Local Expanders

Emanuele Viola* Avi Wigderson†

May 24, 2017

Abstract

A map $f : \{0,1\}^n \to \{0,1\}^n$ has *locality* t if every output bit of f depends only on t input bits. Arora, Steurer, and Wigderson (2009) asked if there exist bounded-degree expander graphs on 2^n nodes such that the neighbors of a node $x \in \{0,1\}^n$ can be computed by maps of constant locality.

We give an explicit construction of such graphs with locality one. We then give three applications of this construction: (1) lossless expanders with constant locality, (2) more efficient error reduction for randomized algorithms, and (3) more efficient hardness amplification of one-way permutations. We also give, for n of the form $n = 4 \cdot 3^t$, an explicit construction of bipartite Ramanujan graphs of degree 3 with $2^n - 1$ nodes in each side such that the neighbors of a node $x \in \{0,1\}^n \setminus \{0^n\}$ can be computed either (1) in constant locality or (2) in constant time using standard operations on words of length $\Omega(n)$.

Our results use in black-box fashion deep explicit constructions of Cayley expander graphs, by Kassabov (2007) for the symmetric group S_n and by Morgenstern (1994) for the special linear group $\mathrm{SL}(2, F_{2^n})$.

1 Introduction and our results

Expander graphs are important objects in theoretical computer science with myriad of applications; for background see e.g. the survey [HLW06]. Some of these applications require the ability to efficiently compute the transition functions, that is, the neighbors of a given n-bit node. Indeed, many algorithms for this task have been devised under various resource constraints, see e.g. [BYGW99], [GV04], [DvM06], and [ASW09]. Still, several natural questions remain open. Here we affirmatively answer a question by [ASW09] who ask if the neighbors can be computed by functions with *constant locality*. A function $f : \{0,1\}^n \to \{0,1\}^n$ has *locality* t if each output bit depends on at most t input bits. The class of functions with constant locality is also known as NC^0.

*Supported by NSF grant CCF-1319206. Work done in part while at the Simons Institute for the Theory of Computing.

†This research was partially supported by NSF grant CCF-1412958.

First, we give a construction of expander graphs where the transition functions have locality one.

Theorem 1. *For every sufficiently large d, and for every n, there exist explicit one-local maps $C_1, C_2, \ldots, C_d$ each mapping n bits to n bits, such that the graph on nodes $\{0,1\}^n$ where node x has the d neighbors $C_1(x), C_2(x), \ldots, C_d(x)$ is an expander graph with second largest eigenvalue at most $d^{-\Omega(1)}$.*

The most interesting setting is when the degree of the graph is $d = O(1)$, but we state a more general tradeoff between degree and eigenvalue bound. We say that a t-local map $f : \{0,1\}^n \to \{0,1\}^n$ is explicit (a.k.a. uniform) if its description can be computed in time polynomial in n. By description we simply mean its graph of connections, and for each output bit a truth table of length 2^t of the function computed at that bit.

In a nutshell, our graph will be a Schreier graph of the semi-direct product of $GF(2)^n$ and the symmetric group S_n. To analyze the semi-direct product we rely on results in [ALW01] which provide an algebraic view of the Zig-Zag product [RVW02]. We crucially use the fact that S_n has a constant number of expanding generators, a result due to Kassabov [Kas07].

Second, we give a construction of bipartite, Ramanujan graphs [LPS88] of degree 3, where the transitions in one direction have constant locality. Let us fix some terminology about bipartite graphs. We think of a bipartite graph as a graph whose vertex set is of the form $V \times \{0,1\}$ and where a vertex (v,b) has neighbors of the form $(w, 1-b)$. We call $V \times \{0\}$ the zero side of the graph, and $V \times \{1\}$ the one side of the graph. Each side of our Ramanujan graph consists of $2^n - 1$ vertices; it can be enlarged to have size 2^n with a slight loss in other parameters.

Theorem 2. *For every n of the form $n = 4 \cdot 3^t$ there exist three explicit constant-locality maps C_1, C_2, and C_3, each mapping n bits to n bits, such that the bipartite graph on the $2(2^n-1)$ vertices $(\{0,1\}^n \backslash \{0^n\}) \times \{0,1\}$ where a node $(v,0)$ has the three neighbors $(C_1(x),1)$, $(C_2(x),1)$, and $(C_3(x),1)$ is a Ramanujan expander graph.*

This theorem uses the Ramanujan graph construction of Morgenstern [Mor94] for a special choice of parameters. Although as we prove the transitions in this graph cannot be computed with constant locality, we show that if we turn this graph into a bipartite one, and permute the vertices on one side appropriately, the necessary computations can be carried out with constant locality. Another benefit of our choice of parameters is that the graph in Theorem 2 ends up having a simple description which does not rely on the structure theory of finite fields.

1.1 Applications

We now describe some applications of the above results.

Error-reduction for free. It is easy to reduce the error of an RP algorithm while increasing the number of random bits used: run the algorithm several times using independent

random bits, and take the AND of the results. For BPP algorithms one can take instead the MAJORITY, but we focus here on the RP setting for simplicity. Obtaining similar results without increasing the number of random bits has received attention since the 80's [RKS85, CW89]. One approach is to replace the independent choices for the random bits with correlated copies, obtained for example by computing the neighbors in an expander graph. Due to the complexity of previous expander constructions, this approach had a non-trivial cost which in particular could not be afforded in restricted computational models.

Using Theorem 1 we eliminate completely the cost of computing the correlated copies in several natural scenarios. We state a result for the computational model of circuits.

Theorem 3. *Given a circuit C and a parameter $p \leq 1/2$, we can construct in polynomial time another circuit D such that:*

1. D is the AND of $\mathrm{poly}(1/p)$ copies of C, where in each copy of C the input variables may be negated or permuted;

2. If C accepts all inputs then so does D;

3. If C accepts at most a 0.5 fraction of inputs then D accepts at most a p fraction.

We note that if C is an unbounded fan-in circuit whose output gate is AND then D has the same depth as C, whereas in previous error-reduction results the depth of D increased.

Proof. Let n be the number of input bits of C. Pick a graph from Theorem 1 with vertices $\{0,1\}^n$, second largest eigenvalue $\lambda \leq p$, and degree $d = \mathrm{poly}(1/p)$. We identify inputs to C with the vertices $\{0,1\}^n$. The circuit D on input $x \in \{0,1\}^n$ consists of the AND of d copies of the circuit C, where copy i gets the i neighbor of x. Items 1. and 2. are immediate. Item 3. follows from the expander mixing lemma. Specifically, let A be the set of inputs which C accepts and let X be the set of inputs which D accepts. Note that X is the set of vertices with all neighbors in A. Suppose X has density q. The probability that a uniformly chosen edge of the expander lands in $X \times A$ is at least q. By the expander mixing lemma, see e.g. [HLW06], $q \leq q/2 + \lambda\sqrt{q/2}$ and so $\sqrt{q/2} \leq \lambda$, implying $q \leq p$. □

Local loss-less expanders. Plugging our expanders in Theorem 7.1 in [CRVW02] we obtain local, bipartite loss-less expanders. A bipartite loss-less expander is a bipartite graph where any small set K of vertices on the zero side has nearly disjoint neighborhoods. Many applications of such graphs are described in [CRVW02]. For simplicity we only state our result for bipartite graphs with two equal sides. (The construction in [CRVW02] allows for the zero side to be smaller than the one side.)

Theorem 4. *For any $\epsilon > 0$ there exists $d = O(1)$ such that for every n there are d explicit local maps $C_1, C_2, \ldots, C_d$, each mapping n bits to n bits, such that the bipartite graph on vertices $\{0,1\}^n \times \{0,1\}$ where a node $(v,0)$ has neighbors $(C_i(x),1)$, for $i = 1,2,\ldots,d$, has the following property: any set K of up to $\Omega(2^n)$ vertices on the zero-side has $\geq (1-\epsilon)d|K|$ neighbors.*

Proof. (Sketch) We follow the proof of Theorem 7.1 in [CRVW02]. The graph constructed there is the zig-zag product [RVW02] of three conductors (an object defined in [CRVW02]). Two of these are of constant size. The other one is an expander graph with degree depending only on ϵ, and hence constant. For this graph we can use Theorem 1. Inspection of the zig-zag product reveals that it preserves constant locality. □

Efficient expanders in the RAM model. Computing neighbors in the graph in Theorem 2 is also efficient in the RAM model and in the C programming language. Specifically, we show how to compute the neighbors of a w-bit node with only a small, constant number of bit-wise AND, SHIFT, and XOR of w-bit words. To our knowledge, the only construction of expander graphs with a comparable efficiency is the one by Margulis [Mar73, GG81, JM87] (which is also not bipartite). A transition in the latter graphs involves only a constant number of w-bit additions. Our construction has the advantage of being Ramanujan of degree three. The constructions in [Mar73, GG81, JM87] would need larger degree to achieve the same eigenvalue bound, and seem to give nothing for degree three.

Hardness amplification of one-way permutations. Call a function $f : \{0,1\}^n \to \{0,1\}^m$ α-one-way for time t if every algorithm running in time t fails to invert f on at least an α fraction of inputs. One-way functions are essential building blocks for cryptography, see e.g. [Gol01]. A problem that has received attention since the 80's is how to "amplify" an α-one-way function f to a new function f' that is α'-one-way for $\alpha' > \alpha$. Yao's classic approach [Yao82] of computing f on k disjoint inputs does achieve this, but has the drawback of blowing up the input length of f by a factor k. This blow-up is unsatisfactory, and one can argue that it makes the new function impractical, see e.g. the discussion in [GIL$^+$90]. The question of whether it can be avoided remains open. However, for the special case of one-way *permutations,* a better approach is known. [GIL$^+$90] essentially show how to get the same improvement on α while only incurring an additive overhead in the input length. Their approach is based on expanders and so, with the expanders provided by this paper, we can afford it even in very restricted computational models. We state one representative result for one-way permutations with constant hardness and computable with constant locality. We note that [AIK06] gives strong evidence that such permutations exist even for hardness α close to 1. However their techniques blow up the input length by a large factor.

Theorem 5. *For any constants $\alpha, \alpha' \in (0,1)$ the following holds. Suppose that there exists an α-one-way permutation $f : \{0,1\}^n \to \{0,1\}^n$ computable with constant locality. Then there exists an α'-one-way permutation $f' : \{0,1\}^{n'} \to \{0,1\}^{n'}$ computable with constant locality and with $n' = n + O(1)$.*

Proof. (Sketch) We use Proposition 1 in [GIL$^+$90]. The input to f' consists of an input $x \in \{0,1\}^n$ for f and $k = O(1)$ edge labels in an expander on vertices $\{0,1\}^n$. The evaluation of f alternates evaluating f and moving to a neighbor in the expander. This is done for k times. Using the expander in Theorem 1, the new function still has constant locality.

We note that to compute f' we have to examine the edge labels, but this is just $O(1)$ bits. Another minor detail is that the expander may have a degree which is not a power of two. But if so we can encode an edge label in a bit string that is longer, but still of constant length, so that the resulting distribution on edge labels is sufficiently close to uniform. □

1.2 Related work and open questions

The study of small locality has received a lot of attention in theoretical computer science. For many tasks that at first sight seem to require large locality, researchers have been able to give implementations in constant locality, and our work makes another contribution in this direction. In the area of pseudorandomness, [Gol00] gives a candidate cryptographic generator computable with constant locality. [MST06] construct a small-bias generator with constant locality, refuting a conjecture in [CM01]. In a phenomenal work, [AIK06] show the existence of cryptographic pseudorandom generators computable with constant locality, assuming the existence of cryptographic generators computable in, say, logarithmic space (for which many candidates are available). Each of these works has been extended and applied in various settings.

Turning to classical reductions, [JMV15] recently show that 3SAT remains NP-complete even if we require that the clauses are computable by a local map of the index, a requirement stronger than what looked "hard (perhaps impossible)" [Wil14]. Our work should be relevant to extending [JMV15] to PCP reductions. The current best result in this direction is [BV14] which achieves locality one but reduces to kSAT for growing k (as opposed to constant k).

The above constructions perhaps explain the difficulty of proving lower bounds for sampling in constant locality. Starting with [Vio12], several papers study such lower bounds, but the whole area is largely uncharted. Closer to the setting of this paper, we can ask if there is a graph property that cannot be realized with constant locality. Rather than making "graph property" precise we mention two specific open questions.

One application of Ramanujan graphs is the construction of unique-neighbor expanders, which in turn have several applications, see [AC02]. However we do not know of local unique-neighbor expanders. The difficulty is that the approach in [AC02] requires Ramanujan graphs with degrees for which we do not know of a local construction.

It is also an open problem to prove a result like Theorem 2 for non-bipartite graphs. For context we note that there are several other cases in the literature where certain good bipartite graphs are constructed, but a corresponding non-bipartite construction is not known. These include the recent construction of bipartite Ramanujan graphs of any degree [MSS15] and the 15-year old construction of bipartite lossless expanders [CRVW02].

Related to expander graphs, another question that remains open is: Can we compute in NC^1 the endpoint of an n-step walk on a constant-degree expander graph with n-bit nodes?

Organization. We begin in Section 2 with some preliminaries on expanders and groups. Then in Section 3 we prove Theorem 1 and in Section 4 we prove Theorem 2.

2 Preliminaries

All the graphs in this paper are connected, undirected, and regular. We allow self loops and multiple edges. We can thus think of a graph as a symmetric non-negative integer adjacency matrix with a fixed row-sum (and, by symmetry, column-sum) called the degree. Alternatively we can think of a graph with degree d on vertices V as a map $f : V \times \{1, 2, \ldots, d\} \to V$ such that for any v and w in V we have $|\{i : f(v, i) = w\}| = |\{i : f(w, i) = v\}|$. We also write f_i for $f(., i)$.

Let G be a d-regular graph with adjacency matrix M', and let $M := M'/d$ be its normalized adjacency matrix. We recall basic facts from spectral graph theory which can be found e.g. in Problem 2.9 in [Vad12]. All eigenvalues are at most 1 in absolute value. The number 1 is an eigenvalue of M, and it has multiplicity one if and only if the graph is connected. The graph is bipartite if and only if -1 is an eigenvalue.

Definition 6. A family of connected graphs is called an expander if all the eigenvalues except 1 and -1 are in absolute value at most $\lambda < 1$ where λ is a universal constant. It is called a Ramanujan expander if $\lambda = 2d^{-1}\sqrt{d-1}$.

We note that this definition of expander graphs allows for the degree to be non-constant. We shall use this flexibility in Section 3.

2.1 Cayley and Schreier graphs

Let H be a group. Given a multiset S of elements from H we form the Cayley graph $Cay(H, S)$ whose vertices are H and where vertex $h \in H$ has neighbors sh for every element $s \in S$. We shall only consider symmetric multisets, that is multisets where the occurrences of s and s^{-1} are the same. These give symmetric graphs.

Further suppose that H is a group acting on a set V, namely there is a homomorphism from H to the group of permutations of V. Then we can form the Schreier graph $Sch(H, S, V)$ whose vertices are V and where $v \in V$ has neighbors sv for every $s \in S$, where $S \subseteq H$ and we wrote sv for the permutation corresponding to s applied to v.

The following lemma – Claim 7.2 in [RSW06] – shows that the expansion of $Sch(H, S, V)$ is at least as good as that of $Cay(H, S)$. For completeness we also include a proof (in a language that is slightly different from [RSW06]).

Lemma 7. *Let λ be an eigenvalue of $Sch(H, S, V)$. Then λ is also an eigenvalue of $Cay(H, S)$.*

Proof. Let $e : V \to \mathbb{C}$ be an eigenvector of $Sch(H, S, V)$ with eigenvalue λ. That is, for any $v \in V$ we have $\mathbb{E}_{s \in S} e(sv) = \lambda e(v)$. Pick any vertex $v_0 \in V$, and define $e' : H \to \mathbb{C}$ as $e'(h) = e(hv_0)$. We claim that e' is an eigenvector of $Cay(H, S)$ with eigenvalue λ. Indeed, $\mathbb{E}_{s \in S} e'(sh) = \mathbb{E}_{s \in S} e(shv_0) = \lambda e(hv_0) = \lambda e'(h)$. □

2.2 Bipartite graphs

Let G be a graph on vertex set V where vertex v has neighbors $f_i(v)$. The double-cover of G is the bipartite graph $V \times \{0,1\}$ where vertex (v,b) has neighbors $(f_i(v), 1-b)$.

Fact 8. *Let G' be the double cover of a graph G. If G' has eigenvalue λ then G has eigenvalue λ or $-\lambda$. In particular, the double cover of a Ramanujan graph is a bipartite Ramanujan graph.*

Proof. Let $e' : V \times \{0,1\} \to \mathbb{C}$ be an eigenvector of G' with eigenvalue λ. Assume that the vectors $e'(.,0)$ and $e'(.,1)$ are different. Then define $e(v) := e'(v,0) - e'(v,1)$ which is not zero. We have $\mathbb{E}_i e(f_i v) = \mathbb{E}_i(e'(f_i v, 0) - e'(f_i v, 1)) = \lambda e'(v,1) - \lambda e'(v,0) = -\lambda e(v)$.

Otherwise, if $e'(.,0)$ and $e'(.,1)$ are equal (and non-zero) define $e(v) := e'(v,0) + e'(v,1)$. We now have $\mathbb{E}_i e(f_i v) = +\lambda e(v)$. □

3 One-local expander

In this section we prove Theorem 1. First we note that the composition of two one-local maps is still one-local. So it suffices to prove the theorem for some $d = O(1)$ with an eigenvalue bound of $1 - \Omega(1)$. To obtain the general theorem one can take the t power of this graph, which has degree d^t and eigenvalue bound $(1-\Omega(1))^t = d^{-\Omega(1)}$.

Background on [ALW01]. By reinterpreting (a variant of) the zig-zag product [RVW02] in group-theoretic terms, [ALW01] give a way to prove that the semi-direct product C of two groups A and B is, with respect to certain generators, a Cayley expander graph. Specifically, assume that B acts on A, namely we can view homeomorphically the elements of B as automorphisms of A. Recall that the semi-direct product C of groups A and B has elements $A \times B$ and multiplication defined as follows:

$$(\hat{a}, \hat{b})(a, b) = (\hat{a}\hat{b}^{-1}(a), \hat{b}b),$$

where $b(a)$ is the image of a under the action b.

Let S and T be sets of generators for A and B, respectively. Further suppose that S is a (disjoint) union of c orbits under B, i.e., $S = \bigcup_{i=1}^{c} B(a_i)$, where $B(a)$ is the orbit of $a \in A$ under B. Then consider the following set U of generators for C:

$$U = \{(1_A, b)(a_i, 1)(1_A, b') : b, b' \in T, i \in [c]\}.$$

The key property is that the size of U is only $c|T|^2$, which can be a constant even if $|S|$ is not. (We note that even if the orbits have different sizes – as will happen to us – they are each picked with the same probability in the random walk induced by this zigzag operation.)

Theorem 9. *[[ALW01]] $Cay(C,U)$ is an expander graph if both $Cay(A,S)$ and $Cay(B,T)$ are.*

Our construction. For the group A we simply pick $GF(2)^n$ equipped with bit-wise xor (namely, addition). For B we take the permutation group S_n on n elements. We let B act on A by permuting coordinates.

Theorem 10. *[[Kas07]] There exists an explicit, constant-size set T of generators such that $Cay(S_n, T)$ is an expander.*

For generators for A we pick the union S of the orbits under B of the following three vectors: $0^n, 10^{n-1}, 1^k0^{n-k}$ where k is the ceiling of $n/2$.

Lemma 11. *$Cay(A, S)$ is an expander graph.*

Proof. It is a standard fact that it suffices to show that for every non-zero $v \in \{0,1\}^n$ the probability that $\langle v, x\rangle = 1$ over x picked uniformly from the multiset S is bounded away from 0 and from 1 (essentially following from the fact that the eigenvalues of the adjacency matrix are the Fourier coefficients of the distribution on generators); see for example the proof of Theorem 3.1 in [ALW01]. ($\langle v, x\rangle$ is the inner product modulo 2 of v and x.) To verify this, note that for any v, the probability that $\langle v, x\rangle = 0$ is $\Omega(1)$ thanks to the vector 0^n. So we just need to show that the probability that $\langle v, x\rangle = 1$ is $\Omega(1)$ as well. If the weight of v is larger than, say, $n/3$ this is true thanks to the vector 10^{n-1}. Now consider a vector v of weight less than $n/3$, and let x be a uniform permutation of 1^k0^{n-k}. Let us think instead of taking a random permutation of v and computing the inner product with the fixed vector $y = 1^k0^{n-k}$. After all but one of the non-zero entries of v have been permuted, we have covered no more than $n/3$ of the coordinates of y. So the last non-zero entry of v has a constant probability of being mapped to a one in y, and a constant probability of being mapped to a zero in y. □

By Theorem 9, the semi-direct product C of A and B with the generators

$$U = \{(1,b)(a,1)(1,b') : b, b' \in T, a \in \{0^n, 10^{n-1}, 1^k0^{n-k}\}\}$$

is an expander graph. Note that $|U| = O(1)$.

Finally, we view C as a group of permutations on $\{0,1\}^n$ as follows. Element (a,b) first permutes the coordinates by b and then xor's by a. To verify that this is a proper definition we need to check that the permutation of $(\hat{a},\hat{b})(a,b) = (\hat{a}\hat{b}^{-1}(a), \hat{b}b)$ is the same as the composition of the permutation of $(\hat{a},\hat{b})$ and the permutation of (a,b), which is true. This gives the Schreier graph $Sch(C, U, \{0,1\}^n)$. This graph is connected and by Lemma 7 is an expander. The transition functions only xor and permute bits, and so they can be implemented by one-local maps.

4 Local Ramanujan

In this section we prove Theorem 2. We make use of the following Ramanujan graph construction of Morgenstern.

Theorem 12. *[Theorem 5.13 in [Mor94]] Let $g(x) \in F_2[x]$ be an irreducible polynomial of even degree n, and represent F_{2^n} as $F_2[x]/g(x)$. Then the Cayley graph of $SL(2, F_{2^n})$ with the three generators zM_1, zM_2, zM_3 is a Ramanujan expander graph, where $L \in F_{2^n}$ satisfies $L^2 + L = 1$ and we define $z = 1/\sqrt{1+x}$, $M_1 = \begin{pmatrix} 1 & L \\ (L+1)x & 1 \end{pmatrix}$, $M_2 = \begin{pmatrix} 1 & 1 \\ x & 1 \end{pmatrix}$, and $M_3 = \begin{pmatrix} 1 & L+1 \\ Lx & 1 \end{pmatrix}$.*

An explicit choice for g and L is made below in Section 4.2.

Before continuing with our proof let us explain how Theorem 12 follows from Theorem 5.13 in [Mor94]. Using the notation in the latter, we pick $q = 2$ and $\epsilon = 1$, and note that $x^2 + x + 1$ is irreducible in $F_2[x]$. Note that the determinants of M_1, M_2, and M_3 are all $1 + x$ because $1 + (L^2 + L)x = 1 + x$. With the normalization $1/\sqrt{1+x}$, the determinants become 1. (The square root of $1+x$ exists because every element is a square in characteristic 2.) Morgenstern does not include this normalization, but we prefer to identify the group $\mathrm{SL}(2, F_{2^n})$ with the 2×2 matrices of determinant 1 over the field F_{2^n}. Finally note that $M_i^2 = \det(M_i)I$ and so each of our three generators is its own inverse.

The graph in Theorem 12 is problematic for us: In section 4.3 below we show that multiplication by z (or by zM_2) is not locally computable.

Our first step is to build the Schreier graph on vertex set $V := (F_{2^n})^2 - (0,0)$, which we view as column vectors, with respect to the generators in Theorem 12. (The permutation on V associated to $h \in SL(2, F_{2^n})$ is simply the matrix-vector multiplication.) We note that this graph is connected: every $(a,b)^T \in V$ equals $h(1,0)^T$ for some $h \in SL(2, F_{2^n})$. Indeed, if $a \neq 0$ we have $(a,b)^T = \begin{pmatrix} a & 0 \\ b & 1/a \end{pmatrix} (1,0)^T$, and similarly if $b \neq 0$ we have $(a,b)^T = \begin{pmatrix} a & 1/b \\ b & 0 \end{pmatrix} (1,0)^T$. By Lemma 7 this Schreier graph is also Ramanujan.

The next step is to take the double cover of this graph. We thus obtain a graph G on $2(2^{2n} - 1)$ vertices which is also Ramanujan by Fact 8. Later we show that we can pick any n of the form $n = 2 \cdot 3^t$, thus obtaining graphs on $2(2^{4 \cdot 3^t} - 1)$ nodes as in Theorem 2.

We still have not fixed the problem mentioned earlier, that multiplication by z (or by zM_2) is not locally computable. The last step is aimed to fix that, and is perhaps the least obvious. We argue that the normalization factor z can be removed from this last graph, and that doing so allows us to compute locally the neighbors of a vertex on the zero side.

4.1 Twisting the graph

Let G be a bipartite graph with vertices $V \times \{0,1\}$, where node (v,b) has neighbors $(f_i(v), 1-b)$. Let π be a permutation of V. We define the $\pi-$twist G' of G as follows. The vertices of G' are again $V \times \{0,1\}$. However vertex $(v,0) \in G'$ has neighbors $(\pi f_i v, 1)$ (and so vertex $(v,1) \in G'$ has neighbors $(f_i \pi^{-1} v, 0)$). We claim that twisting a graph does not affect its spectral expansion.

Lemma 13. *The eigenvalues of G and G' are the same.*

Proof. We show that if λ is an eigenvalue of G' then λ is also an eigenvalue of G. Let $e' : V \times \{0,1\} \to \mathbb{C}$ be an eigenvector of the twisted graph G' with eigenvalue λ. This means that

$$e'(v,0) = \lambda^{-1}\mathbb{E}_i e'(\pi f_i v, 1)$$

and

$$e'(v,1) = \lambda^{-1}\mathbb{E}_i e'(f_i \pi^{-1} v, 0).$$

Define $e(v,0) := e'(v,0)$ and $e(v,1) := e'(\pi v, 1)$. Note that e is non-zero if and only if e' is non-zero. We claim that e is an eigenvector of G with eigenvalue λ. Indeed,

$$e(v,0) = e'(v,0) = \lambda^{-1}\mathbb{E}_i e'(\pi f_i v, 1) = \lambda^{-1}\mathbb{E}_i e(f_i v, 1).$$

Similarly,

$$e(v,1) = e'(\pi v, 1) = \lambda^{-1}\mathbb{E}_i e'(f_i \pi^{-1}\pi v, 0) = \lambda^{-1}\mathbb{E}_i e(f_i v, 0).$$

□

We twist the graph by multiplying a node by $\sqrt{1+x}$. This means that the neighbors of a zero-side vertex $(v,0)$ are simply $(M_i v, 1)$ where the M_i are as in Theorem 12.

4.2 Local computation

We now argue that multiplication by M_i can be done with constant locality. Inspection of the M_i reveals that the only non-trivial steps are multiplication of an arbitrary element of F_{2^n} by x and L, where L is the field element in Theorem 12. Multiplication by x is again simple and works for any irreducible polynomial we choose to define the field. On the other hand, multiplication by L relies on the specific irreducible polynomial $g(x) := x^n + x^{n/2} + 1$ when $n = 2 \cdot 3^t$.

Lemma 14. *[Theorem 1.1.28 in [vL99]] The polynomial $g(x)$ is irreducible.*

Earlier, [HV06] shows that the order of x modulo $g(x)$ is small, and exploits this to compute efficiently the exponentiation of an n-bit field element to an n-bit exponent, for example in space $O(\log n)$.

In this work the critical observations are that L is sparse – in fact, $L = x^{n/2}$ – and that modulo $g(x)$ multiplication by any fixed sparse element can be carried out with constant locality.

Claim 15. The field element $L := x^{n/2}$ satisfies $L^2 + L = 1$.

Proof. We have $L^2 = x^n = x^{n/2} + 1 = L + 1$. □

Claim 16. Let $n = 2 \cdot 3^t$ and represent F_{2^n} as $F_2[x]/g(x)$ where $g(x)$ is the irreducible polynomial $x^{2\cdot 3^t} + x^{3^t} + 1$. For any sparse (i.e., with $O(1)$ monomials) element $a \in F_{2^d}$ there is an explicit local map $C : \{0,1\}^n \to \{0,1\}^n$ such that $C(b) = ab$ for every $b \in F_{2^n}$.

Proof. It is enough to consider the case where a consists of a single monomial x^s. Hence, given as input $\sum_{j<n} c_j x^j$ we have to output the coefficients of the polynomial $\sum_{j<n} c_j x^{j+s}$. For simplicity of notation we only consider the case $s = 3^t$, i.e., multiplication by L, which is all that is needed for the application.

Write an element $y \in F_{2^n}$ as a pair (y_2, y_1) where $|y_2| = |y_1| = n/2$ and y_1 consists of the least significant $n/2$ bits. If $y = \sum_{j<n} c_j x^j$ then we have

$$\begin{aligned} ay = \sum_{j<n} c_j x^{j+n/2} &= \sum_{0\le j<n/2} c_j x^{j+n/2} + \sum_{0\le j<n/2} c_{j+n/2} x^{j+n} \\ &= x^{n/2} \sum_{0\le j<n/2} c_j x^j + (1 + x^{n/2}) \sum_{0\le j<n/2} c_{j+n/2} x^j \\ &= x^{n/2} \sum_{0\le j<n/2} (c_{j+n/2} + c_j) x^j + \sum_{0\le j<n/2} c_{j+n/2} x^j. \\ &= (y_2 + y_1, y_2). \end{aligned}$$

□

Finally, we show that the expander in Theorem 2 is efficiently computable in the RAM model. Bit-wise XOR is clearly efficient. Multiplication by x is simply a cyclic shift plus possibly a bit-wise XOR depending on the most significant bit of x. It only remains to verify that multiplication by $L = x^{n/2}$ is efficient too. Indeed, as already seen, this multiplication has the following simple format. Write an element $y \in F_{2^n}$ as a pair (y_2, y_1) where $|y_2| = |y_1| = n/2$ and y_1 consists of the least significant $n/2$ bits. Then $L \cdot (y_2, y_1) = (y_2 + y_1, y_2)$.

4.3 Negative results for local computation

In this section we make two remarks that aim to give some context for the results in Section 4.2. First, we note that the sparsity of $g(x)$ alone is not sufficient for Claim 16. Specifically we show that the Parity function on n bits can be reduced to multiplication modulo the polynomial $h(x) := x^n + x^{n-1}$. Since Parity requires locality n, the result follows. To see the reduction, first note that for any $j \ge n$, $x^j = x^{n-1}$ modulo h. So if we multiply an n-bit element $\sum_{j=0}^{n-1} c_j x^j$ by x^{n-1} we obtain $x^{n-1} \sum_{j=0}^{n-1} c_j$. Thus, the parity of the input bits is in the most significant bit of the output. In our result we use the stronger property that in the binary representation of g the ones are spaced away by $\Omega(n)$ zeros.

Second, we show that the transitions in Morgenstern's expander in Theorem 12 are not locally computable, for our choice of the underlying field. This justifies twisting the graph. Note that multiplication of an arbitrary vector by zM_2 requires multiplication of an arbitrary field element by the normalization factor $z = 1/\sqrt{1+x}$. We show that parity on $\Omega(n)$ bits reduces to the latter. This also has consequences for the RAM model, because there is no known way to compute parity very efficiently there.

Claim 17. $z = 1 + x + x^2 + \ldots + x^{b-1}$ where $b = (3n/2 + 1)/2$.

Proof. First we note that $\sqrt{1+x} = 1 + x^b$. Indeed, $(1 + x^b)^2 = 1 + x^{3n/2+1} = 1 + x$, because $x^{3n/2} = 1$, a fact also pointed out and used in [HV06]. It remains to prove that $1/(1 + x^b) = 1 + x + x^2 + \ldots + x^{b-1}$, which is equivalent to $1 = 1 + x + \ldots + x^{2b-1}$. Note that $2b - 1 = 3n/2$, and so we want to show that $\sum_{i=0}^{3n/2} x^i = 1$. Indeed,

$$\sum_{i=0}^{3n/2} x^i = \sum_{i=0}^{n-1} x^i + \sum_{i=0}^{n/2-1} (x^i + x^{i+n/2}) + x^{3n/2} = x^{3n/2} = 1.$$

□

Claim 18. Parity on $\Omega(n)$ bits reduces to multiplying by z.

Proof. Note that $b < 3n/4 + 1$ in Claim 17. So, if you multiply $z = 1/\sqrt{1+x} = 1 + x + x^2 + \ldots + x^{b-1}$ by an input y that is zero in all but the least significant $0.2n$ bits, there will be no wrapping around, and what you are doing is plain convolution. Thus, the parity of y will be one of the bits in zy. □

Acknowledgments. We thank the anonymous referees for their useful comments.

References

[AC02] Noga Alon and Michael R. Capalbo. Explicit unique-neighbor expanders. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, page 73, 2002.

[AIK06] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. *SIAM J. on Computing*, 36(4):845–888, 2006.

[ALW01] Noga Alon, Alexander Lubotzky, and Avi Wigderson. Semi-direct product in groups and zig-zag product in graphs: Connections and applications. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 630–637, 2001.

[ASW09] Sanjeev Arora, David Steurer, and Avi Wigderson. Towards a study of low-complexity graphs. In *Coll. on Automata, Languages and Programming (ICALP)*, pages 119–131, 2009.

[BV14] Eli Ben-Sasson and Emanuele Viola. Short PCPs with projection queries. In *Coll. on Automata, Languages and Programming (ICALP)*, 2014.

[BYGW99] Ziv Bar-Yossef, Oded Goldreich, and Avi Wigderson. Deterministic amplification of space bounded probabilistic algorithms. In *IEEE Conf. on Computational Complexity (CCC)*, pages 188–198, 1999.

[CM01] Mary Cryan and Peter Bro Miltersen. On pseudorandom generators in NC^0. In *26th Symposium on Mathematical Foundations of Computer Science (MFCS 01)*, pages 272–284. Springer-Verlag, 2001.

[CRVW02] Michael R. Capalbo, Omer Reingold, Salil P. Vadhan, and Avi Wigderson. Randomness conductors and constant-degree lossless expanders. In *ACM Symp. on the Theory of Computing (STOC)*, pages 659–668, 2002.

[CW89] Aviad Cohen and Avi Wigderson. Dispersers, deterministic amplification, and weak random sources. In *30th Symposium on Foundations of Computer Science*, pages 14–19, Research Triangle Park, North Carolina, 30 October–1 November 1989. IEEE.

[DvM06] Scott Diehl and Dieter van Melkebeek. Time-space lower bounds for the polynomial-time hierarchy on randomized machines. *SIAM J. on Computing*, 36(3):563–594, 2006.

[GG81] Ofer Gabber and Zvi Galil. Explicit constructions of linear size superconcentrators. *J. of Computer and System Sciences*, 22:407–420, 1981.

[GIL+90] Oded Goldreich, Russell Impagliazzo, Leonid A. Levin, Ramarathnam Venkatesan, and David Zuckerman. Security preserving amplification of hardness. In *31st IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 318–326, 1990.

[Gol00] Oded Goldreich. Candidate one-way functions based on expander graphs. Technical report, Electronic Colloquium on Computational Complexity, 2000.

[Gol01] Oded Goldreich. *Foundations of Cryptography: Volume 1, Basic Tools*. Cambridge University Press, 2001.

[GV04] Dan Gutfreund and Emanuele Viola. Fooling parity tests with parity gates. In *8thWorkshop on Randomization and Computation (RANDOM)*, pages 381–392. Springer, 2004.

[HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bull. Amer. Math. Soc. (N.S.)*, 43(4):439–561 (electronic), 2006.

[HV06] Alexander Healy and Emanuele Viola. Constant-depth circuits for arithmetic in finite fields of characteristic two. In *23rd Symp. on Theoretical Aspects of Computer Science (STACS)*, pages 672–683. Springer, 2006.

[JM87] S. Jimbo and A. Maruoka. Expanders obtained from affine transformations. *Combinatorica. An Journal of the János Bolyai Mathematical Society*, 7(4):343–355, 1987.

[JMV15] Hamid Jahanjou, Eric Miles, and Emanuele Viola. Local reductions. In *Coll. on Automata, Languages and Programming (ICALP)*, 2015. Available at http://www.ccs.neu.edu/home/viola/.

[Kas07] Martin Kassabov. Symmetric groups and expander graphs. *Invent. Math.*, 170(2):327–354, 2007.

[LPS88] Alexander Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica. A Journal of the János Bolyai Mathematical Society*, 8(3):261–277, 1988.

[Mar73] G. A. Margulis. Explicit construction of concentrators. *Problems Inform. Transmission*, 9:325–332, 1973.

[Mor94] M. Morgenstern. Existence and explicit constructions of q + 1 regular Ramanujan graphs for every prime power q. *Journal of Combinatorial Theory, Series B*, 62(1):44 – 62, 1994.

[MSS15] Adam W. Marcus, Daniel A. Spielman, and Nikhil Srivastava. Interlacing fami-

lies IV: bipartite Ramanujan graphs of all sizes. In *IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 1358–1377, 2015.

[MST06] Elchanan Mossel, Amir Shpilka, and Luca Trevisan. On epsilon-biased generators in NC^0. *Random Struct. Algorithms*, 29(1):56–81, 2006.

[RKS85] Nicholas Pippenger Richard Karp and Michael Sipser. A time-randomness tradeoff. In *AMS Conference on Probabilistic Computational Complexity*, 1985.

[RSW06] Eyal Rozenman, Aner Shalev, and Avi Wigderson. Iterative construction of cayley expander graphs. *Theory of Computing*, 2(5):91–120, 2006.

[RVW02] Omer Reingold, Salil Vadhan, and Avi Wigderson. Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. of Math. (2)*, 155(1):157–187, 2002.

[Vad12] Salil P. Vadhan. Pseudorandomness. *Foundations and Trends in Theoretical Computer Science*, 7(1-3):1–336, 2012.

[Vio12] Emanuele Viola. The complexity of distributions. *SIAM J. on Computing*, 41(1):191–218, 2012.

[vL99] J. H. van Lint. *Introduction to coding theory*. Springer-Verlag, Berlin, third edition, 1999.

[Wil14] Ryan Williams. Nonuniform ACC circuit lower bounds. *J. of the ACM*, 61(1):2:1–2:32, 2014.

[Yao82] Andrew Yao. Theory and applications of trapdoor functions. In *23rd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 80–91. IEEE, 1982.