Quantum Expanders: Motivation and Constructions

Avraham Ben-Aroya [*] Oded Schwartz [†] Amnon Ta-Shma[‡]

Abstract

We define quantum expanders in a natural way. We give two constructions of quantum expanders, both based on classical expander constructions. The first construction is algebraic, and is based on the construction of Cayley Ramanujan graphs over the group $PGL(2, q)$ given by Lubotzky, Philips and Sarnak [27]. The second construction is combinatorial, and is based on a quantum variant of the Zig-Zag product introduced by Reingold, Vadhan and Wigderson [35]. Both constructions are of constant degree, and the second one is explicit.

Using quantum expanders, we characterize the complexity of comparing and estimating quantum entropies. Specifically, we consider the following task: given two mixed states, each given by a quantum circuit generating it, decide which mixed state has more entropy. We show that this problem is QSZK–complete (where QSZK is the class of languages having a zero-knowledge quantum interactive protocol). This problem is very well motivated from a physical point of view. Our proof resembles the classical proof that the entropy difference problem is SZK–complete, but crucially depends on the use of quantum expanders.

1 Introduction

Expander graphs are graphs of low-degree and high-connectivity. There are several ways to measure the quality of expansion in a graph. One such way measures *set expansion*: given a not too large subset of the vertices S, it measures the size of the set $\Gamma(S)$ of neighbors of S, relative to $|S|$. Another way is *(Rényi) entropic expansion*: given a distribution π on the vertices of the graph, it measures the amount of (Rényi) entropy added in $\pi' = G\pi$. This is closely related to measuring the *algebraic expansion* given by the spectral gap of the adjacency matrix of the graph. See [20] for an excellent survey of the subject.

Pinsker [33] was the first to observe that *non-explicitly*, constant degree random graphs have almost-optimal set expansion. The algebraic measure of expansion naturally led to a series of explicit constructions based on algebraic structures, e.g. [28, 13, 21]. This line of research culminated by the works of Lubotzky, Philips and Sarnak [27], Margulis [29] and Morgenstern [30] who explicitly constructed Ramanujan graphs, i.e., D–regular graphs achieving spectral gap of $1 - 2\frac{\sqrt{D-1}}{D}$. Alon and Boppana (see [32]) showed that Ramanujan graphs achieve almost the best possible algebraic expansion, and Friedman [12] showed that random graphs are almost Ramanujan. Several works [11, 2, 1, 22] showed intimate connections between set expansion and algebraic expansion. We refer the reader, again, to the excellent survey paper [20].

The algebraic definition identifies the graph $G = (V, E)$ with its normalized adjacency matrix A. I.e., one defines a Hilbert space $\mathcal{V}$ of dimension $|V|$, and identifies an element $v \in V$ with a basis vector $|v\rangle \in \mathcal{V}$. A distribution π on V is identified with the vector $|\pi\rangle = \sum_{v \in V} \pi(v) |v\rangle$. G is viewed as a linear operator acting on $\mathcal{V}$, with the action of the normalized adjacency matrix $A : \mathcal{V} \to \mathcal{V}$. A maps probability distributions to probability distributions. Furthermore, this mapping corresponds to taking a random walk on G. Specifically, say one takes a random walk on G starting at time 0 with the distribution π_0 on V. Then, the distribution on the vertices at time k is $A^k |\pi_0\rangle$. Viewing G as a linear operator on the Hilbert space allows one to consider the action of A on arbitrary vectors in $\mathcal{V}$, not necessarily corresponding to probability distributions over V. Nevertheless, they are crucial for understanding the action of A. This is due to the fact that all of A's eigenvectors are such vectors, except for the stationary distribution.

We extend the algebraic definition to the quantum setting. We define *quantum expanders* as *feasible quantum*

[*] Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by the Adams Fellowship Program of the Israel Academy of Sciences and Humanities, by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848 and by USA Israel BSF grant 2004390. Email: abrhambe@post.tau.ac.il.

[†] Institut für Mathematik, MA 4-5, Technische Universität Berlin, 10623 Berlin, Germany. Email: odedsc@math.tu-berlin.de.

[‡] Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by the European Commission under the Integrated Project QAP funded by the IST directorate as Contract Number 015848, by Israel Science Foundation grant 217/05 and by USA Israel BSF grant 2004390. Email: amnon@tau.ac.il.

transformations having a *large spectral gap* and *small degree*. We first explain what a quantum feasible transformation is. Good places to read about the subject are the books [31, 24]. Here, we quickly repeat the essentials we need.

A *feasible classical transformation* is any linear operator that can be implemented by a classical circuit. A feasible classical transformation has the property that it maps probability distributions to probability distributions. A *feasible quantum transformation* is any transformation that can be implemented by a quantum circuit (with unitary operators and measurements). As it turns out, this definition implies that a general quantum state is a *density matrix* (which we explain soon), and a feasible quantum transformation corresponds to a linear operator mapping density matrices to density matrices.

A general classical state is a classical probability distribution over the standard basis $\{|v\rangle\}$ of $\mathcal{V}$, i.e., vectors of the form $\sum_v p_v\,|v\rangle$ as above. A general quantum state is a density matrix $\rho = \sum p_v\,|\psi_v\rangle\langle\psi_v|$, with $0 \leq p_v \leq 1$, $\sum p_v = 1$ and $\{\psi_v\}$ being *some* orthonormal basis of $\mathcal{V}$. Notice that a general classical state resides in the Hilbert space $\mathcal{V}$ of dimension $|V|$, whereas a general quantum state lives in the Hilbert space $L(\mathcal{V}) = \mathrm{Hom}(\mathcal{V},\mathcal{V})$ of linear operators from $\mathcal{V}$ to $\mathcal{V}$, and this Hilbert space has dimension $|V|^2$. A transformation $E : L(\mathcal{V}) \to L(\mathcal{V})$ that can be implemented by a quantum circuit is called an *admissible superoperator*. We define:

Definition 1.1. *An admissible superoperator $E : L(V) \to L(V)$ is $\overline{\lambda}$–expanding if:*

- *$E(\tilde{I}) = \tilde{I}$ and the eigenspace of eigenvalue 1 has dimension 1, where $\tilde{I} = \frac{1}{|V|}\sum_{v\in V}|v\rangle\langle v|$.*
- *For any $A \in L(V)$ that is orthogonal to $\tilde{I}$ (with respect to the Hilbert-Schmidt inner product, i.e. $\mathrm{Tr}(A\tilde{I}) = 0$) it holds that $\|\, E(A)\,\|_2 \leq \overline{\lambda}\,\|\, A\,\|_2$.*

A quantum expander is explicit *if E can be implemented by a polynomial size quantum circuit (i.e. polynomial in $\log(\dim(V))$).*

A few comments are in place. First, we remark that the linear transformation E is not necessarily normal. This already happens in the classical world, when we consider *directed* expanders. In such a case, one usually requires that the graph is *regular*, implying that the largest *singular value* is 1, and furthermore this singular value is obtained with the normalized all-ones vector (that corresponds to the uniform distribution). The two conditions we impose also imply that the largest singular value of E is 1, that this singular value is obtained with the completely mixed state eigenvector $\tilde{I}$, and that all other singular values are bounded by $\overline{\lambda}$.

A crucial property of classical expanders is that they achieve a large spectral gap using only a small degree. The notion of degree is natural when considering graphs, but may seem unnatural for the operators and superoperators algebraic entities. However, one way to look at small degree expanders, is that the operator never adds much entropy to the state it operates on, whereas it always adds some entropy to states that are far away from uniform. Such a view is almost explicit in the work of Capalbo et. al. [8], where they view expanders as entropy conductors.

Definition 1.2. *We say an admissible superoperator $E : L(\mathcal{V}) \to L(\mathcal{V})$ is D–*regular *if $E = \frac{1}{D}\sum_{d=1}^{D} E_d$, and for each $d \in [D]$, $E_d(X) = U_d X U_d^{\dagger}$ for some unitary transformation U_d over $\mathcal{V}$.*

This definition generalizes the classical one. Any D–regular graph can be thought of as a sum of D permutations, and each permutation corresponds to a unitary transformation (and in fact many classical constructions explicitly use this property, e.g. [35, 8]). However, the definition is intuitive in a more basic sense. Unitary transformations (or classically, permutations) are those transformations that do not change the entropy of a state. A transformation that is a linear combination of D such objects, can add at most $\log(D)$ entropy to any state it acts upon.

Definition 1.3. *An admissible superoperator $E : L(V) \to L(V)$ is a $(D, \overline{\lambda})$ quantum expander if E is D–regular and $\overline{\lambda}$–expanding.*

With this definition D–regular quantum expanders can never add more than $\log(D)$ entropy to the state they act on, but always add entropy to states that are far away from the completely-mixed state. This definition can be generalized, but for simplicity we work with Definition 1.3. A similar definition was independently given by Hastings [18].

1.1 Are there any non-trivial quantum expanders?

This is indeed a good question, and a major goal of this paper. A first natural attempt is to directly convert a good classical Cayley expander, to a quantum superoperator. This indeed can be done, and the resulting super operator $T : L(\mathcal{V}) \to L(\mathcal{V})$ is analyzed in Section 3.1. The analysis there shows that T has $|V|$ eigenspaces, each of dimension $|V|$, with eigenvalues $\overrightarrow{\lambda} = (\lambda_1 = 1, \ldots, \lambda_{|V|})$, where $\overrightarrow{\lambda}$ is the spectrum of the Cayley graph. In particular, the eigenspace of eigenvalue 1 has dimension $|V|$ instead of dimension 1.

Nevertheless, Ambainis and Smith [4] obtained the following quantum expander that is implicit in their work:

Theorem 1.4. *([4]) There exists an explicit $(\frac{\log^2 N}{\overline{\lambda}^2}, \overline{\lambda})$ quantum expander $E : L(V) \to L(V)$, where $N = dim(V)$.*

Their quantum expander is based on a Cayley expander over the Abelian group $\mathbb{Z}_2^n$. As explained above, taking the quantum analogue of the classical expander is not enough, and Ambainis and Smith obtain their result using a clever trick, essentially working over $\mathbb{F}_4^n$ rather than $\mathbb{Z}_2^n$.

The main problem with Abelian groups is that it is impossible to get constant degree Cayley expanders over them [25, 3]. This is reflected in the $\log^2 N$ term in Theorem 1.4. There are constant degree, Ramanujan Cayley graphs, i.e., Cayley graphs that achieve the best possible relationship between the degree and the spectral gap, but they are built over non-Abelian groups. If one wants to get constant degree quantum expanders, then he is forced to work over non-Abelian groups.

Our first construction starts with the constant degree Ramanujan expander of [27]. This expander is a Cayley graph over the non-Abelian group $\mathrm{PGL}(2, q)$. We build from it a quantum expander as follows: we take two steps on the classical expander graph, with a basis change between the two steps. The basis change is a carefully chosen refinement of the Fourier transformation that maps the standard basis $|g\rangle$ to the basis of the irreducible, invariant subspaces of $\mathrm{PGL}(2, q)$. Intuitively, in the Abelian case this basis change corresponds to dealing with both the bit and the phase levels, and is similar to the construction of quantum error correcting codes by first applying a classical code in the standard basis and then in the Fourier basis. However, this intuition is not as clear in the non-Abelian case. Furthermore, in the non-Abelian case not every Fourier transform is good. In this work we single out a natural algebraic property we need from the underlying group that is sufficient for proving the spectral gap of the construction. We discuss this in detail in Section 3. We then prove that $\mathrm{PGL}(2, q)$ respects this property. With that we get a $(D = O(\frac{1}{\bar{\lambda}^4}), \bar{\lambda})$ quantum expander.

This construction is not explicit in the sense that it uses the Fourier transform over PGL(2,q), which is not known to have an efficient implementation (see [26] for a non-trivial, but still not fast enough, algorithm). We mention that there are also explicit, constant degree (non-Ramanujan) Cayley expanders over $\mathcal{S}_n$ and $\mathcal{A}_n$ [23]. Also, there is an efficient implementation of the Fourier transform over $\mathcal{S}_n$ [5]. We do not know, however, whether $\mathcal{S}_n$ (or $\mathcal{A}_n$) respect our additional property.

Following our first construction (given in [7]), Hastings [19] showed that quantum expanders cannot be better than Ramanujan, and that non-explicitly, taking D random unitaries gives an almost-Ramanujan graph. However, a random unitary is a highly non-explicit object. It was therefore very natural to ask for an explicit construction of good quantum expanders.

The second construction we present in this paper is explicit, and gives constant degree, constant gap quantum expanders. It works by adapting the classical Zig-Zag construction [35] to the quantum world. We describe it in detail in Section 4.

1.2 What are quantum expanders good for?

Watrous [39] defined the class of quantum statistical zero knowledge languages (QSZK). QSZK is the class of all languages that have a quantum interactive proof system, along with an efficient simulator. The simulator produces transcripts that, for inputs in the language, are statistically close to the correct ones (for the precise details see [39, 40]).

Watrous defined the Quantum State Distinguishability promise problem ($\mathrm{QSD}_{\alpha,\beta}$):

Input: Quantum circuits Q_0, Q_1.
Accept: If $\| \tau_{Q_0} - \tau_{Q_1} \|_{\mathrm{tr}} \geq \beta$.
Reject: If $\| \tau_{Q_0} - \tau_{Q_1} \|_{\mathrm{tr}} \leq \alpha$.

where the notation τ_Q denotes the mixed state obtained by running the quantum circuit Q on the initial state $|0^n\rangle$ and tracing out the non-output qubits,[1] and $\| A \|_{\mathrm{tr}} = \mathrm{Tr}\,|A|$ is the quantum analogue of the classical ℓ_1-norm (and so in particular $\| \rho_1 - \rho_2 \|_{\mathrm{tr}}$ is the quantum analogue of the classical variational distance of two probability distributions).

Watrous showed $\mathrm{QSD}_{\alpha,\beta}$ is complete for honest-verifier-QSZK ($\mathrm{QSZK}_{\mathrm{HV}}$) when $0 \leq \alpha < \beta^2 \leq 1$. He further showed that $\mathrm{QSZK}_{\mathrm{HV}}$ is closed under complement, that any problem in $\mathrm{QSZK}_{\mathrm{HV}}$ has a 2-message proof system and a 3-message public-coin proof system and also that $\mathrm{QSZK} \subseteq \mathrm{PSPACE}$. Subsequently, in [40], he showed that $\mathrm{QSZK}_{\mathrm{HV}} = \mathrm{QSZK}$.

The above results have classical analogues. However, in the classical setting there is another canonical complete problem, the Entropy Difference problem (ED). There is a natural quantum analogue to ED, the Quantum Entropy Difference problem (QED), that we now define:

Input: Quantum circuits Q_0, Q_1.
Accept: If $S(\tau_{Q_0}) - S(\tau_{Q_1}) \geq \frac{1}{2}$.
Reject: If $S(\tau_{Q_1}) - S(\tau_{Q_0}) \geq \frac{1}{2}$.

where $S(\rho)$ is the Von-Neumann entropy of the mixed state ρ.[2] The problem QED is very natural from a physical point of view. It corresponds to the following task: we are given two mixed states, each given by a quantum circuit generating it, and we are asked to decide which mixed state has

[1] Here we assume that a quantum circuit also designates a set of output qubits.

[2] A density matrix ρ is positive semi-definite and has trace 1. Therefore its eigenvalues are all non-negative and sum up to 1, and can be thought of as defining a probability distribution. The Von-Neumann entropy of ρ is the Shannon entropy of the distribution defined by the eigenvalues of ρ.

more entropy. This problem is, in particular, as hard as[3] approximating the amount of entropy in a given mixed state (when again the mixed state is given by a circuit generating it).

We show that QED is QSZK–complete. For that we use *quantum extractors*, which are quantum variants of classical extractors. [4] We build quantum extractors from quantum expanders. It turns out that for the parameters we are interested in, the quantum extractors must have close to optimal entropy loss. The constructions of Theorem 3.1 and Theorem 4.3 are not good enough with that respect. However, the Ambainis-Smith construction has this property. Using it we get that QED is QSZK–complete. This implies that it is not likely that one can estimate quantum entropies in BQP.

Furthermore, a common way of measuring the amount of entanglement between registers A and B in a pure state ψ is by the Von-Neumann entropy of $\mathrm{Tr}_B(|\psi\rangle\langle\psi|)$ [34]. Now suppose we are given two circuits Q_1 and Q_2, both acting on the same initial pure-state $|0^n\rangle$, and we want to know which circuit produces more entanglement between A and B. Our result shows that this problem is QSZK–complete. As before, this also shows that the problem of *estimating* the amount of entanglement between two registers in a given pure-state is QSZK–hard (hence unlikely to be in BQP).

1.3 Summary, related and following work

Quantum expanders were independently defined in our technical report [7] and by Hastings [18]. This work initiated a lot of research on the subject. Hastings [19] proved a lower bound showing that degree D quantum expanders must have spectral gap at most $1 - \frac{2\sqrt{D-1}}{D} + o(1)$, matching the classical situation as proved by Alon and Boppana (see [32]). In the same paper he also showed that, non-explicitly, almost–Ramanujan quantum expanders exist, matching again the classical situation as proved by Friedman [12].

The first explicit construction was implicit in the work of Ambainis and Smith [4], and has poly-logarithmic degree. The construction pre-dates the definition of quantum expanders, and is a main component in solving another problem (which we soon describe).

Our first construction is based on the construction of [27] and was the first to achieve constant spectral gap and constant degree. The construction is not explicit, because currently it is not known how to efficiently implement the quantum Fourier transform over the group we work with. However, there is hope that with progress on the Fourier transform problem, it will become explicit.

Our second construction, that is based on the Zig-Zag construction of [35] is an explicit construction of constant degree, constant gap quantum expanders. This construction first appeared in our technical report [6]. It was followed shortly by two other explicit constructions [17] and [15]. The approaches in both papers is converting a classical expander into a quantum one. [15] shows how to do this for the expander of Margulis [28] while [17] shows how to do this for any classical Cayley graph where the underlying group has an efficient quantum Fourier transform and a large irreducible representation.

Quantum expanders were first used (implicitly) by Ambainis and Smith [4] to construct short quantum one-time pads. Loosely speaking, they showed how two parties sharing a random bit string of length $n + O(\log n)$ can communicate an n qubit state such that any eavesdropper cannot learn much about the transmitted state. A subsequent work [10] showed how to remove the $O(\log n)$ term.

Hastings [18] gave an application from physics. Using quantum expanders, he showed that there exists gapped one-dimensional systems for which the entropy between a given subvolume and the rest of the system is exponential in the correlation length. We add it to the proof that QSD is QSZK-complete.

The paper is organized as follows. After the preliminaries (Section 2), we give our first construction, and its analysis, in Section 3. In Section 4 we describe our second construction. Finally, Section 5 is devoted to proving the completeness of QED in QSZK.

2 Preliminaries

We first define the classical Rényi entropy. Let $P = (p_1, \dots, p_m)$ be a classical probability distribution. The *Shannon entropy* of P is $H(P) = \sum_{i=1}^{m} p_i \lg \frac{1}{p_i}$. The *min-entropy* of P is $H_\infty(P) = \min_i \lg \frac{1}{p_i}$. The *Rényi entropy* of P is $H_2(P) = \lg \frac{1}{\mathrm{Col}(P)}$, where $\mathrm{Col}(P) = \sum p_i^2$ is the collision probability of the distribution defined by $\mathrm{Col}(P) = \mathrm{Pr}_{x,y}[x = y]$ when x, y are sampled from P.

Now let $\rho \in D(V)$ be a density matrix (where V is a Hilbert space, $L(V)$ is the set of linear operators over V and $D(V)$ is the set of positive semi-definite operators in $L(V)$ with trace 1, i.e., all density matrices over V). Let $\alpha = (\alpha_1, \dots, \alpha_N)$ be the set of eigenvalues of ρ. Since ρ is positive semi-definite, all these eigenvalues are non-negative. Since $\mathrm{Tr}(\rho) = 1$ their sum is 1. Thus we can view α as a classical probability distribution. The *von Neumann entropy* of ρ is $S(\rho) = H(\alpha)$. The *min-entropy* of ρ is $H_\infty(\rho) = H_\infty(\alpha)$. The *Rényi entropy* of ρ is $H_2(\rho) = H_2(\alpha)$. The analogue of the collision probability is simply $\mathrm{Tr}(\rho^2) = \sum_i \alpha_i^2 = ||\rho||_2^2$. We remark that for any distribution P, $H_\infty(P) \le H_2(P) \le H(P)$ and

[3] Under Turing reductions.

[4] Quantum extractors should not be mixed with the classical extractors against quantum adversaries.

$2H_\infty(P) \geq H_2(P)$.

The *statistical difference* between two classical distributions $P = (p_1, \ldots, p_m)$ and $Q = (q_1, \ldots, q_m)$ is $\mathrm{SD}(P,Q) = \frac{1}{2}\sum_{i=1}^m |p_i - q_i|$, i.e., half the ℓ_1 norm of $P - Q$. This is generalized to the quantum setting by defining the trace-norm of a matrix $X \in L(V)$ to be $\| X \|_{\mathrm{tr}} = \mathrm{Tr}(|X|)$, where $|X| = \sqrt{XX^\dagger}$, and defining the *trace distance* between density matrices ρ and σ to be $\frac{1}{2} \| \rho - \sigma \|_{\mathrm{tr}}$.

3 Quantum expanders from non-Abelian Cayley graphs

This section is devoted to our first construction, which yields the following:

Theorem 3.1. *There exists a $(D = O(\frac{1}{\overline{\lambda}^4}), \overline{\lambda})$ quantum expander.*

As explained in the introduction, this quantum expander takes two steps on a Cayley expander (over the group PGL(2,q)) with a basis change between each of the steps, and the basis change is a carefully chosen transformation. In Subsection 3.1, we define and analyze taking one step on a (Abelian or non-Abelian) Cayley graph. Then, we study a general template for constructing quantum expanders over non-Abelian groups with a certain property (Subsections 3.2, 3.3) and show that PGL(2,q) has this required property.

In what follows we freely use notions from representation theory. Due to lack of space, for the necessary background we refer the reader to the book of Serre [38], to the book of Fulton and Harris [16], or to our technical report [7].

3.1 A single step on a Cayley graph

We fix an arbitrary (Abelian or non-Abelian) group G of order N. We assume there exists efficient classical algorithms (i.e. running in time $\mathrm{poly}(\log N)$) for multiplying two group elements and for inverting a group element.

Our starting point is generalizing a single step on a Cayley graph to the quantum setting. We fix a subset $\Gamma \subseteq G$ of group elements closed under inverse. The *Cayley graph* associated with Γ, $C(G,\Gamma)$, is a graph over N vertices, with an edge between (g_1, g_2) iff $g_1 = g_2\gamma$ for some $\gamma \in \Gamma$.

We identify the graph $C(G,\Gamma)$ with the operator given by its normalized adjacency matrix, which we denote by M. This is a linear operator over the space $\mathbb{C}[G] = \mathrm{Span}\{|x\rangle : x \in G\}$, and it is given by $M = \frac{1}{|\Gamma|}\sum_{\gamma\in\Gamma} |x\gamma\rangle\langle x|$.[5] Notice that $M = C(G,\Gamma)$ is a symmetric operator, and therefore diagonalizes with real eigenvalues. We denote by $\lambda_1 \geq \ldots \geq \lambda_N$ the eigenvalues of M with orthonormal eigenvectors $v_1, \ldots, v_N$ (i.e., $\| v_i \|_2 = 1$). Since our graph is regular, we have $\lambda_1 = 1$ and $\overline{\lambda} = \max_{i>1} |\lambda_i| \leq 1$.

We now define our basic superoperator $T : L(\mathbb{C}[G]) \to L(\mathbb{C}[G])$. The superoperator has a register R of dimension $|\Gamma|$ that is initialized at $|\overline{0}\rangle$. It does the following:

- It first applies Hadamard on register R (getting into the density matrix $\frac{1}{|\Gamma|}\rho \otimes \sum_{\gamma,\gamma'\in\Gamma} |\gamma\rangle\langle\gamma'|$).
- Then, it applies the unitary transformation $Z : |g,\gamma\rangle \to |g\gamma,\gamma\rangle$. This transformation is a permutation over the standard basis, and hence unitary. It has an efficient quantum circuit since, by our assumption on G, it is classically easy to compute it in both directions.
- Finally, it discards register R.

Thus we have: $T(\rho) = \mathrm{Tr}_R[\, Z(I\otimes H)(\rho\otimes|\overline{0}\rangle\langle\overline{0}|)(I\otimes H)Z^\dagger\,]$. It can be easily checked that over "classical" states (density matrices that are diagonal in the standard basis) T coincides with M. Also, by definition, T is $|\Gamma|$–regular.

We begin by identifying the eigenvectors and eigenvalues of T. We may think of an eigenvector $v_i \in \mathbb{C}^N$ as an element of $\mathbb{C}[G]$, $|v_i\rangle = \sum_g v_i(g)\,|g\rangle$. We also define the linear transformation $R : \mathbb{C}[G] \to L(\mathbb{C}[G])$ by $R\,|g\rangle = |g\rangle\langle g|$. With this notation we define:

$$\mu_{i,g} \quad = \quad \sum_{x\in G} v_i(x)\,|gx\rangle\langle x|$$

It can be verified that:

Lemma 3.2. *The vectors $\{\mu_{i,g} \mid i = 1,\ldots,N, g \in G\}$ form an orthonormal basis of $L(\mathbb{C}[G])$, and $\mu_{i,g}$ is an eigenvector of T with eigenvalue $\lambda_{i,g} = \lambda_i$ (the orthogonality is under the Hilbert-Schmidt inner-product defined by $\langle A|B\rangle = \mathrm{Tr}(AB^\dagger)$).*

Thus, T has N orthogonal eigenspaces, each of dimension N, and the eigenvalues $\lambda_1, \ldots, \lambda_N$ are those of M . In particular, if we start with a good Cayley graph where $\lambda_1 = 1$ and all other eigenvalues have absolute value at most $\overline{\lambda}$, then T has an eigenspace $\mu^{||}$ of dimension N with eigenvalue 1, and all other eigenvalues have absolute value at most $\overline{\lambda}$. The fact that the dimension of $\mu^{||}$ is larger than 1 is not good for us, because it means that T has no spectral gap. However, we know that $\mu^{||} =$

[5] In our definition the generators act from the right. Sometimes the Cayley graph is defined with left action, i.e., g_1 is connected to g_2 iff $g_1 = \gamma g_2$. However, note that if we define the invertible linear transformation P that maps the basis vector $|g\rangle$ to the basis vector g^{-1} , then $PMP^{-1} = PMP$ maps x to $\frac{1}{|\Gamma|}\sum_\gamma (x^{-1}\gamma)^{-1} = \frac{1}{|\Gamma|}\sum_\gamma \gamma^{-1}x = \frac{1}{|\Gamma|}\sum_\gamma |\gamma x\rangle$ and so the right action is M and the left action is PMP^{-1}, and therefore they are similar and in particular have the same spectrum.

$\mathrm{Span}\{\mu_{1,g} \mid g \in G\}$. Let $\mu^{\perp}$ denote the complementary subspace ($\mu^{\perp} = \mathrm{Span}\{\mu_{i,g} \mid i \neq 1, g \in G\}$).

The operator $\mu_{1,g} = \sum_x |gx\rangle\langle x|$ maps x to gx. These operators form what is called the *regular representation* of G. Namely, if we denote $\rho_{\mathrm{reg}}(g) = \mu_{1,g}$, then $\rho_{\mathrm{reg}} : G \to L(\mathbb{C}[G])$ is a group homomorphism (i.e., $\rho_{\mathrm{reg}}(g_1 \cdot g_2) = \rho_{\mathrm{reg}}(g_1) \cdot \rho_{\mathrm{reg}}(g_2)$). Furthermore, a basic theorem of representation theory says that there is a basis change under which all the operators $\mu_{1,g} = \rho_{\mathrm{reg}}(g)$ simultaneously block-diagonalize, with the blocks corresponding to the irreducible representations of G. This (non-unique) basis change is called the Fourier transform of G.

We let U be the Fourier transform over G, and we define the quantum expander to be the superoperator

$$E(\rho) = T(UT(\rho)U^{\dagger}).$$

A simple check shows that E is $|\Gamma|^2$–regular. For explaining the intuition behind our choice of E, we give a rough analysis for the case G is Abelian (and the underlying Cayley graph is a good expander). Suppose σ is some density matrix orthogonal to the completely-mixed state. We can decompose σ according to the basis $\{\mu_{i,g}\}$. For the rough intuition we analyze the action of E on each $\mu_{i,g}$ separately.[6]

Since σ is orthogonal to the completely-mixed state, we need to analyze only $\mu_{i,g} \neq \mu_{1,1}$. Let us analyze the action of E on μ_{i_0,g_0}, for some arbitrary $1 \leq i_0 \leq N$ and $g_0 \in G$ such that $(i_0, g_0) \neq (1,1)$. If $i_0 > 1$ then by Lemma 3.2, the first application of T shrinks μ_{i_0,g_0} (by a factor of λ_i). Hence, the entire operator E shrinks it. On the other hand, if $i_0 = 1$, then by Lemma 3.2 μ_{1,g_0} is unchanged by T. In Abelian groups all the irreducible representations have dimension one, and the Fourier transform U diagonalizes $\mu_{1,g_0} = \rho_{\mathrm{reg}}(g_0)$. Hence, after the basis change, the operator μ_{1,g_0} becomes diagonal, i.e., a "classical" state. Since U is unitary, the resulting diagonal matrix remains orthogonal to the completely-mixed state. Thus, the second application of T shrinks this diagonal matrix.

As explained in the introduction, there are no constant degree expanders over Abelian groups. Therefore, our next goal is analyzing E in the non-Abelian case.

3.2 Template for a quantum expander over a general group

Definition 3.3. *We say U is a* good basis change *if for any $g_1 \neq 1$ it holds that*

$$\mathrm{Tr}(U\rho_{\mathrm{reg}}(g_1)U^{\dagger}\rho_{\mathrm{reg}}(g_2)) \quad = \quad 0. \tag{1}$$

The intuition behind this choice is captured in the following claim:

Claim 3.4. *If U is a good basis change then for any $\rho \in \mathrm{Span}\{\rho_{\mathrm{reg}}(g) : g \neq 1 \in G\}$ we have $U\rho U^{\dagger} \perp \mu^{||}$.*

Proof: $\{\rho_{\mathrm{reg}}(g) : g \in G\}$ is an orthonormal basis for $\mu^{||}$. Therefore, it is enough to verify that for any $g_1 \neq 1$ and for any g_2 it holds that $\mathrm{Tr}(U\rho_{\mathrm{reg}}(g_1)U^{\dagger}\rho_{\mathrm{reg}}(g_2)^{\dagger}) = 0$. Since $\rho_{\mathrm{reg}}(g_2)^{\dagger} = \rho_{\mathrm{reg}}(g_2^{-1})$, this follows directly from Property (1). ■

Thus, intuitively, the analysis is similar to the Abelian case. Vectors from $\mu^{\perp}$ are shrunk by the first T application, and vectors in $\mu^{||}$ are left in place by T, mapped to $\mu^{\perp}$ by U, and then shrunk by the second T application. Indeed, we claim:

Lemma 3.5. *If U is a good basis change then E is a $(|\Gamma|^2, \overline{\lambda})$ quantum expander.*

Proof: The regularity is clear from the way the superoperator E is defined. We turn to the spectral gap. It is easy to check that $E(\tilde{I}) = \tilde{I}$. Furthermore, fix any $X \in L(\mathbb{C}[G])$ that is perpendicular to $\tilde{I}$. Write $X = X^{||} + X^{\perp}$ where $X^{||} \in \mathrm{Span}\{\mu_{1,g} \mid 1 \neq g \in G\}$ and $X^{\perp} \in \mu^{\perp}$. Now it is not true any more that $E(X^{||}) \perp E(X^{\perp})$.

However, $E(X) = T(\sigma^{||} + \sigma^{\perp})$, where $\sigma^{||} = UT(X^{||})U^{\dagger}$ and $\sigma^{\perp} = UT(X^{\perp})U^{\dagger}$. By Claim 3.4, $\sigma^{||} \perp \mu^{||}$. Also, $T(X^{||}) \perp T(X^{\perp})$, and therefore $\sigma^{||} \perp \sigma^{\perp}$. Finally, by Lemma 3.2 we know T is normal. We claim

Lemma 3.6. *Let T be a normal linear operator with eigenspaces $V_1, \ldots, V_n$ and corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$ in descending* absolute *value. Suppose u and w are vectors such that $u \in \mathrm{Span}\{V_2, \ldots, V_n\}$ and $w \perp u$ (w does not necessarily belong to V_1). Then*

$$||(T(u+w))||_2^2 \leq |\lambda_2|^2 ||u||_2^2 + |\lambda_1|^2 ||w||_2^2.$$

We omit the technical verification of the lemma for lack of space (it can be found in our technical report [7]).

Using the lemma it can verified that $||E(X)||_2^2 \leq \overline{\lambda}^2 ||X||_2^2$. ■

Thus, our next goal is checking whether a good basis change actually exists.

3.3 A sufficient condition that guarantees a good basis change

The Fourier transform is a unitary mapping from the standard basis $\{|g\rangle\}$ of $\mathbb{C}[G]$, to the Fourier basis. It can be formally defined as follows. Let $\widehat{G}$ denote the set of all

[6] In the Abelian case this intuition can be made precise, since each $\mu_{i,g}$ is an eigenvector of T, and the Fourier transform is simply a permutation over $\{\mu_{i,g}\}$.

inequivalent irreducible representations of G. For a representing ρ let d_ρ denote the dimension of ρ. We define the transform F by

$$F\left|g\right\rangle \;=\; \sum_{\rho\in\widehat{G}}\sum_{1\le i,j\le d_\rho}\sqrt{\frac{d_\rho}{|G|}}\rho_{i,j}(g)\left|\rho,i,j\right\rangle.$$

It can be checked that F is unitary and that it indeed block-digaonlizes the regular representations, namely,

$$F\rho_{\mathrm{reg}}(g)F^\dagger \;=\; \sum_{\rho\in\widehat{G}}\sum_{1\le i,i',j\le d_\rho}\rho_{i,i'}(g)\left|\rho,i,j\right\rangle\!\left\langle\rho,i',j\right|,$$

i.e., for each $\rho\in\widehat{G}$ and $j\le d_\rho$, we have a $d_\rho\times d_\rho$ block whose entries are $\rho(g)$.

F maps $\mathbb{C}[G]$ to a vector space of the same dimension that is spanned by $\left\{\left|\rho,i,j\right\rangle : \rho\in\widehat{G},\ 1\le i,j\le d_\rho\right\}$. To complete the specification of the Fourier transform we also need to specify a map S between $\{\left|\rho,i,j\right\rangle\}$ and $\{\left|g\right\rangle : g\in G\}$. In the Abelian case there is a canonical map S between $\left\{\left|\rho,i,j\right\rangle : \rho\in\widehat{G},\ i=j=1\right\}$ and $\{\left|g\right\rangle : g\in G\}$, because when G is Abelian $\widehat{G}$ is isomorphic to G. However, when G is not Abelian things are more complicated. It is always true that $\sum_{\rho\in\widehat{G}}d_\rho^2=|G|$, and so there is always a bijection between $\{\left|\rho,i,j\right\rangle\}$ and $\{\left|g\right\rangle : g\in G\}$. However, it is not known, in general, how to find such a natural bijection. The situation, however, is better for *product mappings*.

Definition 3.7. *Let f be a bijection from $\left\{(\rho,i,j)\mid\rho\in\widehat{G},1\le i,j\le d_\rho\right\}$ to G. We say that f is a* product mapping, *if for every $\rho\in\widehat{G}$, $f(\rho,i,j)=f_1(i)\cdot f_2(j)$ for some functions $f_1,f_2:[d_\rho]\times[d_\rho]\to G$ (f_1 and f_2 may depend on ρ).*

For example, for an Abelian group all irreducible representations are of dimension one, and so any such bijection f is a product mapping (since we can just define $f_1(1)=1$, $f_2(1)=f(\rho,1,1)$). Another example is the dihedral group where a simple product mapping exists.

Our claim is that any group that has a product mapping can be used to construct quantum expanders.

Lemma 3.8. *Let G be a group that has a product mapping f, and let F be the Fourier transform over G, $F\left|g\right\rangle=\sum_{\rho\in\widehat{G}}\sum_{1\le i,j\le d_\rho}\sqrt{\frac{d_\rho}{|G|}}\rho_{i,j}(g)\left|\rho,i,j\right\rangle$. Define the unitary mapping $S:\left|\rho,i,j\right\rangle\mapsto\omega_{d_\rho}^{ij}\left|f(\rho,i,j)\right\rangle$, where $\omega_{d_\rho}=e^{2\pi i/d_\rho}$, and set U to be the unitary transformation $U=SF$. Then U has property (1) and is a good basis change.*

Proof:

$$\begin{aligned}
&\mathrm{Tr}\left(U\rho_{\mathrm{reg}}(g_1)U^\dagger\rho_{\mathrm{reg}}(g_2)\right)\\
&= \mathrm{Tr}\left(S\sum_{\rho\in\widehat{G}}\sum_{i,i',j}\rho_{i,i'}(g_1)\left|\rho,i,j\right\rangle\!\left\langle\rho,i',j\right|S^\dagger\sum_x\left|g_2x\right\rangle\!\left\langle x\right|\right)\\
&= \sum_{\rho\in\widehat{G}}\sum_{i,i'}\rho_{i,i'}(g_1)\,\mathrm{Tr}\left(\sum_{j=1}^{d_\rho}S\left|\rho,i,j\right\rangle\!\left\langle\rho,i',j\right|S^\dagger\sum_x\left|g_2x\right\rangle\!\left\langle x\right|\right).
\end{aligned}$$

Therefore, it suffices to show that for any ρ,i,i' we have $\mathrm{Tr}\left(\sum_{j=1}^{d_\rho}S\left|\rho,i,j\right\rangle\!\left\langle\rho,i',j\right|S^\dagger\sum_x\left|g_2x\right\rangle\!\left\langle x\right|\right)=0$. Fix $\rho\in\widehat{G}$ and $i,i'\in\{1,\ldots,d_\rho\}$. Since f is product, $f(\rho,i,j)=f_1(i)\cdot f_2(j)$ for some $f_1,f_2:[d_\rho]\times[d_\rho]\to G$. Denote $h_i=f_1(i)$ and $t_j=f_2(j)$. The sum we need to calculate can be written as

$$\begin{aligned}
&\mathrm{Tr}\left(\sum_{j=1}^{d_\rho}S\left|\rho,i,j\right\rangle\!\left\langle\rho,i',j\right|S^\dagger\sum_x\left|g_2x\right\rangle\!\left\langle x\right|\right)\\
&= \sum_{j=1}^{d_\rho}\omega_{d_\rho}^{(i-i')j}\left\langle g_2\,|\,h_{i'}h_i^{-1}\right\rangle,
\end{aligned}$$

where the last equality is because we get a non-zero value iff $x=h_it_j$ and $h_{i'}t_j=g_2x$, which happens iff $h_it_j=g_2^{-1}h_{i'}t_j$, i.e., $g_2=h_{i'}h_i^{-1}$. However, when $g_2=h_{i'}h_i^{-1}$ we get the sum $\sum_{j=1}^{d_\rho}\omega_{d_\rho}^{(i-i')j}$. This expression itself is zero when $i\neq i'$.

We are therefore left with the case $i=i'$. In this case $g_2=h_{i'}h_i^{-1}=1$. But then,

$$\mathrm{Tr}\left(U\rho_{\mathrm{reg}}(g_1)U^\dagger\rho_{\mathrm{reg}}(g_2)\right) \;=\; \mathrm{Tr}\left(\rho_{\mathrm{reg}}(g_1)\right)=0,$$

where the last equality follows because $g_1\neq 1$. ■

It is not clear at all that for every group G such a product mapping exists. It is trivial for Abelian groups, and simple for the dihedral group.

To complete our construction we need to prove that $\mathrm{PGL}(2,q)$ has a product mapping. We do this using information about its subgroup structure, and its irreducible representations. For lack of space we omit this part from the extended abstract. The details can be found in our technical report [7].

Putting everything together, we get a quantum expander $E(\rho)=T(UT(\rho)U^\dagger)$, with T being a single quantum step on a the Cayley expander, and U being a good basis change. U is obtained from the Fourier transform and the product mapping, as explained in Lemma 3.8. We are now ready to prove Theorem 3.1:

Proof: By Lemma 3.5, Lemma 3.8 and the description of the product mapping above, we know that E is a $(|\Gamma|^2,\overline{\lambda})$

quantum expander. Both $|\Gamma|$ and $\overline{\lambda}$ are determined by the underlying Cayley graph we work with. By the construction of [27] we know that there exists a Cayley graph for $\mathrm{PGL}(2,q)$ with $\overline{\lambda}^2 \leq \frac{4}{|S|}$. Plugging this Cayley graph gives us a $(\frac{16}{\overline{\lambda}^4}, \overline{\lambda})$ quantum expander. ■

4 The Zig-Zag construction

The following section is devoted to our second construction of quantum expanders. The construction uses as building blocks the following operations:

- **Squaring:** For a superoperator $G \in T(\mathcal{V})$ we denote by G^2 the superoperator given by $G^2(X) = G(G(X))$ for any $X \in L(\mathcal{V})$.
- **Tensoring:** For superoperators $G_1 \in T(\mathcal{V}_1)$ and $G_2 \in T(\mathcal{V}_2)$ we denote by $G_1 \otimes G_2$ the superoperator given by $(G_1 \otimes G_2)(X \otimes Y) = G_1(X) \otimes G_2(Y)$ for any $X \in L(\mathcal{V}_1), Y \in L(\mathcal{V}_2)$.
- **Zig-Zag product:** For superoperators $G_1 \in T(\mathcal{V}_1)$ and $G_2 \in T(\mathcal{V}_2)$ we denote by $G_1 ⓩ G_2$ their Zig-Zag product. A formal definition of this is given in Section 4.2. The only requirement is that G_1 is $\dim(\mathcal{V}_2)$–regular.

Proposition 4.1. *If G is a (N, D, λ) quantum expander then G^2 is a (N, D^2, λ^2) quantum expander. If G is explicit then so is G^2. If G_1 is a (N_1, D_1, λ_1) quantum expander and G_2 is a (N_2, D_2, λ_2) quantum expander then $G_1 \otimes G_2$ is a $(N_1 \cdot N_2, D_1 \cdot D_2, \max(\lambda_1, \lambda_2))$ quantum expander. If G_1 and G_2 are explicit then so is $G_1 \otimes G_2$.*

The proof is trivial and is omitted. We next claim:

Theorem 4.2. *If G_1 is a (N_1, D_1, λ_1) quantum expander and G_2 is a (D_1, D_2, λ_2) quantum expander then $G_1 ⓩ G_2$ is a $(N_1 \cdot D_1, D_2^2, \lambda_1 + \lambda_2 + \lambda_2^2)$ quantum expander. If G_1 and G_2 are explicit then so is $G_1 ⓩ G_2$.*

We use Theorem 4.2 to analyze our construction. Later we give a full description of the Zig-Zag product, along with the proof of Theorem 4.2.

Following the iterative construction in [35] we get an explicit quantum expander construction. The construction starts with some constant-degree quantum expander, and iteratively increases its size via alternating operations of squaring, tensoring and Zig-Zag products. The tensoring is used to square the dimension of the superoperator. Then a squaring operation improves the second eigenvalue. Finally, the Zig-Zag product reduces the degree, without deteriorating the second eigenvalue too much.

Suppose H is a (D^8, D, λ) quantum expander. We define a series of superoperators as follows. The first two superoperators are $G_1 = H^2$ and $G_2 = H \otimes H$. For every $t > 2$ we define

$$G_t = \left(G_{\lceil \frac{t-1}{2} \rceil} \otimes G_{\lfloor \frac{t-1}{2} \rfloor} \right)^2 ⓩ H.$$

Similarly to [35], we claim:

Theorem 4.3. *For every $t > 0$, G_t is an explicit (D^{8t}, D^2, λ_t) quantum expander with $\lambda_t = \lambda + O(\lambda^2)$.*

The proof of the equivalent theorem in [35] is based only on the properties of the basic operations. Hence, once Proposition 4.1 and Theorem 4.2 are established, the proof is identical to the one in [35] and we do not repeat it.

4.1 The base superoperator

Theorem 4.3 relies on the existence of a good base superoperator H. In the classical setting, the probabilistic method assures us that a good base graph exists, and so we can use an exhaustive search to find one. The quantum setting exhibits a similar phenomena:

Theorem 4.4. *([19]) There exists a D_0 such that for every $D > D_0$ there exist a (D^8, D, λ) quantum expander for $\lambda = \frac{4\sqrt{D-1}}{D}$.* [7]

We will use an exhaustive search to find such a quantum expander. To do this we first need to transform the searched domain from a continuous space to a discrete one. We do this by using a net of unitary matrices, $S \subset U(\mathcal{H}_{D^8})$. S has the property that for any unitary matrix $U \in U(\mathcal{H}_{D^8})$ there exists some $V_U \in S$ such that

$$\sup_{\| X \| = 1} \left\| UXU^\dagger - V_U X V_U^\dagger \right\| \leq \lambda.$$

It is not hard to verify that indeed such S exists, with size depending only on D and λ. Moreover, we can find such a set in time depending only on D and λ.[8]

Suppose G is a (D^8, D, λ) quantum expander, $G(X) = \frac{1}{D}\sum_{i=1}^{D} U_i X U_i^\dagger$. We denote by G' the superoperator $G'(X) = \frac{1}{D}\sum_{i=1}^{D} V_{U_i} X V_{U_i}^\dagger$. Let $X \in L(\mathcal{H}_{D^8})$ be or-

[7] [19] actually shows that for any D there exist a $(D^8, D, (1 + O(D^{-16/15}\log D))\frac{2\sqrt{D-1}}{D})$ quantum expander.

[8] One way to see this is using the Solovay-Kitaev theorem (see, e.g., [9]). The theorem assures us that, for example, the set of all the quantum circuits of length $O(\log^4 \epsilon^{-1})$ generated only by Hadamard and Tofolli gates give an ϵ-net of unitaries. The accuracy of the net is measured differently in the Solovay-Kitaev theorem, but it can be verified that the accuracy measure we use here is roughly equivalent.

thogonal to $\tilde{I}$. Then:

$$\begin{aligned} \| G'(X) \| &= \left\| \frac{1}{D} \sum_{i=1}^{D} V_{U_i} X V_{U_i}^{\dagger} \right\| \\ &\leq \| G(X) \| + \lambda \| X \| \leq 2\lambda \| X \| . \end{aligned}$$

Hence, G' is a $(D^8, D, \frac{8\sqrt{D-1}}{D})$ quantum expander.[9] This implies that we can find a good base superoperator in time which depends only on D and λ.

4.2 The Zig-Zag product

We now define the Zig-Zag product and prove Theorem 4.2. Suppose G_1, G_2 are two superoperators, $G_i \in T(\mathcal{H}_{N_i})$, and G_i is a (N_i, D_i, λ_i) quantum expander. We further assume that $N_2 = D_1$. G_1 is D_1–regular and so it can be expressed as $G_1(X) = \frac{1}{D_1} \sum_d U_d X U_d^{\dagger}$ for some unitaries $U_d \in U(\mathcal{H}_{N_1})$. We lift the ensemble $\{U_d\}$ to a superoperator $\dot{U} \in L(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ defined by:

$$\dot{U}(|a\rangle \otimes |b\rangle) \quad = \quad U_b \, |a\rangle \otimes |b\rangle \, ,$$

and we define $\dot{G}_1 \in T(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ by $\dot{G}_1(X) = \dot{U} X \dot{U}^{\dagger}$.

Definition 4.5. *Let G_1, G_2 be as above. The Zig-Zag product, $G_1 ⓩ G_2 \in T(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ is defined to be $(G_1 ⓩ G_2)X = (I \otimes G_2)\dot{G}_1(I \otimes G_2^{\dagger})X$.*

We claim:

Proposition 4.6. *For any $X, Y \in L(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ such that X is orthogonal to the identity operator we have:*

$$| \langle \, G_1 ⓩ G_2 X \, , \, Y \rangle \, | \quad \leq \quad f(\lambda_1, \lambda_2) \, \| X \| \cdot \| Y \|$$

where $f(\lambda_1, \lambda_2) = \lambda_1 + \lambda_2 + \lambda_2^2$.

And as a direct corollary we get:

Theorem 4.7. *If G_1 is a (N_1, D_1, λ_1) quantum expander and G_2 is a (D_1, D_2, λ_2) quantum expander then $G_1 ⓩ G_2$ is a $(N_1 \cdot D_1, D_2^2, \lambda_1 + \lambda_2 + \lambda_2^2)$ quantum expander. If G_1 and G_2 are explicit then so is $G_1 ⓩ G_2$.*

Proof: Let X be orthogonal to $\tilde{I}$ and let $Y = (G_1 ⓩ G_2)X$. By Proposition 4.6 $\| Y \|^2 \leq f(\lambda_1, \lambda_2) \| X \| \cdot \| Y \|$. Equivalently, $\| (G_1 ⓩ G_2)X \| \leq f(\lambda_1, \lambda_2) \| X \|$ as required.

The explicitness of $G_1 ⓩ G_2$ is immediate from the definition of the Zig-Zag product. ■

We now turn to the proof of Proposition 4.6. We adapt the proof given in [35] for the classical case to the quantum setting. For that we need to work with linear operators instead of working with vectors. Consequently, we replace the vector inner-product used in the classical proof with the Hilbert-Schmidt inner product on linear operators, and replace the Euclidean norm on vectors, with the $\mathrm{Tr}(XX^{\dagger})$ norm on linear operators. Interestingly, the same proof carries over to this generalized setting. We omit the proof due to lack of space (and also since it is similar to the classical proof). The proof can be found in our technical report [6].

5 The complexity of estimating entropy

In this section we show that the QED problem (as defined in the introduction) is QSZK–complete. We do that by showing that QED reduces to QSD and vice versa, using the already known fact that QSD is QSZK–complete.

For proving QED $\leq$ QSD we first consider a related problem QEA. QEA is the problem of comparing the entropy of a given quantum circuit to some *known* threshold t, instead of comparing the entropies of two quantum circuits as in QED. Specifically, QEA is defined as follows.

Input: Quantum circuit Q, a non-negative integer t.
Accept: If $S(\tau_Q) \geq t + \frac{1}{2}$.
Reject: If $S(\tau_Q) \leq t - \frac{1}{2}$.

We show that QEA reduces to $\overline{\text{QSD}}$, where $\overline{\text{QSD}}$ is the complement promise problem of QSD. Our proof adapts the corresponding classical proof, by using quantum expanders and we discuss it in the next sub-section. Combining QEA $\leq \overline{\text{QSD}}$ with the result of Watrous [39] that $\overline{\text{QSD}} \leq$ QSD, we conclude that QEA $\leq$ QSD.

In the classical setting it is well known that SD is closed under Boolean formula [37]. The quantum analogue is also true, and we will give the straight forward proof in the full version of the paper (the reader can find the proof in our technical report [7]). We can express QED$(Q_0, Q_1) = \bigvee_{t=1} [((Q_0, t) \in \text{QEA}_Y) \wedge ((Q_1, t) \in \text{QEA}_N)]$ and it therefore follows that QED reduces to QSD as desired.

The direction that QSD $\leq$ QED follows the classical reduction, but using the Holevo bound from quantum information theory. The details will be given in the full version of the paper and also appear in our technical report [7].

5.1 QEA $\leq$ QSD

5.1.1 The classical reduction

The classical reduction from EA to SD (where EA is like QEA but with the input being a classical circuit) uses *extractors*. An extractor is a function $E : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$. We say E is a (k, ϵ) extractor if for every distribution X on $\{0,1\}^n$ that has k min-entropy the distribution

[9] We can actually get an eigenvalue bound of $(1 + \epsilon)\frac{2\sqrt{D-1}}{D}$ for an arbitrary small ϵ on the expense of increasing D_0.

$E(X, U_d)$ obtained by sampling $x \in X$, $y \in \{0,1\}^d$ and outputting $E(x,y)$, is ϵ–close to uniform.

We begin with the classical intuition why EA reduces to SD. We are given a circuit C and we want to distinguish between the cases the distribution it defines has substantially more or less than t entropy. First assume that the distribution is flat, i.e., all elements that have a non-zero probability in the distribution, have equal probability. In such a case we can apply an extractor on the n output bits of C, hashing it to about t output bits. If the input distribution has high entropy, it also has high min-entropy (because for flat distributions entropy is the same as min-entropy) and therefore the output of the extractor is close to uniform. If, on the other hand, the circuit entropy is less than $t - d - 1$, where d is the extractor seed length, than even after applying the extractor the output distribution has at most $t - 1/2$ entropy, and therefore it must be far away from uniform. We get a reduction to $\overline{\text{SD}}$.

There are, of course, a few gaps to complete. First, our source is not necessarily flat. This is solved in the classical case by taking many independent copies of the circuit, which makes the output distribution "close" to "nearly-flat". A simple analysis shows that this flattening works also in the quantum setting. Also, we need to amplify the gap we have between entropy $t + 1/2$ and $t - 1/2$ to a gap larger than d (the seed length). This, again, is solved by taking many independent copies of C, because $S(C^{\otimes q}) = qS(C)$.

Thus, before we get to the quantum generalization, we first discuss quantum extractors.

5.1.2 Quantum extractors

Definition 5.1. *Let V be a Hilbert space of dimension N. A superoperator $T : L(V) \to L(V)$ is a (k,d,ϵ)* quantum extractor, *if T is 2^d–regular and for every $\rho \in D(V)$ with $H_\infty(\rho) \geq k$ we have $\left\| T\rho - \tilde{I} \right\|_{\text{tr}} \leq \epsilon$, where $\tilde{I} = \frac{1}{N} I$. We say T is explicit if T can be implemented by a polynomial-size quantum circuit (i.e. polynomial in $\log N$).*

We mention that if T is 2^d–regular (and, in particular, if it is a (k,d,ϵ) quantum extractor) then for any $\rho \in L(V)$ it holds that $S(T\rho) \leq S(\rho) + d$, i.e., the extractor never adds more than d entropy to the quantum state it acts upon. Classically, balanced extractors are closely related to expanders (e.g., [14]). This generalizes to the quantum setting. We prove:

Lemma 5.2. *Let V be a Hilbert space of dimension $N = 2^n$. If $T : L(V) \to L(V)$ is a $(D = 2^d, \bar{\lambda})$ quantum expander, then for every $t > 0$, T is also a $(k = n - t, d, \epsilon)$ quantum extractor with $\epsilon = 2^{t/2} \cdot \bar{\lambda}$.*

Proof: T has a dimension 1 eigenspace W_1 with eigenvalue 1, spanned by the norm 1 eigenvector $v_1 = \frac{1}{\sqrt{N}} I$ (where $dim(V) = N$). Our input ρ is a density matrix and therefore $\langle \rho | v_1 \rangle = \frac{1}{\sqrt{N}} \text{Tr}(\rho) = \frac{1}{\sqrt{N}}$. In particular $\rho - \frac{1}{\sqrt{N}} v_1 = \rho - \tilde{I}$ is perpendicular to W_1. Therefore,

$$||T(\rho) - \tilde{I}||_2^2 \;=\; ||T(\rho - \tilde{I})||_2^2 \leq \bar{\lambda}^2 ||\rho - \tilde{I}||_2^2 \leq \bar{\lambda}^2 ||\rho||_2^2.$$

Plugging $H_2(\rho) \geq H_\infty(\rho) \geq k = n - t$ we see that $||T(\rho) - \tilde{I}||_2^2 \leq \bar{\lambda}^2 2^{-(n-t)}$. Using Cauchy-Schwartz $\left\| T(\rho) - \tilde{I} \right\|_{\text{tr}} \leq \sqrt{N} ||T(\rho) - \tilde{I}||_2 \leq \epsilon$. ■

By using our construction of explicit quantum expanders we get

Corollary 5.3. *For every $n, t, \varepsilon \geq 0$ there exists an explicit $(n - t, d, \epsilon)$ quantum extractor $T : L(V) \to L(V)$ where $n = \log(\dim(V))$, and $d = 2(t + 2\log(\frac{1}{\varepsilon})) + O(1)$.*

Proof: Let $n, t, \varepsilon \geq 0$ and set $\lambda = \varepsilon 2^{-t/2}$. Choose D such that $\frac{1}{2}\lambda \leq \frac{8}{D^{1/4}} \leq \lambda$ and set $d = \log D$. Hence, $d \leq \log(8192 \cdot \varepsilon^{-4} 2^{2t}) = 2(t + 2\log(\frac{1}{\varepsilon})) + O(1)$.

Applying Theorem 4.3 with the base superoperator given by Theorem 4.4, we can get an explicit $(D, \frac{8}{D^{1/4}})$ quantum expander over a space of dimension 2^n. Lemma 5.2 completes the proof. ■

By using the construction of Ambainis and Smith [4] given in Theorem 1.4 we get

Corollary 5.4. *For every $n, t, \varepsilon \geq 0$ there exists an explicit $(n - t, d, \epsilon)$ quantum extractor $T : L(V) \to L(V)$ where $n = \log(\dim(V))$, and $d = t + 2\log(\frac{1}{\varepsilon}) + \log(n) + O(1)$.*

Remark 5.5. *A natural generalization of Definition 5.1 is for a superoperator $T : L(V) \to L(W)$ where V, W are Hilbert spaces of arbitrary dimensions N and M. I.e., here we let W be different than V, and, in particular, the superoperator T can map a large Hilbert space V to a much smaller Hilbert space W. In the classical case this corresponds to hashing a large universe $\{0,1\}^n$ to a much smaller universe $\{0,1\}^m$. In the classical world highly unbalanced extractors exist with a very short seed length d, and these objects have numerous applications. We suspect that no non-trivial unbalanced quantum extractors exist when $dim(W) < dim(V)/2$.*

5.1.3 Preliminaries

A polarization lemma. We need the polarization lemma [39] (that is based on the work of [36]), which is used throughout the section.

Theorem 5.6. ***(Polarization lemma, Theorem 5 at [39])*** *Let α and β satisfy $0 \leq \alpha < \beta^2 \leq 1$. Then there is a deterministic polynomial-time procedure that, on input*

$(Q_0, Q_1, 1^n)$ where Q_0 and Q_1 are quantum circuits, outputs descriptions of quantum circuits (R_0, R_1) (each having size polynomial in n and in the size of Q_0 and Q_1) such that

$$\begin{aligned} \| \tau_{Q_0} - \tau_{Q_1} \|_{\mathrm{tr}} \leq \alpha \quad &\Rightarrow \quad \| \tau_{R_0} - \tau_{R_1} \|_{\mathrm{tr}} \leq 2^{-n}, \\ \| \tau_{Q_0} - \tau_{Q_1} \|_{\mathrm{tr}} \geq \beta \quad &\Rightarrow \quad \| \tau_{R_0} - \tau_{R_1} \|_{\mathrm{tr}} \geq 1 - 2^{-n}. \end{aligned}$$

A flattening lemma. We also need a quantum version of a standard way to flatten distributions:

Definition 5.7. *Let ρ be a density matrix, λ an eigenvalue of ρ and Δ a positive number. We say that λ is Δ-*typical *if $2^{-S(\rho)-\Delta} \leq \lambda \leq 2^{-S(\rho)+\Delta}$. We say ρ is Δ-*flat *if for every $t > 0$, with probability $\geq 1 - 2^{-t^2+1}$, a measurement of ρ in its eigenvector basis results with an eigenvector whose eigenvalue is $t\Delta$-typical .*

A straight forward argument (using the Hoeffding bound) shows that:

Lemma 5.8. *Let ρ be a density matrix and k a positive integer. Suppose that every non-zero eigenvalue of ρ is at least 2^{-m}. Then $\otimes^k \rho$ is Δ-flat for $\Delta = \sqrt{k}m$.*

Finally we relate the distance of a density matrix from the completely-mixed state to its entropy. Consider the following classical random variable X over $\{0,1\}^n$: with probability ϵ, X samples the fixed string 0^n and with probability $1-\epsilon$, X is uniformly distributed over $\{0,1\}^n$. This X has distance about ϵ from uniform ($\epsilon + \frac{1-\epsilon}{2^n} - \frac{1}{2^n}$ to be exact) and its entropy is $S(\rho) \leq (1-\epsilon)n + H(1-\epsilon)$. This is essentially the worst possible:

Lemma 5.9. *Let ρ be a density matrix over n qubits and $\epsilon > 0$. If $S(\rho) \leq (1-\epsilon)n$ then $\| \rho - \frac{1}{2^n} I \|_{\mathrm{tr}} \geq \epsilon - \frac{1}{2^n}$.*

The proof follows from convexity.[10]

5.1.4 The quantum reduction: QEA $\leq \overline{\text{QSD}}$

We are now ready to prove the reduction. In Section 5.1.1 we gave an intuitive explanation of the classical reduction. We follow the same outline in the quantum case. Let (Q, t) be an input to QEA, where Q is a quantum circuit with n input qubits and m output qubits. We first look at the circuit $Q^{\otimes q}$ (for some $q = \mathrm{poly}(n)$ to be specified later). We let E be a $(qt, q(m-t) + 2\log(\frac{1}{\varepsilon}) + \log(qm) + O(1), \epsilon)$ quantum extractor operating on qm qubits, where $\epsilon = 1/\mathrm{poly}(n)$ will be fixed later. Such an extractor exists by Corollary 5.4. Let $\xi = E(\tau_Q^{\otimes q})$ and let $\tilde{I} = 2^{-qm} I$. The output of the reduction is $(\xi, \tilde{I})$.

To show correctness we prove:

Lemma 5.10. *If $(Q,t) \in QEA_Y$ then $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \leq 5\epsilon$. If $(Q,t) \in QEA_N$ then $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \geq \frac{1}{qm} - \frac{1}{2^{qm}}$.*

Proof: The first part: Since Q traces out at most n qubits, the eigenvalues of τ_Q are all at least 2^{-n}, and by Lemma 5.8 we see that $\tau_Q^{\otimes q}$ is Δ-flat for $\Delta = \sqrt{q}n$. Thus, with probability at least $1 - 2^{-r^2+1}$, a measurement of τ_Q in its eigenvector basis results with an eigenvector whose eigenvalue is $r\Delta$-typical. Let Λ denote the set of $r\Delta$-typical eigenvalues of τ_Q, for $r = \sqrt{\log(\frac{1}{\varepsilon})}$. We write $\tau_Q^{\otimes q}$ in its eigenvector basis $\tau_Q^{\otimes q} = \sum_i \lambda_i \left| v_i \right\rangle\!\left\langle v_i \right|$. Let $\sigma_0 = \sum_{\lambda_i \in \Lambda} \lambda_i \left| v_i \right\rangle\!\left\langle v_i \right|$, and let $\sigma_1 = \rho^{\otimes q} - \sigma_0$. Thus, $\mathrm{Tr}(\sigma_0) \geq 1 - 2^{-r^2+1}$. Therefore,

$$\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \quad \leq \quad \left\| E(\frac{1}{\mathrm{Tr}(\sigma_0)} \sigma_0) - \tilde{I} \right\|_{\mathrm{tr}} + 2^{-r^2+2}.$$

Now we use the fact that $\frac{1}{\mathrm{Tr}(\sigma_0)}\sigma_0$ is a density matrix with all its eigenvalues $\leq 2^{-q \cdot S(\rho) + r\Delta} \cdot \frac{1}{\mathrm{Tr}(\sigma_0)} \leq 2^{-q \cdot S(\rho) + r\Delta + 1}$. Thus, $\frac{1}{\mathrm{Tr}(\sigma_0)}\sigma_0$ has min-entropy at least $q \cdot S(\rho) - r\Delta - 1 \geq q \cdot (t+1) - r\Delta - 1$ since we started with a yes instance for QEA_Y. We set the parameters such that $q \geq r\Delta + 1$, and thus our density matrix has min-entropy at least qt and by the guarantee of our quantum extractor we get that $\left\| E(\frac{1}{\mathrm{Tr}(\sigma_0)}\sigma_0) - \tilde{I} \right\|_{\mathrm{tr}} \leq \epsilon$. Therefore, $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \leq \epsilon + 2^{-r^2+2} \leq 5\epsilon$, where the last inequality holds for $r \geq \sqrt{\log(\frac{1}{\varepsilon})}$.

The second part: Suppose that $(Q,t) \in \mathrm{QEA}_N$. Because the extractor is of small degree and does not add much entropy, $S(\xi)$ is bounded by $S(\tau_Q^{\otimes q}) + q(m-t) + 2\log(\frac{1}{\varepsilon}) + \log(qm) + O(1)$. Also, $S(\tau_Q^{\otimes q}) = qS(\tau_Q) \leq q(t-1)$. This is bounded by $qm - 1$ if we choose the parameters such that $q > 2\log(\frac{1}{\varepsilon}) + \log(qm) + O(1)$.

To summarize $S(\xi) \leq qm - 1$. By Lemma 5.9 it follows that $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \geq \frac{1}{qm} - \frac{1}{2^{qm}}$ as required. ■

The constraints we have on the parameters are $q \geq \sqrt{\log(\frac{1}{\varepsilon})}\sqrt{q}n + 1$ and $q > 2\log(\frac{1}{\varepsilon}) + \log(qm) + O(1)$. To this we add $5\epsilon < \left(\frac{1}{qm} - \frac{1}{2^{qm}} \right)^2$. This ensures a gap which can be amplified by Theorem 5.6 to any desired gap, and completes the proof. These constraints can be easily satisfied by choosing q and ϵ^{-1} to be appropriately large polynomials in n.

Acknowledgements

We thank Oded Regev for pointing out [4] to us. We also thank Ashwin Nayak, Oded Regev, Adam Smith and

[10] A variant with slightly different parameters follows from Fannes' inequality.

Umesh Vazirani for helpful discussions about the paper.

References

[1] N. Alon. Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986.

[2] N. Alon and V. Milman. λ_1, isoperimetric inequalities for graphs, and superconcentrators. *Journal of Combinatorial Theory. Series B*, 38(1):73–88, 1985.

[3] N. Alon and Y. Roichman. Random Cayley Graphs and Expanders. *Random Structures and Algorithms*, 5(2):271–285, 1994.

[4] A. Ambainis and A. Smith. Small pseudo-random families of matrices: Derandomizing approximate quantum encryption. In *RANDOM*, pages 249–260, 2004.

[5] R. Beals. Quantum computation of Fourier transforms over symmetric groups. *STOC*, pages 48–53, 1997.

[6] A. Ben-Aroya, O. Schwartz, and A. Ta-Shma. An explicit, constant degree quantum expander. Technical report, arXiv:0709.0911, 2007.

[7] A. Ben-Aroya and A. Ta-Shma. Quantum expanders and the quantum entropy difference problem. Technical report, arXiv:quant-ph/0702129, 2007.

[8] M. Capalbo, O. Reingold, S. Vadhan, and A. Wigderson. Randomness conductors and constant-degree expansion beyond the degree / 2 barrier. In *STOC*, pages 659–668, 2002.

[9] C. M. Dawson and M. A. Nielsen. The Solovay-Kitaev algorithm. *Quantum Information & Computation*, 6(1):81–95, 2006.

[10] P. Dickinson and A. Nayak. Approximate randomization of quantum states with fewer bits of key. In *AIP Conference Proceedings*, volume 864, pages 18–36, 2006.

[11] J. Dodziuk. Difference equations, isoperimetric inequality and transience of certain random walks. *Trans. Amer. Math. Soc.*, 284(2):787–794, 1984.

[12] J. Friedman. A proof of Alon's second eigenvalue conjecture. *Memoirs of the AMS*, to appear.

[13] O. Gabber and Z. Galil. Explicit Constructions of Linear-Sized Superconcentrators. *Journal of Computer and System Sciences*, 22(3):407–420, 1981.

[14] O. Goldreich and A. Wigderson. Tiny families of functions with random properties: A quality-size trade-off for hashing. *Random Structures and Algorithms*, 11(4):315–343, 1997.

[15] D. Gross and J. Eisert. Quantum margulis expanders. Technical report, arXiv:0710.0651, 2007.

[16] J. Harris and W. Fulton. *Representation Theory*. Springer, 1991.

[17] A. Harrow. Quantum expanders from any classical cayley graph expander. Technical report, arXiv:0709.1142, 2007.

[18] M. B. Hastings. Entropy and entanglement in quantum ground states. *Phys. Rev. B*, 76(3):035114, 2007.

[19] M. B. Hastings. Random unitaries give quantum expanders. *Phys. Rev. A*, 76(3):032315, 2007.

[20] S. Hoory, N. Linial, and A. Wigderson. Expander graphs and their applications. *Bulletin of the AMS*, 43(4):439–561, 2006.

[21] S. Jimbo and A. Maruoka. Expanders obtained from affine transformations. *Combinatorica*, 7(4):343–355, 1987.

[22] N. Kahale. Eigenvalues and expansion of regular graphs. *Journal of the ACM*, 42(5):1091–1106, 1995.

[23] M. Kassabov. Symmetric groups and expander graphs. Available at http://arxiv.org/abs/math.GR/0505624.

[24] A. Kitaev, A. Shen, and M. Vyalyi. *Classical and Quantum Computation*. Am. Math. Soc., Providence, Rhode Island, 2002.

[25] M. M. Klawe. Limitations on Explicit Constructions of Expanding Graphs. *SIAM J. Comput.*, 13(1):156–166, 1984.

[26] J. D. Lafferty and D. Rockmore. Fast fourier analysis for SL_2 over a finite field and related numerical experiments. *Experiment. Math.*, 1(2):115–139, 1992.

[27] A. Lubotzky, R. Philips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8:261–277, 1988.

[28] G. A. Margulis. Explicit constructions of expanders. *Problemy Peredaci Informacii*, 9(4):71–80, 1973.

[29] G. 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.

[30] 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.

[31] M. Nielsen and I. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[32] A. Nilli. On the second eigenvalue of a graph. *Discrete Mathematics*, 91(2):207–210, 1991.

[33] M. Pinsker. On the complexity of a concentrator. In *7th Internat. Teletraffic Confer.*, pages 318/1–318/4, 1973.

[34] S. Popescu and D. Rohrlich. Thermodynamics and the measure of entanglement. *Physical Review A*, 56(5):3319–3321, 1997.

[35] O. Reingold, S. Vadhan, and A. Wigderson. Entropy waves, the zig-zag graph product, and new constant degree expanders and extractors. In *FOCS*, pages 3–13, 2000.

[36] A. Sahai and S. Vadhan. A complete promise problem for statistical zero-knowledge. In *FOCS*, pages 448–457, 1997.

[37] A. Sahai and S. Vadhan. Manipulating statistical difference. In *Randomization Methods in Algorithm Design (DIMACS Workshop)*, pages 251–270, 1999.

[38] J. P. Serre. *Linear representations of finite groups*, volume 42 of *Graduate texts in Mathematics*. Springer, 1977.

[39] J. Watrous. Limits on the power of quantum statistical zero-knowledge. In *FOCS*, pages 459–470, 2002.

[40] J. Watrous. Zero-knowledge against quantum attacks. In *STOC*, pages 296–305, 2006.