Deconstructing 1-local expanders

Oded Goldreich

November 29, 2019

Abstract

A 1-local $2d$-regular 2^n-vertex graph is represented by d bijections over $\{0,1\}^n$ such that each bit in the output of each bijection is a function of a single bit in the input. An explicit construction of 1-local expanders was presented by Viola and Wigderson (*ECCC*, TR16-129, 2016), and the goal of the current work is to de-construct it; that is, make its underlying ideas more transparent.

Starting from a generic candidate for a 1-local expander (over $\{0,1\}^n$), we first observe that its underlying bijections consists of pairs of ("relocation") permutations over $[n]$ and offsets (which are n-bit long strings). Next, we formulate a natural problem regarding "coordinated random walks" (CRW) on the corresponding (n-vertex) "relocation" graph, and prove the following two facts:

1. Any solution to the CRW problem yields 1-local expanders.
2. Any constant-size expanding set of generators for the symmetric group (over $[n]$) 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 (a relaxed version of) the CRW problem is equivalent to constructing 1-local expanders.

An early version of this work appeared as TR16-152 of *ECCC*; it was written in a rather laconic style and reflected the train of thoughts of the author. The current version was significantly revised, adding various clarifications and elaborations with the aim of better serving a wider readership.

1 Introduction

Expander graphs are families of regular graphs of fixed degree and constant "expansion" factor (equiv., logarithmic mixing time), where the family consists of graphs for varying number of vertices. Expander graphs exist in abundance (i.e., a random $O(1)$-regular graph is an expander, w.v.h.p.) and have numerous applications in the theory of computation (see, e.g., [3]). Hence, explicit constructions of expanders are of great interest to this field, and the more explicit the construction – the better.

A strong notion of explicitness refers to the computation of the neighborhoods in the expander. Specifically, given the label of a vertex v and an index i, the task is to find the i^{th} neighbor of v. It is also desirable to have $2d$-regular expanders that can be represented by d *simple* permutations of the vertex set, where each permutation corresponds to a collection of disjoint cycles that covers the vertex-set. But, *how simple can these bijections be?*

It turns out that these bijections can be extremely simple. Specifically, considering graphs over the vertex set $\{0,1\}^n$, Viola and Wigderson [6] showed that each bit in the output of each of the corresponding bijections may depend on a single bit in the input (i.e., the label of the vertex); that is, these bijections are 1-local. Recall that a function $f : \{0,1\}^n \to \{0,1\}^n$ is called t-local if each bit in its output depends on at most t bits in its input. The aforementioned result of Viola and Wigderson [6] asserts:

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

Indeed, by association, we refer to a regular 2^n-vertex graph as in Theorem 1 by the term 1-local; that is, a $2d$-regular graph $G = (\{0,1\}^n, E)$ is called 1-local if $E = \bigcup_{i\in[d]}\{\{x, f_i(x)\} : x \in \{0,1\}^n\}$ such that each f_i is a 1-local bijection. (By saying that the foregoing bijections are explicit, we mean that they can be evaluated in poly(n)-time, where throughout the paper we think of n as varying.)

1.1 Initial observations, which should not be skipped

We first observe that each 1-local bijection $f_i : \{0,1\}^n \to \{0,1\}^n$ is determined by a permutation of the bit locations $\pi^{(i)} : [n] \to [n]$, called the relocation (permutation), 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 $s^{(i)}$ (equiv., the j^{th} bit of $f_i(x)$ equals the sum of the ${\pi^{(i)}(j)}^{\text{th}}$ bit of x and the j^{th} bit of $s^{(i)}$).

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 (see Section 1.3)).

Observation 2 (the "shuffle exchange" graph is a 1-local "weak expander"):[1] *Let $f_1(x) = \mathtt{sh}(x)$ and $f_2(x) = x \oplus 0^{n-1}1$, where $\mathtt{sh}(x_1 \cdots x_n) = (x_2 \cdots x_n x_1)$ is a cyclic shift that corresponds to the relocation permutation $\pi(i) = (i \bmod n) + 1$. Then, the 4-regular 2^n-vertex 1-local graph that consists of the vertex set $\{0,1\}^n$ and the edge multi-set $\bigcup_{i\in[2]}\{\{x, f_i(x)\} : x \in \{0,1\}^n\}$ has second* (normalized) *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., $\{x, \mathtt{sh}(x)\}$ and $\{x, \mathtt{sh}^{-1}(x)\}$ for each x) may also be non-distinct.)

[1] A similar result holds for the 4-regular graph that uses the bijections $f_1(x) = \mathtt{sh}(x)$ and $f_2(x) = \mathtt{sh}(x) \oplus 0^{n-1}1$. Note 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 $\mathtt{sh}(x)$ and $\mathtt{sh}(x) \oplus 0^{n-1}1$, 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. For further discussion, see Section 1.4.

Proof Sketch: 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, whereas the rate of convergence (a.k.a. the second eigenvalue) is closely related to the distance from uniformity (see Section 1.3). Specifically, it follows that the convergence rate is at least $2^{-(t'-n)/t}$, which equals $2^{1/O(n^2)} = 1 - \Omega(n^{-2})$ for sufficiently large t' (e.g., $t' > 2n$).

The foregoing claim is proved by observing that during a t-step random 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).[2] ■

The relocation graph. The foregoing argument refers implicitly to a random walk on the n-vertex cycle, which represents the shift relocation permutation ($i \mapsto (i \bmod n) + 1$) used in the 1-local 2^n-vertex graph (considered in Observation 2). In general, we shall consider the n-vertex graph that corresponds to the relocation permutations of the 1-local 2^n-vertex graph that we wish to analyze. Hence, 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 $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ be d permutations and $s^{(1)}, ..., s^{(d)} \in \{0,1\}^n$.*

1. *The* 1-local graph associated with $\pi^{(1)}, ..., \pi^{(d)}$ and $s^{(1)}, ..., s^{(d)}$ *is the $2d$-regular 2^n-vertex graph that consists of the vertex set $\{0,1\}^n$ and the edge multi-set*

$$\bigcup_{i\in[d]} \left\{ \{x, x_{\pi^{(i)}} \oplus s^{(i)}\} : x \in \{0,1\}^n \right\} \qquad (1)$$

where $x_\pi = x_{\pi(1)} \cdots x_{\pi(n)}$.

2. *The* relocation graph associated with $\pi^{(1)}, ..., \pi^{(d)}$ *is the $2d$-regular n-vertex graph that consists of the vertex set $[n]$ and the edge multi-set $\bigcup_{i\in[d]}\{\{j, \pi^{(i)}(j)\} : j \in [n]\}$.*

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 $\pi^{(i)}$.*

Note that $((x_{\pi^{(i)}} \oplus s^{(i)}) \oplus s^{(i)})_{\pi^{(-i)}} = (x_{\pi^{(i)}})_{\pi^{(-i)}} = x$ and $((y \oplus s^{(i)})_{\pi^{(-i)}})_{\pi^{(i)}} \oplus s^{(i)} = (y \oplus s^{(i)}) \oplus s^{(i)} = y$.

The proof of Observation 2 is based on the fact that the corresponding relocation graph (i.e., the n-cycle) has cover time $O(n^2)$. Using an n-vertex expander as the relocation graph, and relying

[2] After i steps, the j^{th} bit in the original string (which is originally located at position j) is located at position $(j - 1 + \sum_{k\in[i]} X_k \bmod n) + 1$, 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 $\texttt{sh}$ (resp., $\texttt{sh}^{-1}$) 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 $t/t' = O(n^2)$ random variables 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 partials sums has value outside $[-n, n]$, then in the corresponding $O(n^2)$ steps each original bit position appeared in the rightmost location.

on the fact that it has cover time $\widetilde{O}(n)$, we may infer that the corresponding 1-local 2^n-vertex graph has spectral gap $1/\widetilde{O}(n)$ (see Appendix). It turns out that we need a relocation graph that has a stronger "mixing" property than a standard expander, and studying this property is the core of the current work. In addition, we need to use offsets that are not of Hamming weight 1. In fact, we need to use offsets that have Hamming weight $\Omega(n)$.

Proposition 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.*

The proof of Proposition 4 appears in the appendix, where it is also shown that using also offsets of Hamming weight $n - o(n)$ does not help. In view of the above, we must use at least one offset that has Hamming weight in $[\Omega(n), n - \Omega(n)]$.

1.2 Our main results

As further discussed in Section 1.3, expander graphs have the salient property by which a random walk of logarithmic length end in a vertex that is almost uniformly distributed on the vertex set, regardless of the start vertex. Our plan is to show that a 2^n-vertex 1-local graph is an expander by showing that a random walk of length $t = \omega(n)$ ends in a vertex that is $\exp(-\Omega(t))$-close to be uniformly distributed in $\{0,1\}^n$.

As stated above, we will show this by relying on the hypothesis that the corresponding n-vertex relocation graph has a property that is stronger than standard expansion. The stronger property that we shall use refers to "coordinated (t-step) random walks" that start at the n different vertices, where the t-step random walks are specified by t indices that determine the choices of neighbors at each step. That is, the sequence $(\sigma_1, ..., \sigma_t)$ corresponds to t-step walks that in the i^{th} step move to the σ_i^{th} neighbor of the current vertex. Hence, by *coordinated random walks* of length t on a $2d$-regular n-vertex graph, we mean selecting uniformly one sequence $(\sigma_1, ..., \sigma_t) \in [2d]^t$, and considering the n corresponding walks (such that the j^{th} walk starts at vertex j).

A standard property of a single random walk refers to the number of times that the walk hits a set of constant density. In a standard n-vertex expander, the fraction of hits in a sufficiently long random walk (i.e., one of $\Omega(\log n)$ length) closely approximates the density of the set (with probability that is exponential in the length of the walk). The property that we shall consider for coordinated random walks of length $t \geq (1 + \Omega(1)) \cdot n$ is that the matrix describing the hitting pattern of the n coordinated walks has full rank with probability $1 - \exp(-\Omega(t - n))$. That is, for a fixed set T (of constant density), we consider a random Boolean t-by-n matrix such that the $(i,j)^{\text{th}}$ entry is 1 if and only if the j^{th} walk hits the set T in the i^{th} step. This property is the pivot of our results. Specifically, we prove the following two facts:

Graphs satisfying the coordinated random walks property yield 1-local expanders (see Theorem 8): *If the n-vertex relocation graph satisfies the foregoing property* (of coordinated random graphs), *then the corresponding 1-local 2^n-vertex graph coupled with adequate offsets is an expander.* (Actually, given a $2d$-regular n-vertex relocation graph, we consider a $8d$-regular 2^n-vertex 1-local graph, where each relocation permutation is coupled with four offsets (and the offsets are easily computed based on the permutations).)

Obtaining graphs that satisfy the coordinated random walks property (see Theorem 9): *Any constant-size expanding set of generators for the symmetric group* (over $[n]$) *yields an n-vertex* (relocation) *graph that satisfies the foregoing property* (of coordinated random graphs).

Lastly, a result of Kassabov [4], which was also used in [6], asserts that the symmetric group has an explicit generating set that is expanding and of constant size.

Combining these three results, we obtain an alternative proof of Theorem 1. In addition, we show that constructing an n-vertex graph that satisfies (a relaxed version of) the foregoing property is equivalent to constructing 1-local 2^n-vertex expanders.

Organization of the rest of this paper. Sections 1.3 and 1.4 provide necessary background and a useful clarification towards the rest of the paper. The main results (i.e., Theorems 8 and 9) are proved in Sections 2 and 3, respectively. Specifically, in Section 2 we present the coordinated random walks (CRW) property and show that satisfying it yields 1-local expanders, and in Section 3 we show that any constant size set of expanding generators for the symmetric group over $[n]$ yields an n-vertex $O(1)$-regular graph that satisfies the CRW property. In Section 4 we prove the aforementioned equivalence (between a relaxed version of the CRW property and constructing 1-local expanders), and in Section 5 we generalize the main results to non-binary alphabets.

1.3 The algebraic definition of expansion and the convergence rate perspective

The combinatorial definition of expansion, which refers to the relative size of neighborhoods of sets of vertices (e.g., the number of vertices that neighbor the set but are not in it as a function of the size of the set), has a strong intuitive appeal. The same holds with respect to the algebraic definition, which refers to the (normalized) second eigenvalue of the corresponding adjacency matrix, provided that one realizes that the (normalized) second eigenvalue represents the rate at which a random walk on a regular graph converges to the uniform distribution. Indeed, in this work we shall use the term convergence rate when referring to this eigenvalue. In an expander the convergence rate is a constant smaller than 1, whereas in a general (regular and non-bipartite) N-vertex graph the convergence rate is upper-bounded by $1 - \frac{1}{\mathrm{poly}(N)}$.

When trying to estimate the convergence rate of a regular graph, it is useful to consider a sufficiently long random walk and relate the convergence rate to the distance of the distribution of its end-vertex from the uniform distribution over the vertex-set. Specifically, consider an N-vertex regular graph, and let λ denote its convergence rate and $\Delta_t^{(p)}$ denote 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. Then, the following two facts relate λ^t and $\Delta_t^{(p)}$ (up-to a slackness of $\mathrm{poly}(N)$):[3]

1. *The distance is upper-bounded in terms of the convergance rate*: $\Delta_t^{(1)} \leq \sqrt{N} \cdot \Delta_t^{(2)} \leq \sqrt{N} \cdot \lambda^t$.
2. *The distance is lower-bounded in terms of the convergance rate*: $N^{-1} \cdot \lambda^t \leq \Delta_t^{(2)} \leq \Delta_t^{(1)}$.

[3] The first inequality (i.e., $\Delta_t^{(2)} \leq \lambda^t$) is well-known and extensively used. It captures the fact that the corresponding linear operator shrinks each vector that is orthogonal to the uniform one. The second inequality (i.e., $\Delta_t^{(2)} \geq \lambda^t/N$) is far less popular. It can be proved by considering a random walk that starts in a probability distribution that is described by the vector $u+v_2$, where $u = (1/N, ..., 1/N)$ is the uniform distribution and v_2 is a vector in the direction of the second eigenvector such that no coordinate in v_2 has value lower than $-1/N$.

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

We shall use $\lambda \leq (N \cdot \Delta_t^{(1)})^{1/t}$ quite extensively. In fact, we have already used it in the proof of Observation 2, where, using $N = 2^n$ and $t = O(t' \cdot n^2)$, we showed that $\Delta_t^{(1)} \leq 2^{-t'}$ and (using $t' \geq 2n$) inferred that $\lambda \leq (2^n \cdot 2^{-t'}) = 2^{1/O(n^2)} = 1 - \Omega(n^{-2})$. In the following sections, for sufficiently large t, we shall show that $\Delta_t^{(1)} = \exp(-\Omega(t))$, and infer that $\lambda = \exp(-\Omega(1)) < 1$ (i.e., the convergence rate is upper-bounded by a constant smaller than 1).

1.4 A technical source of complication

As commented in Footnote 1, the proof of Observation 2 is less simple than it could have been because we have to account for both forward and reverse transitions. Unfortunately, the same phenomenon occurs in the proofs of our main results. In other words, it would have been simpler to analyze a random walk on the directed graph that corresponds to forward transitions only (and, ditto, of course, for the directed graph of reverse transitions). We may refer to such directed walks in the warm-ups, but we cannot reduce the analysis to them. Such a reduction would require a de-composition result that is not true in general but may be true in some special cases (and in particular in those that are of interest to us here).

Open Problem 5 (de-composing random walks on regular graphs): *For a function $f : (0,1) \to (0,1)$, we say that a class of $2d$-regular graphs defined by d bijections on their vertex-sets is f-good if for every graph in the class the following holds: If the convergence rate of the directed graph containing only forward transitions is at most λ, then the convergence rate of the undirected graph is at most $f(\lambda)$. We ask:*

1. *Is the class of 1-local graphs f-good for some f?*

2. *Which natural classes of graphs are f-good for some f?*

Note that Item 1 can only hold with $f(\lambda) = 1 - \Omega(1 - \lambda)^2$. This is the case because for the 1-local 2^n-vertex graph of Observation 2 the convergence rate of the directed graph is $1 - \Theta(1/n)$, whereas the convergence rate of the graph itself is $1 - \Theta(1/n^2)$. (An $3n$-step random walk on the directed graph yields an almost uniformly distributed vertex, whereas an $o(n^2)$-step random walk on the graph itself is unlikely to reach a vertex that is at Hamming distance at most $n/3$ from the start vertex.)

2 A sufficient condition: The coordinated random walks property

As stated upfront, a 1-local $2d$-regular 2^n-vertex graph is associated with d relocation permutations and d offsets, which means that constructing 1-local expander graphs reduces to constructing suitable relocation graphs and offsets. In this section we identify a property of the relocation graph (equiv., of the d relocating permutations) that suffices for showing that the corresponding 1-local graph is an expander. In Section 2.1 we present the basic intuition, while referring to a random walk on the directed graph of forward transitions only. The actual analysis is presented in Section 2.2, and it is intended to be understood also when skipping Section 2.1.

2.1 The intuition

The general problem we face is of finding relocation graphs and offsets that yield 1-local expanders. For sake of simplicity, we consider the case of using a single non-zero offset $s \in \{0,1\}^n$, since it turns out that this suffices. As a warm-up towards the actual problem, we consider a generic 1-local ($4d$-regular) graph with d relocation permutations, $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ and $2d$ forward transitions such that the $2i-1^{\rm st}$ (resp., the $2i^{\rm th}$) transition is $x \mapsto x_{\pi^{(i)}} \oplus s$ (resp., $x \mapsto x_{\pi^{(i)}}$). Indeed, we use each relocation permutation both with the offset s and without it; that is, the set of corresponding bijections is $\{f_{(\sigma,b)} : \{0,1\}^n \to \{0,1\}^n\}_{(\sigma,b)\in[d]\times\{0,1\}}$ such that $f_{(\sigma,1)}(x) = x_{\pi^{(\sigma)}} \oplus s$ and $f_{2\sigma}(x) = x_{\pi^{(\sigma)}}$.

We shall consider a t-step random walk (on the 1-local graph) that uses only the forward transitions (and starts at the vertex 0^n). Such a walk is specified by a sequence of pairs, denoted $((\sigma_1, b_1), ..., (\sigma_t, b_t)) \in ([d] \times \{0,1\})^t$, such that at the $i^{\rm th}$ step the $(2\sigma_i - b_i)^{\rm th}$ forward transition is used (so that the result is moving from the current vertex v to the vertex $v_{\pi^{(\sigma_i)}} \oplus s^{b_i}$, where $s^0 = 0^n$ and $s^1 = s$). Hence, in the $i^{\rm th}$ step, the label of the current vertex is permuted (according to $\pi^{(\sigma_i)}$) and then offset by s if $b_i = 1$ (and is only permuted otherwise).

Actually, it is instructive to say that, in the $i^{\rm th}$ step, the offset $s^{(i)} \stackrel{\rm def}{=} s_{\pi^{(\sigma_i)}\circ\cdots\circ\pi^{(\sigma_1)}}$ is added to the initial label (i.e., 0^n) if and only if $b_i = 1$. Hence, the label of the final vertex in the walk equals $\sum_{i\in[t]:b_i=1} s^{(i)}$ permuted according to $\pi^{(\sigma_t)} \circ \cdots \circ \pi^{(\sigma_1)}$. Note that, for any fixed sequence $(\sigma_1, ..., \sigma_t) \in [d]^t$, we get a distribution that depends only on the random b_i's, since the $s^{(i)}$'s depend only on the sequence $\sigma_1, ..., \sigma_t$ (and on s). The punch-line is that *this distribution* (i.e., $\sum_{i\in[t]:b_i=1} s^{(i)}$ for random b_i's) *is uniform over* $\{0,1\}^n$ *if and only if the* t*-sequence* $(s^{(1)}, ..., s^{(t)})$ *spans* $\{0,1\}^n$.

Focusing on the question of whether or not the t-sequence $(s^{(1)}, ..., s^{(t)})$ spans $\{0,1\}^n$, we observe that this question refers to a property of the corresponding walk on the relocation graph, since the $s^{(i)}$'s are determined by the sequence $\sigma_1, ..., \sigma_t$ (and s). Specifically, for each $j \in [n]$, we consider a walk that starts at vertex j and proceeds according to the sequence $(\sigma_1, ..., \sigma_t) \in [d]^t$. After i steps, this ($j^{\rm th}$) walk is at vertex $\pi^{(\sigma_i)} \circ \cdots \circ \pi^{(\sigma_1)}(j)$. Now, the key observation is that the $j^{\rm th}$ bit of $s^{(i)} = s_{\pi^{(\sigma_i)}\circ\cdots\circ\pi^{(\sigma_1)}}$ is 1 if and only if this walk hits the set of vertices $T = \{k : s_k = 1\}$ (i.e., iff $\pi^{(\sigma_i)} \circ \cdots \circ \pi^{(\sigma_1)}(j) \in T$). Hence, $(s^{(1)}, ..., s^{(t)})$ spans $\{0,1\}^n$ if and only if the t-by-n Boolean matrix that describes the hitting pattern in T if full rank.

That is, for every sequence $\overline{\sigma} = (\sigma_1, ..., \sigma_t)$, we consider a t-by-n Boolean matrix $B = B^{(\overline{\sigma})}$ such that the $(i,j)^{\rm th}$ entry in this matrix is 1 if and only if the walk that starts at j and proceeds according to $\overline{\sigma}$ hits T in the $i^{\rm th}$ step (i.e., iff $\pi^{(\sigma_i)} \circ \cdots \circ \pi^{(\sigma_1)}(j) \in T$). The foregoing observation is that, when letting $s^{(i)} = s_{\pi^{(\sigma_i)}\circ\cdots\circ\pi^{(\sigma_1)}}$, it holds that $(s^{(1)}, ..., s^{(t)})$ spans $\{0,1\}^n$ if and only if $B^{(\overline{\sigma})}$ is full rank.

To summarize, we are looking at coordinated random walks on the relocation graph. These walks start at different vertices $j \in [n]$ but proceed according to a single random sequence $\overline{\sigma} = (\sigma_1, ..., \sigma_t) \in [d]^t$. For a fixed set T, we are interested in the event that the corresponding matrix (i.e., $B^{(\overline{\sigma})}$) has full rank (i.e., its rows span $\{0,1\}^n$). Whenever this happens, the corresponding random walk on the 1-local graph, which is further randomized by the choice of the b_i's, is uniformly distributed over $\{0,1\}^n$. The property that we wish to hold is that, with probability at least $1 - \exp(-\Omega(t))$ over the choice of $\overline{\sigma}$, the matrix $B^{(\overline{\sigma})}$ has full rank.

2.2 The actual analysis

It turns out that it suffices to use a single non-zero offset $s \in \{0,1\}^n$ (of Hamming weight approximately $n/2$), along with the offsets that are derived from it when considering also the reverse transitions. That is, for each relocation permutation $\pi : [n] \to [n]$, 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 $y \mapsto (y \oplus s^c)_{\pi^{-1}} \oplus s^b = y_{\pi^{-1}} \oplus (s_{\pi^{-1}})^c \oplus s^b$.)[4] 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).[5]

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 these t steps, where an initial position is said to be currently active if it currently reside in a location in $\{k : s_k = 1\}$. That is, fixing the choice of relocation permutations (but leaving the choice of the b's and c's undermined), the $(i,j)^{\text{th}}$ 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). Later, when considering the effect of using random b's and c's, this will have the effect of flipping all active locations (together) with probability 1/2.

Note that the foregoing matrix, which is defined based on a fixed 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 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 and c unspecified.)

The punch-line is that, *if the foregoing t-by-n matrix has full rank, then the t random choices of whether to apply the offset s* (which are governed by the random choice of the corresponding bits b and c) *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 resulting n-bit string is added to the initial vertex in the walk). That 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 $\{0,1\}^n$, since it is (essentially) a

[4] In contrast, if we were to use only 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 (i.e., the proof of Theorem 8); see also the following paragraph. Of course, the issue would not have arose if we were analyzing random walks on the directed graph of forward transitions only (see Sections 1.4 and 2.1).

[5] If 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 $b = 1$.

random linear combination of the rows of the matrix (where the randomization is due to the choices of the corresponding b's and c's).

The foregoing motivates the definition of the following property, where we are actually interested in the case that the g_σ's are the relocation permutations and their inverses (in an d-regular relocation (undirected) graph).[6]

Definition 6 (a property of coordinated random walks):[7] *For $d = O(1)$, consider a d-regular n-vertex directed graph such that for every $\sigma \in [d]$ the function $g_\sigma : [n] \to [n]$ that maps each vertex to its $\sigma^{\rm th}$ neighbor is a bijection.*

- *For an integer t, consider a random sequence $\overline{\sigma} = (\sigma_1, ..., \sigma_t) \in [d]^t$ and the n corresponding* coordinate random walks (CRW) *such that the $j^{\rm th}$ walk starts at vertex j and moves in the $i^{\rm th}$ step to the $\sigma_i^{\rm th}$ neighbor of the current vertex.*

- *For a set $T \subseteq [n]$, consider a t-by-n Boolean matrix $B^{(\overline{\sigma})} = B_T^{(\overline{\sigma})}$ such that its $(i,j)^{\rm th}$ entry indicates whether the $j^{\rm th}$ walk passed through T in its $i^{\rm th}$ step; that is, the $(i,j)^{\rm th}$ is 1 if and only if $g_{\sigma_i}(\cdots(g_{\sigma_1}(j)\cdots)) \in T$.*

The desired CRW property *is that for some set T, with probability at least $1 - 2^{-n-\Omega(t)}$ over the choice of $\overline{\sigma} \in [d]^t$, the corresponding matrix $B_T^{(\overline{\sigma})}$ has full rank* (over GF(2)).[8] *In this case we say that the* CRW property holds w.r.t *T.*

Obviously, the CRW property mandates $t \geq n$. Furthermore, the proof of Proposition 4 actually shows that the CRW property mandates that the set T must have size in $[\Omega(n), n - \Omega(n)]$. We are definitely not concerned of these restrictions.

Intuitively, the CRW property postulates that, with extremely high probability, the coordinated random walks are "linearly independent" with respect to hitting the set T. The allowed failure probability is $\exp(-\Omega(t))$, which is extremely low given that the probability space is of size $\exp(O(t))$.

Standard expanders may not satisfy the CRW property. We note that using an arbitrary expander graph and an arbitrary set T of size $\approx n/2$ *will not do*: Indeed, in this case, each column in the matrix corresponding to a random walk has approximately $t/2$ 1-entries (w.v.h.p.), but this matrix may not have full rank. 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

[6] Indeed, Definition 6 is stated in more general terms that fit an arbitrary directed graph that is described in terms of d directed cycle covers; that is, each g_σ describes a collection of directed cycles that cover all the graph's vertices, and the formulation refers to random walks in the direction of the edges. The special case we are interested in refers to the case that $g_{2\sigma'}$ is the inverse of $g_{2\sigma'-1}$; in this case, the directed graph consists of anti-parallel edges that correspond to the forward and reverse transitions, and a random walkmay take forward and reverse transitions (by picking either $g_{2\sigma'}$ or $g_{2\sigma'-1}$).

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

[8] The failure bound is set to $\tau = 2^{-n-\Omega(t)}$ in order to facilitate deriving an upper bound on the convergence rate of the corresponding 1-local graph. Specifically, we shall use $(2^n \cdot \tau)^{1/t} < 1$. An alternative formulation that will support this application is to require error probability at most $\exp(-\Omega(t))$ for some $t = \omega(n)$ (or error probability at most 2^{-ct} for some constant $c > 0$ and some $t \geq \frac{1+c}{c} \cdot n$).

these two expanders. Then, w.r.t this T, coordinated walks on this graph always yields a Boolean matrix of rank at most two, since all the (coordinated) walks that start at vertices in T (resp., in $[n] \setminus T$) always move together to T or to $[n] \setminus T$. (Nevertheless, it may be that for every n-vertex expander, there exists a set T such that the CRW property holds.)

So what does satisfy the CRW property? Indeed, the question we consider is the following.

Open Problem 7 (the CRW problem): *For which graphs and which sets T's does the CRW property* (as in Definition 6) *hold?*

A partial answer to this question is postponed to Section 3, and is revisited in Section 4. But before proceeding there, we establish the usefulness of the CRW property for constructing 1-local expanders.

The CRW property implies 1-local expanders. As outlined above, any $2d$-regular relocation graph that satisfies the CRW property yields an $8d$-regular 1-local 2^n-vertex expander. Let us formally state and prove this claim.

Theorem 8 (graphs satisfying the CRW property yield 1-local expanders): *Let $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ be d permutations and $s \in \{0,1\}^n$. If the $2d$-regular n-vertex graph with the edge multi-set $\bigcup_{i\in[d]}\{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ satisfies the CRW property with respect to the set $\{j \in [n] : s_j = 1\}$, then the $8d$-regular 2^n-vertex 1-local graph with the edge multi-set*

$$\bigcup_{i\in[d],b,c\in\{0,1\}} \left\{\{x, (x \oplus s^b)_{\pi^{(i)}} \oplus s^c\} : x \in \{0,1\}^n\right\} \tag{2}$$

is an expander.

Indeed, the $8d$-regular 1-local graph given by Eq. (2) is not the 1-local graph that correspond to the foregoing $2d$-regular relocation graph (per Definition 3), but it does correspond to the $8d$-regular relocation graph with the edge multi-set $\bigcup_{i\in[d],b,c\in\{0,1\}}\{\{j, \pi^{(i,b,c)}(j)\} : j \in [n]\}$, where $\pi^{(i,b,c)} = \pi^{(i)}$ for every $i \in [d]$ and $b, c \in \{0,1\}$. Needless to say, the difference between these two relocation graphs is immaterial.

Proof: Recall that a t-step random walk on the $2d$-regular relocation graph is specified by a sequence $(\sigma_1, ..., \sigma_t) \in [2d]^t$, whereas a random walk on the $8d$-regular 1-local graph is specified by a sequence $((\sigma_1, b_1, c_1), ..., (\sigma_t, b_t, c_y)) \in ([2d] \times \{0,1\}^2)^t$. By the hypothesis, for some $t = \Omega(n)$, the t-by-n matrix that corresponds to a random walk on the n-vertex relocation graph has full rank with probability at least $1 - 2^{-n-\Omega(t)}$. Fixing an arbitrary walk $\overline{\sigma} = (\sigma_1, ..., \sigma_t) \in [2d]^t$ on this n-vertex graph such that $B^{(\overline{\sigma})} = B^{(\overline{\sigma})}_{\{j\in[n]:s_j=1\}}$ has full rank, for each $i \in [t]$, we consider the residual random choices of $b_i, c_i \in \{0,1\}$ 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 stages.

- In the first stage, 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 stage, 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 stage, the label of the final vertex is a random variable that depends only on the random choices made in the second stage. The key observation is that these random choices have the effect of randomizing the vertex-label by adding to it a corresponding random linear combination of the rows of the matrix $B^{(\overline{\sigma})}$. 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).

> **Detailed (alas tedious) analysis**: Denoting the initial vertex in the walk on the 1-local graph by v_0, the $i^{\rm th}$ 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 $i^{\rm th}$ 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 $i^{\rm th}$ step takes a forward transition and $(x_i, y_i) = (c_i, b_i)$ otherwise. Note that the x_i's were fixed in the first stage, whereas the π_i's were fixed at the onset. In contrast, the y_i's are selected at random in the second stage. In both cases (i.e., regardless if $(x_i, y_i) = (b_i, c_i)$ or $(x_i, y_i) = (c_i, b_i)$), the $i^{\rm th}$ row in the matrix, denoted r_i, equals $s_{(\pi_i \circ \cdots \circ \pi_1)^{-1}}$, where $\pi_i \circ \cdots \circ \pi_1$ is the composition of the relocation permutations applied in the i first steps. Hence,
>
> $$\begin{aligned} (v_i)_{(\pi_i \circ \cdots \circ \pi_1)^{-1}} &= (v_{i-1})_{(\pi_{i-1} \circ \cdots \circ \pi_1)^{-1}} \oplus (s_{(\pi_{i-1} \circ \cdots \circ \pi_1)^{-1}})^{x_i} \oplus (s_{(\pi_i \circ \cdots \circ \pi_1)^{-1}})^{y_i} \\ &= (v_{i-1})_{(\pi_{i-1} \circ \cdots \circ \pi_1)^{-1}} \oplus (s_{(\pi_{i-1} \circ \cdots \circ \pi_1)^{-1}})^{x_i} \oplus r_i^{y_i}. \end{aligned}$$
>
> It follows that $(v_t)_{(\pi_t \circ \cdots \circ \pi_1)^{-1}} = v_0 \oplus w \oplus \bigoplus_{i \in [t]} r_i^{y_i}$, where $w = \bigoplus_{i \in [t]} (s_{(\pi_{i-1} \circ \cdots \circ \pi_1)^{-1}})^{x_i}$. This means that the label of the vertex reached by this random walk is the sum of an already fixed value (i.e., $v_0 \oplus w$) and the random linear combination of the rows of the matrix (i.e., $\bigoplus_{i \in [t]} r_i^{y_i}$, where the y_i's are uniformly distributed).

Hence, in this case (i.e., $B^{(\overline{\sigma})}$ has full rank), 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 $2^{-n-\Omega(t)}$-close to the uniform distribution, which implies that the converagence rate of the 2^n-vertex graph is bounded away from 1 (i.e., it is at most $(2^n \cdot 2^{-n-\Omega(t)})^{1/t} = 2^{-\Omega(1)}$), which means that this 1-local graph is an expander. ■

3 Constructions that satisfy the CRW property

For the benefit of the reader, we distinguish between the main result of this section (presented in Section 3.1) and two secondary comments that follow its presentation (in Section 3.2).

3.1 The main result of this section

Recall that Kassabov's result [4], which was also used in [6], asserts that the symmetric group has an explicit generating set that is expanding and of constant size.[9] We shall show that using this

[9] Indeed, 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 $[n]$). To reduce confusion, in the main text (unlike in footnotes), we shall not

set of permutations (i.e., as our set of relocating permutations) along with the set $T = [n']$ 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 6). (In fact, our result is more general.) Combined with Theorem 8, this yields an alternative proof of Theorem 1.

Theorem 9 (graphs satisfying the CRW property): *Let $\Pi = \{\pi^{(i)} : i \in [d]\}$ be a generating set of the symmetric group of n elements and suppose that Π is expanding.*[10] *Then, the n-vertex graph that consists of the vertex set $[n]$ and the edge multi-set $\bigcup_{i\in[d]}\{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ satisfies the coordinated random walks property of Definition 6 with respect to any set of odd size $n' \approx n/2$.*

Proof: For a sufficiently large t and any set T of size n', consider a random t-by-n Boolean matrix that corresponds to coordinated random walks (from all possible start vertices) on the n-vertex graph; that is, for any $\overline{\sigma} \in [2d]^t$, we consider the matrix $B = B_T^{(\overline{\sigma})}$. 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 of B in positions J is non-zero. (This establishes CRW property for any sufficiently large $t = \Omega(n \log n)$.)[11]

Claim **9.1** (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 choice of a t-step random walk on the n-vertex graph, the sum* (mod 2) *of the corresponding 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.[12] 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 σ mod 2).

In our argument, we shall refer to a set of permutations over $[n]$, denoted $\texttt{Sym}_n$, 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 proving the following two facts:

1. $|W| \approx |\texttt{Sym}_n|/2$.

2. A random walk $\overline{\sigma} \in [2d]^t$ on the relocation graph corresponds to a matrix $B_T^{(\overline{\sigma})}$ with columns in positions J summing up to the all-zero vector if and only if the random walk on $\texttt{Sym}_n$ does not visit W.

refer explicitly to this graph, but rather refer to the generating set of the symmetric group, and refer to its vertices as to states.

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

[11] We comment that the CRW property can be established for any sufficiently large $t = \Omega(n)$; see Claim 9.2..

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

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 $[n] \setminus J$). The easy case is when $|J| = o(n^{1/2})$. In this case, for a uniformly selected $\pi \in \texttt{Sym}_n$, the set $\pi(J) \stackrel{\rm def}{=} \{\pi(j) : j \in J\}$ is close to a sample of $|J|$ independent points in $[n]$, and so the probability that $|\pi(J) \cap T|$ is odd is approximately $1/2$ (since $|T| = n' \approx n/2$). Turning to the case of $|J| \geq 2n^{1/3}$, we let J' be an $n^{1/3}$-subset of J and $J'' = J \setminus J'$. Using $|J| \leq n/2$, it follows that, with very high probability, $\frac{|\pi(J'') \cap T|}{|J''|} \approx \frac{|T|}{n} \approx 1/2$ holds, which implies that $\frac{|T \setminus \pi(J'')|}{|[n] \setminus \pi(J'')|} \approx 1/2$. It follows that conditioned on the values of $\pi(J'')$, the set $\pi(J')$ is close to a sample of $|J'|$ independent points in $[n] \setminus \pi(J'')$, and so the probability that $|\pi(J') \cap T|$ is odd is approximately $1/2$, assuming the foregoing (extremely likely) event (of $\frac{|T \setminus \pi(J'')|}{|[n] \setminus \pi(J'')|} \approx 1/2$). This establishes Fact 1.

We now turn to the proof of Fact 2. The key observation is that the coordinated random walks $\overline{\sigma}$ on the n-vertex graph yield a Boolean matrix $B_T^{(\overline{\sigma})}$ such that the sum of its 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 of permutation, denoted $\pi_1, ..., \pi_t$, that are selected during the random walk (as determined by the sequence $\overline{\sigma} \in [2d]^t$; i.e., π_i is determined by σ_i (as specified above)). Then, in the $i^{\rm th}$ step of the coordinated walks, the $j^{\rm th}$ walk is in position $k = \pi_i(\cdots(\pi_1(j)\cdots))$, whereas the $(i,j)^{\rm th}$ entry in $B_T^{(\overline{\sigma})}$ is 1 if and only if $\pi_i(\cdots(\pi_1(j)\cdots)) \in T$ (i.e., if and only if $k \in T$). Hence, the sum of the entries (of $B_T^{(\overline{\sigma})}$) in row i and columns in J is 1 (mod 2) if and only if $\pi_i \circ \cdots \circ \pi_1 \in W$. This establishes Fact 2.

Having established both facts, we now establish the claim by upper-bounding the probability that a t-step random walk on the set of $n!$ permutations does not pass through states in W. By the expansion property of the generating set for the symmetric group, 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 $O(\log(n!)) = O(n \log n)$ 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 $J \subset [n]$), we conclude that, with probability at least $1 - (2^n - 2) \cdot \exp(-\Omega(t) + O(n \log n))$, 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 8 to the n-vertex graph (and set) analyzed in Theorem 9 implies that any constant-size generating set for the symmetric group that is expanding yields a 1-local expander. Using Kassabov's construction of such a set [4], yields an alternative proof of Theorem 1.

3.2 Secondary comments

In this section, we make two comments about Theorem 9. The first comment refers to its proof, and the second comment refers to its converse (i.e., we show that a converse of a weak form of Theorem 9 fails).

For sake of elegancy: As noted in Footnote 11, the probability bound of Claim 9.1 can be tightened. Specifically, the error bound of $\exp(-\Omega(t) + O(n \log n))$ can be improved to $\exp(-\Omega(t) + O(n))$, which is optimal.

Claim **9.2** (the distribution of linear combinations of the columns, revisited): *Let T be an odd set of size $n' \approx n/2$. Then, with probability at least $1-\exp(-\Omega(t)+O(n))$ over the choice of a t-step random walk $\overline{\sigma}$ on the n-vertex graph, for every non-empty set $J \subseteq [n]$, the sum* (mod 2) *of the columns of $B_T^{\overline{\sigma})}$ in positions J is non-zero.*

Proof Sketch: We proceed as in the proof of Claim 9.1, 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}.

The key observation is that the Boolean matrix that represents coordinated random walks on the n-vertex graph equals (up to a permutation of its columns) the matrix that represents the same walks 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 $\phi : [n] \to [n]$ is a permutation such that $\phi(j) \in T$ for every $j \in T$). That is, *if the matrix B represents coordinated random walks on the original graph and $\phi : [n] \to [n]$ is a permutation that leaves T invariant, then the matrix obtained by permuting the columns of B according to ϕ represents coordinated random walks 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 B 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 B 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 matrix that corresponds to a random 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). That is, the probability that the matrix B (which represents coordinated random walks on the original graph) is full rank equals the probability that a corresponding random walk on the set of permutations misses one of the W_J's (defined as in the proof of Claim 9.1), when starting from a uniformly distributed state in S.

Indeed, for every non-empty $J \subset [n]$, we consider the corresponding set W_J, which is the set of all permutations π such that $|\pi(J) \cap T|$ is odd. By the expansion property of the generating set for the symmetric group and the fact that S has density $\Omega(2^{-n})$, 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 $1 - (2^n - 2) \cdot \exp(-\Omega(t - O(n)))$, 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 $\mathtt{Sym}_n$. Interpreted in terms of sets of permutations over $[n]$, the CRW property asserts that a random walk on this set passes a specific statistical test (which is specified by the corresponding set T). Theorem 9 asserts that if a set of permutations is expanding, then the CRW property is satisfies for *any* set T of odd size $n' \approx n/2$. This holds also if $n' \approx n/4$ (or any odd value in $[0.01n, 0.99n]$).[13] A weaker implication only asserts that if a set of permutations is expanding, then the CRW property is satisfies for *some* set T of odd size $n' \approx n/4$. Here we show that the converse of the latter implication does not hold.[14] In other words, we show that *the fact that a relocation graph satisfies the CRW property* (with respect to some set of vertices) *does not imply that the corresponding set of permutations generates the symmetric group* (let alone in an expanding manner).

Theorem 10 (on the converse of Theorem 9): *There exists a set of permutations, $\{\pi^{(i)} : i \in [3d]\}$, over $[2n]$ that does not generate the symmetric group of $2n$ elements such that the $2n$-vertex graph consisting of the vertex set $[2n]$ and the edge multi-set $\bigcup_{i\in[3d]}\{\{j, \pi^{(i)}(j)\} : j \in [2n]\}$ satisfies the coordinated random walks property* (of Definition 6) *with respect to some set of odd size $n' \approx n/2$.*

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 $b \in \{0, 1\}$. Denoting the resulting set of augmented permutations by Π', we consider the $2n$-vertex $6d$-regular relocation graph G' that corresponds to it. This graph 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' does not generate the symmetric group of $2n$ elements; it rather generates a group of $2 \cdot (n!) \ll (2n)!$ permutations; specifically, a permutation $\pi' : [2n] \to [2n]$ is generated by Π' if and only if for some $\pi \in \mathtt{Sym}_n$ either $\pi'(b \cdot n + j) = b \cdot n + \pi(j)$ or $\pi'(b \cdot n + j) = (1 - b) \cdot n + \pi(j)$ for every $(b, j) \in \{0, 1\} \times [n]$. 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 $n' \approx n/2$).[15] Hence, we focus on proving the following claim.

Claim: *For any set $T \subset [n]$ of odd size $n' \approx n/2$, the $2n$-vertex graph G' satisfies the CRW property with respect to T.*

When analyzing t-step 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. We call the latter steps semi-idle, since they either map each vertex to itself or map each vertex to its sibling (i.e., its other copy).

The key observation is that t-step random walks on G' correspond to t-step *lazy* random walks on G in which the walk stays in the current vertex (i.e., is truly idle) with probability $2/3$. Indeed,

[13] In this case, for any non-empty set $J \subset [n]$, the density of the corresponding set $W = W_J \subseteq \mathtt{Sym}_n$ may reside in $[0.01, 0.99]$, which suffices for showing that this set is hit with probability $1 - \exp(-\Omega(t) + O(n))$.

[14] Indeed, we leave open the possibility that the converse of Theorem 9 holds. We believe that even if the CRW property is satisfies for any set T of odd size $n' \in [0.01n, 0.99n]$, then it does not necessarily hold that the foregoing set of permutations is expanding.

[15] We 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; for example, 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 (whereas, as shown in Proposition 4, the CRW cannot be satisfied with sets of size $o(n)$).

semi-idle steps (on G') correspond to staying in place (on G), whereas in steps that are not semi-idle (on G') the walk moves on the two copies of G in the same manner and identically to the movement of the corresponding walk on G itself. Furthermore, fixing a lazy random walk on G leaves undetermined the type of semi-idle steps taken on the corresponding walk on G': Each of these steps can either be a truly idle step (i.e., stay in place) or a move to the sibling vertex (i.e., the corresponding vertex in the other copy of G).

Turning to the matrices that describe hitting $T \subset [n]$, note that each row in the t-by-n matrix B the describes a walk on G has exactly $|T| = n'$ ones, and the same holds for the matrix B' in walks on G'. Furthermore, each row in the t-by-$2n$ matrix B' (which corresponds to a walk on G') has either the form $0^n r$ or the form $r0^n$, where r is the corresponding row in the matrix B (which represents the hitting pattern in the corresponding walk on G). The choice between $0^n r$ and $r0^n$ is determined by the number of steps (so far) in which the matching permutation (which moves vertices to their sibling in the other copy of G) was selected.

Recall that Claim 9.1 asserts that, for every non-empty $J \subseteq [n]$, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$ over the walks on G, the sum of columns J in the matrix B (which corresponds to a random walk on G) is a non-zero vector. Actually, the same argument applies to a lazy random walk, by focusing on the non-idle steps. Moreover, it can be shown that this vector, denoted $v^{(J)}$, has Hamming weight $t' \stackrel{\text{def}}{=} \Omega(t)$, since expansion implies that, with extremely high probability, sets of constant density are hit with constant frequency (rather than merely hit). Furthermore, with probability at least $1 - \exp(-\Omega(t - O(n \log n)))$, at least $t'/2$ of the 1-entries in $v^{(J)}$ correspond to idle steps. Let us fix a (lazy) random walk on G that has the foregoing properties, and let r_i denote the i^{th} row of the corresponding matrix B.

Turning to the corresponding t-by-$2n$ matrix B' that describes a random walk on G', consider an arbitrary non-empty set of columns $J \subset [2n]$, and let $J' = J \cap [n]$ and $J'' = \{j - n : j \in J \setminus J'\}$. Then, for every $i \in [t]$, if the sum of the entries in r_i and columns J' is 1, then the sum of the entries in row i and columns J of B' is 1 if the i^{th} row of B' equals $r_i 0^n$. Similarly, if the sum of the entries in r_i and columns J'' is 1, then the sum of the entries in row i and columns J of B' is 1 if the i^{th} row of B' equals $0^n r_i$. Recall that the latter event (i.e., where the i^{th} row of B' equals $r_i 0^n$ or $0^n r_i$) depends only on the choices made in the semi-idle steps.

Lastly, we focus on the rows of B that correspond to idle steps and whose sum in columns J' (resp., J'') equals 1. Recalling that if $J'' \neq \emptyset$ (resp., $J'' \neq \emptyset$), then B contains at least $t'/2 = \Omega(t)$ such rows, we note that for each of these rows the sum of the entries in column J of B' is 1 with probability at least $1/2$, since with probability $1/2$ the i^{th} row of B' equals $r_i 0^n$ (resp., $0^n r_i$). Hence, B' is full rank with probability at least $1 - (2^n - 2) \cdot (2^{-\Omega(t - O(n \log n))} + 2^{-t'/2})$, and the claim follows. ■

4 A sufficient and necessary condition: The relaxed CRW property

We now turn back to the relation between the CRW property (of Definition 6) and 1-local expanders. Recall that Theorem 8 asserts that graphs satisfying the CRW property yield 1-local expanders. A natural question is whether this sufficient condition is necessary. Leaving this question open, we shall show that a *relaxed* CRW property (of the n-vertex relocation graph) suffices and is necessary for for obtaining a 1-local 2^n-vertex expander.

The relaxed CRW property uses a generalization of Definition 6 in which several subsets of $[n]$ are considered (rather than one), and at each step of the coordinated random walks hitting are considered with respect to the most beneficial set. That is, given a coordinated random walk, we can select which subset we consider at each step of the walk, and the corresponding row of the matrix is determined accordingly. This freedom of choice is used when proving the "necessity" direction, whereas it can handled in the "sufficiency" direction by using a large number of offsets (i.e., exponential in the number of sets).

Definition 11 (a relaxed property of coordinated random walks): *For $d, d' = O(1)$, consider a d-regular n-vertex graph as in Definition 6, and d' sets $T_1, ..., T_{d'} \subseteq [n]$.*

- *As in Definition 6, for $t \geq n$, consider a random sequence $\overline{\sigma} = (\sigma_1, ..., \sigma_t) \in [d]^t$ 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.*
- *Fixing the random sequence $\overline{\sigma}$, consider an arbitrary sequence $\overline{\tau} = (\tau_1, ..., \tau_t) \in [d']^t$, and let $B^{(\overline{\sigma}, \overline{\tau})}$ be the t-by-n Boolean matrix such that its $(i, j)^{\text{th}}$ entry indicates whether the j^{th} walk passed through T_{τ_i} in its i^{th} step.*

The relaxed CRW property *asserts that, with probability at least $1 - 2^{-n-\Omega(t)}$ over the choice of $\overline{\sigma} \in [d]^t$, there exists $\overline{\tau} \in [d']^t$ such that the Boolean matrix $B^{(\overline{\sigma}, \overline{\tau})}$ has full rank.*

(Indeed, Definition 6 corresponds to the special case of $d' = 1$.)

Theorem 12 (constructing 1-local expanders is equivalent to constructing relocation graphs that satisfy the relaxed CRW property (as in Definition 11)): *Let $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ be permutations.*

1. The relaxed CRW property is neccessary for 1-local expanders: *If the 1-local $2d$-regular 2^n-vertex graph associated with the permutations $\pi^{(1)}, ..., \pi^{(d)}$ and the offsets $s^{(1)}, ..., s^{(d)} \in \{0, 1\}^n$ is an expander, then the corresponding $2d$-regular n-vertex relocation graph satisfies the relaxed CRW property with respect to the sets $T_1, ..., T_{2d}$ such that $T_{2i} = \{j \in [n] : s_j^{(i)} = 1\}$ and $T_{2i-1} = \{\pi^{(i)}(j) : s_j^{(i)} = 1\}$.*
2. The relaxed CRW property is sufficient for 1-local expanders: *Suppose that the $2d$-regular n-vertex relocation graph associated with $\pi^{(1)}, ..., \pi^{(d)}$ satisfies the relaxed CRW property with respect to the sets $T_1, ..., T_{d'}$. Then, the $2^{2d'+1} \cdot d$-regular 2^n-vertex 1-local graph having the edge multi-set*
$$\bigcup_{i \in [d], \beta, \gamma \in \{0,1\}^{d'}} \left\{ \{x, (x \oplus s^{(\beta)})_{\pi^{(i)}} \oplus s^{(\gamma)}\} : x \in \{0, 1\}^n \right\} \tag{3}$$
is an expander, where for every $\alpha \in \{0, 1\}^{d'}$ the string $s^{(\alpha)} \in \{0, 1\}^n$ denotes the indicator sequence of the set $\bigoplus_{i \in [d'] : \alpha_i = 1} T_i \subseteq [n]$; that is, the j^{th} bit of $s^{(\alpha)}$ is 1 if and only if j resides in an odd number of subsets T_i such that $\alpha_i = 1$ (iff $|\{i \in [d'] : \alpha_i = 1 \ \& j \in T_i\}|$ is odd).

Theorem 8 is a special case of Part 2 (i.e., the case of $d' = 1$). Note that the edge multi-set of Eq. (3) may use $(2^{d'} - 1) \cdot d \cdot 2^{d'} + 2^{d'}$ different offsets (i.e., the offsets $s^{(\gamma)}$ and $s_{\pi^{(i)}}^{(\beta)} \oplus s^{(\gamma)}$ for $i \in [d], \gamma \in \{0, 1\}^{d'}$ and $\beta \in \{0, 1\}^{d'} \setminus \{0^n\}$).

Proof: We start with the proof of Part 2, which generalizes the proof of Theorem 8. Specifically, let $\overline{\sigma} = (\sigma_1, ..., \sigma_t) \in [2d]^t$ 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 $1 - \exp(-n - \Omega(t))$ over the choice of $\overline{\sigma}$, there exists $\overline{\tau} = (\tau_1, ..., \tau_t)$ such that $B^{(\overline{\sigma}, \overline{\tau})}$ is full rank. Recall that specifying a random walk on the 1-local graph requires specifying also the random choices of $\beta_i, \gamma_i \in \{0,1\}^{d'}$ for each step $i \in [t]$. We do so depending of the parity of σ_i and the value of $\tau_i \in [d']$. Specifically, we consider the following two-stage process of determining the sequence of auxiliary random choices of $\beta_1, ..., \beta_t \in \{0,1\}^{d'}$ and $\gamma_1, ..., \gamma_t \in \{0,1\}^{d'}$.

1. For every $i \in [t]$ 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 $i \in [d']$ 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 $\overline{\sigma}$ and a corresponding good $\overline{\tau}$ (i.e., choices such that $B^{(\overline{\sigma}, \overline{\tau})}$ is full rank), consider an arbitrary fixing of the choices in Stage 1. Then, the label of the final vertex in the corresponding random walk on the 1-local graph is a fixed string (determined by $\overline{\sigma}$ and the choices made in Stage 1) that is offset by a random linear combination of the rows of $B^{(\overline{\sigma}, \overline{\tau})}$, where the random linear combination is determined in Stage 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 $B^{(\overline{\sigma}, \overline{\tau})}$ 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 $b = 1$ indicates that the original offset is applied; that is, we shall refer to the edge multi-set $\bigcup_{i \in [d], b \in \{0,1\}} \{\{x, x_{\pi^{(i)}} \oplus (s^{(i)})^b\} : x \in \{0,1\}^n\}$. Cosequently, a t-step random walk on the $4d$-regular expander corresponds to a sequence $(\sigma_1, b_1), ..., (\sigma_t, b_t) \in ([2d] \times \{0,1\})^t$, and the sequence $\sigma_1, ..., \sigma_t$ corresponds to a walk on the n-vertex relocation graph.

We shall use our freedom to determine the τ_i's based on the σ_i's, and doing so we shall obtain a matrix as in Definition 11, which we shall show to be of full-rank (with extremely high probability). Specifically, we let $\tau_i = \sigma_i$, while assuming (again, without loss of generality) that an even σ_i (resp., an odd σ_i) indicates a forward (resp., reverse) transition using $\pi^{(\lceil \sigma_i/2 \rceil)}$. This assumption is made only in order to match the $2d$ possible transitions with the $2d$ sets defined in the conclusion of Part 1. Indeed, under this assumption, if $\sigma_i = 2k$ (resp., $\sigma = 2k - 1$), 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 $\pi^{(k)}$, whereas $T_{2k} = \{j \in [n] : s_j^{(k)} = 1\}$ and $T_{2k-1} = \{\pi^{(k)}(j) : s_j^{(k)} = 1\} = \{j : s_{\pi^{(-k)}(j)}^{(k)} = 1\}$. Hence, picking $\tau_i = \sigma_i$, the i^{th} row in $B^{(\overline{\sigma}, \overline{\tau})} = B^{(\overline{\sigma}, \overline{\sigma})}$ indicates hitting the set T_{τ_i}.

As said above, 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* $B^{(\overline{\sigma}, \overline{\sigma})}$ *must have full rank with probability at least* $1 - \exp(-n - \Omega(t))$. This claim is shown as follows.

Let η denote the probability (over the choice of $\overline{\sigma} \in [2d]^t$) that the matrix $B^{(\overline{\sigma},\overline{\sigma})}$ does not have full rank. Such a choice of $\overline{\sigma}$ determines both the permutation $\pi_{\overline{\sigma}}$ that relates the original locations to the final ones (i.e., $\pi_{\overline{\sigma}} = \pi^{((-1)^{\sigma_t}\cdot\lceil\sigma_t/2\rceil)} \circ \cdots \circ \pi^{((-1)^{\sigma_1}\cdot\lceil\sigma_1/2\rceil)}$) and a non-trivial linear combination $J_{\overline{\sigma}}$ of the columns of the matrix that witnesses the hypothesis that the matrix is not full rank. Hence, there exists a non-empty set $J \subseteq [n]$ such that, with probability $\eta' \geq \eta/(2^n - 1)$ over the choice of $\overline{\sigma}$, the sum of the columns indexed by $\pi_{\overline{\sigma}}^{-1}(J)$ (in the matrix $B^{(\overline{\sigma},\overline{\sigma})}$) equals the all-zero vector (e.g., $\pi_{\overline{\sigma}}^{-1}(J) = J_{\overline{\sigma}}$), whereas in the remaining choices (of $\overline{\sigma}$) this sum does not equals the all-zero vector.[16]

Looking at the label of the final vertex $v_{\overline{\sigma}}$ in a random walk $\overline{\sigma}$ on the 1-local 2^n-vertex graph that starts at the vertex 0^n, we observe that $v_{\overline{\sigma}}$ equals a random linear combination of the rows of $B^{(\overline{\sigma},\overline{\sigma})}$ permuted by $\pi_{\overline{\sigma}}$; that is, $(v_{\overline{\sigma}})_{\pi_{\overline{\sigma}}^{-1}}$ equals a random linear combination of the rows of $B^{(\overline{\sigma},\overline{\sigma})}$, where this random linear combination is determined by the sequence $(b_1, ..., b_t)$ of choices of whether or not to apply the original offset. This is the case since the i^{th} row permuted by $\pi^{((-1)^{\sigma_i}\cdot\lceil\sigma_i/2\rceil)} \circ \cdots \circ \pi^{((-1)^{\sigma_1}\cdot\lceil\sigma_1/2\rceil)}$ is the offset that is potentially added in the i^{th} step of the walk, whereas this offset is added if and only if $b_i = 1$.

It follows that the sum of $v_{\overline{\sigma}}$'s bits in locations J (equiv., the sum of the bits of $(v_{\overline{\sigma}})_{\pi_{\overline{\sigma}}^{-1}}$ in locations $\pi_{\overline{\sigma}}^{-1}(J)$) is zero with probability exactly $\eta' + (1-\eta')\cdot 0.5 = 0.5 + 0.5\eta'$, since this sum is 0 whenever the sum of the corresponding columns in $B^{(\overline{\sigma},\overline{\sigma})}$ is the all-zero vector (and is uniformly distributed in $\{0,1\}$ otherwise).[17] Hence, the total variation distance between the distribution of the final vertex and the uniform distribution is at least $0.5\eta'$.

Recalling that the hypothesis (i.e., that the 1-local graph is an expander) implies that $\eta' \leq \exp(-\Omega(t))$, it follows that $\eta < 2^n \cdot \eta' = \exp(-n - \Omega(t))$, for sufficiently large $t = \Omega(n)$. This establishes Part 1. ■

Conclusion. Theorem 12 asserts that constructing graphs that satisfy the relaxed CRW property is equivalent to constructing 1-local expanders. One begging qurestion is whether the relaxed CRW property is easier to achieve that the original CRW property. Lacking a positive answer, the raises the following generalization of Problem 7.

Open Problem 13 (the CRW problem, revised): *For which graphs and which sequences of sets $(T_1, ..., T_{d'})$'s does the relaxed CRW property* (as in Definition 11) *hold?*

An appealing conjecture of Benny Applebaum is that every n-vertex expander graph yield a positive instance of Problem 13; that is, there exists $d' = O(1)$ sets $T_1, ..., T_{d'} \subset [n]$ such that this n-vertex graph satisfies the relaxed CRW property (of Definition 11) with respect to these T_i's.

[16] The issue here is as follows: We know that for η fraction of the $\overline{\sigma}$'s, there exists a $J_{\overline{\sigma}} \neq \emptyset$ such that the sum of these columns is the all-zero vector (and let $J_{\overline{\sigma}} = \emptyset$ otherwise). However, these columns corresponds to locations in the (label of the) initial vertex, whereas we want to analyze locations in the end vertex. Of course, locations $J_{\overline{\sigma}}$ in the initial vertex correspond to locations $\pi_{\overline{\sigma}}(J_{\overline{\sigma}})$ in the final vertex. Hence, there exists a non-empty J (representing locations in final label) such that the sum of the columns in $\pi_{\overline{\sigma}}^{-1}(J)$ (representing locations in initial label) equals the all-zero vector with probability $\eta' \geq \eta/(2^n - 1)$. This lower bound is due to the event $\pi_{\overline{\sigma}}^{-1}(J) = J_{\overline{\sigma}}$, but the sum of these columns may be zero also otherwise. (For this reason, we define η' as the probability that the sum of the columns in $\pi_{\overline{\sigma}}^{-1}(J)$ equals the all-zero vector rather than the probability that $\pi_{\overline{\sigma}}^{-1}(J) = J_{\overline{\sigma}}$.) Needless to say, for the rest of this probability space (of $\overline{\sigma} \in [2d]^t$), this sum is not the all-zero vector.

[17] If the sum of these columns is not the all-zero vector, then a random combination of its entries, as determined by $(b_1, ..., b_t)$, is uniformly distributed in $\{0,1\}$.

5 Generalization to non-binary alphabets

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

Definition 14 (1-local graph, generalized): *For a fixed $d \in \mathbb{N}$ and a fixed prime p, let $\{f_1, ..., f_d : \mathrm{GF}(p)^n \to \mathrm{GF}(p)^n\}_{n\in\mathbb{N}}$ be 1-local bijections. Then, the corresponding $2d$-regular p^n-vertex* 1-local graph *consists of the vertex set $\mathrm{GF}(p)^n$ and the edge multi-set $\bigcup_{i\in[d]}\{\{x, f_i(x)\} : x \in \mathrm{GF}(p)^n\}$.*

Note that each f_i is determined by a permutation on the locations $\pi^{(i)} : [n] \to [n]$, called the relocation, and n bijections denoted $h_1^{(i)}, ..., h_n^{(i)} : \mathrm{GF}(p) \to \mathrm{GF}(p)$ such that, for every $j \in [n]$, the j^{th} bit of $f_i(x)$ equals $h_j^{(i)}(x_{\pi^{(i)}(j)})$. Unlike in the binary case (i.e., $p = 2$), where each $h_j^{(i)}$ is affine (i.e., has the form $h_j^{(i)}(z) = z \oplus s_j^{(i)}$), these bijections are not necessarily affine functions. Still, we shall focus on the case that they are affine. Generalizing Theorems 8 and 9, we obtain.

Theorem 15 (a construction of generalized 1-local expanders): *For every constant prime p, there exists a set of $d = O(p^2)$ explicit 1-local bijections, $\{f_1, ..., f_d : \mathrm{GF}(p)^n \to \mathrm{GF}(p)^n\}_{n\in\mathbb{N}}$, such that the $2d$-regular p^n-vertex graph that consists of the vertex set $\mathrm{GF}(p)^n$ and the edge multi-set $\bigcup_{i\in[d]}\{\{x, f_i(x)\} : x \in \mathrm{GF}(p)^n\}$ is an expander. Furthermore, for each $i \in [d]$, there exists a permutation $\pi^{(i)} : [n] \to [n]$ and an offset $s^{(i)} \in \mathrm{GF}(p)^n$ such that $f_i(x) = x_{\pi^{(i)}} + s^{(i)}$.*

iff $\pi^{(i)}(j) = k$. The expansion feature holds also for varying $p = p(n)$, but in that case the graph is not of constant degree.

Proof: The overall plan is to use a straightforward generalization of the CRW property for rank defined over $\mathrm{GF}(p)$, and prove adequate generalizations of Theorems 8 and 9. Specifically, we first show that any n-vertex graph that satisfies the generalized CRW property yields a 1-local p^n-vertex expander, and then show that any generating set for the symmetric group of n elements that is expanding yields an n-vertex graph that satisfies the generalized CRW property (with respect to any set of size $n' \approx n/2$ such that $n' \not\equiv 0 \pmod p$).

Definition 15.1 (a property of coordinated random walks, generalized): *For a d-regular n-vertex graph as in Definition 6, a set $T \subseteq [n]$ and $t = \Omega(n)$, consider coordinated random walks and Boolean matrices just as in Definition 6. The* generalized CRW property *postulates that, with probability at least $1 - p^{-n} \cdot \exp(-\Omega(t))$, such a random matrix has full rank* when the arithmetics is in $\mathrm{GF}(p)$.

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

Claim 15.2 (Theorem 8, generalized): *Let $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ be d permutations and $s = (s_1, ..., s_n) \in \{0, 1\}^n \subseteq \mathrm{GF}(p)^n$. If the $2d$-regular n-vertex graph with the edge multi-set $\bigcup_{i\in[d]}\{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ satisfies the generalized CRW property* (of Definition 15.1) *with respect to the set $\{j \in [n] : s_j = 1\}$, then the $2p^2d$-regular p^n-vertex graph with the edge multi-set*

$$\bigcup_{i\in[d],b,c\in\mathrm{GF}(p)} \{\{x, (x - b \cdot s)_{\pi^{(i)}} + c \cdot s\} : x \in \mathrm{GF}(p)^n\}$$

is an expander, where $b \cdot (s_1, ..., s_n) = (bs_1, ..., bs_n)$.

Proof Sketch: We mimic the proof of Theorem 8, while noting that in the i^{th} step the vertex's label is randomized by an offset that is a random GF(p)-multiple of the i^{th} row in the corresponding matrix; specifically, in a forward direction the randomization is performed by the value of c (i.e., adding the offset $c \cdot s$), whereas in a reverse direction the randomization is performed by the value of b (i.e., subtracting the offset $-b \cdot s$). Hence, if the matrix has full rank over GF(p), then the label of the final vertex is uniformly distributed in GF$(p)^n$ (since it is randomized by a random linear combination of the rows of the matrix). ■

Claim **15.3** (Theorem 9, 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* $\bigcup_{i\in[d]}\{\{j, \pi^{(i)}(j)\} : j \in [n]\}$ *satisfies the generalized CRW property of Definition 15.1 with respect to any set* T *of size* $n' \approx n/2$ *such that* $n' \not\equiv 0 \pmod p$.

Proof Sketch: Here we mimic the proof of Theorem 9. Specifically, we consider all (non-zero) linear combinations $L : [n] \to \text{GF}(n)$ of the columns of a matrix that corresponds to a random walk, and upper bound the probability that each such linear combination yields the all-zero vector. That is, fixing any set T of size n', for every such linear combination L, we consider the set W_L of permutations $\pi \in \texttt{Sym}_n$ 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 by using $t = \Omega(n \log(np))$, where here we use a union bound on all (non-zero) L's. Hence, we focus on proving that for each non-zero $L : [n] \to \text{GF}(p)$, the set W_L has constant density in $\texttt{Sym}_n$.

The case of non-zero constant functions $L : [n] \to \text{GF}(p)$ is handled by the hypothesis that $n' \not\equiv 0 \pmod p$, which implies that $W_L = \texttt{Sym}_n$, and so we focus on non-constant functions L. In this case, one may show that W_L has density at least $0.5 - o(1)$, but we use a simpler argument to show that it has density at least $0.25 - o(1)$. Specifically, considering any $i_1, i_2 \in [n]$ such that $L(i_1) \neq L(i_2)$, we observe that $\Pr_\pi[|\{\pi(i_1), \pi(i_2)\} \cap T| = 1] > 0.5 - o(1)$. On the other hand, conditioned on the values of π on $I \stackrel{\text{def}}{=} [n] \setminus \{i_1, i_2\}$ and on the foregoing event (i.e., $|\{\pi(i_1), \pi(i_2)\} \cap T| = 1$), the value of $\sum_{i\in[n]:\pi(i)\in T} L(i)$ is a random variable that equals $\sum_{i\in I:\pi(i)\in T} L(i) + L(i_1)$ with probability 1/2 (when $\pi(i_1) \in T$) and equals $\sum_{i\in I:\pi(i)\in T} L(i) + L(i_2)$ otherwise. Recalling that $L(i_1) \neq L(i_2)$, we get

$$\Pr_{\pi:[n]\to[n]}\left[\sum_{i\in[n]:\pi(i)\in T} L(i) \not\equiv 0 \pmod p\right] > (0.5 - o(1)) \cdot \frac{1}{2}$$

and the claim follows. ■

Combining Claims 15.3 and 15.2, we get.

Corollary 15.4 (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*

$$\bigcup_{i\in[d],b,c\in\text{GF}(p)} \left\{\{x, (x - b^{n'}0^{n-n'})_{\pi^{(i)}} + c^{n'}0^{n-n'}\} : x \in \text{GF}(p)^n\right\}$$

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 foregoing generalizes to any finite field; that is, p may be a prime power. For $p = q^e$, 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 15.3 the reductions mod p actually refer to doing arithmetics in $\mathrm{GF}(p)$).

Acknowledgments

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

Appendix: Secondary Observations

A natural way of trying to improve over Observation 2 is to use relocation graphs that have shorter cover time. The natrual choice is to use n-vertex expander graphs.[18]

Observation 16 (using an expander for the relocation graph): *Let* $\pi^{(1)}, ..., \pi^{(d)} : [n] \to [n]$ *be bijections that represent the edges of an* $2d$*-regular expander, and* $\mathtt{id} : [n] \to [n]$ *denote the identity bijection. Then, the 1-local* 2^n*-vertex graph associated with the* $2d$ *relocation permutation* $\pi^{(1)}, ..., \pi^{(d)}, \mathtt{id}, ..., \mathtt{id}$ *and the* $2d$ *offsets* $0^n, ..., 0^n, 0^{n-1}1, ..., 0^{n-1}1$ (i.e,., the i^{th} bijection is $x \mapsto x_{\pi^{(i)}}$ if $i \in [d]$ and $x \mapsto x \oplus 0^{n-1}1$ otherwise) *has second* (normalized) *eigenvalue* $1 - \Theta(1/n \log n)$.

Proof Sketch: In this case, a random walk of length $t = O(t' \cdot n \log n)$ 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").[19] It follows that 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). ■

Proposition 4 (restated): *Consider a* $2d$*-regular* 2^n*-vertex graph as in Definition 3, and suppose that for every* $i \in [d]$ *either* $|s^{(i)}| = o(n)$ *or* $|s^{(i)}| = n - o(n)$*. Then, this 1-local* 2^n*-vertex graph is not an expander.*

Proof: For starters, we assume that $|s^{(i)}| = o(n)$ for every $i \in [d]$. We first 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.

The key observation is that, during a random walk on the 1-local 2^n-vertex graph, bits in the label of the current vertex get randomized by the offsets with too small probability, since at each step only $o(n)$ locations are randomized. Specifically, for a t-step random walk 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 relocation graph that starts at vertex

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

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

$j \in [n]$ does not go through any vertex in the set $S \stackrel{\rm def}{=} \bigcup_{i\in[d]}\{k : s_k^{(i)} = 1\}$. This event occurs with probability at least $\eta = \exp(-o(t))/n$, since the probability that a walk of length t that starts at a random vertex on any regular n-vertex graph misses a set of $o(n)$ vertices is at least $(1-o(1))^t = \exp(-o(t))$.[20]

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 $(1-\eta)\cdot 0.5\cdot n + \eta \cdot 0.5 \cdot (n-1) = (n-\eta)/2$. It follows that the total variation distance between the distribution of the final vertex and the uniform distribution is at least $\eta/2n = \exp(-o(t) - \log n)$.[21] We stress that the foregoing holds for any t, which means that we assume that $n = o(t)$, let alone $\log n = o(t)$. Hence, the convergence rate of the 1-local 2^n-verterx graph is *not* bounded away from 1 (since $\eta = \exp(-o(t))$ whereas the convergence rate λ must satisfy $2^n \cdot \lambda^t > \eta/2n$).[22] Lastly, we note that given that the auxiliary graph is not an expander, the original graph (which is a subgraph of it) is also not an expander.

Turning to the case in which also offsets of Hamming weight $n - o(n)$ exist, we note that this is equivalent to using an offset of weight $o(n)$ and complementing all bits in the resulting label. Hence, such offsets can randomize many individual locations but cannot randomize all pairs of locations (i.e., randomize each location independently of its paired location). Hence, we extend the foregoing argument to pairs of locations.

We first observe that there are two positions $j_1 \neq j_2$ such that with probabiliy $\eta' = \exp(-o(t))$ these position are always randomized together (i.e., in each steps either both j_1 and j_2 are in locations that get randomized by some single offset or both are not in such locations).[23] The argument is completed by considering the expected number of pairs of positions that hold the same value. ■

References

[1] A. Broder and A. Karlin. Bounds on the cover time. *J. of Theoretical Probability*, Vol. 2 (1), pages 101–120, 1989.

[20] Note that here we seek a lower bound on the probability of missing the set S (equiv., staying in $\overline{S} = [n] \setminus S$), whereas the usual 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 $\overline{S}$ (because otherwise, conditioning on not leaving $\overline{S}$ biases the distribution towards vertices that have more neighbors in $\overline{S}$ (equiv., less neighbors in S)). Hence, the probability that the random walk never leaves $\overline{S}$ is at least $(1 - \frac{d|S|}{d\cdot|\overline{S}|})^t$, whereas in our case $|\overline{S}| = (1-o(1))\cdot n$.

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

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

[22] Hence, we have $2^n \cdot \lambda^t > \exp(-o(t))$, which implies $\lambda = \exp(-o(1))$ for $t = \omega(n)$.

[23] To see this, follow the argument in Footnote 20, while defining $S \stackrel{\rm def}{=} \bigcup_{i\in[d]}\{k : s_k^{(i)} = b^{(i)}\}$, where $b^{(i)}$ is the majority value in the string $s^{(i)}$, while noting that the probability that one of the two coordinated random walks does not stay in $\overline{S}$ is only doubled.

[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.