

On approximating the eigenvalues of stochastic matrices in probabilistic logspace

Dean Doron, Amir Sarid, and Amnon Ta-Shma

The Blavatnik School of Computer Science
Tel-Aviv University
Israel 69978

August 1, 2016

Abstract

Approximating the eigenvalues of a Hermitian operator can be solved by a quantum logspace algorithm. We introduce the problem of approximating the eigenvalues of a given matrix in the context of classical space-bounded computation. We show that:

- Approximating the second eigenvalue of stochastic operators (in a certain range of parameters) is BPL-complete, and,
- We show a BPL algorithm that approximates any eigenvalue of a stochastic and *Hermitian* operator with *constant* accuracy.

The last result falls short of achieving the polynomially-small accuracy that the quantum algorithm achieves. Thus, at our current state of knowledge, for stochastic and Hermitian operators we can achieve polynomially-small accuracy with quantum logspace algorithms, constant accuracy with probabilistic logspace algorithms, and no non-trivial result is known for deterministic logspace algorithms. The quantum algorithm also has the advantage of working over arbitrary, possibly non-stochastic Hermitian operators.

Our work raises several challenges. First, a derandomization challenge, trying to achieve a deterministic algorithm approximating eigenvalues with some non-trivial accuracy. Second, a de-quantumization challenge trying to decide whether the quantum logspace model is strictly stronger than the classical probabilistic model or not. It also casts the probabilistic and quantum space-bounded models as problems in linear algebra with differences between operators that are symmetric, stochastic or both. We therefore believe the problem of approximating the eigenvalues of a graph is not only natural and important by itself, but also important for understanding the relative power of deterministic, probabilistic and quantum logspace computation.

1 Introduction

A graph G can be associated with a linear operator A that describes a random walk on G . The operator A takes an especially simple form when G is undirected:

If G is regular and undirected then A is symmetric and has a complete basis of orthonormal eigenvectors with real eigenvalues. In other words, there exists a unitary basis under which A is diagonal with real eigenvalues λ_i on the diagonal.

If G is undirected but not necessarily regular, then A is diagonalizable with real eigenvalues, i.e., the picture is the same as before except that the basis is not necessarily unitary.¹

If G is directed then A does not necessarily have a full basis of eigenvectors. In this case A (like any other linear operator) can be brought to its canonical Jordan Normal Form, where there exists a basis under which A is block-diagonal and each block has an eigenvalue λ on the main diagonal and 1 on the diagonal above it. For each λ -block there exists exactly one eigenvector with eigenvalue λ . The number of blocks with eigenvalue λ is the *geometrical* multiplicity of λ and equals the dimension of the λ -eigenspace. The number of times λ appears on the main diagonal is the *algebraic* multiplicity of λ and equals the multiplicity of λ as a root of the characteristic polynomial of A .

Thus, even though directed graphs do not have a full set of eigenvectors they still have a full set of eigenvalues. The eigenvalues of a stochastic matrix are always bounded by 1, and the all-one vector is a 1 eigenvector. The Perron Frobenius theorem tells us that if the graph is strongly connected and a-periodic then the 1-eigenvector is unique. If the graph is periodic with a period h then there are h complex eigenvalues corresponding to the h different h -roots of unity.

In the paper we show that the problem of estimating the magnitude of the second eigenvalue of a *stochastic* operator is BPL-complete. More precisely, we define the following promise problem:

Definition 1.1 (*SecondEV*). *The input is a stochastic matrix A of dimension n with possibly complex eigenvalues $\lambda_1, \dots, \lambda_n$. We have the promise that $\lambda_1 = 1$ and $\lambda_2 \in \mathbb{R}$. Also, $|\lambda_n| \leq \dots \leq |\lambda_3| \leq 1 - \frac{2}{n}$.*

No instances : $\lambda_2 \geq 1 - \frac{1}{n}$.

Yes instances : $\lambda_2 \leq 1 - \frac{2}{n}$.

And we prove:

Theorem 1.1. *SecondEV is BPL-complete.*

This raises several immediate questions:

- Is it possible to approximate in BPL the second eigenvalue of a general (not necessarily stochastic or non-negative) operator?
- Does the problem stay BPL-hard when the graph is undirected?
- How difficult is it to approximate all the spectrum of an operator? What accuracy can we achieve?

¹If G is undirected and irregular, then the adjacency matrix \tilde{A} is symmetric (because the graph is undirected) but the transition matrix $A = D^{-1}\tilde{A}$, where F is the diagonal out-degree matrix, is not symmetric. Yet, consider the matrix $L = D^{-1/2}\tilde{A}D^{-1/2}$. L is symmetric and thus has an eigenvector basis with real eigenvalues. $A = D^{-1/2}LD^{1/2}$ is conjugate to L and thus is diagonalizable and has the same eigenvalues. As A is stochastic its eigenvalues are in the range $[-1, 1]$.

1.1 The bigger picture

Derandomization is a major challenge of theoretical computer science. In the space-bounded model, Nisan [13] constructed a pseudo-random generator (PRG) against logarithmic space-bounded non-uniform algorithms that uses seed length $O(\log^2 n)$. Using that he showed BPL is contained in the class having simultaneously polynomial time and $O(\log^2 n)$ space. Saks and Zhou [18] showed BPL is contained in $DSPACE(\log^{1.5} n)$. Reingold [14] showed undirected st-connectivity (which was shown to be in RL by [1]) already belongs to L. These results seem to indicate that randomness does not add additional power to the model and many conjecture that in fact $BPL = L$. Yet, we currently do not know a PRG with seed length $o(\log^2 n)$, nor a general derandomization result that simultaneously uses $o(\log^2 n)$ space and polynomial time.

One can look up and ask which upper bounds we know on BPL. We then know the following:

$$NC^1 \subseteq L \subseteq RL \subseteq NL \subseteq DET \subseteq NC^2 \subseteq DSPACE(O(\log^2 n)),$$

where DET is the class of languages that are NC^1 Turing-reducible to the problem *intdet* of computing the determinant of an integer matrix (see [3] for a definition of DET). As it turns out, many important problems in linear algebra, such as inverting a matrix, or equivalently, solving a set of linear equations are in DET, and often complete for it (see, e.g., [3]). The fact that $NL \subseteq DET$ is due to [3] who showed that the directed connectivity problem, STCON is reducible to *intdet*. $DET \subseteq NC^2$ follows from Csanky's algorithm [4] for the parallel computation of the determinant. In addition to the above we also know that $BPL \subseteq DET$ (e.g., using the fact that matrix powering is DET complete).

While matrix powering is complete for DET, *approximating* matrix powering of *stochastic* matrices is in BPL. To see that, assume A represents a stochastic matrix. Then one can approximate $A^k[s, t]$ by estimating the probability a random walk over A starting at s reaches t after k steps.² Conversely, it is possible to convert a BPL machine to a stochastic operator A such that the probability the machine moves from s to t in k steps is $A^k[s, t]$.³ Thus, approximating matrix-powering of stochastic operators is complete for BPL.

We now deviate from the classical picture we had so far and consider a quantum space-bounded model. In 1999, Watrous [22] defined the model of quantum logspace computation, and proved several facts on it. The definition was modified several times, see, [21]. Roughly speaking, a language is in BQL if there exists an L-uniform family of quantum circuits solving the language with only $O(\log n)$ qubits. The quantum circuits are over some universal basis of gates (e.g., CNOT, HAD, T) plus intermediate measurements (that in particular may simulate a stream of random coins). For details we refer the reader to [21, 19]. The works of Watrous, van Melkebeek and Watson showed that BQL is also contained in NC^2 .

Recently, it was shown in [19], building on an earlier work by [10], that it is possible to *approximate* the singular value decomposition (SVD) of a given linear operator in BQL. This also implies that it is possible to approximately invert a matrix in BQL. A natural question left open by this work is:

Open problem: Is it possible to approximate the SVD of an arbitrary linear operator already in BPL? The problem is also open for Hermitian operators where singular values and eigenvalues coincide (up to their sign). In fact, the question is still open even when the operator is the transition matrix arising from a walk on a regular, undirected graph.

²For completeness we include a proof of this in Appendix A. We also extend the class for which this works to matrices with negative or complex entries as long as their infinity norm is at most 1.

³This reduction is standard and appears in many papers, e.g., already in [13]. We also employ such a reduction, see Subsection 3.1 in this paper.

Thus, somewhat surprisingly, we see that the deterministic/probabilistic/quantum space-bounded classes and the class DET are capable of doing some sort of linear algebra on corresponding operators. Namely,

- In DET we can compute exactly the determinant which is the product of all eigenvalues as well as the product of all singular values. We can also solve matrix powering. Both problems are *complete* in DET. With that we can approximately invert an operator or perform the SVD decomposition.

In a sense, DET is an *exact* computation of the spectrum (e.g., in terms of the characteristic polynomial) of an *arbitrary* linear operator.

- BQL is capable of approximating the whole singular value decomposition of any operator. This is somewhat equivalent to saying that BQL is capable of *approximating* the eigenvalues of *Hermitian* operators.
- BPL is capable of *approximating* matrix powering, and second eigenvalue of a *stochastic* operator, and both problems are complete in BPL.

We do not know yet whether we can approximate the whole spectrum even for undirected graphs (that correspond to operators that are both stochastic and Hermitian, in the regular case) and we discuss it next.

- In L we do not know how to do any of the above, but Reingold showed L is capable of solving USTCON, i.e., connectivity on undirected graphs.

Next we will see that on *undirected graphs* we can approximate in BPL arbitrary eigenvalues to within constant accuracy. We remark that in BQL we can do the same to arbitrary Hermitian operators and with polynomially-good accuracy.

1.2 Approximating arbitrary eigenvalues of undirected graphs in BPL

We define the following promise problem:

Definition 1.2 ($EV_{\alpha,\beta}$). *The input is a stochastic, Hermitian matrix A , $\lambda \in [-1, 1]$ and $\alpha < \beta$.*

Yes instances : *There is an eigenvalue λ_i of A such that $|\lambda_i - \lambda| \leq \alpha$.*

No instances : *All eigenvalues of A are β -far from λ .*

One way to design a BPL algorithm for the problem is by “de-quantumizing” the quantum algorithm.⁴ The BQL algorithm solves the above problem for any Hermitian operator A whose eigenvalues are τ -separated, for, say, $\tau = n^{-c}$, $\alpha = \frac{\tau}{4}$ and $\beta = 2\alpha$. That is, the quantum algorithm can handle any polynomially small accuracy. With such accuracy one can turn the solution of the promise problem to a procedure approximating the whole spectrum.

We develop a BPL algorithm that follows the main idea of the quantum algorithm, and in that sense we de-quantumize the quantum algorithm, but we achieve much worse parameters. Specifically, we prove that the promise problem $EV_{\alpha,\beta}$ belongs to BPL, for *constant* parameters $\alpha < \beta$. On the one hand the result is disappointing because the quantum algorithm does so much better and can handle polynomially small gaps. On the other hand, we remark that we do not know how to achieve even constant approximation with a deterministic logspace algorithm.

⁴We remark that Ben-Or and Eldar [2] recently de-quantumized the SVD quantum algorithm and obtained a classical probabilistic algorithm for inverting matrices that achieves the state of the art running time, using a completely new approach that is derived from the quantum algorithm. We would like to do the same in the space-bounded model.

1.3 Our technique

The usual way of describing the quantum algorithm is that it applies quantum phase estimation on the completely mixed state. The completely mixed state is a uniform mixture of the pure states that are formed from the eigenvectors of A , and on each such eigenvector, the quantum phase estimation estimates the corresponding eigenvalue. Thus, if the procedure can be run in (quantum) logarithmic space, we essentially sample a random eigenvector/eigenvalue pair, and from that we can approximately get the SVD decomposition of A .

Another (less standard) way of viewing the quantum algorithm is that it manipulates the eigenvalues of an input matrix A without knowing the decomposition of A to eigenvectors and eigenvalues. This can be done using the simple fact that if $\lambda_1, \dots, \lambda_n$ are the roots of the characteristic polynomial of A , and if p is an arbitrary univariate polynomial, then $p(\lambda_1), \dots, p(\lambda_n)$ are the roots of the characteristic polynomial of the matrix $p(A)$. The probability the algorithm measures λ is proportional to $\text{Tr}(p(A))$, where p is a shift of the Fejér kernel by λ (see, e.g., [11, Chapter 2]). Applying p on A amplifies the eigenvalues that are close to λ to a value close to 1, and damps eigenvalues far from λ close to 0. Thus, $\text{Tr}(p(A))$ approximately counts the number of eigenvalues close to λ .

We would like to follow the same approach but with a probabilistic algorithm rather than a quantum one. We say a matrix A is *simulatable* if a probabilistic logspace algorithm can approximate $A^k[s, t]$ for any k polynomial in n and with polynomially-small accuracy (see Definition 2.2 for the exact details). From the discussion above it is clear that if A is the transition matrix of a (directed or undirected) graph then A is simulatable (see Lemma 2.1). We remark that in the appendix we show that even non-stochastic matrices A with negative or complex entries are simulatable as long as A has infinity norm at most 1, namely, those matrices A for which all rows $i \in [n]$ have ℓ_1 norm at most 1, $\sum_j |A[i, j]| \leq 1$.

If A is simulatable and the coefficients of $p(x) = \sum_i c_i x^i$ are not too large (i.e., only polynomially large in n), then we can approximate in BPL the matrix $p(A) = \sum_i c_i A^i$. In particular, we can also approximate $\text{Tr}(p(A))$. By taking p to be a threshold polynomial with degree logarithmic in n (that guarantees the size of the coefficients c_i is polynomial in n) and a threshold around λ , we can solve $EV_{\alpha, \beta}(A)$ for constants $\alpha < \beta$ (see Section 4).

There are many other possible candidate functions for a threshold polynomial p . However, we prove in Theorem 4.4 that no polynomial can do significantly better than a threshold polynomial. The reason the quantum algorithm works better is because it is able to take p up to some polynomial degree (rather than logarithmic degree) not worrying about the (quite large) size of the coefficients, thus leading to much better accuracy. The quantum algorithm also has the advantage that it works for any normal operator A , not necessarily stochastic or simulatable.

Thus, the algorithm we give for $EV_{\alpha, \beta}$ is simple: Approximate $\text{Tr}(p(A))$ to a simple logarithmic degree polynomial p . Nevertheless, we believe it features a new component that has not been used before by probabilistic space-bounded algorithms. An algorithm that takes a random walk on a graph and takes a decision based on the walk length and connectivity properties of the graph (as, e.g., [1]) works with some power of the input matrix A . More generally, such an algorithm can work with a convex combination of powers of the input matrix (by probabilistically choosing which power to take). The algorithm we present utilizes *arbitrary* (positive or *negative*) combinations of matrix powers and we believe it is a crucial feature of the solution. We are not aware of previous BPL algorithms using such a feature.

The approach above does not work for approximating the eigenvalues of a *directed* graph G . It is still true that the resulting operator A is stochastic and therefore simulatable. Also, it remains true that if λ is an eigenvalue of A (i.e., a root of the characteristic polynomial) then $p(\lambda)$ is a root of $p(A)$. However, since A is not Hermitian, the eigenvalues λ of A may be complex and we do not know how to control $p(\lambda)$ when p may have both negative and positive coefficients. The reason why we manage to prove Theorem

1.1 for arbitrary stochastic operators, is that in the special case of the second eigenvalue (and when the first eigenvalue is known) we can amplify with a polynomial having only positive coefficients, and then we can control the complex eigenvalues.

1.4 A short discussion

At our current state of knowledge, the problem of approximating the eigenvalues of an undirected graph separates deterministic, probabilistic and quantum complexity: In BQL we can solve it with polynomially-small accuracy, in BPL with constant accuracy and in L we do not know how to solve it at all. Thus it poses several challenges:

- First, there is the natural question of whether one can approximate eigenvalues in BPL with better accuracy. A positive answer would imply BPL approximations to many important linear algebra problems that are currently only known to be in NC^2 . A negative answer would imply a separation between BQL and BPL [6].
- Second, it raises the natural question of derandomization. Can one design a *deterministic* algorithm approximating eigenvalues to constant accuracy?

We believe the solution of this problem is not only important by itself, but may also shed new light on the strengths and weaknesses of the space-bounded model, and the relative strengths of the deterministic, probabilistic and quantum models of space-bounded computation.

It is also interesting to compare the BPL-complete promise problem *SecondEV* that we get with the RL-complete promise problem given in [15]. First, following [12, 9] they define $\lambda_\pi(G)$ to be the second largest eigenvalue of $A\tilde{A}$, where $\tilde{A}[i, j] = \frac{\pi(j)A[j, i]}{\pi(i)}$ and π is the stationary distribution of a random walk on G . Then, they define the promise problem POLY-MIXING S-T CONNECTIVITY to be:

Input : $(G, s, t, 1^k)$ where $G = (V, E)$ is an out-regular directed graph, $s, t \in V$, and $k \in \mathbb{N}$.

Yes instances : $\lambda_\pi(G) \leq 1 - \frac{1}{k}$, and $\pi(s), \pi(t) \geq \frac{1}{k}$.

No instances : There is no path from s to t in G .

Notice that the problem is made RL-complete by forcing the condition that in No instances there is no path from s to t . The main difference between the two promise problems is that in Section 3 the spectral gap is defined with respect to the eigenvalues of the transition operator A (i.e., the roots of its characteristic polynomial) whereas in [15], the spectral gap is defined with respect to a *normalized* inner product space, and the normalization is with respect to the unknown stationary distribution π , which is, by itself, a major component of the solution.

2 Preliminaries

2.1 Space bounded probabilistic computation

A deterministic space-bounded Turing machine has three semi-infinite tapes: an *input tape* (that is read-only); a *work tape* (that is read-write) and an *output tape* (that is write-only and uni-directional). The space complexity of the machine is the number of cells on the work tape. The running time of a space-bounded

Turing machine with $s(n) \geq \log n$ space complexity is bounded by $2^{O(s(n))}$ time. A *probabilistic* space-bounded Turing machine is similar to the deterministic machine (and in particular we require it always halts within $2^{O(s(n))}$ time) except that it can also toss random coins. One convenient way to formulate this is by adding a fourth semi-infinite tape, the *random-coins tape*, that is read-only, uni-directional and is initialized with perfectly uniform bits. We are only concerned with bounded-error computation: We say a language is accepted by a probabilistic Turing machine if for every input in the language the acceptance probability is at least $2/3$, and for every input not in the language it is at most $1/3$. As usual, the acceptance probability can be amplified as long as there is some non-negligible gap between the acceptance probability of yes and no instances.

Definition 2.1. *A language is in $\text{BPSPACE}(s(n))$ if it is accepted by a probabilistic space bounded TM with space complexity $s(n)$. $\text{BPL} = \cup_c \text{BPSPACE}(c \log n)$.*

2.2 Simulatable families of matrices

Often we are interested in approximating a *value* (e.g., an entry in a matrix with integer values or the whole matrix) with a probabilistic machine. More precisely, assume there exists some value $u = u(x) \in \mathbb{R}$ that is determined by the input $x \in \{0, 1\}^n$. We say a probabilistic TM $M(x, y)$ (ε, δ) -approximates $u(x)$ if:

$$\forall_{x \in \{0,1\}^n} \Pr_y [|M(x, y) - u(x)| \geq \varepsilon] \leq \delta \quad (1)$$

A random walk on a graph G (or its transition matrix A) can be simulated by a probabilistic logspace machine. As a consequence, a probabilistic logspace machine can approximate powers of A well. Here we try to extend this notion to arbitrary linear operators A , not necessarily stochastic. We say a matrix A is *simulatable* if any power of it can be approximated by a probabilistic algorithm running in small space. Formally:

Definition 2.2. *We say that a family of matrices \mathcal{A} is simulatable if there exists a probabilistic algorithm that on input $A \in \mathcal{A}$ of dimension n with $\|A\| \leq \text{poly}(n)$, $k \in \mathbb{N}$, $s, t \in [n]$, runs in space $O(\log \frac{nk}{\varepsilon\delta})$ and (ε, δ) -approximates $A^k[s, t]$.*

In Appendix A we give for completeness a proof that:

Lemma 2.1. *The family of transition matrices of (directed or undirected) graphs is simulatable.*

We say $\|A\|_\infty \leq c$ if for every $i \in [n]$, $\sum_j |A[i, j]| \leq c$. In the same Appendix we also show:

Lemma 2.2. *The family of real matrices with infinity norm at most 1 is simulatable.*

3 A new canonical problem for probabilistic space-bounded classes

We define the following promise problem:

Definition 3.1 (The promise problem $\text{SecondEV}_{\alpha, \beta}$). *The input is $0 \leq \alpha < \beta < 1$ and a stochastic matrix A of dimension n with possibly complex eigenvalues $\lambda_1, \dots, \lambda_n$. We have the promise that $\lambda_1 = 1$ and $\lambda_2 \in \mathbb{R}$. Also, $|\lambda_n| \leq \dots \leq |\lambda_3| \leq 1 - \beta$.*

No instances : $\lambda_2 \geq 1 - \alpha$.

Yes instances : $\lambda_2 \leq 1 - \beta$.

In what follows, we shall see that this problem captures the hardness of probabilistic computation in logarithmic space.

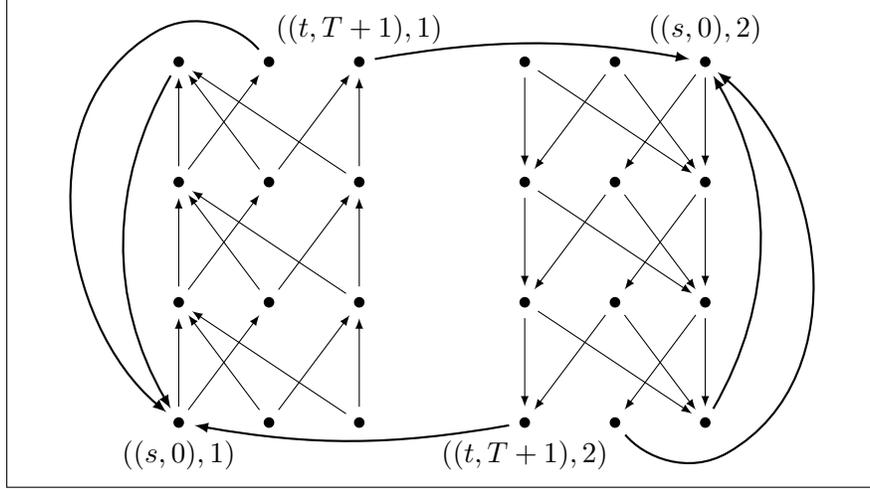


Figure 1: An example for G'' with $T = 3$.

3.1 Reducing Turing machines to graphs

Let M be a probabilistic space-bounded TM that accepts some language L with $c \log n$ space and $T = n^c$ time. Fix an input x . Let $G = (V, E)$ be the configuration graph of G on x , s the initial configuration and t the unique accepting configuration.

Define the *layered configuration graph* $G' = (V', E')$. $V' = V \times \{1, \dots, T+1\}$, i.e., G' has $T+1$ layers, each layer containing $|V|$ vertices. There is an outgoing edge from vertex (i, ℓ) in layer $\ell \leq T$ to vertex $(j, \ell+1)$ in layer $\ell+1$ if and only if $(i, j) \in E$. Without loss of generality, any vertex (i, ℓ) in layer $\ell \leq T$ has at least one neighbor in layer $\ell+1$.

Next, we define a graph $G'' = G''(M, x)$ that contains two copies of G' with some additional edges and has two strongly connected components if and only if s is not connected to t in G . Formally, the vertex set of G'' is $V'' = \{(v', k) \mid v' \in V', k \in \{1, 2\}\}$. The edges E'' of G'' are as follows:

- For every $(a, b) \in E'$ and $k \in \{1, 2\}$, $((a, k), (b, k)) \in E''$.
- For every $r \neq t$, $r \in V$ and $k \in \{1, 2\}$, $((r, T+1), k), ((s, 1), k) \in E''$.
- $((t, T+1), 1), ((s, 1), 2) \in E'$ and $((t, T+1), 2), ((s, 1), 1) \in E''$.

See Figure 1.

Let A be the transition matrix of G'' and note that it is stochastic. We claim:

Lemma 3.1. *Let $p = p(M, x) = A^T[s, t]$ denote the probability that M accepts x . Let ω be a primitive $(T+1)$ -th root of unity. The eigenvalues of A are $\{0, \omega^k, (1-2p)^{\frac{1}{T+1}} \omega^k \mid 0 \leq k \leq T\}$.*

Proof. Denote $W = A^\dagger$. It is known that the eigenvalues of A and W coincide. Let v be any eigenvector of W with eigenvalue $\lambda \neq 0$.

Claim 3.1. *For every $s \neq j \in V$, $v[(j, 1), 1] = v[(j, 1), 2] = 0$*

Proof. The in-degree of $((j, 1), 1), (j, 1), 2)$ is zero and $\lambda \neq 0$. □

Denote $v[((s, 1), 1)] = x$ and $v[((s, 1), 2)] = y$.

Claim 3.2. For every $j \in V$ and $\ell \geq 0$, $W^{\ell-1}[j, s]$ is the probability M reaches j after taking $\ell - 1$ random steps from s . Then,

$$\begin{aligned} v[((j, \ell), 1)] &= \frac{x}{\lambda^{\ell-1}} W^{\ell-1}[j, s] \\ v[((j, \ell), 2)] &= \frac{y}{\lambda^{\ell-1}} W^{\ell-1}[j, s]. \end{aligned}$$

Proof. The proof is by induction. The case of $\ell = 1$ was already covered. Assume the claim holds for ℓ and fix a vertex $((j, \ell + 1), 1)$. Then:

$$\begin{aligned} (Wv)[((j, \ell + 1), 1)] &= \sum_{k:(k,j) \in E} v[((k, \ell), 1)] \cdot \frac{1}{d_{\text{out}}(k)} \\ &= \sum_{k:(k,j) \in E} \frac{x}{\lambda^{\ell-1}} W^{\ell-1}[k, s] \cdot \frac{1}{d_{\text{out}}(k)} \quad (\text{Hypotesis}) \\ &= \frac{x}{\lambda^{\ell-1}} W^{\ell}[j, s]. \end{aligned}$$

As $Wv = \lambda v$, the claim for $((j, \ell + 1), 1)$ follows. The other claim is similar. \square

Finally, the equation $Wv = \lambda v$ for the vertices $((s, 1), 1)$ and $((s, 1), 2)$ give us two equations. The equation for the vertex $((s, 1), 1)$ is

$$(Wv)[((s, 1), 1)] = \sum_{j \neq s} v[((j, T + 1), 1)] + v[((s, T + 1), 2)].$$

As $\sum_{j \in V} v[((j, T + 1), 1)] = \frac{(1-p)x}{\lambda^T}$ and $v[((s, T + 1), 2)] = \frac{py}{\lambda^T}$, we get the equation

$$\lambda x = \frac{(1-p)x}{\lambda^T} + \frac{py}{\lambda^T}. \quad (2)$$

Similarly, the equation for the vertex $((s, 1), 2)$ is

$$\lambda y = \frac{(1-p)y}{\lambda^T} + \frac{px}{\lambda^T}. \quad (3)$$

Adding the two equations we get:

$$\lambda(x + y) = \frac{1}{\lambda^T}(x + y).$$

If $x \neq -y$ we must have $\lambda^{T+1} = 1$. If $x = -y$ Equation (2) and Equation (3) are the same and give $\lambda^{T+1}x = (1 - 2p)x$. Since $x \neq 0$ (otherwise the whole vector v is zero) we have $\lambda^{T+1} = 1 - 2p$. Hence, all the eigenvalues belong to the declared set.

Finally, we notice that for $x = -y$ and λ such that $\lambda^{T+1} = 1 - 2p$ we can build a corresponding eigenvector, and similarly for $x \neq -y$ and λ s.t. $\lambda^{T+1} = 1$ we may take $x = y$ and the equations work. Hence, the eigenvalues are exactly the given set. \square

3.2 An RL-hard problem

In this subsection we show that $SecondEV_{\alpha,\beta}$ with $\alpha = 0$ and $\beta = \frac{1}{n}$, that is – distinguishing between the case that the eigenvalue 1 has algebraic multiplicity 2 and the case that the second eigenvalue is $\frac{1}{n}$ -far from 1, is RL-hard. We stress that we only prove RL-hardness and we do not know how to solve $SecondEV_{0,\frac{1}{n}}$ in RL. We prove:

Theorem 3.2. $SecondEV_{0,\frac{1}{n}}$ is RL-hard.

Proof. Let $L \in \text{RL}$ and let M be the probabilistic space-bounded TM that accepts it. Given an input x , let G, G', G'', ω and p be as above. If $x \in L$ then without loss of generality $\frac{2\pi^2}{T} \leq p < \frac{1}{2}$ and otherwise $p = 0$. The dimension of the matrix is $n = 2(T+1)|V|$.

Let A be the transition matrix of G'' of dimension n and consider $B = \frac{1}{2}A + \frac{1}{2}I_{n \times n}$. By Lemma 3.1, the eigenvalues of B are $\left\{ \frac{1}{2}, \frac{1}{2} + \omega^k, \frac{1}{2} + \frac{1}{2}(1-2p)^{\frac{1}{T+1}}\omega^k \mid 0 \leq k \leq T \right\}$. For $0 \leq k \leq T$, denote

$$\begin{aligned}\lambda_k &= \frac{1}{2} + \frac{1}{2}(1-2p)^{\frac{1}{T+1}}\omega^k \\ \mu_k &= \frac{1}{2} + \frac{1}{2}\omega^k.\end{aligned}$$

First, note that $|\mu_k|^2$ can be expressed as some constant independent of k plus $\frac{1}{2} \cos\left(\frac{2\pi k}{T+1}\right)$. Similarly, $|\lambda_k|^2$ can be expressed as some constant independent of k plus $\frac{(1-2p)^{\frac{1}{T+1}}}{2} \cos\left(\frac{2\pi k}{T+1}\right)$. For $k > 0$, $\cos\left(\frac{2\pi k}{T+1}\right)$ is maximized at the lowest angle, i.e., at $k = 1$. Thus, for $k \neq 0$, $|\mu_k| \leq |\mu_1|$ and $|\lambda_k| \leq |\lambda_1|$. By inspection, $\mu_0 = 1$, $\lambda_0 = \frac{1}{2} + \frac{1}{2}(1-2p)^{\frac{1}{T+1}}$ and $|\lambda_1| \leq |\mu_1| \leq 1 - \frac{\pi^2}{2T^2}$. One can verify that:

- If $p \geq \frac{2\pi^2}{T}$ then $\lambda_0 \leq 1 - \frac{\pi^2}{2T^2} \leq 1 - \frac{1}{n}$, so all eigenvalues but μ_0 are at most $1 - \frac{1}{n}$ in magnitude.
- If $p = 0$ then $\lambda_0 = 1$, so all eigenvalues but μ_0 and λ_0 are at most $1 - \frac{1}{n}$ in magnitude.

B then satisfies the promise of $SecondEV_{0,\frac{1}{n}}$. It is well-known that A can be computed using logarithmic space and so can B . Hence, the reduction is a logspace reduction. \square

3.3 A BPL-complete problem

While we do not know how to solve $SecondEV_{0,\frac{1}{n}}$ in RL, we *do* know how to solve it in BPL. In fact we can do it for even stricter parameters. Specifically, we show that $SecondEV_{\alpha,\beta}$ with $\alpha = \frac{1}{n}$ and $\beta = \frac{2}{n}$ is BPL-complete. We first show:

Lemma 3.3. $SecondEV_{\frac{1}{n},\frac{2}{n}}$ is in BPL.

Proof. Let A be the input to the promise problem and set $m = 4n \log n$. We then have:

Claim 3.3. If A is a Yes instance then $\text{Tr}(A^m) \leq 1 + \frac{1}{n^7}$.

Proof. $\text{Tr}(A^m) = \sum_{k=1}^n \lambda_k^m$.⁵ Then:

$$|\text{Tr}(A^m)| = \left| \sum_{k=1}^n \lambda_k^m \right| \leq 1 + \sum_{k=2}^n |\lambda_k|^m \leq 1 + n \left(1 - \frac{2}{n}\right)^m.$$

⁵This equation is true even when A is not diagonalizable. The λ_k are the roots of the characteristic polynomial. To see that, write A in its Jordan Normal Form, $A = VJV^{-1}$ where the eigenvalues of A , $\lambda_1, \dots, \lambda_n$, lie on the diagonal of J . $A^m = VJ^mV^{-1}$ so $\text{Tr}(A^m) = \text{Tr}(J^m)$ and since J is triangular, the diagonal of J^m is $\lambda_1^m, \dots, \lambda_n^m$.

Using the fact that $(1 - \frac{2}{n})^m \leq e^{-2m/n} = e^{-8 \log n} = \frac{1}{n^8}$, the claim thus follows. \square

Claim 3.4. *If A is a No instance then $\text{Tr}(A^m) \geq 1 + \frac{1}{n^5}$.*

Proof. By the reversed triangle inequality:

$$|\text{Tr}(A^m)| \geq 1 + \lambda_2^m - \left| \sum_{k=3}^n \lambda_k^m \right| \geq 1 + \left(1 - \frac{1}{n}\right)^m - n \left(1 - \frac{2}{n}\right)^m.$$

By the inequality $(1 - \frac{x}{n})^n \geq e^{-x} \left(1 - \frac{x^2}{n}\right)$ we have that $(1 - \frac{1}{n})^m \geq \frac{1}{n^4} \left(1 - \frac{4 \log n}{n}\right)$. Now, since $n \left(1 - \frac{2}{n}\right)^m \leq \frac{1}{n^7}$ we overall conclude that $|\text{Tr}(A^m)| \geq 1 + \frac{1}{n^5}$. \square

Approximating $\text{Tr}(A^m)$ well with high probability can be done in logarithmic space. By Lemma 2.2 and simple composition of logspace reductions, we can $(\frac{1}{n^8}, \frac{1}{3})$ -approximate $\text{Tr}(A^m)$ and decide the promise problem using $O(\log n)$ space. \square

The main result is the converse direction that the problem is BPL-complete.

Lemma 3.4. *Let $L \in \text{BPL}$. Then, L is logspace reducible to $\text{SecondEV}_{\frac{1}{n}, \frac{2}{n}}$.*

Proof. Let $L \in \text{BPL}$ and fix an input x . Let M, G, G', G'', ω, p be as before. Let A and $B = \frac{1}{2}A + \frac{1}{2}I_{n \times n}$ be as in the proof for Theorem 3.2. One can verify that:

- If $p \geq \frac{2\pi^2}{T}$ then $\lambda_0 \leq 1 - \frac{\pi^2}{2T^2} \leq 1 - \frac{2}{n}$, so all eigenvalues but μ_0 are at most $1 - \frac{2}{n}$ in magnitude.
- If $p \leq \frac{1}{4T}$ then $\lambda_0 \geq \frac{1}{2} + \frac{1}{2} \left(1 - \frac{1}{2T}\right)^{\frac{1}{T}} \geq 1 - \frac{1}{3T^2} \geq 1 - \frac{1}{n}$, so all eigenvalues but μ_0 and λ_0 are at most $1 - \frac{2}{n}$ in magnitude. The second inequality is derived from the fact that $(1 - \frac{x}{2})^x \geq 1 - \frac{2x^2}{3}$ for small $x > 0$.

Without loss of generality, p can be taken to be either larger than $\frac{2\pi^2}{T}$ or smaller than $\frac{1}{4T}$ (by simple BPL amplification arguments), so B satisfies the promise of $\text{SecondEV}_{\frac{1}{n}, \frac{1}{n}}$. \square

4 Approximating eigenvalues with constant accuracy

In this section we prove:

Theorem 4.1. *There exists a probabilistic algorithm that on input a stochastic matrix B with real eigenvalues in $[0, 1]$, constants $\beta > \alpha > 0$ and $\lambda \in [0, 1]$ such that:*

- *There are d eigenvalues λ_i satisfying $|\lambda - \lambda_i| \leq \alpha$,*
- *All other eigenvalues λ_i satisfy $|\lambda - \lambda_i| \geq \beta$,*

outputs d with probability at least $2/3$. Furthermore the algorithm runs in probabilistic space $O(\log n)$.

We remark that Theorem 4.1 covers the case of transition matrices of undirected graphs. As mentioned earlier, a transition matrix A of an undirected graph has an eigenvector basis with real eigenvalues in the range $[-1, 1]$. Taking $B = \frac{1}{2}A + \frac{1}{2}I_{n \times n}$ we get a stochastic matrix with eigenvalues in the range $[0, 1]$, and whose eigenvectors are in a natural one-to-one correspondence with A 's eigenvalues.

Proof. (Of Theorem 4.1) The input to the algorithm is $n, B, \lambda, \alpha, \beta$. We assume the existence of a univariate polynomial $p(x) = \sum_{i=0}^M c_i x^i$ with the following properties:

- p has a sharp peak around λ , i.e., $p(x) \geq 1 - \eta$ for $x \in [\lambda - \alpha, \lambda + \alpha]$ and $p(x) \leq \eta$ for $x \in [0, 1] \setminus (\lambda - \beta, \lambda + \beta)$, where $\eta = \eta(n) = n^{-2}$.
- p can be computed in L. Formally, $M = \deg(p)$ and $|c_i|$ are at most $\text{poly}(n)$ and for every i , c_i can be computed (exactly) by a deterministic Turing machine that uses $O(\log n)$ space.

In the next subsection we show how to obtain such a polynomial p with $M = 32(\beta - \alpha)^{-2} \log n$ and $|c_i| \leq 2^{O(M)}$.

Choose $\varepsilon = \frac{1}{n}$ and $\delta = \frac{1}{3}$. Set $\varepsilon' = \varepsilon \cdot 2^{-2M}$ and $\delta' = \delta \cdot 2^{-M}$. The output of the algorithm is the integer closest to

$$R = \sum_{i=0}^M c_i \cdot \text{TP}(B, n, i, \varepsilon', \delta')$$

where TP is the probabilistic algorithm guaranteed by Lemma 2.2 that (ε', δ') -approximates $\text{Tr}(B^i)$.

It is easy to check that:

Claim 4.1. $\Pr[|R - \text{Tr}(p(B))| \geq \varepsilon] \leq \delta$.

As $\text{Tr}(p(B)) = \sum_{i=1}^n p(\lambda_i)$, $\Pr[|R - \sum_{i=1}^n p(\lambda_i)| \geq \varepsilon] \leq \delta$. However, $p(\lambda_i)$ is large when λ_i is α -close to λ and small when it is β -far from λ , and we are promised that *all* eigenvalues λ_i are either α -close or β -far from λ . Thus,

$$|\text{Tr}(p(B)) - d| \leq n\eta.$$

Altogether, except for probability δ , $|R - d| \leq \varepsilon + n\eta \leq \frac{1}{3}$, and the nearest integer closest to R is d . The correctness follows. It is also straightforward to check that the space complexity is $O(\log(n\varepsilon^{-1}\delta^{-1})) = O(\log n)$. \square

The constant accuracy we achieve is far from being satisfying. The matrix B has n eigenvalues in the range $[0, 1]$, so the average distance between two neighboring eigenvalues is $1/n$. Thus, the assumption that there is an interval of length $\beta - \alpha$ with no eigenvalue is often not true. The desired accuracy we would like to get is $o(1/n)$. Having such accuracy would enable outputting an approximation of the whole spectrum of B , using methods similar to those in [19], thus getting a true classical analogue to the quantum algorithm in [19]. However, we do not know how to achieve subconstant accuracy. The question whether better accuracy is possible in BPL is one of the main questions raised by this work.

4.1 Using the symmetric threshold functions

There are several natural candidates for the function p above. In this subsection we use the threshold function to obtain such a function p . For $\lambda = \frac{k}{M}$ for some integers k and M , define:

$$p_\lambda(x) = \sum_{i=k}^M \binom{M}{i} x^i (1-x)^{M-i}.$$

p_λ approximates well the threshold function $\text{Th}_\lambda(x) : [0, 1] \rightarrow \{0, 1\}$ that is one for $x \geq \lambda$ and zero otherwise. Specifically, using the Chernoff bound, we obtain:

Lemma 4.2. Let $x \in [0, 1]$. $p_\lambda(x)$ approximates $\mathbf{Th}_\lambda(x)$ over $[0, 1]$ with accuracy $(\xi(\varepsilon))^{Mx}$, where $\varepsilon = \frac{\lambda-x}{x}$ and $\xi(\varepsilon) = \frac{e^\varepsilon}{(1+\varepsilon)^{1+\varepsilon}}$.

As a polynomial in x , $p_\lambda(x) = \sum_{i=0}^M c_i x^i$ with $c_i = (-1)^i \sum_{j=\lambda M}^i \binom{M}{j} \binom{M-j}{i-j} (-1)^j$ and therefore $|c_i| \leq \sum_{j=\lambda M}^i \binom{M}{j} \binom{M-j}{i-j} \leq M \binom{M}{M/2}^2 = 2^{O(M)}$. Furthermore, c_i can be computed (exactly) by a deterministic Turing machine that uses $O(M)$ space by simply running through the loop over j , each time updating the current result by $(-1)^j \binom{M}{j} \binom{M-j}{i-j}$.

To obtain our polynomial p , define p as the difference between the threshold polynomial around $\lambda + \Delta$ and the threshold polynomial around $\lambda - \Delta$,

$$p(x) = p_{\lambda-\Delta}(x) - p_{\lambda+\Delta}(x)$$

where $M = 32(\beta - \alpha)^{-2} \log n$ and $\Delta = (\alpha + \beta)/2$. It is easy to check that:

Lemma 4.3. $p(x) \geq 1 - n^{-2}$ for every x that is α -close to λ (i.e., $|x - \lambda| < \alpha$) and $p(x) \leq n^{-2}$ for every x that is β -far from λ (i.e., $|x - \lambda| \geq \beta$).

4.2 The limitation of the technique

In this subsection, we prove the accuracy of the above technique cannot be enhanced merely by choosing a different polynomial p . Approximating threshold functions by a polynomial is well-studied and well understood (see, for example, [17, 8, 5] and references therein). However, we need to adapt this work to our needs because we have an additional requirement that the magnitude of the polynomial's coefficients is small.

We start by formalizing the properties of p that were useful to us. We say that $\mathcal{P} = \{p_{\lambda,n}\}_{\lambda \in [0,1], n \in \mathbb{N}}$ is a family of polynomials if for every $\lambda \in [0, 1]$ and $n \in \mathbb{N}$, $p_{\lambda,n}$ is a univariate polynomial with coefficients in \mathbb{R} .

Definition 4.1 (Small family). Let \mathcal{P} be a family of polynomials and fix $\lambda \in [0, 1]$. For every $n \in \mathbb{N}$, write $p_{\lambda,n}(x) = \sum_{i=0}^{\deg(p_{\lambda,n})} c_{\lambda,n,i} x^i$. We say the family is $s(n)$ -small if,

- $\deg(p_{\lambda,n}) \leq 2^{s(n)}$,
- For every $0 \leq i \leq \deg(p_{\lambda,n})$, $|c_{\lambda,n,i}| \leq 2^{s(n)}$, and
- There exists a deterministic Turing machine running in space $s(n)$ that outputs $c_{\lambda,n,0}, \dots, c_{\lambda,n,\deg(p_{\lambda,n})}$.

Definition 4.2 (Distinguisher family). Let \mathcal{P} be a family of polynomials and fix $n \in \mathbb{N}$. Given $\alpha < \beta$ in $(0, 1)$ and $\eta < 1/2$, we say the family is (α, β, η) -distinguisher for $\lambda \in [0, 1]$ if,

- For every $x \in [0, 1]$ that is α -close to λ , $p_{\lambda,n}(x) \in [1 - \eta, 1]$, and
- For every $x \in [0, 1]$ that is β -far from λ , $p_{\lambda,n}(x) \in [0, \eta]$.

Theorem 4.4. Let $\alpha, \beta, \lambda, \eta$ be such that $\alpha \leq \beta$, $\beta = o(1)$, $\eta = o(n^{-1})$ and $\lambda + \beta \leq \frac{1}{2}$. Then there is no (α, β, η) -distinguisher family for λ that is $O(\log n)$ -small.

Proof. Assume there exists such a family $\{p_{\lambda,n}\}_{\lambda \in [0,1], n \in \mathbb{N}}$ with $s(n) = c' \log n$. We first show that we can assume p has logarithmic degree. Let $r_{\lambda,n}(x)$ be the residual error of truncating $p_{\lambda,n}(x)$ after $c \log n$ terms, for c that will soon be determined. Also, without loss of generality, assume $x \in [0, 1)$ is bounded away from 1. Then:

$$r_{\lambda,n}(x) \leq \sum_{i=c \log n+1}^{\deg(p_{\lambda,n})} |c_{\lambda,n,i}| \cdot x^i \leq n^{c'} \cdot \frac{x^{c \log n}}{1-x} \leq \frac{1}{1-x} n^{c'-c \log(1/x)}.$$

So, by taking $c = \lceil \frac{c'+2-\log(1-x)}{\log(1/x)} \rceil$ we obtain $r_{\lambda,n}(x) \leq n^{-2}$.

We now show that $O(\log n)$ -degree polynomials cannot decay around λ fast enough. Assume to the contrary that there exists such a distinguisher family, so $|p_{\lambda,n}(x)| < n^{-1}$ for $x \in [\lambda + \beta, 1]$. The following lemma states that if a function has a small value on an interval, than it cannot be too large outside it. Namely,

Lemma 4.5. [20, Theorem 2.9.11] *Let $T_n(x)$ be the Chebyshev polynomial (of the first kind) of degree n . Then, if the polynomial $P_n(x) = \sum_{i=0}^n c_i x^i$ satisfies the inequality $|P_n(x)| \leq L$ on the segment $[a, b]$ then at any point outside the segment we have*

$$|P_n(x)| \leq L \cdot \left| T_n \left(\frac{2x - a - b}{b - a} \right) \right|.$$

For properties of the Chebyshev polynomials see [16, Chapter 1.1]. We mention a few properties that we use. An explicit representation of $T_n(x)$ is given by $T_n(x) = \frac{(x - \sqrt{x^2 - 1})^n + (x + \sqrt{x^2 - 1})^n}{2}$. $|T_n(-x)| = |T_n(x)|$ and T_n is monotonically increasing for $x > 1$. Also,

$$|T_n(1 + \delta)| \leq \left(1 + \delta + \sqrt{(1 + \delta)^2 - 1} \right)^n \leq \left(1 + 4\sqrt{\delta} \right)^n \leq e^{4n\sqrt{\delta}} \leq 2^{8n\sqrt{\delta}} \quad (4)$$

for $0 \leq \delta \leq 1$. Then:

$$\begin{aligned} |p_{\lambda,n}(\lambda)| &\leq n^{-1} \cdot \left| T_{c \log n} \left(\frac{\lambda - \beta - 1}{-\lambda - \beta + 1} \right) \right| \\ &= n^{-1} \cdot \left| T_{c \log n} \left(1 + \frac{2\beta}{1 - \lambda - \beta} \right) \right| \quad \text{By } |T_n(x)| = |T_n(-x)| \\ &\leq n^{-1} \cdot |T_{c \log n}(1 + 4\beta)| \quad \text{By the monotonicity of } T_n(x) \text{ for } x > 1 \text{ and } \lambda + \beta \leq \frac{1}{2} \end{aligned}$$

By Equation (4) $|p_{\lambda,n}(\lambda)| \leq n^{-1} 2^{32c\sqrt{\beta} \log n} \leq n^{-1+32c\sqrt{\beta}}$. As $\beta = o(1)$ for n large enough we have $|p_{\lambda,n}(\lambda)| \leq n^{-1/2}$, contradicting the fact that $|p_{\lambda,n}(\lambda)| \geq 1 - n^{-1}$. \square

We note that for values very close to 1, polynomials of higher degrees are useful, and indeed better approximations are possible. In particular, one can separate a 1 eigenvalue from $1 - \frac{1}{n}$ by using the polynomial x^{n^2} .

5 Acknowledgements

The first and third authors were supported by the Israel science Foundation grant no. 994/14 and by the United States – Israel Binational Science Foundation grant no. 2010120. The third author was also supported by the Blavatnik Fund.

A preliminary version of this paper appeared in [7] but without the BPL-completeness and RL-hardness of approximating the second eigenvalue of a stochastic operator.

References

- [1] Romas Aleliunas, Richard M. Karp, R.J. Lipton, Laszlo Lovasz, and C. Rackoff. Random walks, universal traversal sequences, and the complexity of maze problems. In *Foundations of Computer Science, 1979., 20th Annual Symposium on*, pages 218–223, Oct 1979.
- [2] Michael Ben-Or and Lior Eldar. Optimal algorithms for linear algebra by quantum inspiration. *CoRR*, abs/1312.3717, 2013.
- [3] Stephen A. Cook. A taxonomy of problems with fast parallel algorithms. *Information and Control*, 64(1–3), 1985. International Conference on Foundations of Computation Theory.
- [4] L. Csanky. Fast parallel matrix inversion algorithms. *SIAM Journal of Computing*, 5(6):618–623, 1976.
- [5] Ilias Diakonikolas, Parikshit Gopalan, Ragesh Jaiswal, Rocco A Servedio, and Emanuele Viola. Bounded independence fools halfspaces. *SIAM Journal on Computing*, 39(8):3441–3462, 2010.
- [6] Dean Doron and Amnon Ta-Shma. On the de-randomization of space-bounded approximate counting problems. *Information Processing Letters*, 2015.
- [7] Dean Doron and Amnon Ta-Shma. On the problem of approximating the eigenvalues of graphs in probabilistic logspace. In *Automata, Languages, and Programming*. Springer, 2015.
- [8] Alexandre Eremenko and Peter Yuditskii. Uniform approximation of $\text{sgn}(x)$ by polynomials and entire functions. *Journal d'Analyse Mathématique*, 101(1):313–324, 2007.
- [9] James Allen Fill. Eigenvalue bounds on convergence to stationarity for nonreversible markov chains, with an application to the exclusion process. *The annals of applied probability*, pages 62–87, 1991.
- [10] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Phys. Rev. Lett.*, 103:150502, Oct 2009.
- [11] K. Hoffman. *Banach Spaces of Analytic Functions*. Dover Books on Mathematics Series. Dover Publications, Incorporated, 2007.
- [12] Milena Mihail. Conductance and convergence of markov chains—a combinatorial treatment of expanders. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 526–531. IEEE Comput. Soc. Press, 1989.
- [13] Noam Nisan. Pseudorandom generators for space-bounded computation. *Combinatorica*, 12(4):449–461, 1992.
- [14] Omer Reingold. Undirected connectivity in log-space. *J. ACM*, 55(4), 2008.
- [15] Omer Reingold, Luca Trevisan, and Salil Vadhan. Pseudorandom walks on regular digraphs and the rl vs. l problem. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 457–466. ACM, 2006.
- [16] T.J. Rivlin. *The Chebyshev polynomials*. Pure and applied mathematics. Wiley, 1974.

- [17] E. B. Saff and V. Totik. Polynomial approximation of piecewise analytic functions. *Journal of the London Mathematical Society*, s2-39(3):487–498, 1989.
- [18] Michael E. Saks and Shiyu Zhou. $\text{BP}_{\text{H}}\text{SPACE}(S) \subseteq \text{DSPACE}(S^{3/2})$. *J. Comput. Syst. Sci.*, 58(2):376–403, 1999.
- [19] Amnon Ta-Shma. Inverting well conditioned matrices in quantum logspace. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, STOC '13, pages 881–890, New York, NY, USA, 2013. ACM.
- [20] A.F. Timan. *Theory of Approximation of Functions of a Real Variable*. Dover books on advanced mathematics. Pergamon Press, 1963.
- [21] Dieter van Melkebeek and Thomas Watson. Time-space efficient simulations of quantum computations. *Electronic Colloquium on Computational Complexity (ECCC)*, 17:147, 2010.
- [22] John Watrous. Space-bounded quantum complexity. *Journal of Computer and System Sciences*, 59(2):281 – 326, 1999.

A Simulatable matrices

Lemma A.1. *The family of transition matrices of (directed or undirected) graphs is simulatable.*

Proof. (of Lemma 2.1) Let $G = (V, E)$ be a graph with n vertices and let A be its transition matrix. Let $k \in \mathbb{N}$, $s, t \in [n]$ and $\delta, \varepsilon > 0$. Consider the algorithm that on input k, s, t , takes T independent random walks of length k over G starting at vertex s . The algorithm outputs the ratio of walks that reach vertex t . Let Y_i be the random value that is 1 if the i -th trial reached t and 0 otherwise. Then, for every i , $\mathbb{E}[Y_i] = A^k[s, t]$. Also, Y_1, \dots, Y_T are independent. By Chernoff,

$$\Pr\left[\left|\frac{1}{T} \sum_{i=1}^T Y_i - A^k[s, t]\right| \geq \varepsilon\right] \leq 2e^{-2\varepsilon^2 T}$$

Taking $T = \text{poly}(\varepsilon^{-1}, \log \delta^{-1})$, the error probability (i.e., getting an estimate that is ε far from the correct value) is at most δ . Altogether, the algorithm runs in space $O(\log(Tnk|E|)) = O(\log(nk\varepsilon^{-1}) + \log \log \delta^{-1})$, assuming $|E| = \text{poly}(n, k)$. \square

We say $\|A\|_{\infty} \leq c$ if for every $i \in [n]$, $\sum_j |A[i, j]| \leq c$. We show:

Lemma A.2. *The family of real matrices with infinity norm at most 1 is simulatable.*

Proof. (of Lemma 2.2) We prove the result to real matrices, with positive or negative entries, as long as they have bounded infinity norm. By generalizing the sign of an entry to its *phase*, the result easily applies to complex matrices as well.

Let A be a real matrix of dimension n such that $\|A\|_{\infty} \leq 1$. Let $d_i(A) = \sum_j |A[i, j]|$. Let $k \in \mathbb{N}$, $s, t \in [n]$ and $\delta, \varepsilon > 0$. Note that:

$$\begin{aligned} A^k[s, t] &= \sum_{i_1=1}^n \sum_{i_2=1}^n \cdots \sum_{i_{k-1}=1}^n A[s, i_1] \cdot A[i_1, i_2] \cdot \dots \cdot A[i_{k-1}, t] \\ &= \sum_{i_1=1}^n \sum_{i_2=1}^n \cdots \sum_{i_{k-1}=1}^n \frac{|A[s, i_1]|}{d_s(A)} \cdot \frac{|A[i_1, i_2]|}{d_{i_1}(A)} \cdot \dots \cdot \frac{|A[i_{k-1}, t]|}{d_{i_{k-1}}(A)} \cdot p(A, \langle s, i_1, i_2, \dots, i_{k-1}, t \rangle), \end{aligned}$$

where

$$p(A, \langle s, i_1, i_2, \dots, i_{k-1}, t \rangle) = \frac{d_s(A) \cdot d_{i_1}(A) \cdot \dots \cdot d_{i_{k-1}}(A)}{\text{sgn}(A[s, i_1] \cdot A[i_1, i_2] \cdot \dots \cdot A[i_{k-1}, i_t])}.$$

Consider the algorithm that on input k, s, t , takes T independent random walks of length k over G starting from vertex s . Iterating over all random walks, the algorithm approximates $\frac{1}{T} \sum_i y(i)$, where $y(i) = p(A, i)$ if the walk i reached t , and 0 otherwise. Correspondingly, let Y_i be the random value that is $p(A, i)$ if the i 'th walk reached t and 0 if it did not. Then,

$$\mathbb{E}[Y_i] = \sum_{i_1=1}^n \sum_{i_2=1}^n \dots \sum_{i_{k-1}=1}^n A[s, i_1] \cdot A[i_1, i_2] \cdot \dots \cdot A[i_{k-1}, t] \cdot p(A, \langle s, i_1, \dots, i_{k-1}, t \rangle) = A^k[s, t].$$

Denote the algorithm's outcome by $M(k, s, t)$. As in Lemma 2.1, and using the fact that $|p(A, i)| \leq 1$, the algorithm can (ε, δ) -approximates $\mathbb{E}[Y_i]$ by choosing T which is $\text{poly}(\varepsilon^{-1}, \log \delta^{-1})$. Following the same analysis as of Lemma 2.1, the algorithm runs in $O(\log nk\varepsilon^{-1} + \log \log \delta^{-1})$ space. We conclude that A is simulatable. \square