THEORY OF COMPUTING, Volume 6 (2010), pp. 47–79
www.theoryofcomputing.org

Quantum Expanders: Motivation and Constructions

Avraham Ben-Aroya* Oded Schwartz Amnon Ta-Shma†

Received: August 3, 2008; published: May 2, 2010.

Abstract: We define quantum expanders in a natural way and 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 $\mathrm{PGL}(2,q)$ given by Lubotzky, Phillips, and Sarnak (1988). The second construction is combinatorial, and is based on a quantum variant of the Zig-Zag product introduced by Reingold, Vadhan, and Wigderson (2000). Both constructions are of constant degree, and the second one is explicit.

Using another construction of quantum expanders by Ambainis and Smith (2004), 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 follows the classical proof structure that the entropy difference problem is SZK–complete, but crucially depends on the use of quantum expanders.

ACM Classification: F.2.0, F.2.3

AMS Classification: 81P68, 68Q17

Key words and phrases: Quantum expanders, quantum entropy difference, QSZK

*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.

†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.

© 2010 Avraham Ben-Aroya and Oded Schwartz and Amnon Ta-Shma
© Licensed under a Creative Commons Attribution License
DOI: 10.4086/toc.2010.v006a003

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 S of the vertices, it measures the size of the set $\Gamma(S)$ of neighbors of S, relative to the size of 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 [22] for an excellent survey on the subject.

Pinsker [34] was the first to observe that constant degree random graphs have almost-optimal set expansion. Explicitly finding such a graph turned out to be a major challenge. One line of research focused on the algebraic measure of expansion, and this led to a series of explicit constructions based on algebraic structures, e. g., [29, 14, 23]. This line of research culminated in the works of Lubotzky, Phillips and Sarnak [28], Margulis [30], and Morgenstern [31] who explicitly constructed Ramanujan graphs, i. e., D-regular graphs achieving spectral gap of $1 - 2\sqrt{D-1}/D$, for D a prime power. Friedman [13] showed that random graphs are "almost Ramanujan" and Alon and Boppana (see [33]) showed Ramanujan graphs have almost the best possible algebraic expansion. Several works [12, 2, 3, 24] showed intimate connections between set expansion and algebraic expansion. We refer the reader, again, to the survey paper [22].

The algebraic definition of expansion views a regular graph $G = (V,E)$ as a linear operator on a Hilbert space $\mathcal{V}$ of dimension $|V|$. In this view an element $v \in V$ is identified with a basis vector $|v\rangle \in \mathcal{V}$, and a distribution π on V corresponds to the vector $|\pi\rangle = \sum_{v\in V} \pi(v)\,|v\rangle$. The action of G on $\mathcal{V}$ is the action of the normalized adjacency matrix $A : \mathcal{V} \to \mathcal{V}$, where the normalization factor is the degree of G, and therefore A maps probability distributions to probability distributions. This mapping corresponds to taking a random walk on G. Specifically, if one takes a random walk on G starting at time 0 with the distribution π_0 on, then the distribution on the vertices at time k is $A^k\,|\pi_0\rangle$. Viewing G as a linear operator allows one to consider the action of A on arbitrary vectors in $\mathcal{V}$, not necessarily corresponding to distributions over V. While such vectors have no combinatorial interpretation, they are crucial for understanding the spectrum of A; none of the non-trivial eigenvectors of A correspond to probability distributions. To summarize: a D-regular expander $G = (V,E)$ is a linear transformation $A : \mathcal{V} \to \mathcal{V}$ that can be implemented by a classical circuit and maps probability distributions to probability distributions. It is a good expander if it has a *large spectral gap* and a *small degree*.

We now want to extend the definition of D-regular expanders to linear operators that map *quantum states* to *quantum states*. A general quantum state is a *density matrix*, which is a trace 1, positive semidefinite operator, i. e., an operator of the form $\rho = \sum p_v\,|\psi_v\rangle\langle\psi_v|$, where $0 \le p_v \le 1$, $\sum p_v = 1$, and $\{\psi_v\}$ is an orthonormal basis of $\mathcal{V}$. Notice that $\rho \in L(\mathcal{V}) \triangleq \mathrm{Hom}(\mathcal{V},\mathcal{V})$.

Among the set of *admissible* quantum transformations $E : L(\mathcal{V}) \to L(\mathcal{V})$, which are those implementable by quantum circuits (allowing both unitary operations and measurements), are those given by the following definition.

Definition 1.1. A superoperator $E : L(\mathcal{V}) \to L(\mathcal{V})$ is a *D-regular admissible superoperator* if

$$E = \frac{1}{D}\sum_{d=1}^{D} E_d\,,$$

where, for each $d \in [D]$, $E_d(X) = U_d X U_d^\dagger$ for some unitary transformation U_d over $\mathcal{V}$.

Note that this definition generalizes the classical one: any D-regular graph can be viewed as a sum of D permutations, each corresponding to a unitary transformation. In fact many classical expander constructions explicitly use this property [36, 9]. The definition is also intuitive in a more basic sense. Unitary transformations (or permutations in the classical setting) are those transformations that do not change the entropy of a state. An operator has small degree if it can never add much entropy to the state it acts upon. Specifically, a degree D operator can never add more than $\log(D)$ entropy. Such a view is almost explicit in the work of Capalbo et al. [9], where they view expanders as entropy conductors.

It is clear that all of the singular values of a D-regular admissible super-operator $E : L(\mathcal{V}) \to L(\mathcal{V})$ are at most 1, and that the completely mixed state $\tilde{I} = I/|V|$ is an eigenvector of any such E, with corresponding eigenvalue 1. We say that such a super-operator E has a $1-\overline{\lambda}$ spectral gap if all the remaining singular values of E are smaller than $\overline{\lambda}$. This is analogous to the way regular, directed expanders are defined, where the regularity implies that the largest eigenvalue is 1, and furthermore this eigenvalue is obtained with the normalized all-ones vector (that corresponds to the uniform distribution). The spectral gap requires that all other singular values are bounded by $\overline{\lambda}$.

Definition 1.2. A D-regular admissible superoperator $E : L(\mathcal{V}) \to L(\mathcal{V})$ is $\overline{\lambda}$–expanding if:

- The eigenspace of E corresponding to the eigenvalue 1 is the one-dimensional space spanned by $\tilde{I}$.
- For any $B \in L(\mathcal{V})$ orthogonal to $\tilde{I}$, it holds that $\|E(B)\|_2 \leq \overline{\lambda}\,\|B\|_2$.

We also say that a superoperator $E : L(\mathcal{V}) \to L(\mathcal{V})$ is a $(\dim(\mathcal{V}), D, \overline{\lambda})$ *quantum expander* if it is D-regular and $\overline{\lambda}$–expanding, and that it is *explicit* if it can be implemented by a quantum circuit of size polynomial in $\log(\dim(\mathcal{V}))$. We sometimes omit the dimension and say that E is a $(D, \overline{\lambda})$ quantum expander.

The orthogonality in the above definition is with respect to the *Hilbert-Schmidt inner product*, defined as $\langle A, B\rangle = \mathrm{Tr}(A^\dagger B)$, and the norm is the one induced by this inner product: $\|B\|_2 = \sqrt{\langle B, B\rangle}$.

Definition 1.2 implies that 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 to the more general class of superoperators that can be expressed as the sum of D Kraus operators, but for simplicity we work only with D-regular admissible superoperators. A similar definition was independently given by Hastings [19].

1.1 Quantum expander constructions

In this paper we give two quantum expander constructions. We give a brief review of all the currently known constructions in the order in which they appeared. All of the constructions are essentially based on classical expanders, with a twist allowing them to work in the quantum setting as well.

The first construction was already implicit in the work of Ambainis and Smith [6] on state randomization:

Theorem 1.3 ([6])**.** *For every $\overline{\lambda} > 0$, there exists an explicit $\left(N, O\left(\log^2(N)/\overline{\lambda}^2\right), \overline{\lambda}\right)$ quantum expander.*

Their quantum expander is based on a Cayley expander over the Abelian group $\mathbb{Z}_2^n$. The main drawback of Cayley graphs over Abelian groups is that [26, 4] showed that such an approach cannot yield constant degree expanders. Indeed, this is reflected in the $\log^2 N$ term in Theorem 1.3. There are constant degree, Ramanujan Cayley graphs, i. e., Cayley graphs that achieve the best possible relationship between the degree and the spectral gap, and in fact the construction in [28] and [30] is such, but they are built over non-Abelian groups.

In order to work with general groups, we describe (in Section 3.2) a natural way to lift a Cayley graph $G = (V,E)$ into a corresponding quantum superoperator T. However, the analysis shows that the spectral gap of T is 0, and more specifically, T has $|V|$ eigenspaces each of dimension $|V|$, with eigenvalues $\vec{\lambda} = (\lambda_1 = 1, \ldots, \lambda_{|V|})$, where $\vec{\lambda}$ is the spectrum of the Cayley graph.

Our first construction starts with the constant degree Ramanujan expander presented in [28]. 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 (by applying the superoperator T twice), with a basis change between the two steps. The basis change is a carefully chosen refinement of the Fourier transform 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 degrees of freedom, 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 ensures that the construction works. In this work we single out a natural algebraic property we need from the underlying group that is sufficient for the existence of a good basis change, and we prove that $\mathrm{PGL}(2,q)$ has this property. This results in a construction of a $(D = O(1/\overline{\lambda}^4), \overline{\lambda})$ quantum expander. We describe this construction in detail in Section 3.

This construction is not explicit in the sense that it uses the Fourier transform over $\mathrm{PGL}(2,q)$, which is not known to have an efficient implementation. (See [27] for a non-trivial, but still not fast enough, algorithm.) We mention that there are also explicit, constant degree (non-Ramanujan) Cayley expanders over the symmetric group S_n and the alternating group A_n [25]. Also, there is an efficient implementation of the Fourier transform over S_n [7]. We do not know, however, whether S_n (or A_n) respects our additional property.

Following the publication of this construction (given in [8]), Hastings [20] showed, using elegant techniques, that quantum expanders cannot be better than Ramanujan, i. e., cannot have spectral gap better than $1 - 2\sqrt{D-1}/D$. Hastings also showed that taking D random unitaries gives an almost-Ramanujan expander. This settles the parameters that can be achieved with a *non-explicit* construction. However, Hastings' work does not give an explicit construction, because a random unitary is a highly non-explicit object.

The second construction presented in this paper adapts the classical Zig-Zag construction [36] to the quantum world. The construction is iterative, starts with a good quantum expander of constant size (that is found with a brute force search), and then builds quantum expanders for larger spaces by repeatedly applying tensoring (which makes the space larger at the expense of the spectral gap), squaring (that improves the spectral gap at the expense of the degree) and a Zig-Zag operation that reduces the degree back to that of the constant-size expander. We again work by lifting the classical operators working over $\mathcal{V}$ to quantum operators working over $L(\mathcal{V})$, and we adapt the analysis along similar lines. The main

issue is generalizing and analyzing the Zig-Zag product. Remarkably, this translation works smoothly and gives the desired quantum expanders with almost the same proof applied over $L(\mathcal{V})$ rather than $\mathcal{V}$. The construction gives explicit, constant degree quantum expanders with a constant spectral gap. We describe this construction in detail in Section 4.

Two other explicit constructions of quantum expanders were published in [18] and [16] shortly after our work first appeared. In [16] it was shown how the expander of Margulis [29] can be twisted to the quantum setting, and in [18] it was shown how any classical Cayley expander can be converted to a quantum expander, provided the underlying group has an efficient quantum Fourier transform and a large irreducible representation. Applying this recipe to the Cayley expanders over S_n of [25] results in another construction of explicit, constant degree quantum expanders. One advantage of our explicit construction is that it achieves a much better relation between the spectral gap and the degree compared to that of the other explicit constructions [29, 18].

The Zig-Zag construction we describe in this paper gives a natural, iterative quantum expander with parameters that are as good as our first construction. However, the Zig-Zag construction is explicit whereas the first construction is not yet explicit (because we do not have an efficient implementation for the Fourier transform of $\mathrm{PGL}(2,q)$). We nevertheless decided to include the first construction in the paper. First, we believe it describes a natural approach, and this can be seen from the various other quantum expander constructions that are based on Cayley graphs. Also, the first construction is appealing in that it has only two stages, and each stage naturally corresponds to a well-known Cayley graph. Finally, and more importantly, in the classical setting there are algebraic constructions of Ramanujan expanders (as opposed to combinatorial constructions). Therefore, we believe our first construction has the potential of being improved to a construction of a quantum Ramanujan expander.

1.2 Applications of quantum expanders

Classical expanders have become well-known and fundamental objects in mathematics and computer science. This is due to the many applications these objects have found and to the intimate relations they have with other central notions in computational complexity. We refer the reader (again) to the survey paper of [22] for a partial list of applications.

While quantum expanders are a natural generalization of classical expanders, they have only recently been defined and it is yet to be seen whether they will be as useful as their classical counterparts. Thus far, the following short list of applications has been identified.

- Quantum one-time pads. Ambainis and Smith [6] implicitly used quantum expanders 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 [11] showed how to remove the $O(\log n)$ term.

- Hastings [19] gave an application from physics. Using quantum expanders, he showed that there exist gapped one-dimensional systems for which the entropy between a given subvolume and the rest of the system is exponential in the correlation length.

- Recently, Hastings and Harrow [21] used specialized quantum expanders (called tensor product expanders) to approximate t-designs as well as to attack a certain open question regarding the Solovay-Kitaev gate approximation.

- In this work we use the quantum expanders constructed by Ambainis and Smith [6] in order to show the problem Quantum Entropy Difference problem (QED) is QSZK-complete.

Let us now elaborate on the last application.

Watrous [42] defined the complexity 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 [42, 43]). 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 $\| \tau_{Q_0} - \tau_{Q_1} \|_{\mathrm{tr}} \leq \alpha$.

Here, 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).

In [42], 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 QSZK $\subseteq$ PSPACE. Subsequently, in [43], 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 promise 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 $S(\tau_{Q_1}) - S(\tau_{Q_0}) \geq \frac{1}{2}$.

Here, $S(\rho)$ is the von Neumann entropy of the mixed state ρ (see Section 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 more entropy. This problem is, in particular, as hard[2] as approximating the amount of entropy in a given mixed state (when again the mixed state is given by a circuit generating it).

We prove that QED is QSZK–complete. The proof follows the classical intuition, which uses classical expanders to convert high entropy states to the completely mixed state, while keeping low-entropy states entropy-deficient. Indeed, our proof is an adaptation of the classical proof to quantum entropies,

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

[2] Under Turing reductions.

but it crucially depends on the use of quantum expanders replacing the classical expanders used in the classical proof.

The proof requires an explicit quantum expander with a near-optimal *entropy loss* (see Section 5.1). As it turns out, the only expander that we currently know of that satisfies this property is the Ambainis-Smith expander. (Indeed it is of non-constant degree but this turns out to be irrelevant in this case.) Using it we obtain that QED is QSZK–complete.

This result 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 $\mathrm{Tr}_B(|\psi\rangle\langle\psi|)$ [35]. 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 under Turing reductions and hence unlikely to be in BQP.

The remainder of this 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 the Zig-Zag construction. Finally, Section 5 is devoted to proving the completeness of QED in QSZK.

2 Preliminaries

For any finite-dimensional Hilbert space $\mathcal{V}$, we write $L(\mathcal{V})$ to denote the set of linear operators over $\mathcal{V}$. The set $L(\mathcal{V})$ is also a Hilbert space, equipped with the inner-product $\langle A,B\rangle = \mathrm{Tr}(A^\dagger B)$ and the norm $\|A\|_2 = \sqrt{\langle A,A\rangle}$.

Let $P = (p_1,\ldots,p_m)$ be a vector with real values $p_i \geq 0$.

- The *Shannon entropy* is $H(P) = \sum_{i=1}^m p_i \log \frac{1}{p_i}$.
- The *min-entropy* is $H_\infty(P) = \min_i \log \frac{1}{p_i}$.
- The *Rényi entropy* is $H_2(P) = \log \frac{1}{\mathrm{Col}(P)}$, where $\mathrm{Col}(P) = \sum p_i^2$ is the collision probability of the distribution defined by $\mathrm{Col}(P) = \Pr_{x,y}[x=y]$ when x,y are sampled independently from P.

(We write $\log(\cdot)$ to denote the base 2 logarithm, and $\ln(\cdot)$ to denote the natural logarithm.)

We have analogous definitions for density matrices. For a density matrix ρ, let $\alpha = (\alpha_1,\ldots,\alpha_N)$ be its set of eigenvalues. Since ρ is a density matrix, all these eigenvalues are non-negative and 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 ρ, $H_\infty(\rho) \le H_2(\rho) \le S(\rho)$.

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 $X \in L(\mathcal{V})$ to be $\|X\|_{\mathrm{tr}} = \mathrm{Tr}(|X|)$, where $|X| = \sqrt{X^\dagger X}$, and by 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

The construction we present in this section constructs a quantum expander by first taking a step on a non-Abelian Cayley expander followed by a Fourier transform and another step on the non-Abelian Cayley expander. It is similar in spirit to the construction of good quantum error correcting codes given by first encoding the input word with a good classical code, then applying a Fourier transform and then encoding it again with a classical code. Technically the analysis here is more complicated because we use a Fourier transform over a non-Abelian group.

We begin this section with some necessary representation theory background. We then describe the construction and we conclude with its analysis.

3.1 Representation theory background

We survey some basic elements of representation theory. For complete accounts, consult the books of Serre [40] or Fulton and Harris [17].

A *representation* ρ of a finite group G is a homomorphism $\rho : G \to \mathrm{GL}(\mathcal{V})$, where $\mathcal{V}$ is a (finite-dimensional) vector space over $\mathbb{C}$ and $\mathrm{GL}(\mathcal{V})$ denotes the group of invertible linear operators on $\mathcal{V}$. Fixing a basis for $\mathcal{V}$, each $\rho(g)$ may be realized as a $d \times d$ matrix over $\mathbb{C}$, where d is the dimension of $\mathcal{V}$. As ρ is a homomorphism, for any $g, h \in G$, $\rho(gh) = \rho(g)\rho(h)$ (the second product being matrix multiplication). The *dimension* d_ρ of the representation ρ is d, the dimension of $\mathcal{V}$.

We say that two representations $\rho_1 : G \to \mathrm{GL}(\mathcal{V})$ and $\rho_2 : G \to \mathrm{GL}(\mathcal{W})$ of a group G are *isomorphic* when there is a linear isomorphism of the two vector spaces $\phi : \mathcal{V} \to \mathcal{W}$ so that for all $g \in G$, $\phi\rho_1(g) = \rho_2(g)\phi$. In this case, we write $\rho_1 \cong \rho_2$.

We say that a subspace $\mathcal{W} \subseteq \mathcal{V}$ is an *invariant subspace* of a representation $\rho : G \to \mathrm{GL}(\mathcal{V})$ if $\rho(g)\mathcal{W} \subseteq \mathcal{W}$ for all $g \in G$. The zero subspace and the subspace $\mathcal{V}$ are always invariant. If no nonzero proper subspaces are invariant, the representation is said to be *irreducible*. Up to isomorphism, a finite group has a finite number of irreducible representations; we let $\widehat{G}$ denote this collection of representations.

If $\rho : G \to \mathrm{GL}(\mathcal{V})$ is a representation, $\mathcal{V} = \mathcal{V}_1 \oplus \mathcal{V}_2$, and each $\mathcal{V}_i$ is an invariant subspace of ρ, then $\rho(g)$ defines two linear representations $\rho_i : G \to \mathrm{GL}(\mathcal{V}_i)$ such that $\rho(g) = \rho_1(g) + \rho_2(g)$. We then write $\rho = \rho_1 \oplus \rho_2$. Any representation ρ can be written as $\rho = \rho_1 \oplus \rho_2 \oplus \cdots \oplus \rho_k$, where each ρ_i is irreducible. In particular, there is a basis in which every matrix $\rho(g)$ is block diagonal, the ith block corresponding to the ith representation in the decomposition. While this decomposition is not, in general, unique, the

number of times a given irreducible representation appears in this decomposition (up to isomorphism) depends only on the original representation ρ.

The *group algebra* $\mathbb{C}[G]$ of a group G is a vector space of dimension $|G|$ over $\mathbb{C}$, with an orthonormal basis $\{|g\rangle \mid g \in G\}$ and multiplication

$$\sum a_g |g\rangle \cdot \sum b_{g'} |g'\rangle = \sum_{g,g'} a_g b_{g'} |g \cdot g'\rangle .$$

This algebra is in bijection with the set $\{f : G \to \mathbb{C}\}$ with the bijection being $f \to \sum_g f(g)|g\rangle$. The inner product in $\mathbb{C}[G]$ translates to the familiar inner product $\langle f,h\rangle = \sum_g \overline{f(g)}h(g)$. The *regular representation* $\rho_{\mathrm{reg}} : G \to \mathrm{GL}(\mathbb{C}[G])$ is defined by $\rho_{\mathrm{reg}}(s) : |g\rangle \mapsto |sg\rangle$, for any $g \in G$. Notice that $\rho_{\mathrm{reg}}(s)$ is a permutation matrix for any $s \in G$.

An interesting fact about the regular representation is that it contains every irreducible representation of G. In particular, if $\rho_1, \ldots, \rho_k$ are the irreducible representations of G with dimensions $d_{\rho_1}, \ldots, d_{\rho_k}$, then

$$\rho_{\mathrm{reg}} = d_{\rho_1}\rho_1 \oplus \cdots \oplus d_{\rho_k}\rho_k ,$$

that is, the regular representation contains each irreducible representation ρ exactly d_ρ times.

The *Fourier transform* over G is the unitary transformation F defined by:

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

where $\rho_{i,j}(g)$ is the (i,j)-th entry of $\rho(g)$ in some predefined basis. In general one has freedom in choosing a basis for each invariant subspace. In this paper we choose an *arbitrary* basis, and later fix this choice by using special properties of the group G.

Fact 3.1. The Fourier transform block-diagonalizes the regular representation, i. e.,

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

This means that when we represent $\rho_{\mathrm{reg}}(g)$ in the basis given by F, we get a block diagonal matrix, with an invariant subspace of dimension d_ρ for each $\rho \in \widehat{G}$, and with $\rho(g)$ as the values of that block.

3.2 The construction

Fix an arbitrary (Abelian or non-Abelian) group G of order N, and a subset Γ of group elements closed under inversion. The *Cayley graph* $C(G,\Gamma)$ associated with Γ is a graph over N vertices, each corresponding to an element of G. This graph contains an edge (g_1, g_2) if and only if $g_1 = g_2\gamma$ for some $\gamma \in \Gamma$. The graph $C(G,\Gamma)$ is a regular undirected graph of degree $|\Gamma|$.

We associate with the graph $C(G,\Gamma)$ the linear operator M over $\mathbb{C}[G]$ whose matrix representation agrees with the normalized adjacency matrix of $C(G,\Gamma)$, i. e.,

$$M = \frac{1}{|\Gamma|} \sum_{\gamma\in\Gamma, x\in G} |x\gamma\rangle\langle x| \,.^{3}$$

(The normalization is such that the operator norm is 1.)

Notice that M is a real and symmetric operator, and therefore diagonalizes with real eigenvalues. We denote by $\lambda_1 \geq \cdots \geq \lambda_N$ the eigenvalues of M with orthonormal eigenvectors $v_1,\ldots,v_N$. As $C(G,\Gamma)$ is regular, we have $\lambda_1 = 1$ and $\overline{\lambda} = \max_{i>1}|\lambda_i| \leq 1$.

We define the superoperator $T : L(\mathbb{C}[G]) \to L(\mathbb{C}[G])$ that corresponds to randomly taking one step on the Cayley graph $C(G,\Gamma)$. More precisely, this superoperator describes the process whereby a register R of dimension $|\Gamma|$ is initialized to $\left|\overline{0}\right\rangle$ and the following steps are taken. First, a transformation H is performed on R that maps $|0\rangle$ to

$$\frac{1}{\sqrt{|\Gamma|}} \sum_{\gamma\in\Gamma} |\gamma\rangle \,,$$

yielding, for an input state ρ, the state

$$\frac{1}{|\Gamma|}\rho \otimes \sum_{\gamma,\gamma'\in\Gamma} |\gamma\rangle\langle\gamma'| \,.$$

Then, the unitary transformation $Z : |g,\gamma\rangle \to |g\gamma,\gamma\rangle$ is applied, and finally the register R is discarded. In more algebraic terms,

$$T(\rho) = \mathrm{Tr}_R\left[Z(I\otimes H)(\rho\otimes\left|\overline{0}\right\rangle\!\left\langle\overline{0}\right|)(I\otimes H)Z^\dagger\right].$$

We note that the transformation Z is a permutation over the standard basis, and is classically easy to compute in both directions, and therefore has an efficient quantum circuit.

We also need the notion of a *good basis change*. We say a unitary transformation U is a good basis change if for any $g_1 \neq e$ (where e denotes the identity element of G) and any g_2 it holds that

$$\mathrm{Tr}\big(U\rho_{\mathrm{reg}}(g_1)U^\dagger\rho_{\mathrm{reg}}(g_2)\big) = 0\,. \tag{3.1}$$

The quantum expander is then defined as

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

[3] 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 if and only if $g_1 = \gamma g_2$. However, note that if we define the invertible linear transformation P that maps the basis vector $|g\rangle$ to the basis vector $\left|g^{-1}\right\rangle$, then $PMP^{-1} = PMP$ maps x to

$$\frac{1}{|\Gamma|}\sum_\gamma \left|(x^{-1}\gamma)^{-1}\right\rangle = \frac{1}{|\Gamma|}\sum_\gamma \left|\gamma^{-1}x\right\rangle = \frac{1}{|\Gamma|}\sum_\gamma |\gamma x\rangle$$

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.

Lemma 3.2. *If U is a good basis change then E is a $(|\Gamma|^2, \overline{\lambda})$ quantum expander for $\overline{\lambda}$ as defined as above.*

The fact that E is $|\Gamma|^2$-regular is immediate and the rest of this section is devoted to proving the claimed spectral gap.

Lubotzky et al. [28] described a constant degree Ramanujan Cayley graph over $\mathrm{PGL}(2,q)$, with degree $|\Gamma|$ and second-largest eigenvalue $\overline{\lambda}$ satisfying $\overline{\lambda}^2 \leq 4/|\Gamma|$. In Section 3.5 we show how to modify the Fourier transform for $\mathrm{PGL}(2,q)$ to obtain a good basis change, and by plugging this basis change into Lemma 3.2 we obtain a $(16/\overline{\lambda}^4, \overline{\lambda})$ quantum expander. The construction is not explicit as it is yet unknown how to efficiently implement the quantum Fourier transform for $\mathrm{PGL}(2,q)$.

3.3 The analysis

First, we fully identify the spectrum of T. We view any eigenvector $v_i \in \mathbb{C}^N$ (of M) as an element of $\mathbb{C}[G]$, $|v_i\rangle = \sum_g v_i(g)\,|g\rangle$. We also define a linear transformation $\mathsf{Diag} : \mathbb{C}[G] \to L(\mathbb{C}[G])$ by $\mathsf{Diag}\,|g\rangle = |g\rangle\langle g|$. Denote

$$\mu_{i,g} \;=\; \rho_{\mathrm{reg}}(g)(\mathsf{Diag}\,|v_i\rangle) \;=\; \sum_{x\in G} v_i(x)\,|gx\rangle\langle x|\,.$$

Then it is easy to see that these matrices form a set of eigenvectors of T.

Lemma 3.3. *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 λ_i.*

Proof. Notice that $T(|g_1\rangle\langle g_2|) = \mathrm{Tr}_R\big[\frac{1}{|\Gamma|}\sum_{\gamma_1,\gamma_2} Z\,|g_1,\gamma_1\rangle\langle g_2,\gamma_2|\,Z^\dagger\big] = \frac{1}{|\Gamma|}\sum_\gamma |g_1\gamma\rangle\langle g_2\gamma|$. Now,

$$\begin{aligned} T(\mu_{i,g}) &= \frac{1}{|\Gamma|}\sum_{x,\gamma} v_i(x)\,|gx\gamma\rangle\langle x\gamma| = \rho_{\mathrm{reg}}(g)\frac{1}{|\Gamma|}\sum_{x,\gamma} v_i(x)\,|x\gamma\rangle\langle x\gamma| \\ &= \rho_{\mathrm{reg}}(g)\mathsf{Diag}\Big(\sum_x v_i(x) M\,|x\rangle\Big) = \rho_{\mathrm{reg}}(g)\mathsf{Diag}(M\,|v_i\rangle) = \lambda_i \rho_{\mathrm{reg}}(g)\mathsf{Diag}(|v_i\rangle) = \lambda_i \mu_{i,g}\,. \end{aligned}$$

To verify orthonormality, notice that $\mathrm{Tr}(\mu_{i,g_1}\mu^\dagger_{i',g_2}) = 0$ for every choice of $g_1 \neq g_2$, as each entry (k,ℓ) must be zero for at least one of the matrices. If $g_1 = g_2 = g$ then

$$\mathrm{Tr}(\mu_{i,g}\mu^\dagger_{i',g}) = \langle v_{i'}|v_i\rangle = \delta_{i,i'}\,.$$

As the number of vectors $\{\mu_{i,g}\}$ is N^2, they form an orthonormal basis for $L(\mathbb{C}[G])$. □

We decompose the space $L(\mathbb{C}[G])$ into three perpendicular spaces:

$$\begin{array}{lcll} \mathrm{Span}\,\{\mu_{1,e}\}\,, & & & \\ W & = & \mathrm{Span}\,\{\mu_{1,g} \mid g \in G, g \neq e\}\,, & \text{and} \\ \mu^\perp & = & \mathrm{Span}\,\{\mu_{i,g} \mid i \neq 1, g \in G\}\,. & \end{array}$$

We also denote $\mu^{||} = \mathrm{Span}\,\{\mu_{1,e}\} + W = \mathrm{Span}\,\{\mu_{1,g} \mid g \in G\}$. Notice that $T(\mu^{||}) = \mu^{||}$ and $T(\mu^\perp) = \mu^\perp$.

Claim 3.4. *If $\rho \in W$ and U is a good basis change then $U\rho U^\dagger \in \mu^\perp$.*

Proof. The set $\{\rho_{\text{reg}}(g) \mid g \in G\}$ is an orthonormal basis for $\mu^{||}$ and hence $\{\rho_{\text{reg}}(g) \mid g \in G,\ g \neq e\}$ is an orthonormal basis for W. Therefore, it is enough to verify that $\text{Tr}(U\rho_{\text{reg}}(g_1)U^\dagger \rho_{\text{reg}}(g_2)^\dagger) = 0$ for any $g_1 \neq e$ and for any g_2. Given that $\rho_{\text{reg}}(g_2)^\dagger = \rho_{\text{reg}}(g_2^{-1})$, this follows directly from (3.1). □

Thus, intuitively speaking, we have a win-win situation when E is applied to ρ. If ρ is in $\mu^\perp$, then the first application of T shrinks its norm, while if ρ is in W, then the first application of T keeps it unchanged, the basis change maps it to $\mu^\perp$, and the last T application shrinks its norm. Indeed, we are now ready to prove Lemma 3.2, namely, that if U is a good basis change then E is a $(|\Gamma|^2, \overline{\lambda})$ quantum expander.

Proof of Lemma 3.2. The regularity of E is clear from its definition. Fix any $X \in L(\mathbb{C}[G])$ that is perpendicular to $\tilde{I} = \mu_{1,e}$, and write $X = X^{||} + X^\perp$ for $X^{||} \in W$ and $X^\perp \in \mu^\perp$. We have

$$E(X) = T(\sigma^{||} + \sigma^\perp),$$

where $\sigma^{||} = UT(X^{||})U^\dagger$ and $\sigma^\perp = UT(X^\perp)U^\dagger$. Observe the following. First, $T(X^{||}) \in W$, so by Claim 3.4, $\sigma^{||} \perp \mu^{||}$. Also, $T(X^{||}) \perp T(X^\perp)$ (as T preserves both $\mu^{||}$ and $\mu^\perp$), and therefore $\sigma^{||} \perp \sigma^\perp$. Moreover, by Lemma 3.3 we know T is normal.

By Lemma 3.5 (stated and proved below) we see that

$$\begin{aligned} \|E(X)\|_2^2 &= \|T(\sigma^{||} + \sigma^\perp)\|_2^2 \leq \overline{\lambda}^2 \|\sigma^{||}\|_2^2 + \|\sigma^\perp\|_2^2 \\ &= \overline{\lambda}^2 \|T(X^{||})\|_2^2 + \|T(X^\perp)\|_2^2 \leq \overline{\lambda}^2 \|X^{||}\|_2^2 + \overline{\lambda}^2 \|X^\perp\|_2^2 = \overline{\lambda}^2 \|X\|_2^2 \end{aligned}$$

as required. □

We are left to prove the following lemma.

Lemma 3.5. *Let T be a normal linear operator with eigenspaces $\mathcal{V}_1, \ldots, \mathcal{V}_n$ and corresponding eigenvalues $\lambda_1, \ldots, \lambda_n$ in descending* absolute *value. Suppose u and w are vectors such that*

$$u \in \text{Span}\{\mathcal{V}_2, \ldots, \mathcal{V}_n\}$$

and $w \perp u$ (and where w does not necessarily belong to $\mathcal{V}_1$). Then

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

Proof. Let $\{v_j\}$ be an eigenvector basis for T with eigenvalues δ_j (from the set $\{\lambda_1, \ldots, \lambda_n\}$). Writing $u = \sum_j \alpha_j v_j$ and $w = \beta v + \sum_j \beta_j v_j$ with $v_j \in \text{Span}\{\mathcal{V}_2, \ldots, \mathcal{V}_n\}$ and $v \in \mathcal{V}_1$, we get:

$$\begin{aligned} \|T(u+w)\|_2^2 &= \left\| \lambda_1 \beta v + \sum_j \delta_j(\alpha_j + \beta_j) v_j \right\|_2^2 \leq |\lambda_1|^2 |\beta|^2 + |\lambda_2|^2 \sum_j |\alpha_j + \beta_j|^2 \\ &= |\lambda_1|^2 |\beta|^2 + |\lambda_2|^2 \Big(\sum_j |\alpha_j|^2 + \sum_j |\beta_j|^2 + \langle u|w\rangle + \langle w|u\rangle \Big) \leq |\lambda_2|^2 \|u\|_2^2 + |\lambda_1|^2 \|w\|_2^2. \end{aligned}$$

□

3.4 A sufficient condition that guarantees a good basis change

So far we have reduced the problem of constructing a quantum expander to that of finding a Cayley graph $C(G,\Gamma)$ and a good basis change for G. We now concentrate on the problem of finding a good basis change for a given group G, and show that if G respects some general condition then one can efficiently construct a good basis change from G from its Fourier transform.

A basic fact of representation theory states that $\sum_{\rho\in\widehat{G}} d_\rho^2 = |G|$. Equivalently, for any group G there is a bijection between

$$\left\{(\rho,i,j)\,\middle|\,\rho\in\widehat{G}, 1\le i,j\le d_\rho\right\}$$

and G. Finding such a natural bijection is a fundamental problem both in mathematics (where it is equivalent to describing the invariant subspaces of the regular representation of G) and in computer science (where it is a main step towards implementing a fast Fourier transform). Indeed, this question was extensively studied. For example, the "Robinson-Schensted" algorithm [37, 39] is a mapping from pairs (P,T) of standard shapes (a shape corresponds to an irreducible representation of S_n, and its dimension is the number of valid fillings of that shape) to S_n.

Here we require more from such a mapping.

Definition 3.6. 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 $\rho\in\widehat{G}$,

$$f(\rho,i,j) \quad = \quad f_1(\rho,i)\cdot f_2(\rho,j) \tag{3.2}$$

for some choice of functions $f_1(\rho,\cdot), f_2(\rho,\cdot) : [d_\rho]\to G$.

The Robinson-Schensted mapping is *not* a product mapping. However, S_n has a product mapping for $n\le 6$, and we think it is a natural question whether product mappings for S_n exist for all n. For some groups it is easy to find a product mapping. For example, in any Abelian group all irreducible representations are of dimension one and so we can define $f_1(\rho,i)=e$ and $f_2(\rho,j)=f(\rho,1,1)$.

Another easy example is the dihedral group D_m of rotations and reflections of a regular polygon with m sides. Its generators are r, the rotation element, and s, the reflection element. This group has $2m$ elements and the defining relations are $s^2=1$ and $srs=r^{-1}$. We shall argue this group has a product mapping for odd m (although it is true for even m as well). The dihedral group has $(m-1)/2$ representations $\{\rho_\ell\}$ of dimension two and two representations $\{\tau_1,\tau_2\}$ of dimension one (see [40, Section 5.3]). A product mapping in this case can be given by defining $f(\rho,i,j)$ as follows:

$$f(\rho,i,j) \quad = \quad \begin{cases} 1 & \text{if } \rho=\tau_1, i=j=1,\\ s & \text{if } \rho=\tau_2, i=j=1,\\ r^{2(\ell-1)+i}s^j & \text{if } \rho=\rho_\ell. \end{cases} \tag{3.3}$$

We now show that if G has a product mapping then G has a good basis change:

Lemma 3.7. *Let G be a group that has a product mapping f, and let F be the Fourier transform over G, that is*

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

Define the unitary mapping

$$S : |\rho, i, j\rangle \mapsto \omega_{d_\rho}^{ij} |f(\rho, i, j)\rangle ,$$

where ω_{d_ρ} is a primitive root of unity of order d_ρ, and set U to be the unitary transformation $U = SF$. Then U is a good basis change.

Proof. Fix $g_1 \neq e$ and g_2. If $g_2 = e$ then

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

where the last equality follows from the assumption that $g_1 \neq e$.

We are left with the case $g_2 \neq e$. By Fact 3.1, it holds that

$$\mathrm{Tr}\left(SF\rho_{\mathrm{reg}}(g_1)F^\dagger S^\dagger \rho_{\mathrm{reg}}(g_2)\right) = \sum_{\rho\in\widehat{G}} \sum_{1\le i,i'\le d_\rho} \rho_{i,i'}(g_1)\, \mathrm{Tr}\left(\sum_{j=1}^{d_\rho} S|\rho,i,j\rangle\langle\rho,i',j| S^\dagger \sum_x |g_2 x\rangle\langle x|\right) .$$

Therefore, it suffices to show that for any ρ, i, i' we have

$$\mathrm{Tr}\left(\sum_{j=1}^{d_\rho} S|\rho,i,j\rangle\langle\rho,i',j| S^\dagger \sum_x |g_2 x\rangle\langle x|\right) = 0 .$$

Fix $\rho \in \widehat{G}$ and $i, i' \in \{1, \ldots, d_\rho\}$. Because f is a product mapping, $f(\rho, i, j) = f_1(\rho, i) \cdot f_2(\rho, j)$ for some choice of functions f_1, f_2. Denote $h_i = f_1(\rho, i)$ and $t_j = f_2(\rho, j)$. The sum we need to calculate can be written as:

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

where the last equality follows from the observation that the sum over x yields a non-zero value if and only if $x = h_i t_j$ and $h_{i'} t_j = g_2 x$. This happens if and only if $h_i t_j = g_2^{-1} h_{i'} t_j$, or equivalently $g_2 = h_{i'} h_i^{-1}$. However, when $g_2 = h_{i'} h_i^{-1}$, we obtain the sum $\sum_{j=1}^{d_\rho} \omega_{d_\rho}^{(i-i')j}$, and because $g_2 \neq e$ it follows that $i \neq i'$. Hence the expression is zero, as required. □

3.5 PGL$(2, q)$ has a product bijection

The group $\mathrm{PGL}(2, q)$ is the group of all 2×2 invertible matrices over $\mathbb{F}_q$ modulo the group center. This group has $(q-3)/2$ irreducible representations of dimension $q+1$, $(q-1)/2$ irreducible representations of dimension $q-1$, 2 irreducible representations of dimension q and 2 irreducible representations of dimension 1 (see [17, Section 5.2] and [1]). We let ρ_x^d denote the xth irreducible representation of dimension d.

We look for a bijection from G to the irreducible representations of G. Our approach is to use a tower of subgroups,

$$G_3 = G > G_2 = D_{2q} > G_1 = Z_q > G_0 = \{e\},$$

with G_2 and G_1 defined as follows. The group G_2 is generated by the equivalence classes of

$$\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \qquad \text{and of} \qquad \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix}.$$

This group is a dihedral subgroup of G with $2q$ elements, i. e., D_q. The first matrix is the reflection, denoted by s, and the second is the rotation, denoted by r. This group has a cyclic subgroup $G_1 \cong Z_q$ (the group generated by r).

Let $T_2 = \{t_1, \ldots, t_\ell\}$ be a transversal for G_2 with

$$\ell = \frac{|G|}{|G_2|} = \frac{(q-1)(q+1)}{2}.$$

For each $\rho \in \widehat{G}$ we let $f_1(\rho, i) \in \{t_1, \ldots, t_\ell\}$ define a coset of G_2, and let $f_2(\rho, j) \in G_2$ define an element in G_2 as follows. The representations of dimension $q+1$ take the first $(q-3)(q+1)/2$ cosets:

$$f_1(\rho_x^{q+1}, i) = \begin{cases} r^{i-1} & \text{if } i = 1, \ldots, q, \\ s & \text{if } i = q+1, \end{cases}$$
$$f_2(\rho_x^{q+1}, j) = t_{(x-1)(q+1)+j},$$

for all $x = 1, \ldots, \frac{q-1}{2} - 1$ and $i, j = 1, \ldots, q+1$. We match them with representations of dimension $q-1$:

$$f_1(\rho_x^{q-1}, i) = sr^i,$$
$$f_2(\rho_x^{q-1}, j) = t_{(x-1)(q-1)+j},$$

for all $x = 1, \ldots, (q-1)/2$ and $i, j = 1, \ldots, q-1$. Notice that so far we have covered the first

$$\frac{(q-3)(q+1)}{2} = \frac{(q-1)(q-1)}{2} - 2$$

cosets without repetitions. Two cosets are partially covered with dimension $q-1$ representations (in each coset $q-1$ elements are covered). We put the dimension 1 representation into these cosets:

$$f_1(\rho_x^1, 1) = s,$$
$$f_2(\rho_x^1, 1) = t_{\frac{(q-3)(q+1)}{2}+x},$$

for $x = 1, 2$. Finally, we fill all the remaining gaps with dimension q representations. The first two fill the partially full cosets, and the rest fill each coset in pairs. Notice that here we use the fact that $G_1 < G_2$. The function f_2 returns an element in the traversal set of G_1 and f_1 returns an element of G_1:

$$f_1(\rho_x^q, i) = r^i,$$
$$f_2(\rho_x^q, j) = \begin{cases} t_{\frac{(q-3)(q+1)}{2}+x} & \text{if } j = q, \\ s^{x-1} t_{\frac{(q-1)(q-1)}{2}+j} & \text{otherwise,} \end{cases}$$

for $x = 1, 2$. One can verify that this product mapping is a bijection as desired.

4 The Zig-Zag construction

We now present our second construction of quantum expanders, following the iterative construction of Reingold et al. [36]. Their starting point is a good expander of constant size, which can be found by an exhaustive search. Then, they construct a series of expanders with an increasing number of vertices by applying a sequence of three basic transformations: tensoring (that squares the number of vertices at the expense of a worse ratio between the spectral gap and the degree), squaring (that improves the spectral gap) and the Zig-Zag product (that reduces the degree to its original size). These three transformations are repeated iteratively, resulting in a good constant-degree expander over many vertices.

The first two transformations have natural counterparts in the quantum setting. For ease of notation, we denote by $T(\mathcal{V})$ the set of superoperators on $L(\mathcal{V})$ (that is, $T(\mathcal{V}) = L(L(\mathcal{V}))$). We also denote by $U(\mathcal{V})$ the set of unitary operators in $L(\mathcal{V})$.

- **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 $X \in L(\mathcal{V}_1), Y \in L(\mathcal{V}_2)$.

In order to define the quantum Zig-Zag product we first recall the classical Zig-Zag product. We have two *graphs* G_1 and G_2. The graph G_1 is a D_1-regular graph over N_1 vertices and the graph G_2 is a D_2-regular graph over $N_2 = D_1$ vertices. We first define the *replacement product* graph, which has $V_1 \times V_2$ as its set of vertices. We refer to the set of vertices $\{v\} \times V_2$ as the *cloud* of v. The replacement product has a copy of G_2 on each cloud, and also *inter-cloud* edges between (v,i) and (w,j) if the i-th neighbor of v is w and the j-th neighbor of w is v in G_1. Thus, the replacement product has the same connected components as the original graph but a much lower degree ($D_2 + 1$ instead of D_1). The Zig-Zag product graph $G_1 \textcircled{z} G_2$ has the same set of vertices as the replacement product, but has an edge between $x = (v,a)$ and $x' = (v',a')$ if and only if in the replacement product graph there is a three step walk from x to x' that first takes a cloud edge, then an inter-cloud edge, and then again a cloud edge. Thus, the graph $G_1 \textcircled{z} G_2$ is D_2^2-regular.

We now define the quantum Zig-Zag transformation. Let $G_1 \in T(\mathcal{H}_{N_1})$ be an N_2-regular operator and $G_2 \in T(\mathcal{H}_{N_2})$, where $\mathcal{H}_N$ denotes the Hilbert space of dimension N. As G_1 is D_1-regular, 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 $\dot{U}(|a\rangle \otimes |b\rangle) = U_b |a\rangle \otimes |b\rangle$. We also define $\dot{G}_1 \in T(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ by $\dot{G}_1(X) = \dot{U} X \dot{U}^\dagger$. The superoperator $\dot{G}_1$ corresponds to the inter-cloud edges in the replacement product. We are now ready to define the quantum Zig-Zag product.

Definition 4.1. Let G_1, G_2 be as above. The Zig-Zag product of G_1 and G_2, denoted by $G_1 \textcircled{z} G_2 \in T(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$, is defined to be $(G_1 \textcircled{z} G_2)X = (I \otimes G_2^\dagger) \dot{G}_1 (I \otimes G_2) X$.

Remark 4.2. Notice that formally $G_1 \textcircled{z} G_2$ depends on the Kraus decomposition of G_1 and the notation should have reflected this. However, we fix this decomposition once and use this simpler notation.

Finally we explain how to find a base quantum operator H that is a (D^8, D, λ) quantum expander. Its existence follows from the following result of Hastings [20].

Theorem 4.3. *There exists an integer D_0 such that for every $D > D_0$ there exists a (D^8, D, λ) quantum expander for $\lambda = 4\sqrt{D-1}/D$.*

Remark 4.4. Hastings actually shows the stronger result that, for any D, there exist a

$$\left(D^8, D, \left(1 + O(D^{-16/15} \log D)\right) \frac{2\sqrt{D-1}}{D}\right)$$

quantum expander.

We use an exhaustive search over a net $S \subset U(\mathcal{H}_{D^8})$ of unitary matrices to find such a quantum expander. The set 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.[4] Suppose that

$$G(X) = \frac{1}{D} \sum_{i=1}^{D} U_i X U_i^\dagger \,.$$

is a (D^8, D, λ) quantum expander, and denote by G' the superoperator

$$G'(X) = \frac{1}{D} \sum_{i=1}^{D} V_{U_i} X V_{U_i}^\dagger \,.$$

For $X \in L(\mathcal{H}_{D^8})$ orthogonal to $\tilde{I}$, it holds that

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

Hence, G' is a $(D^8, D, 8\sqrt{D-1}/D)$ quantum expander. This implies that a brute force search over the net finds a good base superoperator H in time that depends only on D and λ.

Remark 4.5. We can actually get an eigenvalue bound of $(1+\varepsilon)2\sqrt{D-1}/D$ for an arbitrary small ε at the expense of increasing D_0, using the better bound in Remark 4.4.

[4] One way to see this is using the Solovay-Kitaev theorem (see, e. g., [10]). The theorem assures us that, for example, the set of all the quantum circuits of length $O(\log^4 \varepsilon^{-1})$ generated only by Hadamard and Tofolli gates gives 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.

Given all these ingredients we define an iterative process as in [36], composed of a series of superoperators. The first two superoperators are $G_1 = H^2$ and $G_2 = H \otimes H$. For every $t > 2$ we define

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

Theorem 4.6. *For every $t > 0$, G_t is an explicit (D^{8t}, D^2, λ_t) quantum expander with $\lambda_t = \lambda + O(\lambda^2)$, where the constant in the O notation is an absolute constant.*

Thus, G_t is a constant degree, constant gap quantum expander, as desired.

4.1 The analysis

Tensoring and squaring are easy to analyze, and the following proposition is immediate from the definitions of these operations.

Proposition 4.7. *If G is a (N, D, λ) quantum expander then G^2 is a (N, D^2, λ^2) quantum expander. 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.*

We are left to analyze is the quantum Zig-Zag product.

Theorem 4.8. *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 $G_1 \textcircled{z} G_2$ is a $(N_1 \cdot D_1, D_2^2, \lambda_1 + \lambda_2 + \lambda_2^2)$ quantum expander.*

With the above two claims, the proof of Theorem 4.6 is identical to the one in [36] and is omitted. In order to prove Theorem 4.8 we claim the following.

Proposition 4.9. *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*

$$\left| \langle (G_1 \textcircled{z} G_2) X, Y \rangle \right| \leq f(\lambda_1, \lambda_2) \, \| X \| \cdot \| Y \| ,$$

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

Theorem 4.8 follows from this proposition: for a given X orthogonal to $\tilde{I}$ we let $Y = (G_1 \textcircled{z} G_2)X$ and plug X and Y into the proposition. We see that $\| (G_1 \textcircled{z} G_2)X \| \leq f(\lambda_1, \lambda_2) \| X \|$ as required.

The proof of Proposition 4.9 is an adaptation of the proof in [36]. The main difference is that the classical proof works over the Hilbert space $\mathcal{V}$ whereas the quantum proof works over $L(\mathcal{V})$. Remarkably, the same intuition works in both cases.

Proof of Proposition 4.9. We first decompose the space $L(\mathcal{H}_{N_1} \otimes \mathcal{H}_{D_1})$ into

$$\begin{aligned} W^{\|} &= \text{Span}\{ \sigma \otimes \widetilde{I} \mid \sigma \in L(\mathcal{H}_{N_1}) \} \quad \text{and} \\ W^{\perp} &= \text{Span}\{ \sigma \otimes \tau \mid \sigma \in L(\mathcal{H}_{N_1}), \tau \in L(\mathcal{H}_{D_1}), \langle \tau, \widetilde{I} \rangle = 0 \} . \end{aligned}$$

Next, we write X as $X = X^{\|} + X^{\perp}$, where $X^{\|} \in W^{\|}$ and $X^{\perp} \in W^{\perp}$, and similarly $Y = Y^{\|} + Y^{\perp}$. By definition,

$$|\langle (G_1 \textcircled{z} G_2)X, Y \rangle| = |\langle \dot{G}_1 (I \otimes G_2)(X^{\|} + X^{\perp}), (I \otimes G_2)(Y^{\|} + Y^{\perp}) \rangle| .$$

Using linearity and the triangle inequality (and the fact that $I \otimes G_2$ acts trivially on $W^{||}$), we get

$$\begin{aligned} |\langle (G_1 \textcircled{z} G_2)X, Y\rangle| \quad \leq \quad & |\langle \dot{G}_1 X^{||}, Y^{||}\rangle| + |\langle \dot{G}_1 X^{||}, (I\otimes G_2)Y^{\perp}\rangle| + \\ & |\langle \dot{G}_1 (I\otimes G_2)X^{\perp}, Y^{||}\rangle| + |\langle \dot{G}_1 (I\otimes G_2)X^{\perp}, (I\otimes G_2)Y^{\perp}\rangle| . \end{aligned}$$

In the last three terms we have $I \otimes G_2$ acting on an operator from $W^{\perp}$. As expected, when this happens the quantum expander G_2 shrinks the norm of the operator.

Claim 4.10. *For any $Z \in W^{\perp}$ it holds that $\| (I\otimes G_2)Z \| \leq \lambda_2 \| Z \|$.*

Proof. The matrix Z can be written as $Z = \sum_i \sigma_i \otimes \tau_i$, where each τ_i is perpendicular to $\tilde{I}$ and $\{\sigma_i\}$ is an orthogonal set. Hence,

$$\| (I\otimes G_2)Z \|^2 = \left\| \sum_i \sigma_i \otimes G_2(\tau_i) \right\|^2 = \sum_i \| \sigma_i \otimes G_2(\tau_i) \|^2 \leq \sum_i \lambda_2^2 \| \sigma_i \otimes \tau_i \|^2 = \lambda_2^2 \| Z \|^2 .$$

□

To bound the first term, we observe that on inputs from $W^{||}$ the operator $\dot{G}_1$ mimics the operation of G_1 with a random seed.

Claim 4.11. *For any $A, B \in W^{||}$ such that $\langle A, \tilde{I}\rangle = 0$, it holds that $|\langle \dot{G}_1(A), B\rangle| \leq \lambda_1 \| A \| \cdot \| B \|$.*

Proof. Any choice of $A, B \in W^{||}$ may be written as

$$A = \sigma \otimes \tilde{I} = \frac{1}{D_1} \sum_i \sigma \otimes |i\rangle\langle i| ,$$
$$B = \eta \otimes \tilde{I} = \frac{1}{D_1} \sum_i \eta \otimes |i\rangle\langle i| .$$

Moreover, as A is perpendicular to the identity operator, it follows that σ is perpendicular to the identity operator on the space $L(\mathcal{H}_{N_1})$. This means that applying G_1 on σ will shrinks its norm by at a factor of least λ_1.

Considering the inner product

$$\begin{aligned}
|\langle \dot{G_1}A,B\rangle| &= \frac{1}{D_1^2}\left|\sum_{i,j}\mathrm{Tr}\left(\left((U_i\sigma U_i^\dagger)\otimes|i\rangle\langle i|\right)(\eta\otimes|j\rangle\langle j|)^\dagger\right)\right| \\
&= \frac{1}{D_1^2}\left|\sum_{i,j}\mathrm{Tr}\left((U_i\sigma U_i^\dagger\eta^\dagger)\otimes|i\rangle\langle i\,|\,j\rangle\langle j|\right)\right| \\
&= \frac{1}{D_1^2}\left|\sum_{i}\mathrm{Tr}\left((U_i\sigma U_i^\dagger\eta^\dagger)\otimes|i\rangle\langle i|\right)\right| \\
&= \frac{1}{D_1^2}\left|\sum_{i}\mathrm{Tr}\left(U_i\sigma U_i^\dagger\eta^\dagger\right)\right| \\
&= \frac{1}{D_1}\left|\mathrm{Tr}\left(\left(\frac{1}{D_1}\sum_i U_i\sigma U_i^\dagger\right)\eta^\dagger\right)\right| \\
&= \frac{1}{D_1}|\langle G_1(\sigma),\eta\rangle| \le \frac{\lambda_1}{D_1}\|\sigma\|\cdot\|\eta\| = \lambda_1\|A\|\cdot\|B\|\,,
\end{aligned}$$

where the inequality follows from the expansion property of G_1 (and Cauchy-Schwartz). □

With the above claims in hand we see that

$$|\langle(G_1 ⓩ G_2)X,Y\rangle| \quad\le\quad (p_Xp_Y\lambda_1 + p_Xq_Y\lambda_2 + p_Yq_X\lambda_2 + q_Xq_Y\lambda_2^2)\,\|X\|\cdot\|Y\|\,, \tag{4.1}$$

where

$$p_X = \frac{\left\|X^{||}\right\|}{\|X\|} \qquad \text{and} \qquad q_X = \frac{\left\|X^{\perp}\right\|}{\|X\|}\,,$$

and similarly

$$p_Y = \frac{\left\|Y^{||}\right\|}{\|Y\|} \qquad \text{and} \qquad q_Y = \frac{\left\|X^{\perp}\right\|}{\|Y\|}\,.$$

Notice that $p_X^2+q_X^2 = p_Y^2+q_Y^2 = 1$. It is easy to see that $p_Xp_Y, q_Xq_Y \le 1$. Also, by Cauchy-Schwartz, $p_Xq_Y + p_Yq_X \le 1$. Therefore, the right hand side of equation (4.1) is upper bounded by the quantity $f(\lambda_1,\lambda_2)\,\|X\|\cdot\|Y\|$. □

4.2 Explicitness

Recall that a D-regular superoperator

$$E(X) = \frac{1}{D}\sum_i U_iXU_i^\dagger$$

is said to be *explicit* if it can be implemented by an efficient quantum circuit. Now we need a slight refinement of this definition: we say that E is *label-explicit* if each U_i has an efficient implementation. It can be checked that the squaring, tensoring and Zig-Zag operations map label-explicit transformations to label-explicit transformations. Also, our base superoperator is label-explicit (since it is defined over a constant size space). Therefore, the construction is label-explicit (and therefore explicit).

5 The complexity of estimating entropy

In this section we show that the language QED is QSZK–complete. The proof that QSD $\leq$ QED is standard, and is described in Subsection 5.4.

The more challenging direction is the proof that QED is in QSZK, or equivalently that QED $\leq$ QSD. In the classical setting this reduction is proved using extractors. Some parts of our proof of this reduction, for the quantum setting, are also standard. We define the problem QEA (Quantum Entropy Approximation) as follows:

Input: Quantum circuit Q, $t \geq 0$.
Accept: If $S(\tau_Q) \geq t + \frac{1}{2}$.
Reject: If $S(\tau_Q) \leq t - \frac{1}{2}$.

QEA is the problem of comparing the entropy of a given quantum circuit to some *known* threshold t (whereas QED compares two quantum circuits with unknown entropies). One immediately sees that

$$\mathrm{QED}(Q_0,Q_1) = \bigvee_{t=1}^{\max\{\mathsf{out}_1,\mathsf{out}_2\}} [((Q_0,t) \in \mathrm{QEA}_Y) \wedge ((Q_1,t) \in \mathrm{QEA}_N)] \,,$$

where out_i is the number of output qubits of Q_i.

A standard classical reduction can be easily adapted to the quantum setting to show that QEA $\in$ QSZK implies that QED $\in$ QSZK. We describe this part in Section 6. Thus, it is sufficient to prove that QEA $\in$ QSZK. We now focus on this part and the use of quantum expanders in the proof.

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 of the form $E : \{0,1\}^n \times \{0,1\}^d \to \{0,1\}^m$, and we say that such a function is a (k,ε) extractor if, for every distribution X on $\{0,1\}^n$ that has min-entropy k, the distribution $E(X,U_d)$ 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 classical circuit C and we want to distinguish between the cases where the distribution it defines has substantially more than t entropy or substantially 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 to the n output bits of C, hashing it to about t output bits. If the distribution C defines 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 entropy is less than $t-d-1$, where d is the extractor's seed length, then even after applying the extractor the output distribution has at most $t-1$ bits of entropy, and therefore it must be "far away" from uniform. Hence, we get a reduction to $\overline{\mathrm{SD}}$.

There are, of course, a few gaps to fill in. First, the distribution C defines is not necessarily flat. This is solved in the classical case by taking many independent copies of the circuit C, which makes the output distribution "close" to "nearly-flat." A simple analysis shows that this flattening works also in the quantum setting (this is Lemma 5.6). Also, we need to amplify the gap we have between $t+1/2$ and $t-1/2$ to a gap larger than d (the seed length). This, again, is solved by taking many independent copies of C, given that $S(C^{\otimes q}) = qS(C)$.

This section is organized as follows. We first discuss quantum extractors. We then prove the quantum flattening lemma, and prove that QEA $\leq \overline{\text{QSD}}$ through the use of quantum extractors. Together with the closure of QSZK under Boolean formulas, which is proved in Section 6, we have that QED $\in$ QSZK. We conclude this section with a proof that QSD $\leq$ QED, using a simple quantum adaptation of the classical proof.

5.1 Quantum extractors

Definition 5.1. A superoperator $T : L(\mathcal{H}_N) \to L(\mathcal{H}_N)$ is a (k,d,ε) *quantum extractor* if:

- The superoperator T is 2^d-regular.
- For every density matrix $\rho \in L(\mathcal{H}_N)$ with $H_\infty(\rho) \geq k$, it holds that $\left\| T(\rho) - \tilde{I} \right\|_{\text{tr}} \leq \varepsilon$.

We say T is explicit if T can be implemented by a quantum circuit of size polynomial in $\log(N)$. The *entropy loss* of T is $k + d - \log(N)$.

In the classical world balanced extractors are closely related to expanders (see, e. g., [15]). This generalizes to the quantum setting, as we now prove.

Lemma 5.2. *Suppose $T : L(\mathcal{H}_N) \to L(\mathcal{H}_N)$ is a $(N = 2^n, D = 2^d, \overline{\lambda})$ quantum expander. Then for every $t > 0$, T is also a $(k = n - t, d, \varepsilon)$ quantum extractor with $\varepsilon = 2^{t/2} \cdot \overline{\lambda}$. The entropy loss of T is $k + d - n = d - t$.*

Proof. The superoperator T has a one-dimensional eigenspace W_1 with eigenvalue 1, spanned by the unit eigenvector $v_1 = \frac{1}{\sqrt{N}} I$. Our input ρ is a density matrix, and therefore

$$\langle \rho, v_1 \rangle = \frac{1}{\sqrt{N}} \operatorname{Tr}(\rho) = \frac{1}{\sqrt{N}}.$$

In particular, $\rho - \tilde{I} = \rho - \frac{1}{\sqrt{N}} v_1$ is perpendicular to W_1. It follows that

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

where we have used

$$\|\rho - \tilde{I}\|_2^2 = \|\rho\|_2^2 - 2\operatorname{Tr}(\tilde{I}\rho) + \|\tilde{I}\|_2^2 = \|\rho\|_2^2 - \frac{1}{N} \leq \|\rho\|_2^2.$$

Given that $H_2(\rho) \geq H_\infty(\rho) \geq k = n - t$ we see that $\|T(\rho) - \tilde{I}\|_2^2 \leq \overline{\lambda}^2 2^{-(n-t)}$. By the Cauchy-Schwartz inequality, it follows that

$$\left\| T(\rho) - \tilde{I} \right\|_{\text{tr}} \leq \sqrt{N} \|T(\rho) - \tilde{I}\|_2 \leq \varepsilon,$$

which completes the proof. □

Corollary 5.3. *For every $n,t,\varepsilon \geq 0$ there exists an explicit $(n-t,d,\varepsilon)$ quantum extractor $T : L(\mathcal{H}_{2^n}) \to L(\mathcal{H}_{2^n})$, where*

1. *$d = 2(t+2\log\left(\frac{1}{\varepsilon}\right))+O(1)$ and the entropy loss is $t+4\log\left(\frac{1}{\varepsilon}\right)+O(1)$, or*
2. *$d = t+2\log\left(\frac{1}{\varepsilon}\right)+2\log(n)+O(1)$ and the entropy loss is $2\log(n)+2\log\left(\frac{1}{\varepsilon}\right)+O(1)$.*

The first bound on d is achieved using the Zig-Zag quantum expander of Theorem 4.6, and the second bound is achieved using the explicit construction of Ambainis and Smith [6] cited in Theorem 1.3.

One natural generalization of Definition 5.1 is to superoperators of the form $T : L(\mathcal{H}_N) \to L(\mathcal{H}_M)$ where $N = 2^n$ is not necessarily equal to $M = 2^m$. That is, such a superoperator T may map a large Hilbert space $\mathcal{H}_N$ to a much smaller Hilbert space $\mathcal{H}_M$. In the classical case this corresponds to hashing a large universe $\{0,1\}^n$ to a much smaller universe $\{0,1\}^m$. We suspect that unlike the classical case, no non-trivial unbalanced quantum extractors exist when $M < N/2$. Specifically, we suspect that all (k,d,ε) quantum extractors $T : L(\mathcal{H}_N) \to L(\mathcal{H}_M)$ with $k = n-1$ and $d < n-1$ must have error ε close to 1.

5.2 A flattening lemma

We first recall the classical flattening lemma that appears, e. g., in [41, Section 3.4.3].

Lemma 5.4. *Let $\lambda = (\lambda_1,\ldots,\lambda_M)$ be a distribution, let q be a positive integer, and let $\otimes^q\lambda$ denote the distribution composed of q independent copies of λ. Suppose that $\lambda_i \geq \Delta$ for all i. Then for every $\varepsilon > 0$, the distribution $\otimes^q\lambda$ is ε-close to some distribution σ such that*

$$H_\infty(\sigma) \geq qH(\lambda) - O\left(\log\left(\frac{1}{\Delta}\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}\right).$$

One can prove a similar lemma for density matrices.

Lemma 5.5. *Let ρ be a density matrix whose eigenvalues are $\lambda = (\lambda_1,\ldots,\lambda_M)$ and let q a positive integer. Suppose that for all i, $\lambda_i \geq \Delta$. Then for every $\varepsilon > 0$, $\rho^{\otimes q}$ is ε-close to some density matrix σ such that*

$$H_\infty(\sigma) \geq qS(\rho) - O\left(\log\left(\frac{1}{\Delta}\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}\right).$$

Lemma 5.5 follows directly from Lemma 5.4 because $S(\rho) = H(\lambda)$ and the vector of eigenvalues of $\rho^{\otimes q}$ equals $\otimes^q\lambda$.

We also need a way to deal with density matrices that may have arbitrarily small eigenvalues. This is really just a technicality as extremely small eigenvalues hardly affect the von Neumann entropy.

Lemma 5.6. *Let ρ be a density matrix of rank 2^m, let $\varepsilon > 0$ and let q be a positive integer. Then $\rho^{\otimes q}$ is 2ε-close to a density matrix σ, such that*

$$H_\infty(\sigma) \geq qS(\rho) - O\left(m+\log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}.$$

To prove this lemma, we will make use of the following fact [32, Box 11.2].

Fact 5.7 (Fannes' inequality)**.** Suppose ρ and σ are density matrices over a Hilbert space of dimension d. Suppose further that the trace distance between them satisfies $t = \|\rho - \sigma\|_{\mathrm{tr}} \leq 1/e$. Then

$$|S(\rho) - S(\sigma)| \leq t(\ln d - \ln t)\,.$$

Proof of Lemma 5.6. Let $\rho = \sum_{i=1}^{2^m} \lambda_i |v_i\rangle\langle v_i|$ be the spectral decomposition of ρ. Let

$$A = \left\{ i \;\middle|\; \lambda_i < \frac{\varepsilon}{q2^m} \right\}$$

denote the set of indices of "light" eigenvalues and define $\rho_0 = \sum_{i\notin A} \lambda_i |v_i\rangle\langle v_i|$. Observe that

$$\left\| \rho - \frac{\rho_0}{\mathrm{Tr}(\rho_0)} \right\|_{\mathrm{tr}} \leq \frac{\varepsilon}{q}\,.$$

The eigenvalues of the density matrix $\rho_0/\mathrm{Tr}(\rho_0)$ are all at least $\varepsilon/(q2^m)$. Hence, by Lemma 5.5, it holds that $(\rho_0/\mathrm{Tr}(\rho_0))^{\otimes q}$ is ε-close to a density matrix σ such that

$$H_\infty(\sigma) \geq q \cdot S((\rho_0/\mathrm{Tr}(\rho_0))) - O\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}\,.$$

Notice that

$$\left\| \rho^{\otimes q} - \left(\frac{\rho_0}{\mathrm{Tr}(\rho_0)}\right)^{\otimes q} \right\|_{\mathrm{tr}} \leq q \left\| \rho - \frac{\rho_0}{\mathrm{Tr}(\rho_0)} \right\|_{\mathrm{tr}} \leq \varepsilon\,,$$

and therefore $\|\rho^{\otimes q} - \sigma\|_{\mathrm{tr}} \leq 2\varepsilon$. By Fact 5.7,

$$\left| S\left(\frac{\rho_0}{\mathrm{Tr}(\rho_0)}\right) - S(\rho) \right| \leq \frac{\varepsilon}{q}\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\,.$$

Thus,

$$H_\infty(\sigma) \geq q \cdot S(\rho) - O\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}\,,$$

which completes the proof. □

5.3 QEA $\leq \overline{\mathrm{QSD}}$

We follow the outline of the classical reduction described at the beginning of the section. Let (Q,t) be an input to QEA, where Q is a quantum circuit with n input qubits and m output qubits. We consider the circuit $Q^{\otimes q}$ for $q = \mathrm{poly}(n)$ to be specified later, and we let E be a (qt,d,ε) quantum extractor operating on qm qubits, where

$$d = q(m-t) + 2\log(1/\varepsilon) + \log(qm) + O(1)\,,$$

and where $\varepsilon = 1/\mathrm{poly}(n)$ is to be specified later. Such an extractor E exists by Corollary 5.3. We then let $\xi = E(\tau_Q^{\otimes q})$ and $\tilde{I} = 2^{-qm} I$, and take the output of the reduction to be $(\xi, \tilde{I})$.

To prove the correctness of the reduction, consider first a NO-instance $(Q,t) \in \mathrm{QEA}_N$. This implies

$$S(\xi) \leq S(\tau_Q^{\otimes q}) + d \leq q(t-0.5) + d\,.$$

We fix the parameters such that

$$\frac{q}{2} \geq 2\log\left(\frac{1}{\varepsilon}\right) + \log(qm) + O(1) \tag{5.1}$$

and then $S(\xi) \leq qm - 1$. However, for any density matrix ρ over n qubits and $\varepsilon > 0$, if $S(\rho) \leq (1-\varepsilon)n$ then

$$\left\| \rho - \frac{1}{2^n} I \right\|_{\mathrm{tr}} \geq \varepsilon - \frac{1}{2^n}\,.$$

It follows that

$$\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \geq \frac{1}{qm} - \frac{1}{2^{qm}} \triangleq \beta$$

as required.

Now assume $(Q,t) \in \mathrm{QEA}_Y$. By Lemma 5.6, $\tau_Q^{\otimes q}$ is 2ε-close to a density matrix σ such that

$$\begin{aligned} H_\infty(\sigma) &\geq qS(\rho) - O\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)} \\ &\geq q\left(t+\frac{1}{2}\right) - O\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log\left(\frac{1}{\varepsilon}\right)}\,, \end{aligned}$$

and $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \leq \left\| E(\sigma) - \tilde{I} \right\|_{\mathrm{tr}} + 2\varepsilon$. We set the parameters such that $H_\infty(\sigma)$ is larger than qt, that is,

$$\frac{q}{2} \geq O\left(m + \log\left(\frac{q}{\varepsilon}\right)\right)\sqrt{q\log(1/\varepsilon)}\,. \tag{5.2}$$

Now, by the quantum extractor property we obtain $\left\| \sigma - \tilde{I} \right\|_{\mathrm{tr}} \leq \varepsilon$. Therefore, $\left\| \xi - \tilde{I} \right\|_{\mathrm{tr}} \leq 3\varepsilon \triangleq \alpha$.

We set q and ε^{-1} large enough (but still polynomial in n, e. g., $\varepsilon = \Theta(m^{-10})$ and $q = \Theta(m^4)$) such that the constraints (5.1) and (5.2) are satisfied and also that $\alpha \leq \beta^2$. Watrous [42] showed $\overline{\mathrm{QSD}_{\alpha,\beta}} \in \mathrm{QSZK}$ for these values of α, β.

5.4 QSD $\leq$ QED

Watrous [42] showed that $\mathrm{QSD}_{\alpha,\beta}$ is QSZK-complete, even with parameters $\alpha = w(n)$ and $\beta = 1 - w(n)$ where n is the size of the input and $w(n)$ is a function smaller than any inverse polynomial in n. Assume we are given an input to $\mathrm{QSD}_{\alpha,\beta}$, namely, two quantum circuits Q_0, Q_1, and construct quantum circuits Z_0 and Z_1 as follows. The circuit Z_1 outputs

$$\frac{1}{2}\left(|0\rangle\langle 0| \otimes \tau_{Q_0} + |1\rangle\langle 1| \otimes \tau_{Q_1}\right),$$

and the circuit Z_0 is the same as Z_1 except that the first register is traced out. The output of Z_0 is therefore $(1/2)(\tau_{Q_0}+\tau_{Q_1})$.

First consider the case where τ_{Q_0} and τ_{Q_1} are α close to each other, i. e., Q_0 and Q_1 produce almost the same mixed state. In this case $\tau_{Z_0} \approx \tau_{Q_0}$ whereas

$$\tau_{Z_1} \approx \frac{1}{2}(|0\rangle\langle 0| + |1\rangle\langle 1|) \otimes \tau_{Q_0}\,,$$

and therefore τ_{Z_1} has about one bit of entropy more than τ_{Z_0}. On the other hand, when τ_{Q_0} and τ_{Q_1} are very far from each other, $\tau_{Z_0} = (1/2)(\tau_{Q_0}+\tau_{Q_1})$ contains about the same amount of entropy as

$$\tau_{Z_1} = \frac{1}{2}|0\rangle\langle 0| \otimes \tau_{Q_0} + \frac{1}{2}|1\rangle\langle 1| \otimes \tau_{Q_1}\,.$$

Formally, to estimate the entropy of τ_{Z_1} one can use the joint-entropy theorem (see [32, Theorem 11.8]) to get that $S(\tau_{Z_1}) = 1 + (1/2)(S(\tau_{Q_0}) + S(\tau_{Q_1}))$. When τ_{Q_0} and τ_{Q_1} are α close to each other, Fannes' inequality (Fact 5.7) tells us that $S(\tau_{Z_0})$ is close to

$$\frac{1}{2}(S(\tau_{Q_0}) + S(\tau_{Q_1})) \leq S(\tau_{Z_1}) - 0.9\,.$$

When τ_{Q_0} and τ_{Q_1} are β far from each other, there exists a measurement that distinguishes the two with probability $(1+\beta)/2$, so by [5, Lemma 3.2] we have

$$S(\tau_{Z_0}) \geq \frac{1}{2}[S(\tau_{Q_0}) + S(\tau_{Q_1})] + \left(1 - H\left(\frac{1+\beta}{2}\right)\right) \geq S(\tau_{Z_1}) - 0.1\,.$$

The reduction from $\mathrm{QSD}_{\alpha,\beta}$ to QED is therefore as follows. Given an input (Q_0, Q_1) to $\mathrm{QSD}_{\alpha,\beta}$ we reduce it to the pair of circuits $(O_0 = Z_0 \otimes Z_0 \otimes C, O_1 = Z_1 \otimes Z_1)$ where C outputs a qubit in the completely mixed state. If $(Q_0, Q_1) \in (\mathrm{QSD}_{\alpha,\beta})_Y$ then

$$S(\tau_{O_0}) = S(\tau_{Z_0 \otimes Z_0 \otimes C}) = 2S(\tau_{Z_0}) + 1 \leq 2S(\tau_{Z_1}) - 0.8 < S(\tau_{O_1})\,,$$

whereas if $(Q_0, Q_1) \in (\mathrm{QSD}_{\alpha,\beta})_N$ then

$$S(\tau_{O_0}) = S(\tau_{Z_0 \otimes Z_0 \otimes C}) = 2S(\tau_{Z_0}) + 1 \geq 2S(\tau_{Z_1}) + 0.8 = S(\tau_{O_1}) + 0.8\,.$$

6 Closure under Boolean formulas

We have observed that one can express QED as a formula in QEA, namely,

$$\mathrm{QED}(Q_0, Q_1) = \bigvee_{t=1}^{\max\{\mathsf{out}_1, \mathsf{out}_2\}} [((Q_0, t) \in \mathrm{QEA}_Y) \wedge ((Q_1, t) \in \mathrm{QEA}_N)]\,,$$

where out_i is the number of output qubits of Q_i. In the classical setting it is known that SZK is closed under Boolean formulas. We now briefly explain why the same holds for QSZK, and refer the reader

to [38] for further details. We first define what closure under Boolean formulas means. For a promise problem Π, the *characteristic function* of Π is the map $\chi_\Pi : \{0,1\}^* \to \{0,1,\star\}$ given by

$$\chi_\Pi(x) = \begin{cases} 1 & \text{if } x \in \Pi_Y, \\ 0 & \text{if } x \in \Pi_N, \\ \star & \text{otherwise.} \end{cases}$$

A *partial assignment* to variables $v_1,\ldots,v_k$ is a k-tuple $\overline{a} = (a_1,\ldots,a_k) \in \{0,1,\star\}^k$. For a propositional formula ϕ on variables $v_1,\ldots,v_k$ the evaluation $\phi(\overline{a})$ is recursively defined as follows:

$$v_i(\overline{a}) = a_i\,, \qquad (\neg\phi)(\overline{a}) = \begin{cases} 1 & \text{if } \phi(\overline{a}) = 0, \\ 0 & \text{if } \phi(\overline{a}) = 1, \\ \star & \text{otherwise,} \end{cases}$$

$$(\phi \wedge \psi)(\overline{a}) = \begin{cases} 1 & \text{if } \phi(\overline{a}) = 1 \text{ and } \psi(\overline{a}) = 1, \\ 0 & \text{if } \phi(\overline{a}) = 0 \text{ or } \psi(\overline{a}) = 0, \\ \star & \text{otherwise,} \end{cases} \qquad (\phi \vee \psi)(\overline{a}) = \begin{cases} 1 & \text{if } \phi(\overline{a}) = 1 \text{ or } \psi(\overline{a}) = 1, \\ 0 & \text{if } \phi(\overline{a}) = 0 \text{ and } \psi(\overline{a}) = 0, \\ \star & \text{otherwise.} \end{cases}$$

Notice that, e. g., $0 \wedge \star = 0$ even though one of the inputs is "undefined" in Π. This is because one has the evaluation $a \wedge 0 = 0$, irrespective of the value of a. For any promise problem Π, we define a new promise problem $\Phi(\Pi)$, with m instances of Π as input, as follows:

$$\begin{aligned} \Phi(\Pi)_Y &= \{(\phi, x_1,\ldots,x_m) \mid \phi(\chi_\Pi(x_1),\ldots,\chi_\Pi(x_m)) = 1\}\,, \\ \Phi(\Pi)_N &= \{(\phi, x_1,\ldots,x_m) \mid \phi(\chi_\Pi(x_1),\ldots,\chi_\Pi(x_m)) = 0\}\,. \end{aligned}$$

If one can solve $\Phi(\Pi)$ then one can solve any Boolean formula over Π.

Theorem 6.1. *For any promise problem $\Pi \in$ QSZK we have $\Phi(\Pi) \in$ QSZK.*

The proof is identical to the classical proof in [38] except for straightforward adaptations (replacing the variational distance with the trace distance, using the closure of QSZK under complement, using the polarization lemma for QSD, etc.) and we sketch it here for completeness.

Proof. As QSD is QSZK-complete, Π reduces to QSD, inducing a reduction from $\Phi(\Pi)$ to $\Phi(\text{QSD})$. Thus, it suffice to show that $\Phi(\text{QSD})$ reduces to QSD. To this end, let $w = (\phi, (X_0^1, X_1^1), \ldots, (X_0^m, X_1^m))$ be an instance of $\Phi(\text{QSD})$. By applying De Morgan's Laws, we may assume that the only negations in ϕ are applied directly to the variables. (Note that De Morgan's Laws still hold in our extended Boolean algebra.) By the polarization lemma [42] and by the closure of QSZK under complementation [42], we can construct pairs of circuits $(Y_0^1, Y_1^1), \ldots, (Y_0^m, Y_1^m)$ and $(Z_0^1, Z_1^1), \ldots, (Z_0^m, Z_1^m)$ in polynomial time such that:

$$\begin{aligned} (X_0^i, X_1^i) \in \text{QSD}_Y \quad &\Rightarrow \quad \left\| \tau_{Y_0^i} - \tau_{Y_1^i} \right\|_{\text{tr}} \geq 1 - \frac{1}{3|\phi|} \quad \text{and} \quad \left\| \tau_{Z_0^i} - \tau_{Z_1^i} \right\|_{\text{tr}} \leq \frac{1}{3|\phi|}\,, \\ (X_0^i, X_1^i) \in \text{QSD}_N \quad &\Rightarrow \quad \left\| \tau_{Y_0^i} - \tau_{Y_1^i} \right\|_{\text{tr}} \leq \frac{1}{3|\phi|} \quad \text{and} \quad \left\| \tau_{Z_0^i} - \tau_{Z_1^i} \right\|_{\text{tr}} \geq 1 - \frac{1}{3|\phi|}\,. \end{aligned}$$

The reduction outputs the pair of circuits $(\text{BuildCircuit}(\phi,0), \text{BuildCircuit}(\phi,1))$, where BuildCircuit is described by the following recursive procedure:

BuildCircuit(ψ, b)

1. If $\psi = v_i$, output Y_b^i.
2. if $\psi = \neg v_i$, output Z_b^i.
3. If $\psi = \zeta \vee \mu$, output $\text{BuildCircuit}(\zeta,b) \otimes \text{BuildCircuit}(\mu,b)$.
4. If $\psi = \zeta \wedge \mu$, output

$$\frac{1}{2}(\text{BuildCircuit}(\zeta,0) \otimes \text{BuildCircuit}(\mu,b)) + \frac{1}{2}(\text{BuildCircuit}(\zeta,1) \otimes \text{BuildCircuit}(\mu,1-b)).$$

Notice that the number of recursive calls equals the number of sub-formula of ϕ, and therefore the procedure runs in time polynomial in $|\psi|$ and $|X_i^j|$, i. e., polynomial in its input length.

We now turn to proving the correctness of this reduction. The correctness will follow from the claim below, wherein we define

$$\Delta(\zeta) \quad = \quad \frac{1}{2}\left\| (\text{BuildCircuit}(\zeta,0) - \text{BuildCircuit}(\zeta,1))\,|0\rangle \right\|_{\text{tr}}$$

for each sub-formula ζ of ϕ.

Claim 6.2. *Let $\overline{a} = (\chi_{QSD}(X_0^1, X_1^1), \ldots, \chi_{QSD}(X_0^m, X_1^m))$. For every sub-formula ψ of ϕ, we have:*

$$\begin{aligned} \psi(\overline{a}) = 1 \quad &\Rightarrow \quad \Delta(\psi) \geq 1 - \frac{|\psi|}{3|\phi|}, \\ \psi(\overline{a}) = 0 \quad &\Rightarrow \quad \Delta(\psi) \leq \frac{|\psi|}{3|\phi|}. \end{aligned}$$

Proof. The proof is by induction on the sub-formulas ψ of ϕ, and we note that it clearly holds for atomic sub-formulas. The remaining two cases are as follows.

Case 1: $\psi = \zeta \vee \mu$. If $\psi(\overline{a}) = 1$ then either $\zeta(\overline{a}) = 1$ or $\mu(\overline{a}) = 1$. Without loss of generality assume $\zeta(\overline{a}) = 1$. In this case we have for any $i \in \{0,1\}$ that $\text{BuildCircuit}(\zeta, i) = \mathcal{E}(\text{BuildCircuit}(\psi, i))$, where $\mathcal{E}$ is the quantum operation tracing out the registers associated with the μ sub-formula. Thus, by induction,

$$\Delta(\psi) \geq \Delta(\zeta) \geq 1 - \frac{|\zeta|}{3|\phi|} \geq 1 - \frac{|\psi|}{3|\phi|}.$$

If $\psi(\overline{a}) = 0$, then both $\zeta(\overline{a}) = \mu(\overline{a}) = 0$.

Using

$$\begin{aligned} \| \rho_0 \otimes \rho_1 - \sigma_0 \otimes \sigma_1 \|_{\text{tr}} \quad &\leq \quad \| \rho_0 \otimes \rho_1 - \sigma_0 \otimes \rho_1 \|_{\text{tr}} + \| \sigma_0 \otimes \rho_1 - \sigma_0 \otimes \sigma_1 \|_{\text{tr}} \\ &= \quad \| \rho_0 - \sigma_0 \|_{\text{tr}} + \| \rho_1 - \sigma_1 \|_{\text{tr}}, \end{aligned}$$

we obtain

$$\Delta(\psi) \leq \Delta(\zeta) + \Delta(\mu) \leq \frac{|\zeta|}{3|\phi|} + \frac{|\mu|}{3|\phi|} \leq \frac{|\psi|}{3|\phi|}.$$

Case 2: $\psi = \zeta \wedge \mu$. Using

$$\begin{aligned} & \frac{1}{2} \left\| \frac{1}{2}[\rho_0 \otimes \sigma_0 + \rho_1 \otimes \sigma_1] - \frac{1}{2}[\rho_0 \otimes \sigma_1 + \rho_1 \otimes \sigma_0] \right\|_{\mathrm{tr}} \\ = \quad & \frac{1}{4} \| (\rho_0 - \rho_1) \otimes (\sigma_0 - \sigma_1) \|_{\mathrm{tr}} \;=\; \frac{1}{4} \| \rho_0 - \rho_1 \|_{\mathrm{tr}} \| \sigma_0 - \sigma_1 \|_{\mathrm{tr}}, \end{aligned}$$

we obtain $\Delta(\psi) = \Delta(\zeta) \cdot \Delta(\mu)$. If $\psi(\bar{a}) = 1$, then, by induction,

$$\Delta(\psi) \geq \left(1 - \frac{|\zeta|}{3|\phi|}\right)\left(1 - \frac{|\mu|}{3|\phi|}\right) > 1 - \frac{|\zeta| + |\mu|}{3|\phi|} \geq 1 - \frac{|\psi|}{3|\phi|}.$$

If $\psi(\bar{a}) = 0$, then, without loss of generality, we may assume $\zeta(\bar{a}) = 0$. By induction we have

$$\Delta(\psi) = \Delta(\zeta) \cdot \Delta(\mu) \leq \Delta(\zeta) \leq \frac{|\zeta|}{3|\phi|} \leq \frac{|\psi|}{3|\phi|}.$$

Thus, the claim has been proved. □

Let $A_b = \mathrm{BuildCircuit}(\phi, b)$. By the above claim, if $w \in \Phi(\mathrm{QSD})_Y$ then $\| \tau_{A_0} - \tau_{A_1} \|_{\mathrm{tr}} \geq 2/3$ and if $w \in \Phi(\mathrm{QSD})_N$ then $\| \tau_{A_0} - \tau_{A_1} \|_{\mathrm{tr}} \leq 1/3$. This completes the proof of the theorem. □

Acknowledgements

We thank Oded Regev for pointing out [6] to us. We also thank Ashwin Nayak, Oded Regev, Adam Smith and Umesh Vazirani for helpful discussions about the paper. We thank the anonymous referees for many helpful comments.

References

[1] JEFFREY ADAMS: Character tables for $GL(2), SL(2), PGL(2)$ and $PSL(2)$ over a finite field. http://www.math.umd.edu/~jda/characters/characters.pdf, 2002. 60

[2] N. ALON AND V. D. MILMAN: λ_1, isoperimetric inequalities for graphs, and superconcentrators. *J. Combin. Theory Ser. B*, 38(1):73–88, 1985. [Elsevier:10.1016/0095-8956(85)90092-9]. 48

[3] NOGA ALON: Eigenvalues and expanders. *Combinatorica*, 6(2):83–96, 1986. [Springer:90541164662405r4]. 48

[4] NOGA ALON AND YUVAL ROICHMAN: Random Cayley graphs and expanders. *Random Structures Algorithms*, 5(2):271–285, 1994. [Wiley:10.1002/rsa.3240050203]. 50

[5] ANDRIS AMBAINIS, ASHWIN NAYAK, AMNON TA-SHMA, AND UMESH VAZIRANI: Quantum dense coding and quantum finite automata. *J. ACM*, 49(4):496–511, 2002. [JACM:10.1145/581771.581773]. 72

[6] ANDRIS AMBAINIS AND ADAM SMITH: Small pseudo-random families of matrices: Derandomizing approximate quantum encryption. In *Proc. 8th Intern. Workshop Randomization and Comput.*, volume 3122 of *LNCS*, pp. 249–260. Springer-Verlag, 2004. [doi:10.1007/b99805]. 49, 51, 52, 69, 75

[7] ROBERT BEALS: Quantum computation of Fourier transforms over symmetric groups. In *Proc. 29th STOC*, pp. 48–53. ACM Press, 1997. [STOC:10.1145/258533.258548]. 50

[8] AVRAHAM BEN-AROYA AND AMNON TA-SHMA: Quantum expanders and the quantum entropy difference problem. Technical report, arXiv:quant-ph/0702129, 2007. [arXiv:quant-ph/0702129]. 50

[9] MICHAEL CAPALBO, OMER REINGOLD, SALIL VADHAN, AND AVI WIGDERSON: Randomness conductors and constant-degree lossless expanders. In *Proc. 34th STOC*, pp. 659–668. ACM Press, 2002. [STOC:10.1145/509907.510003]. 49

[10] C. M. DAWSON AND M. A. NIELSEN: The Solovay-Kitaev algorithm. *Quantum Inf. Comput.*, 6(1):81–95, 2006. 63

[11] PAUL DICKINSON AND ASHWIN NAYAK: Approximate randomization of quantum states with fewer bits of key. In *Proc. AIP Conf.*, volume 864, pp. 18–36. AIP, 2006. 51

[12] JOZEF DODZIUK: Difference equations, isoperimetric inequality and transience of certain random walks. *Trans. Amer. Math. Soc.*, 284(2):787–794, 1984. http://www.jstor.org/stable/1999107. 48

[13] JOEL FRIEDMAN: A proof of Alon's second eigenvalue conjecture. *Mem. Amer. Math. Soc.*, to appear. 48

[14] OFER GABBER AND ZVI GALIL: Explicit constructions of linear-sized superconcentrators. *J. Comput. System Sci.*, 22(3):407–420, 1981. [JCSS:10.1016/0022-0000(81)90040-4]. 48

[15] ODED GOLDREICH AND AVI WIGDERSON: Tiny families of functions with random properties: A quality-size trade-off for hashing. *Random Structures Algorithms*, 11(4):315–343, 1997. [Wiley:10.1002/(SICI)1098-2418(199712)11:4¡315::AID-RSA3¿3.0.CO;2-1]. 68

[16] D. GROSS AND J. EISERT: Quantum Margulis expanders. *Quantum Inf. Comput.*, 8(8&9):722–733, 2008. 51

[17] JOE HARRIS AND WILLIAM FULTON: *Representation Theory*. Springer, 1991. 54, 60

[18] A.W. HARROW: Quantum expanders from any classical Cayley graph expander. *Quantum Inf. Comput.*, 8(8&9):715–721, 2008. 51

[19] M. B. HASTINGS: Entropy and entanglement in quantum ground states. *Phys. Rev. B*, 76(3):035114, 2007. [PRB:10.1103/PhysRevB.76.035114]. 49, 51

[20] M. B. HASTINGS: Random unitaries give quantum expanders. *Phys. Rev. A*, 76(3):032315, 2007. [PRA:10.1103/PhysRevA.76.032315]. 50, 63

[21] M.B. HASTINGS AND A.W. HARROW: Classical and quantum tensor product expanders. *Quantum Inf. Comput.*, 9(3&4):336–360, 2009. 52

[22] SCHLOMO HOORY, NATHAN LINIAL, AND AVI WIGDERSON: Expander graphs and their applications. *Bull. Amer. Math. Soc.*, 43(4):439–561, 2006. [doi:10.1090/S0273-0979-06-01126-8]. 48, 51

[23] SHUJI JIMBO AND AKIRA MARUOKA: Expanders obtained from affine transformations. *Combinatorica*, 7(4):343–355, 1987. [Springer:6w2126413464n387]. 48

[24] NABIL KAHALE: Eigenvalues and expansion of regular graphs. *J. ACM*, 42(5):1091–1106, 1995. [JACM:10.1145/210118.210136]. 48

[25] MARTIN KASSABOV: Symmetric groups and expanders. *Electron. Res. Announc. Amer. Math. Soc.*, 11, 2005. 50, 51

[26] MARIA KLAWE: Limitations on explicit constructions of expanding graphs. *SIAM J. Comput.*, 13(1):156–166, 1984. [SICOMP:10.1137/0213011]. 50

[27] JOHN D. LAFFERTY AND DANIEL ROCKMORE: Fast Fourier analysis for SL_2 over a finite field and related numerical experiments. *Experiment. Math.*, 1(2):115–139, 1992. `http://www.expmath.org/restricted/1/1.2/lafferty.ps.gz`. 50

[28] A. LUBOTZKY, R. PHILIPS, AND P. SARNAK: Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988. [Springer:k285687344657q53]. 48, 50, 57

[29] G. A. MARGULIS: Explicit constructions of expanders. *Problemy Peredachi Informatsii*, 9(4):71–80, 1973. `http://mi.mathnet.ru/eng/ppi925`. 48, 51

[30] 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. `http://mi.mathnet.ru/eng/ppi686`. 48, 50

[31] MOSHE MORGENSTERN: Existence and explicit constructions of $q+1$ regular Ramanujan graphs for every prime power q. *J. Combin. Theory Ser. B*, 62(1):44–62, 1994. [Elsevier:10.1006/jctb.1994.1054]. 48

[32] MICHAEL A. NIELSEN AND ISAAC L. CHUANG: *Quantum Computation and Quantum Information*. Cambridge University Press, 2000. 70, 72

[33] A. NILLI: On the second eigenvalue of a graph. *Discrete Math.*, 91(2):207–210, 1991. [Elsevier:10.1016/0012-365X(91)90112-F]. 48

[34] M. PINSKER: On the complexity of a concentrator. In *Proc. 7th Intern. Teletraffic Conf.*, pp. 318/1–318/4, 1973. 48

[35] SANDU POPESCU AND DANIEL ROHRLICH: Thermodynamics and the measure of entanglement. *Phys. Rev. A*, 56(5):3319–3321, 1997. [PRA:10.1103/PhysRevA.56.R3319]. 53

[36] OMER REINGOLD, SALIL VADHAN, AND AVI WIGDERSON: Entropy waves, the zig-zag graph product, and new constant-degree expanders. *Ann. of Math.*, 155(1):157–187, 2002. 49, 50, 62, 64

[37] G. DE B. ROBINSON: On the representations of the symmetric group. *Amer. J. Math.*, 60(3):745–760, 1938. http://www.jstor.org/stable/2371609. 59

[38] AMIT SAHAI AND SALIL VADHAN: Manipulating statistical difference. In *Proc. Randomization Methods in Algorithm Design (DIMACS Workshop, December 1997)*, volume 43 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pp. 251–270. American Mathematical Society, 1999. 73

[39] C. SCHENSTED: Longest increasing and decreasing subsequences. *Canad. J. Math.*, 13(2):179–191, 1961. 59

[40] JEAN-PIERRE SERRE: *Linear representations of finite groups*, volume 42 of *Graduate texts in Mathematics*. Springer, 1977. 54, 59

[41] SALIL VADHAN: *A Study of Statistical Zero-Knowledge Proofs*. PhD thesis, Massachusetts Institute of Technology, 1999. 69

[42] JOHN WATROUS: Limits on the power of quantum statistical zero-knowledge. In *Proc. 43rd FOCS*, pp. 459–470. IEEE Comp. Soc. Press, 2002. [FOCS:10.1109/SFCS.2002.1181970]. 52, 71, 73

[43] JOHN WATROUS: Zero-knowledge against quantum attacks. In *Proc. 38th STOC*, pp. 296–305. ACM Press, 2006. [STOC:10.1145/1132516.1132560]. 52

AUTHORS

Avraham Ben-Aroya
student
Tel-Aviv University, Tel-Aviv, Israel
abrhambe@tau.ac.il
http://www.cs.tau.ac.il/~abrhambe

Oded Schwartz
postdoc
Institut für Mathematik, MA 4-5, Technische Universität Berlin, 10623 Berlin, Germany
odedsc@math.tu-berlin.de
http://www.math.tu-berlin.de/numerik/mt/schwartz_de.html

Amnon Ta-Shma
professor
Tel-Aviv University, Tel-Aviv, Israel
amnon@tau.ac.il
http://www.cs.tau.ac.il/~amnon

ABOUT THE AUTHORS

AVRAHAM BEN-AROYA is a graduate student at Tel-Aviv University. His advisors are Oded Regev and Amnon Ta-Shma. His research interests include quantum computation, pseudorandomness and other topics in theoretical computer science. He also enjoys playing tennis, chess and plastic guitars.

ODED SCHWARTZ completed his PhD at Tel-Aviv University in 2007; his advisors were Muli Safra and Amnon Ta-Shma. This is his first paper in Theory of Computing.

AMNON TA-SHMA is a theoretical computer scientist. This is his second paper in Theory of Computing.