Oded: August 29, 2016

Summary. Contemplating the recently announced 1-local expanders of Viola and Wigderson, one may observe that weaker constructs that use logarithmic degree are well-known (e.g., the hypercube). Likewise, one may easily obtain a 4-regular N-vertex graph with spectral gap that is $\Omega(1/\log^2 N)$, and similarly a $O(1)$-regular N-vertex graph with spectral gap $1/\widetilde{O}(\log N)$. Following this line of thought, we formulate a natural problem regarding "coordinated random walks" (CRW), and observe that (1) any solution to the CRW problem yields 1-local expanders, and (2) that any constant-size expanding set of generators for the symmetric group yields a solution to the CRW problem. This yields an arguably simpler construction and a more intuitive analysis than the one used by Viola and Wigderson. Lastly, a modest and natural generalization of the CRW problem is equivalent to the problem of constructing 1-local expanders.

1 The original statement

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. Throughout this note, we view n as varying. We study the following recent result of Viola and Wigderson [5].

Theorem 1 (1-local expanders [5]): *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 in which $x \in \{0,1\}^n$ is connected to the vertex set $\{f_i^\sigma(x) : i \in [d], \sigma \in \{\pm1\}\}$ is an expander.*

Note that each f_i is determined by a permutation on the bit locations $\pi^{(i)} : [n] \to [n]$, called the relocation, and an offset $s^{(i)} \in \{0,1\}$ such that $f_i(x_1 \cdots x_n) = (x_{\pi^{(i)}(1)} \cdots x_{\pi^{(i)}(n)}) \oplus s^{(i)}$.

Recall that the (normalized) second eigenvalue of a regular graph represents the rate at which a random walk on the graph converges to the uniform distribution (hereafter called the convergence rate). In an expander this rate is a constant smaller than 1, whereas in a general (regular) N-vertex graph the rate is upper-bounded by $1 - \frac{1}{\text{poly}(N)}$.

2 Initial thoughts

Obtaining a 1-local expander requires using *both* the offsets (i.e., $s^{(i)}$'s) and the relocation permutations, because without the offsets the f_i's maintain the Hamming weight of the vertex (and so the 2^n-vertex graph is not even connected), whereas without the permutations the 2^n-vertex graph decomposes into even smaller connected components (i.e., each of size 2^d).

Note that using $d = 2$ with $f_1(x) = \mathtt{sh}(x)$ and $f_2(x) = \mathtt{sh}(x) \oplus 0^{n-1}1$, where $\mathtt{sh}(x_1 \cdots x_n) = (x_2 \cdots x_n x_1)$ is a shift that corresponds to the permutation $\pi(i) = (i + 1 \bmod n) + 1$, yields a 2^n-vertex graph with second eigenvalue $1 - \Omega(1/n^2)$, since taking a (lazy) random walk of length $O(t \cdot n^2)$ on this graph yields a distribution that is 2^{-t}-close to uniform.[1] Note that this rate of

[1] The latter assertion is based on the fact that during such a walk, with probability at least $1-2^{-t}$, each position in the original string appeared at the rightmost position at some time during the walk (and at that time the corresponding value is randomized). To see this consider the binary (over $\{\pm1\}$) sequence of decisions describing whether to apply $\mathtt{sh}$ or $\mathtt{sh}^{-1}$ in each of the (non-lazy) random steps, and note that each block of $O(n^2)$ symbols has absolute value of at least $2n$ with probability at least 1/2. Hence, looking at t partial sums that correspond to the endpoints of such t disjoint blocks, we conclude that the probability that all these partial sums are in the interval $[-n, n]$ is at most 2^{-t}.

convergence is bounded away from 1 by the reciprocal of a polylogarithmic function in the size of the graph; specifically, we have rate $1-\Omega(1/\log^2 N)$ for N-vertex graphs.

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

Wishing to use shorter random walks in the rate-convergence analysis, consider the case that the n-vertex graph is a $O(1)$-regular expander graph. In this case, a (lazy) random walk of length $O(t\cdot n\log n)$ 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").[2] It follows that the corresponding 1-local 2^n-vertex graph has second eigenvalue $1-(1/n\log n)$, since taking a (lazy) 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 (which, in turn, follows from the fact that each position in the original n-bit string is mapped to the rightmost position at some time).

There is no hope of getting a constant-degree 2^n-vertex expander when using only offsets of Hamming weight $o(n)$. This is the case because the probability that a walk of length t on any regular n-vertex graph misses a set of $o(n)$ vertices is at least $(1-o(1))^t=\exp(-o(t))$.[3] In that case, there exists a position in the original n-bit string (i.e., in the name of the vertex of the 1-local 2^n-vertex graph) that is not moved to an active location where it may be randomized, where the active locations refer to the 1-entries in the offsets.[4] Using also offsets of Hamming weight $n-o(n)$ does not help, since this is equivalent to adding the all-ones offset, which merely complements the vertex name in the 2^n-vertex graph.[5] In view of the above, we must use at least one offset that has Hamming weight in $[\Omega(n),n-\Omega(n)]$. We shall first consider the case of using a single offset that has weight approximately $n/2$.

3 A sufficient condition

Taking $t=\Theta(n)$ random steps, consider the t-by-n Boolean matrix describing the activity status of the location to which each of the initial positions is moved during the t steps; that is, the

[2] The cover time bound was established in [1, 2, 4].

[3] 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 common focus is on good upper bounds (which exists when the graph is an expander). Letting d denote the degree of the n-vertex graph, we observe that there are at most $d\cdot|S|$ edges incident at S, and the worst case is that their other endpoints are distributed evenly among the vertices in $\overline{S}$ (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 $|\overline{S}|=(1-o(1))\cdot n$.

[4] This rules out not only the line of thinking used above, but also the possibility that the 2^n-vertex graph is an expander. To see this consider a random walk that starts at the vertex 0^n and suppose that with probability at least $\eta=\exp(-o(t))/n$ this walk does not randomize position i. (We stress that randomized bit positions are reset to 1 with probability exactly 1/2, whereas non-randomized positions maintain the value 0.) So in the final vertex of the walk, this (i.e., i^{th}) bit position will be 0 with probability at least $(1-\eta)\cdot 0.5$, which means that the final vertex is η-far from uniform.

[5] In that case, with similar probability, there are two positions in the original string that are not moved through an active location (which implies that their final values are identical). To see this, follow the argument in Footnote 3, while noting that the probability that one of the two coordinated random walks does not stay in $\overline{S}$ is only doubled.

$(i,j)^{\text{th}}$ 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 being used). Using an n-vertex expander, we observe that (w.v.h.p.) each column in this random matrix has approximately $t/2$ 1-entries, but what we need is that (w.v.h.p.) this matrix has rank n.

Note that the matrix that corresponds to a random walk describes n coordinated walks on an n-vertex graph, each starting at a different vertex of the graph and all proceeding according to the same sequence of (random) choices. When this matrix has full rank, the t random choices of whether or not the non-zero offset is applied at each of the t steps correspond to a random linear combination of the t rows of the matrix, which yields a uniformly distributed n-bit long string. In this case, the corresponding random walk on the 2^n-vertex graph yields a uniform distribution (since the latter n-bit string is added to the initial vertex in the walk yielding a uniform distribution on the vertices of the 2^n-vertex graph, regardless of the effect of the relocation permutations).[6] Hence, the question we consider is the following.

Problem 2 (a property of coordinated random walks): *Let $d = O(1)$. For a d-regular n-vertex graph, an integer $t = \Omega(n)$, and a set $T \subseteq [n]$, consider a random sequence $(\sigma_1, ..., \sigma_t) \in [d]^t$ and the n corresponding* coordinate random walks (CRW) *such that the j^{th} walk starts at vertex j and moves in the i^{th} step to the σ_i^{th} neighbor of the current vertex. Now, consider a t-by-n Boolean matrix such that the $(i,j)^{\text{th}}$ entry indicates whether the j^{th} walk passed in T in its i^{th} step; that is, letting $g_\sigma(v)$ denote the σ^{th} neighbor of vertex v, then the $(i,j)^{\text{th}}$ is 1 if and only if $g_{\sigma_i}(\cdots(g_{\sigma_1}(j)\cdots)) \in T$. The desired CRW property is that this random matrix has full rank with probability $1 - \exp(-\Omega(t))$, and the question is for which graphs and which sets T's does this hold?*

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

4 Known constructions that satisfy the CRW property

Recall that Kassabov's result [3], which is used in [5], asserts that the symmetric group has an explicit generating set that is expanding and of constant size.[7] We shall show that using this set of permutations (i.e., as our set of relocating permutations) with the offset $1^{n'}0^{n-n'}$ 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 Problem 2). This yields an alternative proof of Theorem 1.

Consider a random t-by-n Boolean matrix that corresponds to coordinated random walks (from all possible start vertices) on the n-vertex graph (wrt the foregoing offset). We shall show that, for every non-empty set $I \subseteq [n]$, with probability at least $1 - 2^{-3n}$, the sum of columns in position I is non-zero. For $I = [n]$ this follows from the fact that n' is odd. Otherwise (i.e., for $I \subset [n]$), we shall

[6] That is, fixing a random walk on the 2^n-vertex graph, we observe that if the matrix that corresponds to this walk has full rank, then the final vertex in the walk is uniformly distributed in $\{0,1\}^n$.

[7] 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 $[n]$).

prove the claim by using the correspondance between random walks on the n-vertex graph and random walks on the set of all permutations moving according to the selected generators.[8] (That is, selecting the σ^{th} neighbor in the random walk on the n-vertex graph corresponds to selecting the σ^{th} generating permutation (and moving by composing it).)

In our argument, we shall refer to two sets of permutations over $[n]$ (viewed as n-long sequences over $[n]$ with distinct elements):

- The set B of sequences such that locations I hold an odd number of elements of $[n']$; that is, $(e_1, ..., e_n) \in B$ if $|\{i \in I : e_i \in [n']\}|$ is odd. Observe that B has density approximately half within the set of all $n!$ sequences.[9]

  Note that the coordinated random walks on the n-vertex graph yield a Boolean matrix such that the sum of columns in position I is zero if and only if the corresponding walk on the set of $n!$ permutations does not pass through states in B.

- The set S of sequences such that the first n' locations hold all elements of $[n']$; that is, $(e_1, ..., e_n) \in S$ if $\{i \in [n'] : e_i\} = [n']$. Observe that S has density approximately $\frac{(n'!)^2}{n!}$, which is approximately 2^{-n}.

  Note that the Boolean matrix that represents a random walk on the n-vertex graph equals (up to a permutation of its columns) the matrix that represents the same walk on any isomopric copy of that graph that leaves $[n']$ invariant (i.e., rather than walking on an n-vertex graph G, we walk on $\pi(G)$, where $\pi : [n] \to [n]$ is a permutation such that $\pi(j) \in [n']$ for every $j \in [n']$). Hence, we may analyze the corresponding walk (on the set of $n!$ permutations) that starts at a state that is uniformly distributed in S.

Now, by the expansion property of the generating set for the symmetric group, we have that a t-step random walk that starts in uniformly distributed state in S passes via B 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 B. Hence, the corresponding t-by-n Boolean matrix has full rank with probability at least $1 - 2^n \cdot \exp(-\Omega(t) + O(n))$. (Formally, for each I, we consider the corresponding B_I, and observe that a random walk that starts at a state that is uniformly distributed in S avoids none of the B_I's corresponds to a Boolean matrix that is full rank, and that the probability that the complementary event occurs (i.e., there exists an I such that the random walk avoids B_I) is upper bounded by a union bound on all B_I's.)

Theorem 3 (a partial answer to Problem 2): *Let $\Pi = \{\pi_i : i \in [d]\}$ be a generating set of the symmetric group of n elements and suppose that Π is expanding. Consider an n-vertex graph such that, for every $j \in [d]$ and $\sigma \in \{0, 1\}$, the $(2j - \sigma)^{\mathrm{th}}$ neighbor of $i \in [n]$ is $\pi_j^\sigma(i)$. Then, this n-vertex graph combined with any set T of odd size $n' \approx n/2$ satisfies the coordinated random walks property.*

[8] That is, we use the correspondance between random walks on the n-vertex graph and random walks on the $n!$-vertex Cayley graph.

[9] This can be shown by considering, w.l.o.g., the case of $|I| \leq n/2$ (or else consider $[n] \setminus I$). Consider a process of randomly assigning distinct elements to the location in I, and focus on the last assignment in that process. W.v.h.p., before this last assignment, these $|I| - 1 < n/2 \approx n'$ locations were assigned approximately an equal number of elements from $[n']$ and from $[n' + 1, n]$, which means that $n' - (1 \pm o(1)) \cdot |I|/2 = (1 \pm o(1)) \cdot (n - |I|)/2$ elements from each type remain for the last assignment. This means that the parity of elements from $[n']$ is flipped at the last step with probability $(1 \pm o(1))/2 \approx 1/2$.

5 A sufficient and necessary condition

Turning back to Problem 2, we note that the following generalization suffices for obtaining a 1-local expander (with 2^n vertices).

Problem 4 (a relaxed property of coordinated random walks): *Let $d, c = O(1)$. For a d-regular n-vertex graph, an integer $t = O(n)$, and c sets $T_1, ..., T_c \subseteq [n]$, consider a random sequence $(\sigma_1, ..., \sigma_t) \in [d]^t$ and the n corresponding* coordinate random walks *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. Now, consider another random sequence $(\tau_1, ..., \tau_t) \in [c]^t$, and a t-by-n Boolean matrix such that the $(i,j)^{\rm th}$ entry indicates whether the $j^{\rm th}$ walk passed in T_{τ_i} in its $i^{\rm th}$ step. For which graphs and which sequences of sets $(T_1, ..., T_c)$'s does this random matrix have full rank with probability $1 - \exp(-\Omega(t))$?*

On the other hand, we note that any 1-local 2^n-vertex expander yields a positive solution to Problem 4: Firstly, note that w.l.o.g., we may consider a 1-local graph in which each of the $O(1)$ offsets is coupled with each of the $O(1)$ relocation permutations. A random walk on this 2^n-vertex expander yields a matrix as in Problem 4. Now, if a t-step random walk yields a distribution that is $\exp(-\Omega(t))$-close to uniform (and $t = \Omega(n)$ is large enough), then the corresponding matrix must have full rank with probability at least $1 - \exp(-\Omega(t))$. This claim is shown as follows.

Let η denote the probability that the said matrix does not have full rank. Then, with probability η' that is at least $2^{-n} \cdot \eta$ over the choices of the relocation permutations, some linear dependency appears between the n positions in the name of the final vertex, whereas in the remaining walks this dependence does not appear. (That is, we consider the probability distribution over the Boolean matrices, while permuting each matrix according to the final vertex (reached in the walks on the n-vertex graph), and consider a linear dependency among the columns of the resulting random matrix that holds with probability at least $2^{-n} \cdot \eta$.) Hence, this linear dependency holds for the name of the final vertex of a random walk on the 2^n-vertex graph with probability at least $(1 - \eta') \cdot 0.5 + \eta'$, which means that the distribution of the final vertex is η'-far from the uniform distribution. The claim follows, since $\eta' \leq \exp(-\Omega(t))$ implies $\eta \leq 2^n \cdot \exp(-\Omega(t)) = \exp(-\Omega(t))$ for sufficiently large $t = O(n)$.

References

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

[2] A.K. Chandra, P. Raghavan, W.L. Ruzzo, R. Smolensky, and P. Tiwari. The electrical resistance of a graph, and its applications to random walks. In *21st STOC*, 1989.

[3] M. Kassabov. Symmetric groups and expander graphs. *Invent. Math.*, Vol. 170 (2), pages 327–354, 2007.

[4] R. Rubinfeld. The cover time of a regular expander is $O(n \log n)$. *IPL*, Vol. 35, pages 49–51, 1990).

[5] E. Viola and A. Wigderson. Local Expanders. *ECCC*, TR16-129, 2016.