Proving the BQP-Completeness of the Quantum Linear Systems Problem Using a Clock Construction

Tarun Prasad Ashley Zhuang

May 2022

1 Introduction

In virtually any technical field, the classical problem of solving a linear system of equations has a wide range of applications. Specifically, this problem can be formulated as finding $\vec{x}$, given an $N \times N$ matrix A and an N-dimensional vector $\vec{b}$, such that $A\vec{x} = \vec{b}$. The best classical algorithms for solving linear systems of equations scale with time polynomial in N — in fact, even just writing down the full solution $\vec{x}$ takes time $\mathcal{O}(N)$. This may be very slow if N is incredibly big, which is often the case realistically, as the datasets we use to construct linear systems are getting larger over time.

In the quantum analogue to the linear systems problem, we instead take as input the matrix A and the quantum state $|b\rangle = \frac{\sum_{i=1}^{N} b_i |i\rangle}{\| \sum_{i=1}^{N} b_i |i\rangle \|_2}$, where b_i are the entries of $\vec{b}$. The goal is to output a state close to $|x\rangle = \frac{\sum_{i=1}^{N} x_i |i\rangle}{\| \sum_{i=1}^{N} x_i |i\rangle \|_2}$, where x_i are the entries of $\vec{x}$ such that $A\vec{x} = \vec{b}$. One might wonder why we would care about this problem, since the output does not even fully specify the solution $\vec{x}$ to the system of equations. However, it is often the case that one is only interested in a function or summary statistic of the solution $\vec{x}$, which may be efficiently computable from the quantum state $|x\rangle$. For example, one is often interested in the expectation value $\langle x| M |x\rangle$ for some quantum measurement operator M. This technique can be used to extract many properties of $\vec{x}$, such as finding the total weight of certain parts of the state space.

In a groundbreaking paper by Harrow, Hassidim, and Lloyd (HHL), they invent an algorithm that solves this Quantum Linear Systems Problem in poly$(\log N)$ time for matrices satisfying certain conditions, which is an exponential speedup from any existing classical algorithms [3]. Thus, the quantum analogue to the linear systems problem has enormous implications and potential for far-reaching practical use cases. Here, we aim to extend the work done in the HHL paper by expanding on the details of and the motivation behind their proof of the BQP-completeness of matrix inversion.

2 Preliminaries

2.1 Mathematical background

Before we define the Quantum Linear Systems Problem, we'll first need to define a few properties of matrices.

Definition 2.1 (Sparsity)**.** A matrix is *s-sparse* if each row has at most s nonzero elements.

Intuitively, it should be faster to do computations with matrices of low sparsity than with dense matrices.

Definition 2.2 (Condition number)**.** Let A be an $N \times N$ normal matrix, i.e., A such that $A^\dagger A = AA^\dagger$. Let $\lambda_{\max}$ and $\lambda_{\min}$ be the largest and smallest eigenvalues (by moduli), respectively, of A. Then, the *condition number* of A is defined to be $\kappa(A) \equiv \frac{|\lambda_{\max}|}{|\lambda_{\min}|}$.

As the condition number of a matrix grows larger, the matrix becomes more and more "ill-conditioned." This is because it gets closer to a matrix with eigenvalue 0; i.e., a matrix that cannot be inverted. Thus, we will demand that our matrices have bounded condition numbers.

2.2 The Quantum Linear Systems Problem

Definition 2.3 (Quantum Linear Systems Problem (QLSP))**.** An algorithm solves the quantum linear systems problem if it has:

- **Input:** A $\mathcal{O}(1)$-sparse $N \times N$ invertible matrix A specified through an oracle, an N-dimensional quantum state $|b\rangle$, and an error threshold ϵ.

  We also require that A be *well-conditioned*; i.e., the moduli of its eigenvalues lie between $1/\kappa(A)$ and 1.

- **Output:** A quantum state $|\tilde{x}\rangle$ such that $\| |\tilde{x}\rangle - |x\rangle \|_2 \leq \epsilon$, where $|x\rangle$ is a normalized state proportional to $A^{-1}|b\rangle$.

We will show that the Quantum Linear Systems Problem is BQP-complete, where BQP (Bounded-Error Quantum Polynomial Time) is the class of problems solvable in polynomial time by a quantum computer with an error probability of at most 1/3.

3 BQP-completeness of QLSP

To show that QLSP is BQP-hard, we will show that any arbitrary BQP circuit can be simulated by an instance of QLSP. Consider a quantum circuit on $n = \log N$ qubits with input state $|\psi\rangle$ that applies unitaries $U_1, \ldots, U_T$, each of which are two-qubit gates. For notational convenience, define $U_0 = \mathbb{1}$. Note that the input state $|\psi\rangle$ is merely a classical bit-string in $\{0,1\}^n$ encoded as a quantum state. Our objective is to construct an instance of QLSP such that solving it simulates measuring the first (output) qubit of $U_T \ldots U_1 |\psi\rangle$.

3.1 Motivating our QLSP construction

The goal here is to construct a matrix A from $U_1, \ldots, U_T$ such that if QLSP returns the quantum state proportional to A^{-1} applied to the input state $|\psi\rangle$, we can simulate the output state $U_T \ldots U_1 |\psi\rangle$.

One natural way to construct an A such that we can characterize A^{-1} easily would be to define $A = \mathbb{1} - cV$ for some unitary V and constant $0 < c < 1$, because this allows us to extend the geometric series $(1 - cx)^{-1} = \sum_{t=0}^{\infty} c^t x^t$ to obtain

$$A^{-1} = (\mathbb{1} - cV)^{-1} = \sum_{t=0}^{\infty} c^t V^t. \tag{1}$$

In order to simulate the output state $U_T \ldots U_1 |\psi\rangle$ using A^{-1}, we should choose V such that when A^{-1} is applied to $|\psi\rangle$, we obtain a superposition of the intermediate states of our quantum circuit — we call (a normalized version of) this superposition a *history state* since it captures the computation history of the circuit. Then, one term in this superposition will correspond to the output state, so if we measure the history state, we will, with some probability, obtain a measurement of the output state.

The immediate problem with this is that we will not know whether or not the measurement corresponds to a measurement of the output state or of some other intermediate state of the circuit. How can we fix this? One way is to use the idea of a "clock" register from the Feynman-Kitaev clock construction [4]. Introduce an ancillary register C and construct V such that the following invariant (the "clock invariant") is maintained: if register C contains $|t\rangle$, then the second register must contain the intermediate state after t time steps of the circuit, i.e., $U_t \ldots U_1 |\psi\rangle$. Then it suffices to measure the clock register and check whether or not the measured time step corresponds to one in which the second register contains the output state. If so, it is a

successful simulation and we are done; if not, we can simply repeat the whole process with the hope that the next attempt will give us a measurement at an "output" time step. Our modified input state now becomes

$$|in\rangle = |0\rangle_C \otimes |\psi\rangle .$$

Let us next look at how we might define V. Since our expression for A^{-1} as in Equation (1) is a sum of $c^t V^t$, the corresponding expression for $A^{-1}|in\rangle$ will be a superposition of $V^t|in\rangle$. Since we want this to be a superposition of the intermediate states, a natural definition of V would be such that each application of V on an intermediate state results in the next intermediate state, while maintaining the clock invariant (i.e., incrementing the clock register). In other words, we want $V^t|in\rangle = |t\rangle_C \otimes U_t \dots U_1 |\psi\rangle$. But notice that our expression is an *infinite* sum of $c^t V^t$. So what should happen after T applications of V?

To answer this question, we will have to keep two things in mind: 1) in order for $A^{-1}|in\rangle$ to be finite-dimensional, the clock register state must live in a finite-dimensional vector space, so we cannot increment it indefinitely, and 2) we want the probability that a measurement of the history state corresponds to a measurement of the output state to be large, so we ideally want the output state to live for more than just one time step. The first issue can be fixed by performing a "reset" of the clock register back to 0 after a certain time step is reached. The second one can be fixed by maintaining the output state in the second register beyond time step T for, say, T time steps. Finally, we need to maintain the clock invariant while performing the reset, which means we need to return to state $|\psi\rangle$ in the second register: a simple way to do this would be to *un*apply $U_T, \dots, U_1$ by successively applying their conjugate transposes $U_T^\dagger, \dots, U_1^\dagger$.

3.2 Constructing an infinite circuit

To summarize our conclusions from Section 3.1, we want to construct, given a quantum circuit $U_1, \dots, U_T$, an infinite circuit that applies U_1 up to U_T, then applies T instances of $\mathbb{1}$, then applies $U_T^\dagger$ down to $U_1^\dagger$, before repeating this sequence of gates infinitely many times. This achieves our goal of constructing an infinite circuit that remains in the output state $U_T \dots U_1 |\psi\rangle$ for a large fraction of the time, and resets after every $3T$ time steps.[1] The clock register C must hence be of dimension $3T$ to keep track of which time step each intermediate state corresponds to.

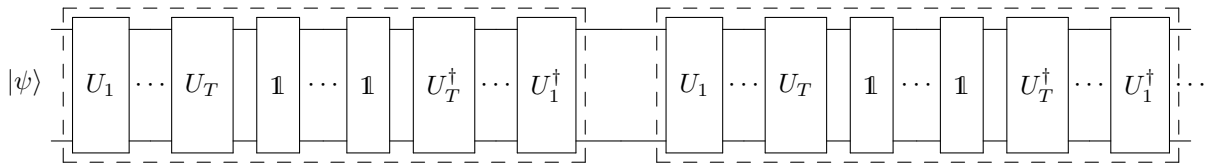

Figure 1: A visualization of our infinite circuit. The dashed box of $3T$ gates is repeated infinitely.

Based on our discussion of what we want V to achieve, the following definition of a "clock unitary" V_{clock} (or simply V for shorthand) suffices:

$$V_{clock} = \underbrace{\sum_{t=0}^{T-1} |t+1\rangle\langle t|_C \otimes U_{t+1}}_{V_{prop}} + \underbrace{\sum_{t=T}^{2T-1} |t+1\rangle\langle t|_C \otimes \mathbb{1}}_{V_{noop}} + \underbrace{\sum_{t=2T}^{3T-1} |t+1 \bmod 3T\rangle\langle t|_C \otimes U_{3T-t}^\dagger}_{V_{deprop}} . \tag{2}$$

We've constructed this unitary such that each application of V_{clock} to an intermediate state both increments the clock register by 1 and applies the next gate to the second register, effectively allowing us to walk

[1] Technically, we could apply all of $U_1, \dots, U_T$ in a single time step and then all of $U_T^\dagger, \dots, U_1^\dagger$ also in a single time step, but we instead apply them one at a time in order to keep A sparse (see Lemma 3.2).

through each time step of our circuit by repeatedly applying V_{clock} to $|in\rangle$. This is because when applied to an intermediate state corresponding to any time-step t, there is exactly one non-zero term in the resulting state. Starting with $|in\rangle$, the non-zero term in each of the first T applications of V_{clock} is within the propagation unitary V_{prop}. After T applications, we are in state $|T\rangle_C \otimes U_T \dots U_1 |\psi\rangle$, and the following T applications continue to increment the clock while leaving the second register untouched (i.e., performing T no-ops as defined by V_{noop}). Finally, the next T (depropagation) applications undo the unitaries until we end up with the original state $|0\rangle_C \otimes |\psi\rangle$, at which point the clock has wrapped back around to 0 and we repeat the process. Let's illustrate this pattern mathematically:

$$
\begin{array}{lll}
V^0 |in\rangle = |0\rangle |\psi\rangle & V^T |in\rangle = |T\rangle U_T \dots U_1 |\psi\rangle & V^{2T} |in\rangle = |2T\rangle U_T \dots U_1 |\psi\rangle \\
V^1 |in\rangle = |1\rangle U_1 |\psi\rangle & V^{T+1} |in\rangle = |T+1\rangle U_T \dots U_1 |\psi\rangle & V^{2T+1} |in\rangle = |2T+1\rangle U_{T-1} \dots U_1 |\psi\rangle \\
V^2 |in\rangle = |2\rangle U_2 U_1 |\psi\rangle & V^{T+2} |in\rangle = |T+2\rangle U_T \dots U_1 |\psi\rangle & \vdots \\
\vdots & \vdots & V^{3T-1} |in\rangle = |3T-1\rangle U_1 |\psi\rangle \\
V^{T-1} |in\rangle = |T-1\rangle U_{T-1} \dots U_1 |\psi\rangle & V^{2T-1} |in\rangle = |2T-1\rangle U_T \dots U_1 |\psi\rangle & V^{3T} |in\rangle = |0\rangle \psi = |in\rangle.
\end{array}
$$

In general, we have

$$
V_{clock}^t |in\rangle = V^t (|0\rangle_C \otimes |\psi\rangle) = \begin{cases} |t \bmod 3T\rangle_C \otimes U_{t \bmod 3T} \dots U_0 |\psi\rangle & \text{if } 0 \le t \bmod 3T < T \\ |t \bmod 3T\rangle_C \otimes U_T \dots U_0 |\psi\rangle & \text{if } T \le t \bmod 3T < 2T \\ |t \bmod 3T\rangle_C \otimes U^\dagger_{3T-t \bmod 3T} \dots U_0 |\psi\rangle & \text{if } 2T \le t \bmod 3T < 3T \end{cases} . \quad (3)
$$

We see from (3) that $V^t |in\rangle$ represents the quantum state of our infinite circuit after t time steps with the appropriate value for our clock register, as desired. Therefore, our history state is simply

$$
|history\rangle \propto \sum_{t=0}^{\infty} c^t V^t |in\rangle ,
$$

as required. As for the choice of constant for c, setting $c = e^{-1/T}$ ensures that we have bounds on the condition number of A, as we shall see in Lemma 3.1. Thus, we've constructed matrices $A = \mathbb{1} - Ve^{-1/T}$ and A^{-1} such that $A^{-1} |in\rangle \propto |history\rangle$.

3.3 Concluding that QLSP is BQP-complete

Before we can argue that our reduction is valid, we must check that the matrix A we constructed satisfies the requirements specified by QLSP, as we shall do in the following lemmas.

Lemma 3.1. *The condition number $\kappa(A)$ of $A = \mathbb{1} - Ve^{-1/T}$ is $\mathcal{O}(T)$.*

Proof. Let $\lambda_{\max}$ and $\lambda_{\min}$ be the largest and smallest eigenvalues, by moduli, of A. Then, using a similar approach as Lemma 7 in [2], first note that

$$
|\lambda_{\max}| \le \max_{\|\mathbf{x}\|_2 = 1} \left| \mathbf{x}^\dagger A \mathbf{x} \right| = \max_{\|\mathbf{x}\|_2 = 1} \left| \mathbf{x}^\dagger \left(\mathbb{1} - Ve^{-1/T} \right) \mathbf{x} \right| \le 1 + \max_{\|\mathbf{x}\|_2 = 1} \left| \mathbf{x}^\dagger V \mathbf{x} \right| \cdot e^{-1/T} = 1 + e^{-1/T}. \quad (4)
$$

We obtain the first inequality above from the fact that $\mathbf{x}^\dagger A \mathbf{x} = \lambda$ for any unit eigenvector $\mathbf{x}$ of A with corresponding eigenvalue λ. The second inequality follows from the triangle inequality. The last equality is due to the fact that V is unitary, which means that $V\mathbf{x}$ is a unit vector, so $|\mathbf{x}^\dagger V \mathbf{x}| \le 1$ for all $\mathbf{x}$, with equality attained if $\mathbf{x}$ is a unit eigenvector.

Similarly,

$$
|\lambda_{\min}| \ge \min_{\|\mathbf{x}\|_2 = 1} \left| \mathbf{x}^\dagger \left(\mathbb{1} - Ve^{-1/T} \right) \mathbf{x} \right| \ge 1 - \max_{\|\mathbf{x}\|_2 = 1} \left| \mathbf{x}^\dagger V \mathbf{x} \right| \cdot e^{-1/T} = 1 - e^{-1/T}, \quad (5)
$$

Now, note that

$$e^{-1/T} \leq 1 - \frac{1}{2T}$$

because $e^{-x} \leq 1 - \frac{x}{2}$ for all $x \in (0, 1]$, since e^{-x} is convex and intersects the line $1 - \frac{x}{2}$ at $x = 0$ and $x \approx 1.6$. Thus, dividing Equation (4) by (5) gives us

$$\kappa(A) = \frac{|\lambda_{\max}|}{|\lambda_{\min}|} \leq \frac{1 + e^{-1/T}}{1 - e^{-1/T}} \leq \frac{1+1}{1 - \left(1 - \frac{1}{2T}\right)} = 4T.$$

Since $\kappa(A) \leq 4T$, we can conclude that $\kappa(A) = \mathcal{O}(T)$. □

Lemma 3.2. *$A = \mathbb{1} - Ve^{-1/T}$ is $\mathcal{O}(1)$-sparse.*

Proof. From the definition in Equation (2), we can represent V in matrix form by noticing that each $|t+1\rangle\langle t|$ term is a $3T \times 3T$ matrix consisting only of 0s with exactly one 1 in the $(t+1)$th row and tth column. V is therefore a $3TN \times 3TN$ matrix that looks like

$$V = \begin{bmatrix} & & & & & & & & & & & U_1^\dagger \\ U_1 & & & & & & & & & & & \\ & \ddots & & & & & & & & & & \\ & & U_T & & & & & & & & & \\ & & & \mathbb{1} & & & & & & & & \\ & & & & \ddots & & & & & & & \\ & & & & & \mathbb{1} & & & & & & \\ & & & & & & U_T^\dagger & & & & & \\ & & & & & & & \ddots & & & & \\ & & & & & & & & U_2^\dagger & & & \end{bmatrix},$$

where the blank entries above are zeroes. Since each submatrix U_t, $\mathbb{1}$, and $U_t^\dagger$ in V is $\mathcal{O}(1)$-sparse (as the U_t's are 2-qubit gates), and the submatrices themselves lie within distinct rows, V must also be sparse. Since A is defined to be $\mathbb{1} - Ve^{-1/T}$, at most one additional non-zero term is introduced in each row, implying that A is also $\mathcal{O}(1)$-sparse. □

Theorem 3.3. *The Quantum Linear Systems Problem is BQP-hard.*

Proof. To simulate our original circuit, we construct our matrix V_{clock} according to (2) from $U_1, \ldots, U_T$. Next, we construct a QLSP instance with inputs $A = \mathbb{1} - V_{clock}e^{-1/T}$, $|b\rangle = |0\rangle_C \otimes |\psi\rangle$, and some small constant ϵ. Note that A is a valid input since it is $\mathcal{O}(1)$-sparse by Lemma 3.2. Moreover, the condition number is small since it is bounded by $\mathcal{O}(T)$ from Lemma 3.1, so we can rescale A to be well-conditioned.

The output of this QLSP instance will be an ϵ-approximation of the history state $|history\rangle \propto A^{-1}|b\rangle$. It is easy to verify that AA^{-1} does in fact equal $\mathbb{1}$:

$$AA^{-1} = (\mathbb{1} - Ve^{1/T}) \sum_{t=0}^{\infty} V^t e^{-t/T} = \sum_{t=0}^{\infty} V^t e^{-t/T} - \sum_{t=0}^{\infty} V^{t+1} e^{-(t+1)/T} = V^0 e^{-0/T} = \mathbb{1}.$$

By our infinite circuit construction, if we measure the clock register of $|history\rangle$ and obtain $|t\rangle$ such that $T \leq t < 2T$, we know that the quantum state has collapsed to $|t\rangle \otimes U_T \ldots U_1 |\psi\rangle$. Thus, we have successfully simulated the original circuit's output. By (3), the terms in $|history\rangle \propto \sum_{t=0}^{\infty} e^{-t/T} V^t$ with the desired clock

register values have coefficients $e^{-(3T\cdot k+t)/T}$ for $k \in \mathbb{Z}^{\geq 0}$ and $t \in \{T, \ldots, 2T-1\}$. Hence, the probability of success is

$$\frac{\sum_{k=0}^{\infty}\sum_{t=T}^{2T-1}(e^{-(3T\cdot k+t)/T})^2}{\sum_{t=0}^{\infty}(e^{-t/T})^2} = \frac{\sum_{k=0}^{\infty}e^{-6k}\sum_{t=T}^{2T-1}e^{-2t/T}}{1/(1-e^{-2/T})} = e^{-2}(1-e^{-2})\sum_{k=0}^{\infty}e^{-6k} = \frac{e^{-2}}{1+e^{-2}+e^{-4}} > \frac{1}{10},$$

where the simplifications above follow from the sum of finite and infinite geometric series.

If we instead measure the clock register and obtain $|t\rangle$ such that $t \notin [T, 2T)$, we can just repeat the QLSP simulation until we are successful, since the probability of success is $> 1/10$ (we can do this since A is specified via an oracle and $|\psi\rangle$ is classical).[2] Thus, our reduction is complete, so QLSP is in fact BQP-hard. □

The number of identities chosen in our infinite circuit construction may have seemed somewhat arbitrary initially, but we see from our proof above that the choice of performing T identities is ideal since it makes the probability of success independent of T.

We have successfully shown that QLSP is BQP-hard above; moreover, the existence of the HHL algorithm (see [3]) that actually solves the Quantum Linear Systems Problem in time $\mathcal{O}\left(\log(N)\kappa^2/\epsilon\right)$ illustrates to us that QLSP is in BQP as well for $1/\epsilon \in \text{poly}(N)$. Therefore, QLSP is BQP-complete, and HHL gives us an exponential speedup compared to the classical case when $1/\epsilon$ and κ are $\text{poly}(\log N)$. Note that HHL typically also assumes that A is Hermitian, but this condition can be relaxed to match our definition of QLSP by using the Hermitian matrix $\begin{pmatrix} 0 & A \\ A^\dagger & 0 \end{pmatrix}$ and vector $\begin{pmatrix} \vec{b} \\ 0 \end{pmatrix}$ in HHL instead for non-Hermitian A.

4 Discussion

Informally, the BQP-completeness of QLSP means that it is in the set of hardest problems that can be solved in polynomial time using a quantum algorithm. This fact can be used to make some claims about the optimality of the HHL algorithm.

For instance, consider the complexity class BPP, which is the classical analogue of BQP: the set of problems solvable by a classical probabilistic algorithm with error probability bounded by 1/3. Note that BPP $\subseteq$ BQP holds trivially since quantum algorithms can simulate classical probabilistic ones. Now if there were to exist a *classical* algorithm that could solve QLSP in $\text{poly}(\kappa, \log N)$ time, then an implication would be that BPP = BQP, because any quantum circuit could be reduced in polynomial time to an instance of QLSP, which could then be solved in polynomial time using the classical algorithm whose existence we assumed.

But this is considered a highly unlikely result since classical probabilistic algorithms would then be able to solve all problems that a quantum algorithm can in at most polynomially worse time. Among other things, this would imply the existence of a classical polynomial time randomized factoring algorithm, the lack of which the security of several modern cryptosystems depends on.

The authors of [3] (Appendix A, Theorem 5) also show that if the κ dependence in the runtime of a QLSP algorithm can be brought down to sub-linear, i.e., if there exists a quantum algorithm that solves QLSP in $\mathcal{O}\left(\kappa^{1-\delta}\text{poly}\log N\right)$ for $\delta > 0$, then BQP = PSPACE (the set of problems solvable in polynomial space classically), which is also a very unlikely result.

Some future directions for our work include 1) optimizing the coefficients chosen in the history state in order to maximize the success probability, while still ensuring that A remains well-conditioned, and 2) motivating analogous clock constructions used in completeness proofs of other problems — such as the MA-completeness proof of "stoquastic 6-SAT" as defined in [1] — with the goal of creating a general framework for proving completeness for problems in certain complexity classes.

[2] Also note that we are working with an ϵ-approximation of $A^{-1}|b\rangle$ (normalized) rather than the precise state, but we can repeat the whole process to ensure that the error probability of the simulation is within a desired range. See Appendix A, Section 5 of [3] for a more rigorous explanation.

5 Acknowledgements

We'd like to thank Anurag for inspiring the topic of this paper and for his guidance and suggestions on our draft. We'd also like to thank Chi-Ning for reviewing our draft and helping us through some mathematical details of our proofs. Finally, we'd like to again thank the CS231 teaching staff Anurag, Chi-Ning, and Prayaag for all their instruction and support throughout this wonderful semester in the course.

An additional thank you to the markers in SEC Classroom 1.414 for enduring several hours of scribbling bras and kets on the whiteboards. :)

6 References

[1] Sergey Bravyi, David P. DiVincenzo, Roberto I. Oliveira, and Barbara M. Terhal. The complexity of stoquastic local hamiltonian problems, 2007.

[2] Danial Dervovic, Mark Herbster, Peter Mountney, Simone Severini, Naïri Usher, and Leonard Wossnig. Quantum linear systems algorithms: a primer, 2018.

[3] Aram W. Harrow, Avinatan Hassidim, and Seth Lloyd. Quantum algorithm for linear systems of equations. *Physical Review Letters*, 103(15), October 2009.

[4] A. Yu. Kitaev, A. H. Shen, and M. N. Vyalyi. *Classical and quantum computation*. AMS, Providence, RI, 2002.