THEORY OF COMPUTING, Volume 2 (2006), pp. 91–120
http://theoryofcomputing.org

Iterative Construction of Cayley Expander Graphs

Eyal Rozenman* Aner Shalev† Avi Wigderson‡

Received: August 4, 2005; published: April 25, 2006.

Abstract: We construct a sequence of groups G_n, and explicit sets of generators $Y_n \subset G_n$, such that all generating sets have bounded size, and the associated Cayley graphs are all expanders. The group G_1 is the alternating group A_d, the set of even permutations on the elements $\{1,2,\ldots,d\}$. The group G_n is the group of all even symmetries of the rooted d-regular tree of depth n. Our results hold for any large enough d.

We also describe a finitely generated infinite group G_∞ with generating set Y_∞, given with a mapping f_n from G_∞ to G_n for every n, which sends Y_∞ to Y_n. In particular, under the assumption described above, G_∞ has property (τ) with respect to the family of subgroups $\ker(f_n)$.

The proof is elementary, using only simple combinatorics and linear algebra. The recursive structure of the groups G_n (iterated wreath products of the alternating group A_d) allows for an inductive proof of expansion, using the group theoretic analogue (of Alon et

*The Hebrew university, Jerusalem. E-mail: eyalroz@cs.huji.ac.il. Part of this research was performed while visiting the Institute for Advanced Study, Princeton, NJ.

†The Hebrew university, Jerusalem. E-mail: shalev@math.huji.ac.il. Partially supported by BSF grant 2000-53 and a grant from the Israel Science Foundation.

‡Institute for Advanced Study, Princeton. E-mail: avi@ias.edu. Partially supported by NSF grant CCR-0324906.

ACM Classification: G.2.2, G.3

AMS Classification: 05C25, 37A30

Key words and phrases: Cayley graphs, Expanders, Expander graphs, Wreath products

Authors retain copyright to their work and grant Theory of Computing unlimited rights to publish the work electronically and in hard copy. Use of the work is permitted as long as the author(s) and the journal are properly acknowledged. For the detailed copyright statement, see http://theoryofcomputing.org/copyright.html.

© 2006 Eyal Rozenman, Aner Shalev, and Avi Wigderson
DOI: 10.4086/toc.2006.v002a005

al., 2001) of the zig-zag graph product (Reingold et al., 2002). The basis of the inductive proof is a recent result by Kassabov (2005) on expanding generating sets for the group A_d.

Essential use is made of the fact that our groups have the *commutator property:* every element is a commutator. We prove that direct products of such groups are expanding even with highly correlated tuples of generators. Equivalently, highly dependent random walks on several copies of these groups converge to stationarity on all of them essentially as quickly as independent random walks. Moreover, our explicit construction of the generating sets Y_n above uses an efficient algorithm for solving certain equations over these groups, which relies on the work of Nikolov (2003) on the commutator width of perfect groups.

1 Introduction

1.1 Expander graphs

Expanders are graphs which are sparse but nevertheless highly connected. Expanders graphs have been used to solve many fundamental problems in computer science, in topics including network design (e.g. [40, 41, 1]), complexity theory ([49, 44, 48]), derandomization ([36, 18, 19]), coding theory ([45, 46]), and cryptography ([15]). Expander graphs have also found some applications in various areas of pure mathematics, such as topology, measure theory, game theory and group theory (e.g. [21, 30, 16, 31]).

Standard probabilistic arguments ([39]) show that almost every constant-degree (≥ 3) graph is an expander. However, most applications demand explicit constructions. Here we take the most stringent definition of explicitness of an infinite family of graphs, requiring that a deterministic polynomial time algorithm can compute the neighbors of any given vertex, from the vertex name and the index of the graph in the family. This challenge of explicit construction led to an exciting and extensive body of research.

Most of this work was guided by the algebraic characterization of expanders, developed in [47, 5, 2]. They showed the intimate relation of (appropriate quantitative versions of) the combinatorial (isoperimetric) notion of expansion above, to the spectral gap in the adjacency matrix (or, almost equivalently, the Laplacian) of the graph. This relationship is tight enough for almost all applications (but there are some exceptions, e.g. see [50, 10]).

Using this connection, an infinite family of regular graphs is defined to be an expander family if for all of them the second largest eigenvalue of the normalized adjacency (i.e. random walk) matrix is bounded above by the same constant that is smaller than 1.

This algebraic definition of expanders by eigenvalues naturally led researchers to consider algebraic constructions where this eigenvalue can be estimated. The celebrated sequence of papers [32, 14, 5, 3, 20, 29, 33, 35] provided such highly explicit families of constant-degree expanders. All of these constructions are based on groups, and their analysis often appeals to deep results in mathematics.

The algebraic mould was broken recently by [42], where a simple, combinatorial construction of constant-degree expander graphs was presented. The construction is iterative, generating the next graph in the family from two previous ones via a novel graph product, the *zig-zag* product. This product was proved (using simple linear algebra) to simultaneously keep the degree small, and retain expansion.

Thus the iteration process need only be provided with an initial, fixed size expander "seed" graph, from which all others are generated. The required parameters of the seed graph are easily shown to hold for a random graph (which suffices for explicitness - it is of constant size), but it is also easy to construct one explicitly.

Our main result in this paper is a similar iterative construction of expanding *Cayley* graphs (which we turn to define next) from one initial "seed" Cayley graph. In our case, the seed Cayley graph is based on the group A_d, the group of even permutations on the set $\{1,2,\dots,d\}$. In a recent breakthrough, Kassabov [22] explicitly constructed a bounded-size, expanding generating set for A_d, which yields the seed expander Cayley graph we need.

Our construction may be seen as another step in exploring this fundamental notion of expansion, and its relations to yet unexplored mathematical structures. It also further explores the power of the zig-zag product in constructing even stronger expanders. It was already shown [10] that it can yield expansion beyond the eigenvalue bound, and is shown here to yield Cayley expanders.

1.2 Expanding Cayley graphs

For a finite group H and a (symmetric) set of elements T in it, the Cayley graph $C(H;T)$ has the elements of H as vertices, and edges connect a pair of vertices g,h if their "ratio" gh^{-1} is in T. We remark that while most applications do not require the expanders to be "Cayley", the recent paper [7] seems to essentially require Cayley expanders to achieve nearly linear-sized locally testable codes (LTCs) and probabilistically checkable proof (PCPs).

Many of the algebraic expander constructions mentioned above are Cayley graphs. In all of these, the groups in question are linear matrix groups over finite fields, and their expansion follows from celebrated results in mathematics, including Kazhdan's work on Property T [25], Selberg's 3/16 theorem [43], and the resolution of the Ramanujan conjecture of Eichler, Deligne and Igusa (starting in [12]). It should be noted that for some of the other algebraic constructions elementary proof of expansion exist, using only a discrete Fourier transform [20].

For other natural families of groups the question was considered both by mathematicians and computer scientists. For example, for Abelian groups it is easy to see that any set of expanding generators has to be at least logarithmic in the size of the group. Thus they cannot provide expanding Cayley graphs of constant degree (a more general result appears in [26]). Lubotzky and Weiss generalized this negative result for all solvable groups of bounded derived length [28].

Understanding which natural families of groups can be made expanding (with a fixed size generating set) is a basic question, and little progress was made over the foundational results above in the last 15 years. However, in the last year several breakthroughs were made by Kassabov and Nikolov [22, 23, 24]. These results suggest that all the simple groups may have fixed size expanding generating sets. Of particular interest to our work is the family of symmetric groups (of all permutations). Much work has been devoted to analyzing the expansion of this group under a variety of generating sets in the context of card shuffling (e.g. see [11, 27]). However, in all these papers the generating sets are huge, and did not provide a clue to the status of this problem. In a recent breakthrough, Kassabov [22] showed that the symmetric groups indeed have explicit, fixed-size expanding generators, independent on the group size.

The possibility that the zig-zag product and iterative construction may be used for Cayley expanders was first revealed in [4]. They discovered that the well-known *semi-direct* product on groups may be viewed (roughly speaking) as a special case of the zig-zag product of graphs. More precisely, the zig-zag product of two Cayley graphs, with certain important restrictions on the structure of their generating sets, is a Cayley graph of the semi-direct product of the associated groups. Thus one can generate larger Cayley expanders of small degree from smaller ones. This observation was used to show that expansion is *not* a group property – in some groups certain constant size sets will expand, while others will not.

This Cayley graph version of the zig-zag theorem raises the hope that, given a "seed" expander Cayley graph, one can obtain a sequence of expander Cayley graphs via an iterative process using the zig-zag theorem. However, unlike the case of unstructured graphs, the restrictions on generators alluded to above for applying the zig-zag product on Cayley graphs, make iterations a highly nontrivial (and illuminating) task. In [34] such a construction was given, which falls short of the task at hand on two counts. First, the generating sets (and hence the degrees) of the groups in the family are not of constant size, but rather grow slowly (roughly like $\log^*$ of the group size). Second, these generating sets are shown to exist via a probabilistic argument, hence the resulting family is not explicit. Still, this construction makes no assumptions, as the seed Cayley expander for the iteration is easily seen to exist.

In this paper we fix both problems. We give a sequence of groups G_n, and explicit generating sets Y_n for each G_n, such that the Cayley graphs $C(G_n,Y_n)$ are expanding. Moreover, Y_n as bounded size, independent of n. Actually, we will later on see that the generators Y_n are consistent with each other: In the natural projection of G_{n+1} to G_n the set Y_{n+1} projects to the set Y_n.

The technique developed yields some results which do not require a seed Cayley graph at all. We show how to obtain an explicit sequence of expanding *Schreier graphs*. (The novelty is in the explicitness, since by [17] every regular graph with even degree is a Schreier graph). We then use the Schreier graph sequence to construct a sequence of expanders $\mathcal{X}_n$ in which each graph $\mathcal{X}_{n+1}$ is a lift of $\mathcal{X}_n$, by noticing that in our Schreier graph sequence each graph is actually a lift of its predecessor (lifts are defined in Section 9).

1.3 Our construction

Our groups are completely different from most groups previously used in this area. Indeed, they are very natural combinatorial objects. Let $T(d,n)$ denote the d-regular tree of depth n. The group of symmetries of this tree allows permuting the children of every internal node arbitrarily. Thus every element of this group may be described by a mapping of the internal nodes to the symmetric group S_d, describing how to permute the children of every such node. Group product of two such elements is simply performing the first set of permutations at every node, and then the next set. Our groups G_n are subgroups of all symmetries, allowing only *even* permutations at every internal node of $T(d,n)$. This natural restriction avoids a huge Abelian quotient that would have rendered expansion impossible, as the group would not even be generated by a constant number of generators. Our method of proof (sketched below) is elementary, using only linear algebra. All other known proofs use representation theory of the groups involved, and in most cases much deeper results as well.

There is a very natural inductive definition of the groups G_n. G_1 is the alternating group A_d of all even permutations on d elements (and is essentially the "seed group" of our construction). G_{n+1} can be obtained from d copies of G_n, and one copy of A_d acting on them simply by permuting the copies. Formally, this is called a wreath product, denoted $G_{n+1} = G_n \wr A_d$, and is a special case of a semidirect product, giving equivalently $G_{n+1} = (G_n)^d \rtimes A_d$. Our assumption gives a small expanding set of generators for A_d, and by induction we have such a set for G_n.

How does induction proceed? Naturally, we would like to use the zig-zag theorem for the semi-direct product [4, 42]. The technical requirement alluded to above is simply that we find an expanding generating set for $(G_n)^d$, which need not be small, but must be an orbit under the action of A_d, given a (small) expanding generating set Y_n for G_n. A natural candidate for such an orbit is all (even) permutations of the *balanced d-vector*, one which has every one of the elements of Y_n occurring the same number of times (if $|Y_n|$ divides d). It is the largest possible orbit, and the projection of a random element of the orbit to any small subset of the coordinates is (almost) a random independent element of Y_n in each coordinate.

We now turn to study the second eigenvalue of the Cayley graph of $(G_n)^d$ under these generators. The associated linear operator acts on the space of real functions on $(G_n)^d$. Luckily, this space of functions is simple to describe - it is the d-fold tensor product of the same space for G_n. What is not so lucky is the dependence between the coordinates of a balanced vector. Indeed, had G_n been Abelian, this orbit would not even be generating (i.e. the graph would not be connected). Here our special group structure is important. A key fact (proved by Nikolov [37]) is that every element in G_n is a commutator. Construct a new generating set $\tilde{Y}_n$ by adding to Y_n, for each of its elements, the constituents of its representation as a commutator. We use Nikolov's proof to actually give a polynomial time algorithm for finding this representation. Now take the orbit of all balanced vectors over $\tilde{Y}_n$ to be the generating set for $(G_n)^d$.

How can this revision take care of the dependencies? A simpler setting, to which we reduce our analysis, is the following Cayley graph. The group is simply $(G_n)^2$, namely only two copies of G_n. The generators are all pairs (g, g^{-1}) for all $g \in \tilde{Y}_n$. Thus, there is *complete* correlation between the two coordinates. The key point is that, using the special structure of $\tilde{Y}_n$, with positive probability a short word in one of the two components will vanish, while in the second it will give an original generator of Y_n, thereby decoupling the dependence of the two components. So, quite surprisingly, this Cayley graph on two copies is expanding despite the complete correlation (it is a nontrivial exercise to even establish connectivity of this graph – note that it would *not* be connected had G_n been Abelian, or if we took instead the pairs (g, g) for any group G_n). This construction (which we feel is of independent interest) is quite special and mysterious, and naturally the description above hides many essential details. Still, it is the heart of the matter.

For $m \geq n$ there is a natural restriction map $G_m \to G_n$ - given a symmetry of the tree with depth m consider its action on the subtree with depth n with the same root. As we shall see, the generating set Y_m is mapped to Y_n under this restriction. This gives rise to an infinite "limit group" G_∞ given with a generating set Y_∞ and restriction maps $f_n : G_\infty \to G_n$, where $f_n(Y_\infty) = Y_n$. In particular, under the assumption on A_d, the group G_∞ has property (τ) with respect to the family of subgroups $\ker(f_n)$ (Lubotzky's property (τ), a "baby" version of Kazhdan's property (T), is defined in Section 8).

1.4 Organization of this paper

In Section 2 we define expander graphs and Cayley graphs, and present some useful results. In Section 3 we define the sequence of groups we use. In Section 4 we describe the expanding generating sets, and prove the main Theorem 4.1 - that they are indeed expanding - by induction. The proof is based on a main lemma (Theorem 4.6). The lemma gives an expanding generating set for the group G^d given an expanding generating set for G (under certain conditions on G). Finally, in Section 6 we present an algorithmic version of Nikolov's theorem, that every element in our family of groups has a commutator representation that can be found efficiently.

We then turn to some corollaries of the main theorem. In Section 7 we explicitly construct a sequence of expanding Schreier graphs, free from the seed graph on the alternating group. In Section 8 we give generators for a subgroup of the symmetry group of the *infinite* rooted d-regular tree. These generators, when restricted to the *finite* rooted d-regular tree with depth n, generate an expander Cayley graph on the (alternating) group of symmetries of this finite tree. As a corollary we deduce that this infinite group has Lubotzky's Property (τ) with respect to a natural infinite family of normal subgroups. Then in Section 9 we combine the previous two results to obtain a sequence of expanding graphs each of which is a lift of the previous one.

2 Preliminaries

2.1 Graphs, eigenvalues and adjacency matrices

All graphs discussed in this paper are undirected, regular graphs. We allow multiple edges and self loops, so graphs are best understood as symmetric nonnegative integer matrices with a fixed row-sum, called the *degree*. For a graph $\mathcal{X}$, we let $V(\mathcal{X})$ denote its set of vertices and $E(\mathcal{X})$ its (multiset of) edges.

Let $\mathcal{X}$ be a k-regular graph, and $M = M_{\mathcal{X}}$ its normalized adjacency matrix (divide the adjacency matrix by the degree k to make it stochastic). We denote by $\lambda(\mathcal{X})$ the second largest (in absolute value) eigenvalue of M. The *spectral gap* of the graph is $1-\lambda(\mathcal{X})$.

Let W be the vector space of real functions on the set $V(\mathcal{X})$, with its standard L_2 inner product. $M_{\mathcal{X}}$ defines a linear operator on W: For $f \in W$, the value of the function $M_{\mathcal{X}}(f) \in W$ on a vertex x is the average value of f on all the neighbors of x (counted with multiplicities).

Let $W_{||}$ be the one-dimensional subspace consisting of the constant functions, and let $W_{\perp}$ be the orthogonal complement. Since the constant functions are eigenvectors of M corresponding to the (largest) eigenvalue 1, then

$$\lambda(\mathcal{X}) = \max_{w \in W_{\perp}} \|Mw\|/\|w\|$$

where $\|w\|$ is the L_2 norm of w.

Definition 2.1. An infinite family of graphs $\mathcal{X}_n$ is called an **expander family** if $\lambda(\mathcal{X}_n) \leq \mu$ for some $\mu < 1$ independent of n. The family is said to be (strongly) **explicitly described**, if there is a polynomial time algorithm which, on input n and the name of a vertex v in G_n (in binary), outputs the neighbors of v in G_n.

We will use the following two simple results, which describe how taking the tensor power of a graph, and taking the power of a graph, affect the 2nd eigenvalue λ:

Claim 2.2. *Let $\mathcal{X} = (V,E)$ be a graph, and let $M_{\mathcal{X}}$ be the normalized adjacency matrix. Let $M_{\mathcal{Y}} = (M_{\mathcal{X}})^{\otimes d}$, and define $\mathcal{Y}$ to be the graph (on the vertex set V^d) with normalized adjacency matrix $M_{\mathcal{Y}}$. Then $\lambda(\mathcal{Y}) = \lambda(\mathcal{X})$.*

Observation 2.3. Let $\mathcal{X} = (V,E)$ be a graph, $M_{\mathcal{X}}$ the normalized adjacency matrix and $M_{\mathcal{Y}} = (M_{\mathcal{X}})^k$. Let $\mathcal{Y}$ be the graph (on vertex set V) with normalized adjacency matrix $M_{\mathcal{Y}}$. Then $\lambda(\mathcal{Y}) = \lambda(\mathcal{X})^k$.

We will use the following convexity result later: If the spectral gap $(1-\lambda(\mathcal{Y}))$ of a graph $\mathcal{Y}$ is not too small, and $\mathcal{Y}$ is a large subgraph of $\mathcal{X}$ (on the same vertex set) then the spectral gap of $\mathcal{X}$ is also not too small.

Claim 2.4. *Let $\mathcal{Y}=(V,E_1)\subset\mathcal{X}=(V,E_2)$ (i.e. $E_1\subset E_2$) be s- and t-regular graphs respectively on the same vertex set V. Then*

$$1-\lambda(\mathcal{X})\geq\frac{s}{t}(1-\lambda(\mathcal{Y}))\ .$$

We will later need the following result on vectors

Claim 2.5. *If for some vectors $w_0,w_1,\ldots,w_L$, all with norm 1,*

$$(1/L)\cdot\|\sum_{i=1}^{L}w_i\|\leq 1-\varepsilon$$

then

$$(1/L)\cdot\sum_{i=1}^{L}\|w_0+w_i\|/2\leq 1-\varepsilon/4\ .$$

2.2 Groups and the wreath product

2.2.1 Cayley graphs

Let G be a finite group. We will represent groups multiplicatively, and 1 will denote the identity element of the group. Let Y be a multi-subset of G. We will always use *symmetric* sets Y, namely the number of occurrences of x and x^{-1} in Y is the same for every $x\in G$. $|Y|$ will denote the size of the multiset (counting multiplicities).

The *Cayley graph* $\mathrm{C}(G,Y)$ has vertex set G, and for every vertex $g\in G$ and $x\in Y$ there is an edge (g,gx). The graph $\mathrm{C}(G,Y)$ is undirected (as Y is symmetric) and is $|Y|$-regular. For $x\in G$ let P_x be the permutation matrix corresponding to $g\to gx$ in G. The normalized adjacency matrix of $\mathrm{C}(G,Y)$ is $\sum_{x\in Y}P_x/|Y|$. We will also use the notation $\mathbb{E}_{x\in Y}[P_x]$ to denote this average of operators.

Let $W=W(G)$ be the vector space of functions $G\to\mathbb{R}$ as in the previous section. We will be interested in the expansion properties of Cayley graphs on the group G^d, the Cartesian product of d copies of G. Note that $W(G^d)=W^{\otimes d}$.

Observation 2.6. Let $W_{||}$ be the space of constant functions on the vertices of G, and let $W_{\perp}$ be its orthogonal complement. Let $\bar{b}=(b_1,\ldots,b_d)$ be a length-d vector in the alphabet $\{||,\perp\}$, and let $W_{\bar{b}}$ be the vector space $\otimes_{i=1}^{d}W_{b_i}$. Consider the space $W^{\otimes d}$, the d-th tensor power of W. The space $W^{\otimes d}$ inherits an inner product structure from W, where the inner product of two pure tensors is the product of the inner products of the components of the tensors. The orthogonal decomposition $W=W_{||}+W_{\perp}$ induces an orthogonal decomposition

$$W^{\otimes d}=\sum_{\bar{b}\in\{||,\perp\}^d}W_{\bar{b}}$$

to 2^d subspaces by the distributive law for tensor products. For any $x\in G^d$ the operator P_x preserves the decomposition.

Corollary 2.7. *For any Cayley graph $C(G^d, \bar{Y})$, the normalized adjacency operator $\mathbb{E}_{x\in\bar{Y}}[P_x]$ preserves the given decomposition of $W^{\otimes d}$, so*

$$\lambda(G^d, \bar{Y}) = \max_{\bar{b}\neq ||^d} \max_{w\in W_{\bar{b}}} \| \mathop{\mathbb{E}}_{x\in\bar{Y}} [P_x(w)] \| / \|w\| \ .$$

That is, it suffices to bound $\|\mathbb{E}_{x\in\bar{Y}}[P_x(w)]\|$ from above, for vectors w that are purely in one of these $2^d - 1$ subspaces.

The following two observations describe cases where we can ignore part of the coordinates of $x \in G^d$ when trying to estimate $\|\mathbb{E}_{x\in\bar{Y}}[P_x(w)]\|$.

Observation 2.8. Let $\bar{b} = b_1, \dots, b_d$ where $b_i = \perp$ for $i \leq r$ and $b_i = ||$ otherwise. For $w \in W_{\bar{b}} = W_{\perp}^{\otimes r} \otimes W_{||}^{\otimes(d-r)}$, the value of $P_x(w)$ does not depend on $x_{r+1}, \dots, x_d$.

Proof. w is a real function on G^d. The statement that $w \in W_{\bar{b}}$ and $b_i = ||$ means that w does not depend on the i-th coordinate of its input. □

Observation 2.9. Let $\bar{X} \subset G^d$ be a set of group elements whose last $d-r$ coordinates constitute some fixed vector $\bar{x} \in G^{d-r}$. Then for every $w \in W^{\otimes d}$ the value of

$$\| \mathop{\mathbb{E}}_{x\in\bar{X}} [P_x(w)] \|$$

does not depend on $\bar{x}$.

Observation 2.3 from Section 2.1 translates nicely to the Cayley graph world

Observation 2.10. Let G be a group, $Y \subset G$. Define Z to be the set of all words of length k in Y. Then $\lambda(G,Z) = \lambda(G,Y)^k$.

We end with an observation which simplifies the proof of explicitness for families of Cayley graphs.

Observation 2.11. A family of Cayley graphs $C(G_n, Y_n)$ is explicit if there are polynomial time algorithms in $\log|G_n|$ for

- performing group multiplication in G_n,
- computing inverses in G_n, and
- computing the set Y_n.

2.2.2 Wreath products and the zig-zag product

Let A and B be finite groups. Assume that $B \subset S_d$, that is, it acts by permutations on the set $[d] = \{1, \dots, d\}$. Define the *wreath product* [1] $A \wr B$ of A and B to be the group whose elements are vectors $(a_1, \dots, a_d, \sigma)$, where $a_i \in A$ for all i, and $\sigma \in B$. The group multiplication rule is

$$(a_1, \dots, a_d, \sigma) \cdot (\tilde{a}_1, \dots, \tilde{a}_d, \tau) = (a_{\tau(1)}\tilde{a}_1, \dots, a_{\tau(d)}\tilde{a}_d, \sigma\tau) \ .$$

[1] More precisely, this is referred to as the *permutational* wreath product in the literature.

One can check that this defines a group structure on $A \wr B$. The wreath product is a special case of a more general construction - it is the *semi-direct product* of A^d and B, where A^d is the Cartesian product of d copies of A. The groups A^d, B are naturally embedded in $A \wr B$, and we will sometimes refer to elements of A^d and B as elements of $A \wr B$.

Let $\alpha \subset A^d, \beta \subset B$ be sets of generators. Suppose α has a special structure: it is a *single B-orbit*. This means that for some arbitrary $\bar{a} \in \alpha$, the set α consists of all vectors obtained from $\bar{a}$ by permuting its coordinates by a permutation in B. We now define a set γ in $A \wr B$ by $\gamma = \{x\bar{a}y \mid x, y \in \beta\}$. One can check that γ generates $A \wr B$. The following theorem from [4], following the zig-zag theorem of [42], shows that if α, β are sufficiently good expanding generators then so is γ.

Theorem 2.12. [4] *If α is a single B-orbit then $\lambda(A \wr B, \gamma) \leq \lambda(A^d, \alpha) + \lambda(B, \beta)$.*

Note that $|\gamma| = |\beta|^2$ depends only on the size of β, while α could be large (it could be as large as $|B|$). Also, it is easy to compute γ given α and β, as multiplications in $A \wr B$ can be computed efficiently.

2.2.3 The commutator property

Let A be a group. For $g, h \in A$ define the *commutator* $[g, h]$ to be $ghg^{-1}h^{-1}$. A has the *commutator property* if for every element $a \in A$ there is a solution in the variables x, y to the equation $a = [x, y]$. (Note that this is a stronger property than just the commutator subgroup $[A, A]$ being equal to A.) Nikolov [37] proves

Theorem 2.13. [37] *Let A be a group, and $B \subset S_d$ a group of permutations. If A, B have the commutator property then so does $A \wr B$.*

We shall need an algorithmic version of this theorem. For a group A, a *commutator representation algorithm* gives, for an input $a \in A$, some pair $x, y \in A$ such that $a = [x, y]$.

Theorem 2.14. *Let A, B be as in Theorem 2.13. Suppose we are given commutator representation algorithms for the groups A, B. Then we obtain such an algorithm for $A \wr B$. This algorithm calls the algorithm on B one time, and the algorithm on A at most d times, and uses at most $O(d)$ extra multiplication operations on A, B. (The description of the algorithm appears in the proof of the theorem.)*

We prove the theorem in Section 6.

3 Overview of the construction

In Section 3.1 we will define our sequence of groups G_n. In Section 4 we will show how to find generating subsets $Y_n \subset G_n$ that give $\lambda(G_n, Y_n) < 1/1000$ with bounded size $|Y_1|^4$. This will be based on the assumption that there exists a small enough subset Y_1 of the alternating group A_d such that $\lambda(A_d, Y_1) < 1/1000$.

3.1 The family of groups

Definition 3.1. The groups in our construction are defined by $G_1 = A_d$ and, inductively, $G_{n+1} = G_n \wr A_d$.

Another way to view the group G_n is as a subgroup of the full group of symmetries of the d-regular, depth n tree (by d-regularity here we mean that each inner vertex has d descendants). Each element in the group of symmetries is uniquely defined by writing a permutation on each internal node of the tree, indicating how the children of this vertex are permuted. In the subgroup G_n all these permutations should be *even*. The representation of an element of G_n as a list of even permutations is polynomial in $\log|G_n|$. Multiplying two elements and inverting an element can be done in time which is polynomial in the size of this representation

The following important corollary of Theorem 2.14 shows that for our groups G_n there is an efficient commutator representation algorithm.

Lemma 3.2. *If $d \geq 5$ then the groups G_n have the commutator property of Section 2.2.3. Moreover, G_n has a commutator representation algorithm that runs in time polynomial in $\log|G_n|$.*

Proof. $G_1 = A_d$, and by [38] it has the commutator property. By induction, using Theorem 2.13, every G_n has the commutator property. The existence of an efficient commutator representation algorithm follows from Theorem 2.14. Full details are given in Section 6. □

4 Main theorem

Theorem 4.1. *Suppose that for some d there exists a set of generators $Y_1 \subset A_d$ such that $\lambda(A_d, Y_1) < 1/1000$ and $|Y_1| \leq d^{1/28}/10^{40}$. Then there exist sets $Y_n \subset G_n$ such that $\lambda(G_n, Y_n) < 1/1000$ and $|Y_n| \leq d^{1/7}/10^{40}$. Furthermore, Y_n can be computed in time polynomial in $\log|G_n|$.*

The graphs $C(G_n, Y_n)$ are the required sequence of Cayley graphs. The sets Y_n can be computed efficiently, and we saw in Section 3.1 that group operations in G_n can also be computed efficiently, so by Observation 2.11 this is an explicit family of Cayley graphs.

The assumption of the theorem is true for very large d:

Theorem 4.2 ([22]). *For every integer $d \geq 0$ there exists a subset U_d of the symmetric group S_d such that $|U_d| \leq 10^{10^7}$ and $\lambda(S_d, U_d) \leq 1/1000$.*

Corollary 4.3. *If $d \geq 10^{10^9}$ Then the conditions of Theorem 4.1 hold.*

We will construct the expanding generators $Y_n \subset G_n$ inductively. The basis of the induction is the assumption in the theorem about $G_1 = A_d$.

Let $G = G_n$. We are given $Y \subset G$ such that $\lambda(G, Y) < 1/1000$ and $|Y| \leq d^{1/7}/10^{40}$. We want to find a set $Y' \subset G \wr A_d$ such that $\lambda(G \wr A_d, Y') < 1/1000$ and $|Y'| \leq d^{1/7}/10^{40}$. We will use Theorem 2.12. The theorem requires an expanding generating set for A_d (which we already have), and an expanding generating set $T \subset G^d$ which is a single A_d-orbit. Given any element of such T, Theorem 2.12 produces (explicitly) an expanding generating set for $G \wr A_d = G_{n+1}$.

Can we find an expanding, single-orbit, generating set for G^d? Here is a simple attempt that fails. Take $T = Y^d$. The set Y^d is expanding, as $\lambda(G^d, Y^d) = \lambda(G, Y)$ by Claim 2.2. Unfortunately, Y^d contains exponentially many A_d-orbits. Another natural set to consider in G^d is the set of *balanced vectors*:

Definition 4.4. Let G be a group, and $Y \subset G$. For $d > |Y|$, define $Y^{(d)}$ to be the vectors in Y^d in which every $u \in Y$ appears exactly $\lfloor d/|Y| \rfloor$ times, and the rest of the elements are $1 \in G$. We call these vectors **balanced vectors**. Every two elements in the set $Y^{(d)}$ are equal up to a permutation of the coordinates. Since $d > |Y|$ we may assume that the permutation is even. In other words, the set $Y^{(d)}$ is a single A_d-orbit.

The set $Y^{(d)}$ looks promising, but is it expanding? Not always. If G is Abelian, $Y^{(d)}$ does not even *generate* G^d, since every element in $Y^{(d)}$ has product of coordinates equal to 1 (Y is symmetric, and every element of Y appears the same number of times in $Y^{(d)}$). The groups G_n are far from being Abelian. Indeed, every element of G_n has a representation as a commutator. It turns out that this property, along with the existence of a small generating set Y for G (assumed by induction) enables us to find a good generating set for G^d. We will enlarge Y somewhat to a set $X \supset Y$, and see that $X^{(d)}$ is expanding for G^d.

Definition 4.5. Let G be a group, and let $Y \subset G$. Suppose every element $y \in Y$ can be written as a commutator in G, namely $y = a_y b_y a_y^{-1} b_y^{-1}$ for some $a_y, b_y \in G$. Define

$$Y^* = \bigcup_{y \in Y} \{a_y, b_y, a_y^{-1}, b_y^{-1}, a_y^{-1} b_y^{-1}, b_y a_y\} \cup \{1\} \ .$$

Y^* is symmetric, and $|Y^*| \leq 7|Y|$.

Theorem 4.6. *Let G be a group. Suppose that every element of Y is a commutator in G. Let $c, k \in \mathbb{N}$ be constants (to be chosen later). Define $c \cdot Y \subset G$ to be the multi-subset where every element of Y appears c times. Define $X = (c \cdot Y) \cup Y^*$, and $\lambda = \lambda(G, Y)$. If $d \geq k^2 \cdot |X|^7$ then*

$$\lambda(G^d, X^{(d)}) < 0.01 + \max\left\{(\lambda + 7/c), e^{-kc(1-\lambda)/10^6}\right\}$$

where $X^{(d)}$ is the set of balanced vectors.

The proof is given in Section 5. To get a feeling for the constants, note that the larger k and c are, the better inequality we get in the theorem. k is large when X is small. c is large when X is much larger that Y, so k gets smaller when c gets larger. Nevertheless, it is not difficult to make both of them large enough for our purposes.

Theorem 4.6 is the required result for the inductive step - it remains to show that we can choose c, k properly such that $\lambda(G^d, X^{(d)})$ is small enough for Theorem 2.12.

We proceed with the inductive step. We are given a set $Y_n \subset G_n$ of size at most $|Y_1|^4$ such that $\lambda(G_n, Y_n) < 1/1000$. Apply Theorem 4.6 (with $c = 10^3, k = 10^5$). Then the conditions of Theorem 4.6 hold, and we obtain a set $X^{(d)} \subset G^d$ such that $\lambda(G, X^{(d)}) < 1/50$ (just substitute our k, c in the theorem to see this). Apply Theorem 2.12 to obtain a subset $P \subset G_{n+1}$ of size $|Y_1|^2$, and $\lambda(G_{n+1}, P) < 1/1000 + 1/50$. Define Y_{n+1} to be the set of all words of length 2 in P. This is a set of size $|Y_1|^4$ and (by Observation 2.10) $\lambda(G_{n+1}, Y_{n+1}) < (1/1000 + 1/50)^2 < 1/1000$. This completes the inductive step. $\square$

5 Proof of Theorem 4.6

The theorem appears in Section 4. Let G,Y,X,λ be as defined in Theorem 4.6. We will use the notation $W = W(G)$ and $W(G^d), W_{\bar{b}}$ defined in Section 2.2.1. We need to prove that for every $w \in W(G^d)_\perp$ such that $\|w\| = 1$, at least one of the following upper bounds holds:

$$\| \mathop{\mathbb{E}}_{x \in X^{(d)}} [P_x(w)] \| \leq 0.01 + \lambda + \frac{7}{c} \quad \text{, or} \tag{5.1}$$

$$\| \mathop{\mathbb{E}}_{x \in X^{(d)}} [P_x(w)] \| \leq 0.01 + e^{-kc(1-\lambda)/10^6} \quad . \tag{5.2}$$

We saw in Section 2.2.1 that it is enough to prove this for $w \in W_{\bar{b}}$ when $\bar{b} \neq \{||\}^d$. Since $X^{(d)}$ is invariant under permutation of the coordinates it is enough to prove the inequality for every $w \in W_\perp^{\otimes r} \otimes W_{||}^{\otimes (d-r)}$ where $1 \leq r \leq d$ (this is $W_{\bar{b}}$ for $b_i = \perp$ for $1 \leq i \leq r$ and $b_i = ||$ for $r < i \leq d$).

We split the proof to small and large r cases. For small r we will prove inequality (5.1), and for large r we will prove inequality (5.2).

Small r case: When $r \leq 0.1\sqrt{d/|X|}$, the first r coordinates of a random element in $X^{(d)}$ are very closely a random element in X^r. By Observation 2.8 $P_x(w)$ only depends on the first r coordinates of x, so it is enough to bound $\| \mathbb{E}_{x \in X^r} [P_x(w)] \|$ for $w \in W_\perp^{\otimes r}$. By Claim 2.2 $\| \mathbb{E}_{x \in X^r} [P_x(w)] \| \leq \lambda(G,X)^r$. The worst case is when $r = 1$. As $Y \subset X$ we can use Claim 2.4 to give an upper bound to $\lambda(G,X)$, and we obtain inequality (5.1). This part is relatively easy, and we will not give a more detailed proof. Notice however that the argument for small r works for *any* group G, not only for our special sequence of groups, and from the generating set X we only used the Y part - not the Y^* part.

Large r case: When r is large the result is no longer true for any group (for any Abelian group there exists an $f \in W^{\otimes d}$ such that $P_y(f) = f$ for all $y \in Y^{(d)}$). We will need the Y^* part of the generating set X (recall that it is only defined when every element of G is a commutator). We will start with the analysis of a different graph - the Cayley graph $C(G \times G, \{(y,y^{-1}) | y \in Y^*\})$. We give a lower bound of $(1 - \lambda(G,Y))/21|Y^*|^2$ on the spectral gap of this graph in Section 5.1. Afterward, in Section 5.2, we will give an upper bound on $\| \mathbb{E}_{x \in X^{(d)}} [P_x(w)] \|$ using the spectral gap of this graph on $G \times G$. This part is again true for every group G, not only our groups.

Notice that the spectral gap bound we get in the $G \times G$ case is rather weak - much smaller than the spectral gap of the original graph $C(G,Y)$. When r is large enough we are able to apply the $G \times G$ result many times in parallel, amplifying the weaker upper bound in $G \times G$. We will obtain the upper bound (5.2).

5.1 Expansion of $G \times G$ with correlated generators

Definition 5.1. Let G be a group, and let $Y \subset G$ be a subset of G. Define

$$\widetilde{Y} = \{(y,y^{-1}) \mid y \in Y\} \quad .$$

Theorem 5.2. *Suppose $\lambda(G,Y) < 1 - \varepsilon$ for some ε, and that every element of Y is a commutator in G. Then*

$$\lambda(G \times G, \widetilde{Y^*}) \leq 1 - \frac{\varepsilon}{21|Y^*|^2} \quad .$$

We find Theorem 5.2 quite surprising. In the set $\widetilde{Y}$ there is *complete correlation* between the two coordinates, and it would seem that this correlation would prevent the graph from being an expander. For example, if G is Abelian and Y generates G then $\widetilde{Y}$ does not even generate $G \times G$, but only the subgroup $\{(g, g^{-1}) \mid g \in G\}$. Also, for any group G the set $\{(y, y) \mid y \in Y\}$ only generates the subgroup $\{(g, g) \mid g \in G\}$. In both cases the correlation in the generating set prevents the graph from being an expander. We manage to decouple this correlation in the case of the special generating set Y^*, whose existence relies on the commutator property of G.

The key observation is that we can represent the element $(y, 1)$ for any $y \in Y$ as a word of length 3 in $\widetilde{Y^*}$. We prove this in the following observation.

Observation 5.3. Let Z be the set of words of length 3 in the set $\widetilde{Y^*}$. Then

$$C(G \times G, \{(Y, 1) \cup (1, Y)\}) \subset C(G \times G, Z) \enspace .$$

Proof. Recall that for every $y \in Y$ the set Y^* contains the elements $a_y, b_y, a_y^{-1} b_y^{-1}$ where $y = a_y b_y a_y^{-1} b_y^{-1}$. Observe that

$$(a_y, a_y^{-1}) \cdot (b_y, b_y^{-1}) \cdot ((a_y^{-1} b_y^{-1}), (a_y^{-1} b_y^{-1})^{-1}) = (y, 1) \enspace .$$

This gives the required representation of $(y, 1)$. We can obtain $(1, y)$ similarly. ☐

It is easy to see that if $C(G, Y)$ has spectral gap ε then the graph $C(G \times G, \{(Y, 1) \cup (1, Y)\})$ has spectral gap $\varepsilon/2$. We now have the decoupling we were looking for - the correlated generating set Z contains the uncorrelated one $(Y, 1) \cup (1, Y)$. More precisely, apply Claim 2.4 to observation 5.3, and deduce that

Observation 5.4. $C(G \times G, Z)$ has spectral gap at least $\varepsilon/7|Y^*|^2$.

Recall that Z consists of all words of length 3 in the $\widetilde{Y^*}$. By Observation 2.10, the spectral gap of $C(G \times G, \widetilde{Y^*})$ is at most 3 times smaller than the spectral gap of $C(G \times G, Z)$, and the theorem is proved.

5.2 Reduction to $G \times G$

We bound the average $\| \mathbb{E}_{x \in X^{(d)}} [P_x(w)] \|$ from above in terms of $\lambda(G \times G, \widetilde{Y^*})$ from Section 5.1.

For $x \in X^d$ write $x = (x_1, x_2, \bar{x})$ where $x_1, x_2 \in G$ and $\bar{x} \in G^{d-2}$. By the triangle inequality,

Claim 5.5. *For every $w \in W^{\otimes d}$*

$$\| \mathop{\mathbb{E}}_{x \in X^{(d)}} [P_x(w)] \| \le \mathop{\mathbb{E}}_{x \in X^{(d)}} \| (P_{x_1, x_2, \bar{x}} + P_{x_2, x_1, \bar{x}})(w)/2 \| \enspace .$$

By Observation 2.9 the value of $\| (P_{x_1, x_2, \bar{x}} + P_{x_2, x_1, \bar{x}})(w)/2 \|$ only depends on the first two coordinates of x. We therefore group together all the x with equal x_1, x_2, replacing $\bar{x}$ by $\bar{1}$, a $(d-2)$-length vector of 1's, and it is enough to bound $\mathbb{E}_{x \in X^{(d)}} \| (P_{x_1, x_2, \bar{1}} + P_{x_2, x_1, \bar{1}})(w)/2 \|$. The number of times each pair x_1, x_2 appears in the average above is proportional to the number of extensions of x_1, x_2 to a vector $(x_1, x_2, \bar{x}) \in X^{(d)}$. As d is much larger than 2, the number of such extensions is nearly equal for every pair x_1, x_2, and we obtain the following:

Claim 5.6. *If $d \geq 100|X|$ then for every $w \in W^{\otimes d}$*

$$\mathop{\mathbb{E}}_{x \in X^{(d)}} \|(P_{x_1,x_2,\bar{x}} + P_{x_2,x_1,\bar{x}})(w)/2\| \leq \mathop{\mathbb{E}}_{y \in X^2} \|(P_{y_1,y_2,\bar{1}} + P_{y_2,y_1,\bar{1}})(w)/2\| + 0.01\|w\| \ .$$

The 0.01 above pays for the fact that the number of extensions is only nearly equal.

The following lemma bounds the RHS of Claim 5.6.

Lemma 5.7. *If $\lambda(G,Y) < 1-\varepsilon$ and $r \geq 2$ then for every $w \in W_\perp^{\otimes r} \otimes W^{\otimes(d-r)}$*

$$\mathop{\mathbb{E}}_{y \in X^2} \|(P_{y_1,y_2,\bar{1}} + P_{y_2,y_1,\bar{1}})(w)/2\| \leq (1 - \frac{c\varepsilon}{2 \cdot 10^4 |X|^3})\|w\| \stackrel{\text{def}}{=} \Delta\|w\| \ .$$

We prove the lemma in Section 5.2.1.

Combining Claim 5.6 and Lemma 5.7 we obtain

$$\left\| \mathop{\mathbb{E}}_{x \in X^{(d)}} [P_x(w)] \right\| \leq (\Delta + 0.01)\|w\|$$

but Δ is too close to 1. The problem originates from Claim 5.5, where we partitioned the set $X^{(d)}$ into pairs based on the value of the first 2 coordinates. This partition turns out to be too coarse. We will use a finer partition of $X^{(d)}$ by looking at the first t pairs of coordinates, for some properly chosen $t \leq r$. This will amplify the bound to Δ^t.

We now define this finer partition precisely. Let $H_t < S_d$ be the subgroup (of size 2^t) generated by the transpositions $(2k-1, 2k)$ for $1 \leq k \leq t$, and group together the elements $\{\sigma(x) \mid \sigma \in H_t\}$. When $t = 1$ we get the grouping into pairs discussed above. The argument leading to Claim 5.6 shows the following:

Claim 5.8. *If $2t \leq 0.1\sqrt{d/|X|}$ then for every $w \in W^{\otimes d}$*

$$\left\| \mathop{\mathbb{E}}_{x \in X^{(d)}} [P_x(w)] \right\| \leq \mathop{\mathbb{E}}_{y \in X^{2t}} \left\| \mathop{\mathbb{E}}_{\sigma \in H_t} \left[P_{\sigma(y,\bar{1})}(w) \right] \right\| + 0.01 \ .$$

The case $t = 1$ is Claim 5.6. However, the weak upper bound Δ we had for $t = 1$ amplifies to Δ^t.

Claim 5.9. *Suppose that for every $w \in W_\perp^{\otimes 2} \otimes W^{\otimes d-2}$*

$$\mathop{\mathbb{E}}_{y \in X^2} \left\| \frac{1}{2}(P_{y_1,y_2,\bar{1}} + P_{y_2,y_1,\bar{1}})(w) \right\| \leq \Delta\|w\| \ .$$

Then for every $w \in W_\perp^{\otimes 2t} \otimes W^{\otimes d-2t}$

$$\mathop{\mathbb{E}}_{y \in X^{2t}} \left\| \mathop{\mathbb{E}}_{\sigma \in H_t} \left[P_{\sigma(y,\bar{1})}(w) \right] \right\| \leq \Delta^t \|w\| \ .$$

The notation $\bar{1}$ denotes a vector of length $d-2$ in the first inequality, and a vector of length $d-2t$ in the second one.

Proof. The proof is by induction on t. The case $t=1$ is the assumption of the claim. For general t

$$\mathop{\mathbb{E}}_{y\in X^{2t}}\left\|\mathop{\mathbb{E}}_{\sigma\in H_t}\left[P_{\sigma(y,\bar{1})}(w)\right]\right\| = \mathop{\mathbb{E}}_{\substack{z\in X^2\\ y\in X^{2(t-1)}}}\left\|\mathop{\mathbb{E}}_{\sigma\in H_{t-1}} P_{\sigma(1,1,y,\bar{1})}\left[P_{z_1,z_2,\bar{1}}+P_{z_2,z_1,\bar{1}}(w)\right]\right\|$$
$$\leq \Delta^{t-1}\mathop{\mathbb{E}}_{z\in X^2}\left\|(P_{z_1,z_2,\bar{1}}+P_{z_2,z_1,\bar{1}})(w)\right\| \leq \Delta^t\|w\| \ .$$

Note that in the second line above $\sigma\in H_{t-1}$ acts on the vector y - not on the first $2t-2$ coordinates. The first inequality follows from the induction hypothesis for H_{t-1}. The second inequality follows from the induction hypothesis for H_1. □

We can now complete the proof using $\lambda(G,Y)<1-\varepsilon$. Pick an integer t satisfying

$$0.05\sqrt{d/|X|}\leq 2t\leq 0.1\sqrt{d/|X|}\leq r \ .$$

Then by the claims in this section, for $w\in W_{\perp}^{\otimes r}\otimes W^{\otimes d-r}$ of norm 1,

$$\left\|\mathop{\mathbb{E}}_{x\in X^{(d)}}[P_x(w)]\right\| \leq 0.01+\left(1-\frac{c\varepsilon}{2\cdot 10^4|X|^3}\right)^t \leq 0.01+\exp\left(\frac{-ct\varepsilon}{2\cdot 10^4|X|^3}\right)\leq 0.01+\exp\left(\frac{-kc\varepsilon}{10^6}\right) \ .$$

We plugged in $2t\geq 0.05\sqrt{d/|X|}\geq 0.05k|X|^3$. This concludes the proof of Theorem 4.6 for large r.

5.2.1 Proof of Lemma 5.7

Let τ be the spectral gap of $C(G\times G,\{(y,y^{-1})\mid y\in Y^*\})$. From Theorem 5.2 we have for every $u\in W_{\perp}\otimes W$

$$\left\|\mathop{\mathbb{E}}_{y\in Y^*}\left[P_{y,y^{-1}}(u)\right]\right\|\leq(1-\tau)\|u\| \ . \tag{5.3}$$

In Lemma 5.7 we want to bound

$$\mathop{\mathbb{E}}_{y\in X^2}\|(P_{y_1,y_2,\bar{1}}+P_{y_2,y_1,\bar{1}})(w)/2\| \tag{5.4}$$

from above, for every $w\in W_{\perp}^{\otimes r}\otimes W^{\otimes(d-r)}$.

We will start with the case $d=2$. We will bound (5.4) in terms of the LHS of (5.3). In order to do that, we will have to deal with the fact that the norm in (5.3) appears outside the expectation, while in (5.4) it appears inside the expectation (see Claim 5.10). Also, the average in (5.4) is over $y\in X^2$, while in (5.3) the average is over $y\in Y^*$ (see Claim 5.11). After completing the proof in the case $d=2$, we turn to prove the lemma for general d (Claim 5.12).

Claim 5.10. *For every $u\in W_{\perp}\otimes W$*

$$\mathop{\mathbb{E}}_{y\in Y^*}\left\|\frac{1}{2}(P_{y,1}+P_{1,y})(u)\right\|\leq(1-\tau/4)\|u\| \ .$$

Proof. From Claim 2.5 and (5.3)

$$\mathop{\mathbb{E}}_{y\in Y^*}\left\|\frac{1}{2}(P_{y,y^{-1}}+I)(u)\right\| \le (1-\tau/4)\|u\| \ .$$

Applying the unitary operator $P_{1,y}$ to each element above proves the claim. □

Claim 5.11. *For every $u \in W_\perp \otimes W$*

$$\mathop{\mathbb{E}}_{y\in X^2}\left\|\frac{1}{2}(P_{y_1,y_2}+P_{y_2,y_1})(u)\right\| \le (1-\frac{\tau}{8c|X|})\|u\| \ .$$

Proof. Let p be the probability that for a random $y \in X^2$ we have $y_1 \in Y^*$ and $y_2 = 1$. Then $p \ge (1/2c)\cdot 1/|X|$ (as $X = c\cdot Y \cup Y^*$ and Y^* is larger than Y). Using a convexity argument similar to Claim 2.4 we see that

$$\begin{aligned}\mathop{\mathbb{E}}_{y\in X^2}\left\|\frac{1}{2}(P_{y_1,y_2}+P_{y_2,y_1})(u)\right\| &\le p\cdot \mathop{\mathbb{E}}_{y\in Y^*}\left\|\frac{1}{2}(P_{y,1}+P_{1,y})(u)\right\| + (1-p)\cdot\|u\| \\ &\le p\cdot(1-\tau/4)\|u\| + (1-p)\|u\| \le (1-p\tau)\|u\| \le (1-\frac{\tau}{8c|X|})\|u\|\end{aligned}$$

which proves Claim 5.11. □

We have shown that for every $u \in W_\perp \otimes W$

$$\mathop{\mathbb{E}}_{y\in X^2}\left\|\frac{1}{2}(P_{y_1,y_2}+P_{y_2,y_1})(u)\right\| \le \left(1-\frac{\tau}{8c|X|}\right)\|u\| \le \left(1-\frac{\varepsilon}{21\cdot 8|Y^*|^2\cdot|X|}\right)\|u\| \le \left(1-\frac{c\varepsilon}{2\cdot 10^4|X|^3}\right)\|u\| \ .$$

The last step follows from $|Y^*| \le 10|X|/c$ (which is true since $X = cY \cup Y^*$ and $|Y^*| \le 10|Y|$).

We have almost completed proving the lemma. We have the right upper bound, but for $u \in W^{\otimes 2}$ instead of in $W^{\otimes d}$.

Claim 5.12. *If there exists $\lambda > 0$ such that for every $u \in W_\perp^{\otimes 2}$*

$$\mathop{\mathbb{E}}_{y\in X^2}\left\|[\frac{1}{2}(P_{y_1,y_2}+P_{y_2,y_1})(u)]\right\| \le \lambda\|u\|$$

then for every $w \in W_\perp^{\otimes 2}\otimes W^{\otimes(d-2)}$

$$\mathop{\mathbb{E}}_{y\in X^2}\left\|[\frac{1}{2}(P_{y_1,y_2,\bar{1}}+P_{y_2,y_1,\bar{1}})(w)]\right\| \le \lambda\|w\| \ .$$

Proof. Write $w \in W_\perp^{\otimes r}\otimes W^{\otimes(d-r)}$ as $w = \sum u_i \otimes v_i$ where $u_i \in W_\perp^{\otimes 2}$ and $v_i \in W^{\otimes(d-2)}$, such that the v_i are orthogonal and $\|v_i\| = 1$. We have

$$\mathop{\mathbb{E}}_{y}\left\|[\frac{1}{2}(P_{y_1,y_2,\bar{1}}+P_{y_2,y_1,\bar{1}})(w)]\right\|^2 = \mathop{\mathbb{E}}_{y}\sum_i\left\|[\frac{1}{2}(P_{y_1,y_2}+P_{y_2,y_1})(u_i)]\right\|^2 \le \lambda^2\|w\|^2$$

and the result follows since $\mathbb{E}(X)^2 \le \mathbb{E}(X^2)$ for any random variable X. □

6 Proof of Theorem 2.14

The theorem appears in Section 2.2.3.

Remark 6.1. This section contains equations in groups. Constants in the equations will be written in Greek letters. Variables will be written in small Latin letters. Vectors of length d are underlined.

Let $C = A \wr B$, where A is any group and $B \subset S_d$. Given an element $\gamma \in C$ we look for a "commutator representation algorithm" that solves the equation $\gamma = [c_1, c_2] := c_1 c_2 {c_1}^{-1} {c_2}^{-1}$. By assumption we have such an algorithm for A and B. The proof below extends Nikolov's proof in [37].

Any element $\gamma \in A \wr B$ has a unique representation $c = \beta \cdot \underline{\alpha}$ with $\beta \in B$, $\underline{\alpha} \in A^d$, so it is enough to solve, for every pair $(\beta \in B, \underline{\alpha} \in A^d)$, the equation $\beta\underline{\alpha} = [b_1\underline{x}, b_2\underline{y}]$. Now

$$[b_1\underline{x}, b_2\underline{y}] = [b_1, b_2] \cdot \underline{x}^{b_2 b_1^{-1} b_2^{-1}} \underline{y}^{b_1^{-1} b_2^{-1}} \underline{x}^{-b_1^{-1} b_2^{-1}} \underline{y}^{-b_2^{-1}}$$

where $\underline{x}^b = b^{-1}\underline{x}b$. In our case $\underline{x}^b$ is simply a permutation of the coordinates of $\underline{x}$ by $b \in B \subset S_d$.

We obtain a pair of equations:

$$\beta = [b_1, b_2], \text{ and}$$
$$\underline{\alpha} = \underline{x}^{b_2 b_1^{-1} b_2^{-1}} \underline{y}^{b_1^{-1} b_2^{-1}} \underline{x}^{-b_1^{-1} b_2^{-1}} \underline{y}^{-b_2^{-1}} \quad .$$

By assumption there is an algorithm that solves $\beta = [b_1, b_2]$. Fix some solution $b_1 = \beta_1, b_2 = \beta_2$. It remains to solve

$$\underline{\alpha} = \underline{x}^{\beta_2 \beta_1^{-1} \beta_2^{-1}} \underline{y}^{\beta_1^{-1} \beta_2^{-1}} \underline{x}^{-\beta_1^{-1} \beta_2^{-1}} \underline{y}^{-\beta_2^{-1}} \quad .$$

Since $\underline{x}^\beta$ is a permutation (depending on β) of the coordinates of $\underline{x}$, the following lemma solves a more general system of equations.

Lemma 6.2. *For any four permutations $\sigma_1, \sigma_2, \sigma_3, \sigma_4 \in S_d$ and for any $\underline{\alpha} = \alpha_1, \ldots, \alpha_d \in A^d$, the following system of d equations, one for each $1 \le i \le d$:*

$$\alpha_i = x_{\sigma_1(i)} y_{\sigma_2(i)} x^{-1}_{\sigma_3(i)} y^{-1}_{\sigma_4(i)}$$

has a solution algorithm that calls the commutator representation algorithm on A at most d times, and does at most $O(d)$ operations in the group A.

The rest of this section is devoted to the proof of this lemma.

Definition 6.3. We shall refer to the α_i as **constants** and to the $x_i, y_i, x_i^{-1}, y_i^{-1}$ as **literals**.

There are d constants and $4d$ literals in our system. An important fact is that each literal appears *exactly once* in the system.

Let us solve first in the case that all four σ_i are the identity permutation. The system in this case is:

$$\alpha_1 = [x_1, y_1]$$
$$\alpha_2 = [x_2, y_2]$$
$$\ldots$$
$$\alpha_d = [x_d, y_d]$$

In this case the equations are independent (no variable appears in more than one equation). Each equation asks for a commutator representation for $\alpha_i \in A$. We solve the system of equations by calling the commutator representation algorithm for A for each equation separately.

The solution for general σ_i is by reduction to a system similar to the one we obtained for the $\sigma_i = 1$ case. As long as there are variables that appear in more than one equation, we will remove equations by "Gaussian elimination," until we obtain a system of independent equations. We will then translate each equation to a commutator representation equation like the ones above.

As mentioned, each literal appears exactly once in the system. If x_i, x_i^{-1} do not both appear in the same equation, then we can eliminate x_i, x_i^{-1} from the system by substitution (paying $O(1)$ multiplications in A). This reduces the number of equations in the system by 1. Repeat the substitution operation until it is no longer possible. Notice that the property that each literal appears exactly once is preserved along the way.

Claim 6.4. *The substitution process ends with $L \le d$ equations*

$$\delta_l = W_l \quad \forall l \in \{1, \ldots, L\}$$

where W_l is some word in literals and constants. The equations are now independent - every literal appears in the same equation as its inverse, or they both do not appear in the system. □

We will now reduce this system to L commutator representation problems in the group A. The following lemma finds a "hidden commutator" in each of the words W_l:

Lemma 6.5. *[37] In every W_l there exist $g, h \in \{1, 2, \ldots, d\}$ depending on l, such that*

$$W_l = Z_1 x_g Z_2 y_h Z_3 x_g^{-1} Z_4 y_h^{-1} Z_5$$

where the Z_i are words in literals and constants from the word W_l (they do not contain $x_g^{\pm 1}, x_h^{\pm 1}$ since each literal appears at most once in the system of equations).

The proof is in [37]. Given that such a hidden commutator exists, it is easy to find one in time polynomial in d by looking at all the literals appearing in W_l (there are at most $2d$ of those). Substitute every variable appearing in the Z_i by 1. This does not affect any other equation - the equations are independent at this point. We obtain a new equation

$$\delta_l = \zeta_1 x_g \zeta_2 y_h \zeta_3 x_g^{-1} \zeta_4 y_h^{-1} \zeta_5 \ .$$

This is now an equation in two variables x_g, x_h - all the other words are constants. This is almost a "commutator representation" equation. Indeed, if the five ζ_i are all equal 1, we obtain the equation

$$\delta_l = [x_g, y_h]$$

which is solved by calling the commutator algorithm on A. For general ζ_i we transform the "hidden" commutator to a "real" commutator by changing variables. Define $\tilde{x}_g = \zeta_3 x_g \zeta_4$ and $\zeta_h = y_h \zeta_2^{-1} \zeta_3^{-1}$. Observe that

$$\delta_l = \zeta_1 \zeta_4 [\tilde{x}_g, \tilde{y}_h] \zeta_3 \zeta_2 \zeta_5 \ .$$

Rewrite this equation as

$$(\zeta_1\zeta_4)^{-1}\delta_l(\zeta_3\zeta_2\zeta_5)^{-1} = [\tilde{x}_g, \tilde{y}_h] \ .$$

The LHS is some constant element in A, and the equation requests a representation of this element as a commutator. We can find a solution by calling the commutator representation algorithm on A. The solution is in the variables $\tilde{x}_g, \tilde{y}_h$, but this is easily translated to a solution in our original variables x_g, y_h.

How many operations did we use? We called the commutator representation algorithm in A at most d times (one call for each final equation $v_l = W_l$). We called the commutator representation algorithm on B one time. We used $O(1)$ multiplications in B, and $O(d)$ multiplications in A (there were $O(1)$ per either removing an equation or solving a final equation).

We can now deduce Lemma 3.2. Define $m(n)$ to be the cost (in bit operations) of multiplication in G_n, and define $c(n)$ to be the cost of computing the commutator representation of an element in G_n. As $m(n+1) < (d+1)m(n)$ and $m(1) = O(d^2)$ we deduce that $m(n) < (d+1)^{n+2} \cdot O(1)$. From the discussion above we see that $c(n+1) < (d+1)c(n) + m(n) \cdot O(d) < (d+1)c(n) + d^{n+3} \cdot O(1)$. This implies that $c(n) < d^{4n} \cdot O(1)$ for large enough d. Finally, as $\log|G_n| > d^n$, Lemma 3.2 follows.

7 Expanding Schreier graphs

For a finite group H, a subgroup $H' < H$, and a (symmetric) set of elements U in it, the *Schreier graph* $\mathrm{Sch}(H, H', U)$ has vertex set H/H' , and edges (gH', ugH') for every $u \in U$, resulting in a $|U|$-regular graph. If $H' = \{1\}$ then $\mathrm{Sch}(H, \{1\}, U)$ is simply the Cayley graph $\mathrm{C}(H, U)$.

In this section we prove an analogue of Theorem 4.1 for Schreier graphs. In Theorem 4.1 we constructed a sequence of expanding Cayley graphs assuming the existence of a good “seed” Cayley graph. Here we do the same for Schreier graphs. The difference here is that the “seed” Schreier graph is known to exist by elementary arguments, and we do not rely on the strong theorem of [22]. By [17], *every* $2d$-regular graph is a Schreier graph, so a sequence of expanding Schreier graphs is implicit in any sequence of (even degree regular) expander graphs. However, it is generally hard to compute a Schreier graph representation of a given d-regular graph. In this section we explicitly provide the Schreier graph representation of our graphs.

There is another way to describe Schreier graphs. Let H be a group acting transitively on a set E. Define a graph $\mathrm{Sch}(H, E, U)$ whose vertices are E, and whose edges are (e, ue) for all $u \in U$ and $e \in E$. Pick a vertex $e_0 \in E$, and define $H' = \{h \in H \mid he_0 = e_0\}$, the stabilizer of e_0. The graph we defined on E is isomorphic to $\mathrm{Sch}(H, H', U)$ by taking he_0 to hH'. The following definition gives an example of groups acting on sets. This example will be the basis of a construction of expander Schreier graphs. To fix notation, we redefine our basic objects:

Definition 7.1. Let $T_{n,d}$ be the rooted d-regular tree with depth n, let $\mathrm{Sym}(n, d)$ be its group of symmetries, and let E_n be the set of leaves of $T_{n,d}$, on which $\mathrm{Sym}(n, d)$ acts naturally.

Expansion of a Cayley graph implies the expansion of all its Schreier graphs:

Claim 7.2. *Let H be a group , let $H' < H$ be a subgroup and let $U \subset G$ be a subset. Then $\lambda(Sch(H, H', U)) \leq \lambda(C(H, U))$.*

Proof. Let $v : H/H' \to \mathbb{C}$ be an eigenvector of the Schreier graph. Define $\hat{v} : H \to \mathbb{C}$ by $\hat{v}(h) = v(hH')$. Then $\hat{v}$ is an eigenvector of $\mathrm{C}(H,U)$ with the same eigenvalue as v. □

In Theorem 4.1 we constructed a sequence of Cayley graphs $\mathrm{C}(G_n,Y_n)$ where G_n is a subgroup of $\mathrm{Sym}(n,d)$, and showed that it is an expander family under some assumption on the symmetric group A_d, which is true for very large d. In light of Claim 7.2, the family $\mathrm{Sch}(G_n,E_n,Y_n)$ is also a sequence of expander graphs, under the same assumption. Below we construct expanding generating sets for G_n, which are both simpler than Y_n and do not require any assumptions (and work for much smaller d).

Reminder: for two groups G,K, such that $K < S_d$, the *wreath product* $G \wr K$ has elements $G^d \times K$ and multiplication rule

$$(g_1,\ldots,g_d,\sigma)\cdot(\tilde{g}_1,\ldots,\tilde{g}_d,\tau) = (g_{\tau(1)}\tilde{g}_1,\ldots,g_{\tau(d)}\tilde{g}_d,\sigma\tau) \quad .$$

Elements of G^d are naturally embedded in $G \wr K$ by setting the K coordinate to be 1. The group K is embedded in $G \wr K$ similarly by setting the G^d coordinates to be 1.

Definition 7.3. Given a group $K < S_d$, define a sequence of groups inductively by $K_1 = K$ and $K_{n+1} = K_n \wr K$ (The groups G_n of Theorem 4.1 are such groups with $K = A_d$). Recall that each element in $\mathrm{Sym}(n,d)$ is uniquely defined by writing a permutation in S_d on each internal node of $T_{n,d}$, indicating how the children of this vertex are permuted. The group K_n is the subgroup of $\mathrm{Sym}(n,d)$ where the permutation written on every internal vertex is an element of K. The group K_n acts on the set E_n (the leaves of the d-regular depth n tree) via its embedding in $\mathrm{Sym}(n,d)$.

The following theorem is the Schreier graph analogue of Theorem 4.1.

Theorem 7.4. *If there is a generating set $Q \subset K$ of size at most $(d^{1/4}/2)$ with $\lambda(K,[d],Q) \le 1/4$, then there exist $Q_n \subset K_n$ of size $|Q|^4$ such that $\lambda(K_n,E_n,Q_n) \le 1/4$, and Q_n can be computed in time polynomial in* $\log|E_n|$.

The main difference from Theorem 4.1 is that in the Schreier case the set Q is known to exist for many groups K. The claim below shows the existence of such Q for $K = S_d$ (for d large enough).

Claim 7.5. *Let $d \ge 100$, and let U be 100 permutations in S_d chosen randomly uniformly. Then $\lambda(S_d,[d],U) \le 1/4$ with probability larger than 1/2.*

For proofs see [13] (or [9] for a weaker result which would result in a larger required d).

Corollary 7.6. *For every $d \ge 4\cdot 100^4$ there is a sequence of subsets $U_n \subset \mathrm{Sym}(n,d)$ of size 100^4 such that $\lambda(\mathrm{Sym}(n,d),E_n,U_n) < 1/4$. Furthermore, U_n is computable in time polynomial in* $\log|E_n|$.

Proof of Theorem 7.4: We will assume that $|Q|^4$ divides d. The divisibility condition is not crucial, but it simplifies the proof. We proceed by induction - the case $n = 1$ is the assumption of the theorem. Assume the theorem holds for some n. We show that it holds for $n+1$.

Claim 7.7. *Let $Q_n^{(d)}$ be the vectors in Q_n^d in which every element in Q_n appears exactly $d/|Q_n|$ times (see Definition 4.4). Let $\underline{x} = (x_1,\ldots,x_d)$ be an element of $Q_n^{(d)}$. Define $U = \{y\underline{x}z \mid y,z \in Q\} \subset K_{n+1}$. Then $\lambda(K_{n+1},E_{n+1},U) \le \lambda(K_n,E_n,Q_n) + \lambda(K,[d],Q)$.*

We prove the claim later, and now proceed with the proof of the theorem. Define Q_{n+1} to be the set of words of length 2 in the set U given by Claim 7.7. Then

$$\lambda(K_{n+1},E_{n+1},Q_{n+1}) = \lambda(K_{n+1},E_{n+1},U)^2 \leq \left[\lambda(K_n,E_n,Q_n)+\lambda(K,[d],Q))\right]^2 \leq (1/4+1/4)^2 \leq 1/4$$

where the equality follows from Observation 2.3, the first inequality is Claim 7.7 and the second inequality is the induction assumption. By definition $|Q_{n+1}| = |Q|^4$. This concludes the proof of Theorem 7.4 □

Proof of Claim 7.7. The proof uses the zig-zag theorem [42]. Here is a quick definition of the zig-zag product:

Definition 7.8. Let $\mathcal{X},\mathcal{Y}$ be regular graphs such that the degree of $\mathcal{X}$ is equal the size of $\mathcal{Y}$. For every $v \in \mathcal{X}$ write the list of neighbors of v as an array $v[i]$ for $i \in \mathcal{Y}$ (the ordering of the list of neighbors is arbitrary, and different lists may lead to different graphs). Define a graph $\mathcal{Z}$ whose vertices are pairs (v,i), with $v \in \mathcal{X}$ and $i \in \mathcal{Y}$. The neighbors of a vertex (v,i) are the vertices reached by making the following three steps:

- Step 1: Walk from (v,i) to (v,j) where (i,j) is an edge of $\mathcal{Y}$.
- Step 2: Walk from (v,j) to $(v[j],j)$ where $v[j]$ is the j-th neighbor of v in the graph $\mathcal{X}$.
- Step 3: Walk from $(v[j],j)$ to $(v[j],k)$ where (j,k) is an edge of $\mathcal{Y}$.

$\mathcal{Z}$ has degree $(\deg \mathcal{Y})^2$. It is called the *zig-zag product* of $\mathcal{X}$ and $\mathcal{Y}$, and we write $\mathcal{Z} = \mathcal{X} ⓩ \mathcal{Y}$.

Theorem 7.9 ([42]). *If* $\mathcal{Z} = \mathcal{X} ⓩ \mathcal{Y}$ *then* $\lambda(\mathcal{Z}) \leq \lambda(\mathcal{X}) + \lambda(\mathcal{Y})$.

Define $\widetilde{Q_n}$ to be the multiset of size d obtained by duplicating every element of Q_n exactly $d/|Q_n|$ times. Notice that the vector $\underline{x}$ is simply a list of the elements of $\widetilde{Q_n}$. Let $\mathcal{X} = \mathrm{Sch}(K_n,E_n,\widetilde{Q_n})$, $\mathcal{Y} = \mathrm{Sch}(K,[d],Q)$, and $\mathcal{Z} = \mathrm{Sch}(K_{n+1},E_{n+1},U)$. We claim that $\mathcal{Z} = \mathcal{X} ⓩ \mathcal{Y}$. The proof of Claim 7.7 then follows from Theorem 7.9 (notice that $\lambda(\mathcal{X}) = \lambda(K_n,E_n,Q_n)$). The first requirement is that the degree of $\mathcal{X}$ is equal to the size of $\mathcal{Y}$, and indeed they are both d. It remains to verify that edges of $\mathcal{Z}$ are the walks of length 3 of the zig-zag product. For every $v \in \mathcal{X}$ and $i \in \mathcal{Y}$ define $v[i] = x_i(v)$ (the element $x_i \in K_n$ acts on $v \in E_n$). The array $v[i]$ is the list of neighbors of v in $\mathcal{X}$. An edge of $\mathcal{Z}$ connects (v,i) to $y\underline{x}z(v,i)$ (embedded in K_{n+1} as $y = (1,1,\dots,1,y)$, $z = (1,1,\dots,1,z)$ and $\underline{x} = (x_1,x_2,\dots,x_d,1)$). Let (v,i) be a vertex of $\mathcal{Z}$, and let $j = z(i)$, and $k = z(j) = yz(i)$. Then

$$y\underline{x}z(v,i) = y\underline{x}(v,z(i)) = y\underline{x}(v,j) = y(x_j(v),z(i)) = y(v[j],j) = (v[j],k) \quad .$$

This is exactly the definition of an edge in the zig-zag product, and we have proved Claim 7.7. □

8 Generators for an infinite group with property (τ)

There are natural mappings, for $k \le n$ between $\mathrm{Sym}(n,d)$ and $\mathrm{Sym}(k,d)$. The *embedding map* sends an element $\sigma \in \mathrm{Sym}(k,d)$ to the element of $\mathrm{Sym}(n,d)$ which acts on the first k levels of the tree by σ. The *restriction map* sends $\tau \in \mathrm{Sym}(n,d)$ to its restriction to the first k levels of the tree.

In Theorem 4.1 we constructed subsets $Y_n \subset G_n < \mathrm{Sym}(n,d)$ that generated G_n as expanders. In this section we will prove that the sets Y_n are consistent: the restriction of Y_n to $\mathrm{Sym}(k,d)$ is exactly Y_k. This implies that there is a set Y_∞ of symmetries of the infinite rooted d-regular tree, which restricts to Y_n for any n. A nice corollary is that the infinite group generated by Y_∞ has property τ (defined below) with respect to some sequence of subgroups. In the next section we will use this consistency to construct a sequence of expander graphs each of which is a lift of its predecessor.

Theorem 8.1. *Let $Y_n \subset G_n$ be the groups and generating (multi)sets of Theorem 4.1. Then for every $n \ge k \ge 2$ the restriction of Y_n to $\mathrm{Sym}(k,d)$ is equal Y_k. The same holds for the sets Q_n of Theorem 7.4.*

Corollary 8.2. *Define Y_∞ to be the set of symmetries of the infinite tree whose restriction to $\mathrm{Sym}(n,d)$ is Y_n. The set Y_∞ generates an infinite subgroup G_∞ of the symmetries of the infinite tree, and the restriction of G_∞ to $\mathrm{Sym}(n,d)$ is G_n for all $n \ge 2$. The same holds for Q_n of Theorem 7.4.*

The following definition of property (τ) is from [30], page 49.

Definition 8.3. Let G be a finitely generated group, and let Y be a finite symmetric generating set for G. Let $\mathcal{L} = \{N_n\}_{n\in\mathbb{N}}$ be a family of finite index normal subgroups in G. Then **G has property (τ) with respect to** $\mathcal{L}$ if the family $\mathrm{C}(G/N_n, Y/N_n)$ is an expander family.

Corollary 8.4. *Let N_n be the kernel of the restriction function from G_∞ to G_n. Then, under the assumption on the alternating group described in Theorem 4.1, the group G_∞ has property (τ) with respect to the family $\{N_n\}_{n=2}^{\infty}$.*

Proof of Theorem 8.1. The proof will only deal with the (harder) case of Y_n. Recall that elements in $\mathrm{Sym}(n,d)$ are represented by writing a permutation on each internal vertex of $T_{n,d}$. Define the *k-th level* of an element $u \in \mathrm{Sym}(n,d)$ to be the permutations written on the k-th level of $T_{n,d}$ in this representation of u.

The following claim is somewhat complicated to state, but its proof is an easy induction.

Claim 8.5. *Let $F_{i,j}$ be sequence of functions $F_{i,j} : \mathrm{Sym}(\infty,d)^q \to S_d$, where $1 \le i \le q$ and j is an internal vertex of $T_{\infty,d}$. Suppose that for vertices j in the k-th level of $T_{\infty,d}$, the output of $F_{i,j}$ only depends on levels 1 up to $(k-1)$ of its inputs (in particular $F_{i,1}$ is a constant function). Define $U_1 \subset \mathrm{Sym}(n,d)$ to be the set $F_{i,1}()$ for $1 \le i \le q$, and inductively, given the set $U_n = u_1^n, u_2^n, \dots, u_q^n$ in $\mathrm{Sym}(n,d)$ define $u_i^{n+1} \in \mathrm{Sym}(n+1,d)$ by writing the permutation $F_{i,j}(u_1^n, u_2^n, \dots, u_q^n)$ in internal vertex number j of $T_{n+1,d}$. Then the restriction of u_i^{n+1} to $\mathrm{Sym}(n,d)$ is u_i^n.* $\square$

Theorem 8.1 now follows by observing that the sets Y_n are indeed constructed by the procedure in Claim 8.5. (The only exception is Y_1 which was constructed differently, so the theorem's statement holds only for $n \ge 2$). To show this, recall briefly how we constructed Y_{n+1} given the set Y_n.

- Construct the set $Y_n{}^*$ defined in Definition 4.5.

- Write $X = c \cdot Y_n \cup Y_n{}^*$.
- Pick an element $\underline{x} \in X^{(d)} \subset \mathrm{Sym}(n,d)^d$, and embed it in $\mathrm{Sym}(n+1,d)$.
- Define $Z = Y_1 \underline{x} Y_1$ by regarding the elements of Y_1 as elements in $\mathrm{Sym}(n+1,d)$.
- Define Y_{n+1} to be the set of words of length 2 in the set Z.

We will now verify that the $(k+1)$-level an element in Y_{n+1} is a function of levels 1 up to k of the elements in Y_n. We will also verify that the first level of elements in Y_{n+1} is indeed a constant independent of n. We leave to the reader to verify that the conditions of Claim 8.5 hold precisely, which we feel is rather too technical.

Observation 8.6. Let g be an element of G_n. Let $g = [x,y]$ be the commutator representation derived in Section 6. Then level k of x and y depends only on levels 1 up to k of g.

The observation follows by following the construction of the commutator representation, which is simply induction on the level. We conclude that for elements in $Y_n{}^*$, and therefore in X, the k-th level depends only on levels 1 up to k of the elements in Y_n.

Observation 8.7. Let g,h be elements in $\mathrm{Sym}(\infty,d)$. Then level k of gh depends only on levels 1 up to k of g and h.

Observation 8.8. Let $\underline{x} = (x_1,\ldots,x_d)$ be an element of $\mathrm{Sym}(n,d)^d$. Embed $\underline{x}$ in $\mathrm{Sym}(n+1,d)$ as $(x_1,\ldots,x_d,1)$, represented by writing a permutation on every internal vertex of $T_{n+1,d}$. The permutation written on the root is the identity, and the permutations written on level $k+1$ are permutations written on level k of $x_1,\ldots,x_d$.

The two observations above imply that level $k+1$ of elements in Z depend only on levels 1 up to k of X. Also, level 1 of elements in Z is independent of n, since it depends on level 1 of elements in Y_1 and level 1 of $\underline{x}$ which is the identity permutation. The same holds for Y_{n+1} as elements there are products of elements of Z (we use Observation 8.7 again). This concludes the proof of the theorem. □

9 A sequence of expanding lifts of graphs

Definition 9.1. Given a graph $\mathcal{X}$, on n vertices $v_1,\ldots,v_n$, a *d-lift* of $\mathcal{X}$ is a graph $\mathcal{Y}$ on nd vertices $w_{i,k}$ where $i \in [n], k \in [d]$. For each edge $e = (v_i,v_j)$ of $\mathcal{X}$ choose a permutation $\sigma_e \in S_d$, and connect $w_{i,k}$ with $w_{j,\sigma_e(k)}$ for all $k \in [d]$. The vertices $w_{i,k}$ for fixed i and $k \in [d]$ are called the *fiber* above v_i. The fibers above an edge $e = (v_i,v_j)$ are connected by a perfect matching defined by σ_e. There are many non-isomorphic lifts of a graph $\mathcal{X}$ depending on the choice of the permutations σ_e. For more information on lifts see [6].

In this section we show how to obtain an explicit sequence of expander graphs, each of which is a d-lift of its predecessor for any (large enough) d. Actually, the sequence of Schreier graphs constructed in the previous sections do.

Here are some basic properties of lifts which are not hard to prove:

- The degree of a vertex v is equal to the degree of all the vertices in the fiber above v, so a lift of a regular graph is regular with the same degree.
- The definition of a lift works fine for parallel edges and loops (where the loop counts as two edges when computing the degree of a vertex).
- Lifting is transitive: If $\mathcal{Y}$ is a lift of $\mathcal{X}$ and $\mathcal{Z}$ is a lift of $\mathcal{Y}$ then $\mathcal{Z}$ is a lift of $\mathcal{X}$.
- If $\mathcal{Y}$ is a lift of $\mathcal{X}$ then $\lambda(\mathcal{Y}) \geq \lambda(\mathcal{X})$.

As an example, consider the graph $\mathcal{X}_0$ which consists of a single vertex with q loops on it. A lift $\mathcal{X}_1$ of $\mathcal{X}_0$ is encoded by q permutations $\sigma_1, \ldots, \sigma_q \in S_d$. The graph $\mathcal{X}_1$ has vertex set $[d]$ and edges $(i, \sigma_l(i)), i \in [d], l \in [q]$, making it a $2q$-regular graph.

Linial raised the following conjecture:

Conjecture 9.2 (Linial). For every graph $\mathcal{X}$ and every d there exists a d-lift $\mathcal{Y}$ of $\mathcal{X}$ such that $\lambda(\mathcal{Y}) \leq \max(\lambda(\mathcal{X}), O(\sqrt{d}))$.

For $d = 2$ a slightly weaker version of the conjecture was proved in [8].

The conjecture yields a method to construct a sequence of expander graphs each of which is a lift of its predecessor. Pick any regular graph $\mathcal{X}_1$ with $\lambda(\mathcal{X}_1) = 1/2$. Now choose a sequence of graphs $\mathcal{X}_n$ such that $\lambda(\mathcal{X}_{n+1}) \leq \lambda(\mathcal{X}_n)$ and $\mathcal{X}_{n+1}$ is a lift of $\mathcal{X}_n$ (we need the degree of the initial graph to be large enough for this to work).

Theorem 9.3. *Let $\mathcal{X}_n = Sch(K_n, E_n, Q_n)$ be the family of graphs of Theorem 7.4. Then $\mathcal{X}_{n+1}$ is a d-lift of $\mathcal{X}_n$ for all $n \geq 1$.*

By Corollary 7.6 we obtain the required sequence of expanding lifts:

Corollary 9.4. *Let $K = S_d$ with $d \geq 4 \cdot 100^4$, and let $Q \subset K_n$ be the generating set given in 7.6. Let $\mathcal{X}_n$ be the sequence constructed in Theorem 9.3. Then $\lambda(\mathcal{X}_n) \leq 1/4$ for all n and $\mathcal{X}_{n+1}$ is a d-lift of $\mathcal{X}_n$ for all $n \geq 1$.*

The proof of Theorem 9.3 is by induction. The following two claims show how to construct a Schreier graph of a wreath product $G \wr H$ which is naturally a lift of a Schreier graph of H. These two claims will be used in the induction step.

Claim 9.5. *Let G, H be groups acting on E_G, E_H respectively. H is a subgroup of the symmetric group on E_H, so the group $G \wr H$ is defined, and its elements are written as $(\underline{g}, h)$ where $\underline{g} = (g_y)_{y \in E_H}$ and $h \in H$. Then $G \wr H$ acts on $E_G \times E_H$ by $(\underline{g}, h)(x, y) = (g_y(x), h(y))$.*

Proof. We need to show that for two elements $(\underline{g}, h), (\underline{\tilde{g}}, \tilde{h}) \in G \wr H$ and an element $(x, y) \in E_G \times E_H$

$$(\underline{g}, h)\big[(\underline{\tilde{g}}, \tilde{h})(x, y))\big] = \big[(\underline{g}, h) \cdot (\underline{\tilde{g}}, \tilde{h})\big](x, y) \quad .$$

And indeed,

$$(\underline{g}, h)\big[(\underline{\tilde{g}}, \tilde{h})(x, y))\big] = (\underline{g}, h)(\tilde{g}_y(x), \tilde{h}(y)) = (g_{\tilde{h}(y)} \cdot \tilde{g}_y, h \cdot \tilde{h})(x, y) = \big[(\underline{g}, h) \cdot (\underline{\tilde{g}}, \tilde{h})\big](x, y) \quad .$$

□

Claim 9.6. *Let G, H be as in Claim 9.5, and let U be a subset of $G \wr H$. Then $Sch(G \wr H, E_G \times E_H, U)$ is a $|E_G|$-lift of $Sch(H, E_H, U)$. (Notice that we have identified U with its restriction to H).*

Proof. The vertices of $\mathrm{Sch}(G \wr H, E_G \times E_H, U)$ are pairs (x, y) with $x \in E_G$ and $y \in E_H$. Partition these vertices to subsets $S_y = \{(x, y) \mid x \in E_G\}$. We will show that $\mathrm{Sch}(G \wr H, E_G \times E_H, U)$ is a $|E_G|$-lift of $\mathrm{Sch}(H, E_H, U)$ where the fiber above $y \in E_H$ is S_y. In order to prove this, we need to show that for every edge $e = (y_1, y_2)$ of $\mathrm{Sch}(H, E_H, U)$ there corresponds a perfect matching between S_{y_1} and S_{y_2}.

Edges in $\mathrm{Sch}(H, E_H, U)$ are of the form (y, uy), for $y \in E_H$ and $u \in U$. Write $u = (\underline{g}, h)$ in $G \wr H$, so $uy = h(y)$. In $\mathrm{Sch}(G \wr H, E_G \times E_H, U)$, a vertex (x, y) is connected to $u(x, y) = (g_y(x), h(y))$. This is a perfect matching between S_y and $S_{h(y)}$ since g_y is a permutation of E_G for y fixed. □

Can we use Claim 9.6 to obtain a sequence of expanding lifts? In Section 8 we constructed an expander sequence $\mathcal{X}_n = \mathrm{Sch}(K_n, Q_n, E_n)$ where each Q_n is the restriction of a single set Q_∞. Since $K_{n+1} = K_n \wr K$ we deduce by Claim 9.6 that $\mathrm{Sch}(K_{n+1}, Q_\infty, E_{n+1})$ is a lift of $\mathrm{Sch}(K, Q_\infty, [d]) = \mathcal{X}_1$, while we wanted $\mathcal{X}_{n+1}$ to be a lift of $\mathcal{X}_n$. The following observation comes to the rescue (notice the change of order in the wreath product).

Observation 9.7. Let K_n be the sequence of groups defined in 7.3. Consider K_n as a subset of the permutation group on E_n. Then $K_{n+1} = K \wr K_n$.

We can now use Claim 9.6 to conclude that $\mathrm{Sch}(K_{n+1}, Q_\infty)$ is a d-lift of $\mathrm{Sch}(K_n, Q_\infty)$, which proves Theorem 9.3. □

Acknowledgments

We are grateful to Alex Lubotzky for many insightful conversations, in part supplying the group theoretic tools we ended up needing for the proof. We thank Yair Glasner and Shachar Mozes for useful remarks. We thank the anonymous referee for many helpful comments.

References

[1] * MIKLÓS AJTAI, JÁNOS KOMLÓS, AND ENDRE SZEMERÉDI: Sorting in $c \log n$ parallel steps. *Combinatorica*, 3(1):1–19, 1983. 1.1

[2] * NOGA ALON: Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. 1.1

[3] * NOGA ALON, ZVI GALIL, AND VITALI D. MILMAN: Better expanders and superconcentrators. *Journal of Algorithms*, 8(3):337–347, 1987. [JAlg:10.1016/0196-6774(87)90014-9]. 1.1

[4] * NOGA ALON, ALEXANDER LUBOTZKY, AND AVI WIGDERSON: Semi-direct product in groups and zig-zag product in graphs: connections and applications (extended abstract). In *Proc. of the 42nd Annual Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)*, pp. 630–637.

IEEE Computer Soc., Los Alamitos, CA, 2001. [FOCS:10.1109/SFCS.2001.959939]. 1.2, 1.3, 2.2.2, 2.12

[5] * NOGA ALON AND VITALI D. MILMAN: λ_1, isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory. Series B*, 38(1):73–88, 1985. [JCombThB:10.1016/0095-8956(85)90092-9]. 1.1

[6] * ALON AMIT AND NATHAN LINIAL: Random graph coverings. I. General theory and graph connectivity. *Combinatorica*, 22(1):1–18, 2002. [Combinatorica:er6qljcyq3v8pyd2]. 9.1

[7] * ELI BEN-SASSON, MADHU SUDAN, SALIL VADHAN, AND AVI WIGDERSON: Randomness-efficient low degree tests and short PCPs via epsilon-biased sets. In *Proc. of the 35th Annual ACM Symposium on Theory of Computing*, pp. 612–621. ACM Press, 2003. [STOC:780542.780631]. 1.2

[8] * YONATAN BILU AND NATI LINIAL: Constructing expander graphs by 2-lifts and discrepancy vs. spectral gap. In *Proc. of the 45th Annual Symposium on Foundations of Computer Science*, pp. 404–412, Rome, Italy, 17-19 October 2004. IEEE. [FOCS:10.1109/FOCS.2004.19]. 9

[9] * ANDREI BRODER AND ELI SHAMIR: On the second eigenvalue of random regular graphs (preliminary version). In *Proc. of the 28th Annual Symposium on Foundations of Computer Science*, pp. 286–294, Los Angeles, California, 12–14 October 1987. IEEE. 7

[10] * MICHAEL CAPALBO, OMER REINGOLD, SALIL VADHAN, AND AVI WIGDERSON: Randomness conductors and constant-degree lossless expanders. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing*, pp. 659–668. ACM Press, 2002. [STOC:509907.510003]. 1.1

[11] * PERSI DIACONIS AND MEHRDAD SHAHSHAHANI: On the eigenvalues of random matrices. *J. Appl. Probab.*, 31A:49–62, 1994. 1.2

[12] * MARTIN EICHLER: Quaternäre quadratische Formen und die Riemannsche Vermutung für die Kongruenzzetafunktion. *Arch. Math.*, 5:355–366, 1954. 1.2

[13] * JOEL FRIEDMAN: A proof of Alon's second eigenvalue conjecture. *Memoirs of the AMS*, to appear. [arXiv:cs.DM/0405020]. 7

[14] * OFER GABBER AND ZVI GALIL: Explicit constructions of linear-sized superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, June 1981. [JCSS:10.1016/0022-0000(81)90040-4]. 1.1

[15] * ODED GOLDREICH, RUSSELL IMPAGLIAZZO, LEONID LEVIN, RAMARATHNAM VENKATESAN, AND DAVID ZUCKERMAN: Security preserving amplification of hardness. In *Proc. of the 31st Annual Symposium on Foundations of Computer Science*, volume I, pp. 318–326, St. Louis, Missouri, 22–24 October 1990. IEEE. 1.1

[16] * MISHA GROMOV: Spaces and questions. *Geometric and Functional Analysis*, pp. 118–161, 2000. Part I of Special Volume on GAFA 2000 (Tel Aviv, 1999). 1.1

[17] * JONATHAN L. GROSS: Every connected regular graph of even degree is a Schreier coset graph. *Journal of Combinatorial Theory. Series B*, 22(3):227–232, 1977. [JCombThB:10.1016/0095-8956(77)90068-5]. 1.2, 7

[18] * RUSSELL IMPAGLIAZZO, NOAM NISAN, AND AVI WIGDERSON: Pseudorandomness for network algorithms. In *Proc. of the 26th Annual ACM Symposium on the Theory of Computing*, pp. 356–364, Montréal, Québec, Canada, 23–25 May 1994. [STOC:195058.195190]. 1.1

[19] * RUSSELL IMPAGLIAZZO AND AVI WIGDERSON: $P = BPP$ if E requires exponential circuits: Derandomizing the XOR lemma. In *Proc. of the 29th Annual ACM Symposium on Theory of Computing*, pp. 220–229, El Paso, Texas, 4–6 May 1997. [STOC:258533.258590]. 1.1

[20] * SHUJI JIMBO AND AKIRA MARUOKA: Expanders obtained from affine transformations. *Combinatorica*, 7(4):343–355, 1987. 1.1, 1.2

[21] * NIGEL J. KALTON AND JAMES W. ROBERTS: Uniformly exhaustive submeasures and nearly additive set functions. *Transactions of the American Mathematical Society*, 278(2):803–816, 1983. 1.1

[22] * MARTIN KASSABOV: Symmetric groups and expander graphs. arxiv:math.GR/0505624. [arXiv:math.GR/0505624]. 1.1, 1.2, 4.2, 7

[23] * MARTIN KASSABOV: Universal lattices and unbounded rank expanders. arxiv:math.GR/0502237. [arXiv:math.GR/0502237]. 1.2

[24] * MARTIN KASSABOV AND NIKOLAY NIKOLOV: Universal lattices and property τ. arxiv:math.GR/0502112. [arXiv:math.GR/0502112]. 1.2

[25] * DAVID KAZHDAN: On the connection of the dual space of a group with the structure of its closed subgroups (Russian). *Funkcional. Anal. i Prilozh.*, 1:71–74, 1967. 1.2

[26] * MARIA KLAWE: Limitations on explicit constructions of expanding graphs. *SIAM J. Comput.*, 13(1):156–166, 1984. [SICOMP:13/0213011]. 1.2

[27] * LÁSZLÓ LOVÁSZ AND PETER WINKLER: Mixing times. In *Microsurveys in discrete probability (Princeton, NJ, 1997)*, volume 41 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pp. 85–133. Amer. Math. Soc., Providence, RI, 1998. 1.2

[28] * A. LUBOTZKY AND B. WEISS: Groups and expanders. In *Expanding graphs (Princeton, NJ, 1992)*, volume 10 of *DIMACS Ser. Discrete Math. Theoret. Comput. Sci.*, pp. 95–109. Amer. Math. Soc., Providence, RI, 1993. 1.2

[29] * ALEX LUBOTZKY, RALPH PHILLIPS, AND PETER SARNAK: Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. [Combinatorica:k285687344657q53]. 1.1

[30] * ALEXANDER LUBOTZKY: *Discrete groups, expanding graphs and invariant measures*. Birkhäuser Verlag, Basel, 1994. 1.1, 8

[31] * ALEXANDER LUBOTZKY AND IGOR PAK: The product replacement algorithm and Kazhdan's property (T). *Journal of the American Mathematical Society*, 14(2):347–363 (electronic), 2001. [JAMS:2001-14-02/S0894-0347-00-00356-8]. 1.1

[32] * GREGORY A. MARGULIS: Explicit constructions of expanders. *Problemy Peredači Informacii*, 9(4):71–80, 1973. 1.1

[33] * GREGORY A. MARGULIS: Explicit group-theoretic constructions of combinatorial schemes and their applications in the construction of expanders and concentrators. *Problemy Peredachi Informatsii*, 24(1):51–60, 1988. 1.1

[34] * ROY MESHULAM AND AVI WIGDERSON: Expanders from symmetric codes. In *Proc. of the 34th Annual ACM Symposium on Theory of Computing*, pp. 669–677. ACM Press, 2002. [STOC:509907.510004]. 1.2

[35] * MOSHE 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. [JCombThB:10.1006/jctb.1994.1054]. 1.1

[36] * JOSEPH NAOR AND MONI NAOR: Small-bias probability spaces: Efficient constructions and applications. *SIAM Journal on Computing*, 22(4):838–856, August 1993. [SICOMP:22/0222053]. 1.1

[37] * NIKOLAY NIKOLOV: On the commutator width of perfect groups. *Bull. London Math. Soc.*, 36(1):30–36, 2004. [BLMS:10.1112/S0024609303002601]. 1.3, 2.2.3, 2.13, 6, 6.5, 6

[38] * OYSTEIN ORE: Some remarks on commutators. *Proc. Amer. Math. Soc.*, 2:307–314, 1951. 3.1

[39] * MARK S. PINSKER: On the complexity of a concentrator. In *Proc. of the 7th Annual Teletraffic Conference*, pp. 318/1–318/4, Stockholm, 1973. 1.1

[40] * NICHOLAS PIPPENGER: Sorting and selecting in rounds. *SIAM Journal on Computing*, 16(6):1032–1038, December 1987. [SICOMP:16/0216066]. 1.1

[41] * NICHOLAS PIPPENGER AND ANDREW C. YAO: Rearrangeable networks with limited depth. *SIAM Journal on Algebraic and Discrete Methods*, 3(4):411–417, 1982. [SIMAX:03/0603041]. 1.1

[42] * 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. [arXiv:math.CO/0406038]. 1.1, 1.3, 2.2.2, 7, 7.9

[43] * ATLE SELBERG: On the estimation of Fourier coefficients of modular forms. In *Proc. of the Sympos. Pure Math.*, volume VIII, pp. 1–15. Amer. Math. Soc., Providence, R.I., 1965. 1.2

[44] * MICHAEL SIPSER: Expanders, randomness, or time versus space. *Journal of Computer and System Sciences*, 36(3):379–383, June 1988. [JCSS:10.1016/0022-0000(88)90035-9]. 1.1

[45] * MICHAEL SIPSER AND DANIEL A. SPIELMAN: Expander codes. *IEEE Transactions on Information Theory*, 42(6, part 1):1710–1722, 1996. [TIT:10.1109/18.556667]. 1.1

[46] * DANIEL A. SPIELMAN: Linear-time encodable and decodable error-correcting codes. *IEEE Transactions on Information Theory*, 42(6, part 1):1723–1731, 1996. [TIT:10.1109/18.556668]. 1.1

[47] * MICHAEL R. TANNER: Explicit concentrators from generalized n-gons. *SIAM Journal on Algebraic Discrete Methods*, 5(3):287–293, 1984. [SIMAX:05/0605030]. 1.1

[48] * ALASDAIR URQUHART: Hard examples for resolution. *Journal of the Association for Computing Machinery*, 34(1):209–219, 1987. [JACM:7531.8928]. 1.1

[49] * LESLIE G. VALIANT: Graph-theoretic arguments in low-level complexity. In *Proc. of the 6th Symposium on Mathematical Foundations of Computer Science*, volume 53 of *Lecture Notes in Comput. Sci.*, pp. 162–176. Springer, Berlin, 1977. 1.1

[50] * AVI WIGDERSON AND DAVID ZUCKERMAN: Expanders that beat the eigenvalue bound: explicit construction and applications. *Combinatorica*, 19(1):125–138, 1999. [Combinatorica:wcjlnyjmdxf30b9x]. 1.1

AUTHORS

Eyal Rozenman
Postdoc
School of Mathematics
Institute for Advanced Study, Princeton
eyal@ias.edu
http://www.math.ias.edu/~eyal

Aner Shalev
Professor
Einstein Institute of Mathematics
The Hebrew University of Jerusalem
shalev@math.huji.ac.il
http://www.ma.huji.ac.il

Avi Wigderson
Professor
School of Mathematics
Institute for Advanced Study, Princeton
avi@ias.edu
http://www.math.ias.edu/~avi

ABOUT THE AUTHORS

EYAL ROZENMAN is currently a member in the Institute for Advanced Study, Princeton. He grew up near Tel Aviv, Israel, went to Tel Aviv University for his B. Sc. degree, and received his Ph. D. from the Hebrew University under the guidance of Nathan Linial. Eyal's research interests include groups, expanders and their interaction, and pseudo-randomness in general. He loves to read and write papers which are elementary, but he likes them better if they contain fancy mathematical notions. Eyal's non-research interests include hiking, reading, and avoiding driving.

ANER SHALEV was born in the previous century in a kibbutz on the Sea of Galilee. His first love (at 10) was experimental chemistry, which proved too dangerous when a test tube full of hydrogen exploded in his hands. He wisely moved to pure mathematics and writing fiction.

In 1988 he finished his Ph. D. in mathematics at the Hebrew University of Jerusalem, supervised by Shimshon Amitsur and Avinoam Mann, and published his first book, *Opus 1* - a collection of stories with a musical structure. Other literary works followed: *Overtures* (1996) - seventy openings of stories without ends, and the recent novel *Dark Matter* (2004), a complex love story with astrophysical parallels, alternating between prose and email, which will also appear in some European languages.

The mathematical work of Aner Shalev is in algebra. He enjoys studying groups using other disciplines, such as Lie methods and probabilistic methods. He was an invited ICM speaker (1998), and chaired the Institute of Mathematics at the Hebrew University (1999-2001). He used to enjoy playing chess with his daughter Ariella until he started losing.

AVI WIGDERSON was born in Haifa, Israel in 1956, and received his Ph. D. in 1983 at Princeton University under Dick Lipton. He enjoys and is fascinated with studying the power and limits of efficient computation, and the remarkable impacts of this field on understanding our world. Avi's other major source of fascination and joy are his three kids, Eyal, Einat, and Yuval.