 1-local expanders

Oded Goldreich

Department of Computer Science
Weizmann Institute of Science, Rehovot, ISRAEL.
oded.goldreich@weizmann.ac.il

September 27, 2016

Abstract

Contemplating the recently announced 1-local expanders of Viola and Wigderson (*ECCC*, TR16-129, 2016), one may observe that weaker constructs are well known. For example, one may easily obtain a 4-regular N -vertex graph with spectral gap that is $\Omega(1/\log^2 N)$, and similarly a $O(1)$ -regular N -vertex graph with spectral gap $1/\tilde{O}(\log N)$. Starting from a generic candidate for a 1-local expander, we formulate a natural problem regarding “coordinated random walks” (CRW) on the corresponding “relocation” graph (which has size that is logarithmic in the size of the candidate 1-local graph), and observe that

1. any solution to the CRW problem yields 1-local expanders, and
2. any constant-size expanding set of generators for the symmetric group yields a solution to the CRW problem.

This yields an alternative construction and different analysis than the one used by Viola and Wigderson. Furthermore, we show that solving the CRW problem is equivalent to constructing 1-local expanders.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 1 |
| 2 | Initial thoughts | 1 |
| 3 | A sufficient condition | 4 |
| 4 | Known constructions that satisfy the CRW property | 6 |
| 5 | A sufficient and necessary condition | 11 |
| 6 | An afterthought: Generalization to non-binary alphabets | 13 |

1 Introduction

A function $f : \{0, 1\}^n \rightarrow \{0, 1\}^n$ is called t -local if each bit in its output depends on at most t bits in its input. We study the following recent result of Viola and Wigderson [6], where (throughout this text) we view n as varying.

Theorem 1 (a construction of 1-local expanders [6]): *There exists a constant d and a set of d explicit 1-local bijections, $\{f_1, \dots, f_d : \{0, 1\}^n \rightarrow \{0, 1\}^n\}_{n \in \mathbb{N}}$, such that the $2d$ -regular 2^n -vertex graph that consists of the vertex set $\{0, 1\}^n$ and the edge multiset $\cup_{i \in [d]} \{\{x, f_i(x)\} : x \in \{0, 1\}^n\}$ is an expander.*

Note that each f_i is determined by a permutation on the bit locations $\pi^{(i)} : [n] \rightarrow [n]$, called the relocation, and an offset $s^{(i)} \in \{0, 1\}^n$ such that $f_i(x) = x_{\pi^{(i)}} \oplus s^{(i)}$, where $x_{\pi^{(i)}} = x_{\pi^{(i)}(1)} \cdots x_{\pi^{(i)}(n)}$; that is, $f_i(x)$ is the string obtained by relocating the bits of x according to $\pi^{(i)}$ and offsetting the result by $s^{(i)}$. Indeed, by association, we refer to a $2d$ -regular graph with an edge multi-set that is described by 1-local bijections by the term 1-local.

Recall that the (normalized) second eigenvalue of a regular graph represents the rate at which a random walk on the graph converges to the uniform distribution (hereafter called the convergence rate). In an expander this rate is a constant smaller than 1, whereas in a general (regular and non- k -partite) N -vertex graph the rate is upper-bounded by $1 - \frac{1}{\text{poly}(N)}$. When trying to estimate the convergence rate, denoted λ , of an N -vertex regular graph it is useful to recall the following facts, where $\Delta_t^{(p)}$ denotes the distance (in norm L_p) of the uniform distribution from the distribution of the final vertex in a t -step random walk that starts at the worst possible vertex:¹

1. $\Delta_t^{(1)} \leq \sqrt{N} \cdot \Delta_t^{(2)} \leq \sqrt{N} \cdot \lambda^t$.
2. $N^{-1} \cdot \lambda^t \leq \Delta_t^{(2)} \leq \Delta_t^{(1)}$.

Hence, for sufficiently large t , it holds that $\lambda \approx (\Delta_t^{(1)})^{1/t}$.

2 Initial thoughts

Obtaining a 1-local expander requires using *both* the offsets (i.e., $s^{(i)}$'s) and the relocation permutations, because without the offsets the f_i 's maintain the Hamming weight of the vertex (and so the 2^n -vertex graph is not even connected), whereas without the permutations the 2^n -vertex graph decomposes into even smaller connected components (i.e., each of size at most 2^d). On the other hand, using both offsets and relocations, it is quite easy to obtain 1-local 4-regular graphs with polylogarithmic mixing time (equiv., the rate of convergence is bounded away from 1 by the reciprocal of a polylogarithmic function in the size of the graph).

¹The first inequality is well-known and captures the fact that the corresponding linear operator shrinks each vector that is orthogonal to the uniform one. The second inequality can be proved by considering a random walk that starts in a probability distribution that is described by the vector $u + v_2$, where $u = (1/N, \dots, 1/N)$ is the uniform distribution and v_2 is a vector in the direction of the second eigenvector (such that no coordinate has value lower than $-1/N$).

Observation 2 (the “shuffle exchange” graph is a 1-local weak expander):² Let $f_1(x) = \mathbf{sh}(x)$ and $f_2(x) = x \oplus 0^{n-1}1$ (or, alternatively, $f_2(x) = \mathbf{sh}(x) \oplus 0^{n-1}1$), where $\mathbf{sh}(x_1 \cdots x_n) = (x_2 \cdots x_n x_1)$ is a cyclic shift that corresponds to the relocation permutation $\pi(i) = (i \bmod n) + 1$. Then, the 4-regular 2^n -vertex graph that consists of the vertex set $\{0, 1\}^n$ and the edge multiset $\cup_{i \in [2]} \{x, f_i(x)\} : x \in \{0, 1\}^n$ has second eigenvalue $1 - \Theta(1/n^2)$.

(Indeed, in this graph, x is connected to $x \oplus 0^{n-1}1$ by two parallel edges, and the other pairs of edges (i.e., $\{x, \mathbf{sh}(x)\}$ and $\{x, \mathbf{sh}^{-1}(x)\}$ for each x) may also be non-distinct.)

Proof: We claim that taking a random walk of length $t = O(t' \cdot n^2)$ on this graph yields a distribution that is $2^{-t'}$ -close to uniform. The claim is proved by observing that during such a walk, with probability at least $1 - 2^{-t'}$, each position in the original string appeared at the rightmost position at some time during the walk (and that at the next step the corresponding value is randomized, since at that step f_2 is applied with probability one half).³ ■

The foregoing argument refers implicitly to a random walk on the n -vertex cycle, which represents the shift relocation permutation used in the 1-local 2^n -vertex graph that consists of the relocation permutation \mathbf{sh} and the offset $0^{n-1}1$. In general, we shall be discussing two graphs: The 2^n -vertex graph with transitions that are 1-local, and an n -vertex graph that describes the relocation permutations used in the 1-local graph.

Definition 3 (a generic 1-local graph and the corresponding relocation graph): Let $\pi^{(1)}, \dots, \pi^{(d)} : [n] \rightarrow [n]$ be d permutations and $s^{(1)}, \dots, s^{(d)} \in \{0, 1\}^n$.

1. The 1-local graph associated with $\pi^{(1)}, \dots, \pi^{(d)}$ and $s^{(1)}, \dots, s^{(d)}$ is the $2d$ -regular 2^n -vertex graph that consists of the vertex set $\{0, 1\}^n$ and the edge multi-set $\cup_{i \in [d]} \{x, x_{\pi^{(i)}} \oplus s^{(i)}\} : x \in \{0, 1\}^n$, where $x_\pi = x_{\pi(1)} \cdots x_{\pi(n)}$.
2. The relocation graph associated with $\pi^{(1)}, \dots, \pi^{(d)}$ is the $2d$ -regular n -vertex graph that consists of the vertex set $[n]$ and the edge multi-set $\cup_{i \in [d]} \{j, \pi^{(i)}(j)\} : j \in [n]$.

The mapping $x \mapsto x_{\pi^{(i)}} \oplus s^{(i)}$ (resp., $j \mapsto \pi^{(i)}(j)$) is called a forward transition, whereas the reverse mapping $y \mapsto (y \oplus s^{(i)})_{\pi^{(-i)}}$ (resp., $k \mapsto \pi^{(-i)}(k)$) is called a reverse transition, where $\pi^{(-i)}$ denotes the inverse of $\pi^{(i)}$.

Wishing to use shorter random walks in the rate-convergence analysis, consider the case that the n -vertex relocation graph is a $O(1)$ -regular expander graph.⁴ In this case, a random walk of length

²Note that when taking an n -step random walk on the 2-regular directed graph in which edges are directed from each vertex x to the vertices $\mathbf{sh}(x)$ and $\mathbf{sh}(x) \oplus 0^{n-1}1$ the final vertex is uniformly distributed (regardless of the start vertex). However, there is a fundamental difference between random walks on directed graphs and random walks on the underlying undirected graphs.

³The location of the j^{th} bit in the original string after i steps is determined by $j + \sum_{k \in [i]} X_k \bmod n$, where the X_k 's are the $\{0, \pm 1\}$ -indicators of the chosen transitions (i.e., $X_k = 1$ (resp. $X_k = -1$) if the transition \mathbf{sh} (resp., \mathbf{sh}^{-1}) was taken in the k^{th} step and $X_k = 0$ otherwise (i.e., if the offset $0^{n-1}1$ was applied)). Note that each block of $O(n^2)$ symbols has absolute value of at least $2n$ with probability at least $1/2$. Hence, looking at t' partial sums that correspond to t' such disjoint blocks, we observe that the probability that all these partial sums are in the interval $[-n, n]$ is at most $2^{-t'}$. Finally, note that if any of these partial sums has value outside $[-n, n]$, then at the corresponding $O(n^2)$ steps each original bit position appeared in the rightmost location.

⁴We assume that the edges of this $2d$ -regular expander can be represented by d permutations, as in the foregoing definition of a relocation graph.

$t = O(t' \cdot n \log n)$ on the n -vertex graph visits all vertices with probability at least $1 - 2^{-t'}$ (since its cover time is $O(n \log n)$ and we have t' “covering attempts”).⁵ It follows that the corresponding 1-local 2^n -vertex graph (in which half of the edges use the corresponding relocation permutations and the other half use the offset $0^{n-1}1$) has second eigenvalue $1 - \Omega(1/n \log n)$. This is the case because taking a random walk of length $O(t' \cdot n \log n)$ on the 1-local graph yields a distribution that is $2^{-t'}$ -close to uniform, since (with probability $1 - 2^{-t'}$) each position in the original n -bit string is mapped to the rightmost position at some time, and at the next step the corresponding value is “randomized” (since the offset is applied with probability $1/2$).

We observe that there is no hope of getting a constant-degree 2^n -vertex expander when using only offsets of Hamming weight $o(n)$. This is the case because the probability that a walk of length t on any regular n -vertex graph misses a set of $o(n)$ vertices is at least $(1 - o(1))^t = \exp(-o(t))$.⁶ In that case, there exists a position in the original n -bit string (i.e., in the label of the vertex of the 1-local 2^n -vertex graph) that is not moved to an active location where it may be randomized, where the active locations refer to the 1-entries in the offsets. This rules out not only the proof strategy used above, but also the possibility that the 2^n -vertex graph is an expander.

Observation 4 (using only light offsets can not yield an expander): *Consider a $2d$ -regular 2^n -vertex graph as in Definition 3, and suppose that $|s^{(i)}| = o(n)$ for all $i \in [d]$. Then, this 1-local 2^n -vertex graph is not an expander.*

Proof: For starters, consider an auxiliary $4d$ -regular 2^n -vertex graph in which, for each $i \in [d]$, the i^{th} relocation permutation (i.e., $\pi^{(i)}$) is coupled both with the offset $s^{(i)}$ and with the all-zero offset. Now, for a t -step random walk (on this 2^n -vertex graph) that starts at the vertex 0^n , consider the event this walk does not randomize position $j \in [n]$ (in the initial n -bit string); that is, the corresponding walk on the n -vertex graph that starts at vertex $j \in [n]$ does not go through any vertex in the set $\cup_{i \in [d]} \{k : s_k^{(i)} = 1\}$. This event occurs with probability at least $\eta = \exp(-o(t))/n$, since (as argued above) a t -step random walk that starts at the uniform probability misses the said set with probability at least $\exp(-o(t))$.

Note that randomized bit positions are reset to 1 with probability exactly $1/2$ (by virtue of the auxiliary construction performed upfront), whereas non-randomized positions maintain the value 0. Considering the expected number of ones in the label of the final vertex of a t -step random walk (on the 2^n -vertex graph), observe that if some bit is not randomized with probability η , then the expected number of ones is at most $(1 - \eta) \cdot 0.5n + \eta \cdot 0.5(n - 1) = (n - \eta)/2$. It follows that the distribution of the final vertex is $\eta/2n$ -far from uniform⁷, which implies that the convergence rate

⁵The cover time bound was established in [1, 2, 5].

⁶Note that here we seek a lower bound on the probability of missing the set S (equiv., staying in $\bar{S} = [n] \setminus S$), whereas the common focus is on good upper bounds (which exists when the graph is an expander). Letting d denote the degree of the n -vertex graph, we observe that there are at most $d \cdot |S|$ edges incident at S , and the worst case is that their other endpoints are distributed evenly among the vertices in \bar{S} (because otherwise, conditioning on not leaving \bar{S} biases the distribution towards vertices that have more neighbors in \bar{S} (equiv., less neighbors in S)). Hence, the probability that the random walk never leaves \bar{S} is at least $(1 - \frac{d|S|}{d|\bar{S}|})^t$, whereas in our case $|\bar{S}| = (1 - o(1)) \cdot n$.

⁷We use the fact that if $E[X] < E[Y] - \epsilon$ and $X, Y \in [0, 1]$, then there exists a set of values S such that $\Pr[X \in S] < \Pr[Y \in S] - \epsilon$. This can be proved by taking $S = \{v : \Pr[X = v] < \Pr[Y = v]\} \subseteq [0, 1]$ and using

$$\Pr[Y \in S] - \Pr[X \in S] = \sum_{v \in S} (\Pr[Y = v] - \Pr[X = v]) \geq \sum_{v \in [0, 1]} (\Pr[Y = v] - \Pr[X = v]) \cdot v = E[Y] - E[X] > \epsilon.$$

is not bounded away from 1 (since $(2^{-n} \cdot \eta/2n)^{1/t} \approx \exp(-o(1) - (n/t)) = \exp(-o(1))$ for sufficiently large t).⁸ Since the auxiliary graph is not an expander, the original graph (which is a subgraph of it) is also not an expander. ■

We note that using also offsets of Hamming weight $n - o(n)$ does not help, since this is equivalent to adding the all-ones offset, which merely complements the vertex label in the 2^n -vertex graph.⁹ In view of the above, we must use at least one offset that has Hamming weight in $[\Omega(n), n - \Omega(n)]$.

3 A sufficient condition

We now identify a property of 1-local 2^n -vertex $O(1)$ -regular graphs that suffices for showing that they are expanders. For simplicity, we consider the case of using a single non-zero offset $s \in \{0, 1\}^n$ (along with the offsets that are derived from it when considering also the reverse transitions). Actually, for each relocation permutation $\pi : [n] \rightarrow [n]$, we consider the four transitions $x \mapsto (x \oplus s^b)_\pi \oplus s^c$, where $b, c \in \{0, 1\}$ and $s^0 = 0^n$ (and $s^1 = s$). (Note that such a generic transition can be viewed as $x \mapsto x_\pi \oplus (s_\pi)^b \oplus s^c$, and that the reverse transition has the form $y \mapsto (y \oplus s^c)_{\pi^{-1}} \oplus s^b = y_{\pi^{-1}} \oplus (s_{\pi^{-1}})^c \oplus s^b$.)¹⁰ In other words, referring to Definition 3 and assuming that d is a multiple of 4, we postulate that for some $s \in \{0, 1\}^n \setminus \{0^n\}$ and every $i \in [d/4]$ and $b, c \in \{0, 1\}$ it holds that $\pi^{(4i-2b-c)} = \pi^{(4i)}$ and $s^{(4i-2b-c)} = (s_{\pi^{(4i)}})^b \oplus s^c$. Note that in this case, for every i , taking at random one of the four corresponding (forward) transitions has the effect of randomizing the vertex label by the offset s (by virtue of the random value of $c \in \{0, 1\}$), and the same holds when taking the reverse transition (by virtue of the random value of $b \in \{0, 1\}$). When taking a random walk on this graph, we consider only the randomizing effect of this offset (i.e., of the choice of c in a forward move, and the choice of b in a reverse move).¹¹

To clarify the above and motivate the following property, suppose that we take $t = \Omega(n)$ random steps on the 1-local graph, and consider the t -by- n Boolean matrix describing the activity status of the location to which each of the initial positions is moved during the t steps, where an initial position is said to be active if it currently reside in location in $\{k : s_k = 1\}$. That is, the (i, j) th entry in the matrix indicate whether or not, in the i th step of the fixed random walk being considered, the j th initial location is mapped to an active location (i.e., a 1-entry in the offset s being used). Using an n -vertex expander and s of weight approximately $n/2$, we observe that (w.v.h.p.) each column in this random matrix has approximately $t/2$ 1-entries, but as we shall see what we need is that (w.v.h.p.) this matrix has full rank.

Note that the foregoing matrix, which is defined based on a random walk on the 1-local 2^n -vertex graph, describes n coordinated walks on the n -vertex relocation graph, each starting at a different

⁸Since the convergence rate λ must satisfy $2^n \cdot \lambda^t > \eta/2n$.

⁹In that case, with similar probability, there are two positions in the original string that are not moved through an active location (which implies that their final values are identical). To see this, follow the argument in Footnote 6, while noting that the probability that one of the two coordinated random walks does not stay in \bar{S} is only doubled. The argument is completed by considering the expected number of pairs of positions that hold the same value.

¹⁰In contrast, if we were only to use the transitions $x \mapsto x_\pi \oplus s^c$, then the reverse transitions would have had the form $y \mapsto (y \oplus s^c)_{\pi^{-1}} = y_{\pi^{-1}} \oplus (s_{\pi^{-1}})^c$, which would have hindered the argument that follows.

¹¹If we are currently at vertex x and take the forward transition associated with (π, b, c) , then we move to vertex $x_\pi \oplus (s_\pi)^b \oplus s^c$, and the foregoing randomization effect refers to the addition of the offset s (to $(x \oplus s^b)_\pi$), which occurs if and only if $c = 1$. Likewise, if we are currently at vertex y and take the reverse transition associated with (π, b, c) , then we move to vertex $(y \oplus s^c)_{\pi^{-1}} \oplus s^b$, and the foregoing randomization effect refers to the addition of the offset s (to $(y \oplus s^c)_{\pi^{-1}}$), which occurs if and only if $b = 1$.

vertex of the graph and all proceeding according to the same sequence of (random) choices. (Note that each step on the n -vertex relocation graph, which has degree $2d/4$, only determines $i \in [d/4]$ and the direction of the transition (i.e., forward or backward), while leaving the choice of the corresponding bits b, c unspecified.) When the foregoing t -by- n matrix has full rank, the t random choices of whether to apply the offset s correspond to a random linear combination of the t rows of the matrix, which yields a uniformly distributed n -bit long string. In this case, the corresponding random walk on the 2^n -vertex graph yields a uniform distribution (since the latter n -bit string is added to the initial vertex in the walk).¹² This motivates the definition of the following property.

Definition 5 (a property of coordinated random walks):¹³ For $d = O(1)$, consider a d -regular n -vertex graph such that for every $\sigma \in [d]$ the function $g_\sigma : [n] \rightarrow [n]$ that maps each vertex to its σ^{th} neighbor is a bijection. For a set $T \subseteq [n]$ and an integer $t = \Omega(n)$, consider a random sequence $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in [d]^t$ and the n corresponding coordinate random walks (CRW) such that the j^{th} walk starts at vertex j and moves in the i^{th} step to the σ_i^{th} neighbor of the current vertex, and consider a t -by- n Boolean matrix $B^{(\bar{\sigma})}$ such that its $(i, j)^{\text{th}}$ entry indicates whether the j^{th} walk passed in T in its i^{th} step; that is, the $(i, j)^{\text{th}}$ is 1 if and only if $g_{\sigma_i}(\dots(g_{\sigma_1}(j)\dots)) \in T$. The desired CRW property is that, with probability at least $1 - \exp(-\Omega(t))$ over the choice of $\bar{\sigma} \in [d]^t$, the matrix $B^{(\bar{\sigma})}$ has full rank (over $\text{GF}(2)$).

We have already noted that for this property to hold, the set T must have size in $[\Omega(n), n - \Omega(n)]$. We now note that using an arbitrary expander graph and an arbitrary set T of any predetermined size (e.g., $|T| \approx n/2$) will not do: For example, consider an n -vertex expander that consists of two $n/2$ -vertex expanders that are connected by a matching, and let T be the set of vertices in one of these two expanders. Then, coordinated walks on this graph (w.r.t this T) always yields a Boolean matrix of rank at most two, since the coordinated walks that start at vertices in T (resp., in $[n] \setminus T$) always move together to T or to $[n] \setminus T$. Hence, the question we consider is the following.

Problem 6 (the CRW problem): For which graphs and which sets T 's does the property in Definition 5 hold?

As outlined above, any $2d$ -regular relocation graph that satisfies this property yields an $8d$ -regular 1-local 2^n -vertex expander.

Theorem 7 (solutions to the CRW problem yield 1-local expanders): Let $\pi^{(1)}, \dots, \pi^{(d)} : [n] \rightarrow [n]$ be d permutations and $s \in \{0, 1\}^n$. If the $2d$ -regular n -vertex graph with the edge multi-set $\cup_{i \in [d]} \{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ along with the set $\{j \in [n] : s_j = 1\}$ satisfies the CRW property, then the $8d$ -regular 2^n -vertex graph with the edge multi-set $\cup_{i \in [d], b, c \in \{0, 1\}} \{\{x, (x \oplus s^b)_{\pi^{(i)}} \oplus s^c\} : x \in \{0, 1\}^n\}$ is an expander.

Proof: By the hypothesis, for some $t = \Omega(n)$, the relevant t -by- n matrix has full rank with probability at least $1 - \exp(-\Omega(t))$. Fixing an arbitrary walk $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in [2d]^t$ on the n -vertex relocation graph such that $B^{(\bar{\sigma})}$ has full rank, for each $i \in [t]$, we consider the corresponding

¹²That is, fixing a random walk on the n -vertex relocation graph, we observe that if the matrix that corresponds to this walk has full rank, then the final vertex in the corresponding random walk on the 1-local 2^n -vertex graph is uniformly distributed in $\{0, 1\}^n$, regardless of the effect of the relocation permutations.

¹³An alternative way of defining the matrix $B^{(\bar{\sigma})}$ consists of considering a sequence of permutations over $[n]$, denoted $\pi_0, \pi_1, \dots, \pi_t$, such that π_0 is the identity permutation, and $\pi_i(j) = g_{\sigma_i}(\pi_{i-1}(j))$. The i^{th} row of $B^{(\bar{\sigma})}$ is then defined as the T -indicator of π_i ; that is, the $(i, j)^{\text{th}}$ entry in the matrix is 1 if and only if $\pi_i(j) \in T$.

choices of the $b_i, c_i \in \{0, 1\}$ for the i^{th} step of the corresponding random walk on the 2^n -vertex graph. Specifically, we consider a random process that selects these bits uniformly, in two steps.

- In the first step, for every $i \in [t]$, if the i^{th} transition is in the forward direction, we select b_i at random, otherwise we select c_i at random.
- In the second step, we make the remaining choices; that is, for every $i \in [t]$, if the i^{th} transition is in the forward direction, we select c_i at random, otherwise we select b_i at random.

Fixing any sequence of choices for the first step, the label of the final vertex is a random variable that depends only on the random choices made in the second step, but such random choices have the effect of randomizing the vertex label by adding to it a corresponding linear combination of the rows of the matrix $B^{(\bar{\sigma})}$. Specifically, row i is taken to this linear combination if and only if the relevant c_i or b_i equals 1 (where for a forward direction c_i determines whether the current label is offset by s , and for the reverse direction this choice is determined by b_i).¹⁴ Hence, in this case, the corresponding random walk on the 2^n -vertex graph yields a uniform distribution (regardless of the start vertex). It follows that the distribution of the label of the final vertex is $\exp(-\Omega(t))$ -close to the uniform distribution, which implies that the convergence rate of the 2^n -vertex graph is bounded away from 1 (i.e., this 1-local graph is an expander). ■

4 Known constructions that satisfy the CRW property

Recall that Kassabov's result [4], which is used in [6], asserts that the symmetric group has an explicit generating set that is expanding and of constant size.¹⁵ We shall show that using this set of permutations (i.e., as our set of relocating permutations) along with the set $[n']$ such that $n' \approx n/2$ is odd (e.g., odd $n' \in \{\lfloor n/2 \rfloor, \lfloor n/2 \rfloor + 1\}$) yields an n -vertex graph that satisfies the coordinated random walks property (of Definition 5). This yields an alternative proof of Theorem 1.

Theorem 8 (a positive answer to Problem 6): *Let $\Pi = \{\pi^{(i)} : i \in [d]\}$ be a generating set of the symmetric group of n elements and suppose that Π is expanding.¹⁶ Then, the n -vertex graph that consists of the vertex set $[n]$ and the edge multi-set $\cup_{i \in [d]} \{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ combined with any set of odd size $n' \approx n/2$ satisfies the coordinated random walks property of Definition 5.*

Proof: For a sufficiently large t , consider a random t -by- n Boolean matrix that corresponds to coordinated random walks (from all possible start vertices) on the n -vertex graph (wrt the foregoing

¹⁴Formally, denoting the initial vertex in the walk on the 1-local graph by v_0 , the i^{th} vertex in the walk, denoted v_i , satisfies $v_i = (v_{i-1} \oplus s^{b_i})_{\pi} \oplus s^{c_i}$ (resp., $v_i = (v_{i-1} \oplus s^{c_i})_{\pi^{-1}} \oplus s^{b_i}$) if σ_i indicates a forward (resp., reverse) transition according to π . Denoting by π_i the relocation permutation applied in the i^{th} step of the walk (i.e., $\pi_i = \pi$ (resp., $\pi_i = \pi^{-1}$) if σ_i indicates a forward (resp., reverse) transition according to π), note that $v_i = (v_{i-1})_{\pi_i} \oplus (s_{\pi_i})^{x_i} \oplus s^{y_i}$, where $(x_i, y_i) = (b_i, c_i)$ if the i^{th} step takes a forward transition and $(x_i, y_i) = (c_i, b_i)$ otherwise. In both cases, the i^{th} row in the matrix, denoted r_i , equals $s_{(\pi_i \circ \dots \circ \pi_1)^{-1}}$, where $\pi_i \circ \dots \circ \pi_1$ is the composition of the relocation permutations applied in the i first steps. Hence, $(v_i)_{(\pi_i \circ \dots \circ \pi_1)^{-1}} = (v_{i-1})_{(\pi_{i-1} \circ \dots \circ \pi_1)^{-1}} \oplus (s_{(\pi_{i-1} \circ \dots \circ \pi_1)^{-1}})^{x_i} \oplus r_i^{y_i}$. It follows that $(v_i)_{(\pi_i \circ \dots \circ \pi_1)^{-1}} = v_0 \oplus \bar{s} \oplus \bigoplus_{i \in [t]} r_i^{y_i}$, where $\bar{s} = \bigoplus_{i \in [t]} (s_{(\pi_{i-1} \circ \dots \circ \pi_1)^{-1}})^{x_i}$.

¹⁵Indeed, this refers to a third graph, which is the corresponding Cayley graph with $n!$ vertices (i.e., the vertices are all the possible permutations over $[n]$).

¹⁶That is, letting Sym_n denote the symmetric group of n elements, we consider the Cayley graph consisting of the vertex set Sym_n and the edge multi-set $\cup_{i \in [d]} \{\{\pi, \pi^{(i)} \circ \pi\} : \pi \in \text{Sym}_n\}$, where \circ denote composition of permutations. The hypothesis postulates that this Cayley graph is an expander.

set T of size n'). We shall show that, for every non-empty set $J \subseteq [n]$, with probability at least $1 - \exp(-\Omega(t) + O(n \log n))$, the sum of columns in positions J is non-zero. (This establishes CRW property for any sufficiently large $t = \Omega(n \log n)$.)¹⁷

Claim 8.1 (the distribution of a specific linear combination of the columns): *For every non-empty set $J \subseteq [n]$, with probability at least $1 - \exp(-\Omega(t) + O(n \log n))$ over the matrix that corresponds to a random walk, the sum (mod 2) of the matrix's columns in positions J is non-zero.*

Proof: For $J = [n]$ this follows from the fact that n' is odd. Otherwise (i.e., for $J \subset [n]$), we shall prove the claim by using the correspondence between random walks on the n -vertex graph and random walks on the set of all permutations where in a random step the current permutation is composed with the selected generator.¹⁸ That is, selecting the σ^{th} neighbor in the random walk on the n -vertex graph, a choice that determines a transition (i.e., $\lceil \sigma/2 \rceil \in [d]$) as well as the direction (i.e., forward or reverse) in which the transition is applied, corresponds to selecting the $\lceil \sigma/2 \rceil^{\text{th}}$ generating permutation and moving by composing it or its inverse (according to the value of $\sigma \bmod 2$).

In our argument, we shall refer to a set of permutations over $[n]$, denoted Sym_n , and consider the set of permutation, denoted W , consisting of permutations having an J -image that contains an odd number of elements of T ; that is, $\pi \in W$ if and only if $|\{j \in J : \pi(j) \in T\}|$ is odd. The claim will follow by showing that (1) $|W| \approx |\text{Sym}_n|/2$, and (2) a random walk on the relocation graph corresponds to a matrix with columns in positions J summing up to the all-zero vector if and only if the random walk on Sym_n does not visit W .

We first show that W has density approximately half within the set of all $n!$ permutations over $[n]$. This can be shown by considering, w.l.o.g., the case of $|J| \leq n/2$ (or else consider $[n] \setminus J$). To estimate the probability that a random permutation is in W , consider the process of selecting uniformly $\pi \in \text{Sym}_n$ by randomly assigning distinct elements to the location in J , and ignore the residual assignment of elements to $[n] \setminus J$. Now, focus on the last assignment in that process (i.e., the assignment of the $|J|^{\text{th}}$ element). Using the hypothesis that $|T| = n' \approx n/2$, with probability $1 - o(1)$, before this last assignment, the previous $|J| - 1 < n/2 \approx n'$ locations were assigned approximately an equal number of elements from T and from $[n] \setminus T$, which means that $n' - (1 \pm o(1)) \cdot |J|/2 = (1 \pm o(1)) \cdot (n - |J|)/2$ elements of each type remain for the last assignment, where $n - |J| = \Omega(n)$. This implies that the parity of elements from T is flipped at the last step with probability $(1 \pm o(1))/2 \approx 1/2$.

The key observation is that the coordinated random walks on the n -vertex graph yield a Boolean matrix such that the sum of columns in positions J is zero (mod 2) if and only if the corresponding walk on the set of $n!$ permutations does not pass through states in W , where the latter walk starts at the identity permutation. (To see this, consider the sequence, denoted π_1, \dots, π_t , of permutations that are selected during the random walk (as determined by the sequence $\bar{\sigma} \in [2d]^t$). Then, in the i^{th} step of the coordinated walks, the j^{th} walk is in position $k = \pi_i(\dots(\pi_1(j)\dots))$, whereas the $(i, j)^{\text{th}}$ entry in the matrix is 1 if and only if $\pi_i(\dots(\pi_1(j)\dots)) \in T$ (i.e., if and only if $k \in T$). Hence, the sum of the entries in row i and columns in J is one (mod 2) if and only if $\pi_i \circ \dots \circ \pi_1 \in W$.)

Finally, consider a t -step random walk on the set of permutations that starts at the identity permutation. By the expansion property of the generating set for the symmetric group,

¹⁷We comment that the CRW property can be established for any sufficiently large $t = \Omega(n)$; see Claim 8.2..

¹⁸That is, we use the correspondence between (coordinated) random walks on the n -vertex graph and random walks on the $n!$ -vertex Cayley graph.

the probability that this walk does not pass through a fixed set of constant density is at most $\exp(-\Omega(t - O(n \log n)))$, where the first $O(n \log n)$ steps are taken for convergence to the uniform distribution and the remaining steps are used for hitting attempts. ■

Using a union bound (over all non-empty sets J), we conclude that, with probability at least $1 - (2^n - 1) \cdot \exp(-\Omega(t) + O(n \log n))$, the corresponding t -by- n Boolean matrix has full rank. Taking $t = \Omega(n \log n)$, the theorem follows. ■

Conclusion: Indeed, as stated upfront, applying Theorem 7 to the n -vertex graph (and set) analyzed in Theorem 8 (and using [4]) yields an alternative proof of Theorem 1.

For sake of elegance: As noted in Footnote 17, the bound of Claim 8.1 can be tightened.

Claim 8.2 (the distribution of linear combinations of the columns, revisited): *With probability at least $1 - \exp(-\Omega(t) + O(n))$ over the matrix that corresponds to a random walk, for every non-empty set $J \subseteq [n]$, the sum of the matrix's columns in positions J is non-zero.*

Proof: We proceed as in the proof of Claim 8.1, but consider random walks (on the set of all permutations) that start at a state that is uniformly distributed in a specific set S (rather than start at the identity permutation). The set S is the set of all permutations such that each location in T holds an element of T ; that is, $\pi \in S$ if and only if $\{i \in T : \pi(i)\} = T$. Using $|T| = n' \approx n - |T|$, observe that S has density approximately $\frac{(n')^2}{n!}$, which is approximately 2^{-n} .

Note that the Boolean matrix that represents a random walk on the n -vertex graph equals (up to a permutation of its columns) the matrix that represents the same walk on any isomorphic copy of that graph that leaves T invariant (i.e., rather than walking on an n -vertex graph $G = ([n], E)$, we walk on its isomorphic copy $\phi(G) = ([n], \{\{\phi(i), \phi(j)\} : \{i, j\} \in E\})$, where $\phi : [n] \rightarrow [n]$ is a permutation such that $\phi(j) \in T$ for every $j \in T$). That is, if the matrix M represents a random walk on the original graph and $\phi : [n] \rightarrow [n]$ is a permutation that leaves T invariant, then the matrix obtained by permuting the columns of M according to ϕ represents a random walk on the isomorphic copy of the original graph obtained by relabeling its vertices according to ϕ . (This is the case because the j^{th} column in M indicates whether the walk on G that starts at vertex j hits T in each of the t steps, but this column also indicates whether the same walk on $\phi(G)$ that starts at $\phi(j)$ hits $\phi(T) = T$ in each of the t steps.) Now, since M is full rank if and only if permuting its columns yields a full rank matrix, we may consider random walks on such random isomorphic copies of the original graph (i.e., copies obtained by relabeling it using a random permutation that leaves T invariant). Hence, we may analyze the corresponding walk (on the set of $n!$ permutations) that starts at a state that is uniformly distributed in S (rather than starting at the identity permutation).

Now, for every non-empty $J \subset [n]$, we consider the corresponding set W_J (as defined in the proof of Claim 8.1). By the expansion property of the generating set for the symmetric group, we have that a t -step random walk that starts in uniformly distributed state in S passes via W_J with probability at least $1 - \exp(-\Omega(t - O(n)))$, where the first $O(n)$ steps are taken for convergence to the uniform distribution and the remaining steps are used for hitting W_J . Hence, with probability at least $1 - (2^n - 1) \cdot \exp(-\Omega(t - O(n)))$, a random walk that starts at a state that is uniformly distributed in S avoids none of the W_J 's. In this case, for every non-empty set $J \subseteq [n]$, the sum of columns (of the corresponding matrix) in positions J is non-zero. ■

The CRW property does not imply that the set of relocations is an expanding set of generators for Sym_n . Interpreted in terms of sets of permutations over $[n]$, the CRW property asserts that a random walk on the corresponding Cayley graph passes a specific statistical test (which is specified by the corresponding set T). Theorem 8 asserts that if this Cayley graph is an expander, then the CRW property is satisfied for any set T of odd size $n' \approx n/2$. (One may observe that this holds also if $n' \approx n/4$.) This implies that if this Cayley graph is an expander, then the CRW property is satisfied for some set T . Here we show that the converse of the latter conditional statement does not hold.¹⁹

Theorem 9 (on the converse of Theorem 8): *There exists a set of permutations, $\{\pi^{(i)} : i \in [3d]\}$, over $[2n]$ that does not generate the symmetric group of $2n$ elements (let alone in an expanding manner) such that the $2n$ -vertex graph consisting of the vertex set $[2n]$ and the edge multi-set $\cup_{i \in [3d]} \{j, \pi^{(i)}(j)\} : j \in [2n]$ combined with some set of odd size $n' \approx n/2$ satisfies the coordinated random walks property of Definition 5.*

Proof Sketch: We start with a set of permutations $\Pi = \{\pi^{(i)} : i \in [d]\}$ that generates the symmetric group of n elements and is expanding. We first extend each $\pi^{(i)} \in \Pi$ to the domain $[2n]$ such that $\pi^{(i)}(n+j) = n + \pi^{(i)}(j)$ for every $j \in [n]$ (and $i \in [d]$). Next we add d copies of the identity permutation and d copies of the permutation that switches $[n]$ and $[2n] \setminus [n]$; that is, for every $i \in [d]$, we have $\pi^{(d+i)}(b \cdot n + j) = b \cdot n + j$ and $\pi^{(2d+i)}(b \cdot n + j) = (1-b) \cdot n + j$ for every $j \in [n]$ and $b \in \{0, 1\}$. The $2n$ -vertex $6d$ -regular relocation graph G' that corresponds to the augmented set of permutations Π' consists of two copies of the $2d$ -regular n -vertex graph G that corresponds to Π augmented by d self-loops on each vertex (where each self-loop contributing two units to the vertex's degree) and $2d$ copies of a perfect matching that matches the two copies of each original vertex.

Note that Π' (i.e., the new set of permutations over $[2n]$) does not generate the symmetric group of $2n$ elements; it rather generates a group of $2 \cdot (n!) \ll (2n)!$ permutations. The theorem follows by showing that the ($2n$ -vertex) relocation graph G' satisfies the CRW property (with any set $T \subset [n]$ of odd size $n' \approx n/2$).²⁰ When analyzing random walks on G' , we distinguish steps in which one of the first d permutations is employed from steps in which one of the last $2d$ permutations is employed, calling the latter steps semi-idle, since they either map each vertex to itself or map each vertex to its sibling (i.e., the other copy). The key observation is that a t -step random walk has the following two properties:

1. With probability at least $1 - \exp(-\Omega(t))$, at least $t/2$ of the steps in the walk are semi-idle, since at each time a semi-idle step is selected with probability $2/3$. Furthermore, for every constant c , with probability at least $1 - \exp(-\Omega(t))$, there are at most t/c^2 time-intervals of the form $[(i-1) \cdot c + 1, i \cdot c]$ that do not contain a semi-idle step.
2. Fixing a (non-empty) set $J \subseteq [n]$ and a random walk on G' , we call a vertex **good (for J)** if the sum of the columns $J \cup \{n+j : j \in J\}$ in the matrix that corresponds to the random

¹⁹Indeed, we leave open the possibility that the converse of Theorem 8 holds. We believe that even if the CRW property is satisfied for any set T of odd size $n' \approx n/2$, then it does not necessarily hold that the foregoing Cayley graph is an expander.

²⁰We stress that T is an arbitrary subset of size n' of $[n]$, whereas the vertex set is $[2n]$. Indeed, picking T of size n' arbitrarily in $[2n]$ will fail (e.g., if $T = T' \cup (n+T') \cup \{n\}$, for any $T' \subseteq [n-1]$, then, for every non-empty $J' \subseteq [n]$, the sum of matrix's columns with indices in $J' \cup (n+J')$ is exactly as in the case of $T = \{n\}$, since the contributions of T' and $n+T'$ cancel out).

walk yields a vector that has at least $\Omega(t)$ ones in rows that correspond to steps that are not semi-idle. (Recall that the $(i, j)^{\text{th}}$ entry in this matrix indicates whether the j^{th} walk hits T in its i^{th} step, and that we consider here the parity of the entries in the columns $J \cup \{n + j : j \in J\}$ restricted to rows that are not semi-idle.) The observation is that *every non-empty set $J \subseteq [n]$, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$ over the random walks on G' , it holds that J is good.*

The foregoing property of the walk on (the $2n$ -vertex graph) G' is proven by considering an analogous lazy walk on (the n -vertex) graph G such that the walk on G takes an idle step (i.e., stays in place) if and only if the walk on G' takes a semi-idle step. First note that each row of the foregoing matrix (which corresponds to a walk on G') has 1-entries either in columns in $[n]$ or in columns in $[2n] \setminus [n]$; hence the sum of the entries in columns $J \cup \{n + j : j \in J\}$ (in this row) is either due to J or to $n + J \stackrel{\text{def}}{=} \{n + j : j \in J\}$. Recall that Claim 8.1 asserts that, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$, at least one row in the matrix that corresponds to a non-lazy walk on G has an odd sum of the entries in columns J , but the argument extends to the matrix that corresponds to a lazy random walk (on G) and to rows that do not correspond to idle steps.²¹ Hence, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$, at least one of the rows in a matrix that corresponds to a random walk on G' is good (for J). Furthermore, the same argument extends to showing that the number of good (for J) rows is $\Omega(t)$, by using a ‘‘Chernoff bound’’ for random walks on expanders.²²

(Note that we use the correspondence between columns $J \cup (n + J)$ in the matrix B' that describes a random walk on G' and columns J in the matrix B that describes a corresponding lazy random walk on G . The hypothesis was used to analyze the random walks on G , and the inference to G' holds because each row in B appears either in the first n entries of the corresponding row in B' or in its last n entries (while the other n entries of this row are all zero).)

Hence, a random walk on G' satisfies both the foregoing properties with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$. Fixing such a walk on G' , note that satisfying these properties is independent of the choices made in the semi-idle steps of this walk (i.e., whether to stay idle or move to the sibling vertex). This fact will allow us to re-randomize these choices later.

By the choice of this walk, the corresponding matrix has $t' = \Omega(t)$ good rows. Letting $c = 10t/t'$, we infer that for at most $2c \cdot t/c^2 = 0.2t'$ of the good rows i there is no semi-idle step in the time-interval $[i - c, i - 1]$. Hence, there are at least $t'' = 0.8t'/c = \Omega(t)$ good (for J) rows that are at distance c apart from one another such that each such row i is preceded by a semi-idle step in the time-interval $[i - c, i - 1]$. Calling these rows **very good (for J)**, note that we can match each very good row i to a semi-idle step (in $[i - c, i - 1]$) that precedes row i but does not precede any very good row that precedes i . Now, if we re-randomize the choices made in all the semi-idle steps, then with probability at least $1 - 2^{-t''}$ one of the rows that is (very) good for J has all its ($|T|$) ones in the columns $[n]$. In this case, the sum of the columns in $J \subseteq [n]$ yields a non-zero vector. The same argument applies to $n + J$.

²¹This can be seen by mapping the lazy random walk to a non-lazy random walk.

²²Alternatively, we may just start with a graph G that satisfies a strong CRW property with respect to T , where this strong property asserts that *for sufficiently large t , with probability at least $1 - \exp(-\Omega(t))$, a t -step random walk yields a matrix with columns that span an n -dimensional linear code of distance $\Omega(t)$* (whereas the CRW property only asserts that the columns of this matrix span an n -dimensional linear space).

Recall that in order to argue that the matrix has full rank we have to show that for every non-empty $J' \cup J''$ such that $J' \subseteq [n]$ and $J'' \subset [2n] \setminus [n]$ it holds that the sum of the columns in $J' \cup J''$ yields a non-zero vector. Assuming, w.l.o.g., that J' is non-empty, we have already established that, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$, the sum of the columns in J' yields a non-zero vector. But the sum of the columns in J'' cannot cancel non-zero entries of that vector, since each row of the matrix has 1-entries either in columns $[n]$ or in columns $[2n] \setminus [n]$. Hence, with probability at least $1 - (2^{2n} - 1) \cdot \exp(-\Omega(t - O(n \log n)))$, the matrix has full rank. ■

5 A sufficient and necessary condition

Turning back to the relation between the CRW property (of Definition 5) and 1-local expander, we shall show that the following generalization of Definition 5 suffices and is necessary for obtaining a 1-local expander (with 2^n vertices).

Definition 10 (a relaxed property of coordinated random walks): *For $d, d' = O(1)$, consider a d -regular n -vertex graph as in Definition 5, and d' sets $T_1, \dots, T_{d'} \subseteq [n]$. As in Definition 5, for $t = \Omega(n)$, consider a random sequence $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in [d]^t$ and the n corresponding coordinate random walks such that the j^{th} walk starts at vertex j and moves in the i^{th} step to the σ_i^{th} neighbor of the current vertex. Now, fixing the random sequence $\bar{\sigma}$, consider an arbitrary sequence $\bar{\tau} = (\tau_1, \dots, \tau_t) \in [d']^t$, and let $B^{(\bar{\sigma}, \bar{\tau})}$ be the t -by- n Boolean matrix such that its $(i, j)^{\text{th}}$ entry indicates whether the j^{th} walk passed in T_{τ_i} in its i^{th} step. The relaxed CRW property asserts that, with probability at least $1 - \exp(-\Omega(t))$ over the choice of $\bar{\sigma} \in [d]^t$, there exists $\bar{\tau} \in [d']^t$ such that the Boolean matrix $B^{(\bar{\sigma}, \bar{\tau})}$ has full rank.*

(Indeed, Definition 5 corresponds to the special case of $d' = 1$.)

Theorem 11 (constructing 1-local expanders is equivalent to constructing relocation graphs along with sets that satisfy Definition 10): *Let $\pi^{(1)}, \dots, \pi^{(d)} : [n] \rightarrow [n]$ be permutations.*

1. *If the 1-local $2d$ -regular 2^n -vertex graph associated with $\pi^{(1)}, \dots, \pi^{(d)}$ and $s^{(1)}, \dots, s^{(d)} \in \{0, 1\}^n$ is an expander, then the corresponding $2d$ -regular n -vertex relocation graph along with the sets T_1, \dots, T_{2d} such that $T_{2i} = \{j \in [n] : s_j^{(i)} = 1\}$ and $T_{2i-1} = \{\pi^{(i)}(j) : s_j^{(i)} = 1\}$ satisfies Definition 10.*
2. *Suppose that the $2d$ -regular n -vertex relocation graph associated with $\pi^{(1)}, \dots, \pi^{(d)}$ along with the sets $T_1, \dots, T_{d'}$ satisfies Definition 10, and for every $\alpha \in \{0, 1\}^{d'}$ let $s^{(\alpha)} \in \{0, 1\}^n$ denote the indicator string of the set $\bigoplus_{i:\alpha_i=1} T_i \subseteq [n]$; that is, the j^{th} bit of $s^{(\alpha)}$ is 1 if and only if $|\{i \in [d'] : \alpha_i = 1 \ \& \ j \in T_i\}|$ is odd. Then, the $2^{2d'+1} \cdot d$ -regular 2^n -vertex graph with the edge multi-set $\cup_{i \in [d], \beta, \gamma \in \{0, 1\}^{d'}} \{x, (x \oplus s^{(\beta)})_{\pi^{(i)}} \oplus s^{(\gamma)}\} : x \in \{0, 1\}^n$ is an expander.*

Proof: We start with the proof of Part 2, which generalizes the proof of Theorem 7. Specifically, let $\bar{\sigma} = (\sigma_1, \dots, \sigma_t) \in [2d]^t$ be a random walk on the relocation graph such that an even σ_i (resp., an odd σ_i) indicates a forward (resp., reverse) transition using $\pi^{(\lceil \sigma_i / 2 \rceil)}$. Then, by the hypothesis, with probability at least $1 - \exp(-\Omega(t))$ over the choice of $\bar{\sigma}$, there exists $\bar{\tau} = (\tau_1, \dots, \tau_t)$ such that $B^{(\bar{\sigma}, \bar{\tau})}$ is full rank. When analyzing a corresponding random walk on the 1-local graph, consider the following process of determining the sequence of auxiliary random choices of $\beta_1, \dots, \beta_t \in \{0, 1\}^{d'}$ and $\gamma_1, \dots, \gamma_t \in \{0, 1\}^{d'}$.

1. For every i such that the i^{th} step is a forward (resp., reverse) transition,
 - (a) select β_i (resp., γ_i) uniformly in $\{0, 1\}^{d'}$, and,
 - (b) for every $k \in [d'] \setminus \{\tau_i\}$, select the bit $\gamma_{i,k}$ (resp., $\beta_{i,k}$) uniformly in $\{0, 1\}$.
2. For every i such that the i^{th} step is a forward (resp., reverse) transition, select γ_{i,τ_i} (resp., β_{i,τ_i}) uniformly in $\{0, 1\}$.

Fixing a good $\bar{\sigma}$ and a corresponding good $\bar{\tau}$ (i.e., choices such that $B^{(\sigma,\tau)}$ is full rank), consider an arbitrary fixing of the choices in Step 1. Then, the label of the final vertex in the corresponding random walk on the 1-local graph is a fixed string that is offset by a random linear combination of the rows of $B^{(\bar{\sigma},\bar{\tau})}$, where the random linear combination is determined in Step 2. (Specifically, if the i^{th} step is a forward (resp., reverse) transition, then the i^{th} row is included in this offset if and only if $\gamma_{i,\tau_i} = 1$ (resp., $\beta_{i,\tau_i} = 1$.) Thus, when $B^{(\bar{\sigma},\bar{\tau})}$ has full rank, the label of the final vertex is uniformly distributed in $\{0, 1\}^n$, and Part 2 follows.

Turning to the proof of Part 1, we start by considering the $4d$ -regular 2^n -vertex 1-local expander obtained from the given $2d$ -regular 1-local expander by augmenting each transition of the form $x \mapsto x_\pi \oplus s$ with the transition $x \mapsto x_\pi$. (The auxiliary graph is an expander because it contains an expander as a subgraph.) Hence, a step on this auxiliary graph is specified by a pair $(\sigma, b) \in [2d] \times \{0, 1\}$, where σ specifies a step on the original 1-local graph and b specifies whether the original offset is applied (i.e., we shall refer to the edge multi-set $\cup_{i \in [d], b \in \{0, 1\}} \{x, x_{\pi^{(i)}} \oplus (s^{(i)})^b\} : x \in \{0, 1\}^n$). Consequently, a t -step random walk on the $4d$ -regular expander corresponds to a sequence $(\sigma_1, b_1), \dots, (\sigma_t, b_t) \in ([2d] \times \{0, 1\})^t$, and the sequence $\sigma_1, \dots, \sigma_t$ corresponds to a walk on the n -vertex relocation graph. Determining the τ_i 's based on the σ_i 's yields a matrix as in Definition 10, Specifically, we shall determine the τ_i 's so that they fit the σ_i 's transition, while recalling that σ_i determines both the edge used and the direction in which it is traversed. (Suppose, again, without loss of generality, that an even σ_i (resp., an odd σ_i) indicates a forward (resp., reverse) transition using $\pi^{\lceil \sigma_i/2 \rceil}$). Then, we let $\tau_i = \sigma_i$.²³

Now, we claim that if a t -step random walk on the $4d$ -regular 1-local graph yields a distribution that is $\exp(-\Omega(t))$ -close to uniform (and $t = \Omega(n)$ is large enough), then the matrix $B^{(\bar{\sigma},\bar{\sigma})}$ must have full rank with probability at least $1 - \exp(-\Omega(t))$. This claim is shown as follows.

Let η denote the probability (over the choice of $\bar{\sigma} \in [2d]^t$) that the matrix $B^{(\bar{\sigma},\bar{\sigma})}$ does not have full rank. Such a choice of $\bar{\sigma}$ determines both the permutation $\pi_{\bar{\sigma}}$ that relates the original locations to the final ones (i.e., $\pi_{\bar{\sigma}} = \pi^{((-1)^{\sigma_t} \cdot \lceil \sigma_t/2 \rceil)} \circ \dots \circ \pi^{((-1)^{\sigma_1} \cdot \lceil \sigma_1/2 \rceil)}$) and a non-trivial linear combination $J_{\bar{\sigma}}$ of the columns of the matrix that witnesses that the matrix is not full rank. Hence, with probability $\eta' \geq \eta/(2^n - 1)$ over the choice of $\bar{\sigma}$, there exists a non-empty set $J \subseteq [n]$ such that the sum of the columns indexed by $\pi_{\bar{\sigma}}^{-1}(J)$ (in the matrix $B^{(\bar{\sigma},\bar{\sigma})}$) equals the all-zero vector, whereas in the remaining choices (of $\bar{\sigma}$) this sum does not equal the all-zero vector.²⁴ Looking

²³Note that if $\sigma_i = 2k$ (resp., $\sigma_i = 2k - 1$), then the i^{th} step applied the forward (resp., reverse) transition $x \mapsto x_{\pi^{(k)}} \oplus s^{(k)}$ (resp., $y \mapsto (y \oplus s^{(k)})_{\pi^{(-k)}}$), where $\pi^{(-k)}$ denotes the inverse of $\pi^{(k)}$. Recall that $T_{2k} = \{j \in [n] : s_j^{(k)} = 1\}$ and $T_{2k-1} = \{\pi^{(k)}(j) : s_j^{(k)} = 1\} = \{j : s_{\pi^{(-k)}(j)}^{(k)} = 1\}$.

²⁴The issue here is as follows: We know that for η fraction of the $\bar{\sigma}$'s, there exists a $J_{\bar{\sigma}}$ such that the sum of these columns is the all-zero vector. However, these columns correspond to locations in the (label of the) initial vertex, whereas we want to analyze locations in the end vertex. Still, there exists a non-empty J (representing locations in final label) such that the sum of the columns in $\pi_{\bar{\sigma}}^{-1}(J)$ (representing locations in initial label) equals the all-zero vector with probability $\eta' > \eta/(2^n - 1)$. Needless to say, for the rest of this probability space (of $\bar{\sigma} \in [2d]^t$), the sum is not the all-zero vector.

at the label of the final vertex $v_{\bar{\sigma}}$ in a random walk $\bar{\sigma}$ on the 1-local 2^n -vertex graph that starts at the vertex 0^n , we observe that $v_{\bar{\sigma}}$ equals a random linear combination of the rows of $B^{(\bar{\sigma}, \bar{\sigma})}$ permuted by $\pi_{\bar{\sigma}}$ (i.e., $(v_{\bar{\sigma}})_{\pi_{\bar{\sigma}}^{-1}}$ equals a random linear combination of the rows of $B^{(\bar{\sigma}, \bar{\sigma})}$, where this random linear combination is determined by the sequence (b_1, \dots, b_t)).²⁵ It follows that the sum of $v_{\bar{\sigma}}$'s bits in locations J is zero with probability exactly $\eta' + (1 - \eta') \cdot 0.5 = 0.5 + 0.5\eta'$, since this sum is 0 if the sum of the corresponding columns in $B^{(\bar{\sigma}, \bar{\sigma})}$ is the all-zero vector (and is uniformly distributed in $\{0, 1\}$ otherwise). Hence, the distribution of the final vertex is $0.5\eta'$ -far from the uniform distribution. The claim follows, since $\eta' \leq \exp(-\Omega(t))$ by the hypothesis, whereas this implies $\eta \leq 2^n \cdot \exp(-\Omega(t)) = \exp(-\Omega(t))$ for sufficiently large $t = \Omega(n)$. ■

Problem 12 (the CRW problem, revised): *For which graphs and which sequences of sets $(T_1, \dots, T_{d'})$'s does the random matrix considered in Definition 10 have full rank with probability at least $1 - \exp(-\Omega(t))$?*

An appealing conjecture of Benny Applebaum is that every n -vertex expander graph yield a positive answer to Problem 12 (i.e., there exists $d' = O(1)$ sets $T_1, \dots, T_{d'} \subset [n]$ such that this n -vertex graph combined with these sets satisfies the relaxed CRW property of Definition 10).

6 An afterthought: Generalization to non-binary alphabets

We generalize the basic definitions to an arbitrary alphabet of prime size, which is identified with the field $\text{GF}(p)$. A function $f : \text{GF}(p)^n \rightarrow \text{GF}(p)^n$ is called t -local if each symbol in its output depends on at most t symbol in its input. This yields to a generalized notion of a 1-local expander.

Definition 13 (1-local expanders, generalized): *For a fixed $d \in \mathbb{N}$ and a fixed prime p , let $\{f_1, \dots, f_d : \text{GF}(p)^n \rightarrow \text{GF}(p)^n\}_{n \in \mathbb{N}}$ be 1-local bijections. Then, the corresponding $2d$ -regular p^n -vertex graph consists of the vertex set $\text{GF}(p)^n$ and the edge multiset $\cup_{i \in [d]} \{\{x, f_i(x)\} : x \in \text{GF}(p)^n\}$.*

Note that each f_i is determined by a permutation on the bit locations $\pi^{(i)} : [n] \rightarrow [n]$, called the **relocation**, and n bijections denoted $h_1^{(i)}, \dots, h_n^{(i)} : \text{GF}(p) \rightarrow \text{GF}(p)$. Unlike in the binary case, where each $h_j^{(i)}$ is linear (i.e., has the form $h_j^{(i)}(z) = z \oplus s_j^{(i)}$), these bijections are not necessarily linear functions. Still, we shall focus on the case that they are linear. Generalizing Theorems 7 and 8, we obtain.

Theorem 14 (a construction of generalized 1-local expanders): *There exists a set of $d = O(1)$ explicit 1-local bijections, $\{f_1, \dots, f_d : \text{GF}(p)^n \rightarrow \text{GF}(p)^n\}_{n \in \mathbb{N}}$, such that, for every prime p , the $2d$ -regular p^n -vertex graph that consists of the vertex set $\text{GF}(p)^n$ and the edge multiset $\cup_{i \in [d]} \{\{x, f_i(x)\} : x \in \text{GF}(p)^n\}$ is an expander. Furthermore, the f_i 's are linear mappings.*

Proof: The overall plan is to use a straightforward generalization of the CRW property for rank defined over $\text{GF}(p)$, show that any generating set for the symmetric group of n elements that is expanding (along with any set of size $n' \approx n/2$ such that $n' \not\equiv 0 \pmod{p}$) satisfies this property, and that this yields a 1-local p^n -vertex expander.

²⁵This is the case since the i^{th} row permuted by $\pi^{((-1)^{\sigma_i \cdot \lceil \sigma_i / 2 \rceil}} \circ \dots \circ \pi^{((-1)^{\sigma_1 \cdot \lceil \sigma_1 / 2 \rceil}}$ is the offset that is potentially added in the i^{th} step of the walk, whereas this offset is added if and only if $b_i = 1$.

Definition 14.1 (a property of coordinated random walks, generalized): *For a d -regular n -vertex graph as in Definition 5, a set $T \subseteq [n]$ and $t = \Omega(n)$, consider coordinated random walks and Boolean matrices just as in Definition 5. The generalized CRW property postulates that, with probability at least $1 - \exp(-\Omega(t))$, such a random matrix has full rank when the arithmetics is in $\text{GF}(p)$.*

We stress that although these random matrices have entries in $\{0, 1\}$, we consider their rank over $\text{GF}(p)$.

Claim 14.2 (Theorem 7, generalized): *Let $\pi^{(1)}, \dots, \pi^{(d)} : [n] \rightarrow [n]$ be d permutations and $s = (s_1, \dots, s_n) \in \{0, 1\}^n \subseteq \text{GF}(p)^n$. If the $2d$ -regular n -vertex graph with the edge multi-set $\cup_{i \in [d]} \{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ along with the set $\{j \in [n] : s_j = 1\}$ satisfies the generalized CRW property (of Definition 14.1), then the $2p^2d$ -regular p^n -vertex graph with the edge multi-set $\cup_{i \in [d], b, c \in \text{GF}(p)} \{\{x, (x - b \cdot s)_{\pi^{(i)}} + c \cdot s\} : x \in \text{GF}(p)^n\}$ is an expander, where $b \cdot (s_1, \dots, s_n) = (bs_1, \dots, bs_n)$.*

Proof Sketch: We mimic the proof of Theorem 7, while noting that in the i^{th} step the vertex's label is randomized by an offset that is a random $\text{GF}(p)$ -multiple of the i^{th} row in the corresponding matrix. Hence, if the matrix has full rank over $\text{GF}(p)$, then the label of the final vertex is uniformly distributed in $\text{GF}(p)^n$ (since it is randomized by a random linear combination of the rows of the matrix). ■

Claim 14.3 (Theorem 8, generalized): *Let $\Pi = \{\pi^{(i)} : i \in [d]\}$ be a generating set of the symmetric group of n elements and suppose that Π is expanding. Then, the n -vertex graph that consists of the vertex set $[n]$ and the edge multi-set $\cup_{i \in [d]} \{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ combined with any set of size $n' \approx n/2$ such that $n' \not\equiv 0 \pmod{p}$ satisfies the generalized CRW property of Definition 14.1.*

Proof Sketch: Here we mimic the proof of Theorem 8. Specifically, we consider all (non-zero) linear combinations $L : [n] \rightarrow \text{GF}(p)$ of the columns of a random matrix, and upper bound the probability that each such linear combination yields the all-zero vector. That is, fixing an set T of size n' , for every such linear combination L , we consider the set W_L of permutations $\pi \in \text{Sym}_n$ such that $\sum_{i \in [n]: \pi(i) \in T} L(i) \not\equiv 0 \pmod{p}$. Once we show that each W_L has constant density, the claim follows as in the binary case (where here we use a union bound on all L 's).

The case of constant function $L : [n] \rightarrow \text{GF}(p)$ is handled by the hypothesis that $n' \not\equiv 0 \pmod{p}$ (which implies that $W_L = \text{Sym}_n$), and so we focus on non-constant functions L . We shall show that, for every value $v \in \text{GF}(p)$, the fraction of permutations π such that $\sum_{i \in [n]: \pi(i) \in T} L(i) \equiv v \pmod{p}$ is at most $0.5 + o(1)$. We first reduce the general case to the case that L has at least n/p zero entries: Given an arbitrary (non-constant) $L' : [n] \rightarrow \text{GF}(p)$, let $w \in \text{GF}(p)$ be an element that appears at least n/p times in L (i.e., $|\{j \in [n] : L'(j) = w\}| \geq n/p$), and consider the (non-zero) function $L(j) = L'(j) - w$.

Next, letting $J = \{i \in [n] : L(i) \neq 0\}$, suppose that we generate a random permutation π by first assigning elements to J , and consider the situation before the last assignment (i.e., after assigning $|J| - 1$ elements). Using the hypothesis that $|T| = n' \approx n/2$, w.v.h.p., before this last assignment, these $|J| - 1 < n - n/p$ locations were assigned approximately an equal number of elements from T and from $[n] \setminus T$, which means that $n' - (1 \pm o(1)) \cdot |J|/2 = (1 \pm o(1)) \cdot (n - |J|)/2$ elements from each type remain for the last assignment, where $n - |J| = \Omega(n)$. This means that the value of $\sum_{i \in [n]: \pi(i) \in T} L(i) \pmod{p}$ changes at the last assignment with probability $(1 \pm o(1))/2 \approx 1/2$ (i.e., if $i_{|J|}$ is the last element in J being assigned an element, then $L(i_{|J|})$ is added to the current

sum with probability $\approx 1/2$). The claim follows (since if the partial sum was v before the last assignment then it changes with probability at least $0.5 - \epsilon$, whereas if the partial sum was not v then it becomes v with probability at most $0.5 + o(1)$). ■

Combining Claims 14.3 and 14.2, we get.

Corollary 14.4 (obtaining generalized 1-local expanders): *Let $\Pi = \{\pi^{(i)} : i \in [d]\}$ be a generating set of the symmetric group of n elements and suppose that Π is expanding. Then, for any $n' \approx n/2$ such that $n' \not\equiv 0 \pmod{p}$, the $2p^2d$ -regular p^n -vertex graph with the edge multi-set $\cup_{i \in [d], b, c \in \text{GF}(p)} \{x, (x - b^{n'} 0^{n-n'})_{\pi^{(i)}} + c^{n'} 0^{n-n'}\} : x \in \text{GF}(p)^n\}$ is an expander.*

Using Kassabov's result [4] (which asserts that the symmetric group has an explicit generating set that is expanding and of constant size), the theorem follows. ■

Comment: The generalizes to any finite field; that is, p may be a prime power. For $p = q^\ell$, where q is prime, we select $n' \approx n/2$ such that $n' \not\equiv 0 \pmod{q}$, and proceed as above (while noting that in the proof of Claim 14.3 the reductions mod p actually refer to doing the arithmetics in $\text{GF}(p)$).

Acknowledgments

I wish to thank Benny Applebaum for helpful discussions and for permission to include his conjecture in this text. I am also grateful to Roei Tell for commenting on a draft of this text.

This work was partially supported by the Minerva Foundation with funds from the Federal German Ministry for Education and Research.

References

- [1] A. Broder and A. Karlin. Bounds on the cover time. *J. of Theoretical Probability*, Vol. 2 (1), pages 101–120, 1989.
- [2] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph, and its applications to random walks. In *21st STOC*, 1989.
- [3] S. Horry, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bull. (new series) of the AMS*, Vol. 43 (4), pages 439–561, 2006.
- [4] M. Kassabov. Symmetric groups and expander graphs. *Invent. Math.*, Vol. 170 (2), pages 327–354, 2007.
- [5] R. Rubinfeld. The cover time of a regular expander is $O(n \log n)$. *IPL*, Vol. 35, pages 49–51, 1990).
- [6] E. Viola and A. Wigderson. Local Expanders. *ECCC*, TR16-129, 2016.