How to Sample From The Limiting Distribution of a Continuous-Time Quantum Walk

Javad Doliskani[*]

Abstract

We introduce ε-projectors, using which we can sample from limiting distributions of continuous-time quantum walks. The standard algorithm for sampling from a distribution that is close to the limiting distribution of a given quantum walk is to run the quantum walk for a time chosen uniformly at random from a large interval, and measure the resulting quantum state. This approach usually results in an exponential running time.

We show that, using ε-projectors, we can sample exactly from the limiting distribution. In the black-box setting, where we only have query access to the adjacency matrix of the graph, our sampling algorithm runs in time proportional to Δ^{-1}, where Δ is the minimum spacing between the distinct eigenvalues of the graph. In the non-black-box setting, we give examples of graphs for which our algorithm runs exponentially faster than the standard sampling algorithm.

Contents

---

[*]Department of Computer Science, Ryerson University, `javad.doliskani@ryerson.ca`.

1 Introduction

Continuous-time quantum walks, first considered by Farhi and Gutmann [21], are quantum ana-
logues of continuous-time classical random walks. The dynamics of a continuous-time classical walk
on an undirected graph Γ is described by the differential equation

$$\frac{d}{dt}p(t) = Lp(t) \tag{1}$$

where $p(t)$ is the state of the walk at time t and L is the Laplacian of Γ. The entries of $p(t)$ are
indexed by the set of vertices of Γ. In the quantum walk on Γ, equation (1) is replaced by the
Schrödinger equation

$$i\frac{d}{dt}|\psi_t\rangle = H|\psi_t\rangle \tag{2}$$

where the Hamiltonian $H = L$, and $|\psi_t\rangle$ is a quantum state whose amplitudes encode a probability
distribution. Quantum walks have found many applications in quantum computing and quantum
information. It was shown by Childs [10] that universal quantum computation can be implemented
using quantum walks on low degree graphs. There are many algorithms based on quantum walk that
achieve polynomial speedup over classical algorithms, e.g., [13, 20, 3, 4]. There are also black-box
problems for which quantum walks achieve exponential speedup over classical algorithms [12, 14].

A classical continuous-time random walk has a unique stationary distribution, assuming its
underlying Markov chain is irreducible [34]. Regardless of the initial state, the walk converges to
this stationary distribution as $t \to \infty$, and therefore, the stationary distribution is the same as the
limiting distribution. This, however, does not hold for quantum walks, since quantum evolutions are
unitary and preserve distance. Nevertheless, one can define a time-averaged probability distribution
of a quantum walk by choosing a time $t \in [0, T]$ uniformly at random, running the walk for a total
time t and measuring the resulting state. When $T \to \infty$, this distribution converges to a limiting
distribution which is what we will consider in this paper. This limiting distribution generally
depends on the initial state of the walk.

If the graph Γ is connected and simple, i.e., has no self loops and multiple edges, then the
limiting distribution of the classical walk on Γ is always the uniform distribution. In contrast,
the limiting distribution of the quantum walk on Γ depends on the initial state and is often not
uniform. For example, the quantum walk on the hypercube [40] or the Symmetric group [24], or
more generally G-circulant graphs [1] does not converge to the uniform distribution.

The mixing time M_δ of a quantum (or classical) walk is the minimum time after which the
distribution of the walk is within distance δ of the limiting distribution. The mixing time of a
classical walk depends inversely on the spectral gap of the transition matrix of the walk, while for a
quantum walk, the mixing time depends inversely on the minimum gap between all pairs of distinct
eigenvalues of the Hamiltonian H. The mixing time of quantum walks have been studied for specific
graphs such as hypercubes, cycles and lattices [40, 22, 41], and for Erdős-Rényi random graphs [9, 8].
The bounds on the quantum mixing time for some graphs imply quadratic speedup over classical
walks, while for some other graphs these bounds are larger than their classical counterparts.

The problem of sampling from the limiting distribution of a classical walk is an important
problem and has been the focus of much research. The underlying randomness strategy in many
algorithms of practical interest reduces to sampling from a limiting distribution. Much like in the
classical case, sampling from the limiting distribution of a quantum walk is of both practical and
theoretical interest. For example, Richter [41] proposed a "double-loop" quantum walk algorithm
for sampling from a distribution that is close to the uniform distribution over a given graph Γ. The
inner loop in Richter's algorithm samples from a distribution that is close to the limiting distribution

Π of the quantum walk. Therefore, an efficient algorithm for sampling from Π results in an efficient algorithm for sampling uniformly from the vertex set of Γ. The black-box graph problem proposed by Childs et al. [12], which is interesting from a theoretical perspective, is based on sampling from the limiting distribution of a quantum walk. They prove that no classical algorithm can efficiently solve the proposed problem.

Given the importance of sampling from the limiting distribution of a quantum walk, it is natural to ask whether there are efficient algorithms for sampling from such distributions, at least for specific graphs. There has not been much research explicitly addressing this question. The standard way of sampling from the limiting distribution of a quantum walk is by mixing: set a large value for T, run the walk for a uniformly random time $t \in [0, T]$, and measure. Chakraborty et al. [8] considered sampling over the Erdős-Rényi graphs by mixing. They obtained an upper bound on the mixing time of these graphs through analyzing their spectrum. Their bound implies an exponential time sampling algorithm.

1.1 Sampling without mixing

In this work, we propose an algorithm for sampling from the limiting distribution of a given continuous-time quantum walk, that is not based on mixing. The idea behind our algorithm is to uniquely "tag" the eigenspaces of the adjacency matrix using polynomially long binary strings. More precisely, give a graph Γ with N vertices and adjacency matrix A, let $|\phi_j\rangle$ be an eigenstate of A that belongs to an eigenspace $\mathcal{X}_j$ of A. Then, we will see that sampling from the limiting distribution Π on Γ reduces to performing the transform

$$|0\rangle|\phi_j\rangle \mapsto |t_j\rangle|\phi_j\rangle \tag{3}$$

where t_j is a string of length $\text{poly}(\log N)$ that uniquely identifies $\mathcal{X}_j$.

To perform the transform (3), we introduce the general idea of ε-projectors. Informally, an ε-projector for the adjacency matrix A is a set of hermitian matrices that have the same eigenspaces as A, and can be efficiently simulated as Hamiltonians. Moreover, the set of matrices in an ε-projector have to satisfy a separation condition with respect to their eigenvalues. A set of matrices satisfying such a separation condition is called an ε-separated set. A specific ε-projector was first implicitly used by Kane, Sharif and Silverberg [29] for constructing a quantum money scheme based on quaternion algebras. Their ε-projector is a set of sparse Brandt matrices, which they use to verify an alleged bill.

Technique. Given an ε-projector $\mathcal{A}$ for A, we use phase estimation to store, in a separate register, an estimate of the eigenvalues of each operator in $\mathcal{A}$. More precisely, we perform the transform

$$|0\rangle|\phi_j\rangle \mapsto |\tilde{\lambda}_{1,j}\rangle|\tilde{\lambda}_{2,j}\rangle \cdots |\tilde{\lambda}_{r,j}\rangle|\phi_j\rangle$$

where the $\tilde{\lambda}_{k,j}$ are approximate eigenvalues corresponding to the eigenstate $|\phi_j\rangle$. It follows from the ε-separatedness of $\mathcal{A}$ that the vectors $\boldsymbol{\lambda}_j = (\tilde{\lambda}_{1,j}, \ldots, \tilde{\lambda}_{r,j})$ uniquely identify the eigenspaces of A. This means the binary representation of the $\boldsymbol{\lambda}_j$ can be used as the binary strings t_j in (3). Therefore, efficient sampling from the limiting distribution Π reduces to finding a good ε-projector for A. In general, a good ε-projector for A is an ε-projector for which ε^{-1} and r are both at most $\text{poly}(\log N)$. In specific cases where the operators in the ε-projector can be simulated efficiently even for large powers, it is not necessary for ε^{-1} to be bounded by $\text{poly}(\log N)$.

Black-box vs non-black-box. When the graph Γ is given as a black-box, we can only see the structure of Γ locally. In other words, we can only access the nonzero entries of each row of the adjacency matrix A of Γ through queries to an oracle. When we are restricted to query access to A, the global structure of Γ is not known, and we do not have any other information on A as an operator. In this case, we are essentially left with one choice for an ε-projector for A: A itself. Consequently, we always have $\varepsilon \leq \Delta$, where Δ is the minimum distance between any two distinct eigenvalues of A. The complexity of our sampling algorithm will then be bounded below by a multiple of Δ^{-1}.

In the non-black-box setting, we often have some knowledge of the global structure of Γ that enables us to find nontrivial ε-projectors for A. In this paper, we give two examples of graphs for which we can find good ε-projectors, Winnie Li graphs and Supersingular Isogeny graphs.

Winnie Li graphs [35] are special cases of quasi-abelian graphs which are a subclass of Cayley graphs. We give a general strategy for finding ε-projectors for quasi-abelian graphs, and apply it to Winnie Li graphs. Without the use of an ε-projector, one could sample from the limiting distribution of quantum walks on Winnie Li graphs using two different methods. The first method is by mixing, which takes exponential time because the eigenvalues of the (normalized) adjacency matrix are very close. The second method is to use the quantum Fourier transform. For that, we need to be able to approximate the eigenvalues of A efficiently. When the dimension of the underlying space is odd, these eigenvalues are multiples of some exponential sums called Kloosterman sums. There is no known efficient classical or quantum algorithm for approximating these sums. We will see that using a specific ε-projector, we can efficiently sample from the limiting distributions on these graphs.

A supersingular isogeny graphs is a regular graph in which the set of vertices is the set of supersingular elliptic curves and the edges are isogenies between these curves. Isogeny graphs have found many applications in cryptography [36, 37]. The adjacency matrices of these graphs are called Hecke operators. The minimum distance between the eigenvalues of a Hecke operator is exponentially small, so the quantum mixing time for these graphs is exponentially large. It is known that the set of Hecke operators form a commutative algebra over $\mathbb{C}$. Using this fact, and assuming some standard heuristics, we will see that a small set of these operators form an ε-projector with high probability. Using this ε-projector, we can efficiently sample from the limiting distribution on these graphs. As an application, the sampling algorithm can be used to generate *honest hard curves*. There is no known classical algorithm for efficiently generating such curves.

2 Preliminaries

2.1 Continuous-time quantum walk

Let $\Gamma = (V, E)$ be an undirected graph with $N = |V|$ vertices, and let $\mathcal{X} = \mathbb{C}^V$ be the complex euclidean space with basis V. We will refer to this basis as the vertex basis and denote its elements by $|v\rangle$, $v \in V$. Let A be the adjacency matrix of Γ. The continuous-time quantum walk on Γ is described by the differential equation (2) where the Hamiltonian is the Laplacian of Γ. Another common choice (which we also use in this paper) for the Hamiltonian of the walk is the adjacency matrix A. Then, the continuous-time quantum walk on G at time t is defined by the operator

$$W(t) = e^{iAt}$$

on $\mathcal{X}$. For an initial quantum state $|\psi_0\rangle$ and a real number $T > 0$, define the following probability distribution on V: choose $t \in [0, T]$ uniformly at random, evolve the state $|\psi_0\rangle$ under $W(t)$, i.e.,

compute $W(t)|\psi_0\rangle$, and measure the resulting state in the vertex basis. The probability of measuring a vertex $v \in V$ is

$$P_T(v|\psi_0) = \frac{1}{T}\int_0^T |\langle v|W(t)|\psi_0\rangle|^2 dt. \tag{4}$$

Let $\{|\phi_j\rangle\}_{1\leq j\leq N}$ be a set of eigenstates of A that form an orthonormal basis for $\mathcal{X}$, and let $\{\lambda_j\}_{1\leq j\leq N}$ be the set of corresponding eigenvalues. Let $\{\mathcal{X}_j\}_{1\leq j\leq M}$, where $M \leq N$, be the set of eigenspaces of A, and define $I_j = \{k : |\phi_k\rangle \in \mathcal{X}_j\}$. Therefore, I_j is the set of indices k for which the eigenstates $|\phi_k\rangle$ correspond to the same eigenvalue. A straightforward calculation shows that

$$P_T(v|\psi_0) = \sum_{j=1}^M \left|\sum_{k\in I_j}\langle v|\phi_k\rangle\langle\phi_k|\psi_0\rangle\right|^2 +$$
$$\sum_{k,\ell=1:\lambda_k\neq\lambda_\ell}^N \langle v|\phi_k\rangle\langle\phi_k|\psi_0\rangle\langle\psi_0|\phi_\ell\rangle\langle\phi_\ell|v\rangle\frac{e^{i(\lambda_k-\lambda_\ell)T}-1}{i(\lambda_k-\lambda_\ell)T} \tag{5}$$

Letting $T \to \infty$, the second term in the above expansion will vanish, and we get the distribution

$$P_\infty(v|\psi_0) := \lim_{T\to\infty} P_T(v|\psi_0) = \sum_{j=1}^M \left|\sum_{k\in I_j}\langle v|\phi_k\rangle\langle\phi_k|\psi_0\rangle\right|^2. \tag{6}$$

This is called the *limiting distribution* of the quantum walk $W(t)$. Given a real number $\delta \geq 0$, the mixing time M_δ of the walk $W(t)$, with respect to the initial state $|\psi_0\rangle$, is defined as

$$M_\delta = \min\{T' : \|P_T(\cdot|\psi_0) - P_\infty(\cdot|\psi_0)\|_1 \leq \delta, \forall T \geq T'\}, \tag{7}$$

where $P_T(\cdot|\psi_0)$ and $P_\infty(\cdot|\psi_0)$ are the probability vectors defined by (4) and (6), respectively. Denote by Δ the minimum distance between all pairs of distinct eigenvalues of A, i.e.,

$$\Delta = \min_{\lambda_k\neq\lambda_\ell} |\lambda_k - \lambda_\ell|, 1 \leq k, \ell \leq N. \tag{8}$$

Using the same analysis as in [2], it can be shown that

$$\|P_T(\cdot|\psi_0) - P_\infty(\cdot|\psi_0)\|_1 \leq \frac{2\ln M + 2}{T\Delta}. \tag{9}$$

A proof of this bound is given in Appendix A for completeness. From the definition of M_δ and the bound (9), we see that

$$M_\delta \leq \frac{2\ln M + 2}{\delta\Delta}. \tag{10}$$

2.2 Representation theory

For an introduction to representation theory see [15, 45]. Let V be a $\mathbb{C}$-vector space of finite dimension, and let $\mathrm{GL}(V)$ be the group of automorphisms of V. Let G be a finite group. A linear representation of G in V is a homomorphism of groups $\rho : G \to \mathrm{GL}(V)$. The degree of ρ, denoted by d_ρ, is the dimention of V as a $\mathbb{C}$-vector space. The character of ρ is the function $\chi_\rho : G \to \mathbb{C}$ defined by $\chi_\rho(a) = \mathrm{Tr}\,\rho(a)$. A morphism of representations $\rho_1 : G \to \mathrm{GL}(V_1)$ and $\rho_2 : G \to \mathrm{GL}(V_2)$ is a $\mathbb{C}$-linear map $\phi : V_1 \to V_2$ such that for every $a \in G$ the diagram

$$V_1 \xrightarrow{\rho_1(a)} V_2$$

$$\downarrow \phi \qquad\qquad \downarrow \phi$$

$$V_2 \xrightarrow{\rho_2(a)} V_2$$

commutes. The representations ρ_1 and ρ_2 are said to be isomorphic if ϕ is an isomorphism. A subrepresentation of ρ is a representation $\rho^W : G \to \mathrm{GL}(W)$ where $W \subseteq V$ is linear subspace. A representation that has no subrepresentations except for $W = 0, V$ is called an irreducible representation. We denote by $\widehat{G}$ the set of isomorphism classes of irreducible representations of G. Any representation ρ of G can be decomposed as a direct sum of irreducible representations: If $\varrho_1, \ldots, \varrho_k$ is a complete set of irreducible representations of G then $\rho = n_1 \varrho_1 \oplus \cdots \oplus n_k \varrho_k$ for some integers $n_j \geq 0$. Here, $n_j \varrho_j$ means a direct sum of n_j copies of ϱ_j.

Let $V = \mathbb{C}^G$, i.e., the space of functions $f : G \to \mathbb{C}$. The (left) regular representation of G is a representation ρ_{reg} in V defined by $\rho_{\mathrm{reg}}(s)f(a) = f(s^{-1}a)$ for any function $f \in V$. The regular representation decomposes as

$$\rho_{\mathrm{reg}} \cong \bigoplus_{\varrho \in \widehat{G}} d_\varrho \varrho, \tag{11}$$

which also shows that $\sum_\varrho d_\varrho^2 = |G|$.

A representation ρ of G is called unitary if $\rho(a)$ is a unitary matrix for all $a \in G$. Given any representation ρ of G, there always exists an inner product on V with respect to which ρ is unitary. Therefore, in this paper, we assume that all representations are unitary. In particular, any unitary representation can be decomposed as a sum of unitary irreducible representations.

The Fourier transform of a function $f : G \to \mathbb{C}$ at a representation $\varrho \in \widehat{G}$ is defined by

$$\widehat{f}(\varrho) = \sqrt{\frac{d_\varrho}{|G|}} \sum_{a \in G} \varrho(a) f(a).$$

The Fourier transform of f is given by $\oplus_\varrho \widehat{f}(\varrho)$. The quantum Fourier transform of a state $|\psi\rangle = \sum_{x \in G} \alpha_x |x\rangle$ is given by

$$|\widehat{\psi}\rangle = \sum_{\varrho \in \widehat{G}} \sum_{1 \leq j,k \leq d_\varrho} \widehat{\alpha}(\varrho)_{j,k} |\varrho, j, k\rangle \tag{12}$$

where $\alpha : G \to \mathbb{C}$ is defined by $\alpha(x) = \alpha_x$, and $\widehat{\alpha}(\varrho)_{j,k}$ is the (i, j) entry of the matrix $\widehat{\alpha}(\varrho)$.

3 Sampling Using ε-Projectors

Let $\Gamma = (V, E)$ be a graph with N vertices and let A be the adjacency matrix of Γ. Assume the same notation as in Section 2.1. A closer look at the sum in (6) suggests the following simple approach to sampling from the limiting distribution P_∞. Since $\{|\phi_j\rangle\}$ is an orthonormal basis for $\mathcal{X}$, given any initial state $|\psi_0\rangle$, we can always write

$$|\psi_0\rangle = \sum_{j=1}^{N} \langle \phi_j | \psi_0 \rangle |\phi_j\rangle.$$

Suppose we have a quantum algorithm Q that can uniquely "tag" the eigenspaces of A in the above superposition, using an extra register. More precisely, Q performs the following operation

$$\sum_{j=1}^{N} \langle \phi_j | \psi_0 \rangle |0\rangle |\phi_j\rangle \longmapsto \sum_{j=1}^{N} \langle \phi_j | \psi_0 \rangle |t_j\rangle |\phi_j\rangle = \sum_{j=1}^{M} \sum_{k \in I_j} \langle \phi_k | \psi_0 \rangle |t_j\rangle |\phi_k\rangle \tag{13}$$

6

where the strings t_j are unique with respect to the eigenspaces of A. If we measure the second register in the vertex basis, we obtain a vertex $v \in V$ with the probability given by (6).

A naive choice for the tags t_j are the eigenvalues of A. These eigenvalues can be approximated using phase estimation on the walk operator $W(t)$. However, to be able to uniquely identify the eigenspaces of A, one might need to compute the eigenvalues with exponential accuracy. Any such computation generally takes exponential time unless $W(t)$ can be applied efficiently for exponentially large t. In particular, if we treat A as a black-box, the complexity of performing (13) is going to be exponential in t. Therefore, any successful attempt at efficiently performing (13) will require some extra information or assumptions on A.

In the following we present the main idea of the paper, an algorithm for performing (13) that uses a specific set of operators that commute with A. We call such a set of operators an ε-*projector*. For many classes of graphs, we can find ε-projectors that enable us to efficiently perform (13). We need to adapt the definition of an ε-separated set from [29] to a set of operators.

Definition 3.1. For an integer $r > 0$, let $\mathcal{A} = \{A_j\}_{1 \leq j \leq r}$ be a set of hermitian operators, acting on $\mathcal{X}$, that have the same eigenspaces. For an eigenstate $|\phi_j\rangle$, let $\lambda_{1,j}, \lambda_{2,j}, \ldots, \lambda_{r,j}$ be the eigenvalues of the operators $A_1, A_2, \ldots, A_r$ associated with $|\phi_j\rangle$, respectively. Define the vector $\boldsymbol{\lambda}_j = (\lambda_{1,j}, \lambda_{2,j}, \ldots, \lambda_{r,j})$ for each $j = 1, \ldots, N$. For a real number $\varepsilon > 0$, the set of operators $\mathcal{A}$ is said to be ε-separated if

$$\|\boldsymbol{\lambda}_j - \boldsymbol{\lambda}_k\|_2 \geq \varepsilon, \text{ for all } \boldsymbol{\lambda}_j \neq \boldsymbol{\lambda}_k, 1 \leq j, k \leq N.$$

Definition 3.2. Let A be a hermitian operator on $\mathcal{X}$. An ε-projector for A is an ε-separated set $\mathcal{A} = \{A_j\}_{1 \leq j \leq r}$ such that for all $j = 1, \ldots, r$

- the walk $e^{iA_j t}$ can be performed in $F(t) \operatorname{poly}(\log N)$ operations, where $F(t) \in O(t)$, and

- A_j has the same eignenspaces as A.

The function $F(t)$ in Definition 3.2 determines how efficient the walks $e^{iA_j t}$ can be performed for different values of t. For an ε-projector we require that $F(t)$ be bounded above by a linear function in t. Also, here operations refer to elementary quantum gate operations. We now give an algorithm for sampling from the limiting distribution of the quantum walk on Γ. The algorithm takes as input an ε-projector for the adjacency matrix A.

Algorithm 1 (Sampling).
Input: An adjacency matrix A of a graph $\Gamma = (V, E)$, an ε-projector $\mathcal{A} = \{A_j\}_{1 \leq j \leq r}$ for A, an initial state $|\psi_0\rangle \in \mathcal{X}$.
Output: A sample from the limiting distribution of the walk $W(t) = e^{iAt}$ on Γ.

1. Perform phase estimation on the unitaries $e^{iA_1}, \ldots, e^{iA_r}$ and the input state $|\psi_0\rangle$ with accuracy $\varepsilon/2\sqrt{r}$, and store the approximate phases in extra registers. Denote by $\tilde{\lambda}_{k,j}$ the approximation of the eigenvalue $\lambda_{k,j}$ of A_k corresponding to the eigenstate $|\phi_j\rangle$. Then the resulting state of this step is

$$\sum_{j=1}^{N} \langle \phi_j | \psi_0 \rangle |\tilde{\lambda}_{1,j}\rangle |\tilde{\lambda}_{2,j}\rangle \cdots |\tilde{\lambda}_{r,j}\rangle |\phi_j\rangle. \tag{14}$$

where $|\lambda_{i,j} - \tilde{\lambda}_{i,j}| < \varepsilon/2\sqrt{r}$ for all $i = 1, \ldots, r$. If we group the content of the first r registers as a vector $\tilde{\boldsymbol{\lambda}}_j$ then the state (14) can be written as

$$\sum_{j=1}^{N} \langle \phi_j | \psi_0 \rangle |\tilde{\boldsymbol{\lambda}}_j\rangle |\phi_j\rangle \tag{15}$$

7

2. Measure the last register in the vertex basis.

3. Return the measured vertex.

Algorithm 1 proposes to use the (binary representations) of the phase vectors $\tilde{\boldsymbol{\lambda}}_j$ as the tags t_j in (13). The correctness of the algorithm follows from the next lemma.

Lemma 3.3. *The vectors $\tilde{\boldsymbol{\lambda}}_j$ uniquely determine the eigenspaces of A. More precisely, $\tilde{\boldsymbol{\lambda}}_j = \tilde{\boldsymbol{\lambda}}_k$ if and only if j and k both belong to I_ℓ for some ℓ.*

Proof. Let $\boldsymbol{\lambda}_j = (\lambda_{1,j}, \lambda_{2,j}, \ldots, \lambda_{r,j})$ be the vectors of the exact eigenvalues of $A_1, \ldots, A_r$ corresponding to the eigenstate $|\phi_j\rangle$. Then for all $j = 1, \ldots, N$ we have

$$\|\boldsymbol{\lambda}_j - \tilde{\boldsymbol{\lambda}}_j\|_2^2 = \sum_{i=1}^{r} |\lambda_{i,j} - \tilde{\lambda}_{i,j}|^2 < \frac{\varepsilon^2}{4} \tag{16}$$

where the last inequality follows from the bound $|\lambda_{i,j} - \tilde{\lambda}_{i,j}| < \varepsilon/2\sqrt{r}$. Now, suppose $\tilde{\boldsymbol{\lambda}}_j = \tilde{\boldsymbol{\lambda}}_k$ where $k \in I_h$ and $j \in I_\ell$ and $h \neq \ell$. Then

$$\begin{aligned}
\|\boldsymbol{\lambda}_j - \boldsymbol{\lambda}_k\|_2 &= \|(\boldsymbol{\lambda}_j - \tilde{\boldsymbol{\lambda}}_j) - (\boldsymbol{\lambda}_k - \tilde{\boldsymbol{\lambda}}_k)\|_2 \\
&\leq \|\boldsymbol{\lambda}_j - \tilde{\boldsymbol{\lambda}}_j\|_2 + \|\boldsymbol{\lambda}_k - \tilde{\boldsymbol{\lambda}}_k\|_2 \\
&< \varepsilon && \text{by (16)}
\end{aligned}$$

which contradicts the ε-separatedness of $\mathcal{A}$. $\qquad\square$

The following theorem records the main result of this section.

Theorem 3.4. *Let $\Gamma = (V, E)$ be a graph with N vertices and adjacency matrix A. Given an ε-projector $\mathcal{A} = \{A_j\}_{1 \leq j \leq r}$ for A such that the walks $e^{iA_j t}$ can be performed in $F(t)\,\mathrm{poly}(\log N)$ operations, there is a quantum algorithm that can sample from the limiting distribution of the walk $W(t) = e^{iAt}$ in $O(rF(2\sqrt{r}\varepsilon^{-1})\,\mathrm{poly}(\log N))$ operations.*

Proof. The time consuming part of the algorithm is the phase estimation for the operators $W_j = e^{iA_j}$ for $j = 1, \ldots, r$. Each of these phase estimations is done with accuracy $\varepsilon/2\sqrt{r}$, and therefore requires $O(F(2\sqrt{r}\varepsilon^{-1})\,\mathrm{poly}(\log N))$ operations [31, Chapter 7]. Since there are r phase estimations, the claimed running time follows. $\qquad\square$

From Theorem 3.4 we see that the complexity of Algorithm 1 is mostly determined by the "quality" of the given ε-projector $\mathcal{A}$. If the number r of the operators in $\mathcal{A}$ and the separation parameter ε are $\mathrm{poly}(\log N)$ and $1/\mathrm{poly}(\log N)$, respectively, then the algorithm is efficient, i.e., runs in $\mathrm{poly}(\log N)$ operations. Otherwise, there is a natural trade-off between the sizes of the two parameters. Also note that, by definition, we always have $F(t) \in O(t)$ for any ε-projector, so the running time in Theorem 3.4 is always upper bounded by $O(r^{3/2}\varepsilon^{-1}\,\mathrm{poly}(\log N))$.

An immediate special case of Theorem 3.4 is when the ε-projector is a singleton set $\mathcal{A} = \{B\}$, that is, when $r = 1$. If we only have black-box access to B we can only perform the walk e^{iBt} with a running time that scales linearly in t. In this case, the running time of Algorithm 1 scales linearly in ε^{-1}. On the other hand, if we can perform the walk e^{iBt} with a running time that scales polynomially in $\log t$, then the running time of Algorithm 1 scales polynomially in $\log(\varepsilon^{-1})$. A lower bound for the running time of Algorithm 1 can be obtained using the fact that for any such ε-projector $\mathcal{A}$ we must have $\varepsilon \leq \Delta$, where Δ, defined in (8), is the minimum spacing between the distinct eigenvalues of A. Let us record these observations for the sake of referencing.

Corollary 3.5. *Let $\Gamma = (V, E)$ be a graph with N vertices and adjacency matrix A. Given an ε-projector $\mathcal{A} = \{B\}$ for A, we have the following:*

(a) *If the unitary e^{iB} can be applied in $\mathrm{poly}(\log N)$ operations, then there is a quantum algorithm that can sample from the limiting distribution of the walk $W(t) = e^{iAt}$ in $O(\varepsilon^{-1} \mathrm{poly}(\log N))$ operations.*

(b) *If the unitary e^{iBt} can be applied in $O((\log t) \mathrm{poly}(\log N))$ operations, then there is a quantum algorithm that can sample from the limiting distribution of the walk $W(t) = e^{iAt}$ in $O(\log(\varepsilon^{-1}) \mathrm{poly}(\log N))$ operations.*

Moreover, if Δ is known or can be approximated efficiently, then the above running times can be improved to $O(\Delta^{-1} \mathrm{poly}(\log N))$ in case (a) and $O(\log(\Delta^{-1}) \mathrm{poly}(\log N))$ in case (b).

In the black-box setting, we are usually given access to the adjacency matrix A of Γ such that we can apply the unitary e^{iA} in time $\mathrm{poly}(\log N)$. In this setting, we can just take the ε-projector $\mathcal{A} = \{A\}$ for a small enough ε. This makes the complexity of sampling from the limiting distribution of the walk $W(t)$ fundamentally dependent on Δ. If we use the naive way of running $W(t)$ for a large t and measuring the resulting state, the bound (10) suggests that we should take T proportional to $(2 \ln M + 2)/\delta\Delta$ to be within distance δ of the limiting distribution. In comparison, Corollary 3.5 says we only need to run $W(t)$ for $t \approx 1/\Delta$ (and perform some other negligible operations) to sample exactly from the limiting distribution.

In the non-black-box setting, we have the opportunity to exploit some extra information on A to find ε-projectors that allow us to efficiently sample from the limiting distribution of $W(t)$. In the following sections, we give examples of graphs for which we can find such ε-projectors.

4 Quasi-Abelian Graphs

As a first application of our sampling algorithm we consider a class of graphs, called quasi-abelian graphs, in this section. Quasi-abelian graphs, as we will see, are a potential source of concrete examples for which Algorithm 1 runs in polynomial time. In the following, we first review some general properties of quasi-abelian graphs, and then look more closely at a specific example called Winnie Li graphs.

Let G be a finite group of size N, and let the subset $S \subseteq G$ be such that $1 \notin S$. The Cayley digraph $\Gamma = \Gamma(G, S)$ of the pair (G, S) is a directed graph in which the vertex set is the set of elements in G and the edge set is $\{(a, as) : a \in G, s \in S\}$. If S is symmetric, i.e., $s \in S$ if and only if $s^{-1} \in S$, then Γ is an undirected graph called the Cayley graph. In this paper, we always assume that S generates the entire group G, which means that Γ is connected. Denote by $f_S : G \to \{0, 1\}$ the characteristic function of S defined by $f_S(a) = 1$ if $a \in S$ and $f_S(a) = 0$ if $a \notin S$. We will also denote by $A(\Gamma)$ the adjacency matrix of Γ.

Definition 4.1. A quasi-abelian graph is a Cayley graph $\Gamma(G, S)$ in which S is the union of some conjugacy classes of G.

A class function is a function $f : G \to \mathbb{C}$ that is constant on conjugacy classes of G. It is not hard to show that the Fourier transform of a class function f is a diagonal matrix, e.g., see [17, Chapter 2]. From the definition of quasi-abelian graphs we see that f_S is always a class function. For any irreducible representation $\varrho \in \widehat{G}$ we obtain

$$\widehat{f}_S(\varrho) = \frac{1}{\sqrt{d_\varrho |G|}} \sum_{s \in S} \chi_\varrho(s) \mathbb{1}_{d_\varrho} \tag{17}$$

9

where d_ϱ is the dimension of ϱ and χ_ϱ is the character of ϱ. The following proposition is partially proved in [17] and [42] with different notations. Here, we give a short proof consistent with our notations.

Theorem 4.2. *Let* $\Gamma(G, S)$ *be a quasi-abelian graph on a finite group* G. *Denote by* F_G *the Fourier transform over* G. *Then*

(a) *The adjacency matrix* $A(\Gamma)$ *is diagonalized by* F_G,

(b) *The eigenvectors of* $A(\Gamma)$ *are given by* $F_G^*|\varrho, j, k\rangle$ *for* $\varrho \in \widehat{G}$ *and* $1 \leq j, k \leq d_\varrho$,

(c) *The eigenvalue corresponding to the eigenvector* $F_G^*|\varrho, j, k\rangle$ *is given by*

$$\lambda_\varrho = \frac{1}{d_\varrho} \sum_{s \in S} \chi_\varrho(s).$$

From part (c) of the theorem we see that the eigenvalues of $A(\Gamma)$ are only determined by the irreducible representations of G and the set S. Each irreducible representation ϱ corresponds to d_ϱ^2 eigenvectors. This means an eigenvalue λ_ϱ has multiplicity at least d_ϱ^2. If G is abelian, we always have $d_\varrho = 1$, but that does not mean the eigenvalues λ_ϱ are distinct for different $\varrho \in \widehat{G}$.

Proof of Theorem 4.2. Let ρ_{reg} be the regular representation of G. Then we can write

$$A(\Gamma) = \sum_{s \in S} \rho_{\text{reg}}(s) = \sum_{a \in G} f_S(a) \rho_{\text{reg}}(a).$$

The Fourier transform F_G decomposes ρ_{reg} as

$$F_G \rho_{\text{reg}}(a) F_G^* = \bigoplus_{\varrho \in \widehat{G}} (\varrho(a) \otimes \mathbb{1}_{d_\varrho}).$$

It follows from this decomposition that

$$F_G A(\Gamma) F_G^* = \sum_{a \in G} \bigoplus_{\varrho \in \widehat{G}} (f_S(a) \varrho(a) \otimes \mathbb{1}_{d_\varrho})$$

$$= \bigoplus_{\varrho \in \widehat{G}} \sqrt{\frac{|G|}{d_\varrho}} (\widehat{f}_S(\varrho) \otimes \mathbb{1}_{d_\varrho})$$

$$= \bigoplus_{\varrho \in \widehat{G}} (\lambda_\varrho \mathbb{1}_{d_\varrho} \otimes \mathbb{1}_{d_\varrho}),$$

where the last equality follows from the fact that f_S is a class function. This proves (a). Parts (b) and (c) follow from the definition of the quantum Fourier transform (12) and the identity (17) for the characteristic function f_S. □

According to Theorem 4.2, for a quasi-abelian graph $\Gamma(G, S)$ we can write the adjacency matrix $A = A(\Gamma)$ as

$$A = \sum_{\varrho \in \widehat{G}} \sum_{1 \leq j, k \leq d_\varrho} \lambda_\varrho F_G^* |\varrho, j, k\rangle \langle \varrho, j, k| F_G. \tag{18}$$

Therefore,

$$e^{iAt} = \sum_{\varrho \in \widehat{G}} \sum_{1 \leq j, k \leq d_\varrho} e^{i\lambda_\varrho t} F_G^* |\varrho, j, k\rangle \langle \varrho, j, k| F_G. \tag{19}$$

The expansion (19) suggest that we could perform the walk e^{iAt} using the following three steps:

1. Apply the quantum Fourier transform F_G,

2. Apply the phase operator $U_\varrho : |\varrho, j, k\rangle \mapsto e^{i\lambda_\varrho t}|\varrho, j, k\rangle$,

3. Apply the inverse quantum Fourier transform F_G^*.

Of course, this is only efficient if both the Fourier transform F_G and the phase operator U_ϱ can be applied efficiently. Suppose we can apply F_G efficiently for a given group G. Then a general strategy for constructing ε-projectors for A is as follows. For any function $g : \widehat{G} \to \mathbb{C}$ the operator

$$A_g = \sum_{\varrho \in \widehat{G}} \sum_{1 \le j,k \le d_\varrho} g(\varrho) F_G^* |\varrho, j, k\rangle\langle\varrho, j, k| F_G$$

commutes with A. Suppose that g satisfies the following condition:

$$g(\varrho_1) = g(\varrho_2) \text{ if and only if } \lambda_{\varrho_1} = \lambda_{\varrho_2} \text{ for all } \varrho_1, \varrho_2 \in \widehat{G} \tag{20}$$

Then A_g has the same eigenspaces as A. If we can efficiently approximate g with exponential accuracy, then $\mathcal{A} = \{A_g\}$ is an ε-projector that satisfies the conditions of Corollary 3.5 (b). If we can only approximate g with polynomial accuracy, then we might need many more of these functions g that satisfy (20). Ideally, we need to find $g_1, g_2, \ldots, g_r : \widehat{G} \to \mathbb{C}$, with $r = \text{poly}(\log N)$, such that $\mathcal{A} = \{A_{g_j}\}_{1 \le j \le r}$ is an ε-projector for A for some $\varepsilon = 1/\text{poly}(\log N)$. In this case, $\mathcal{A}$ satisfies the conditions of Theorem 3.4, and we can efficiently sample from the limiting distribution of the walk $W(t) = e^{iAt}$ using $\mathcal{A}$.

4.1 Winnie Li graphs

Let $\mathbb{F}_p$ be a finite field of characteristic $p \ge 3$. For an extension $F/\mathbb{F}_p$ of degree n, the norm map $N_{F/\mathbb{F}_p} : F \to \mathbb{F}_p$ is defined by $N_{F/\mathbb{F}_p}(a) = a^{(p^n-1)/(p-1)}$, which is a homomorphism between the multiplicative groups $F^\times$ and $\mathbb{F}_p^\times$. Let $S = \ker(N_{F/\mathbb{F}_p})$, i.e., the set of elements of F of norm 1. The Winnie Li graph of F over $\mathbb{F}_p$ is the Cayley digraph $\Gamma(F, S)$ where the vertex set is F and the edge set is $\{(a, a+s) : a \in F, s \in S\}$.

Before we get into the specifics of the structures of these graphs, recall the quantum Fourier transform over F. The set of additive characters of F is given by

$$\chi_{F,a}(x) = \omega_p^{\text{Tr}_{F/\mathbb{F}_p}(ax)}, \quad a \in F,$$

where $\text{Tr}_{F/\mathbb{F}_p}(x) = x + x^p + \cdots + x^{p^{n-1}}$ is the trace map from F to $\mathbb{F}_p$. The quantum Fourier transform of the basis element $|a\rangle$, where $a \in F$, is given by

$$|\hat{a}\rangle = \frac{1}{\sqrt{|F|}} \sum_{x \in F} \chi_{F,a}(x)|x\rangle$$

For even n, the graph Γ is undirected, since $N_{F/\mathbb{F}_p}(-a) = N_{F/\mathbb{F}_p}(a)$. Let us look at the simple case where $n = 2$. The extension $F/\mathbb{F}_p$ is a quadratic extension, and if we assume that -1 is a quadratic nonresidue in $\mathbb{F}_p$, we can take $F = \mathbb{F}_p(i)$ where $i^2 = -1$. The elements of $\mathbb{F}_p(i)$ can be written in the form $x + iy$ for $x, y \in \mathbb{F}_p$, so the norm map takes the simple form

$$N_{\mathbb{F}_q(i)/\mathbb{F}_p}(x + iy) = x^2 + y^2.$$

Therefore, the set S of elements of norm 1 is the set of $\mathbb{F}_p$-points of the circle $x^2 + y^2 = 1$. The Winnie Li graph $\Gamma(\mathbb{F}_p(i), S)$ is a $(p+1)$-regular graph of p^2 vertices. It follows from (18) that the adjacency matrix of Γ can be written as

$$A(\Gamma) = \sum_{a \in F} \lambda_a |\hat{a}\rangle\langle\hat{a}|,$$

For an element $a = u + iv \in \mathbb{F}_p(i)$, Theorem 4.2 (c) gives

$$\lambda_a = \sum_{\mathrm{N}_{F/\mathbb{F}_p}(b)=1} \chi_{F,a}(b) = \sum_{x^2+y^2=1} \omega_p^{2(ux-vy)} = \begin{cases} q+1 & \text{if } a = 0, \\ -K(1, u^2 + v^2) & \text{if } a \neq 0. \end{cases} \tag{21}$$

Here, $K(a, b)$ is the Kloosterman sum with parameters a, b, which we will briefly talk about next.

Kloosterman sums. For $a, b \in \mathbb{F}_p$, the exponential sum

$$K(a, b) = \sum_{x \in \mathbb{F}_p^\times} \omega_p^{ax+bx^{-1}} \tag{22}$$

is called the Kloosterman sum with respect to a, b. The last equality in (21) follows from the definition (22). Since $\overline{K(a,b)} = K(a,b)$, these sum are real numbers. A well known result of Weil [50] gives the bound $|K(a, b)| \leq 2\sqrt{p}$. When p is large, estimating $K(a, b)$ is an open problem; there are no known classical or quantum algorithms that can efficiently estimate $K(a, b)$. For a multiplicative character χ of $\mathbb{F}_p^\times$, the χ-twisted Kloosterman sum is defined by

$$K(\chi, a, b) = \sum_{x \in \mathbb{F}_p^\times} \chi(x)\omega_p^{ax+bx^{-1}}. \tag{23}$$

Interestingly, when χ is the quadratic character of $\mathbb{F}_p^\times$, the sum (23) has a closed form, and is easy to estimate [43, 7].

Euclidean graphs. The construction of the Winnie Li graph $\Gamma(\mathbb{F}_p(i), S)$ can be directly generalized to obtain the so called Euclidean graphs [39]. Let $n > 0$ be an integer and $b \in \mathbb{F}_p$. Define the quadratic form $Q(x) = x_1^2 + x_2^2 + \cdots + x_n^2$ over $\mathbb{F}_p$. An Euclidean graph for n and b is the Cayley graph $\Gamma(\mathbb{F}_p^n, S_b)$ where S_b is the set of solutions of $Q(x) = b$ in $\mathbb{F}_p^n$. For simplicity, assume $b = 1$. The adjacency matrix of Γ is

$$A(\Gamma) = \sum_{a \in \mathbb{F}_p^n} \lambda_a |\hat{a}\rangle\langle\hat{a}|$$

where

$$|\hat{a}\rangle = \frac{1}{\sqrt{p^n}} \sum_{x \in \mathbb{F}_p^n} \omega_p^{\langle a,x\rangle} |x\rangle$$

is the quantum Fourier transform of $|a\rangle$, and $\langle a, x\rangle = a_1 x_1 + \cdots + a_n x_n$ for $a = (a_1, \ldots, a_n)$ and $x = (x_1, \ldots, x_n)$. The eigenvalues λ_a are [39]

$$\lambda_a = \sum_{x \in S_1} \omega_p^{\langle a,x\rangle} = \begin{cases} |S_1| & \text{if } a = 0, \\ \frac{G_1^n}{p} K\left(\chi^n, 1, \frac{Q(a)}{4}\right) & \text{if } a \neq 0, \end{cases} \tag{24}$$

12

where χ is the quadratic character of $\mathbb{F}_p^\times$, and where

$$G_1 = \begin{cases} \sqrt{p} & \text{if } p \equiv 1 \bmod 4, \\ i\sqrt{p} & \text{if } p \equiv 3 \bmod 4. \end{cases}$$

When n is odd, it is easy to approximate the eigenvalues λ_a with exponential accuracy, since it is easy to approximate (23) when χ is the quadratic character. So, it is easy to perform the walk

$$e^{iA(\Gamma)t} = \sum_{a \in \mathbb{F}_p^n} e^{i\lambda_a t} |\hat{a}\rangle\langle\hat{a}|$$

for exponentially large t. Therefore, we can easily sample from the limiting distribution of $W(t)$ by running the walk for random values of $t \in [0, T]$ for a large T. However, when n is even, we do not know how to perform $W(t)$ for large t. In fact, there is no known way to even perform $W(1) = e^{iA(\Gamma)}$ efficiently. We now use the general strategy introduced at the beginning of Section 4 to efficiently sample from the limiting distribution of $W(t)$ when n is even. Note that since $G = \mathbb{F}_p^n$ is an abelian group, we have $\widehat{G} \cong G$. Therefore, we need to find functions $g : G \to \mathbb{C}$ that satisfy the condition in (20).

Proposition 4.3. *Define the function $g : G \to \mathbb{C}$ as $g(x) = (Q(x) \bmod p)/p$. Then the operator*

$$A_g = \sum_{a \in \mathbb{F}_p^n} g(a)|\hat{a}\rangle\langle\hat{a}|$$

is an ε-projector for $A(\Gamma)$ for $\varepsilon = 1/p$.

Proof. We first need to show that A_g has the same eigenspaces as $A(\Gamma)$. It is known that the Kloosterman sums $K(1, b)$, $b \in \mathbb{F}_p$, are distinct [23]. Therefore, according to (24), two eigenvalues λ_a and λ_b of $A(\Gamma)$ are the same if and only if $Q(a) = Q(b)$. But we also trivially have $g(a) = g(b)$ if and only if $Q(a) = Q(b)$. Note that the eigenspace corresponding to λ_a is the set of all vectors $|\hat{x}\rangle$ such that $Q(x) = Q(a)$.

Now, for $\lambda_a \neq \lambda_b$ we have $Q(a) \neq Q(b)$, hence

$$|g(a) - g(b)| = \frac{1}{p}\big|(Q(a) - Q(b)) \bmod p\big| \geq \frac{1}{p},$$

and since we can efficiently compute g with exponential accuracy, it follows that A_g is a $\frac{1}{p}$-projector for $A(\Gamma)$. $\square$

With g as in Proposition 4.3, we can perform the walk

$$e^{iA_g t} = \sum_{a \in \mathbb{F}_p^n} e^{ig(a)t}|\hat{a}\rangle\langle\hat{a}|$$

in $O((\log t)\,\mathrm{poly}(\log N))$ operations. Corollary 3.5 (b) and Proposition 4.3 now give

Theorem 4.4. *Let $\Gamma(\mathbb{F}_p^n, S)$ be an Euclidean graph with adjacency matrix $A(\Gamma)$. For a given initial state $|\psi_0\rangle$ we have*

(a) *The limiting distribution of the walk $W(t) = e^{iA(\Gamma)t}$ on Γ is given by*

$$P_\infty(v|\psi_0) = \sum_{j=0}^{p-1} \left| \sum_{k \in S_j} \langle v|\hat{k}\rangle\langle\hat{k}|\psi_0\rangle \right|^2.$$

(b) *There is a quantum algorithm that can sample from $P_\infty(\cdot|\psi_0)$ in $\mathrm{poly}(n \log p)$ operations.*

13

5 Isogeny Graphs

As another application of Algorithm 1 we consider a class of graphs, called isogeny graphs, in this section. Isogeny graphs have attracted much attention in last two decade mainly because of their applications in cryptography [36, 37]. We will see in the following how the beautiful theory of Hecke algebras give us natural ε-projectors that make it possible to efficiently sample from the limiting distributions of quantum walks on these graphs.

The ε-projectors considered in this section, were first implicitly used by Kane, Sharif and Silverberg [29]. They considered the same set of operators we use here but in the context of quaternion algebras. In here, we consider the space of supersingular elliptic curves over the finite fields $\mathbb{F}_{p^2}$, whereas in [29] they considered the space of ideals in an ideal class of the quaternion algebra $B_{p,\infty}$. These two spaces are mathematically essentially the same, in a precise sense, but they are vastly different from a cryptographic perspective. In particular, the isogeny problem is easy in the latter space, but believed to be hard in the former [32]. Working with the ε-projectors in the space of elliptic curves yields a potential solution to the problem of generating an honest curve, as we will explain in Section 5.3.

Let $\mathbb{F}_q = \mathbb{F}_{p^2}$ be a finite field of characteristic $p \geq 5$. An elliptic curve E over $\mathbb{F}_q$ is a projective smooth curve of genus one. For the definition of these terms and an extensive introduction to elliptic curve see [28]. The affine version of E is usually written as the cubic $y^2 = x^3 + ax + b$, $a, b \in \mathbb{F}_q$, known as the Weierstrass equation of E. The set of points on E in any extension of $\mathbb{F}_p$ form an abelian group. An elliptic curve E over $\mathbb{F}_q$ is called supersingular if it has no nontrivial points of order p. It can be shown that any supersingular elliptic curves over the algebraic closure $\overline{\mathbb{F}}_p$ can be defined over $\mathbb{F}_q$, or, more precisely, is isomorphic to a curve over $\mathbb{F}_q$. Therefore, we always assume that any supersingular elliptic curve E has its coefficients a, b in $\mathbb{F}_q$.

An isogeny $\phi : E_1 \to E_2$ between two elliptic curves E_1 and E_2 is a rational function that is also a homomorphism of groups of points on E_1 and E_2. An isogeny ϕ induces an embedding of function fields $\phi^* : K(E_2) \to K(E_1)$ defined by $\phi^*(f) = f \circ \phi$. The degree of ϕ is the degree of the extension $K(E_1)/\phi^* K(E_2)$. An isogeny of degree ℓ is called an ℓ-isogeny. For any isogeny ϕ there is a unique isogeny $\widehat{\phi} : E_2 \to E_1$ called the dual of ϕ. Let ℓ be a prime different from the characteristic p. Define a graph G_ℓ with vertices the set of all $\overline{\mathbb{F}}_p$-isomorphism classes of supersingular elliptic curves, and edges the set of ℓ-isogenies between the curves. G_ℓ is called the supersingular ℓ-isogeny graph. Since the dual of an ℓ-isogeny is again an ℓ-isogeny but in the opposite direction, we usually consider G_ℓ as an undirected graph. For simplicity, assume that $p \equiv 1 \mod 12$. Then G_ℓ is an $(\ell + 1)$-regular graph with $N = \lfloor p/12 \rfloor$ vertices and no self loops. The adjacency matrix of G_ℓ is a symmetric matrix denoted by T_ℓ and is called the Hecke operator.

5.1 Simulating the Hecke operators

Let S be the set of vertices of G_ℓ, i.e., the set of isomorphism classes of supersingular elliptic curves in characteristic p. The Hecke operator T_ℓ acts on the formal abelian group

$$M = \bigoplus_{E \in S} \mathbb{Z}E$$

by sending each curve E to a sum of its neighbours in G_ℓ. In the quantum setting, we consider the action of T_ℓ on the complex Euclidean space $\mathcal{X} = M \otimes_{\mathbb{Z}} \mathbb{C}$ with the basis $\{|E\rangle\}_{E \in S}$. The operators T_ℓ, for different values of ℓ, are closely related to the Hecke operators acting on the space of modular forms [19, 33, 5], so the terminology we use here is mostly adopted from the theory of

modular forms. For example, the trivial eigenvector

$$|\mathcal{E}\rangle = \frac{1}{\sqrt{N}} \sum_{E \in S} |E\rangle$$

of T_ℓ is called the Eisenstein eigenform and corresponds to the eigenvalue $\lambda_{\mathcal{E}} = \ell+1$. Deligne's proof of the Riemann hypothesis for function fields [16, 30] implies that the nontrivial eigenvalues of T_ℓ are contained in the interval $[-2\sqrt{\ell}, 2\sqrt{\ell}]$. We make the heuristic assumption that the eigenvalues of T_ℓ are distinct. This assumption is in fact a consequence of (the well known) Maeda's conjecture [26] which states that the characteristic polynomial of T_ℓ is irreducible over $\mathbb{Q}$. We refer the reader to [38, 46, 25], and the references therein, for results on the computational verification of Maeda's conjecture.

An ℓ-isogeny can be computed in $O(\ell)$ operations over $\mathbb{F}_q$ using the Vélu formulas [47]. When ℓ is small, i.e., $\ell = \mathrm{poly}(\log N)$, it is easy to compute the list of all the neighbours of a given curve E in G_ℓ. In particular, we can efficiently implement the isometry

$$T : |E\rangle \rightarrow \frac{1}{\sqrt{\ell+1}} \sum_{j=0}^{\ell} |E\rangle |E_j\rangle$$

from $\mathcal{X}$ to $\mathcal{X} \otimes \mathcal{X}$, where $E_1, \ldots, E_\ell$ are the $\ell+1$ neighbours of E. Therefore, we can use the existing Hamiltonian simulation techniques [11, 49] to efficiently approximate the unitary e^{iT_ℓ}.

Proposition 5.1. *For any prime $\ell = \mathrm{poly}(\log N)$, the walk $W(t) = e^{iT_\ell t}$ on the ℓ-isogeny graph G_ℓ can be performed in $O(t \, \mathrm{poly}(\log N))$ operations.*

5.2 Distribution of eigenvalues

For different primes ℓ we have different isogeny graphs G_ℓ over $\mathbb{F}_q$; the set of vertices is always the same but the set of edges change with ℓ. Therefore, we get different Hecke operators T_ℓ acting on the same space $\mathcal{X}$. More generally, Hecke operators can be defined for any integer $n > 0$. The operator T_n represent the adjacency matrix of the n-isogeny graph G_n, although for non-prime n one needs to be more careful about some definitions. The algebra $\mathbb{T}_\mathbb{Z} = \mathbb{Z}[\{T_n\}_{n \in \mathbb{Z}}]$, generated by all Hecke operators acting on $\mathcal{X}$, is called the Hecke algebra. Let $\mathbb{T} = \mathbb{T}_\mathbb{Z} \otimes_\mathbb{Z} \mathbb{C}$ be the Hecke algebra over $\mathbb{C}$. It can be shown that $\mathbb{T}$ is a commutative ring [48, Chapter 41]. In particular, for any m, n, the operators T_n and T_m commute. This means all Hecke operators are simultaneously diagonalizable.

It was proved by Serre [44] that for large p, the eigenvalues of the normalized Hecke operator $T_\ell/\sqrt{\ell}$ are equidistributed in $[-2, 2]$ with respect to the measure

$$\mu_\ell = \frac{\ell+1}{\pi} \frac{(1 - x^2/4)^{1/2} dx}{(\ell^{1/2} + \ell^{-1/2})^2 - x^2}.$$

Let $|\phi_1\rangle, |\phi_2\rangle, \ldots, |\phi_N\rangle \in \mathcal{X}$ be a simultaneous basis for all Hecke operators, and let $\ell_1, \ell_2, \ldots, \ell_r$ be a set of distinct primes each bounded by $\mathrm{poly}(\log N)$. For all $1 \leq k \leq r$ and $1 \leq j \leq N$ we have

$$\frac{1}{\sqrt{\ell_k}} T_{\ell_k} |\phi_j\rangle = \lambda_{j,k} |\phi_j\rangle$$

for some $\lambda_{j,k} \in [-2, 2]$. Define $\boldsymbol{\lambda}_j = (\lambda_{j,1}, \lambda_{j,2}, \ldots, \lambda_{j,r})$. It was also proved in [44] that for large p, the vectors $\boldsymbol{\lambda}_j$ are equidistributed in $[-2, 2]^r$ with respect to the product measure $\boldsymbol{\mu} = \prod_{k=1}^{r} \mu_{\ell_k}$.

This means that asymptotically we can treat the $\boldsymbol{\lambda}_j$ as independent samples from the distribution given by the measure $\boldsymbol{\mu}$.

For large N, μ_ℓ approaches the measure $\mu = \frac{1}{2\pi}\sqrt{4-x^2}dx$, the semicircle distribution on $[-2,2]$. So, when the characteristic p is large enough, it is natural to assume that the vectors $\boldsymbol{\lambda}_j$ are independent random samples from μ^r. We now prove that for $r \in O(\log N)$ and $\varepsilon = 1/\sqrt{\log N}$, the set of Hecke operators $\mathcal{T} = \{T_{\ell_k}/\sqrt{\ell_k}\}_{1\leq k\leq r}$ is an ε-projector for any Hecke operator T_ℓ with ℓ a prime number. The following proves the ε-separatedness of $\mathcal{T}$.

Lemma 5.2. *Let* $r \geq 32\log N$, *let* $\ell_1, \ell_2, \ldots, \ell_r$ *be a set of distinct primes each bounded by* poly$(\log N)$, *and let* $\varepsilon = 1/\sqrt{\log N}$. *Then the set of operators* $\mathcal{T} = \{T_{\ell_k}/\sqrt{\ell_k}\}_{1\leq k\leq r}$ *is ε-separated with overwhelming probability.*

Proof. Let $Z = (X - Y)^2$ where X and Y are distributed according to μ. We have $\mathrm{Var}(X) = \mathrm{Var}(Y) = 1$ and $\mathbb{E}[X] = \mathbb{E}[Y] = 0$. So

$$\mathbb{E}[Z] = \mathrm{Var}(X - Y) + \mathbb{E}[X - Y]^2 = 2.$$

Now define $Z_k = (\lambda_{j,k} - \lambda_{l,k})^2$ for $k = 1, \ldots, r$, and $W = Z_1 + Z_2 + \cdots + Z_r$. Then $\mathbb{E}[W] = 2r$. We have

$$
\begin{aligned}
\Pr[\|\boldsymbol{\lambda}_j - \boldsymbol{\lambda}_l\|_2 \leq \varepsilon] &= \Pr[W \leq \varepsilon^2] \\
&= \Pr[W - 2r \leq \varepsilon^2 - 2r] \\
&\leq \exp\left(\frac{-2(2r - \varepsilon^2)^2}{256r}\right) \qquad (25)\\
&\leq \frac{1}{N^{1.4}}
\end{aligned}
$$

where the inequality (25) is the Hoeffding's inequality [27]. $\qquad\square$

Lemma 5.2 and Theorem 3.4 now give

Theorem 5.3. *Let* $\ell \neq p$ *be prime, and let* G_ℓ *be the ℓ-isogeny graph with vertices the set of supersingular elliptic curves over* $\mathbb{F}_{p^2}$. *For a given initial state* $|\psi_0\rangle$ *we have*

(a) The limiting distribution of the walk $W(t) = e^{iT_\ell t}$ *on* G_ℓ *is given by*

$$P_\infty(E|\psi_0) = \sum_{j=1}^{N}|\langle E|\phi_j\rangle\langle\phi_j|\psi_0\rangle|^2 \qquad (26)$$

(b) There is a quantum algorithm that can sample from $P_\infty(\cdot|\psi_0)$ *in* poly$(\log p)$ *operations.*

5.3 Honest hard curves

A hard problem in isogeny based cryptography is to compute the endomorphism ring $\mathrm{End}(E)$ of a given supersingular elliptic E. The majority of other computational assumptions can be reduced to the endomorphism ring problem [51]. Informally, a *hard curve* is a random curve on G_ℓ with an unknown endomorphism ring. To classically generate a random curve on G_ℓ, one can do the following:

1. Start with a known curve E_0 on G_ℓ,
2. Take a random walk of length at least $2\log p$, and

16

3. Return the final curve

For small ℓ, taking classical random walks on G_ℓ can be done very efficiently [18]. Also, it is well known that G_ℓ has a small diameter, and taking a walk of length $\approx 2 \log p$ is enough to get close to the uniform distribution on G_ℓ.

At the first glance, it seems that the above classical random walk on G_ℓ can be used to efficiently generate a hard curve. However, there is an issue with this approach that seems to be unavoidable: the random walk explicitly generates a path on G_ℓ. This path can be used to compute the endomorphism ring of the returned curve. More precisely, if the endomorphism ring of the initial curve E_0 is known and we are given a path $\phi : E_0 \to E$ (which is an isogeny) then we can use ϕ to compute $\mathrm{End}(E)$. Any such path ϕ is called a backdoor for E. A hard curve without a backdoor is called an honest hard curve.

There is no known efficient classical solution for generating an honest hard curve. A potential solution using quantum walks was first discussed in [6]. The idea is to sample from the limiting distribution of the walk $W(t) = e^{iT_\ell t}$ on G_ℓ. According to Theorem 5.3, this can be done efficiently, and in contrast to the classical walk, the quantum walk does not generate any path on G_ℓ. Two important questions about the quantum walk solution in [6] remained unresolved. We briefly address those questions here.

The first question is weather the distribution (26) is close to uniform. Intuitively, there is no reason to believe that the eigenvectors $|\phi_j\rangle$ are localized[1]. In particular, the action of the Hecke operator T_ℓ on $\mathcal{X}$ implies that each entry of $|\phi_j\rangle$ is an average of a set of $\ell+1$ other entries. Since $|\phi_j\rangle$ is an eigenvector for all Hecke operators, these averages involve all the other entries for large enough ℓ. Therefore, one would expect that the entries of $|\phi_j\rangle$ are not too small or too large, and that the distribution $P_\infty(\cdot|\psi_0)$ is not concentrated on a small subset of vertices. In fact, it is not hard to show, using general techniques in the theory of Markov chains, that $P_\infty(E_1|E_2) \geq N^{-2}$ for every two curves E_1, E_2 [41]. Notwithstanding, a rigorous proof that $P_\infty(\cdot|\psi_0)$ is close to the uniform distribution does not seem straightforward. Maybe there are ways to analyze the amplitudes of the vectors $|\phi_j\rangle$ through their close connection to complex modular forms, but we are not aware of any work regarding this in the literature. Instead, this question can be approached algorithmically using the double-loop technique of [41]. The double-loop algorithm for our case is:

1. Set E to a known curve E_0 on G_ℓ
2. Repeat for k times
   (a) Run Algorithm 1 with initial state $|E\rangle$ to obtain a curve E_1
   (b) Set $E \leftarrow E_1$
3. Return E

Assuming that the maximum column distance

$$\alpha = \max_{E_1, E_2 \in S} \|P_\infty(\cdot|E_1) - P_\infty(\cdot|E_2)\|_1$$

of P_∞ is bounded by a constant smaller than 1, we need to only set $k = \lceil \log_{1/\alpha} \delta^{-1} \rceil$ for the distribution of the final curve E to be within distance δ of the uniform distribution. Again, the assumption that α is always bounded by a constant $c < 1$ is not rigorously proved, but it is more plausible than the assumption that $P_\infty(\cdot|\phi_0)$ is close to uniform.

The second question is only specific to the sampling algorithm proposed in [6]. In that algorithm, after preparing the state (15), the first register is measured. The measurement outcome is a random vector $\tilde{\boldsymbol{\lambda}}_j$, and the post-measurement state is the corresponding eigenstate $|\phi_j\rangle$. The state $|\phi_j\rangle$ is

---

[1]A unit vector is localized if the mass of the vector is concentrated on a small subset of entries.

then measured in the vertex basis to obtain a curve E. The question is whether the vector $\tilde{\boldsymbol{\lambda}}_j$ reveals any information about the endomorphism ring $\text{End}(E)$ of the curve E. This situation is entirely avoided in Algorithm 1; we never measure the first register. The post-measurement state of Algorithm 1 is a superposition of all the vectors $\tilde{\boldsymbol{\lambda}}_j$, $j = 1, \ldots, N$, which does not seem to provide any useful information about $\text{End}(E)$.

Remark 1. In a cryptography scenario, if Alice presents Bob with an alleged hard curve E, there is no way for Bob to know whether E is an honest hard curve, or that Alice is in possession of a backdoor for E. In other words, Bob is unable to determine which algorithm Alice has used to generate E. This kind of trust issue is normally solved by a higher level cryptographic construction.

References

[1] William Adamczak, Kevin Andrew, Peter Hernberg, and Christino Tamon. A note on graphs resistant to quantum uniform mixing. *arXiv preprint quant-ph/0308073*, 2003.

[2] Dorit Aharonov, Andris Ambainis, Julia Kempe, and Umesh Vazirani. Quantum walks on graphs. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 50–59, 2001.

[3] Andris Ambainis, Andrew M Childs, Ben W Reichardt, Robert Špalek, and Shengyu Zhang. Any and-or formula of size n can be evaluated in time $n^{1/2+o(1)}$ on a quantum computer. *SIAM Journal on Computing*, 39(6):2513–2530, 2010.

[4] Simon Apers, Shantanav Chakraborty, Leonardo Novo, and Jérémie Roland. Quadratic speedup for spatial search by continuous-time quantum walk. *arXiv preprint arXiv:2112.12746*, 2021.

[5] Andriyan Bilyk, Javad Doliskani, and Zhiyong Gong. Cryptanalysis of three quantum money schemes. *arXiv preprint arXiv:2205.10488*, 2022.

[6] Jeremy Booher, Ross Bowden, Javad Doliskani, Tako Boris Fouotsa, Steven D Galbraith, Sabrina Kunzweiler, Simon-Philipp Merz, Christophe Petit, Benjamin Smith, Katherine E Stange, et al. Failing to hash into supersingular isogeny graphs. *arXiv preprint arXiv:2205.00135*, 2022.

[7] Leonard Carlitz. Weighted quadratic partitions over a finite field. *Canadian Journal of Mathematics*, 5:317–323, 1953.

[8] Shantanav Chakraborty, Kyle Luh, and Jérémie Roland. Analog quantum algorithms for the mixing of markov chains. *Physical Review A*, 102(2):022423, 2020.

[9] Shantanav Chakraborty, Kyle Luh, and Jérémie Roland. How fast do quantum walks mix? *Physical review letters*, 124(5):050501, 2020.

[10] Andrew M Childs. Universal computation by quantum walk. *Physical review letters*, 102(18):180501, 2009.

[11] Andrew M Childs. On the relationship between continuous-and discrete-time quantum walk. *Communications in Mathematical Physics*, 294(2):581–603, 2010.

[12] Andrew M Childs, Richard Cleve, Enrico Deotto, Edward Farhi, Sam Gutmann, and Daniel A Spielman. Exponential algorithmic speedup by a quantum walk. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 59–68, 2003.

[13] Andrew M Childs and Jeffrey Goldstone. Spatial search by quantum walk. *Physical Review A*, 70(2):022314, 2004.

[14] Andrew M Childs, Leonard J Schulman, and Umesh V Vazirani. Quantum algorithms for hidden nonlinear structures. In *48th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 395–404. IEEE, 2007.

[15] Charles W Curtis and Irving Reiner. *Representation theory of finite groups and associative algebras*, volume 356. American Mathematical Soc., 1966.

[16] Pierre Deligne. La conjecture de Weil. I. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 43(1):273–307, 1974.

[17] Persi Diaconis. Group representations in probability and statistics. *Lecture notes-monograph series*, 11:i–192, 1988.

[18] Javad Doliskani, Geovandro CCF Pereira, and Paulo SLM Barreto. Faster cryptographic hash function from supersingular isogeny graphs. *Cryptology ePrint Archive*, 2017.

[19] Martin Eichler. The basis problem for modular forms and the traces of the hecke operators. In *Modular functions of one variable I*, pages 75–152. Springer, 1973.

[20] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. A quantum algorithm for the hamiltonian nand tree. *arXiv preprint quant-ph/0702144*, 2007.

[21] Edward Farhi and Sam Gutmann. Quantum computation and decision trees. *Physical Review A*, 58(2):915, 1998.

[22] Leonid Fedichkin, Dmitry Solenov, and Christino Tamon. Mixing and decoherence in continuous-time quantum walks on cycles. *arXiv preprint quant-ph/0509163*, 2005.

[23] Benji Nichols Kaplan Fisher. *Distinctness of Kloosterman sums*. Princeton University, 1990.

[24] Heath Gerhardt and John Watrous. Continuous-time quantum walks on the symmetric group. In *Approximation, Randomization, and Combinatorial Optimization.. Algorithms and Techniques*, pages 290–301. Springer, 2003.

[25] Alexandru Ghitza and Angus McAndrew. Experimental evidence for maeda's conjecture on modular forms. *Tbilisi Mathematical Journal*, 5(2):55–69, 2012.

[26] Haruzo Hida and Yoshitaka Maeda. Non-abelian base change for totally real fields. *pacific journal of mathematics*, 181(3):189–217, 1997.

[27] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. In *The collected works of Wassily Hoeffding*, pages 409–426. Springer, 1994.

[28] D. Husemoller. *Elliptic Curves*. Graduate Texts in Mathematics. Springer New York, 2013.

[29] Daniel M. Kane, Shahed Sharif, and Alice Silverberg. Quantum money from quaternion algebras. arXiv:2109.12643, 2021.

[30] N Katz. An overview of deligne's proof of the riemann hypothesis for varieties over finite fields. *Mathematical developments arising from Hilbert problems (Proc. Sympos. Pure Math., XXVIII, Northern Illinois Univ., De Kalb, Ill., 1974)*, pages 275–305, 1976.

[31] Phillip Kaye, Raymond Laflamme, and Michele Mosca. *An introduction to quantum computing*. OUP Oxford, 2006.

[32] David Kohel, Kristin Lauter, Christophe Petit, and Jean-Pierre Tignol. On the quaternion-isogeny path problem. *LMS Journal of Computation and Mathematics*, 17(A):418–432, 2014.

[33] David R Kohel. Computing modular curves via quaternions, 1999.

[34] David A Levin and Yuval Peres. *Markov chains and mixing times*, volume 107. American Mathematical Soc., 2017.

[35] W-C Winnie Li and Keqin Feng. Character sums and abelian ramanujan graphs. *Journal of Number Theory*, 41(2):199–217, 1992.

[36] Luca De Feo. Mathematics of Isogeny Based Cryptography. arXiv:1711.04062, 2017.

[37] Luca De Feo. *Exploring Isogeny Graphs*. Habilitation, 2018.

[38] Kimball Martin. An on-average maeda-type conjecture in the level aspect. *Proceedings of the American Mathematical Society*, 149(4):1373–1386, 2021.

[39] Archie Medrano, Perla Myers, Harold M Stark, and Audrey Terras. Finite analogues of euclidean space. *Journal of Computational and Applied Mathematics*, 68(1-2):221–238, 1996.

[40] Cristopher Moore and Alexander Russell. Quantum walks on the hypercube. In *International Workshop on Randomization and Approximation Techniques in Computer Science*, pages 164–178. Springer, 2002.

[41] Peter C Richter. Almost uniform sampling via quantum walks. *New Journal of Physics*, 9(3):72, 2007.

[42] Dan Rockmore, Peter Kostelec, Wim Hordijk, and Peter F Stadler. Fast fourier transform for fitness landscapes. *Applied and Computational Harmonic Analysis*, 12(1):57–76, 2002.

[43] Hans Salié. Über die kloostermanschen summen $S(u, v; q)$. *Mathematische Zeitschrift*, 34(1):91–109, 1932.

[44] Jean-Pierre Serre. Répartition asymptotique des valeurs propres de l'opérateur de hecke tp. *Journal of the American Mathematical Society*, pages 75–102, 1997.

[45] Jean-Pierre Serre et al. *Linear representations of finite groups*, volume 42. Springer, 1977.

[46] Panagiotis Tsaknias. A possible generalization of maeda's conjecture. In *Computations with modular forms*, pages 317–329. Springer, 2014.

[47] Jacques Vélu. Isogénies entre courbes elliptiques. *CR Acad. Sci. Paris, Séries A*, 273:305–347, 1971.

[48] John Voight. *Quaternion algebras*. Springer Nature, 2021.

[49] John Watrous. Quantum simulations of classical random walks and undirected graph connectivity. *Journal of computer and system sciences*, 62(2):376–391, 2001.

[50] André Weil. On some exponential sums. *Proceedings of the National Academy of Sciences*, 34(5):204–207, 1948.

[51] Benjamin Wesolowski. Orientations and the supersingular endomorphism ring problem. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 345–371. Springer, 2022.

A   Bound on the mixing time

In this section, we bound the mixing time M_δ defined in (7). Our analysis closely follows that of [2]. We have

$$|P_T(v|\psi_0) - P_\infty(v|\psi_0)| = \left| \sum_{k,\ell=1:\lambda_k \neq \lambda_\ell}^{N} \langle v|\phi_k\rangle\langle\phi_k|\psi_0\rangle\langle\psi_0|\phi_\ell\rangle\langle\phi_\ell|v\rangle \frac{e^{i(\lambda_k - \lambda_\ell)T} - 1}{i(\lambda_k - \lambda_\ell)T} \right|$$

$$\leq \sum_{k,\ell=1:\lambda_k \neq \lambda_\ell}^{N} \frac{2}{|\lambda_k - \lambda_\ell|T} |\langle v|\phi_k\rangle\langle\phi_\ell|v\rangle| \cdot |\langle\phi_k|\psi_0\rangle\langle\psi_0|\phi_\ell\rangle|. \qquad (27)$$

Using the inequality $2|ab| \leq |a|^2 + |b|^2$ we can bound (27) by

$$\sum_{k,\ell=1:\lambda_k \neq \lambda_\ell}^{N} \frac{1}{|\lambda_k - \lambda_\ell|T} (|\langle v|\phi_k\rangle|^2 + |\langle\phi_\ell|v\rangle|^2) \cdot \frac{1}{2}(|\langle\phi_k|\psi_0\rangle|^2 + |\langle\psi_0|\phi_\ell\rangle|^2).$$

Summing over all $v \in V$ we obtain

$$\|P_T(\cdot|\psi_0) - P_\infty(\cdot|\psi_0)\|_1 \leq \sum_{k,\ell=1:\lambda_k \neq \lambda_\ell}^{N} \frac{|\langle\phi_k|\psi_0\rangle|^2 + |\langle\psi_0|\phi_\ell\rangle|^2}{|\lambda_k - \lambda_\ell|T}$$

$$= \sum_{h,j:h \neq j}^{M} \sum_{k \in I_h, \ell \in I_j} \frac{|\langle\phi_k|\psi_0\rangle|^2 + |\langle\psi_0|\phi_\ell\rangle|^2}{|\lambda_k - \lambda_\ell|T}$$

Define $\beta_j = \sum_{k \in I_j} |\langle\phi_k|\psi_0\rangle|^2$ for all $j = 1, \ldots, M$. Then we have $\sum_{j=1}^{M} \beta_j = 1$. Also, without loss of generality, assume that $\lambda_1 \geq \lambda_2 \geq \cdots \geq \lambda_N$. Then for $h \neq j$ and any $k \in I_h$ and $\ell \in I_j$ we have $|\lambda_k - \lambda_\ell| \geq |h - j|\Delta$ where Δ is defined in (8). Putting the above together we have

$$\|P_T(\cdot|\psi_0) - P_\infty(\cdot|\psi_0)\|_1 = \sum_{v \in V} |P_T(v|\psi_0) - P_\infty(v|\psi_0)|$$

$$\leq \sum_{h,j=1:h \neq j}^{M} \frac{\beta_h + \beta_j}{|h - j|T\Delta}$$

$$= \sum_{r=1}^{M-1} \frac{1}{rT\Delta} \sum_{h,j:|h-j|=r} \beta_h + \beta_j$$

$$\leq \frac{2 \ln M + 2}{T\Delta},$$

where in the last inequality we have used the bound $\sum_{i=1}^{n} 1/n \leq \ln n + 1$ for harmonic sums, and the fact that the inner sum is always ≤ 2.