

Quantum Algorithm Implementations for Beginners

ABHIJITH J.*, ADETOKUNBO ADEDOYIN, JOHN AMBROSIANO, PETR ANISIMOV, WILLIAM CASPER, GOPINATH CHENNUPATI, CARLETON COFFRIN, HRISTO DJIDJEV, DAVID GUNTER, SATISH KARRA, NATHAN LEMONS, SHIZENG LIN, ALEXANDER MALYZHENKOV, DAVID MASCARENAS, SUSAN MNISZEWSKI, BALU NADIGA, DANIEL O’MALLEY, DIANE OYEN, SCOTT PAKIN, LAKSHMAN PRASAD, RANDY ROBERTS, PHILLIP ROMERO, NANDAKISHORE SANTHI, NIKOLAI SINITSYN, PIETER J. SWART, JAMES G. WENDELBERGER, BORAM YOON, RICHARD ZAMORA, WEI ZHU, STEPHAN EIDENBENZ*, ANDREAS BÄRTSCHI*, PATRICK J. COLES*, MARC VUFFRAY*, and ANDREY Y. LOKHOV*,

Los Alamos National Laboratory, Los Alamos, New Mexico 87545, USA

As quantum computers become available to the general public, the need has arisen to train a cohort of quantum programmers, many of whom have been developing classical computer programs for most of their careers. While currently available quantum computers have less than 100 qubits, quantum computing hardware is widely expected to grow in terms of qubit count, quality, and connectivity. This review aims to explain the principles of quantum programming, which are quite different from classical programming, with straightforward algebra that makes understanding of the underlying fascinating quantum mechanical principles optional. We give an introduction to quantum computing algorithms and their implementation on real quantum hardware. We survey 20 different quantum algorithms, attempting to describe each in a succinct and self-contained fashion. We show how these algorithms can be implemented on IBM’s quantum computer, and in each case, we discuss the results of the implementation with respect to differences between the simulator and the actual hardware runs. This article introduces computer scientists, physicists, and engineers to quantum algorithms and provides a blueprint for their implementations.

CONTENTS

Abstract	1
Contents	1
1 Introduction	3
1.1 The quantum computing programming model	4
1.2 Implementations on a real quantum computer	11
1.3 Classes of quantum algorithms	15
2 Grover’s Algorithm	16
2.1 Problem definition and background	16
2.2 Algorithm description	17
2.3 Algorithm implemented on IBM’s 5-qubit computer	18
3 Bernstein-Vazirani Algorithm	19
3.1 Problem definition and background	19
3.2 Algorithm description	19
3.3 Algorithm implemented on IBM’s 5-qubit and 16-qubit computers	20

*abhijithj@lanl.gov; eidenben@lanl.gov; baertschi@lanl.gov; pcoles@lanl.gov; vuffray@lanl.gov; lokhov@lanl.gov. LA-UR-20-22353

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2022 Copyright held by the owner/author(s).

2643-6817/2022/3-ART

<https://doi.org/10.1145/3517340>

4	Linear Systems	22
4.1	Problem definition and background	22
4.2	Algorithm description	23
4.3	Phase estimation	23
4.4	Algorithm implemented on IBM's 5 qubit computer	26
5	Shor's Algorithm for Integer Factorization	27
5.1	Problem definition and background	27
5.2	Algorithm description	29
5.3	Algorithm implemented on IBM's 5-qubit computer	30
6	Matrix Elements of Group Representations	32
6.1	Problem definition and background	32
6.2	Algorithm description	35
6.3	Algorithm implemented on IBM's 5-qubit computer	36
7	Quantum Verification of Matrix Products	37
7.1	Problem definition and background	37
7.2	Algorithm description	37
8	Group Isomorphism	39
8.1	Problem definition and background	39
8.2	Algorithm description	39
8.3	Algorithm implemented using Qiskit	41
9	Quantum Random Walks	42
9.1	Problem definition and background	42
9.2	Example of a quantum random walk	43
9.3	Algorithm implementation using Qiskit on IBM Q	44
10	Quantum Minimal Spanning Tree	45
10.1	Problem definition and background	45
10.2	Algorithm description	45
11	Quantum Maximum Flow Analysis	49
11.1	Problem definition and background	49
11.2	Algorithm description	51
12	Quantum Approximate Optimization Algorithm	52
12.1	Problem definition and background	52
12.2	Algorithm description	54
12.3	QAOA MaxCut on ibmqx2	57
12.4	A proof-of-concept experiment	58
13	Quantum Principal Component Analysis	61
13.1	Problem definition and background	61
13.2	Algorithm description	62
13.3	Algorithm implemented on IBM's 5-qubit computer	63
14	Quantum Support Vector Machine	65
15	Quantum Simulation of the Schrödinger Equation	66
15.1	Problem definition and background	66
15.2	Algorithm description	67
15.3	Algorithm implemented on IBM's 5-qubit computer	68
16	Ground State of the Transverse Ising Model	69
16.1	Variational quantum eigenvalue solver	69

16.2	Simulation and results	71
17	Quantum Partition Function	72
17.1	Background on the partition function	72
17.2	A simple example	75
17.3	Calculating the quantum partition function	75
17.4	Implementation of a quantum algorithm on the IBM Quantum Experience	76
18	Quantum State Preparation	77
18.1	Single qubit state preparation	77
18.2	Schmidt decomposition	78
18.3	Two-qubit state preparation	80
18.4	Two-qubit gate preparation	80
18.5	Four qubit state preparation	81
19	Quantum Tomography	81
19.1	Problem definition and background	81
19.2	Short survey of existing methods	84
19.3	Implementation of the Maximum Likelihood method on 5-qubit IBM QX	84
20	Tests of Quantum Error Correction in IBM Q	86
20.1	Problem definition and background	87
20.2	Test 1: errors in single qubit control	88
20.3	Test 2: errors in entangled 3 qubits control	89
20.4	Discussion	89
	Acknowledgments	90
	References	90

1 INTRODUCTION

Quantum computing exploits quantum-mechanical effects—in particular superposition, entanglement, and quantum tunneling—to more efficiently execute a computation. Compared to traditional, digital computing, quantum computing offers the potential to dramatically reduce both execution time and energy consumption. These potential advantages, steady advances in nano-manufacturing, and the slow-down of traditional hardware scaling laws (such as Moore’s Law) have led to a substantial commercial and national-security interest and investment in quantum computing technology in the 2010s. Recently, Google announced that it has reached a major milestone known as quantum supremacy—the demonstration that a quantum computer can perform a calculation that is intractable on a classical supercomputer [9]. The problem tackled here by the quantum computer is not one with any direct real-world application. Nonetheless, this is a watershed moment for quantum computing and is widely seen as an important step on the road towards building quantum computers that will offer practical speedups when solving real-world problems [100]. (See [3] for a precise technical definition of quantum supremacy.)

While the mathematical basis of quantum computing, the programming model, and most quantum algorithms have been published decades ago (starting in the 1990s), they have been of interest only to a small dedicated community. We believe the time has come to make quantum algorithms and their implementations accessible to a broad swath of researchers and developers across computer science, software engineering, and other fields. The quantum programming model is fundamentally different from traditional computer programming. It is also dominated by physics and algebraic notations that at times present unnecessary entry barriers for mainstream computer scientists and other more mathematically trained scientists.

In this review, we provide a self-contained, succinct description of quantum computing and of the basic quantum algorithms with a focus on implementation. Since real quantum computers, such as IBM Q [69], are now available as a cloud service, we present results from simulator and actual hardware experiments for smaller input data sets. Other surveys of quantum algorithms with a different target audience and also without actual implementations include [11, 30, 72, 90, 91, 106]. Other cloud service based quantum computers are also available from Rigetti and IonQ, but in this review we will focus solely on IBM's quantum computing ecosystem. The code and implementations accompanying the paper can be found at https://github.com/lanl/quantum_algorithms.

1.1 The quantum computing programming model

Here we provide a self-contained description of the quantum computing programming model. We will define the common terms and concepts used in quantum algorithms literature. We will not discuss how the constructs explained here are related to the foundations of quantum mechanics. Interested readers are encouraged to take a look at Ref. [92] for a more detailed account along those lines. Readers with a computer science background are referred to Refs. [82, 103, 136], for a more comprehensive introduction to quantum computing from a computer science perspective.

Quantum computing basically deals with the manipulation of quantum systems. The physical details of this is dependent on the quantum computer's hardware design. Here we will only talk about the higher level abstractions used in quantum computing: a typical programmer will only be exposed to these abstractions. The state of any quantum system is always represented by a vector in a complex vector space (usually called a Hilbert space). Quantum algorithms are always expressible as transformations acting on this vector space. These basic facts follow from the axioms of quantum mechanics. Now we will explain some of the basic concepts and terminology used in quantum computing.

1.1.1 The qubit. The qubit (short for 'quantum bit') is the fundamental information carrying unit used in quantum computers. It can be seen as the quantum mechanical generalization of a bit used in classical computers. More precisely, a qubit is a two dimensional quantum system. The state of a qubit can be expressed as,

$$|\phi\rangle = \alpha |0\rangle + \beta |1\rangle . \quad (1)$$

Here α and β are complex numbers such that, $|\alpha|^2 + |\beta|^2 = 1$. In the *ket-notation* or the *Dirac notation*, $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$ are shorthands for the vectors encoding the two basis states of a two dimensional vector space. So according to this notation, Eq. (1) expresses the fact that the state of the qubit is the two dimensional complex vector $\begin{pmatrix} \alpha \\ \beta \end{pmatrix}$. Unlike a classical bit the state of a qubit cannot be measured without changing it. Measuring a qubit, whose state given by Eq. (1), will yield the classical value of either zero ($|0\rangle$) with probability $|\alpha|^2$ or one ($|1\rangle$) with probability $|\beta|^2$. Qubit implementations and technologies are a very active area of research that is not the focus of our review, an interested reader is referred to [80] for a survey.

1.1.2 System of qubits. The mathematical structure of a qubit generalizes to higher dimensional quantum systems as well. The state of any quantum system is a normalized vector (a vector of norm one) in a complex vector space. The normalization is necessary to ensure that the total probability of all the outcomes of a measurement sum to one.

A quantum computer contains many number of qubits. So it is necessary to know how to construct the combined state of a system of qubits given the states of the individual qubits. The joint state of a system of qubits is described using an operation known as the *tensor product*, \otimes . Mathematically, taking the tensor product of two

states is the same as taking the Kronecker product of their corresponding vectors. Say we have two single qubit states $|\phi\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ and $|\phi'\rangle = \begin{pmatrix} \alpha' \\ \beta' \end{pmatrix}$. Then the full state of a system composed of two independent qubits is given by,

$$|\phi\rangle \otimes |\phi'\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} \otimes \begin{pmatrix} \alpha' \\ \beta' \end{pmatrix} = \begin{pmatrix} \alpha\alpha' \\ \alpha\beta' \\ \beta\alpha' \\ \beta\beta' \end{pmatrix} \quad (2)$$

Sometimes the \otimes symbol is dropped all together while denoting the tensor product to reduce clutter. Instead the states are written inside a single ket. For example, $|\phi\rangle \otimes |\phi'\rangle$ is shortened to $|\phi\phi'\rangle$, and $|0\rangle \otimes |0\rangle \otimes |0\rangle$ is shortened to $|000\rangle$. For larger systems the Dirac notation gives a more succinct way to compute the tensor product using the distributive property of the Kronecker product. For a system of, say, three qubits with each qubit in the state $|y_j\rangle = \alpha_j |0\rangle + \beta_j |1\rangle$, for $j = 1, 2, 3$, the joint state is,

$$|y_1 y_2 y_3\rangle = |y_1\rangle \otimes |y_2\rangle \otimes |y_3\rangle \quad (3)$$

$$\begin{aligned} &= \alpha_1 \alpha_2 \alpha_3 |000\rangle + \alpha_1 \alpha_2 \beta_3 |001\rangle + \alpha_1 \beta_2 \alpha_3 |010\rangle + \alpha_1 \beta_2 \beta_3 |011\rangle \\ &\quad + \beta_1 \alpha_2 \alpha_3 |100\rangle + \beta_1 \alpha_2 \beta_3 |101\rangle + \beta_1 \beta_2 \alpha_3 |110\rangle + \beta_1 \beta_2 \beta_3 |111\rangle \end{aligned} \quad (4)$$

A measurement of all three qubits could result in any of the eight (2^3) possible bit-strings associated with the eight basis vectors. One can see from these examples that the dimension of the state space grows exponentially in the number of qubits n and that the number of basis vectors is 2^n .

1.1.3 Superposition and entanglement. *Superposition* refers to the fact that any linear combination of two quantum states, once normalized, will also be a valid quantum state. The upshot to this is that any quantum state can be expressed as a linear combination of a few basis states. For example, we saw in Eq. (1) that any state of a qubit can be expressed as a linear combination of $|0\rangle$ and $|1\rangle$. Similarly, the state of any n qubit system can be written as a normalized linear combination of the 2^n bit-string states (states formed by the tensor products of $|0\rangle$'s and $|1\rangle$'s). The orthonormal basis formed by the 2^n bit-string states is called the *computational basis*.

Notice that Eq. (3) described a system of three qubits whose complete state was the tensor product of three different single qubit states. But it is possible for three qubits to be in a state that cannot be written as the tensor product of three single qubit states. An example of such a state is,

$$|\psi\rangle = \frac{1}{\sqrt{2}}(|000\rangle + |111\rangle). \quad (5)$$

States of a system of which cannot be expressed as a tensor product of states of its individual subsystems are called *entangled states*. For a system of n qubits, this means that an entangled state cannot be written a tensor product of n single qubit states. The existence of entangled states is a physical fact that has important consequences for quantum computing, and quantum information processing in general. In fact, without the existence of such states quantum computers would be no more powerful than their classical counterparts [128]. Entanglement makes it possible to create a complete 2^n dimensional complex vector space to do our computations in, using just n physical qubits.

1.1.4 Inner and outer products. We will now discuss some linear algebraic notions necessary for understanding quantum algorithms. First of these is the *inner product* or *overlap* between two quantum states. As we have seen before, quantum states are vectors in complex vectors spaces. The overlap between two states is just the inner product between these complex vectors. For example, take two single qubit states, $|\phi\rangle = \alpha |0\rangle + \beta |1\rangle$ and

$|\psi\rangle = \gamma |0\rangle + \delta |1\rangle$. The overlap between these states is denoted in the ket notation as $\langle\psi|\phi\rangle$. And this is given by,

$$\langle\psi|\phi\rangle = \gamma^* \alpha + \delta^* \beta, \quad (6)$$

where $*$ denotes the complex conjugate. Notice that, $\langle\psi|\phi\rangle = \langle\phi|\psi\rangle^*$. The overlap of two states is in general a complex number. The overlap of a state with a bit-string state will produce the corresponding coefficient. For instance from Eq. (1), $\langle 0|\phi\rangle = \alpha$ and $\langle 1|\phi\rangle = \beta$. And from Eq. (3), $\langle 001|\gamma_1\gamma_2\gamma_3\rangle = \alpha_1\alpha_2\beta_3$. Another way to look at overlaps between quantum states is by defining what is called a *bra* state. The states we have seen so far are ket states, like $|\phi\rangle$, which represented column vectors. A bra state corresponding to this ket state, written as $\langle\phi|$, represents a row vector with complex conjugated entries. For instance $|\phi\rangle = \begin{pmatrix} \alpha \\ \beta \end{pmatrix}$ implies that $\langle\phi| = (\alpha^* \ \beta^*)$. The overlap of two states is then the matrix product of a row vector with a column vector, yielding a single number. The reader must have already noticed the wordplay here. The overlap, with its closing angled parenthesis, form a 'bra-ket'!

The *outer product* of two states is an important operation that outputs a matrix given two states. The outer product of the two states we defined above will be denoted by, $|\psi\rangle\langle\phi|$. Mathematically the outer product of two states is a matrix obtained by multiplying the column vector of the first state with the complex conjugated row vector of the second state (notice how the ket is written before the bra to signify this). For example,

$$|\psi\rangle\langle\phi| = \begin{pmatrix} \alpha \\ \beta \end{pmatrix} (\gamma^* \ \delta^*) = \begin{pmatrix} \alpha\gamma^* & \alpha\delta^* \\ \beta\gamma^* & \beta\delta^* \end{pmatrix} \quad (7)$$

In this notation any matrix can be written as a linear combination of outer products between bit-string states. For a 2×2 matrix,

$$A = \begin{pmatrix} A_{00} & A_{01} \\ A_{10} & A_{11} \end{pmatrix} = A_{00} |0\rangle\langle 0| + A_{01} |0\rangle\langle 1| + A_{10} |1\rangle\langle 0| + A_{11} |1\rangle\langle 1|. \quad (8)$$

Acting on a state with a matrix then becomes just an exercise in computing overlaps between states. Let us demonstrate this process:

$$\begin{aligned} A|\phi\rangle &= A_{00} |0\rangle\langle 0|\phi\rangle + A_{01} |0\rangle\langle 1|\phi\rangle + A_{10} |1\rangle\langle 0|\phi\rangle + A_{11} |1\rangle\langle 1|\phi\rangle, \\ &= (A_{00}\alpha + A_{01}\beta) |0\rangle + (A_{10}\alpha + A_{11}\beta) |1\rangle = \begin{pmatrix} A_{00}\alpha + A_{01}\beta \\ A_{10}\alpha + A_{11}\beta \end{pmatrix}. \end{aligned} \quad (9)$$

This notation might look tedious at first glance but it makes algebraic manipulations of quantum states easily understandable. This is especially true when we are dealing with large number of qubits as otherwise we would have to explicitly write down exponentially large matrices.

The outer product notation for matrices also gives an intuitive input-output relation for them. For instance, the matrix $|0\rangle\langle 1| + |1\rangle\langle 0|$ can be read as "output 0 when given a 1 and output 1 when given a 0". Likewise, the matrix, $|00\rangle\langle 00| + |01\rangle\langle 01| + |10\rangle\langle 11| + |11\rangle\langle 10|$ can be interpreted as the mapping $\{ "00" \rightarrow "00", "01" \rightarrow "01", "11" \rightarrow "10", "10" \rightarrow "11" \}$. But notice that this picture becomes a bit tedious when the input is in a superposition. In that case the correct output can be computed like in Eq. (9).

1.1.5 Measurements. Measurement corresponds to transforming the quantum information (stored in a quantum system) into classical information. For example, measuring a qubit typically corresponds to reading out a classical bit, i.e., whether the qubit is 0 or 1. A central principle of quantum mechanics is that measurement outcomes are *probabilistic*.

Using the aforementioned notation for inner products, for the single qubit state in Eq. (1), the probability of obtaining $|0\rangle$ after measurement is $|\langle 0|\phi\rangle|^2$ and the probability of obtaining $|1\rangle$ after measurement is $|\langle 1|\phi\rangle|^2$.

So measurement probabilities can be represented as the squared absolute values of overlaps. Generalizing this, the probability of getting the bit string $|x_1 \dots x_n\rangle$ after measuring an n qubit state, $|\phi\rangle$, is then $|\langle x_1 \dots x_n | \phi \rangle|^2$.

Now consider a slightly more complex case of measurement. Suppose we have a three qubit state, $|\psi\rangle$ but we only measure the first qubit and leave the other two qubits undisturbed. What is the probability of observing a $|0\rangle$ in the first qubit? This probability will be given by,

$$\sum_{(x_2 x_3) \in \{0,1\}^2} |\langle 0x_2x_3 | \phi \rangle|^2. \quad (10)$$

The state of the system after this measurement will be obtained by normalizing the state,

$$\sum_{(x_2 x_3) \in \{0,1\}^2} \langle 0x_2x_3 | \phi \rangle |0x_2x_3\rangle. \quad (11)$$

Applying this paradigm to the state in Eq. (5) we see that the probability of getting $|0\rangle$ in the first qubit will be 0.5, and if this result is obtained, the final state of the system would change to $|000\rangle$. On the other hand, if we were to measure $|1\rangle$ in the first qubit we would end up with a state $|111\rangle$. Similarly we can compute the effect of subsystem measurements on any n qubit state.

In some cases we will need to do measurements on a basis different from the computational basis. This can be achieved by doing an appropriate transformation on the qubit register before measurement. Details of how to do this is given in a subsequent section discussing observables and expectation values.

The formalism discussed so far is sufficient to understand all measurement scenarios in this paper. We refer the reader to Ref. [92] for a more detailed and more general treatment of measurement.

1.1.6 Unitary transformations and gates. A qubit or a system of qubits changes its state by going through a series of *unitary transformations*. A unitary transformation is described by a matrix U with complex entries. The matrix U is called unitary if

$$UU^\dagger = U^\dagger U = I, \quad (12)$$

where U^\dagger is the transposed, complex conjugate of U (called its *Hermitian conjugate*) and I is the identity matrix. A qubit state $|\phi\rangle = \alpha |0\rangle + \beta |1\rangle$ evolves under the action of the 2×2 matrix U according to

$$|\phi\rangle \rightarrow U|\phi\rangle = \begin{pmatrix} U_{00} & U_{01} \\ U_{10} & U_{11} \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \end{pmatrix} = \begin{pmatrix} U_{00}\alpha + U_{01}\beta \\ U_{10}\alpha + U_{11}\beta \end{pmatrix}. \quad (13)$$

Operators acting on different qubits can be combined using the Kronecker product. For example, if U_1 and U_2 are operators acting on two different qubits then the full operator acting on the combined two qubit system will be given by $U_1 \otimes U_2$.

For an n qubit system the set of physically allowed transformations, excluding measurements, consists of all $2^n \times 2^n$ unitary matrices. Notice that the size of a general transformation increases exponentially with the number of qubits. In practice a transformation on n qubits is effected by using a combination of unitary transformations that act only on one or two qubits at a time. By analogy to classical logic gates like NOT and AND, such basic unitary transformations, which are used to build up more complicated n qubit transformations, are called *gates*. Gates are unitary transformations themselves and from Eq. (12) it is clear that unitarity can only be satisfied if the number of input qubits is equal to the number of output qubits. Also, for every gate U it is always possible to have another gate U^\dagger that undoes the transformation. So unlike classical gates quantum gates have to be reversible. *Reversible* means that the gate's inputs can always be reconstructed from the gate's outputs. For instance, a classical NOT gate, which maps 0 to 1 and 1 to 0 is reversible because an output of 1 implies the input was 0 and vice versa. However, a classical AND gate, which returns 1 if and only if both of its inputs are 1, is not reversible.

An output of 1 implies that both inputs were 1, but an output of 0 provides insufficient information to determine if the inputs were 00, 01, or 10.

But this extra restriction of reversibility does not mean that quantum gates are ‘less powerful’ than classical gates. Even classical gates can be made reversible with minimal overhead. Reversibility does not restrict their expressive power [105]. Quantum gates can then be seen as a generalization of classical reversible gates.

The most common gates are described in Table 1. The X gate is the quantum version of the NOT gate. The CNOT or “controlled NOT” negates a target bit if and only if the control bit is 1. We will use the notation CNOT_{ij} for a CNOT gate controlled by qubit i acting on qubit j . The CNOT gate can be expressed using the outer product notation as,

$$\text{CNOT} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes X = |00\rangle\langle 00| + |01\rangle\langle 01| + |10\rangle\langle 11| + |11\rangle\langle 10|. \quad (14)$$

The Toffoli gate or “controlled-controlled NOT” or CCNOT, is a three qubit gate that is essentially the quantum (reversible) version of the AND gate. It negates a target bit if and only if both control bits are 1. In the outer product notation,

$$\text{CCNOT} = |11\rangle\langle 11| \otimes X + (I - |11\rangle\langle 11|) \otimes I. \quad (15)$$

Another way to look at the CCNOT gate is as a CNOT gate with an additional control qubit,

$$\text{CCNOT} = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes \text{CNOT}. \quad (16)$$

In general, one can define controlled versions of any unitary gate U as,

$$CU = |0\rangle\langle 0| \otimes I + |1\rangle\langle 1| \otimes U. \quad (17)$$

CU applies U to a set of qubits only if the first qubit (called the control qubit) is $|1\rangle$.

A set of gates that together can execute all possible quantum computations is called a *universal gate set*. Taken together, the set of all unary (i.e., acting on one qubit) gates and the binary (i.e., acting on two qubits) CNOT gate form a universal gate set. More economically, the set $\{H, T, \text{CNOT}\}$ (Refer Table 1 for definitions of these gates) forms a universal set. Also, the Toffoli gate by itself is universal [92].

1.1.7 Observables and expectation values. We have seen that experiments in quantum mechanics are probabilistic. Often in experiments we will need to associate a real number with a measurement outcome. And quantities that we measure in quantum mechanics will always be statistical averages of these numbers. For instance, suppose we do the following experiment on many copies of the single qubit state in Eq. (1): We measure a copy of the state and if we get $|0\rangle$ we record 1 in our lab notebook, otherwise we record -1 . While doing this experiment we can never predict the outcome of a specific measurement. But we can ask statistical questions like: “What will be the average value of the numbers in the notebook?” From our earlier discussion on measurement we know that the probability of measuring $|0\rangle$ is $|\alpha|^2$ and the probability of measuring $|1\rangle$ is $|\beta|^2$. So the average value of the numbers in the notebook will be,

$$|\alpha|^2 - |\beta|^2 \quad (18)$$

In quantum formalism, there is neat way to express such experiments and their average outcomes, without all the verbiage, using certain operators. For the experiment described above the associated operator would be the Z gate,

$$Z = |0\rangle\langle 0| - |1\rangle\langle 1| = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \quad (19)$$

By associating this operator with the experiment we can write the average outcome of the experiment, on $|\phi\rangle$, as the overlap between $|\phi\rangle$ and $Z|\phi\rangle$,

$$\langle \phi | Z | \phi \rangle = \langle \phi | 0 \rangle \langle 0 | \phi \rangle - \langle \phi | 1 \rangle \langle 1 | \phi \rangle = |\alpha|^2 - |\beta|^2. \quad (20)$$

One-qubit gates	Multi-qubit gates
Hadamard = $H = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$	CNOT = $CX = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}$
$I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}$, $S = \begin{pmatrix} 1 & 0 \\ 0 & i \end{pmatrix}$	$CZ = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}$
$T = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\pi/4} \end{pmatrix}$	Controlled- $U = CU = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & U_{00} & U_{01} \\ 0 & 0 & U_{10} & U_{11} \end{pmatrix}$
NOT = $X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$	SWAP = $\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$
$Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}$, $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$	Toffoli (CCNOT) = $\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}$
$R(\theta) = P(\theta) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\theta} \end{pmatrix}$	

Table 1. Commonly used quantum gates.

The operator Z is called the *observable* associated with this experiment. And the quantity $\langle \phi | Z | \phi \rangle$ is called its *expectation value*. The expectation value is sometimes denoted by $\langle Z \rangle$, when there is no ambiguity about the state on which the experiment is performed.

Here we saw an experiment done in the computational basis. But this need not always be the case. Experiments can be designed by associating real numbers to measurement outcomes in any basis. What would be the observable for such an experiment? For an experiment that associates the real numbers $\{a_i\}$ to a measurement onto a basis set $\{|\Phi_i\rangle\}$, the observable will be,

$$O \equiv \sum_i a_i |\Phi_i\rangle \langle \Phi_i|. \quad (21)$$

This observable will reproduce the correct expectation value for this experiment done on any state $|\psi\rangle$,

$$\langle\psi|O|\psi\rangle = \sum_i a_i \langle\psi|\Phi_i\rangle \langle\Phi_i|\psi\rangle = \sum_i a_i |\langle\Phi_i|\psi\rangle|^2. \quad (22)$$

Because the states $\{|\Phi_i\rangle\}$ are orthonormal, we can see that O obeys the following eigenvalue equation,

$$O|\Phi_j\rangle = \sum_i a_i |\Phi_i\rangle \langle\Phi_i|\Phi_j\rangle = a_j |\Phi_j\rangle. \quad (23)$$

So O is an operator that has complete set of orthogonal eigenvectors and real eigenvalues. Such operators are called *Hermitian operators*. Equivalently, these operators are equal to their Hermitian conjugates ($O = O^\dagger$). In quantum mechanics, any Hermitian operator is a valid observable. The eigenvectors of the operator give the possible outcomes of the experiment and the corresponding eigenvalues are the real numbers associated with that outcome.

But can all valid observables be measured in practice? The answer to this depends on the quantum system under consideration. In this tutorial, the system under consideration is an IBM quantum processor. And in these processors only measurements onto the computational basis are supported natively. Measurements to other basis states can be performed by applying an appropriate unitary transformation before measurement. Suppose that the hardware only lets us do measurements onto the computational basis $\{|i\rangle\}$ but we want to perform a measurement onto the basis set $\{|\Phi_i\rangle\}$. This problem can be solved if we can implement the following unitary transformation,

$$U = \sum_i |i\rangle \langle\Phi_i|. \quad (24)$$

Now measuring $U|\psi\rangle$ in the computational basis is the same as measuring $|\psi\rangle$ in the $\{|\Phi_i\rangle\}$ basis. This can be seen by computing the outcome probabilities on $U|\psi\rangle$,

$$|\langle j|U|\psi\rangle|^2 = \left| \sum_i \langle j|i\rangle \langle\Phi_i|\psi\rangle \right|^2 = |\langle\Phi_j|\psi\rangle|^2. \quad (25)$$

So once U is applied, the outcome $|j\rangle$ becomes equivalent to the outcome $|\Phi_j\rangle$ in the original measurement scenario. Now, not all such unitary transformations are easy to implement. So if a quantum algorithm requires us to perform a measurement onto some complicated set of basis states, then the cost of implementing the corresponding U has to be taken into account.

1.1.8 Quantum circuits. Quantum algorithms are often diagrammatically represented as circuits in literature. Here we will describe how to construct and read quantum circuits. In the circuit representation, qubits are represented by horizontal lines. Gates are then drawn on the qubits they act on. This is done in sequence from left to right. The initial state of the qubit is denoted at the beginning of each of the qubit lines. Notice that when we write down a mathematical expression for the circuit, the gates are written down from right to left in the order of their application.

These principles are best illustrated by an example. Given in Fig. 1 is a circuit to preparing an entangled two qubit state called a Bell state from $|00\rangle$.




Fig. 1. Quantum circuit for preparing a Bell state

The circuit encodes the equation,

$$\text{CNOT}_{12} (H \otimes I) |00\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

Let us now carefully go over how the circuit produces the Bell state. We read the circuit from left to right. The qubits are numerically labelled starting from the top. First the H gate acts on the top most qubit changing the state of the system to,

$$H \otimes I |00\rangle = (H |0\rangle) \otimes (I |0\rangle) = \left(\frac{|0\rangle + |1\rangle}{\sqrt{2}} \right) \otimes |0\rangle = \frac{1}{\sqrt{2}}(|00\rangle + |10\rangle).$$

Then CNOT_{12} acts on both of these qubits. The blackened dot on the first qubit implies that this qubit is the control qubit for the CNOT. The \oplus symbol on the second qubit implies that this qubit is the target of the NOT gate (controlled by the state of the first qubit). The action of the CNOT then gives,

$$\text{CNOT}_{12} \left(\frac{1}{\sqrt{2}}(|00\rangle + |10\rangle) \right) = \frac{1}{\sqrt{2}}(\text{CNOT}_{12} |00\rangle + \text{CNOT}_{12} |10\rangle) = \frac{1}{\sqrt{2}}(|00\rangle + |11\rangle).$$

The measurement of a qubit is also denoted by a special gate with a meter symbol on it, given in Fig 2. The presence of this gate on a qubit means that the qubit must be measured in the computational basis.




Fig. 2. The measurement gate

1.1.9 Quantum algorithms. We have now introduced all the basic elements needed for the discussion of practical quantum algorithms. A quantum algorithm consists of three basic steps:

- Encoding of the data, which could be classical or quantum, into the state of a set of input qubits.
- A sequence of quantum gates applied to this set of input qubits.
- Measurements of one or more of the qubits at the end to obtain a classically interpretable result.

In this review, we will describe the implementation of these three steps for a variety of quantum algorithms.

1.2 Implementations on a real quantum computer

1.2.1 The IBM quantum computer. In this article, we consider IBM's publicly available quantum computers. In most cases, we specifically consider the `ibmqx4`, which is a 5-qubit computer, although in some cases we also consider other quantum processors freely accessible through the IBM Quantum Experience platform. These processors can be accessed by visiting the IBM Quantum Experience website (<https://quantum-computing.ibm.com/>)

There are several issues to consider when implementing an algorithm on real quantum computers, for example:

- (1) What is the available gate set with which the user can state their algorithm?
- (2) What physical gates are actually implemented?
- (3) What is the qubit connectivity (i.e., which pairs of qubits can two-qubit gates be applied to)?
- (4) What are the sources of noise (i.e., errors)?

We first discuss the available gate set. In IBM's graphical interface to the `ibmqx4`, the available gates include:

$$\{I, X, Y, Z, H, S, S^\dagger, T, T^\dagger, U_1(\lambda), U_2(\lambda, \phi), U_3(\lambda, \phi, \theta), \text{CNOT}\}. \quad (26)$$




Fig. 3. The connectivity diagram of ibmqx4. The circles represent qubits and the arrows represent the ability to apply a physical CNOT gate between the qubits.

The Graphical User Interface (GUI) also provides other controlled gates and operations like measurement and reset. Most of these gates appear in our Table 1. The gates $U_1(\lambda)$, $U_2(\lambda, \phi)$, and $U_3(\lambda, \phi, \theta)$ are continuously parameterized gates, defined as follows:

$$U_1(\lambda) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\lambda} \end{pmatrix}, \quad U_2(\lambda, \phi) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -e^{i\lambda} \\ e^{i\phi} & e^{i(\lambda+\phi)} \end{pmatrix}, \quad U_3(\lambda, \phi, \theta) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\lambda+\phi)} \cos(\theta/2) \end{pmatrix}. \quad (27)$$

Note that $U_3(\lambda, \phi, \theta)$ is essentially an arbitrary one-qubit gate.

The gates listed in Eq. (26) are provided by IBM for the user's convenience. However these are not the gates that are physically implemented by their quantum computer. IBM has a compiler that translates the gates in (26) into products of gates from a physical gate set. The physical gate set employed by IBM is essentially composed of three gates [1]:

$$\{U_1(\lambda), R_X(\pi/2), \text{CNOT}\}. \quad (28)$$

Here, $R_X(\pi/2)$ is a rotation by angle $\pi/2$ of the qubit about its X -axis, corresponding to a matrix similar to the Hadamard:

$$R_X(\pi/2) = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ -i & 1 \end{pmatrix}. \quad (29)$$

The reason why it could be important to know the physical gate set is that some user-programmed gates may need to be decomposed into multiple physical gates, and hence could lead to a longer physical algorithm. For example, the X gate gets decomposed into three gates: two $R_X(\pi/2)$ gates sandwiching one $U_1(\lambda)$ gate.

The connectivity of the computer is another important issue. Textbook algorithms are typically written for a fully-connected hardware, which means that one can apply a two-qubit gate to any two qubits. In practice, real quantum computers may not have full connectivity. In the ibmqx4, which has 5 qubits, there are 6 connections, i.e., there are only 6 pairs of qubits to which a CNOT gate can be applied (Fig.3). In contrast a fully connected 5-qubit system would allow a CNOT to be applied to 20 different qubit pairs. In this sense, there are 14 "missing connections". Fortunately, there are ways to effectively generate connections through clever gate sequences. For example, a CNOT gate with qubit j as the control and qubit k as the target can be reversed (such that j is the target and k is the control) by applying Hadamard gates on each qubit both before and after the CNOT, i.e.,

$$\text{CNOT}_{kj} = (H \otimes H) \text{CNOT}_{jk} (H \otimes H). \quad (30)$$

Similarly, there exists a gate sequence to make a CNOT between qubits j and l if one has connections between j and k , and k and l , as follows:

$$\text{CNOT}_{jl} = \text{CNOT}_{kl} \text{CNOT}_{jk} \text{CNOT}_{kl} \text{CNOT}_{jk}. \quad (31)$$




Fig. 4. The quantum circuit to prepare a Bell state and measure it in the IBM quantum experience GUI

Hence, using (30) and (31), one can make up for lack of connectivity at the expense of using extra gates.

Finally, when implementing a quantum algorithm it is important to consider the sources of noise in the computer. The two main sources of noise are typically gate *infidelity* and *decoherence*. Gate infidelity refers to the fact that the user-specified gates do not precisely correspond to the physically implemented gates. Gate infidelity is usually worse for multi-qubit gates than for one-qubit gates, so typically one wants to minimize the number of multi-qubit gates in one’s algorithm. Decoherence refers to the fact that gradually over time the quantum computer loses its “quantumness” and behaves more like a classical object. After decoherence has fully occurred, the computer can no longer take advantage of quantum effects. This introduces progressively more noise as the quantum algorithm proceeds in time. Ultimately this limits the depth of quantum algorithms that can be implemented on quantum computers. It is worth noting that different qubits decohere at different rates, and one can use this information to better design one’s algorithm. The error rates for individual qubits in the IBM processors are listed in the IBM Quantum Experience website. In this tutorial, we will show in many cases how infidelity and decoherence affect the algorithm performance in practice.

A simple example of programming the IBM quantum computer is given in Fig. 4, which shows the Bell state preparation circuit Fig.1 compiled using the IBM quantum experience GUI. Extra measurement operations at the end serve to verify the fidelity of the implementation.

1.2.2 Programming the IBM quantum computer: Qiskit library. Qiskit [4] is an open-source quantum computing library developed under the aegis of IBM. Qiskit allows users to write and run programs on either IBM’s quantum processors or on a local simulator, without the use of the graphical interface. This is an important feature because the graphical interface becomes impractical as the number qubits become large. At the time of writing, users can use Qiskit to access quantum processors with up to 16 qubits. Smaller processors are also accessible. Qiskit is a very powerful software development kit (SDK) which has multiple elements in it that tackle a variety of problems associated with practical quantum computing. Qiskit is further split into four modules called: Terra, Aer, Aqua, and Ignis. Each of these modules deal with a specific part of quantum software development. In this section we will only give a brief overview of programming simple quantum circuits with Qiskit. For a comprehensive overview of Qiskit and its various capabilities, the reader is encouraged to visit the official website (www.qiskit.org) [4].

For our purposes, Qiskit can be viewed as a Python library for quantum circuit execution. A basic Qiskit code has two parts, designing the circuit and running it. In the circuit design phase, we create an instance of `QuantumCircuit` with the required number of qubits and classical bits. Then gates and measurements are added to this blank circuit. Gates and measurements are implemented in Qiskit as methods of the `QuantumCircuit` class. After the circuit has been designed we need to choose a backend to run the circuit. This can be either be a simulator called the `qasm_simulator` or it can be one of IBM’s quantum processors. To use a quantum processor, you will need to load your IBM Q account information into Qiskit. Given in Fig. 5 is a simple code to construct

the Bell state. This is the Qiskit version of the circuit in Fig. 1 with measurement added at the end to verify our results.

```

#### Quantum circuit for preparing the Bell state ####

import numpy as np
from qiskit import QuantumCircuit, execute, Aer

# Create a Quantum Circuit with two qbits and 2 classical bits
circuit = QuantumCircuit(2,2)

# Add a H gate on qubit 0
circuit.h(0)

# Add a CX (CNOT) gate on control qubit 0 and target qubit 1
circuit.cx(0,1)

# Map the quantum measurement to the classical bits
circuit.measure([0,1],[0,1])

# Use Aer's qasm_simulator
simulator = Aer.get_backend('qasm_simulator')

# Execute the circuit on the qasm simulator
job = execute(circuit, simulator, shots=1000)

# Grab results from the job
result = job.result()

# Returns counts
counts = result.get_counts(circuit)
print("\nTotal count for 00 and 11 are:",counts)

```

Fig. 5. Qiskit code to create and measure a Bell state. Source: www.qiskit.org

In Fig.5 we are running the circuit on the simulator for 1000 independent runs. The final output was {'11': 493, '00': 507}. This is what we expect from measuring the Bell state ($\frac{|00\rangle+|11\rangle}{\sqrt{2}}$), up to statistical fluctuations. While running the same code on the 14 qubit `ibmq_16_melbourne` processor for 1024 runs gave |11> with probability 0.358 and |00> with probability 0.54. The remaining probability was distributed over 01 and 10, which should not be a part of the Bell state. As we discussed before, this phenomenon is due to errors inherent to the quantum processor. As the backend technology improves we expect to get better results from these trials. Often, we will also present a circuit using OpenQASM (Open Quantum Assembly Language). OpenQASM provides an intermediate representation of a program in the form of a quantum circuit, that is neither the actual program written by the programmer nor the machine instructions seen by the processor. OpenQASM 'scores' we

Class	Problem/Algorithm	Paradigms used	Hardware	Simulation Match
Inverse Function Computation	Grover's Algorithm	GO	QX4	med
	Bernstein-Vazirani	n.a.	QX4, QX5	high
Number-theoretic Applications	Shor's Factoring Algorithm	QFT	QX4	med
Algebraic Applications	Linear Systems	HHL	QX4	low
	Matrix Element Group Representations	QFT	ESSEX	low
	Matrix Product Verification	GO	n.a.	n.a.
	Subgroup Isomorphism	QFT	none	n.a.
Graph Applications	Quantum Random Walk	n.a.	VIGO	med-low
	Minimum Spanning Tree	GO	QX4	med-low
	Maximum Flow	GO	QX4	med-low
	Approximate Quantum Algorithms	SIM	QX4	high
Learning Applications	Quantum Principal Component Analysis (PCA)	QFT	QX4	med
	Quantum Support Vector Machines (SVM)	QFT	none	n.a.
	Partition Function	QFT	QX4	med-low
Quantum Simulation	Schrödinger Equation Simulation	SIM	QX4	low
	Transverse Ising Model Simulation	VQE	none	n.a.
Quantum Utilities	State Preparation	n.a.	QX4	med
	Quantum Tomography	n.a.	QX4	med
	Quantum Error Correction	n.a.	QX4	med

Table 2. Overview of studied quantum algorithms. Paradigms include Grover Operator (GO), Quantum Fourier Transform (QFT), Harrow-Hassidim-Lloyd (HHL), Variational Quantum Eigenvalue solver (VQE), and direct Hamiltonian simulation (SIM). The simulation match column indicates how well the hardware quantum results matched the simulator results

show in this paper will be simple sequence of gates and measurements, with the corresponding registers that they act on. The syntax of these scores will be self explanatory.

1.3 Classes of quantum algorithms

In this review, we broadly classify quantum algorithms according to their area of application. We will discuss quantum algorithms for graph theory, number theory, machine learning and so on. The complete list of algorithms discussed in this paper, classified according to their application areas, can be found in Table 2. The reader is also encouraged to take a look at the excellent Quantum Algorithm Zoo website [72] for a concise and comprehensive list of quantum algorithms.

In classical computing, algorithms are often designed by making use of one or more algorithmic paradigms like dynamic programming or local search, to name a few. Most known quantum algorithms also use a combination of algorithmic paradigms specific to quantum computing. These paradigms are the Quantum Fourier Transform (QFT), the Grover Operator (GO), the Harrow-Hassidim-Lloyd (HHL) method for linear systems, variational quantum eigenvalue solver (VQE), and direct Hamiltonian simulation (SIM). The number of known quantum algorithmic paradigms is much smaller compared to the number of known classical paradigms. The constraint of unitarity on quantum operations and the impossibility of non-intrusive measurement make it difficult to design quantum paradigms from existing classical paradigms. But researchers are constantly in search for new paradigms and we can expect this list to get longer in the future. Table 2 also contains information about the paradigms used by the algorithms in this article.

The rest of the paper presents each of the algorithms shown in Table 2, one after the other. In each case, we first discuss the goal of the algorithm (the problem it attempts to solve). Then we describe the gate sequence

required to implement this algorithm. Finally, we show the results from implementing this algorithm on IBM’s quantum computer¹.

The list of algorithms in Table 2 is by no means exhaustive. These algorithms have been chosen due to their relative importance and to provide an overview of the field. Many interesting quantum algorithms like those for topological data analysis [83], spatial search [123], supervised learning [108], etc., have not been covered in this review. Nevertheless the tools and ideas elucidated in this paper will help the reader understand and implement many quantum algorithms that are not included here.

2 GROVER’S ALGORITHM

2.1 Problem definition and background

Grover’s algorithm as initially described [63] enables one to find (with probability $> 1/2$) a specific item within a randomly ordered database of N items using $O(\sqrt{N})$ operations. By contrast, a classical computer would require $O(N)$ operations to achieve this. Therefore, Grover’s algorithm provides a quadratic speedup over an optimal classical algorithm. It has also been shown [15] that Grover’s algorithm is optimal in the sense that no quantum Turing machine can do this in less than $O(\sqrt{N})$ operations.

While Grover’s algorithm is commonly thought of as being useful for searching a database, the basic ideas that comprise this algorithm are applicable in a much broader context. This approach can be used to accelerate search algorithms where one could construct a “quantum oracle” that distinguishes the needle from the haystack. The needle and hay need not be part of a database. For example, it could be used to search for two integers $1 < a < b$ such that $ab = n$ for some number n , resulting in a factoring algorithm. Grover’s search in this case would have worse performance than Shor’s algorithm [113, 114] described below, which is a specialised algorithm to solve the factoring problem. Implementing the quantum oracle can be reduced to constructing a quantum circuit that flips an ancillary qubit, q , if a function, $f(\mathbf{x})$, evaluates to 1 for an input \mathbf{x} . We use the term *ancilla* or *ancillary qubit* to refer to some extra qubits that are used by the algorithm.

The function $f(\mathbf{x})$ is defined by

$$f(\mathbf{x}) = \begin{cases} 1 & \text{if } \mathbf{x} = \mathbf{x}^* \\ 0 & \text{if } \mathbf{x} \neq \mathbf{x}^* \end{cases} \quad (32)$$

where $\mathbf{x} = x_1x_2 \dots x_n$ are binary strings and \mathbf{x}^* is the specific string that is being sought. It may seem paradoxical at first that an algorithm for finding \mathbf{x}^* is needed if such a function can be constructed. The key here is that $f(\mathbf{x})$ need only recognize \mathbf{x}^* – it is similar to the difference between writing down an equation and solving an equation. For example, it is easy to check if the product of a and b is equal to n , but harder to factor n . In essence, Grover’s algorithm can invert an arbitrary function with binary outputs, provided we have a quantum oracle that implements the function. Grover’s algorithm has been used, with appropriate oracles, to solve problems like finding triangles in a graph [87], finding cycles [32], and finding maximal cliques [131]. For the analysis of Grover’s algorithm, the internals of the oracle is typically considered a black-box. Often, the oracle operator for the problem at hand has to be constructed as a quantum circuit. But, keep in mind that an inefficient oracle construction can nullify any practical advantages gained by using Grover’s search.

Here we implement a simple instance of Grover’s algorithm. That is, the quantum oracle we utilize is a very simple one. Let $\mathbf{x} = x_1x_2$ and we wish to find \mathbf{x}^* such that $x_1^* = 1$ and $x_2^* = 1$. While finding such an \mathbf{x}^* is trivial, we don a veil of ignorance and proceed as if it were not. This essentially means that our function $f(\mathbf{x})$ is an AND gate. But AND gate is not reversible and cannot be a quantum gate. However the Toffoli gate, that was introduced in the previous section, is a reversible version of the classical AND gate. The Toffoli gate takes three bits as input and outputs three bits. The first two bits are unmodified. The third bit is flipped if the first two bits

¹The code and implementations for most of the algorithms can be found at https://github.com/lanl/quantum_algorithms.

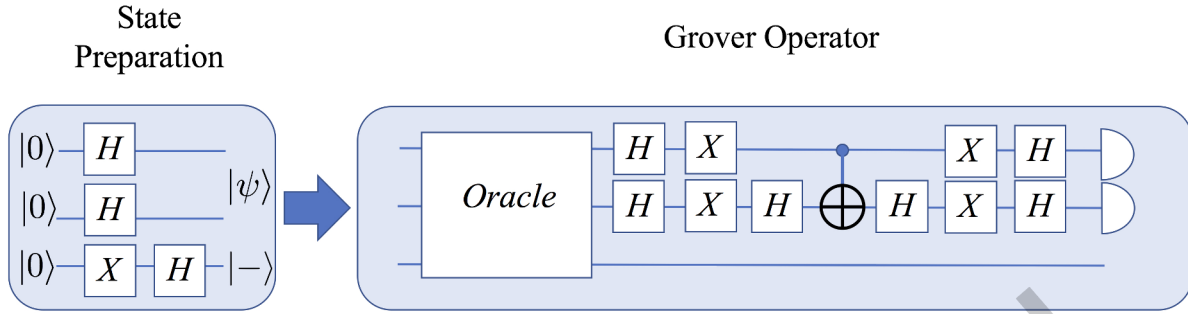


Fig. 6. A schematic diagram of Grover's algorithm is shown. Note that in this case, one application of the Grover operator is performed. This is all that is necessary when there are only two bits in \mathbf{x} , but the Grover operator should be applied repeatedly for larger problems.

are 1. The unitary matrix corresponding to the Toffoli gate can be found in Table 1. In other words, the Toffoli gate implements our desired quantum oracle where the first two inputs are x_1 and x_2 and the third bit is the ancillary bit, q . The behavior of the oracle in general is $|\mathbf{x}\rangle|q\rangle \rightarrow |\mathbf{x}\rangle|f(\mathbf{x}) \oplus q\rangle$, where \oplus is the XOR operation. Here we will only discuss the case where \mathbf{x}^* is unique. Grover's algorithm can also be used to search for multiple items in a database.

2.2 Algorithm description

Here we present a brief introduction to Grover's algorithm. A more detailed account can be found in Nielsen and Chuang [92]. Let N be the number of items (represented as bit strings) amongst which we are performing the search. This number will also be equal to the dimension of the vector space we are working with. An operator, called the Grover operator or the diffusion operator, is the key piece of machinery in Grover's algorithm. This operator is defined by

$$G = (2|\psi\rangle\langle\psi| - I)O \quad (33)$$

where $|\psi\rangle = \frac{1}{\sqrt{N}} \sum_i |i\rangle$ is the uniform superposition over all the basis states and O is the oracle operator (see Fig. 6 for a representation of this operator in the case where \mathbf{x} consists of 2 bits). The action of $(2|\psi\rangle\langle\psi| - I)$ on an arbitrary state, given by $\sum_i a_i |i\rangle$, when decomposed over the basis states is,

$$(2|\psi\rangle\langle\psi| - I) \sum_i a_i |i\rangle = \sum_i (2\langle a\rangle - a_i) |i\rangle \quad (34)$$

where $\langle a\rangle = \frac{\sum_i a_i}{N}$ is the average amplitude in the basis states. From Eq. (34) one can see that the amplitude of each $|i\rangle$ -state (a_i) is flipped about the mean amplitude.

In order to use the Grover operator to successfully perform a search, the qubit register must be appropriately initialized. The initialization is carried out by applying a Hadamard transform to each of the the main qubits ($H^{\otimes n}$) and applying a Pauli X transform followed by a Hadamard transform (HX) to the ancilla. This leaves the main register in the uniform superposition of all states, $|\psi\rangle$, and the ancilla in the state $\frac{|0\rangle - |1\rangle}{\sqrt{2}}$. After performing these operations, the system is in the state $|\psi\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}}$. Using Eq. (34), we can now understand how the Grover

operator works. The action of the oracle operator on $|x^*\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}}$ reverses the amplitude of that state

$$O|x^*\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} \rightarrow |x^*\rangle \frac{|f(x^*) \oplus 0\rangle - |f(x^*) \oplus 1\rangle}{\sqrt{2}} = |x^*\rangle \frac{|1\rangle - |0\rangle}{\sqrt{2}} = -|x^*\rangle \frac{|0\rangle - |1\rangle}{\sqrt{2}} \quad (35)$$

A similar argument shows that all other states are unmodified by the oracle operator. Combining this with Eq. (34) reveals why the Grover operator is able to successfully perform a search. Consider what happens on the first iteration: The oracle operator makes it so that the amplitude of $|x^*\rangle$ is below $\langle a \rangle$ (using the notation of Eq. (34)) while all the other states have an amplitude that is slightly above $\langle a \rangle$. The effect of applying $2|\psi\rangle\langle\psi| - I$ is then to make $|x^*\rangle$ have an amplitude above the mean while all other states have an amplitude below the mean. The desired behavior of the Grover operator is to increase the amplitude of $|x^*\rangle$ while decreasing the amplitude of the other states. If the Grover operator is applied too many times, this will eventually stop happening. The Grover operator should be applied exactly $\left\lceil \frac{\pi\sqrt{N}}{4} \right\rceil$ times after which a measurement will reveal x^* with probability close to 1. In the case where x has two bits, a single application of Grover's operator is sufficient to find x^* with certainty (in theory). Below is a high level pseudocode for the algorithm.

Algorithm 1 Grover's algorithm

Input:

- An Oracle operator effecting the transformation $|x\rangle|q\rangle \rightarrow |x\rangle|q \oplus f(x)\rangle$.

Output:

- The unique bit string x^* satisfying Eq. (32)

Procedure:

Step 1. Perform state initialization $|0 \dots 0\rangle \rightarrow |\psi\rangle \left(\frac{|0\rangle - |1\rangle}{\sqrt{2}} \right)$

Step 2. Apply Grover operator $\left\lceil \frac{\pi\sqrt{N}}{4} \right\rceil$ times

Step 3. Perform measurement on all qubit except the ancillary qubit.

2.3 Algorithm implemented on IBM's 5-qubit computer

Fig. 7 shows the circuit that was designed to fit the `ibmqx4` quantum computer. The Toffoli gate is not available directly in `ibmqx4` so it has to be constructed from the available set of gates given in Eq. 26.

The circuit consists of state preparation (first two time slots), a Toffoli gate (the next 13 time slots), followed by the $2|\psi\rangle\langle\psi| - I$ operator (7 time slots), and measurement (the final 2 time slots). We use $q[0]$ (in the register notation from Fig. 7) as the ancillary qubit, and $q[1]$ and $q[2]$ as x_1 and x_2 respectively. Note that the quantum computer imposes constraints on the possible source and target of CNOT gates.

Using the simulator, this circuit produces the correct answer $x = (1, 1)$ every time. We executed 1,024 shots using the `ibmqx4` and $x = (1, 1)$ was obtained 662 times with $(0, 0)$, $(0, 1)$, and $(1, 0)$ occurring 119, 101, and 142 times respectively. This indicates that the probability of obtaining the correct answer is approximately 65%. The deviation between the simulator and the quantum computer is due to the inherent errors in `ibmqx4`. This deviation will get worse for circuits of larger size.

We also ran another test using CNOT gates that did not respect the underlying connectivity of the computer. This resulted in a significantly deeper circuit and the results were inferior to the results with the circuit in Fig. 7.

This implementation used a Toffoli gate with a depth of 23 (compared to a depth of 13 here) and obtained the correct answer 48% of the time.




Fig. 7. The circuit that was executed on IBM’s 5-qubit quantum computer. The first two time slots correspond to the state preparation. The next 13 time slots implement a Toffoli gate. The next 7 time slots implement the $2|\psi\rangle\langle\psi| - I$ operator, and the final two time slots are used for observing x_1 and x_2 .

3 BERNSTEIN-VAZIRANI ALGORITHM

3.1 Problem definition and background

Suppose we are given a classical Boolean function, $f : \{0, 1\}^n \mapsto \{0, 1\}$. It is guaranteed that this function can always be represented in the form, $f_s(\mathbf{x}) = \bigoplus_i s_i x_i \equiv \langle \mathbf{s}, \mathbf{x} \rangle$. Here, \mathbf{s} is an unknown bit string, which we shall call a *hidden string*. Just like in Grover’s algorithm we assume that we have a quantum oracle that can compute this function.

The Bernstein-Vazirani (BV) algorithm then finds the hidden string with just a single application of the oracle. The number of times the oracle is applied during an algorithm algorithm is known as its *query complexity*. The BV algorithm has a query complexity of one. From our earlier discussions we saw that Grover’s algorithm has a query complexity of $O(\sqrt{N})$.

In the classical case each call to $f_s(\mathbf{x})$ produces just 1 bit of information, and since an arbitrary hidden string \mathbf{s} has n -bits of information, the classical query complexity is seen to be n . Even with bounded error, there is no way that this classical complexity can be brought down, as can be seen using slightly more rigorous information-theoretic arguments.

The quantum algorithm to solve this problem was developed by Bernstein and Vazirani [16] building upon the earlier work of Deutsch and Jozsa [41]. Their contribution was a quantum algorithm for the hidden string problem, which has a non-recursive quantum query complexity of just 1. This constitutes a polynomial $O(n)$ query-complexity separation between classical and quantum computation. They also discovered a less widely known recursive hidden-string query algorithm, which shows a $O(n^{\log n})$ separation between classical and quantum query-complexities. These developments preceded the more famous results of Shor and Grover, and kindled a lot of early academic interest in the inherent computational power of quantum computers.

One thing to note about the BV algorithm is that the *black-box* function $f_s(\cdot)$ can be very complex to implement using reversible quantum gates. For an n -bit hidden string, the number of simple gates needed to implement $f_s(\cdot)$ scales typically as $O(4^n)$ [92]. Since the black box is a step in the algorithm, its serial execution time could in the worst-case even scale exponentially. The real breakthrough of this quantum algorithm lies in speeding up the query complexity and not the execution time per se.

3.2 Algorithm description

Let us explore the BV algorithm in more detail. Let U_s be the oracle for the function $f_s(\mathbf{x})$. It acts in the usual way and computes the value of the function onto an ancilla qubit,

$$U_s |\mathbf{x}\rangle |q\rangle = |\mathbf{x}\rangle |q \oplus \langle \mathbf{s}, \mathbf{x} \rangle\rangle \tag{36}$$

Denoting $|-\rangle = (|0\rangle - |1\rangle)/\sqrt{2}$, we can easily verify from Eq. (35) that,

$$U_s |x\rangle |-\rangle = (-1)^{\langle s, x \rangle} |x\rangle |-\rangle. \quad (37)$$

Also, note that the n -qubit Hadamard operator, which is just n single qubit H operators applied in parallel, can be expanded as,

$$H^{\otimes n} = \frac{1}{\sqrt{2^n}} \sum_{x, y \in \{0, 1\}^n} (-1)^{\langle x, y \rangle} |y\rangle \langle x| \quad (38)$$

The reader may verify this identity by applying $H^{\otimes n}$ to the computational basis states.

U_s and $H^{\otimes n}$ are the only two operators needed for the BV algorithm. The pseudocode for the algorithm is given in Algorithm 2. Notice that the initialization part is identical to that of Grover's algorithm. This kind of initialization is a very common strategy in quantum algorithms.

Algorithm 2 Bernstein-Vazirani algorithm

Input:

- An oracle operator, U_s , effecting the transformation $|x\rangle |q\rangle \rightarrow |x\rangle |q \oplus \langle s, x \rangle\rangle$.

Output:

- The hidden string s .

Procedure:

Step 1. Perform state initialization on $n + 1$ qubits, $|0 \dots 0\rangle \rightarrow |\psi\rangle |-\rangle$

Step 2. Apply U_s .

Step 3. Apply $H^{\otimes n}$ to the first n qubits.

Step 4. Measure all qubits except the ancillary qubit.

The final measurement will reveal the hidden string, s , with probability 1. Let us now delve into the algorithm to see how this result is achieved. The entire circuit for the BV algorithm is represented in Figure 8. This circuit can be analyzed as follows,

$$\begin{aligned} |0\rangle^n |1\rangle &\xrightarrow{H^{\otimes(n)} \otimes H} \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} |x\rangle \otimes |-\rangle \xrightarrow{U_s} \frac{1}{\sqrt{2^n}} \sum_{x=0}^{2^n-1} (-1)^{\langle s, x \rangle} |x\rangle \otimes |-\rangle \\ &\xrightarrow{H^{\otimes n}} \frac{1}{\sqrt{2^n}} \sum_{x, y=0}^{2^n-1} (-1)^{\langle s, x \rangle \oplus \langle x, y \rangle} |y\rangle \otimes |-\rangle \equiv |s\rangle \otimes |-\rangle. \end{aligned} \quad (39)$$

Here we have crucially used the identity for $H^{\otimes n}$ given in Eq.(38).

3.3 Algorithm implemented on IBM's 5-qubit and 16-qubit computers

From the BV algorithm description in the previous section, we see that in any practical implementation of this algorithm, the main ingredient is the construction of the oracle U_s given a binary hidden string s . Let us see how this is done using an example binary hidden string "01". Equation (40) below shows how the 3-qubit operator maps the $2^3 = 8$ basis vectors onto themselves. The first line is the input binary vector (in the order x_1, x_0, q), and the second line is the output binary vector.

$$U_{01} = \begin{pmatrix} 000 & 010 & 100 & 110 & 001 & 011 & 101 & 111 \\ 000 & 011 & 100 & 111 & 001 & 010 & 101 & 110 \end{pmatrix} \quad (40)$$

This mapping, $U_{01} : |x\rangle |q\rangle \mapsto |x\rangle |\langle 01, x \rangle \oplus q\rangle$, is unitary. The next task in implementation is to lower the unitary matrix operator U_{01} to primitive gates available in the quantum computer's architecture given in Eq (26). The

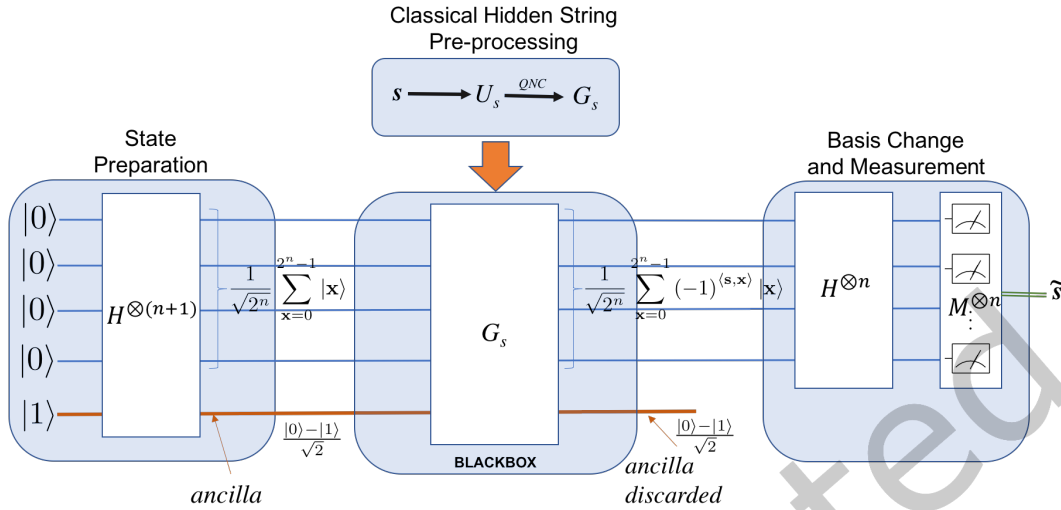


Fig. 8. Bernstein-Vazirani hidden string discovery quantum algorithm. The hidden string s is discovered with just a single query. The measurement result \tilde{s} gives the hidden string.

time cost of applying these gates can be accessed from IBM’s published calibration models [125] for the primitive hardware gates.

In order to decompose arbitrary unitary matrices to the primitive gates, we need to first perform a unitary diagonalization of the $2^{(n+1)} \times 2^{(n+1)}$ matrix using multi-qubit-controlled single-qubit unitary Given’s rotation operations. Such multi-qubit-controlled single-qubit operations can be decomposed further to primitive gates using standard techniques [92] to the hardware primitive gates. Even after this step we will be left with arbitrary CNOT gates that do not respect the topology of the underlying quantum processor. Since both `ibmqx4`, `ibmqx5` computers have restricted CNOT connectivity between qubits, we will need to decompose the CNOT gates further into available CNOT gates using the method discussed in the introductory section. As we saw in the Grover’s algorithm section, such decompositions will further degrade the quality of our results. As the overall primitive gate counts scale as $\mathcal{O}(4^n)$ for arbitrary n -qubit unitary operators, these decompositions quickly becomes hard to do by hand. To address this we wrote a piece of software called *Quantum Netlist Compiler (QNC)* [107] for performing these transformations automatically. QNC can do much more than convert arbitrary unitary operators to OpenQASM-2.0 circuits—it has specialized routines implemented to generate circuits to do state-preparations, permutation operators, Gray coding to reduce gate counts, mapping to physical machine-topologies, as well as gate-merging optimizations. Applying QNC tool to the unitary matrix U_s gives us a corresponding quantum gate circuit G_s as shown in Figure 8 for a specific bit-string s .

QNC generated black-box circuits with following gate-counts for the non-trivial 2-bit hidden-strings: “01”: 36, “11”: 38, “10”: 37, with estimated execution time² for critical path $\sim 17\mu s$ on an ideal machine with all-to-all connection topology. For the 5-qubit `ibmqx4` machine the corresponding gate-counts where: “01”: 42, “11”: 43, “10”: 41, with estimated execution time for critical path $\sim 15\mu s$, and for the 16-qubit `ibmqx5`, they were: “01”: 66, “11”: 67, “10”: 67, with estimated execution time for critical path $\sim 28\mu s$. In all these cases, QNC used a specialized decomposition of U_{01} , considering its permutation matrix nature, and therefore was able to reduce gate-counts

² These times are estimated using the data available from IBM at the time of writing. These values will change as the hardware improves.

by $5\times$ over the case when this special structure was ignored. Considering that the machines' observed coherence times are of the order of $\sim 60\mu\text{s}$, these QNC optimizations were crucial to the feasibility of the resulting score. The quantum score (circuit) generated by QNC for U_{01} for ibmqx4 is shown in Figure 9. A similarly prepared score for 3-bit hidden-string "111" had a gate-count of 428 in the ibmqx4 architecture with an estimated execution time of $153\mu\text{s}$ which was well above the machines' coherence times.




Fig. 9. Quantum circuit for BV algorithm with hidden string "01" targeting the ibmqx4 architecture.

We tested the QNC generated quantum scores for all non-trivial 1-qubit, 2-bit and 3-bit strings using the IBM-Qiskit based local simulator. In all cases, the simulator produced the exact hidden-string as the measurement result, 100% of the trials. We then tested all 1-bit and 2-bit strings on both the 5-qubit ibmqx4 and the 16-qubit ibmqx5 machines. The results are shown in Figure 10. For 2-bit strings, the worst case noise was observed for the string "01" on ibmqx4 when the qubits q_0, q_1, q_2 were used for x_0, x_1, y respectively. Since the estimated critical path times exceeded the machines' coherence times for 3-bit strings, we did not run those scores on the physical machines. Even for 2-bit strings, the scores were quite long, and the results were quite noisy even with 8192 machine-shots.




Fig. 10. Results from running the BV algorithm for 8192 shots on 2-bit hidden-strings "01", "10" and "11" respectively (left to right) on ibmqx4 . The y-axis here is the probability of obtaining the hidden string, which theoretically should be 1.

4 LINEAR SYSTEMS

4.1 Problem definition and background

Solving linear systems is central to a majority of science, engineering, finance and economics applications. For example, one comes across such systems while solving differential or partial differential equations or while performing regression. The problem of solving a system of linear equations is the following: Given a system $A\vec{x} = \vec{b}$, find \vec{x} for a given matrix A and vector \vec{b} . Here we assume that A is a Hermitian matrix, in that it is self-adjoint. To represent \vec{x}, \vec{b} as quantum states $|x\rangle, |b\rangle$, respectively, one has to rescale them as unit vectors, such that $\|\vec{x}\| = \|\vec{b}\| = 1$. Thus, one can pose the problem as finding $|x\rangle$ such that

$$A|x\rangle = |b\rangle, \quad (41)$$

with the solution $|x\rangle$ being

$$|x\rangle = \frac{A^{-1}|b\rangle}{\|A^{-1}|b\rangle\|}. \quad (42)$$

4.2 Algorithm description

The quantum algorithm for the linear system was first proposed by Harrow, Hassidim, and Lloyd (HHL) [66]. The HHL algorithm has been implemented on various quantum computers in [12, 27, 137]. The problem of solving for \vec{x} in the system $A\vec{x} = \vec{b}$ is posed as obtaining expectation value of some operator M with \vec{x} , $\vec{x}^\dagger M \vec{x}$, instead of directly obtaining the value of \vec{x} . This is particularly useful when solving on a quantum computer, since one usually obtains probabilities with respect to some measurement, typically, these operators are Pauli's operators X, Y, Z . These probabilities can then be translated to expectation values with respect to these operators.

The user has to keep in mind certain caveats while using the HHL algorithm. The algorithm requires that the elements of \vec{b} be accessible in superposition. Also, the solution vector is given as a quantum state which collapses after every measurement. This means that the HHL algorithm involves additional overheads for loading and reading data from a quantum computer [2]. Recently, classical algorithms inspired by HHL have been developed that, while having assumptions similar to HHL, considerably reduce the complexity of solving linear systems on classical computers [29]

The main idea of the HHL algorithm is as follows. Let $\{|u_j\rangle\}$ and $\{\lambda_j\}$ be the eigenvectors and eigenvalues of A , respectively, with the eigenvalues rescaled such that $0 < \lambda_j < 1$. Then the state $|b\rangle$, can be written as a linear combination of the eigenvectors $\{|u_j\rangle\}$, $|b\rangle = \sum_{j=1}^N \beta_j |u_j\rangle$. The goal of the HHL algorithm is to obtain $|x\rangle$ in the form $|x\rangle = \sum_{j=1}^N \beta_j \frac{1}{\lambda_j} |u_j\rangle$. By decomposing $A = R^\dagger \Lambda R$, the HHL algorithms in a nutshell involves performing a set of operations that essentially performs the three steps:

$$R^\dagger \Lambda R |x\rangle = |b\rangle \xrightarrow{\text{Step1}} \Lambda R |x\rangle = R |b\rangle \xrightarrow{\text{Step2}} R |x\rangle = \Lambda^{-1} R |b\rangle \xrightarrow{\text{Step3}} |x\rangle = R^\dagger \Lambda^{-1} R |b\rangle \quad (43)$$

This procedure requires us to find the eigenvalues of A . This can be done using a quantum subroutine called *phase estimation*. We will discuss this subroutine in some detail as it is a common ingredient in many quantum algorithms.

4.3 Phase estimation

Phase estimation is a quantum subroutine that lets us find the eigenvalues of a unitary matrix U given the ability to apply it to a quantum register as a controlled gate. Let $|u\rangle$ be an eigenvector of U such that, $U |u\rangle = e^{2\pi i \lambda_u} |u\rangle$. Then the phase estimation subroutine effects the following transformation,

$$|0\rangle |u\rangle \rightarrow |\tilde{\lambda}_u\rangle |u\rangle. \quad (44)$$

Here $\tilde{\lambda}_u$ is an estimate for λ_u . This subroutine makes use of an important transformation called the Quantum Fourier Transform (QFT)

Quantum Fourier Transform. The Discrete Fourier Transform (DFT) takes as an input a vector X of size N and outputs vector $Y = WX$ where the *Fourier matrix* W is defined by

$$W = \frac{1}{\sqrt{N}} \begin{bmatrix} 1 & 1 & 1 & \dots & 1 \\ 1 & \omega & \omega^2 & \dots & \omega^{N-1} \\ 1 & \omega^2 & \omega^4 & \dots & \omega^{2(N-1)} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & \omega^{N-1} & \omega^{2(N-1)} & \dots & \omega^{(N-1)(N-1)} \end{bmatrix},$$

where the ij -th element of the matrix is $W_{ij} = \omega^{ij}$ and ω is a primitive N -th root of one ($\omega^N = 1$). A straightforward implementation of the matrix-vector multiplication takes $O(N^2)$ operations, but, by using the special structure of the matrix, the Fast Fourier Transform (FFT) does the multiplication in only $O(N \log N)$ time. The algorithm is recursive and is illustrated on Figure 11. The Quantum Fourier Transform (QFT) is defined as a transformation




Fig. 11. Fast Fourier Transform circuit, where j denotes a row from the top half of the circuit and ω^j denotes that the corresponding value is multiplied by ω^j . The plus and minus symbols indicate that the corresponding values have to be added or subtracted, respectively.

between two quantum states that are determined using the values of DFT (FFT). If W is a Fourier matrix and $X = \{x_i\}$ and $Y = \{y_i\}$ are vectors such that $Y = WX$, then the QFT is defined as the transformation

$$QFT \left(\sum_{k=0}^{N-1} x_k |k\rangle \right) = \sum_{k=0}^{N-1} y_k |k\rangle. \tag{45}$$

The implementation of the QFT mimics the stages (recursive calls) of the FFT, but implements each stage using only $n + 1$ additional gates per stage. A single Hadamard gate on the last (least significant) bit implements the additions/subtractions of the outputs from the recursive call and the multiplications by ω^j are done using n controlled phase gates. The circuit for $n = 5$ is shown on Figure 12.

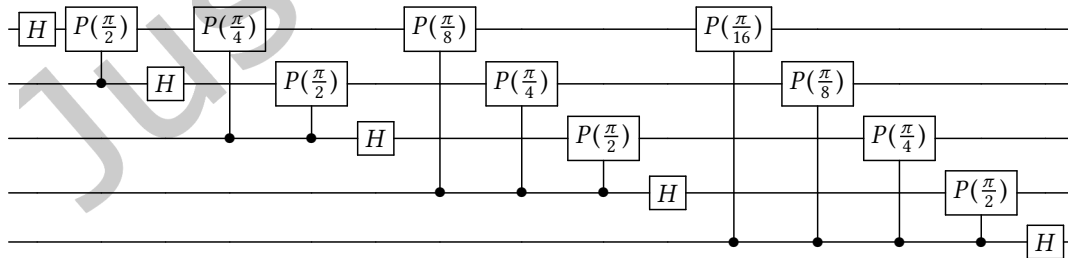


Fig. 12. A Quantum Fourier Transform circuit for five qubits ($n = 5$).

The phase estimation procedure cleverly uses the QFT operator to estimate the eigenphases of the operator U . The circuit for performing phase estimation given in Fig. 13. Notice that the QFT is applied in reverse.




Fig. 13. Quantum circuit for phase estimation.

The pseudocode for phase estimation is given in Algorithm 3. Notice that the algorithm also works if the input state is not an eigenstate. The output in this case can be determined by expanding the input state in terms of the eigenstates and then applying the linearity of quantum operations. In the code, we have numbered the ancillary qubits from the top and $C_i U$ denotes the unitary controlled by the i^{th} ancilla qubit acting on the main n qubit register.

Algorithm 3 Phase estimation subroutine

Input:

- Controlled unitaries $C_i U$
- An n qubit input state $|\psi\rangle = \sum_u \psi_u |u\rangle$, where $U |u\rangle = e^{2\pi i \lambda_u} |u\rangle$.

Output:

- $\sum_u \psi_u |\tilde{\lambda}_u\rangle |u\rangle$

Procedure:

Step 1. Take t ancillary qubits initialized to zero and perform $H^{\otimes t}$ on them to produce the uniform superposition state over them.

for $0 \leq i < t$ **do**

Step 2. Apply $C_{t-i-1} U^{2^i}$

end for

Step 3. Apply QFT^\dagger .

Optional Measure the ancillary qubits to get $|\tilde{\lambda}_u\rangle |u\rangle$ with probability $|\psi_u|^2$

The number of ancillary qubits used in the phase estimation algorithm will determine both its run-time and its accuracy. On the accuracy front, the number of ancillary qubits used is equal to the bit precision of $\tilde{\lambda}_u$ as the answer is stored in this register. The exact complexity of this subroutine is discussed in Ref. [92].

Now we can discuss the HHL algorithm which makes use of the phase estimation procedure to perform a matrix inversion. The HHL algorithm requires three sets of qubits: a single ancilla qubit, a register of n qubits used to store the eigenvalues of A in binary format with precision up to n bits, and a memory of $O(\log(N))$ that initially stores $|b\rangle$ and eventually stores $|x\rangle$. Start with a state $|0\rangle_a |0\rangle_r |b\rangle_m$, where the subscripts a, r, m , denote the sets of ancilla, register and memory qubits, respectively. This subscript notation was used in [137], and we

found it to be most useful in keeping things clear. The HHL algorithm requires us to run the phase estimation procedure on the unitary operator e^{iA} . The phases estimated would be approximations to the eigenvalues of A . The problem of applying the unitary operation e^{iA} given the matrix A is called *quantum simulation*. There are many algorithms in literature that tackle the problem of quantum simulation [17] [56] and that will not be our focus in this section. We will explain the steps of the HHL algorithm below assuming that the quantum simulation part is taken care of. We will also include some mathematical details in the pseudocode given in Algorithm 4 .

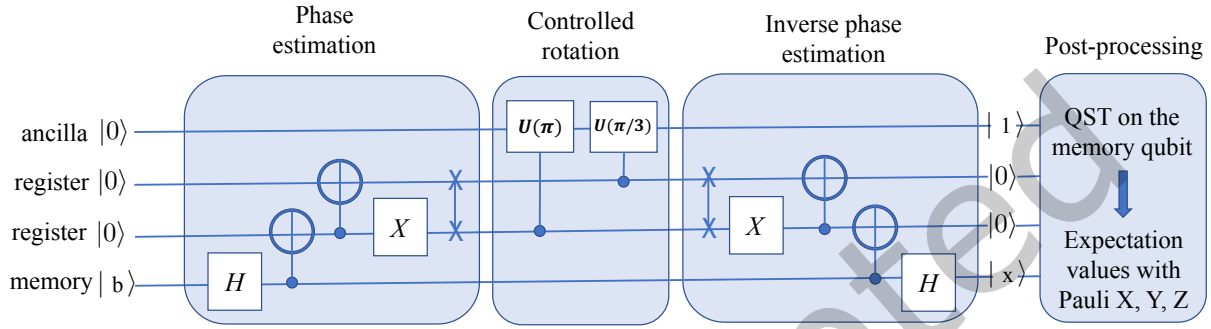


Fig. 14. Schematic of the circuit for the quantum algorithm for solving a 2×2 linear system. The first step involves phase estimation, which maps the eigenvalues λ_j of A into the register in the binary form. The second step involves controlled rotation of the ancilla qubit, so that the inverse of the eigenvalues $\frac{1}{\lambda_j}$ show up in the state. The third step is the inverse phase estimation to disentangle the system, and restores the registers to $|0\rangle$. The memory qubit now stores $|x\rangle$, which is then post-processed to get the expectation values with respect to the Pauli operators X , Y and Z .

These three steps are equivalent to the three steps shown in Eq. (43). The algorithm is probabilistic, we get $|x\rangle$ only if the final measurement gives $|1\rangle$. But this probability can be boosted using a technique called *amplitude amplification* [23]. This technique is explained in detail in Section VII.

4.4 Algorithm implemented on IBM's 5 qubit computer

Now we implement the HHL algorithm on a 2×2 system. For this, we chose $A = \begin{pmatrix} 1.5 & 0.5 \\ 0.5 & 1.5 \end{pmatrix}$. We use four qubits for solving the system – one ancilla, one memory and two register qubits. For this case, the eigenvalues of A are $\lambda_1 = 1$ and $\lambda_2 = 2$ with the eigenvectors being $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix} \equiv |-\rangle$ and $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix} \equiv |+\rangle$, respectively. For this system, the three steps of the HHL algorithm, can be performed by the operations shown in Fig. 14. For the controlled rotation, we use a controlled U rotation with $\theta = \pi$ for λ_1 and $\theta = \pi/3$ for λ_2 . This is done by setting $C = 1$ in the Eq. (47). Both λ and ϕ are set to zero in these controlled U rotations. Although the composer on Quantum Experience does not have this gate, in IBM Qiskit-sdk-py, we use `cu3` function for this purpose. Three cases are used for b : $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ -1 \end{pmatrix}$ and $\frac{1}{\sqrt{2}} \begin{pmatrix} 1 \\ 1 \end{pmatrix}$. We post selected the states with $|1\rangle$ in the ancilla qubit. The probabilities of these states are normalized such that their sum is one. Measurements with respect to $\langle X \rangle$, $\langle Y \rangle$, $\langle Z \rangle$ can then be performed to obtain the expectation values. QASM code is output from Qiskit-sdk-py and then uploaded on to IBM Quantum Experience. Figure 15 shows the equivalent composer circuit generated from QASM for the measurement in the computational basis (Z measurement).

To first test our implementation of the algorithm, we ran nine cases on the local simulator provided by Qiskit-sdk-py – three b cases and three measurements with respect to the operators X , Y , Z , for each b case. The

Algorithm 4 HHL algorithm**Input:**

- The state $|b\rangle = \sum_j \beta_j |u_j\rangle$
- The ability to perform controlled operations with unitaries of the form e^{iAt}

Output:

- The quantum state $|x\rangle$ such that $A\vec{x} = \vec{b}$.

Procedure:

Step 1. Perform quantum phase estimation using the unitary transformation e^{iA} . This maps the eigenvalues λ_j into the register in the binary form to transform the system,

$$|0\rangle_a |0\rangle_r |b\rangle_m \rightarrow \sum_{j=1}^N \beta_j |0\rangle_a |\lambda_j\rangle_r |u_j\rangle_m. \quad (46)$$

Step 2. Rotate the ancilla qubit $|0\rangle_a$ to $\sqrt{1 - \frac{C^2}{\lambda_j^2}} |0\rangle_a + \frac{C}{\lambda_j} |1\rangle_a$ for each λ_j . This is performed through controlled rotation on the $|0\rangle_a$ ancilla qubit. The system will evolve to

$$\sum_{j=1}^N \beta_j \left(\sqrt{1 - \frac{C^2}{\lambda_j^2}} |0\rangle_a + \frac{C}{\lambda_j} |1\rangle_a \right) |\lambda_j\rangle_r |u_j\rangle_m. \quad (47)$$

Step 3. Perform the reverse of Step 1. This will lead the system to

$$\sum_{j=1}^N \beta_j \left(\sqrt{1 - \frac{C^2}{\lambda_j^2}} |0\rangle_a + \frac{C}{\lambda_j} |1\rangle_a \right) |0\rangle_r |u_j\rangle_m. \quad (48)$$

Step 4. Measuring the ancilla qubit will give ,

$$|x\rangle \approx \sum_{j=1}^N C \left(\frac{\beta_j}{\lambda_j} \right) |u_j\rangle, \quad (49)$$

if the measurement outcome is $|1\rangle$

comparison between the theoretical expectation values $\langle X \rangle$, $\langle Y \rangle$, $\langle Z \rangle$ and the simulator values are shown in Table 3. The simulator expectation values and the theoretical values match well. This shows that the implementation of the algorithm gives expected results. Similar expectation values were also seen using the simulator on IBM Quantum Experience instead of the local simulator. We then ran the circuit on the quantum computer `ibmqx4`. Fig. 16 shows a comparison between the simulator results and the results from the `ibmqx4` with Z measurement on the circuit. As can be seen from Fig. 16, the results from the actual run do not give the expected answer as seen in the simulator results. We remark that recent modifications to the algorithm [24, 121] can in some cases allow for larger scale and more accurate implementations on noisy quantum computers.

5 SHOR'S ALGORITHM FOR INTEGER FACTORIZATION

5.1 Problem definition and background

The integer factorization problem asks, given an integer N as an input, to find integers $1 < N_1, N_2 < N$ such that $N = N_1 N_2$. This problem is hardest when N_1 and N_2 are primes with roughly the same number of bits. If n denotes the number of bits of N , no algorithm with polynomial in n time complexity is known. The straightforward




Fig. 15. Circuit implemented on IBM's 5-qubit `ibmqx4` quantum computer for the case with $|b\rangle$ set to $|0\rangle$ and with $\langle Z \rangle$ measurement. After implementing the circuit in Fig. 14 and setting the coupling map of the `ibmqx4` architecture, Qiskit-sdk-py re-arranges the qubits to fit the mapping. This circuit represents the outcome of the re-arrangement which was implemented on the `ibmqx4` quantum computer.

Table 3. Comparison between theoretical and simulator values for the expectation values $\langle X \rangle$, $\langle Y \rangle$, $\langle Z \rangle$. T stands for theoretical and S stands for simulator.

$ b\rangle$	T $\langle X \rangle$	S $\langle X \rangle$	T $\langle Y \rangle$	S $\langle Y \rangle$	T $\langle Z \rangle$	S $\langle Z \rangle$
$ 0\rangle$	-0.60	-0.60	0.00	-0.027	0.80	0.81
$ +\rangle$	1.00	1.00	0.00	-0.06	0.00	0.02
$ -\rangle$	-1.00	-1.00	0.0060	0.000	-0.02	0.00

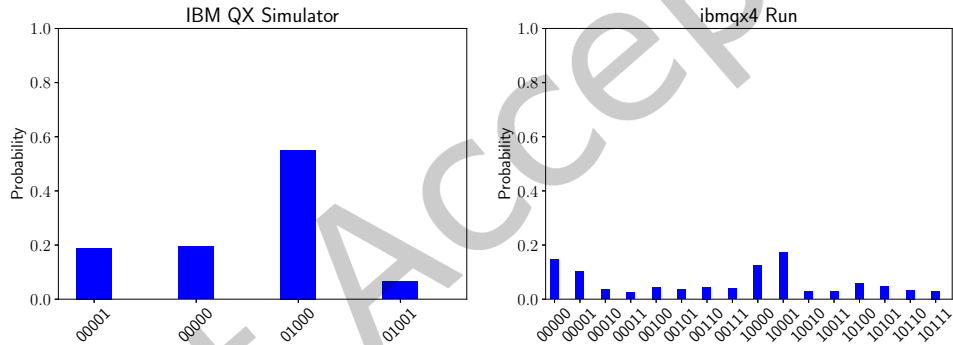


Fig. 16. Results of the circuit with Z measurement (computational basis measurement) from the actual run and the simulator on a `ibmqx4`. 4096 shots were used for both the cases.

algorithm that tries all factors from 2 to \sqrt{N} takes time polynomial in N , but exponential in n . The most efficient known classical algorithm has running time $O\left(\exp\left(\sqrt[3]{\frac{64}{9}n(\log n)^2}\right)\right)$ [99]. In practice, integers with 1000 or more bits are impossible to factor using known algorithms and classical hardware. The difficulty of factoring big numbers is the basis for the security of the RSA cryptosystem [104], one of the most widely used public-key cryptosystems.

One of the most celebrated results in quantum computing is the development of a quantum algorithm for factorization that works in time polynomial in n . This algorithm, due to Peter Shor and known as Shor's algorithm [113], runs in $O(n^3 \log n)$ time and uses $O(n^2 \log n \log \log n)$ gates. The first experimental implementation of this algorithm on a quantum computer was reported in 2001, when the number 15 was factored [127]. The largest integer factored by Shor's algorithm so far is 21 [88].

In this section we describe Shor's algorithm and its implementation on `ibmqx4`

5.2 Algorithm description

Reducing factorization to period finding. One way to factor an integer is by using modular exponentiation. Specifically, let an odd integer $N = N_1N_2$ be given, where $1 < N_1, N_2 < N$. Pick any integer $k < N$ such that $\gcd(k, N) = 1$, where \gcd denotes the greatest common divisor. One can show that there exists an exponent $p > 0$ such that $k^p \equiv 1 \pmod{N}$. Recall that, by definition, $x \equiv y \pmod{m}$ if and only if m divides $x - y$. Assume that p is the smallest such number. If we find such p and p is even, then, by the definition of the modulo operation, N divides

$$k^p - 1 = (k^{p/2} - 1)(k^{p/2} + 1).$$

But since the difference between $n_1 = k^{p/2} + 1$ and $n_2 = k^{p/2} - 1$ is 2, n_1 and n_2 have no common factor greater than 2. Moreover, both numbers are nonzeros by the minimality of p . Since $N = N_1N_2$ was assumed to be odd, then N_1 is a factor of either n_1 or n_2 . Assume N_1 is a factor of n_1 . Since N_1 is also a factor of N , then N_1 divides both n_1 and N and one can find N_1 by computing $\gcd(n_1, N)$. Hence, if one can compute such a p , one can find the factors of N efficiently as \gcd can be computed in polynomial time.

In order to find p , consider the modular exponentiation sequence $A = a_0, a_1, \dots$, where $a_i = k^i \pmod{N}$. Each a_i is a number from the finite set $\{0, \dots, N - 1\}$, and hence there exists indices q and r such that $a_q = a_r$. If q and r are the smallest such indices, one can show that $q = 0$ and A is periodic with period r . For instance, for $N = 15$ and $k = 7$, the modular exponentiation sequence is $1, 7, 4, 13, 1, 7, 4, 13, 1, \dots$ with period 4. Since the period 4 is an even number, we can apply the above idea to find

$$7^4 \pmod{15} \equiv 1 \Rightarrow 7^4 - 1 \pmod{15} \equiv 0 \Rightarrow (7^2 - 1)(7^2 + 1) \pmod{15} \equiv 0 \Rightarrow 15 \text{ divides } 48 \cdot 50,$$

which can be used to compute the factors of 15 as $\gcd(48, 15) = 3$ and $\gcd(50, 15) = 5$.

Finding the period of the sequence A is, however, not classically easier than directly searching for factors of N , since one may need to check as many as \sqrt{N} different values of A before encountering a repetition. However, with quantum computing, the period can be found in polynomial time using the Quantum Fourier Transform (QFT). The QFT operation was introduced earlier during our discussion of phase estimation.

The property of the QFT that is essential for the factorization algorithm is that it can “compute” the period of a periodic input. Specifically, if the input vector X is of length M and period r , where r divides M , and its elements are of the form

$$x_i = \begin{cases} \sqrt{r/M} & \text{if } i \pmod{r} \equiv s \\ 0 & \text{otherwise} \end{cases}$$

for some offset $s < r$, and $QFT\left(\sum_{i=0}^M x_i |i\rangle\right) = \sum_{i=0}^M y_i |i\rangle$, then

$$y_i = \begin{cases} 1/\sqrt{r} & \text{if } i \pmod{M/r} \equiv 0 \\ 0 & \text{otherwise} \end{cases}$$

i.e., the output has nonzero values at multiples of M/r (the values $\sqrt{r/M}$ and $1/\sqrt{r}$ are used for normalization). Then, in order to factor an integer, one can find the period of the corresponding modular exponentiation sequence using QFT, if one is able to encode its period in the amplitudes of a quantum state (the input to QFT).

A period-finding circuit for solving the integer factorization problem is shown in Fig 17 [39]. The first QFT on register A produces an equal superposition of the qubits from A , i.e., the resulting state is

$$\frac{1}{\sqrt{M}} \sum_{i=0}^M |i, 0\rangle.$$




Fig. 17. Illustration of the period-finding circuit, where $m = 2n$ and $M = 2^m$.

Next is a modular exponentiation circuit that computes the function $f(i) = x^i \pmod{N}$ on the second register. The resulting state is

$$\frac{1}{\sqrt{M}} \sum_{i=0}^{M-1} |i, f(i)\rangle.$$

Before we apply the next QFT transform, we do a measurement of register B . (By the principle of deferred measurement [92] and due to the fact that register A and B don't interact from that point on, we don't have to actually implement the measurement, but it will help to understand the final output.) If the value measured is s , then the resulting state becomes

$$\frac{1}{\sqrt{M/r}} \sum_{\substack{i=0 \\ f(i)=s}}^{M-1} |i, s\rangle,$$

where r is the period of $f(i)$. In particular, register A is a superposition with equal non-zero amplitudes only of $|i\rangle$ for which $f(i) = s$, i.e., it is a periodic superposition with period r . Given the property of QFT, the result of the transformation is the state

$$\frac{1}{\sqrt{r}} \sum_{i=0}^{r-1} |i(M/r), s\rangle.$$

Hence, the measurement of register A will output a multiple of M/r . If the simplifying assumption that r divides M is not made, then the circuit is the same, but the classical postprocessing is a bit more involved [92].

Period finding can also be viewed as a special case of phase estimation. The reader may refer Nielsen and Chuang [92] for this perspective on period finding.

5.3 Algorithm implemented on IBM's 5-qubit computer

We implemented the algorithm on `ibmqx4`, a 5-qubit quantum processor from the IBM Quantum Experience, in order to factor number 15 with $x = 11$. The circuit as described on Figure 17 requires 12 qubits and 196 gates, too large to be implemented on `ibmqx4`. Hence, we used an optimized/compiled version from [127] that uses 5 qubit and 11 gates (Fig 18).

The results from the measurements are shown on Figure 19.




Fig. 18. Circuit for Shor's algorithm for $N = 15$ and $x = 11$.





Fig. 19. Output from the circuit from Figure 18 implemented on the simulator (left) and ibmqx4 (right).

The periods found by the simulator are $p = 0$, which is ignored as a trivial period, and $p = 4$, which is a good one. Since $M = 8$, we can conclude that r divides $M/p = 8/4 = 2$, hence $r = 2$. Then 15 divides

$$(x^r - 1) = (11^2 - 1) = (11 - 1)(11 + 1) = 10 \cdot 12.$$

By computing $\gcd(15, 10) = 5$ and $\gcd(15, 12) = 3$, we find the factors of 15.

The output from `ibmqx4` finds the same periods 0 and 4 with the highest probabilities, but contains much more noise.

6 MATRIX ELEMENTS OF GROUP REPRESENTATIONS

6.1 Problem definition and background

In this section we will discuss another quantum algorithm that makes use of the QFT operation. In this section we will also introduce a subroutine called the *Hadamard test*, which lets us compute matrix elements of unitary operators. But first, we will require some knowledge of group theory to understand the problem being tackled here. This section follows the work of Jordan in Ref. [73].

A *Group* (G, \cdot) or (G) is a mathematical object defined by its elements (g_1, g_2, \dots) and an operation between elements (\cdot) , such that these four properties are satisfied.

- (1) Closure: for any two group elements, the defined group operation produces another element, which belongs to the group (for $\forall g_i, g_j \in G, g_i \cdot g_j = g_k \in G$).
- (2) Associativity: for $\forall g_i, g_j, g_m \in G, g_i \cdot (g_j \cdot g_m) = (g_i \cdot g_j) \cdot g_m$.
- (3) Identity element: $e \in G$, such that $e \cdot g_i = g_i \cdot e = g_i$.
- (4) Inverse element: for $\forall g_i \in G$, there exists g_p , such that $g_i \cdot g_p = g_p \cdot g_i = e$.

A group with a finite amount of elements n is called a finite group with order n , while a group with an infinite amount of elements is an infinite group. In this section, we will discuss quantum algorithms to solve certain problems related to finite groups. As before, we will also implement them on the IBM machines. Some examples of groups are given below.

Example 1A. Abelian group A_n with n elements: $0, 1, \dots, n-1$, and the group operation addition modulo n : $g_i \cdot g_j = (i + j) \bmod(n)$. For instance, for $n = 3$: $a_0 = 0, a_1 = 1, a_2 = 2$. Then, $a_2 \cdot a_2 = 4 \bmod(3) = 1 = a_1$, $a_2 \cdot a_1 = 3 \bmod(3) = 0 = a_0$, etc. The identity element is $a_0 = 0$ and its inverse is itself. For all other elements the inverse element is, $a_i^{-1} = a_{n-i}$. This group is called Abelian or commutative, because in addition to the four group properties, it has a property of commutativity: $a_i \cdot a_j = a_j \cdot a_i$ for $\forall a_i, a_j \in A_n$.

Example 1S. Symmetry group S_n with $n!$ group elements, each is a permutation of n objects: $[1, 2, \dots, n], [2, 1, \dots, n], \dots, [n, n-1, \dots, 2, 1]$. Consequent application of two permutations is a group operation. For instance, for group S_2 : (e, p) we have two objects a and b . The identity element e is no permutation: $ab \rightarrow ab$, while one permutation p is the second group element: $ab \rightarrow ba$. Then, $p \cdot p = e$, and $p^{-1} = p$. Only S_1 and S_2 are Abelian groups. For $n \geq 3$, S_n are not commutative. Let us write elements of group S_3 as a permutation of elements 123 in the next order: $[123] \rightarrow [123], [231], [312], [213], [132], [321]$. Then $s_4 \cdot s_2 = s_6$, while $s_2 \cdot s_4 = s_5$.

While group definition is quite simple, it is not straightforward how to operate with group elements in general, especially when defined operations between them is not trivial and/or the group order, n , is large. In this case, it is helpful to apply the representation theory to the group. The idea is simple: if we can map a group of unknown objects with nontrivial operations to the group of known objects with some trivial operations, we can gain some information about the unknown group. In general, we introduce a function applied to a group element: $\rho(g_i)$, which does this mapping between two groups. Such function defines the group representation of G if for $\forall g_i, g_j \in G, \rho(g_i) * \rho(g_j) = \rho(g_i \cdot g_j)$, where $(*)$ can be a different operation from (\cdot) .

Example 2A. Representation of Abelian group A_n : $a_j \rightarrow \rho(a_j) = e^{i2\pi j/N}$, where the original operation $(+\bmod(n))$ is substituted by the new operation of multiplication. Note that the group S_2 can be represented in the same way as A_2 .

Example 2S. Representation of group S_3 : $s_j \rightarrow \rho(s_j) = 1$, where the original operation is again substituted by the new operation of multiplication. Such representation of the group S_3 is trivial, since it does not carry any

information about the group, however it satisfies the definition of the group representation. Moreover, $[1, 1, \dots]$ is a trivial representation for any group. Another representation of group S_3 is, $[1, 1, 1, -1, -1, -1] \rightarrow [s_1, s_2, \dots, s_n]$, where we map odd permutations to -1 and even permutations to 1 . While it carries more information about the initial group than the trivial representation, it does not imply that the group S_3 is not Abelian. One cannot construct a one-dimensional representation for group S_3 which would retain all its properties. The smallest equivalent representation for S_3 is two-dimensional. The multidimensional representations can be easily understood when represented by matrices.

Most useful representations are often ones which map a group to a set of matrices. When $\rho(g)$ is a $d_\rho \times d_\rho$ matrix, the representation is referenced as a matrix representation of the order d_ρ , while $(*)$ is the operation of matrix multiplication. All representations of finite group can be expressed as unitary matrices given an appropriate choice of basis. To prove the last fact, we introduce a particular representation called the *regular representation*.

The regular representation of a group of N elements is a matrix representation of order N . We will explain the construction of the regular representation using the Dirac notation. First, we associate with each element of the group g_i a ket $|g_i\rangle$. This ket could simply be the basis state $|i\rangle$, since the elements of the group are numbered. This ensures that the kets associated with different group elements are orthonormal by construction, $\langle g_i | g_j \rangle = \delta_{ij}$. This also ensures that the identity operator can be expressed as $\sum_{i=1}^N |g_i\rangle \langle g_i|$. The regular representation of g_k is then given by,

$$R(g_k) = \sum_{j=1}^N |g_k \cdot g_j\rangle \langle g_j|. \quad (50)$$

The matrix elements of this representation are, $R_{ij}(g_k) \equiv \langle g_i | R(g_k) | g_j \rangle = \langle g_i | g_k \cdot g_j \rangle$. From the defining properties of a group it can be easily seen that multiplying every element in the group by the same element just permutes the elements of the group. This means that $R(g_k)$ matrices are always permutation matrices and are hence unitary. We can prove that the regular representation is a representation using simple algebra,

$$\begin{aligned} R(g_k) \cdot R(g_m) &= \sum_{i=1}^N \sum_{j=1}^N |g_k \cdot g_i\rangle \langle g_i | g_m \cdot g_j \rangle \langle g_j|, \\ &= \sum_{i=1}^N \sum_{j=1}^N |g_k \cdot g_m \cdot g_j\rangle \langle g_i | g_m \cdot g_j \rangle \langle g_j|, \\ &= \sum_{j=1}^N |g_k \cdot g_m \cdot g_j\rangle \langle g_j| = R(g_k \cdot g_m). \end{aligned} \quad (51)$$

Here we used orthogonality: $\langle g_i | g_m \cdot g_j \rangle = 1$ only if $|g_i\rangle = |g_m \cdot g_j\rangle$ and 0 otherwise, which allowed us to swap these two states. Then, we used the same fact to calculate the sum over i . Below we give some explicit examples of regular representations.

Example 3A. Regular representation of the Abelian group A_4 , where each matrix element is calculated using the result derived above $R_{ij}(a_k) = \langle a_i | a_k \cdot a_j \rangle$:

$$R(a_0) = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad R(a_1) = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad R(a_2) = \begin{pmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \quad R(a_3) = \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \end{pmatrix}. \quad (52)$$

Commutative property is conserved: $R(a_i) \cdot R(a_j) = R(a_j) \cdot R(a_i)$.

Example 3S. Regular representation of the group S_3 , where we use the same order of permutations introduced above ($[123] \rightarrow [123], [231], [312], [213], [132], [321]$)

$$R(s_1) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix}, R(s_2) = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \end{pmatrix}, R(s_3) = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 & 0 \end{pmatrix}, \quad (53)$$

$$R(s_4) = \begin{pmatrix} 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{pmatrix}, R(s_5) = \begin{pmatrix} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix}, R(s_6) = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (54)$$

Now we can finally explain the problem of calculating matrix elements of the group representations, which is equivalent to the problem of calculating an expectation value of an operator A in respect to the state $|\psi\rangle$ in quantum mechanics: $\langle A \rangle = \langle \psi | A | \psi \rangle$.

Example 4A. Calculating matrix elements of the regular representation of the element a_2 from the Abelian group A_4 with respect to the state ψ_{13} which is the equal superposition of $|a_1\rangle$ and $|a_3\rangle$. In operator form we find:

$$\langle \psi_{12} | a_2 | \psi_{12} \rangle = \frac{\langle a_1 | + \langle a_3 |}{\sqrt{2}} \left(\sum_{i=0}^{N-1} |a_2 \cdot a_i\rangle \langle a_i| \right) \frac{|a_1\rangle + |a_3\rangle}{\sqrt{2}} = \frac{\langle a_3 | a_2 \cdot a_1 \rangle \langle a_1 | a_1 \rangle}{2} + \frac{\langle a_1 | a_2 \cdot a_3 \rangle \langle a_3 | a_3 \rangle}{2} = 1. \quad (55)$$

It is quite obvious that if a quantum computer is capable of finding expectation values of a unitary operator, it will be able to solve the problem of finding the matrix elements of the regular representation of a group element. This will consist of, at least, two stages: the first stage is the state preparation, and the second is applying the unitary operator of the regular representation to that state. The unitary operator of the regular representation of an element of any group G_n can be created using a combination of only two type of operations: qubit flip ($|0\rangle \rightarrow |1\rangle$) and qubit swap ($|q_j q_i\rangle \rightarrow |q_i q_j\rangle$).

Up to this point, we have only talked about the regular representation. The regular representation is quite convenient, it is straightforward to find for any group, it carries all the information about the group, and a corresponding unitary operator is easy to construct using standard quantum circuits. However, for groups with a large number of elements, it requires matrix multiplication between large matrices. So for many applications, instead of regular representations one is interested in what are known as *irreducible representations*, which are matrix representations that cannot be decomposed into smaller representations. Or in other words, every matrix representation (including the regular representation) can be shown to be equivalent to a direct sum of irreducible representations, up to a change of basis. This lets us reduce the representation theory of finite groups into the study of irreducible representations. The importance of irreducible representations in group theory cannot be overstated. The curious reader may refer these notes by Kaski [74].

A result from group theory ensures that the direct sum of all irreducible representations (each has different dimensions d_ρ in general) where each irreducible representation appears exactly d_ρ times is a block diagonal $N \times N$ matrix (the group has N elements). The Fourier transform pair over this group representation can be introduced by decomposing each irreducible representation over the group elements and *vice versa*. Moreover,

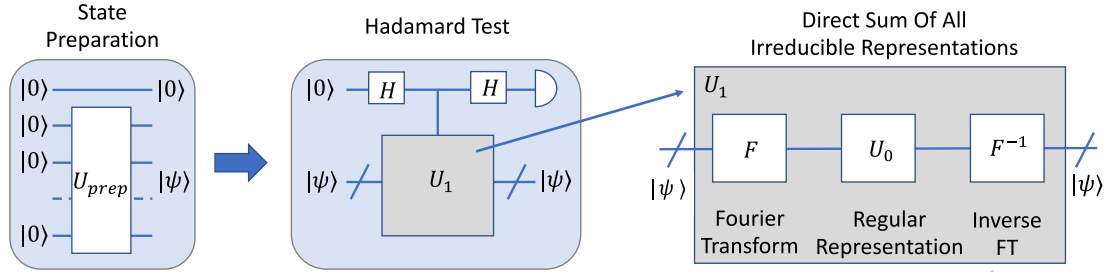


Fig. 20. Schematic diagram for the quantum algorithm

the above defined direct sum of all irreducible representations can be decomposed as a regular representation conjugated by the direct and inverse Fourier transform operators [73]. This result lets us find the the matrix elements of the irreducible representations given the ability to implement the regular representation.

6.2 Algorithm description

In this section we will describe an algorithm to find the matrix elements of irreducible representations of a group given the ability to apply its regular representations to a quantum register in a controlled fashion. The quantum algorithm calculating matrix elements $\langle \psi | U_1 | \psi \rangle$ of a unitary operator U_1 is known as the Hadamard test, which is illustrated on Fig. 20.

Algorithm 5 Hadamard test

Input:

- The controlled unitary CU .
- Input state $|0\rangle|\psi\rangle$.

Output:

- An estimate for the real part of $\langle \psi | U | \psi \rangle$

Procedure:

Step 1. Apply H to the ancilla. This produces the state,

$$\frac{|0\rangle + |1\rangle}{\sqrt{2}} |\psi\rangle$$

Step 2. Apply CU controlled on the ancilla. This produces the state,

$$\frac{|0\rangle |\psi\rangle + |1\rangle U |\psi\rangle}{\sqrt{2}}$$

Step 3. Apply H to the ancilla again. This gives,

$$\frac{|0\rangle (|\psi\rangle + U |\psi\rangle) + |1\rangle (|\psi\rangle - U |\psi\rangle)}{\sqrt{2}}$$

Step 4. Measure the ancillary qubit. Repeat to estimate the probability of obtaining $|0\rangle$ and $|1\rangle$.




Fig. 21. Actual circuit implemented on IBM's 5-qubit computer for calculating matrix elements of the regular representation for the second element of the group S_2 and A_2 in respect to the state $|0\rangle$ on the left and $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ on the right. The expected probabilities to find a final state in the ground state are $(1+0)/2 = 0.5$ and $(1+1)/2 = 1$ respectively. The results of the 1024 runs on the actual chip (on the top) and the simulator (on the bottom) are presented on the right side of each circuit.

The ancilla qubit should be prepared as $\frac{|0\rangle-i|1\rangle}{\sqrt{2}}$ to calculate the imaginary parts of the matrix element. From the pseudocode, we can see that the probability of measuring $|0\rangle$ is $P_0 = \|\frac{|\psi\rangle+U|\psi\rangle}{\sqrt{2}}\|^2 = \frac{1+Re\langle\psi|U|\psi\rangle}{2}$. Hence, we find: $Re\langle\psi|U|\psi\rangle = 2P_0 - 1$. The reader is encouraged to work out the same steps for the imaginary part as well.

With the Hadamard test algorithm, the problem of calculating matrix elements of an arbitrary unitary operator is reduced to the problem of effectively implementing it as a controlled gate. For the regular representation of any group U_0 , where unitary operator is an $N \times N$ square matrix with only one non-zero element equal to 1 in each row, this implementation can be done for any group as a combination of *CNOT* and *Z* gates.

At the same time solutions for the direct sum of all irreducible representations U_1 , which can be decomposed as $U_1(g) = F_1 U_0(g^{-1}) F_1^{-1}$, exists for any group whose Fourier transform over that group can be effectively implemented using quantum circuits. Quantum circuits for the Fourier transform are already known for the symmetric group $S(n)$ [13], the alternating group A_n , and some Lie groups: $SU(n)$, $SO(n)$ [10], while solutions for other groups, hopefully, will be found in the future. For Abelian groups this Fourier transform implementation can be efficiently done using the QFT circuit that was discussed in the earlier sections. For non-Abelian groups the implementation is trickier and efficient implementations are not universally known.

6.3 Algorithm implemented on IBM's 5-qubit computer

The actual gate sequence that we implemented on IBM's 5-qubit computer (ibmq_essex) and IBM's quantum simulator to find matrix elements of the regular representation of the second element of the group S_2 is shown in Fig. 21. The matrix for this representation is simply a *X* gate. Hence, we have to use one *CNOT* gate and two Hadamard gates, plus some gates to prepare state $|\psi\rangle$ from the state $|00\rangle$. We mapped the ancilla qubit to the actual machine q_1 qubit instead of q_0 , because of the machine architecture, where the first qubit can control the zero qubit but not *vice versa*. We could have used the original qubit sequence as in Fig. 20, by realizing the *CNOT* gate as a swapped *CNOT* and four Hadamard gates, but this would add more gates to the circuit and potentially more computational errors rather than just a virtual swap of the qubits.

For the irreducible representation of the same element of the group A_2 , the element is represented by the *Z* gate. Hence the Hadamard test requires implementing a controlled-*Z* gate, which is not available as an actual gate on the IBM Quantum Experience. However, it can be constructed using two Hadamard and one *CNOT* gates as shown in Fig. 22. Notice that the Hadamard gate is actually the Fourier transform operator over group S_2 and

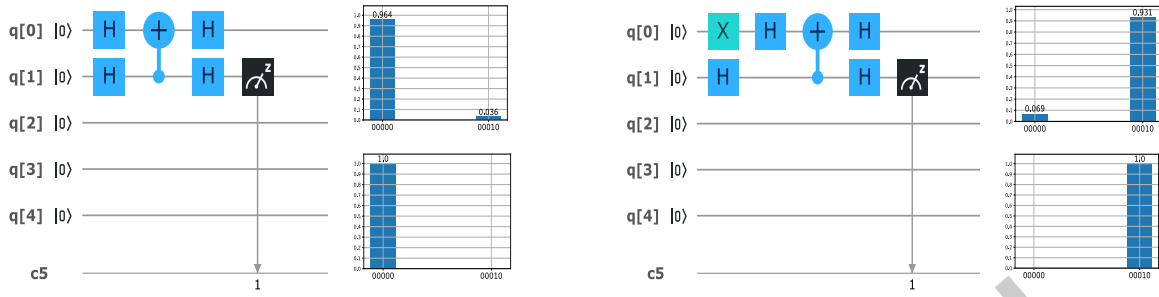


Fig. 22. Actual circuit implemented on IBM’s 5-qubit computer for calculating matrix elements of the direct sum of the irreducible representations for the second element of the group S_2 and A_2 with respect to the state $|0\rangle$ on the left and $|1\rangle$ on the right. The expected probabilities to find a final state in the ground state are $(1 + 1)/2 = 1$ and $(1 - 1)/2 = 0$ respectively. The results of the 1024 runs on the actual chip (on the top) and the simulator (on the bottom) are presented on the right side of each circuit.

A_2 , while the X gate is a regular representation operator, as we mentioned earlier. Hence, such controlled- Z gate representation is in fact the decomposition of the irreducible representation to the regular representation using Fourier transform over that group.

7 QUANTUM VERIFICATION OF MATRIX PRODUCTS

7.1 Problem definition and background

Matrix multiplication is one of the most important linear algebra subroutines. Most scientific computing algorithms use matrix multiplication in one form or another. Therefore, the computational complexity of matrix multiplication is a subject of intense study. For two $n \times n$ matrices the computational complexity of the naive matrix multiplication algorithm is $O(n^3)$. A faster algorithm for matrix multiplication implies a considerable performance improvement for a variety of computational tasks. Strassen [120] first showed that two $n \times n$ matrices can be multiplied in time $O(n^{2+\alpha})$ ($\alpha < 1$). The best known algorithm to date with $\alpha \approx 0.376$ was found by Coppersmith and Winograd [37]. Despite that, it remains an open problem to determine the optimal value of α . The so-called problem of matrix verification is defined as, verifying whether the product of two $n \times n$ matrices is equal to a third one. So far the best classical algorithm can do this with high probability in time proportional to n^2 [53].

Ref. [6] was the first to study matrix verification for quantum computation. The authors use a quantum algorithm based on Grover’s algorithm to verify whether two $n \times n$ matrices equal a third in time $O(n^{7/4})$, thereby improving the optimal classical bound of Ref. [53]. Ref. [26] presents a quantum algorithm that verifies a product of two $n \times n$ matrices over any integral domain with bounded error in worst-case time $O(n^{5/3})$ and expected time $O(n^{5/3}/\min(w, \sqrt{n})^{1/3})$, where w is the number of wrong entries. This further improves the time performance $O(n^{7/4})$ from Ref. [6].

7.2 Algorithm description

We briefly sketch the quantum algorithm from Ref. [6]. The presentation here follows from Ref. [119]. Before we discuss this algorithm we introduce the concept of *amplitude amplification*.

Many real world algorithms are probabilistic, i.e., independent runs of the algorithm on the same input will not necessarily give the same output. This is because the algorithm uses some source of randomness during its

execution. Most quantum algorithms are probabilistic owing to the inherent randomness present in quantum mechanics.

Suppose that the job of our probabilistic classical/quantum algorithm is to return one of a specific set of states. Assume that we also have at our disposal an oracle that can identify the members of this set from other states. An example of this would be polynomial root finding. The set of states in this case would correspond to the roots of the polynomial. Our algorithm should return one of the roots of the polynomial and we can verify if an output is a root by plugging it in to the polynomial.

Obviously the algorithm is good only if it can return a state that is a member of this set with high probability. But how high of a success probability is good enough? For practical reasons we would like the probability of success to be a constant. That is, it should be a value independent of the problem size and other parameters in the problem. Any constant value between 0 and 1 would work here. The value $\frac{2}{3}$ is usually used in literature.

But often algorithms won't succeed with constant probability and their success provability will diminish with growing input size. In that case, how can we boost the success probability to the desired level? The classical answer to this question is to repeatedly run the algorithm until we succeed, i.e., till the algorithm outputs a state from the specific set of states that we want. If the algorithm initially had a success probability of p , after $O(\frac{1}{p})$ repetitions we are guaranteed to find the desired state with constant probability.

For quantum algorithms we can do something better. Let U be a quantum algorithm and suppose that we want this algorithm to return a state from the subspace spanned by the orthogonal states, $\{|u_i\rangle\}$. Let P be the projection operator onto this subspace, $P = \sum_i |u_i\rangle\langle u_i|$. The oracle we have is then, $O = I - 2P$. This oracle will mark the states in the desired subspace. The success probability of our algorithm is $p = \langle 0 \dots 0 | U^\dagger P U | 0 \dots 0 \rangle$. In this scenario we can use amplitude amplification to boost the success probability to a constant with only $O(\frac{1}{\sqrt{p}})$ repetitions. This is a quadratic speedup over the classical strategy.

Essentially, amplitude amplification is a generalization of Grover search described in Section II. In Grover search we repeatedly apply the Grover operator, $G = (2|\psi\rangle\langle\psi| - I)O$, where $|\psi\rangle$ is the uniform superposition state. Amplitude amplification uses a more general operator,

$$G_U = U(2|0\rangle\langle 0| - I)U^\dagger O. \quad (56)$$

To get the desired result we apply this to the $U|0 \dots 0\rangle$ state $O(\frac{1}{\sqrt{p}})$ times. Notice that the original Grover search is a specific case of amplitude amplification with $U = H \otimes \dots \otimes H$. In that case, the probability of getting the marked state in $|\psi\rangle$ is $\frac{1}{N}$ so we run the algorithm for $O(\sqrt{N})$ steps. The reader is referred to Ref. [23] for more details on amplitude amplification.

The matrix product verification procedure uses amplitude amplification as its outer loop. The algorithm first splits the full matrix verification problem into smaller matrix verification problems. Then it uses amplitude amplification to search if one of these verifications fail. Each of these smaller verification steps also use a Grover search to look for disagreements. So the complete algorithm uses one quantum search routine nested inside another quantum search routine. This is a common strategy used while designing quantum algorithms to improve query complexity. The full algorithm is sketched below.

The number of qubits and the circuit depth required for this algorithm is too large for it to be successfully implemented on the IBM machines. But at the heart of this algorithm is the Grover search procedure, which we have already discussed and implemented in Section II

Algorithm 6 Matrix product verification [6] [119]**Input:**

- $n \times n$ matrices A, B, C .

Output:

- Verifies if $AB = C$

Procedure:

Step 1. Partition B and C into \sqrt{n} submatrices of size $n \times \sqrt{n}$. Call these B_i and C_i respectively. $AB = C$ if and only if $AB_i = C_i$ for all i .

Step 2. Use amplitude amplification over i on these steps:

Step 2a. Choose a random vector x of dimension \sqrt{n} .

Step 2b. Compute $y = B_i x$ and $z = C_i x$ classically

Step 2c. Verify equation $Ay = z$ by Grover search. Search for a row j such that $(Ay - z)_j \neq 0$

8 GROUP ISOMORPHISM

8.1 Problem definition and background

The *group isomorphism* problem, originally identified by Max Dehn in 1911 [40], is a well-known decision problem in abstract algebra. Simply stated, it asks whether there exists an isomorphism between two finite groups, G and G' . Which, according to the standpoint of group theory, means that they are equivalent (and need not be distinguished). At the end of Section 5 we saw an example of two isomorphic groups, S_2 and A_2 . These two are the same group in terms of how the group operation works on the group elements, but are defined in different ways. More precisely, two groups, (G_1, \cdot) and $(G_2, *)$ are called isomorphic if there is a bijection, $f : G_1 \rightarrow G_2$, between them such that, $f(g_1 \cdot g_2) = f(g_1) * f(g_2)$.

To solve this problem using a quantum algorithm, we assume that each element can be uniquely identified by an arbitrary bit-string label. We also assume that a so-called group *oracle* can be used to return the product of multiple elements. That is, given an ordered list of group-element labels, the oracle will return the product label. In practice, this means that we must be able to construct a quantum circuit to implement $U_a : |y\rangle \rightarrow |ay\rangle$, for any $a \in G$.

In this section, we will focus our attention on the abelian group isomorphism problem, because it can be solved using a generalization of Shor's algorithm [114]. As we saw before, abelian simply means that the operation (\cdot) used to define the group is commutative, such that $a \cdot b = b \cdot a$, for $a, b \in G$. Although Shor's approach is specifically intended to leverage a quantum period-finding algorithm to reduce the time-complexity of factoring, the procedure effectively solves the group isomorphism problem over cyclic groups. Using this relationship, Cheung and Mosca [28] have developed a theoretical quantum algorithm to solve the abelian group isomorphism problem by computing the decomposition of a given group into a direct product of cyclic subgroups.

8.2 Algorithm description

The procedure presented in Algorithm 7 assumes the fundamental theorem of finite abelian groups, that they can be decomposed as a direct sum of cyclic subgroups of prime power order. This decomposition can then be used to test if an isomorphism exists between two groups.

Since the procedure in Algorithm 7 is mostly classical, we shall treat the task of finding the generators of the hidden subgroup in **Step 1** as the most critical for us to explore. This task is commonly referred to as the hidden subgroup problem (HSP). This means that, given a function g that maps a finite group A onto a finite set X , we are asked to find a generating set for the subgroup K . For K to be the so-called *hidden subgroup* of A , we require that g is both constant and distinct on the cosets of K . On a quantum computer, this problem can

Algorithm 7 Decompose(a_1, \dots, a_k, q), of Cheung and Mosca [28]**Input:**

- A generating set $\{a_1, \dots, a_k\}$ of G .
- The maximum order, q , of the generating set.

Output:

- The set of elements g_1, \dots, g_l from group G , with $l \leq k$.

Procedure:

Step 1. Define $g : \mathbb{Z}_q^k \rightarrow G$ by mapping $(x_1, \dots, x_k) \rightarrow g(x) = a_1^{x_1} \cdots a_k^{x_k}$.

Find generators for the hidden subgroup K of \mathbb{Z}_q^k as defined by function g .

Step 2. Compute a set $y_1, \dots, y_l \in \mathbb{Z}_q^k / K$ of generators for \mathbb{Z}_q^k / K .

Step 3. Output the set $\{g(y_1), \dots, g(y_l)\}$.

be solved using a number of operations that is polynomial in $\log|A|$, in addition to one oracle evaluation of the unitary transform $U |a\rangle |h\rangle = |a\rangle |h \oplus g(a)\rangle$. The general procedure needed to complete **Step 1** of algorithm 7 is described in algorithm 8.

Algorithm 8 Solution to the hidden subgroup problem (for finite abelian groups). Based on Ref. [92]**Input:**

- Two quantum registers.
- Elements of the finite abelian group A (or the generating set).
- A function g , such that $g : A \rightarrow X$, with $a \in A$ and $h \in X$.

Output:

- The generating set for the hidden subgroup K .

Procedure:

Step 1. Create initial state.

Step 2. Create superposition between registers.

Step 3. Apply unitary operation (U) for function $g(a)$.

$$\rightarrow \frac{1}{\sqrt{|A|}} \sum_{a \in A} |a\rangle |g(a)\rangle \quad (57)$$

Step 4. Apply inverse Fourier transform.

$$\rightarrow \frac{1}{\sqrt{|A|}} \sum_{l=0}^{|A|-1} e^{2\pi i l a / |A|} |\hat{g}(l)\rangle \quad (58)$$

Step 5. Measure the phase from first register.

$$\rightarrow l / |A| \quad (59)$$

Step 6. Sample K from $l / |A|$.

Like the period-finding approach used in quantum factorization in Section V, Algorithm 8 is heavily based on the concept of phase estimation. Note that the Fourier transform in Eq. 58 represents $a \in A$ indexed by l . The key concept of the procedure is that $|\hat{g}(l)\rangle$ has nearly zero amplitude for all values of l , except those which satisfy




Fig. 23. Basic phase-estimation quantum circuit needed to solve the general hidden subgroup problem in algorithm 8. Here, $|u\rangle$ is an eigenstate of the unitary operator U .

$$|K| = \sum_{h \in K} e^{-2\pi i h l / |A|}, \quad (60)$$

and that knowledge of l can be used to determine both the elements and generating set of K . As discussed by Nielsen and Chuang [92], the final step in algorithm 8 can be accomplished by expressing the phase as

$$\rightarrow e^{2\pi i l a / |A|} = \prod_{i=1}^M e^{2\pi i l_i a_i / p_i}, \quad (61)$$

for $a_i \in \mathbb{Z}_{p_i}$, where p_i are primes, and \mathbb{Z}_{p_i} is the group containing integers $\{0, 1, \dots, p_i - 1\}$ with the operator being addition modulo p_i .

The quantum circuit needed to solve the HSP is schematically illustrated in Fig. 23. This simplified circuit includes steps 1-5 of algorithm 8, and makes it clear that all forms of the HSP (order-finding, period-finding, discrete logarithm, etc.) are extensions of quantum phase estimation.

8.3 Algorithm implemented using Qiskit

Since the generalized group isomorphism problem is somewhat complex, we will focus here on the implementation of the HSP circuit fragment illustrated in Fig. 23. We also chose a specific instance of the HSP: the problem of finding the period of $a \bmod n$. In Fig. 24, the basic outline of the code needed for this specific problem is illustrated using the python-based Qiskit interface.

Like most instances of the HSP, one of the most challenging practical tasks of finding the period of $a \bmod n$ on a quantum computer is the implementation of the oracle. The details of the oracle are not explicitly shown in the Qiskit snippet, but for the required $Ca \bmod 15$ operations, one can simply use the circuits developed by Markov and Saeedi [105]. The code in Fig. 24 also assumes that a function `qft_inv()` will return the gates for an inverse quantum Fourier transform, and that a classical *continued fractions* algorithm can be used to convert the end result (a phase) to the desired integer period.

Although the specific procedure outlined in Fig. 24 can be directly implemented using the IBM Qiskit interface, the resulting QASM code is not expected to lead to accurate results on the IBMX4 (or IBMX5). This is because the generated circuit is long enough for decoherence error and noise to ultimately dominate the measured state. In

```

#####
#----- Finding period (r) of a % N, with N=15 -----#
#####
def findperiod(a, N=15, nqubits1, nqubits2):

    # Create QuantumProgram object, and define registers and circuit
    Q_program = QuantumProgram()
    qr1 = Q_program.create_quantum_register("qr1", nqubits1)
    qr2 = Q_program.create_quantum_register("qr2", nqubits2)
    cr1 = Q_program.create_classical_register("cr1", nqubits1)
    cmod15 = Q_program.create_circuit("cmod15", [qr1, qr2], [cr1])

    # Apply a hadamard to each qubit in register 1
    # and prepare state |1> in register 2
    for j in range(nqubits1): cmod15.h(qr1[j])
    cmod15.x(qr2[nqubits2-1])

    # Loop over qubits in register 1
    for p in range(nqubits1):

        # Calculate next 'b' in the Ub to apply
        # ( Note: b = a^(2^p) % N ).
        # Then apply Ub
        b = pow(a,pow(2,p),N)
        CxModM(cmod15, qr1, qr2, p, b, N, nqubits1, nqubits2)

    # Perform inverse QFT on first register
    qft_inv(cmod15, qr1, nqubits1)

    # Measure each qubit, storing the result in the classical register
    for i in range(n_qr1): cmod15.measure(qr1[i], cr1[i])

```

Fig. 24. Simple implementation of the quantum period-finding algorithm in Qiskit

other words, the physical hardware requires further optimization to reduce the number of gates used between the initial state preparation and the final measurement.

9 QUANTUM RANDOM WALKS

9.1 Problem definition and background

Quantum algorithms for graph properties using the adjacency matrix (as part of an oracle) have been published for minimum spanning tree, shortest path, deciding if a graph is bipartite, detecting cycles, finding subgraphs (such as a triangle), maximal clique, and many more. Each typically involves the use of Grover's search [63] with an oracle constructed from the adjacency matrix.

But for some problems Grover's algorithm is insufficient to achieve optimal query complexity. In such cases, a quantum random walk can sometimes be useful in reducing the query complexity of an algorithm further. An example of this is the quantum algorithm for element distinctness by Ambainis [5]. Additionally, quantum walk algorithms can also be used to search and find graph properties [33, 43, 48, 75, 76, 86]. Quantum random walks can be seen as a quantum mechanical generalization of classical random walks. Quantum random walk algorithms come in two forms, discrete time quantum walks and continuous time quantum walks [75]. The discrete form operates in a step-wise fashion, requiring multiple copies of a set of gates per step. The continuous form uses a transition matrix that is expressed as a Hamiltonian, whose time evolution is then simulated. Quantum random walks can be used to walk a graph [43, 76], search for marked vertices [48], and to solve s-t connectivity [76]. An excellent survey of this approach to quantum search can be found in Ref. [106].

Most quantum algorithms that solve graph problems requires an oracle that knows the properties of the underlying graph. A graph properties oracle can be assembled as a circuit based on the adjacency matrix of the graph and linear algebra transformations. For example, a quantum circuit for finding maximal cliques in a graph with n nodes, requires an oracle workspace of n^2 data qubits and n^2 ancilla qubits (see [131]). Each oracle call requires execution of $6n^2$ Toffoli gates and $2n$ CNOT gates. An oracle such as this can be run on a simulator, but requires too many qubits to run on actual qubit hardware. Quantum algorithms for finding a triangle, quadrilateral, longer cycles, and arbitrary subgraphs [32] typically use the adjacency matrix to create the oracles. Here we will not get into using quantum random walks to solve such problems. Instead we will demonstrate how to implement a simple quantum random walk on a quantum computer.

9.2 Example of a quantum random walk

Quantum random walks or simply quantum walks are quantum analogues of classical random walks and Markov chains. Unlike the continuous time quantum walk, the discrete time quantum walk algorithm requires the use of one or more coin qubits representing the number of movement choices from each graph vertex. These extra coin degrees of freedom are necessary to ensure unitarity of the quantum walk. An appropriate unitary transformation on these coin qubits then acts like the quantum version of a random coin toss, to decide the next vertex for the walker.

Intuitively, the quantum walk is very similar to its classical cousin. In a classical walk, the walker observes some random process, say a coin toss, and decides on his next step conditioned on the output of this random process. So for a classical random walk, the walker is given a probability to make a transition. In a quantum walk, on the other hand, the random process is replaced by a quantum process. This quantum process is the application of the coin operator, which is a unitary matrix. So the next step of the walker is controlled by a complex amplitude rather than a probability. This generalization, from positive numbers to complex numbers, makes quantum walks more powerful than classical random walks.

The full Hilbert space for the discrete quantum walk on a cycle with $N = 2^n$ nodes can then be constructed as follows. We use an n qubit register to represent the nodes of the graph as bit strings. For the cycle every node has only two neighbours, so the coin space only needs a dimension of 2. Hence, only one extra coin qubit is required. The basis vectors of the coin ($|0\rangle$ and $|1\rangle$) will denote the right and left neighbours. So a basis state in the full Hilbert space will have the form $|k, q\rangle$, where k is represents a node on the cycle and q is a single bit representing the coin state.

The quantum walk is then a product of two operators, the shift operator (S) and the coin operator (C). As we mentioned before the coin operator only acts on the coin qubit. The coin operator can be in principle any unitary that mixes the coin states, but here we will use the Hadamard coin which is just the H gate on the coin qubit,

$$C |k, q\rangle = I \otimes H |k, q\rangle = \frac{|k, 0\rangle + (-1)^q |k, 1\rangle}{\sqrt{2}}. \quad (62)$$

The shift operator acts on both the registers. It moves the walker to the left or right depending on the coin state and then flips the coin state,

$$S |k, q\rangle = |k + (-1)^q, q \oplus 1\rangle \quad (63)$$

The quantum walk then proceeds by applying these two operators in alternation. A p step quantum walk is just the operator $(SC)^p$. This type of a walk was first introduced in Ref. [112] and is sometimes referred to as a ‘flip-flop’ quantum walk.

The definition of these operators can change for different types of quantum walk. The coin operator can be a Hadamard gate or a sub-circuit that results in mixing the coin states. The shift operator can be simple as described above or can be a more complicated circuit that selects the next vertex in the path based on the state of the coin. A simple pseudo-code for implementing the quantum walk is given in Algorithm 9.

Algorithm 9 Discrete time quantum walk

Input:

- Two quantum registers. The coin register and the position register.
- Number of steps, T .

Output:

- State of the quantum walk after T steps.

Procedure:

Step 1. Create the initial state. The initial state depends on the application. For instance, in quantum search algorithms, the initial state is the uniform superposition state.

for $0 \leq k < T$ **do**

Step 2a. Apply the coin operator, C , to the coin register.

Step 2b. Apply the shift operator, S . This shifts the position of the walker controlled on the coin state.

end for

Step 3. (Optional) Measure the final state.

9.3 Algorithm implementation using Qiskit on IBM Q

In this section we will implement a simple quantum walk on Qiskit and execute it on both the simulator and `ibmq_vigo`, which is a 5 qubit machine available on IBM Q. We will test the quantum walk on a simple 4 vertex cycle with the vertices labels as given in Fig. 25.




Fig. 25. A graph of 4 nodes in the form of a square is used for the random walk algorithm. The starting vertex is labeled 0. The vertex labels are converted to binary for input into the quantum circuit. The quantum walk algorithm will walk around the graph.

The coin operator in Eq. (62) is just the H gate acting on the coin qubit. The shift operator defined in Eq. (63) is more non-trivial. We can implement it by the circuit given in Fig. 26.




Fig. 26. Quantum circuit for the shift operation on the 4 vertex cycle. The top qubit is the coin qubit.

Running the walk for multiple steps requires us to apply the shift operator circuit many times. So it would be tedious to implement the quantum walk on the IBM Q graphical interface. Instead we can use Qiskit to design the shift operator as a user defined gate and then run the walk for multiple steps using a simple for loop. The Qiskit code for this is given in Fig. 27.

We ran this Qiskit code for 4 steps of the quantum walk. We chose 4 steps since, a simple calculation shows that, starting from $|000\rangle$ and applying $(SC)^4$ will concentrate all the probability to the state $|100\rangle$. This is confirmed by running the Qiskit code on the simulator. But running the same code on `ibm_vigo` gave $|100\rangle$ with only 21.7% probability. The rest of the probability was distributed among the other basis states, but $|100\rangle$ was still the state with the largest probability. This poor performance is due to the circuit having large depth. We can expect to get better results by running the quantum walk for a single step. After a single step, starting from $|000\rangle$, the state of the system is $\frac{|111\rangle + |010\rangle}{\sqrt{2}}$. This is again confirmed by the simulator. Running on `ibm_vigo`, we got $|111\rangle$ with 33.5% probability and $|010\rangle$ with 28.5% probability.

10 QUANTUM MINIMAL SPANNING TREE

10.1 Problem definition and background

A common problem in network design is to find a minimum spanning tree. Suppose we are responsible for maintaining a simple network of roads. Unfortunately, each segment needs repair and our budget is limited. What combination of repairs will guarantee the network remains connected? Fig 28 shows a model of a simple road network as a graph, together with a minimal spanning tree.

Formally, a graph $G = (V, E)$ consists of a set V (the nodes) and a set E consisting of pairs of nodes. A graph is connected if between any two nodes there exists a path. A spanning tree of a connected graph $G = (V, E)$ is the graph $T = (V, E_T)$ where $E_T \subset E$ and T contains no cycles (i.e., there is exactly one path between any two vertices). It is not hard to see that a graph T is a spanning tree if and only if T is connected and has n nodes and $n - 1$ edges. A weighted graph is a graph $G = (V, E, w)$ where w is a map on the edges $w : E \rightarrow \mathbb{R}$. A minimal spanning tree of a graph G is then a spanning tree $T = (V, E_T)$ which minimizes

$$\sum_{e \in E_T} w(e). \quad (64)$$

10.2 Algorithm description

Algorithmically, a graph is usually presented in one of two ways: either as a list of edges or as an adjacency matrix. We consider the case where G is presented as a list of edges. A quantum algorithm for finding a minimal spanning tree of an input graph is given in [45]. This algorithm requires only $O(\sqrt{nm})$ queries where n is the number of nodes and m the number of edges in the graph. Classically, the best algorithms run in time $O(m \log n)$. In particular, this is the time complexity of Borůvka's algorithm [20]. The quantum algorithm combines Borůvka's algorithm together with the quantum search algorithm of Grover [63].

```

from qiskit import QuantumCircuit, QuantumRegister, ClassicalRegister
from qiskit import Aer, execute

n_steps = 4          #Number of steps

#Defining the shift gate

shift_q = QuantumRegister(3)          #3 qubit register
shift_circ = QuantumCircuit (shift_q, name='shift_circ') #Circuit for shift operator
shift_circ.ccx (shift_q[0], shift_q[1], shift_q[2]) #Toffoli gate
shift_circ.cx ( shift_q[0], shift_q[1] ) #CNOT gate
shift_circ.x ( shift_q[0] )
shift_circ.x ( shift_q[1] )
shift_circ.ccx (shift_q[0], shift_q[1], shift_q[2])
shift_circ.x ( shift_q[1] )
shift_circ.cx ( shift_q[0], shift_q[1] )

shift_gate = shift_circ.to_instruction() #Convert the circuit to a gate

q = QuantumRegister (3, name='q') #3 qubit register
c = ClassicalRegister (3, name='c') #3 bit classical register
circ = QuantumCircuit (q,c) #Main circuit
for i in range(n_steps):
    circ.h (q[0]) #Coin step
    circ.append (shift_gate, [q[0],q[1],q[2]]) #Shift step

circ.measure ([q[0],q[1],q[2]], [c[0],c[1],c[2]])

```

Fig. 27. Qiskit code to implement the quantum walk on a 4 vertex cycle.

Borůvka's algorithm builds a collection of disjoint trees (i.e., a forest) and successively merges by adding minimal weight edges. The first two steps of the algorithm are shown in Fig 29. Formally, we have

- (1) Let \mathcal{T} be the collection of n disjoint trees, each consisting of exactly one node from the graph G .
- (2) Repeat:
 - (a) For each tree T_i in \mathcal{T} find the minimal weighted edge, e_i , of G between T_i and the other trees of \mathcal{T} .
 - (b) Merge the trees $\{T_i \cup \{e_i\}\}$ so that they are disjoint: set this new collection to \mathcal{T} .

If there are k trees in \mathcal{T} at a given iteration of Step (2), then the algorithm performs k searches for the respective minimal weighted edges. As the trees are disjoint, we can perform the k searches in one sweep by inspecting

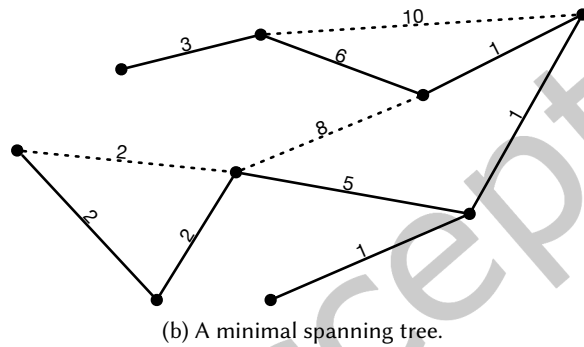



Fig. 28. A graph modeling repair costs of a simple transportation network (a) together with (b) its minimal spanning tree (the solid edges). The sum of the weights of the edges in the minimal spanning tree is 21.




Fig. 29. The first two steps of Borůvka's algorithm. Starting with each node as a distinct tree, find the minimal weighed edge between each tree and the rest of the trees. The direction of the solid edges indicates the edge is the minimal weighted edge for the source node. The components connected by solid edges (disregarding the directions) will form the trees at the start of the second run of step (2) of Borůvka's algorithm

each of the m edges of G once. As there will be at most $\log n$ iterations of Step (2), this results in a running time of $O(m \log n)$. The quantum algorithm takes advantage of the Grover search algorithm, to speed up the searches in Step (2).

In the previous sections we used Grover search to look for a single item in a list of N elements. But the search algorithm will work even if there are M elements in the list that are marked by the oracle. One of these marked elements can then be found using $O(\sqrt{\frac{N}{M}})$ queries to the oracle.

In the algorithm above, we need to find the minimal element of an appropriate list. Clearly this can not be implemented directly as an oracle without actually inspecting each of the list elements. Luckily, there is a simple work around given by Durr et al [45] which involves multiple calls to the Grover algorithm as described in Algorithm 10.

Algorithm 10 Minima finding algorithm

Input:

- A unitary implementation a function F on a list of N elements,

$$U_F |x\rangle |y\rangle = |x\rangle |y \oplus F(x)\rangle.$$

Output:

- $|x^*\rangle$ such that $F(x^*)$ is the minimum of the function over the list.

Procedure:

Step 1. Pick a random j from the list.

for $0 \leq k < T$ **do**

Step 2a. Do Grover search [21] with the oracle for function f_j such that,


$$f_j(i) = \begin{cases} 1 & \text{if } F(i) \leq F(j) \\ 0 & \text{if } F(i) > F(j) \end{cases}$$

Step 2b. Update j with the result of Grover search.

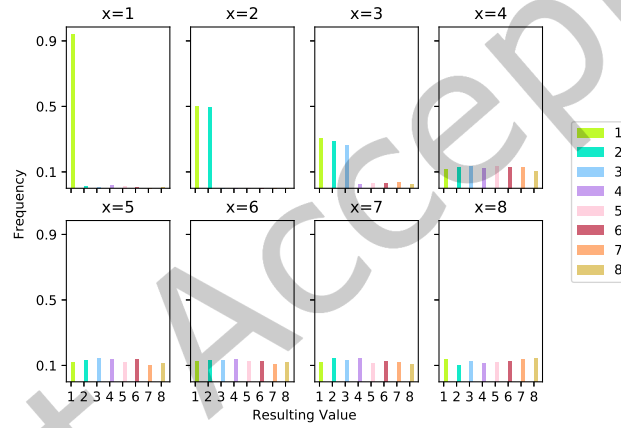
end for

A probabilistic analysis shows that $T = 22.5\sqrt{N} + 1.4 \log_2^2(N)$ suffices to find the minimal element with high probability [46]. The inner loop of the algorithm uses a Grover search routine with potentially multiple marked items. But the number of marked items is not known beforehand. This poses problem as Grover search being a unitary algorithm needs to be stopped exactly at the right number of iterations to give the correct answer. Running the procedure for longer deteriorates the quality of the answer. If the number of marked items is unknown the stopping criterion of the algorithm is also unknown. This problem can be rectified using some extra steps by a technique given in Boyer et al [21]. We have to use this modified version of Grover search in the inner loop.

We did not implement the full algorithm due to space constraints on the IBM computer. Even to successfully implement a minima finding algorithm, at least 6 qubits would be necessary to compare two 3-bit numbers. Therefore we implemented the minima finding algorithm by hard coding the oracle for each of eight possible functions $f_x: \{f_x(i) = 1 \text{ if } F(i) \leq F(x)\}$. The results are shown in Figure 30. The QASM code for implementing $f_2(i) = 1 \text{ if } F(i) \leq F(2)$ required just under 100 lines of code (more than 100 individual gates.) The results, even when using the simulator are not good when $k \geq N/4$ elements are marked. A typical way to get around this is to double the length of the list by adding N extra items which will evaluate to 0 under f_x , which however requires an extra qubit.



(a) IBM Q Implementation



(b) Simulator Implementation

Fig. 30. The results of running 1000 trials of the minima finding algorithm on both (a) the ibmqx4 chip and (b) the IBM simulator to find values less than or equal to the input x .

11 QUANTUM MAXIMUM FLOW ANALYSIS

11.1 Problem definition and background

Network flow problems play a major role in computational graph theory and operations research (OR). Solving the max-flow problem is the key to solving many important graph problems, such as finding a minimum cut set, and finding a maximal graph matching. The Ford-Fulkerson algorithm [52] is a landmark method that defines key heuristics for solving the max flow problem. The most important of these heuristics include the construction of a residual graph, and the notion of augmenting paths. For integer-capacity flows, Ford-Fulkerson has complexity $O(fm)$ for m edges and max flow f . The Edmonds-Karp variant has complexity $O(nm^2)$ for n vertices and m edges. The quantum-accelerated classical algorithm discussed here [7] claims complexity $O(n^{7/6}\sqrt{m})$.




Fig. 31. A simple directed graph representing flows and capacities. Conventionally, the state of the flow problem is indicated by the current flow relative to the capacity on any directed link using the notation f/c .




Fig. 32. The Ford-Fulkerson solution to the max-flow problem in three steps. Each step represents the application of an augmenting path to the previous flow state.

The best classical implementations of the max-flow solver involve several important improvements [47], especially that of using breadth-first search to find the shortest augmenting path on each iteration. This is equivalent to constructing layered subgraphs for finding augmenting paths.

An illustration of the essential method introduced by Ford and Fulkerson can be described using Figures 31 and 32. At each link in the network, the current flow f and the capacity c are shown. Typically, the state of flow on the graph is designated by f/c , with the residual capacity implicitly given by $c - f$. In Figure 31, the initial flow has been set to zero.

The basic steps in the solution to the max-flow problem are illustrated by Figure 32. The algorithm begins on the left by considering the path $[s, v, t]$. Since 2 is the maximum capacity allowed along that path, all the flows on the path are tacitly set to that value. Implicitly, a reverse flow of -2 is also assigned to each edge so that the tacit flow may be “undone” if necessary. Next, the middle of the figure examines the lower path $[s, w, t]$. This path is constrained by a maximum capacity on the edge $[s, w]$ of again 2. Finally, the path $[s, v, w, t]$ is the only remaining path. It can only support the residual capacity of 1 on edge $[s, v]$. We can then read off the maximum flow result at the sink vertex t since the total flow must end there. The maximum flow is seen to be 5.

While this method seems straightforward, without the efficiencies provided by the improvements of Edmonds and Karp, convergence might be slow for integer flows on large graphs, and may not converge at all for real-valued flows. The modification of always choosing the shortest next path in the residual network to augment, is what

makes the algorithm practical. To see this, consider what would have happened if the path $[s,v,w,t]$ had been chosen first. Since augmenting that path blocks the remaining paths, flows would have to be reversed before the algorithm could proceed.

Choosing the shortest path requires performing a breadth-first search whenever new flow values have been assigned to the residual graph. This is equivalent to building a layered set of subgraphs to partition the residual graph. This is the step that leads to the m^2 complexity of Edmonds-Karp, and it is this step that is speeded up in the “quantized” version of the algorithm, leading to a complexity term of \sqrt{m} instead of m^2 .

11.2 Algorithm description

The Quantum algorithm described by Ambainis and Spalek is a “quantized” version of the Edmonds-Karp algorithm, that is, the classical algorithm with quantum acceleration. The key quantum component is a generalized version of Grover’s search algorithm that finds k items in an unsorted list of length L [21]. The algorithm is used in creating a layered subgraph data structure that is subsequently used to find the shortest augmenting path at a given iteration. Like in Section XI, we will be oblivious to the number of marked items Grover’s algorithm is searching for. So once again we have to use techniques from Ref.[21] while performing the search.

Here we will describe how to build a layered graph partition. In a layered graph partition each vertex in the graph is assigned to the i -th layer such that edges of the graph only connect between i -th and $(i + 1)$ -th layers. The key to “quantization” lies in using Grover’s search to build a layered graph partition by computing layer numbers for all vertices. The layers are represented by an array \mathcal{L} indexing the vertices of the graph, and assigning to each element a subgraph layer number. The sink vertex at vertex zero is set to zero. The the algorithm proceeds according to the following pseudo-code described in Algorithm 11.

Algorithm 11 Layered graph partitioning

Input:

- Adjacency information of the graph (Adjacency matrix, list of edges,etc.)
- Source vertex s .

Output:

- \mathcal{L} such that $\mathcal{L}[i]$ is the layer number of the i -th vertex.

Procedure:

Step 1. Set $\mathcal{L}[s] = 0$ and $\mathcal{L}[x] = \infty$ for $x \neq 0$

Step 2. Create a one-entry queue $W = \{s\}$ ($x = 0$)

while $W \neq \phi$ **do**

Step 3a. Take the first vertex x from W .

Step 3b. Find by Grover search all its neighbors y with $\mathcal{L}[y] = \infty$.

Step 3c. Set $\mathcal{L}(y) = \mathcal{L}[x] + 1$, append y into W , and remove x from W

end while

Notice that the oracle for Grover search required for this algorithm is one that marks all the neighbours of x whose layer number is currently set to ∞ . Grover’s search speeds up the layers assignment of the vertices by quickly finding all the entries in the layer array \mathcal{L} that contain the value ∞ . In practical terms, ∞ might simply be the largest value reachable in an n -qubit machine. The generalized Grover search would look for all such values without a priori knowing the number of such values. The size of a circuit required to do a full layered graph partitioning makes it impractical to implement it on the IBM machine. But the heart of the algorithm is Grover search, which we have already implemented earlier.




Fig. 33. An illustration of the MaxCut problem: (left) input Graph G , (middle) a solution of maximum value 5, (right) values of all possible cuts; note that swapping 0/1 for all variables would result in the same cut sizes.

12 QUANTUM APPROXIMATE OPTIMIZATION ALGORITHM

12.1 Problem definition and background

Combinatorial optimization problems are pervasive and appear in applications such as hardware verification, artificial intelligence and compiler design, just to name a few. Some examples of combinatorial optimization problems include Knapsack, Traveling Saleman, Vehicle Routing, and Graph Coloring problems. A variety of combinatorial optimization problems including MaxSat, MaxCut, and MaxClique can be characterized by the following generic unconstrained discrete maximization problem,

$$\begin{aligned} &\text{maximize: } \sum_{\alpha=1}^m C_{\alpha}(x) \\ &x_i \in \{0, 1\} \quad \forall i \in \{1, \dots, n\} \end{aligned} \quad (65)$$

In this generic formulation, there are n binary variables, x , and m binary functions of those variables, $C(x)$, called clauses. The challenge is to find the assignment of x that maximizes the number of clauses that can be satisfied, i.e. that can be evaluated to 1. In case each clause is an OR of literals (positive or negated variables), this is the so-called MaxSat problem, which is NP-Hard in general [79], and is an optimization variant of the well-known satisfiability problem, which is NP-Complete [36]. Hence, solving an instance of Eq. (65) in practice can be computationally challenging, meaning that there is no algorithm which can solve all instances of the problem in time polynomial in their input size (n, m) , unless $P=NP$.

The Maximum Cut Problem. To provide another concrete example of Eq. (65), let us consider the MaxCut problem. As input, the MaxCut problem takes a graph $G = (V, E)$, which is characterized by a set of n nodes V and a set of m undirected edges E . The task is to partition the nodes into two sets, such that the number of edges crossing these sets is maximized. Figure 33 provides an illustrative example, in which a graph with five nodes and six edges is partitioned into two sets that result in a cut of size five. In general, the MaxCut problem is characterized by the following unconstrained discrete maximization problem,

$$\begin{aligned} &\text{maximize: } \sum_{\{u,v\} \in E} \text{XOR}(x_u, x_v) = \sum_{\{u,v\} \in E} x_u + x_v - 2x_u x_v \\ &x_i \in \{0, 1\} \quad \forall i \in n \end{aligned} \quad (66)$$

It is clear that Eq. (66) conforms to the structure of Eq. (65): There is one binary variable $x_i \in \{0, 1\}$ for each node in the graph, indicating which set it belongs to. The objective function consists of one term for each edge in the graph. This term is 0 if the nodes of that edge take the same value and 1 otherwise. Consequently, the optimal solution of (66) will be a maximal cut of the graph G . In foresight, we also reformulate Eq. (66) in terms of spin variables $z_i \in \{-1, +1\}$ as a sum $\sum_{\alpha} C_{\alpha}(z)$, using the linear transformation $x_i = (z_i + 1)/2$:

$$\begin{aligned} \text{maximize: } & \sum_{\{u,v\} \in E} \frac{1}{2}(1 - z_u z_v) = \frac{m}{2} - \sum_{\{u,v\} \in E} \frac{z_u z_v}{2} \\ & z_i \in \{-1, 1\} \quad \forall i \in n \end{aligned} \quad (67)$$

Interestingly, the form of Eq. (67) also highlights that finding a maximal cut of G is equivalent to finding a ground state of the antiferromagnet of G in an Ising model interpretation. We will use Eq. (67) later in the next subsection to formulate a quantum problem Hamiltonian by replacing spin variables z_i with Pauli Z -operators acting on qubit i , $Z_i = \underbrace{Id \otimes \dots \otimes Id}_{i-1} \otimes \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \otimes \underbrace{Id \otimes \dots \otimes Id}_{n-i}$.

Heuristics and approximation algorithms. Given that the decision version of the Maximum Cut problem is NP-hard [54], the Maximum Cut optimization problem is often approached by *heuristic algorithms* [44]. These are algorithms for which one cannot provide guarantees on their performance, but which are often observed to perform well on typical instances. A simple example would be the search for local improvements, where nodes are moved from one set of the cut to the other, if this strictly (or monotonically) increased the cut size. While heuristics may perform very well on most instances, it is generally difficult to tell when they get stuck in a local optimum (or end up in a loop, respectively). As heuristics do not always achieve the perfect solution, one can measure their performance on an individual instance I via an approximation ratio, defined as the algorithm's solution value $A(I)$ versus the optimal value $\text{OPT}(I)$, $\frac{A(I)}{\text{OPT}(I)}$.

A different approach is the design of a polynomial-time r -approximation algorithm A , for which one proves an approximation ratio r over *all* possible input instances I : $r = \min_I \frac{A(I)}{\text{OPT}(I)} < 1$ for maximization problems (or $r = \max_I \frac{A(I)}{\text{OPT}(I)} > 1$ for minimization problems). In the case of randomized (classical or quantum) algorithms, the deterministic value $A(I)$ in the definition of the approximation ratio is replaced by the expected value $\mathbb{E}[A(I)]$ of the solution given by the algorithm. A prominent example for the Maximum Cut problem is the 0.878..-approximation algorithm by Goemans and Williamson [61], which first solves a semi-definite programming relaxation of the problem followed by a randomized hyperplane rounding of the SDP solution. Assuming the unique games conjecture, this is the best possible polynomial-time approximation ratio [77] for MaxCut, which on the other hand is only known to be NP-hard to approximate within $0.941.. + \varepsilon$ [68]. Interestingly, for graphs of maximum degree 3, it is possible to improve on Goemans-Williamson by combining SDP relaxation & rounding with the aforementioned local improvement heuristic, tailored to low-degree nodes, achieving a polynomial-time approximation ratio of 0.932.. [64].

When we assume – as is commonly believed – that quantum computers cannot solve NP-hard problems in polynomial time (i.e. $\text{NP} \not\subseteq \text{BQP}$), then there are two routes for the design of quantum algorithms for combinatorial optimization problems: polynomial speedups for existing approximation algorithms or new quantum heuristics. The Quantum Approximate Optimization Algorithm (QAOA) studied in this section is such a quantum heuristic. Only in rare cases an approximation ratio was shown for QAOA, among them a 0.692..-approximation for the MaxCut problem on 3-regular graphs (where all nodes have degree 3) [50].

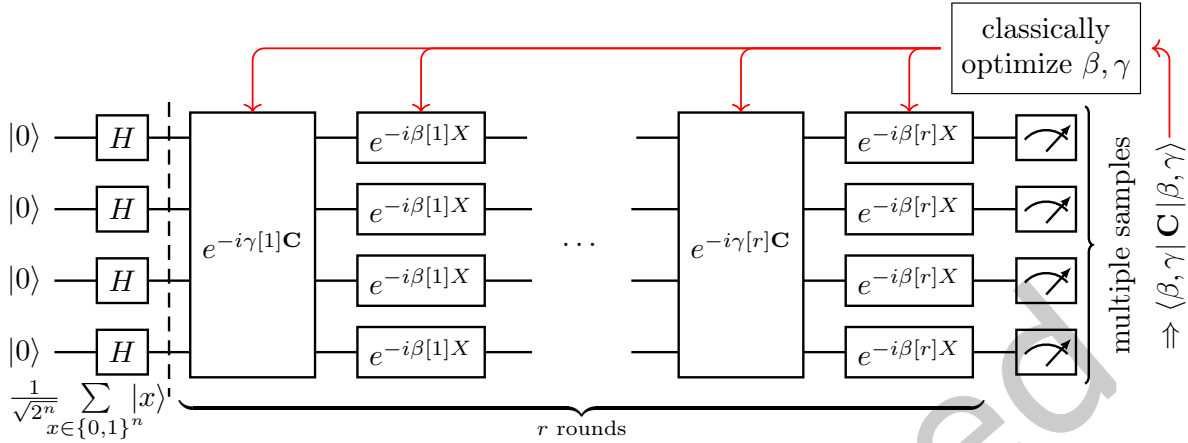


Fig. 34. A high-level view of the hybrid quantum-classical Quantum Approximate Optimization Algorithm: Starting from a uniform superposition over all computational basis states, the quantum subroutine alternately applies the quantum problem Hamiltonian C and a transverse field $B = \sum_j X_j$ for times $\gamma[1 : r]$ and $\beta[1 : r]$, respectively, to prepare the state $|\beta, \gamma\rangle$. Collecting multiple samples from $|\beta, \gamma\rangle$, one can estimate the expectation $\langle \beta, \gamma | C | \beta, \gamma \rangle$ and use a classical optimizer to adjust the parameters β, γ . This eventually results in a state $|\beta, \gamma\rangle$ of high expectation value, from which one can sample solutions of high objective value.

12.2 Algorithm description

The Quantum Approximate Optimization Algorithm (QAOA) as proposed in [50] is a hybrid quantum-classical heuristic algorithm. It leverages gate-based quantum computing for finding candidate solutions to combinatorial optimization problems that have the form of Eq. (65), using a variational circuit with parameters tuned in a classical outer loop (Fig. 34).

To apply the quantum subroutine, the user first translates the clause functions $C_\alpha(z)$ into equivalent quantum clause Hamiltonians C_α , which give rise to the quantum problem Hamiltonian $C = \sum_\alpha C_\alpha$. Additionally, one defines a transverse field mixing Hamiltonian $B = \sum_j X_j$ with Pauli X -operators $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ acting on qubits j . Finally, the user defines a number of rounds $r \geq 1$ and two angles per round, $0 \leq \beta[k] \leq \pi$ and $0 \leq \gamma[k] \leq 2\pi$ for the k -th round. Starting in a uniform superposition $|+\rangle^{\otimes n}$ of all n -bit computational basis states, the quantum subroutine then prepares a state $|\beta, \gamma\rangle$ by alternately applying the Hamiltonians C and B for times $\gamma[k], \beta[k]$. The former operation corresponds to a unitary $e^{-i\gamma[k]C}$ which phases each basis state by an angle proportional to its objective value and proportional to $\gamma[k]$. It can be implemented by applying each clause Hamiltonian C_α on its own. The second operation corresponds to unitaries $e^{-i\beta[k]X_j}$ applied to each qubit j (see pseudocode in Algorithm 12). The goal is to prepare a state $|\beta, \gamma\rangle$ such that one can indeed sample candidate solutions of objective value as close to the optimum value as possible. This, however, heavily depends on the choice of the parameters r, β, γ .

Parameter finding and relation to the adiabatic theorem. To find parameters β, γ that enable sampling of good candidate solutions to Eq. (65) (for a given round r), one usually resorts to optimize β, γ for a large expectation value $\langle \beta, \gamma | C | \beta, \gamma \rangle$, where the optimization is done through a classical optimizer outer loop, see Fig. 34. The reason for this is two-fold: Assuming every solution has an objective value in the range $0, \dots, m$, with high probability $\geq 1 - \frac{1}{m}$ the number of samples sufficient to (i) sample at least one solution of value at least $(\langle \beta, \gamma | C | \beta, \gamma \rangle - 1)$ is $m \log m$ [50], and to (ii) precisely estimate $\langle \beta, \gamma | C | \beta, \gamma \rangle$ is m^3 [35], which can be further reduced if the distribution

Algorithm 12 Quantum subroutine of the Quantum Approximate Optimization Algorithm**Input:**

- Number of rounds of optimization r
- Two size r array of angles, γ and β .
- Hamiltonians C_α corresponding to the clauses of the optimization problem.

Output:

- An approximation to the solution of problem in Eq. (65).

Procedure:

Step 1. Construct the n -qubit uniform superposition state by applying $H^{\otimes n}$ to $|0 \dots 0\rangle$

for $1 \leq k \leq r$ **do**

Step 2a. Apply $\prod_{\alpha=1}^m e^{-i\gamma[k]C_\alpha}$

Step 2b. Apply $\prod_{j=1}^n e^{-i\beta[k]X_j}$

end for

Step 3. We call the state so constructed $|\beta, \gamma\rangle$.

Preparing and measuring $|\beta, \gamma\rangle$ multiple times allows to both estimate the expectation value $\sum_{\alpha=1}^m \langle \beta, \gamma | C_\alpha | \beta, \gamma \rangle$, and to sample an approximate solution to the problem of objective value of at least the expectation minus 1 (with high probability).

of C is concentrated around the expectation, e.g. to m^2 for MaxCut on bounded-degree graphs with a small number r of QAOA rounds [50].

On an ideal noise-free quantum computer, an increase in the number of rounds r should lead to a monotonic increase in the quality of the expectation value $\langle \beta, \gamma | C | \beta, \gamma \rangle$, provided that the chosen angles β, γ are optimal. In fact, any r -round QAOA with $\beta[r] = \gamma[r] = 0$ corresponds simply to a $(r - 1)$ -round QAOA with parameters $\beta[1 : r - 1], \gamma[1 : r - 1]$, hence the optimal expectation value is non-decreasing in the number of rounds. One can also show that with an increasing number of rounds, in the limit $r \rightarrow \infty$ with suitably chosen angles β, γ , the expectation value will converge to the optimum value: $\lim_{r \rightarrow \infty} \max_{\beta, \gamma} \langle \beta, \gamma | C | \beta, \gamma \rangle = \max C(z)$. To this end, we consider a quantum adiabatic algorithm [51] running for time T with time-dependent Hamiltonian $H(t) = (1 - t/T)\mathbf{B} + (t/T)\mathbf{C}$. The starting state of QAOA, $|+\rangle^{\otimes n}$, is also the starting state and unique highest energy eigenstate of $H(0) = \mathbf{B}$. Running the quantum adiabatic algorithm sufficiently slow ($T \gg \text{poly}(n)$) thus results in a highest energy eigenstate of $H(T) = \mathbf{C}$, provided the energy difference between the highest and the second highest eigenstate of $H(t)$ is strictly positive for all $t < T$. This is the case by the Perron-Frobenius theorem whenever \mathbf{C} has only non-negative entries since \mathbf{B} is non-negative irreducible matrix, in particular for combinatorial optimization problems. Quickly alternating between \mathbf{B} and \mathbf{C} with suitable angles β, γ such that $\sum_r \beta[r] + \gamma[r] = T$ gives a discretization (or so-called Trotterization) of the adiabatic algorithm, with improving precision and improving expectation value $\max_{\beta, \gamma} \langle \beta, \gamma | C | \beta, \gamma \rangle$ for increasing r .

However, while QAOA may certainly be looked at as inspired by the quantum adiabatic algorithm, there are problems for which the latter fails for subexponential runtimes while QAOA succeeds even in a single round [50]. Furthermore, for MaxCut on 3-regular graphs, 1-round QAOA was shown to give a 0.692...-approximation. No approximation ratios have been shown for more than 1 round and thus in this regime QAOA is purely heuristic. However, there are known limitations for small/constant-round QAOA approaches: for larger node degrees and sublogarithmic number of rounds $r \in o(\log n)$ the approximation ratio is limited by ≈ 0.834 [25].

So, what strategies of classical optimizers are there to tune the angles β, γ for a given r -round QAOA subroutine? In the original QAOA proposal, an exhaustive search over a fine cartesian grid is suggested for small constant r [50], where the number of grid points is polynomial in n . This works because $\langle \beta, \gamma | C | \beta, \gamma \rangle$ does not have narrow peaks that fall between grid points. Other approaches are necessary for larger r , but we have to be aware

that the parameter landscape is non-convex and thus most classical optimization techniques cannot provide a guarantee for convergence to optimum parameters. Possible optimizers for angle-finding have been extensively studied and are based on techniques such as gradient descent [59, 129], optimal control [22, 135], interpolation of angles for a r -round QAOA based on good angles for a $(r - 1)$ rounds [138], and basin-hopping [94, 130]. For example, basin-hopping starts with random angles, locally optimizes the solution and then randomly perturbs the found angles more significantly to explore a new basin to try to find a better local optima. The use of a quantum-variational-eigensolver is also possible [89, 96].

A closer look at the quantum Hamiltonians. We discuss the translation of the clauses of a combinatorial problem in the form of Equation (65) to quantum clause Hamiltonians by the MaxCut and note generalizations along the way. Recall that we have already discussed how to transform a combinatorial problem $\max \sum_{\alpha} C_{\alpha}(x)$ with binary variables $x_i \in \{0, 1\}$ (see Eq. (66)) into a combinatorial problem $\max \sum_{\alpha} C_{\alpha}(z)$ with spin variables $z_i \in \{-1, 1\}$ (see Eq. (67)) using the linear transformation $x_i = (z_i + 1)/2$. To formulate a quantum clause Hamiltonian C_{α} , we replace in the clause $C_{\alpha}(z)$ each constant 1 with the identity $Id^{\otimes n} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}^{\otimes n}$ and each spin variable z_i with Pauli Z -operators $\begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ acting on qubit i , $Z_i = Id^{\otimes i-1} \otimes Z \otimes Id^{\otimes n-i}$.

We note that the MaxCut problem is particularly advantageous for QAOA for the following reasons: (1) all of the clauses in the objective function have the same structure, hence a circuit implementation has only to be found for one unitary $e^{-iyC_{\alpha}}$; (2) each clause only involves two decision variables, which keeps the structure of C_{α} relatively simple. We note that we have $Z_i \cdot Z_j = Id^{\otimes i-1} \otimes Z \otimes Id^{\otimes j-i} \otimes Z \otimes Id^{\otimes n-j}$. As an example, for the MaxCut problem on the 2-edge path ①–②–③, we get a maximization function $C(z)$ and a quantum problem Hamiltonian C :

$$\begin{aligned} C(z) &= \frac{1}{2}(1 - z_1z_2) + \frac{1}{2}(1 - z_2z_3) = C_{12}(z) + C_{23}(z) \\ \Rightarrow C &= \frac{1}{2}(Id - Z_1Z_2) + \frac{1}{2}(Id - Z_2Z_3) = C_{12} + C_{23} \end{aligned} \quad (68)$$

$$= \frac{1}{2} \left(\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix} \right) + \frac{1}{2} \left(\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} - \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & -1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & -1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & -1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -1 \end{pmatrix} \right)$$

$$C = \begin{matrix} & |000\rangle & |001\rangle & |010\rangle & |011\rangle & |100\rangle & |101\rangle & |110\rangle & |111\rangle \\ \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 2 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 2 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} & \begin{matrix} |000\rangle \\ |001\rangle \\ |010\rangle \\ |011\rangle \\ |100\rangle \\ |101\rangle \\ |110\rangle \\ |111\rangle \end{matrix} \end{matrix} \quad (69)$$

Above, Eq. (68) shows the transformation of the maximization function $C(z)$ into the quantum Hamiltonian C , and Eq. (69) illustrate how the quantum Hamiltonian encodes the inputs and outputs of the different cuts. We note that both clause terms $C_{12} = \frac{1}{2}(Id - Z_1Z_2)$ and $C_{23} = \frac{1}{2}(Id - Z_2Z_3)$ are diagonal matrices, hence they commute and we have $e^{-iy(C_{12}+C_{23})} = e^{-iyC_{12}} \cdot e^{-iyC_{23}}$. Similarly, for the mixing Hamiltonian $B = \sum_j X_j$ the X_j pairwise commute, and we have $e^{-i\beta B} = \prod_j e^{-i\beta X_j}$ as used in Algorithm 12.




Fig. 35. The CNOT connectivity and error rates of the ibmqx2 Computer (left) followed by the Single Edge (center left), Triangle (center right) and four edge Triangle+Edge (right) graphs considered in the proof-of-concept experiments.

Finally, we discuss the influence of the forms of \mathbf{C} and \mathbf{B} on the bounds on the angles γ, β . Since \mathbf{C} is a diagonal matrix with integer eigenvalues, we get that $e^{-i\gamma\mathbf{C}}$ is 2π -periodic in γ and we get $0 \leq \gamma \leq 2\pi$. For \mathbf{B} we have $e^{-i(\pi+\beta)\mathbf{X}} = e^{-i\pi\mathbf{X}} \cdot e^{-i\beta\mathbf{X}} = -\text{Id} \cdot e^{-i\beta\mathbf{X}}$, hence $e^{-i\beta\mathbf{B}}$ is (up to a global phase) π -periodic in β and we get $0 \leq \beta \leq \pi$.

By looking specifically at MaxCut, we can further narrow the angle bounds: we have $e^{-i\pi/2\mathbf{X}} = \begin{pmatrix} 0 & -i \\ -i & 0 \end{pmatrix}$, i.e. increasing β in $e^{-i\beta\mathbf{B}}$ by $\pi/2$ only adds a global phase and swaps 0/1 values of all variables and thus the two cut sets. Hence $\langle \beta, \gamma | \mathbf{C} | \beta, \gamma \rangle$ is even only $\frac{\pi}{2}$ -periodic in β . Using these periodicities, we observe

$$\left\langle \frac{\pi}{2} - \beta, 2\pi - \gamma \middle| \mathbf{C} \middle| \frac{\pi}{2} - \beta, 2\pi - \gamma \right\rangle = \langle -\beta, -\gamma | \mathbf{C} | -\beta, -\gamma \rangle = \langle \beta, \gamma | \mathbf{C} | \beta, \gamma \rangle, \quad (70)$$

where the last equality comes from the fact that \mathbf{C} and \mathbf{B} are real valued and satisfy time reversal symmetry. This cuts our angle search space further in halves, and we can restrict our angle search space to $0 \leq \beta[k] \leq \frac{\pi}{2}$ and $0 \leq \gamma[k] \leq \pi$.

12.3 QAOA MaxCut on ibmqx2

This section investigates the implementation of the QAOA MaxCut algorithm on the ibmqx2 quantum computer (Figure 35). The first challenge is to transform the QAOA algorithm from its mathematical form into a sequence of operations that are available in the IBM Quantum Experience platform. For the sake of convenience we will mention here the gates we will use in the ensuing discussion,

$$U_1(\lambda) = \begin{pmatrix} 1 & 0 \\ 0 & e^{i\lambda} \end{pmatrix}, \quad U_3(\theta, \phi, \lambda) = \begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i(\lambda+\phi)} \cos(\theta/2) \end{pmatrix}, \quad \text{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}.$$

The inner loop of the algorithm first requires the application of the γ angle with the clause Hamiltonians. For the MaxCut Hamiltonian, this can be expanded as follows,

$$e^{-i\frac{\gamma}{2}(Id-Z_1Z_2)} = e^{-i\gamma \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i\gamma} & 0 & 0 \\ 0 & 0 & e^{-i\gamma} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad (71)$$

$$= \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{\text{CNOT}_{12}} \cdot \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & e^{-i\gamma} & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & e^{-i\gamma} \end{pmatrix}}_{Id \otimes U_1(-\gamma)} \cdot \underbrace{\begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}}_{\text{CNOT}_{12}}$$

i.e., we observe that this gate can be implemented as a combination of two CNOT gates and one $U_1(-\gamma)$ gate, as indicated in Figure 36. It is also interesting to note the alternate implementation of this gate in [34], which leverages a different variety of gate operations [116]. We also remark here that the CNOT gates can be interpreted as computing and uncomputing the parity of qubits 1 and 2 inline, with the phase shift of the $U_1(-\gamma)$ gate applied to odd parities. For higher than quadratic terms of Z -operators (such as $Z_iZ_jZ_k$ terms in the MaxE3Lin2 problem [49]), parities can be computed by CNOT gates from all other qubits into a central qubit.



Fig. 36. An IBM Quantum Experience score illustrating an implementation of the MaxCut edge gate (71).

The next term in the loop is the application of the β angle, which is expanded as follows,

$$e^{-i\beta X} = e^{-i\beta \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}} = \begin{pmatrix} \cos(\beta) & -i \sin(\beta) \\ -i \sin(\beta) & \cos(\beta) \end{pmatrix} \quad (72)$$

Careful inspection of the IBM Quantum Experience gates reveals that this operation is implemented by $U_3(2\beta_k, -\pi/2, \pi/2)$. So we need to apply this gate to every qubit in the register. Putting all of these components together, Figure 37 presents an IBM Quantum Experience circuit for implementing the quantum subroutine of QAOA for MaxCut on the “Triangle plus Edge” graph from Figure 35 with parameters,

$$r = 2: \gamma_1 = 0.2 \cdot \pi = 0.628\dots, \quad \beta_1 = 0.15 \cdot \pi = 0.471\dots, \\ \gamma_2 = 0.4 \cdot \pi = 1.256\dots, \quad \beta_2 = 0.05 \cdot \pi = 0.157\dots$$

12.4 A proof-of-concept experiment

With a basic implementation of QAOA for MaxCut in qiskit, a preliminary proof-of-concept study is conducted to investigate the effectiveness of QAOA for finding high-quality cuts in the a) Single Edge, b) Triangle and c) Triangle plus Edge graphs presented in Figure 35. For both a) and b), we ran a numerical grid search for a for a 1-round QAOA with resolution $\frac{1}{8}\pi$ and $\frac{1}{10}\pi$, respectively, using qiskit’s `statevector_simulator`. The statevector simulator allows for exact evaluation of the expectation value $\langle \beta, \gamma | C | \beta, \gamma \rangle$. In both cases, we found parameters β, γ resulting in a state representing an optimum cut. For c), we ran a grid search on the statevector




Fig. 37. An IBM Quantum Experience circuit for 2-round QAOA of MaxCut on the “Triangle plus Edge” graph. We can see the edge gate from Figure 36 replicated for the edges of the triangle between qubits $q[2]$, $q[3]$, $q[4]$ as well as the edge between qubits $q[1]$, $q[2]$. Similarly, the vertical layers of U_3 gates implement $e^{-i\beta B}$.

simulator for both 1- and 2-round QAOA with resolution $\frac{1}{1000}\pi$ and $\frac{1}{20}\pi$ respectively. The parameter landscape for 1 round is given in Fig. 38 and overlaid with a basin-hopping approach.

We then executed the QAOA subroutine for the best-found parameters on *Hardware* by executing the IBM Quantum Experience circuit on the *ibmqx2* device using 4096 shots, and compared the results to a *Simulation* with the same number of shots using the *qasm_simulator*. For both computations we give the expectation/mean of the returned solutions and the probability to sample the maximum cut and contrast these with the values for a *Random* cut. The simulation computation serves to demonstrate the mathematical correctness of the proposed QAOA circuit. The hardware computation demonstrates the viability of the circuit in a deployment scenario where environmental noise, intrinsic bias, and decoherence can have a significant impact on the results. The




Fig. 38. (Heatmap) Parameter landscape for a 1-round QAOA for MaxCut on the Triangle+Edge instance. We observe the symmetry $\langle \beta, \gamma | C | \beta, \gamma \rangle = \langle \frac{\pi}{2} - \beta, 2\pi - \gamma | C | \frac{\pi}{2} - \beta, 2\pi - \gamma \rangle$ derived in Eq. (70). (Overlay) Using a basin-hopping optimizer: A random initialization first explores a basin finding a local maximum; a random perturbation next hops to worse-than-random parameters, but the exploration then finds a global optimum.

random computation serves to demonstrate that the hardware results are better than what one would expect by chance from pure noise.

a) The first shot experiment considers the Single Edge graph from Figure 35 (center left) and implements a 1-round QAOA with the parameters

$$r = 1: \gamma_1 = 0.5 \cdot \pi, \quad \beta_1 = 0.125 \cdot \pi.$$

The results are summarized in Table 4. The simulation results indicate that the proposed score is mathematically sound and the hardware results indicate similar performance to the simulation, with a few additional errors. The random results indicate that both the simulation and hardware perform significantly better than random chance.

Table 4. MaxCut QAOA with one round on a Single Edge.

	Random	Simulation	Hardware
Expected Size of a sampled cut	0.500	1.000	0.950
Probability of sampling a maximum cut	0.500	1.000	0.950

b) The second shot experiment considers the Triangle graph from Figure 35 (center right) with parameters

$$r = 1: \gamma_1 = 0.8 \cdot \pi, \quad \beta_1 = 0.4 \cdot \pi.$$

The results are summarized in Table 5. The simulation results indicate that the proposed circuit is mathematically sound. Even though the QAOA circuit for a Triangle is longer than the QAOA circuit for a Single Edge, the Hardware performance is better, most likely due to the more favourable distribution of the cuts, also notable in Random.

Table 5. MaxCut QAOA with one round on a Triangle.

	Random	Simulation	Hardware
Expected Size of a sampled cut	1.500	1.999	1.904
Probability of sampling a maximum cut	0.750	1.000	0.952

c) The third shot experiment considers the Triangle plus Edge graph from Figure 35 (right). We run QAOA both in a 1-round and a 2-round scenario, implemented with the following parameters, found through numerical grid searches with a resolution of $\pi/1000$ (1-round) and $\pi/20$, respectively (2-round):

$$r = 1: \gamma_1 = 0.208 \cdot \pi, \quad \beta_1 = 0.105 \cdot \pi. \quad r = 2: \gamma_1 = 0.2 \cdot \pi, \quad \beta_1 = 0.15 \cdot \pi, \\ \gamma_2 = 0.4 \cdot \pi, \quad \beta_2 = 0.05 \cdot \pi.$$

The results are summarized in Table 6, the 2-round circuit is shown in Figure 37. Simulation and Hardware outperform Random both on 1-round and 2-round QAOA. However, the gains made by Simulation in 2-round over 1-round QAOA almost vanish on the Hardware. This degradation in performance is likely due to the double length in the circuit, making the experiment more susceptible to gate errors, environmental noise and qubit decoherence.

Table 6. MaxCut QAOA with several rounds on a Triangle plus Edge graph.

	Random	1-round QAOA		2-round QAOA	
		Simul.	Hardw.	Simul.	Hardw.
Expected Size of a sampled cut	2.000	2.720	2.519	2.874	2.570
Probability of sampling a maximum cut	0.375	0.744	0.652	0.895	0.727

Towards practical relevance. When going from the toy problems studied in this subsection towards relevant problems, many problems and issues present themselves. A recent study explored these in detail [65], here we just note a few of these questions:

For example, the studied graphs would likely not fit the (planar) hardware connectivity of an actual device; instead they might come from a random Erdős–Rényi, a random D -regular graph, or even a dense graph such as in the Sherrington-Kirkpatrick model (a MaxCut problem for complete graphs with edge weights). In the case of sparse graphs, one could resort to heuristic compilers implementing SWAP operations to bring qubits representing adjacent nodes together. For dense graphs (or hypergraphs), one can use (generalized) swap networks instead, which swap qubits such that all pairs (sets) are adjacent at some point during its execution [93]. All of these choices affect the quality of sampled solutions in the presence of noise.

This gives rise to the question of up to which problem size and number of rounds the quality of sampled solutions can be increased without gains disappearing due to noise. (A similar question also arises even for idealized quantum devices, if one limits the number of evaluations a classical optimizer can take; in this case a QAOA with fewer rounds might perform better due to the smaller and thus better explored parameter search space [35].)

As of now, quantum advantage for QAOA has not been achieved yet. Current devices are still too noisy and the possible instances too small. Whether QAOA can achieve a significant speedup for combinatorial optimization problems, or whether it enables better provable approximation ratios, remains open.

13 QUANTUM PRINCIPAL COMPONENT ANALYSIS

13.1 Problem definition and background

In data analysis, it is common to have many features, some of which are redundant or correlated. As an example, consider housing prices, which are a function of many features of the house, such as the number of bedrooms, number of bathrooms, square footage, lot size, date of construction, and the location of the house. Often, one is interested in reducing the number of features to the few, most important features. Here, by important, we mean features that capture the largest variance in the data. For example, if one is only considering houses on one particular street, then the location may not be important, while the square footage may capture a large variance.

Determining which features capture the largest variance is the goal of Principal Component Analysis (PCA) [95]. Mathematically, PCA involves taking the raw data (e.g., the feature vectors for various houses) and computing the covariance matrix, Σ . For example, for two features, X_1 and X_2 , the covariance is given by

$$\Sigma = \begin{pmatrix} \mathbf{E}(X_1 * X_1) & \mathbf{E}(X_1 * X_2) \\ \mathbf{E}(X_2 * X_1) & \mathbf{E}(X_2 * X_2) \end{pmatrix}, \quad (73)$$

where $\mathbf{E}(A)$ is the expectation value of A , and we have assumed that $\mathbf{E}(X_1) = \mathbf{E}(X_2) = 0$. Next, one diagonalizes Σ such that the eigenvalues $e_1 \geq e_2 \geq e_3 \geq \dots$ are listed in decreasing order. Again, for the two-feature case, this

becomes

$$\Sigma = \begin{pmatrix} e_1 & 0 \\ 0 & e_2 \end{pmatrix}. \quad (74)$$

Once Σ is in this form, one can choose to keep the features with n -largest eigenvalues and discard the other features. Here, n is a free parameter that depends on how much one wants to reduce the dimensionality. Naturally, if there are only two features, one would consider $n = 1$, i.e., the single feature that captures the largest variance.

As an example, consider the number of bedrooms and the square footage of several houses for sale in Los Alamos. Here is the raw data, taken from www.zillow.com, for 15 houses:

$$\begin{aligned} X_1 &= \text{number of bedrooms} = \{4, 3, 4, 4, 3, 3, 3, 3, 4, 4, 4, 5, 4, 3, 4\} \\ X_2 &= \text{square footage} = \{3028, 1365, 2726, 2538, 1318, 1693, 1412, 1632, 2875, 3564, 4412, 4444, 4278, 3064, 3857\}. \end{aligned} \quad (75)$$

Henceforth, for scaling purposes, we will divide the square footage by 1000 and subtract off the mean of both features. Classically, we compute the covariance matrix and its eigenvalues to be the following:

$$\Sigma = \begin{pmatrix} 0.380952 & 0.573476 \\ 0.573476 & 1.29693 \end{pmatrix}, \quad e_1 = 1.57286, \quad e_2 = 0.105029. \quad (76)$$

We now discuss the quantum algorithm for doing the above calculation, i.e., for finding the eigenvalues of Σ .

13.2 Algorithm description

Before we discuss the algorithm we will provide a quick introduction to the concept of a *density matrix*. Density matrices are used to represent probabilistic mixtures of quantum states. Suppose that there is a quantum system whose state is not known, rather we know that it can be in one of M states, $|\psi_i\rangle$, each occurring with probability p_i . The state of this system is then represented by a density matrix ρ , defined as

$$\rho = \sum_{i=1}^M p_i |\psi_i\rangle \langle \psi_i|. \quad (77)$$

If the state of a system is known (with probability 1) to be $|\psi\rangle$, then the density matrix would just be $|\psi\rangle \langle \psi|$ and the system is said to be in a *pure state*. Otherwise, the system is said to be in a *mixed state*. So the density matrix can be seen as a generalization of the usual state representation with the extra ability to represent a probabilistic mixture of quantum states. From the definition of the density matrix it can be seen that it is a positive semi-definite matrix with unit trace. In fact, any matrix that satisfies these two properties can be interpreted as a density matrix. More details on the definition and interpretation of density matrices are given in the quantum tomography section (Section 19).

Density matrices are clearly more expressive than state vectors as state vectors can only represent pure states. But, even a system in a mixed state can be seen as a part of a larger system that is in a pure state. This process of converting a mixed state into a pure state of an enlarged system is called *purification*. A mixed state of an n qubit system can be purified by adding n more qubits and working with the $2n$ qubit system. Once purified, the joint system of $2n$ qubits will be in a pure state while the first n qubits will still be in the original mixed state. We will not discuss the transformations required to purify a state. Interested readers are referred to Ref. [92] for a complete discussion.

The quantum algorithm for performing PCA presented in Ref. [85] uses the density matrix representation. The algorithm discussed there has four main steps: (1) encode Σ in a quantum density matrix ρ (exploiting the fact that Σ is a positive semi-definite matrix), (2) prepare many copies of ρ , (3) perform the exponential SWAP operation on each copy and a target system, and (4) perform quantum phase estimation to determine the eigenvalues. For

an implementation of this quantum PCA algorithm on a noisy simulator, we refer the reader to Ref. [81], which also gives a short-depth compilation of the exponential SWAP operation.

However, given the constraint of only 5 qubits on IBM's computer, preparing many copies of ρ is not possible. Hence, we consider a simpler algorithm as follows. In the special case where there are only two features, Σ and ρ are 2×2 matrices (one qubit states), and ρ can be purified to a pure state $|\psi\rangle$ on two qubits. Suppose one prepares two copies of $|\psi\rangle$, which uses a total of 4 qubits, then the fifth qubit (on IBM's computer) can be used as an ancilla to implement an algorithm that determines the purity $P := \text{Tr}(\rho^2)$ of ρ . This algorithm was discussed, e.g., in Ref. [71]. It is then straightforward to calculate the eigenvalues of Σ from P , as follows:

$$e_1 = \text{Tr}(\Sigma) * (1 + \sqrt{1 - 2(1 - P)})/2 \quad (78)$$

$$e_2 = \text{Tr}(\Sigma) * (1 - \sqrt{1 - 2(1 - P)})/2. \quad (79)$$

We remark that recently (after completion of this review article), a simpler algorithm for computing purity P was given in Ref. [31]. While the results presented in what follows use the approach in Ref. [71], the approach in Ref. [31] could lead to more accurate results.

As depicted in Fig. 39, this simple algorithm is schematically divided into four steps: (1) classical pre-processing, (2) state preparation, (3) quantifying the purity, and (4) classical post-processing.

In the first step, the classical computer converts the raw data vectors into a covariance matrix Σ , then normalizes this matrix to form $\rho = \Sigma/\text{Tr}(\Sigma)$, then purifies it to make a pure state $|\psi\rangle$, and finally computes the unitary U_{prep} needed to prepare $|\psi\rangle$ from a pair of qubits each initially in the $|0\rangle$ state.

In the second step, the quantum computer actually prepares the state $|\psi\rangle$, or in fact, two copies of $|\psi\rangle$, using U_{prep} , which can be decomposed as follows:

$$U_{\text{prep}} = (U_A \otimes U_B) \text{CNOT}_{AB} (U'_A \otimes \mathbb{1}_B). \quad (80)$$

Note that U_{prep} acts on two qubits, denoted A and B , and CNOT_{AB} is a CNOT gate with A the control qubit and B the target. The single qubit unitaries U_A , U_B , and U'_A can be written in IBM's standard form:

$$\begin{pmatrix} \cos(\theta/2) & -e^{i\lambda} \sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i\lambda+\phi} \cos(\theta/2) \end{pmatrix}, \quad (81)$$

where the parameters θ , λ , and ϕ were calculated in the previous (classical pre-processing) step.

The third step is purity calculation, which makes up the bulk of the quantum algorithm. As shown in Fig. 39, first one does a Hadamard on an ancilla. Let us denote the ancilla as C , while the other four qubits are denoted A , B , A' , and B' . During the state preparation step, qubits A and B were prepared in state $|\psi\rangle$ with the state of A being ρ . Likewise we have the state of A' to be ρ . Next, qubit C is used to control a controlled-SWAP gate, where the targets of the controlled-SWAP are qubits A and A' . Then, another Hadamard is performed on C . Finally, C is measured in the Z basis. One can show that the final expectation value of Z on qubit C is precisely the purity of ρ , i.e.,

$$\langle Z \rangle_C = p_0 - p_1 = \text{Tr}(\rho^2) = P, \quad (82)$$

where p_0 (p_1) is the probability for the zero (one) outcome on C .

The fourth step is classical post-processing, where one converts P into the eigenvalues of Σ using Eqs. (78) and (79).

13.3 Algorithm implemented on IBM's 5-qubit computer

The actual gate sequence that we implemented on IBM's 5-qubit computer is shown in Fig. 40. This involved a total of 16 CNOT gates. The decomposition of controlled-SWAP into one- and two-qubit gates is done first by

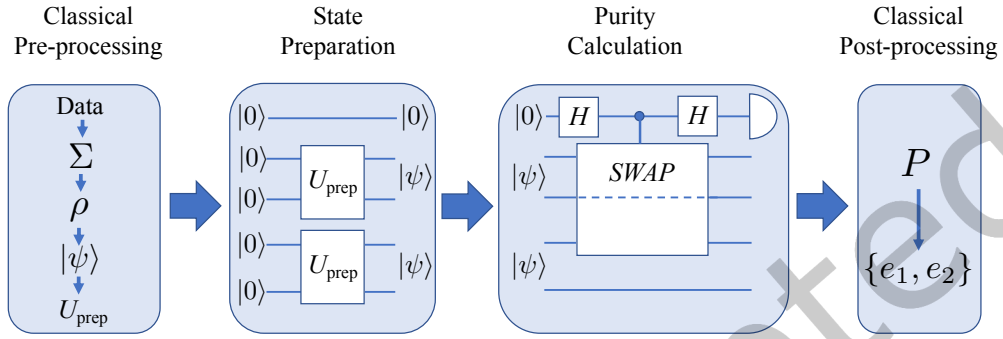


Fig. 39. Schematic diagram for the quantum algorithm for PCA, in the special case of only two features. The first step is classical pre-processing: transforming the raw data into a covariance matrix Σ , then normalizing to compute $\rho = \Sigma/\text{Tr}(\Sigma)$, then purifying ρ to a two-qubit pure state $|\psi\rangle$, and finally determining the unitary U_{prep} needed to prepare $|\psi\rangle$. The second step is to prepare two copies of $|\psi\rangle$ by implementing U_{prep} on a quantum computer. The third step is purity calculation, which is the bulk of the quantum algorithm. This involves doing a Hadamard on an ancilla, which then is used to implement a controlled-SWAP gate on two qubits (from different copies of $|\psi\rangle$), and then another Hadamard on the ancilla. Finally measuring $\langle Z \rangle$ on the ancilla gives the purity $P = \text{Tr}(\rho^2)$. The last step is to classically compute the eigenvalues using Eqs. (78)-(79).

relating it to the Toffoli gate:

$$\text{controlled-SWAP}_{CAB} = (\mathbb{1}_C \otimes \text{CNOT}_{BA})\text{Toffoli}_{CAB}(\mathbb{1}_C \otimes \text{CNOT}_{BA}) \quad (83)$$

and then decomposing the Toffoli gate, as in Ref. [110].

We note that the limited connectivity of IBM's computer played a significant role in determining the algorithm. For example, we needed to implement a CNOT from q[1] to q[2], which required a circuit that reverses the direction of the CNOT from q[2] to q[1]. Also, we needed a CNOT from q[3] to q[1], which required a circuit involving a total of four CNOTs (from q[3] to q[2] and from q[2] to q[1]).

Our results are as follows. For the example given in Eq. (75), IBM's 5-qubit simulator with 40960 trials gave:

$$e_1 = 1.57492, \quad e_2 = 0.102965 \quad (\text{IBM's simulator}). \quad (84)$$




Fig. 40. Actual circuit for quantum PCA implemented on IBM’s 5-qubit simulator and quantum computer. The first three time slots in the score correspond to the state preparation step of the algorithm, and the subsequent time slots correspond to the purity calculation step. Due to connectivity reasons, we chose qubit q[3] as the ancilla and qubits q[1] and q[2] as the targets of the controlled-SWAP operation. We decomposed the controlled-SWAP operation into CNOT gates by first relating it to the Toffoli gate via Eq. (83), and then decomposing the Toffoli gate into CNOT gates [110].

A comparison with Eq. (76) shows that IBM’s simulator essentially gave the correct answer. On the other hand, IBM’s 5-qubit quantum computer with 40960 trials gave:

$$e_1 = 0.838943 + 0.45396i, \quad e_2 = 0.838943 - 0.45396i \quad (\text{IBM's Quantum Computer}). \quad (85)$$

This is a non-sensical result, since the eigenvalues of a covariance matrix must be (non-negative) real numbers. So, unfortunately IBM’s quantum computer did not give the correct answer for this problem.

14 QUANTUM SUPPORT VECTOR MACHINE

Support Vector Machines (SVM) are a class of supervised machine learning algorithms for binary classifications. Consider M data points of $\{(\vec{x}_j, y_j) : j = 1, 2, \dots, M\}$. Here \vec{x}_j is a N -dimensional vector in data feature space, and y_j is the label of the data, which is $+1$ or -1 . SVM finds the hyperplane $\vec{w} \cdot \vec{x} + b = 0$ that divides the data points into two categories so that $\vec{w} \cdot \vec{x}_j + b \geq 1$ when $y_j = +1$ and $\vec{w} \cdot \vec{x}_j + b \leq -1$ when $y_j = -1$, and that is maximally separated from the nearest data points on each category. Least Squares SVM (LS-SVM) is a version of SVM [122]. It approximates the hyperplane finding procedure of SVM by solving the following linear equation:

$$\begin{bmatrix} 0 & \vec{1}^T \\ \vec{1} & \mathbf{K} + \gamma^{-1}\mathbf{1} \end{bmatrix} \begin{bmatrix} b \\ \vec{\alpha} \end{bmatrix} = \begin{bmatrix} 0 \\ \vec{y} \end{bmatrix}. \quad (86)$$

Here \mathbf{K} is called the kernel matrix of dimension $M \times M$, γ is a tuning parameter, and $\vec{\alpha}$ forms the normal vector \vec{w} where $\vec{w} = \sum_{j=1}^M \alpha_j \vec{x}_j$. Various definitions for the kernel matrix are available, but the quantum SVM [102] uses linear kernel: $K_{ij} = \vec{x}_i \cdot \vec{x}_j$. Classically, the complexity of the LS-SVM is $\mathcal{O}(M^2(M + N))$.

The quantum version of SVM performs the LS-SVM algorithm using quantum computers [102]. It calculates the kernel matrix using the quantum algorithm for inner product [84] on quantum random access memory [60], solves the linear equation using a quantum algorithm for solving linear equations [60], and performs the classification of a query data using the trained qubits with a quantum algorithm [102]. The overall complexity of the quantum SVM is $\mathcal{O}(\log NM)$.

The algorithm is summarized below:

Algorithm 13 Quantum SVM [102]**Input:**

- Training data set $\{(\vec{x}_j, y_j) : j = 1, 2, \dots, M\}$.
- A query data \vec{x} .

Output:

- Classification of \vec{x} : +1 or -1.

Procedure:

Step 1. Calculate kernel matrix $K_{ij} = \vec{x}_i \cdot \vec{x}_j$ using quantum inner product algorithm [84].

Step 2. Solve linear equation Eq. (86) and find $|b, \vec{\alpha}\rangle$ using a quantum algorithm for solving linear equations [60] (training step).

Step 3. Perform classification of the query data \vec{x} against the training results $|b, \vec{\alpha}\rangle$ using a quantum algorithm [102].

The inner product calculation to compute the kernel matrix cannot be done reliably in the currently available quantum processors. The other important part of the algorithm, which is linear system solving, can be quantized and has been dealt with in Section IV.

15 QUANTUM SIMULATION OF THE SCHRÖDINGER EQUATION

15.1 Problem definition and background

The Schrödinger's equation describes the evolution of a wave function $\psi(x, t)$ for a given Hamiltonian \hat{H} of a quantum system:

$$i\hbar \frac{\partial}{\partial t} \psi(x, t) = \hat{H} \psi(x, t) = \left[\frac{\hbar^2 \hat{k}^2}{2m} + V(\hat{x}) \right] \psi(x, t), \quad (87)$$

where the second equality illustrates the Hamiltonian of a particle of mass m in a potential $V(x)$. Simulating this equation starting with a known wave function $\psi(x, 0)$ provides knowledge about the wave function at a given time t_f and allows determination of observation outcomes. For example, $|\psi(x, t_f)|^2$ is the probability of finding a quantum particle at a position x at time t_f .

Solving the Schrödinger's equation numerically is a common approach since analytical solutions are only known for a handful of systems. On a classical computer, the numerical algorithm starts by defining a wave function on a discrete grid $\psi(x_i, 0)$ with a large number of points $i \in [1, N]$. The form of the Hamiltonian, Eq. (87), allows one to split the system's evolution on a single time step Δt in two steps, which are easy to perform:

$$\psi(x_i, t_{n+1}) = e^{-iV(x_i)\Delta t} QFT^\dagger e^{-ik^2\Delta t} QFT \psi(x_i, t_n), \quad (88)$$

where we have assumed that $\hbar = 1$ and $m = \frac{1}{2}$. And QFT and QFT^\dagger are the quantum Fourier transform and its inverse. The quantum state evolution thus consists of alternating application of the phase shift operators in the coordinate and momentum representations. These two representation are linked together by the Fourier Transformation as in the following example of a free space evolution of a quantum particle:

$$\psi(x_i, t_f) = QFT^\dagger e^{-ik^2 t_f} QFT \psi(x_i, 0), \quad (89)$$

where $V(x) = 0$ for a free particle.

We now discuss the quantum simulation of the Schrödinger's equation similar to the one discussed in [14], [118] that provides the wave function of the system at a given time t_f . Finding a proper measurement on a quantum simulator that reveals information about the quantum system will however be left out of the discussion. $|\psi(x, t_f)|^2$ will be the only information we will be interested in finding out.

15.2 Algorithm description

A quantum algorithm that performs a quantum simulation of one dimensional quantum systems was presented in [14]. The procedure is outlined in Algorithm 14.

Algorithm 14 Quantum simulation of Schrödinger equation [118], [14]

Input:

- Initial wave function
- Time step size, Δt , and the number of time steps, T .
- The ability to apply phase shifts in the computational basis.
- The potential function V .

Output:

- Final wave function at time $t_f = T\delta t$ when evolved using the Schrödinger equation with the potential V .

Procedure:

Step 1. Encode the wave function on a N -point grid in a quantum state of $n = \log_2(N)$ qubits. The value of this discretized wavefunction on a grid point is equal to the value of the original wave function at the same point. The constant of proportionality must then be calculated by renormalizing the discretized wavefunction.

for $1 \leq j \leq T$ **do**

Step 2a. Apply the Quantum Fourier Transform (QFT) to go to the momentum representation.

Step 2b. Apply a diagonal phase shift of the form $|x\rangle \rightarrow e^{-ix^2\Delta t} |x\rangle$ in the computational basis.

Step 2c. Apply the inverse Quantum Fourier Transform to come back to the position representation.

Step 2d. Apply a phase shift of the form $|x\rangle \rightarrow e^{-iV(x)\Delta t} |x\rangle$.

end for

Step 3. Measure the state in the computational basis.

Figure 41 shows the following stages of the algorithm. The implementation of QFT was discussed in Section IV. Implementing phase shifts corresponding to arbitrary functions can be done using a series of controlled Z gates or CNOT gates [14]. Repeating the final measurement step over many independent runs will let us estimate the probabilities $|\psi(x, t_f)|^2$. We will now consider a 2-qubit example of the quantum simulation algorithm in the case of a free particle, $V(x) = 0$.

Our initial wave function is a Π -function (a rectangular wave), which has $\{0, 1, 1, 0\}$ representation on a 2^n -point grid for $n = 2$ qubits. Its representation by the state of the qubits is proportional to $|0, 1\rangle + |1, 0\rangle$, which can be prepared by constructing the Bell state (see Fig. 1) and applying the X gate to the first qubit.

We define the 2-qubit QFT as $QFT = \text{SWAP}_{12} H_2 C_2 \left[\mathcal{P}_1 \left(\frac{\pi}{2} \right) \right] H_1$, where $C_2 \mathcal{P}$ is a phase operator controlled from the second qubit. This transformation applies phase shifts to the probability amplitudes of the qubit states similar to the ones applied by the classical FFT to the function values. Hence, the resulting momentum representation is identical to the classical one in a sense that it is not centered around $k = 0$, which can be easily remedied by a single X_1 gate.

The momentum encoding adopted in this discussion is $k = -\frac{1}{2} \sqrt{\frac{\phi}{\Delta t}} \left(1 + \sum_{k=1}^n 2^{n-k} Z_k \right)$, where ϕ is a characteristic phase shift experienced by the state on a time step Δt . In this representation $-ik^2\Delta t$ phase shift contains one and two qubit contributions that commute with each other and can be individually implemented. The one qubit phase shift gate has a straightforward implementation but the two qubit phase shift gate requires an ancillary qubit according to Ref. [92], which results in a three qubit implementation on a quantum computer.

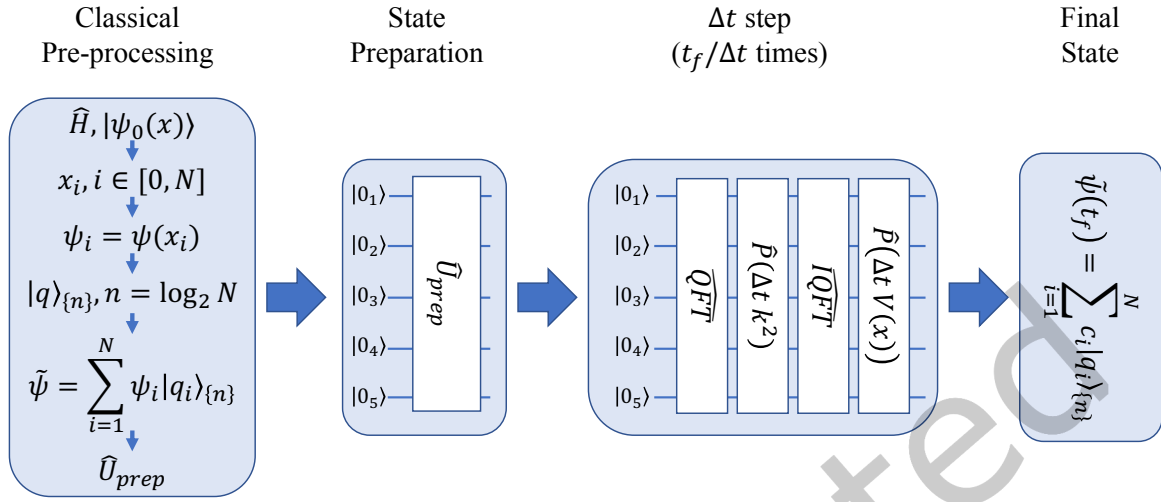


Fig. 41. The quantum simulation of the Schrödinger's equation. The first stage is a classical pre-processing that encodes the wave function to available qubits and derives a state preparation operator that takes an all-zero state of a quantum computer to a desired state. The second stage prepares an initial state by implementing the state preparation operator \hat{U}_{prep} on a quantum computer. The third stage is an iterative update looped over Δt steps based on the operator splitting method.




Fig. 42. The quantum circuit implementation of a 2-qubit algorithm that solves the Schrödinger's equation on a quantum computer. The initial state preparation is followed by the Quantum Fourier Transform and centering of the momentum representation. The single qubit phase shift transformations are followed by the two-qubit phase shift transformation that uses an ancillary qubit $q[0]$. The inverse Quantum Fourier Transform preceded by removing the centering operation completes the circuit and returns the wave function to the coordinate representation.

This implementation is captured in Fig 42 where removing the centering of the momentum representation and the inverse QFT have been added in order to return to the coordinate representation.

15.3 Algorithm implemented on IBM's 5-qubit computer

The implementation in Fig. 43 takes into account the topology of the chip and the availability of the gates such as $U1$ and $U2$. Finally, it performs a consolidation of the single qubit gates in order to reduce the number of physical operations on the qubits.

The circuit in Fig. 43 was run on the `ibmqx4` quantum chip, where the maximum number of executions in a single run is 2^{10} . The probabilities of observing qubit states in the computational basis was measured for $\phi = 0$, $\phi = \pi/2$, $\phi = \pi$, $\phi = 3\pi/2$ and $\phi = 2\pi$. We expect that as ϕ increases from 0 to π the wave function evolves from




Fig. 43. The quantum circuit implementation of a 2-qubit algorithm that solves the Schrödinger’s equation on the ibmqx4 quantum computer.

a Π -function to a uniform function to a function peaked at the ends of the interval. The consecutive increase returns the wave function back to the Π -function.

We started with the $\phi = 0$ case that should have reproduced our initial state with ideal probabilities of $\{0, 0.5, 0.5, 0\}$. However, the observed probabilities were $\{0.173, 0.393, 0.351, 0.084\}$. Thus it was surprising to see that the $\phi = \pi/2$ case was very close to expected probability of 0.25 with the observed values of $\{0.295, 0.257, 0.232, 0.216\}$. This surprise was however short lived as the $\phi = \pi$ case has reverted back large errors for observed probabilities: $\{0.479, 0.078, 0.107, 0.335\}$. The final two case had the following observed probabilities $\{0.333, 0.248, 0.220, 0.199\}$ and $\{0.163, 0.419, 0.350, 0.068\}$ respectively.

16 GROUND STATE OF THE TRANSVERSE ISING MODEL

In this section the ground state of the transverse Ising model is calculated using the variational quantum eigenvalue solver, and the result is compared to the exact results. This is a hybrid method that uses alternating rounds of classical and quantum computing.

In the previous section we saw how to simulate the evolution of a single quantum particle. But often, real world phenomena are dependent on the interactions between many different quantum systems. The study of many-body Hamiltonians that model physical systems is the central theme of condensed matter physics (CMP).

Many-body Hamiltonians are inherently hard to study on classical computers as the dimension of the Hilbert space grows exponentially with the number of particles in the system. But using a quantum computer we can study these many-body systems with less overhead as the number of qubits required only grows polynomially.

16.1 Variational quantum eigenvalue solver

A central task in CMP is finding the ground state (lowest energy eigenstate) of a given Hamiltonian, \mathcal{H} ,

$$\mathcal{H}|\Psi\rangle = E_g|\Psi\rangle. \tag{90}$$

Studying the ground state gives us information about the low temperature properties of the system. Once we know $|\Psi\rangle$, we can deduce the physical properties from the wave function. In this section, we will describe how to use IBM Q to find the ground state energy of the transverse Ising model. We will not be using the ibmqx4 in this section. This is because the algorithm we use will require many rounds of optimization. Each round requires us to run a circuit on the quantum computer followed by a classical optimization step on a local machine. This process can be automated easily using Qiskit. But the long queuing times in IBM Q makes repeated runs on the quantum processor impractical.

To find the eigenvalue of a Hamiltonian, we could use the quantum phase estimation algorithm that was discussed in Section IV. To do this we need the ability to perform controlled operations with the unitary $U = \exp(-i\mathcal{H}\delta t/\hbar)$, where δt is the time step. Then, by preparing different initial states $|\psi_i\rangle$ and repeating the phase estimation many times one can obtain, in principle, the whole spectrum of the eigenvalues and the

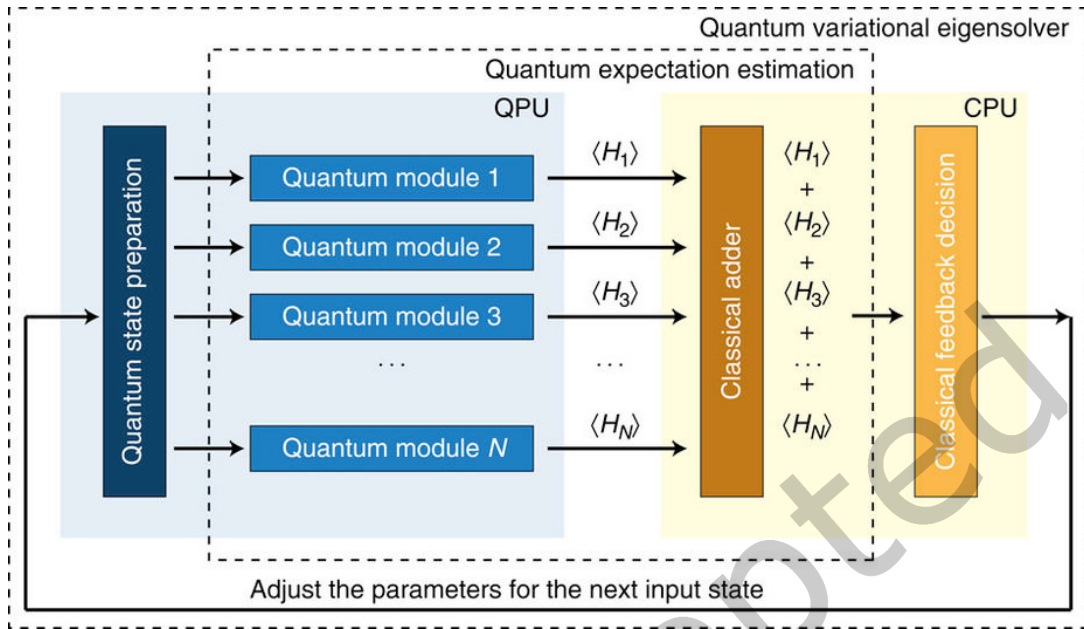


Fig. 44. Schematic view of the implementation of the variational quantum eigensolver using a hybrid classical and quantum circuit. The figure is adopted from Ref. [96].

corresponding eigenwave functions. For a general Hamiltonian, however, the implementation of a controlled U may be not straightforward. For realistic problems, the quantum phase estimation circuits have large depth. This requires qubits with long coherence times, which are not available at the time of writing. For CMP problems, we are mainly interested in the lowest eigenvalue for most cases.

To overcome these limitations, we use the recently developed variational quantum eigensolver (VQE) [89, 96]. The basic idea is to take the advantages of both quantum and classical computers, as shown in Fig. 44. It allocates the classically easy tasks to classical computers and the other tasks to quantum computers. The algorithm is summarized as follows:

- (1) Prepare a variational state $|\psi(\theta_i)\rangle$ with parameters θ_i . For an efficient algorithm, the number of variational parameters should grow linearly with the system size.
- (2) Calculate the expectation value of the Hamiltonian using a quantum computer, $E = \langle \psi | \mathcal{H} | \psi \rangle / \langle \psi | \psi \rangle$.
- (3) Use classical nonlinear optimizer algorithms to find new optimal θ_i . In this report, we will use the relaxation method $\tau_0 \partial_t \theta_i = -\partial E / \partial \theta_i$, where τ_0 is a parameter to control the relaxation rate.
- (4) Iterate this procedure until convergence.

VQE has the following advantage: For most CMP problems, where the interaction is local, we can split the Hamiltonian into a summation over many terms. This means that we can parallelize the algorithm to speed up the computation. The quantum expectation calculations for one term in the Hamiltonian are relatively simple, thus no long coherence times are not required. On the other hand, VQE also has limitations. Because of its variational nature, the trial wave function needs to be prepared carefully. This usually requires physical insights into the problem. The ground state eigenvalue and eigenwave function are biased by the choice of the trial wave functions. In addition, VQE requires communications between classical and quantum computers, which could be a bottleneck for the performance.




Fig. 45. Schematic view of the quantum phases described by the transverse Ising model. The arrows represent the spin configuration in the ordered and disordered phases.

16.2 Simulation and results

We use VQES to find the ground state of the transverse Ising model (TIM) defined by

$$\mathcal{H} = -\sum_i \sigma_i^z \sigma_{i+1}^z - h \sum_i \sigma_i^x, \quad (91)$$

where σ^z , σ^x are Pauli matrices and h is the external magnetic field. Let us first review briefly the physical properties of this Hamiltonian. This Hamiltonian is invariant under the global rotation of spin along the x axis by π , $R_x \mathcal{H} R_x^\dagger = \mathcal{H}$, where $R_x(\pi)$ is the rotation operator

$$R_x \sigma^x R_x^\dagger = \sigma^x, \quad R_x \sigma^z R_x^\dagger = -\sigma^z. \quad (92)$$

The TIM has two phases: When the transverse field h is small, the spins are ordered ferromagnetically and the rotational symmetry associated with R_x is broken. In the ordered phase, the quantum expectation value $\langle \sigma^z \rangle \neq 0$. As h is increased, there is a quantum phase transition from the ordered phase to the disordered phase where $\langle \sigma^z \rangle = 0$, as the rotational symmetry is restored. The phase diagram is shown schematically in Fig. 45.

Using the phase diagram as a guide, first we propose a product state as a trial wave function. The wave function can be written as

$$|\psi_i(\theta_i)\rangle = \prod_i U(\theta_i) |0_i\rangle. \quad (93)$$

Here $U(\theta_i)$ is the unitary operation which describes the spin rotation along the y axis by an angle θ_i ,

$$U(\theta_i) = \begin{pmatrix} \cos(\theta_i/2) & -\sin(\theta_i/2) \\ \sin(\theta_i/2) & \cos(\theta_i/2) \end{pmatrix},$$

where θ_i are the variational parameters. Here we have used the Bloch sphere representation for a qubit state. For the TIM, we calculate the expectation value of

$$E_{J,i} = -\langle \psi | \sigma_i^z \sigma_{i+1}^z | \psi \rangle, \quad E_{Z,i} = -\langle \psi | \sigma_i^x | \psi \rangle. \quad (94)$$

The quantum circuit to perform the preparation of the state and calculation of the expectations value are shown in Fig. 46(a) and Fig. 46(b). We have

$$E_{J,i} = -[P(q_i = 0) - P(q_i = 1)][P(q_{i+1} = 0) - P(q_{i+1} = 1)], \quad (95)$$

$$E_{Z,i} = -[P(q_i = 0) - P(q_i = 1)], \quad (96)$$

where $P(q_i = 0, 1)$ is the measured probability for the qubit q_i in the $|0\rangle$ or $|1\rangle$ state. As we mentioned before, the communication bottleneck prevented us from implementing this on `ibmqx4`. We ran the code using the quantum simulator in Qiskit. The comparison of the results obtained from quantum simulation and analytical results are shown in Fig. 47. Our trial wave function works very well in the ordered phase, but the simulation results deviate




Fig. 46. Quantum circuits to prepare the trial wave-functions. The single qubit unitaries in the text can be implemented using available gates in IBM Q. The first two circuits prepare unentangled trial states. Circuit (a) can be used to measure $\langle \psi | \sigma_2^z \sigma_3^z | \psi \rangle$. Circuit (b) can be used to measure the $\langle \psi | \sigma_3^x | \psi \rangle$. Circuit (c) prepares the entangled trial state.

from the exact solution in the quantum phase transition region. This discrepancy is caused by the fact that we have neglected the quantum entanglement in our trial wave function.

In a second set of experiments, we use a trial wave function that includes quantum entanglement. Because of the symmetry, $|\Psi_i(\theta_i)\rangle$ and $R_x(\pi)|\Psi_i(\theta_i)\rangle$ are two degenerate wave functions with the same energy. The trial wave function can be written as a linear superposition of these two degenerate wave functions

$$|\psi_i(\theta_i)\rangle = \alpha|\Psi_i(\theta_i)\rangle + \beta R_x(\pi)|\Psi_i(\theta_i)\rangle. \quad (97)$$

The first step is to prepare $|\psi_i(\theta_i)\rangle$ using quantum circuit. To prepare an arbitrary state in a quantum circuit is not trivial as it requires of the order of 2^n CNOT gates, where n is the number of qubits [98]. The state in Eq. (97) can be prepared easily using the circuit in Fig. 46(c). Here we consider 4 spins. The first $U_0(\theta, \phi)$ operation transforms the state into

$$|0000\rangle \rightarrow e^{i\phi} \sin(\theta/2)|1000\rangle + \cos(\theta/2)|0000\rangle.$$

The first CNOT transforms the state into

$$e^{i\phi} \sin(\theta/2)|1100\rangle + \cos(\theta/2)|0000\rangle.$$

The second CNOT transforms the state into

$$e^{i\phi} \sin(\theta/2)|1110\rangle + \cos(\theta/2)|0000\rangle.$$

The third CNOT transforms the state into

$$e^{i\phi} \sin(\theta/2)|1111\rangle + \cos(\theta/2)|0000\rangle.$$

Finally we apply $U(\theta_i)$ rotation and we obtain the desired state in Eq. (97). Here

$$U_0(\theta, \phi) = \begin{pmatrix} \cos(\theta/2) & -\sin(\theta/2) \\ e^{i\phi} \sin(\theta/2) & e^{i\phi} \cos(\theta/2) \end{pmatrix}.$$

We then use VQES to find the ground state energy. As can be seen in Fig. 48, the new trial function nearly reproduces the exact results in the whole magnetic field region and improves upon the product state trial function.

17 QUANTUM PARTITION FUNCTION

17.1 Background on the partition function

Calculation or approximation of the partition function is a sub-step of inference problems in Markov networks [78]. Even for small networks, this calculation becomes intractable. Therefore an efficient quantum algorithm for the




Fig. 47. Comparison of the ground state energy (a) and average magnetization (b) $M_x = \langle \psi | \sum_i \sigma_i^x | \psi \rangle / N$ obtained by using the trial wave functions in Eq. (93) and the exact results. Here we have used the periodic boundary condition. The simulations are run both on the quantum simulator (black symbols) and classical computers (red symbols). The mean-field results (blue line) are also displayed for comparison.

partition function would make many problems in graphical model inference and learning tractable and scalable; the same holds for other problems in computational physics [8, 55, 57, 58].

The partition function is of particular interest for calculating probabilities from graphical models such as Markov random fields [78]. For this article, we consider the graphical model form known as the Potts model. Let $\Gamma = (E, V)$ be a weighted graph with edge set E and vertex set V and $n = |V|$. In the q -state Potts model, each vertex can be in any of q discrete states. The Potts model is a generalization of the classical Ising model. In the classical Ising model $q = 2$, whereas in the Potts model $q \geq 2$. The edge connecting vertices i and j has a weight J_{ij} which is also known as the interaction strength between corresponding states. The Potts model Hamiltonian

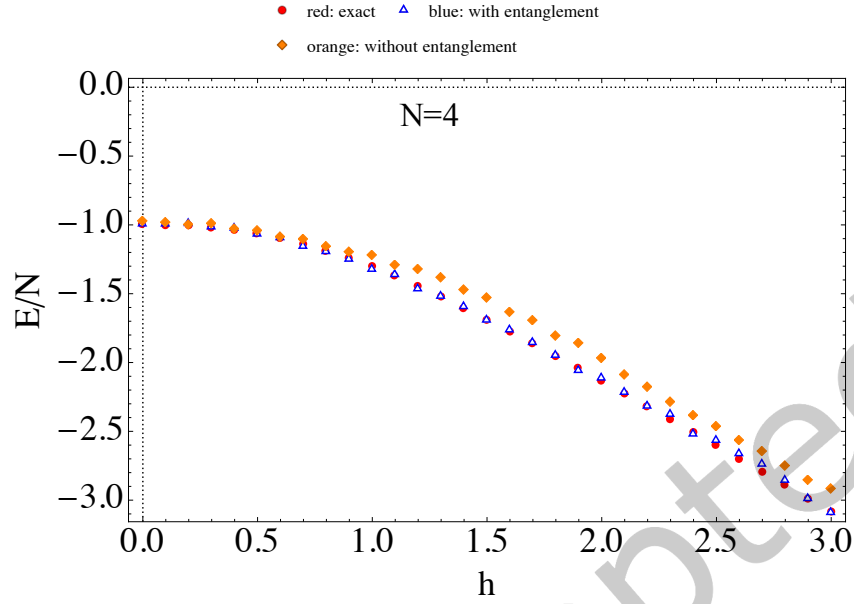


Fig. 48. (color online) Comparison of the ground state energy obtained by using the trial wave functions in Eqs. (93) and (97) and the exact result. Here we have used the periodic boundary condition. The number of spins is 4.

for a particular state configuration $\sigma = (\sigma_1, \dots, \sigma_n)$ is

$$H(\sigma) = - \sum_{i \sim j} J_{ij} \delta_{\sigma_i, \sigma_j}, \quad (98)$$

where $i \sim j$ indicates that there exists an edge between vertices i and j ; and where $\delta_{\sigma_i, \sigma_j} = 1$ if $\sigma_i = \sigma_j$ and 0 otherwise.

The probability of any particular configuration being realized in the Potts model at a given temperature, T , is given by the Gibbs distribution:

$$P(\sigma) = \frac{1}{Z} e^{-\beta H(\sigma)}, \quad (99)$$

where $\beta = 1/(k_B T)$ is the inverse temperature in energy units and k_B is the Boltzmann constant. The normalization factor, Z , is also known as the *partition function*:

$$Z = \sum_{\{\sigma\}} e^{-\beta H(\sigma)}, \quad (100)$$

where $\{\sigma\}$ means the full set of all possible state configurations. There are q^n possible state configurations, and so this is a sum over a large number of items and is generally intractable as well as difficult to approximate. The calculation of the partition function is #P-hard (i.e., it is a counting problem which is at least as hard as the NP-hard class of decision problems). There is no known fully polynomial randomized approximation scheme (fpras), and it is unlikely that there exists one [58].




Fig. 49. (a) Simple example with (b) the enumeration of state configurations and the value of the Hamiltonian for a fully-connected 3-vertex Ising model ($q = 2$ Potts model)

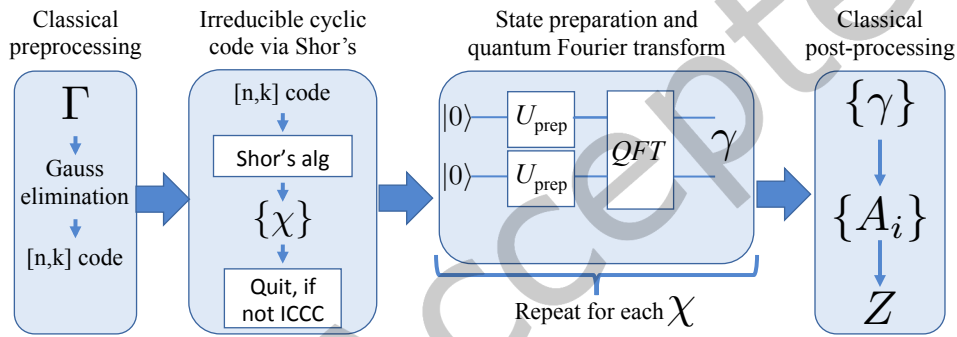


Fig. 50. Overview of the quantum partition function algorithm.

17.2 A simple example

We give a small example with a graph of $n = 3$, $V = \{a, b, c\}$, with edges between all pairs of vertices for three total edges, pictured in Figure 49a, and we use $q = 2$ for binary states on each vertex. To demonstrate the calculation of the partition function, we first enumerate the configurations as shown in Fig. 49b.

We plug the value of the Hamiltonian for each of the q^n configurations into the partition function given in Eq. (100) to get the normalization constant:

$$Z = 2e^{\beta(J_{ab}+J_{bc}+J_{ac})} + 2e^{\beta J_{ab}} + 2e^{\beta J_{bc}} + 2e^{\beta J_{ac}}. \quad (101)$$

Letting $J_{ij} = 1$ for all $i \sim j$, gives:

$$Z = 2e^{3\beta} + 6e^{\beta}. \quad (102)$$

17.3 Calculating the quantum partition function

An efficient quantum algorithm for the partition function is given by [58] for Potts models whose graph, Γ , has a topology such that it can be represented with an irreducible cyclic cocycle code (ICCC). This stipulation is non-intuitive and it takes a quantum algorithm to efficiently determine if a given graph meets this requirement. From the graph, Γ , calculate a cyclic code $C(\Gamma)$ that represents the generating structure of the graph by using Gaussian elimination on the incidence matrix of the graph, and then use Shor's algorithm to determine the

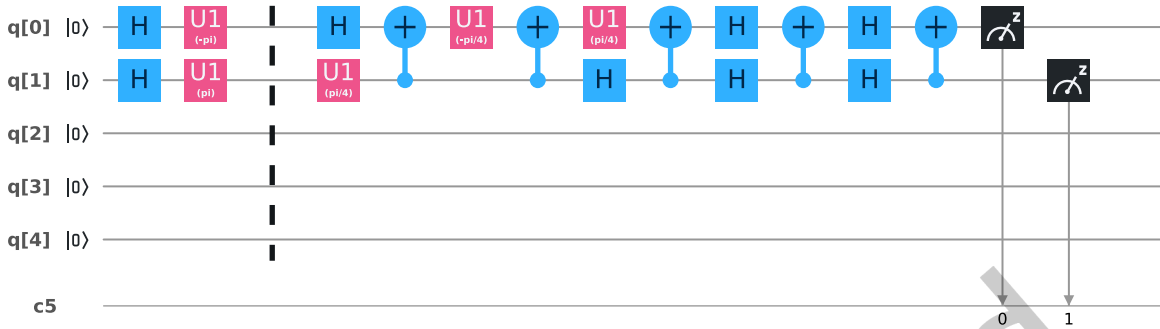


Fig. 51. Circuit for preparing the first two qubits and quantum Fourier transform on 2 qubits.

irreducible set of code words χ . If the code is not irreducible, then we will not be able to efficiently calculate the partition function for this graph.

Assuming that the given graph is ICCC, the first step in the partition function algorithm is to calculate the Gauss sum of $G_{F_{q^k}} = \sqrt{q^k} e^{i\gamma}$, where γ is a function of χ . The difficult part is to calculate γ , which can be done efficiently using the quantum Fourier transform (QFT). Using the set of values, $\{\gamma\}$ for all of the words, $\{\chi\}$ in the code; we calculate the weight spectrum $\{A_i\}$ of the code representing Γ . From this weights spectrum, the partition function Z can be efficiently calculated using classical computing.

17.4 Implementation of a quantum algorithm on the IBM Quantum Experience

We implemented one step of the full partition function algorithm using the IBM Quantum Experience. The implemented algorithm is the 2-qubit quantum Fourier transform (QFT2), as the first step in actual calculation of the partition function. The input to this step is the irreducible cocyclic code. The irreducible cyclic code for the example problem of a 3-vertex Ising model is $[1, -1]$ with $n = |V| = 3$ and $k = |E| - c(\Gamma) = 2$, where $c(\Gamma)$ is the number of connected components in the graph Γ . This small example does meet the ICCC requirement (as checked through classical calculation), so we will continue with the calculation of the partition function of the example without implementing the quantum algorithm for checking this requirement. In the fully-connected 3-vertex Ising model example given, the input to QFT2 is $q[0] = |+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}}$ and $q[1] = |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}$. In the sample score shown in Fig. 51, these qubits are prepared before the barrier. The QFT2 algorithm, as given by the Qiskit Tutorial provided by IBM[4], is the rest of the code. The output bits should be read in reverse order. Some gates could be added at the end of the QFT2 algorithm to read the gates in the same order as the input.

The result from simulating 1000 shots gives $P(\gamma = 1) = 0.47$ and $P(\gamma = 3) = 0.53$. The results from running on the actual hardware are, $P(\gamma = 0) = 0.077$, $P(\gamma = 1) = 0.462$, $P(\gamma = 2) = 0.075$, and $P(\gamma = 3) = 0.386$. We can threshold the low-probability values of gamma, ensuring that no more than the maximum number (as given in [58]) of distinct values of gamma remain. These gammas are then plugged into the calculation of the weight spectrum and the partition function.

18 QUANTUM STATE PREPARATION

The problem of preparing an n -qubit state consists first of finding the unitary transformation that takes the N -dimensional vector $(1, 0, \dots, 0)$ to the desired state $(\alpha_1, \dots, \alpha_N)$, where $N = 2^n$, and then rendering the unitary transformation into a sequence of gates.

18.1 Single qubit state preparation

As discussed before, a single qubit quantum state $|\psi\rangle$ is represented as a superposition of $|0\rangle$ and $|1\rangle$ states $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where $|\alpha|^2 + |\beta|^2 = 1$. The sizes $|\alpha|^2$ and $|\beta|^2$ represent the probability of $|\psi\rangle$ being $|0\rangle$ or $|1\rangle$. Up to a non-observable global phase, we may assume that α is real, so that $|\psi\rangle = \cos\theta|0\rangle + e^{i\phi}\sin\theta|1\rangle$ for some angles θ, ϕ . In this way, we can represent the state as a point on the unit sphere with θ the co-latitude and ϕ the longitude. This is the well-known *Bloch sphere representation*. In this way, the problem of 1-qubit state preparation consists simply of finding the unitary transformation that takes the North pole to (α, β) . In practice, this amounts to finding a sequence of available gates on actual hardware that will leave the qubit in the desired state, to a specified desired accuracy.

To prepare a specified state $|\psi\rangle$, we must find a 2×2 unitary matrix U taking the vector $|0\rangle$ to $|\psi\rangle$. An obvious simple choice for U is

$$U = \begin{pmatrix} \cos\theta & -\sin\theta e^{-i\phi} \\ \sin\theta e^{i\phi} & \cos\theta \end{pmatrix}$$

This gate is directly available in IBM Q and is implemented in a composite fashion on `ibmqx4` at the hardware level. If our goal is to initialize a base state with the fewest possible standard gates, this may not be the best choice. Instead, it makes sense to consider a more general possible unitary operator whose first column is our desired base state, and then determine the requisite number of standard gates to obtain it.

Any 2×2 unitary matrix may be obtained by means of a product of three rotation matrices, up to a global phase

$$U = e^{i\alpha} R_z(\beta) R_y(\gamma) R_z(\delta)$$

where here $R_z(\beta) = \text{diag}(e^{i\beta/2}, e^{-i\beta/2})$ and $R_y(\gamma)$ is related to $R_z(\gamma)$ by $R_y(\gamma) = SHR_z(\gamma)HSZ$. The rotation matrices $R_y(\gamma)$ and $R_z(\beta)$ correspond to the associated rotations of the unit sphere under the Bloch representation. In this way, the above decomposition is a reiteration of the standard Euler angle decomposition of elements of $SO(3)$. Thus the problem of approximating an arbitrary quantum state is reduced to the problem of finding good approximations of $R_z(\gamma)$ for various values of γ .

There has been a great deal of work done on finding efficient algorithms for approximating elements $R_z(\gamma)$ using universal gates to a specified accuracy. However, these algorithms tend to focus on the asymptotic efficiency: specifying approximations with the desired accuracy which are the generically optimal in the limit of small errors. From a practical point of view, this is an issue on current hardware, since representations tend to involve hundreds of standard gates, far outside the realm of what may be considered practical. For this reason, it makes sense to ask the question of how accurately one may initialize an arbitrary qubit with a specified number of gates.

We empirically observe that the maximum possible chordal distance from a point on the Bloch sphere to the set of exact states decreases exponentially with the number of gates. With 30 gates, every point is within a distance of 0.024 of a desired gate. Thus, to within an accuracy of about 2.5%, we can represent any base state as a product of about 30 states. We do so by preserving the states generated by 30 gates, and then for any point finding the closest exact point.




Fig. 52. Possible exact state initializations using 10, 15, and 20 gates. With 20 gates, every point on the sphere is within a distance of approximately 0.072 of an exactly obtainable state. With 30 gates, every point is within 0.024




Fig. 53. Schmidt decomposition.

18.2 Schmidt decomposition

The initialization of qubit states using more than one qubit is aided by the so-called Schmidt decomposition, which we now introduce. Specifically, the Schmidt decomposition allows one to initialize a $2n$ -qubit state by initializing a single n -qubit state, along with two specific n -qubit gates, combined together with n CNOT gates.

Mathematically, an arbitrary $2n$ -qubit state $|\psi\rangle$ may be represented as a superposition

$$|\psi\rangle = \sum_{i_1, \dots, i_n \in \{0,1\}} \sum_{j_1, \dots, j_n \in \{0,1\}} a_{i_1, \dots, i_n, j_1, \dots, j_n} |i_1 i_2 \dots i_n j_1 j_2 \dots j_n\rangle.$$

In a Schmidt decomposition, we obtain such a state by strategically choosing two orthonormal bases $|\xi_j\rangle$, $|\varphi_j\rangle$ for $j = 1, \dots, 2^n$ of the Hilbert space of n -qubit states and then writing $|\psi\rangle$ as the product

$$|\psi\rangle = \sum_{i=1}^{2^n} \lambda_i |\xi_i\rangle |\varphi_i\rangle,$$

for some well-chosen λ_i 's.

The bases $|\xi_j\rangle$ and $|\varphi_j\rangle$ may be represented in terms of two unitary matrices $U, V \in U(2^n)$, while the λ_i 's may be represented in terms of a single n -qubit state. We represent this latter state as $B|00 \dots 0\rangle$ for some $B \in U(2^n)$. Then from a quantum computing perspective, the product in the Schmidt decomposition may be accomplished by a quantum circuit combining U, V , and B with n CNOT gates as shown below for $n = 6$.

Let C_i^j denote the CNOT operator with control j and target i . Algebraically, the above circuit may be written as a unitary operator $T \in U(2^{2n})$ of the form

$$T = (U \otimes V)(C_{n+1}^1 \otimes C_{n+2}^2 \otimes \cdots \otimes C_{2n}^n)(B \otimes I).$$

We will use $|e_1\rangle, \dots, |e_{2^n}\rangle$ to denote the standard computational basis for the space of n -qubit states, in the usual order. We view each of the elements e_j as a vector in $\{0, 1\}^n$. In this notation, the formation of CNOT gates above acts on simple tensors by sending

$$C_{n+1}^1 \otimes C_{n+2}^2 \otimes \cdots \otimes C_{2n}^n : |e_i\rangle |e_j\rangle \mapsto |e_i\rangle |e_i + e_j\rangle, \quad e_i, e_j \in \{0, 1\}^n,$$

where addition in the above is performed modulo 2. Therefore the action of the operator T associated to the above circuit on the basis vector $|00 \dots 0\rangle$ is

$$\begin{aligned} T|00 \dots 0\rangle &= (U \otimes V)(C_{n+1}^1 \otimes C_{n+2}^2 \otimes \cdots \otimes C_{2n}^n)(B \otimes I)|00 \dots 0\rangle \\ &= (U \otimes V)(C_{n+1}^1 \otimes C_{n+2}^2 \otimes \cdots \otimes C_{2n}^n) \sum_{i=1}^{2^n} b_{i1} |e_i\rangle |e_1\rangle \\ &= (U \otimes V) \sum_{i=1}^{2^n} b_{i1} |e_i\rangle |e_i\rangle \\ &= \sum_{i=1}^{2^n} b_{i1} (U |e_i\rangle)(V |e_i\rangle) = |\psi\rangle. \end{aligned}$$

Thus we see that the above circuit performs precisely the sum desired from the Schmidt decomposition.

To get the precise values of U , V , and B , we write $|\psi\rangle = \sum_{i,j=1}^{2^n} a_{ij} |e_i\rangle |e_j\rangle$ for some constants $a_{ij} \in \mathbb{C}$ and define A to be the $2^n \times 2^n$ matrix whose entries are the a_{ij} 's. Then comparing this to our previous expression for $|\psi\rangle$, we see

$$\sum_{i,j=1}^{2^n} a_{ij} |e_i\rangle |e_j\rangle = \sum_{k=1}^{2^n} b_{k1} (U |e_k\rangle)(V |e_k\rangle).$$

Multiplying on the left by $\langle e_i | \langle e_j |$ this tells us

$$a_{ij} = \sum_{k=1}^{2^n} b_{k1} u_{ik} v_{jk},$$

where here $u_{ik} = \langle e_i | U |e_k\rangle$ and $v_{jk} = \langle e_j | V |e_k\rangle$ are the i, k 'th and j, k 'th entries of U and V , respectively. Encoding this in matrix form, this tells us

$$V \text{diag}(b_{i1}, \dots, b_{in}) U^T = A.$$

Then to calculate the value of U , V and the b_{i1} 's, we use the fact that V is unitary to calculate:

$$A^\dagger A = U^{T\dagger} \text{diag}(|b_{i1}|^2, \dots, |b_{in}|^2) U^T.$$

Thus if we let $|\lambda_1|^2, \dots, |\lambda_n|^2$ be the eigenvalues of $A^\dagger A$, and let U to be a unitary matrix satisfying

$$U^T A^\dagger A U^{T\dagger} = \text{diag}(|\lambda_1|^2, \dots, |\lambda_n|^2),$$

let $b_{i1} = \lambda_i$ for $i = 1, \dots, n$ and let

$$V = A U^{T\dagger} \text{diag}(\lambda_1, \dots, \lambda_n)^{-1}.$$

The matrix U is unitary, and one easily checks that V is therefore also unitary. Moreover $\sum_i |b_{i1}|^2 = \text{Tr}(A^\dagger A) = \sum_i |a_{ij}|^2 = 1$, and so the b_{i1} 's are representative of an n -qubit state and can be taken as the first column of B . Readers familiar with singular value decompositions (SVD) will recognize that the Schmidt decomposition of a bipartite state is essentially the SVD of the coefficient matrix A associated with the state. The λ_i coefficients being the singular values of A .

18.3 Two-qubit state preparation

An arbitrary two-qubit state $|\psi\rangle$ is a linear combination of the four base states $|00\rangle, |01\rangle, |10\rangle, |11\rangle$ such that the square sum of the magnitudes of the coefficients is 1. In terms of a quantum circuit, this is the simplest case of the circuit defined above in the Schmidt decomposition, and may be accomplished with three 1-qubit gates and exactly 1 CNOT gate, as featured in Fig. 54.




Fig. 54. Circuit for two qubit-state preparation. The choice of U , V , and B are covered comprehensively in the Schmidt decomposition description.

18.4 Two-qubit gate preparation

In order to initialize a four-qubit state, we require the initialization of arbitrary two-qubit gates. A two-qubit gate may be represented as an element U of $SU(4)$. As it happens, any element of $U(4)$ may be obtained by means of precisely 3 CNOT gates, combined with 7 1-qubit gates arranged in a circuit of the form given in Fig. 55.




Fig. 55. Circuit implementation of an arbitrary two qubit gate.

The proof of this is nontrivial and relies on a characterization of the image of $SU(2)^{\otimes 2}$ in $SU(4)$ using the Makhlin invariants [111]. We do not aim to reproduce the proof here. Instead, we merely aim to provide a recipe by which one may successfully obtain any element of $SU(4)$ via the above circuit and an appropriate choice of the one-qubit gates.

Let $U \in SU(4)$ be the element we wish to obtain. To choose A, B, C, D and the R_i 's, let C_j^i denote the CNOT gate with control on qubit i and target qubit j and define α, β, δ by

$$\alpha = \frac{x+y}{2}, \quad \beta = \frac{x+z}{2}, \quad \delta = \frac{y+z}{2}$$

for e^{ix}, e^{iy}, e^{iz} the eigenvalues of the operator $U(Y \otimes Y)U^\dagger(Y \otimes Y)$. Then set

$$R_1 = R_z(\delta), R_2 = R_y(\beta), R_3 = R_y(\alpha), E = C_1^2(S_z \otimes S_x)$$

and also

$$V = e^{i\pi/4}(Z \otimes I)C_1^2(I \otimes R_3)C_2^1(R_1 \otimes R_2)C_1^2(I \otimes S_z^\dagger),$$




Fig. 56. Circuit for four qubit-state preparation. The four phases of the circuit are indicated in dashed boxes.

where S_z is the single qubit $\pi/2$ rotation around the z axis. Define \tilde{U}, \tilde{V} by $\tilde{U} = E^\dagger U E$ and $\tilde{V} = E^\dagger V E$. Let \tilde{A}, \tilde{B} be the real, unitary matrices diagonalizing the eigenvectors of $\tilde{U}\tilde{U}^T$ and $\tilde{V}\tilde{V}^T$, respectively. Set $X = \tilde{A}^T \tilde{B}$ and $Y = \tilde{V}^T \tilde{B}^T \tilde{A} U$. Then EXE^\dagger and EYE^\dagger are in $SU(2)^{\otimes 2}$ and we choose A, B, C, D such that

$$(AS_z^\dagger) \otimes (Be^{i\pi/4}) = EXE^\dagger \quad \text{and} \quad C \otimes (S_z D) = EYE^\dagger.$$

By virtue of this construction, the above circuit is algebraically identical to U .

18.5 Four qubit state preparation

For efficient four qubit state preparation we use the recipe in Ref.[97]. Results in the previous sections show that any two-qubit state requires 1 CNOT gate, any two-qubit operator requires three CNOT gates, and the Schmidt decomposition of a four qubit state requires two CNOT gates. From this we see that we should be able to write a circuit initializing any four-qubit state with only 9 CNOT gates in total, along with 17 one-qubit gates. This represents the second most simple case of the Schmidt decomposition, which we write in combination with our generic expression for 2-qubit gates as shown in Fig. 56. The above circuit naturally breaks down into four distinct stages, as shown by the separate groups surrounded by dashed lines. During the first stage, we initialize the first two qubits to a specific state relating to a Schmidt decomposition of the full 4 qubit state. Stage two consists of two CNOT gates relating the first and last qubits. Stages three and four are generic circuits representing the unitary operators associated to the orthonormal bases in the Schmidt decomposition.

The results of this circuit implemented on a quantum processor are given in Fig. 57. While the results when implemented on a simulator are given in Fig. 58.

19 QUANTUM TOMOGRAPHY

19.1 Problem definition and background

Quantum state estimation, or tomography, deals with the reconstruction of the state of a quantum system from measurements of several preparations of this state. In the context of quantum computing, imagine that we start with the state $|00\rangle$, and apply some quantum algorithm (represented by a unitary matrix U) to the initial state, thus obtaining a state $|\psi\rangle$. We can measure this state in the computational z basis, or apply some rotation (represented by V) in order to perform measurements in a different basis. Quantum state tomography aims to answer the following question: is it possible to reconstruct the state $|\psi\rangle$ from a certain number of such measurements? Hence, quantum tomography is not a quantum algorithm *per se*, but it is an important procedure for certifying the performance of quantum algorithms and assessing the quality of the results that can be corrupted by decoherence, environmental noise, and biases, inevitably present in analogue machines. Moreover similar procedures can be used for certifying the initial state, as well as for measuring the fidelity of gates.




Fig. 57. Verification of 4 qubit state preparation on `ibmqx2` which is a 5 qubit machine. The last qubit is not used in the circuit. The above histogram shows that, the state prepared in `ibmqx2` has nonzero overlaps with basis states that are orthogonal to the target state to be prepared.




Fig. 58. Verification of the quantum circuit for four qubit-state preparation. The differences in the exact and the simulator results are due to statistical fluctuations arising from the probabilistic nature of quantum measurement. They will become closer to each other when the number of samples are increased.

Given a single copy of the state, it is impossible to reconstruct $|\psi\rangle$: for example, there is no quantum measurement that can even distinguish non-orthogonal quantum states, such as $|0\rangle$ and $(|0\rangle + |1\rangle)/\sqrt{2}$, with certainty. However, it is possible to perform quantum tomography when multiple copies of the state is available. It means that one needs to run the quantum algorithm (i.e., apply U to initial state) many times to produce many copies of the quantum state to be able to characterize $|\psi\rangle$. Unfortunately, because of the noise, in practice it is impossible to obtain the exact same state $|\psi\rangle$ every time; instead, one should see a mixture of different states: $|\psi_1\rangle, |\psi_2\rangle, \dots, |\psi_k\rangle$. In general, there does not exist a single $|\psi\rangle$ describing this mixture. Therefore, we need a to use the *density*

matrix representation of quantum states. We briefly discussed this representation in the context of quantum principal component analysis in Section XIV.

Let us denote p_i the probability of occurrence of the state $|\psi_i\rangle$. The density matrix of this ensemble is given by,

$$\rho = \sum_i p_i |\psi_i\rangle\langle\psi_i|. \quad (103)$$

Using this more general definition of the state, the expected value of an observable A is given by $\langle A \rangle = \sum_i p_i \langle\psi_i|A|\psi_i\rangle = \text{Tr}(A\rho)$. The density matrix has the following properties:

- $\text{Tr } \rho = 1$, i.e., probabilities sum to one;
- $\rho = \rho^\dagger$, and $\rho \geq 0$, i.e., all eigenvalues are either positive or zero.

The goal of quantum state tomography is to reconstruct the quantum state ρ from many repeated runs of the quantum algorithm, using a set of measurements on ρ . Quantum measurements are described by a collection of *measurement operators* M_i that satisfy the completeness relation $\sum_i M_i^\dagger M_i = I$. This condition ensures that the probabilities of measurement outcomes $p_i = \text{Tr}(\rho M_i^\dagger M_i)$ sum to one for any state ρ . In a popular setting for quantum tomography that is closely related to what is currently achievable with modern quantum computers, the measurement operators are chosen in a special class of *projective measurements* – projectors P_i that satisfy $P_i P_j = \delta_{i,j} P_i$ and $\sum_i P_i = I$ given that in this case $P_i^\dagger P_i = P_i$, and hence the measurement probability is given by the relation $p_i = \text{Tr}(\rho P_i)$. This choice represents a single instance of a more general measurement formalism that deals with situations when we are only interested in the probabilities of the measurement outcomes, called *Positive Operator-Valued Measures* (POVM). In this introduction we will only deal with projective measurements described above which will be sufficient for our purpose, and refer the reader to [92] for details on POVMs and general measurement formalism.

For a single qubit, examples of measurement projectors in the computational basis are given by $P_0 = |0\rangle\langle 0|$ and $P_1 = |1\rangle\langle 1|$, and in the x -basis by $P_\pm = \frac{1}{\sqrt{2}}(|0\rangle \pm |1\rangle) \otimes \frac{1}{\sqrt{2}}(\langle 0| \pm \langle 1|)$. At this point, one could ask: what is the set of projectors that represent a *quorum*, i.e. provides sufficient information to identify the state of the system in the limit of a large number of observations? The answer to this question is important for the quantum state tomography task, as it allows to determine an *informationally complete* set of measurements that is necessary to reconstruct the state. For a single-qubit example, the density matrix ρ_{qubit} can be decomposed as

$$\rho_{\text{qubit}} = \frac{\text{Tr}(\rho_{\text{qubit}})I + \text{Tr}(\rho_{\text{qubit}}X)X + \text{Tr}(\rho_{\text{qubit}}Y)Y + \text{Tr}(\rho_{\text{qubit}}Z)Z}{2}, \quad (104)$$

where X , Y , and Z are Pauli matrices, that can be interpreted as projectors in the x -, y -, and z -basis. From this expression, it is easy to see that $I/\sqrt{2}$, $X/\sqrt{2}$, $Y/\sqrt{2}$, $Z/\sqrt{2}$ forms the set of measurement operators that provide sufficient information for reconstructing the state ρ_{qubit} from measurements $p_i = \text{Tr}(\rho_{\text{qubit}}P_i)$. Using this example, a set of measurements for a general system ρ is said to be *informationally complete* if the relations $p_i = \text{Tr}(\rho P_i)$ can be inverted to unambiguously reconstruct the state ρ . The state ρ can be uniquely expressed using the obtained measurements whenever the matrix $\text{Tr}(P_i P_j)$ is invertible. For a multi-qubit state on n qubits, a simple example of an informationally complete set of measurements is given by a set of tensor products of all possible combinations of Pauli matrices. Notice, however, that a smaller number of measurement operators may be sufficient; the necessary number of measurement operators is related to the number of independent parameters in the density matrix ρ .

In real experiment, a finite number of measurements is collected for each measurement operator. Given the measurement occurrences m_i for each projector P_i , we define the associated empirical frequency as $\omega_i = m_i/m$. Then the quantum tomography problem can be stated as follows: reconstruct ρ from the informationally complete set of couples of projectors and measurement frequencies $\{P_i, \omega_i\}$. In other words, we would like to “match”

$\text{Tr}(P_i\rho)$ and ω_i . The next section presents a short overview of most popular general methods for the quantum state estimation.

19.2 Short survey of existing methods

Most popular methods for quantum tomography in the general case include:

- (1) **Linear inversion.** In this method, we simply aim at inverting the system of equations $\text{Tr}(P_i\rho) = \omega_i$. Although being fast, for a finite number of measurements thus obtained estimation $\hat{\rho}$ does not necessarily satisfy $\hat{\rho} \succeq 0$ (i.e., might contain negative eigenvalues): this happens due to experimental inaccuracies and statistical fluctuations of coincidence counts, which leads to differences between the empirical measurement frequencies ω_i and the calculated values $\text{Tr}(P_i\rho)$ [70].
- (2) **Linear regression.** This method corrects for the disadvantages of the linear inversion by solving a constrained quadratic optimization problem [101]:

$$\hat{\rho} = \underset{\rho}{\text{argmin}} \sum_i [\text{Tr}(P_i\rho) - \omega_i]^2 \quad \text{s.t. } \text{Tr } \rho = 1 \text{ and } \rho \succeq 0.$$

However, this objective function implicitly assumes that the residuals are Gaussian-distributed, which may hold in a limit of a large number of independent measurements due to the central limit theorem, but does not necessarily apply in practice for a finite number of measurements where deviations from normal distribution can be important.

- (3) **Maximum likelihood.** In this by far most popular algorithm for quantum state estimation, one aims at maximizing the log-probability of observations [67, 70]:

$$\hat{\rho} = \underset{\rho}{\text{argmax}} \sum_i \omega_i \ln \text{Tr}(P_i\rho) \quad \text{s.t. } \text{Tr } \rho = 1 \text{ and } \rho \succeq 0.$$

This is a convex problem that outputs a positive semidefinite (PSD) solution $\hat{\rho} \succeq 0$. However, it is often stated that the maximum likelihood (ML) method is slow, and several recent papers attempted to develop faster methods of gradient descent with projection to the space of PSD matrices, see e.g. [109]. Among other common criticisms of this method one can name the fact that ML might yield rank-deficient solutions, which results in an infinite conditional entropy that is often used as a metric of success of the reconstruction.

- (4) **Bayesian methods.** This is a slightly more general approach compared to the ML method which includes some prior [19], or corrections to the basic ML objective, see e.g., the so-called Hedged ML [18]. However, it is not always clear how to choose these priors in practice. Markov Chain Monte Carlo Methods that are used for general priors are known to be slow.

Let us mention that there exist other state reconstruction methods that attempt to explore a particular known structure of the density matrix, such as compressed-sensing methods [62] in the case of low-rank solutions, and matrix product states [38] or neural networks based approaches [126] for pure states with limited entanglement, *etc.* One of the points we can conclude from this section is that the ultimately best general method for the quantum state tomography is not yet known and certainly depends on the applications. However, it seems that maximum likelihood is still the most widely discussed method in the literature; in what follows, we implement and test ML approach to quantum tomography on the IBM quantum computer.

19.3 Implementation of the Maximum Likelihood method on 5-qubit IBM QX

We present an efficient implementation of the ML method using a fast gradient descent with an optimal 2-norm projection [117] to the space of PSD matrices³. In what follows, we apply quantum tomography to study the performance of the IBM Q.

³The julia implementation of the algorithm is available at <http://gitlab.lanl.gov/QuantumProgramming2017/QuantumTomography>




Fig. 59. Left: measurements of the single qubit state after the application of the Hadamard gate, in z , y and x basis. Right: experimental setup for testing the effects of decoherence.

19.3.1 Warm-up: Hadamard gate. Let us start with a simple one-qubit case of the Hadamard gate, see Fig. 59, Left. This gate transforms the initial qubit state $|0\rangle$ as follows: $H : |0\rangle \rightarrow |+\rangle_x = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$, so that the density matrix should be close to $\rho = |+\rangle_x\langle +|_x$. As discussed in section 19.1, for performing quantum tomography in the single-qubit case, it is sufficient to collect measurements in the x , y , and z basis. In the limit of a large number of measurements, we expect to see the following frequencies in the z , y , and x basis (all vector expressions are given in the computational basis):

$$\begin{bmatrix} 1 \\ 0 \end{bmatrix} \rightarrow \frac{1}{2}, \quad \begin{bmatrix} 0 \\ 1 \end{bmatrix} \rightarrow \frac{1}{2}, \quad \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ i \end{bmatrix} \rightarrow \frac{1}{2}, \quad \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -i \end{bmatrix} \rightarrow \frac{1}{2}, \quad \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ 1 \end{bmatrix} \rightarrow 1, \quad \frac{1}{\sqrt{2}} \begin{bmatrix} 1 \\ -1 \end{bmatrix} \rightarrow 0.$$

We learn the estimated density matrix $\hat{\rho}$ from measurements in each basis using the maximum likelihood method, and look at the decomposition:

$$\hat{\rho} = \lambda_1 |\psi_1\rangle\langle\psi_1| + \lambda_2 |\psi_2\rangle\langle\psi_2|,$$

which would allow us to see what eigenstates contribute to the density matrix, and what is their weight. Indeed, in the case of ideal observations we should get $\lambda_1 = 1$, with $|\psi_1\rangle = \left[\frac{1}{\sqrt{2}} \quad \frac{1}{\sqrt{2}}\right]^T$, and $\lambda_2 = 0$ with $|\psi_2\rangle = \left[\frac{1}{\sqrt{2}} \quad -\frac{1}{\sqrt{2}}\right]^T$, corresponding to the original pure state associated with $|+\rangle_x$.

Instead, we obtain the following results for the eigenvalues and associated eigenvectors after 8152 measurements (the maximum number in one run on IBM QX) in each basis (z , y , x):

$$\lambda_1 = 0.968 \rightarrow \begin{bmatrix} 0.715 - 0.012i \\ 0.699 \end{bmatrix}, \quad \lambda_2 = 0.032 \rightarrow \begin{bmatrix} 0.699 - 0.012i \\ -0.715 \end{bmatrix},$$

i.e., in 96% of cases we observe the state close to $|+\rangle_x$, and the rest corresponds to the state which is close to $|-\rangle_x$. Note that the quantum simulator indicates that this amount of measurements is sufficient to estimate matrix elements of the density matrix with an error below 10^{-3} in the ideal noiseless case. In order to check the effect of decoherence, we apply a number of identity matrices (Fig. 59, Right) which forces an additional waiting on the system, and hence promotes decoherence of the state. When applying 18 identity matrices, we obtain the following decomposition for $\hat{\rho}$

$$\lambda_1 = 0.940 \rightarrow \begin{bmatrix} 0.727 - 0.032i \\ 0.686 \end{bmatrix}, \quad \lambda_2 = 0.060 \rightarrow \begin{bmatrix} 0.685 - 0.030i \\ -0.728 \end{bmatrix},$$

while application of 36 identity matrices results in

$$\lambda_1 = 0.927 \rightarrow \begin{bmatrix} 0.745 - 0.051i \\ 0.664 \end{bmatrix}, \quad \lambda_2 = 0.073 \rightarrow \begin{bmatrix} 0.663 - 0.045i \\ -0.747 \end{bmatrix}.$$

The effect of decoherence is visible in the degradation of the eigenstates, as well as in a more frequent occurrence of the eigenstate close to $|-\rangle_x$.




Fig. 60. Left: example of a measurement of the two-qubit maximally entangled state created with the combination of H , X and $CNOT$ gates in the yz basis. Right: experimental setup for testing the effects of decoherence.

19.3.2 Maximally entangled state for two qubits. Let us now study the two-qubits maximally entangled state, which is an important part of all quantum algorithms achieving quantum speed-up over their classical counterparts. The state $\frac{1}{\sqrt{2}}(|10\rangle + |01\rangle)$ we are interested in is produced by the combination of H , X and $CNOT$ gates as shown in Fig. 60, Left. We follow the same procedure as in the case of the Hadamard gate, described above, and first estimate the density matrix $\hat{\rho}$ using 8152 measurements for each of the zz , yy , xx , zx and yz basis, and then decompose it as $\hat{\rho} = \sum_{i=1}^4 \lambda_i |\psi_i\rangle\langle\psi_i|$. Once again, ideally we should get $\lambda_1 = 1$ associated with $|\psi_1\rangle = \left[0 \quad \frac{1}{\sqrt{2}} \quad \frac{1}{\sqrt{2}} \quad 0\right]^T$. Instead, the analysis of the leading eigenvalues indicates that the eigenstate which is close (although significantly distorted) to the theoretical “ground truth” $|\psi_1\rangle$ above occurs in the mixture only with probability 0.87:

$$\lambda_1 = 0.871 \rightarrow \begin{bmatrix} -0.025 - 0.024i \\ 0.677 \\ 0.735 \\ -0.029 - 0.017i \end{bmatrix}, \quad \lambda_2 = 0.059 \rightarrow \begin{bmatrix} 0.598 \\ 0.123 + 0.468i \\ -0.075 - 0.445i \\ 0.454 - 0.022i \end{bmatrix}.$$

Our test of decoherence implemented using 18 identity matrices (see Figure 60, Right) shows that the probability of the “original” entangled state decreases to 0.79:

$$\lambda_1 = 0.793 \rightarrow \begin{bmatrix} -0.025 - 0.012i \\ 0.664 \\ 0.747 \\ -0.017 - 0.008i \end{bmatrix}, \quad \lambda_2 = 0.111 \rightarrow \begin{bmatrix} 0.997 \\ -0.002 - 0.058i \\ 0.035 + 0.036i \\ 0.006 + 0.007i \end{bmatrix}.$$

Interestingly enough, the second most probable eigenstate changes to the one that is close to $|00\rangle$. This might serve as an indication of the presence of biases in the machine.

The application of the quantum tomography state reconstruction to simple states in the IBM QX revealed an important level of noise and decoherence present in the machine. It would be interesting to check if the states can be protected by using the error correction schemes, which is the subject of the next section.

20 TESTS OF QUANTUM ERROR CORRECTION IN IBM Q

In this section, we study whether quantum error correction (QEC) can improve computation accuracy in `ibmqx4`. The practical answer to this question seems to be “No”. Although some error correction effects are observed in `ibmqx4`, improvements are not exponential and get completely spoiled by errors induced by extra gates and qubits needed for the error correction protocols.

20.1 Problem definition and background

As we have seen throughout this review, the quality of computation on actual quantum processors is degraded by errors in the system. This is because currently available chips are not *fault tolerant*. It is widely believed that once the inherent error rates of a quantum processor is sufficiently lowered, fault tolerant quantum computation will be possible using quantum error correction (QEC). The current error rates of the IBM Q machines are not small enough to allow fault tolerant computation. We refer the reader to a survey and introduction on QEC [42], while at the same time offering an alternative point of view that we support with a few experiments on the IBM chip. Detailed studies of QEC using more sophisticated error correcting schemes have been performed on IBM hardware [124, 132, 134]. Qiskit also provides some tools that can be used for QEC. The recent work in Ref. [133] introduces some of these capabilities of Qiskit with example code.

The central idea of QEC is to use entanglement to encode quantum superposition in a manner which is robust to errors. The exact encoding depends upon the kind of errors we want to protect against. In this section we will look at a simple encoding that will protect against bit flip errors. Here we encode a single qubit state,

$$|\psi\rangle = C_0|0\rangle + C_1|1\rangle, \quad (105)$$

using an entangled state, such as

$$|\psi\rangle = C_0|0\rangle^{\otimes n_q} + C_1|1\rangle^{\otimes n_q}, \quad (106)$$

where n_q is the number of qubits representing a single qubit in calculations.

The assumption is that small probability errors will likely lead to unwanted flips of only one qubit (in case when $n_q > 3$ this number can be bigger but we will not consider more complex situations here). Such errors produce states that are essentially different from those described by Equation (106). Measurements can then be used to fix a single qubit error using, for instance, a majority voting strategy. More complex errors are assumed to be exponentially suppressed, which can be justified if qubits experience independent decoherence sources.

We question whether QEC can work to protect quantum computations that require many quantum gate operations for the following reason. The main source of errors then is not spontaneous qubit decoherence but rather the finite fidelity of quantum gates. When quantum gates are applied to strongly entangled states, such as (106), they lead to *highly correlated* dynamics of *all* entangled qubits. We point out that errors introduced by such gates have essentially different nature from random uncorrelated qubit flips. Hence, gate-induced errors may not be treatable by standard error correction strategies when transitions are made between arbitrary unknown quantum state.

To explore this point, imagine that we apply a gate that rotates a qubit by an angle $\pi/2$. It switches superposition states $|\psi_{\pm}\rangle = (|0\rangle \pm |1\rangle)/\sqrt{2}$ into, respectively, $|0\rangle$ or $|1\rangle$ in the measurement basis. Let the initial state be $|\psi_{+}\rangle$ but we do not know this before the final measurement. Initially, we know only that initial state can be either $|\psi_{+}\rangle$ or $|\psi_{-}\rangle$. To find what it is, we rotate qubit to the measurement basis. The gate is not perfect, so the final state after the gate application is

$$|u\rangle = \cos(\delta\phi)|0\rangle + \sin(\delta\phi)|1\rangle, \quad (107)$$

with some error angle $\delta\phi \ll 1$. Measurement of this state would produce the wrong answer 1 with probability

$$P \approx (\delta\phi)^2. \quad (108)$$

The value $1 - P$ is called the fidelity of the gate. In IBM chip it is declared to be 0.99 at the time of writing, which is not much. It means that after about 30 gates we should loose control. Error correction strategies can increase the number of allowed gates by an order of magnitude even at such a fidelity if we encode one qubit in three.




Fig. 61. Quantum circuit that creates the state $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ then applies 16 T-gates that are equivalent to the identity operation, and then applies the gate that transforms the entangled state into the trivial state $|0\rangle$.

In order to reduce this error, we can attempt to work with the 3-qubit version of the states in Eq. (106). For example, let us consider the desired gate that transfers states

$$|\pm\rangle = (|000\rangle \pm |111\rangle)/\sqrt{2}, \quad (109)$$

into states $|000\rangle$ and $|111\rangle$ in the measurement basis, respectively. This gate is protected in the sense that a single unwanted random qubit flip leads to final states that are easily corrected by majority voting.

However, this is not enough because now we have to apply the gate that makes a rotation by $\pi/2$ in the basis (109). The error in this rotation angle leads to the final state

$$|u\rangle = \cos(\delta\phi)|000\rangle + \sin(\delta\phi)|111\rangle, \quad (110)$$

i.e., this particular error cannot be treated with majority voting using our scheme because it flips all three qubits. On the other hand, this is precisely the type of errors that is most important when we have to apply many quantum gates because basic gate errors are mismatches between desired and received qubit rotation angles irrespectively of how the qubits are encoded. With nine qubits, we could protect the sign in Eq. 110 but this was beyond our hardware capabilities.

Based on these thoughts, traditional QEC may not succeed in achieving exponential suppression of errors related to non-perfect quantum gate fidelity. The latter is the main source of decoherence in quantum computing that involves many quantum gates. As error correction is often called the only and first application that matters before quantum computing becomes viable at large scale, this problem must be studied seriously and expeditiously. In the following subsection we report on our experimental studies of this problem with IBM's 5-qubit chip.

20.2 Test 1: errors in single qubit control

First, let us perform trivial operation shown in Fig. 61: we create a superposition of two qubit states

$$|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}, \quad (111)$$

then apply many gates that altogether do nothing, i.e., they just bring the qubit back to the superposition state (111). We need those gates just to accumulate some error while the qubit's state is not trivial in the measurement basis. Finally, we apply the gate that transforms its state back to $|0\rangle$.

Repeated experiments with measurements then produced wrong answer 13 times from 1000 samples. Thus, we estimate the error of the whole protocol, which did not use QEC, as

$$P_1 = 0.013,$$

or 1.6%. This is consistent and even better than declared 1% single gate fidelity because we applied totally 18 gates.




Fig. 62. Quantum circuit that creates state $|-\rangle = (|000\rangle - |111\rangle)/\sqrt{2}$ then applies 16 T -gates that are equivalent to the identity operation, and then applies the gate that transforms the entangled GHZ state back into the trivial state $|000\rangle$. Measurements that return 1 for only one of the three qubits are interpreted as the $|000\rangle$ state at the end, while outcomes with two or three units are interpreted as the final state $|111\rangle$.

20.3 Test 2: errors in entangled 3 qubits control

Next, we consider the circuit in Fig. 62 that initially creates the GHZ state $|-\rangle = (|000\rangle - |111\rangle)/\sqrt{2}$, then applies the same number, i.e. 16, of T -gates that lead to the same GHZ state. Then we apply the sub-circuit that brings the whole state back to $|000\rangle$.

Our goal is to quantify the precision of identifying the final result with the state $|000\rangle$. If a single error bit flip happens, we can interpret results $|100\rangle$, $|010\rangle$ and $|001\rangle$ as $|000\rangle$ using majority voting. If needed, we can then apply a proper pulse to correct for this. So, in such cases we can consider the error treatable. If the total sum of probabilities of the final state $|000\rangle$ and final states with a single bit flipped is larger than P_1 from the previous single-qubit test, then we say that the quantum error correction works, otherwise, it doesn't. Our experiments showed that probabilities of events that lead to wrong final interpretation are as follows:

$$P_{110} = 0.006, \quad P_{101} = 0.02, \quad P_{011} = 0.016, \quad P_{111} = 0.005.$$

Thus, the probability to get the wrong interpretation of the result as the final state $|111\rangle$ of the encoded qubit is

$$P_3 = P_{110} + P_{101} + P_{011} + P_{111} = 0.047,$$

while the probability to get any error $1 - P_{000} = 0.16$.

20.4 Discussion

Comparing results of the tests without and with QEC, we find that the implementation of a simple version of QEC does not improve the probability to interpret the final outcome correctly. The error probability of calculations without QEC gives a smaller probability of wrong interpretation, $P_1 = 1.3\%$, while the circuit with QEC gives an error probability $P_3 = 4.7\%$, even though we used majority voting that was supposed to suppress errors by about an order of magnitude.

More importantly, errors that lead to more than one qubit flip are not exponentially suppressed. For example, the probability $P_{101} = 0.02$ is close to the probability of a single bit flip event $P_{010} = 0.029$. We interpret this to mean that errors are not the results of purely random bit flip decoherence effects but rather follow from correlated errors induced by the finite precision of quantum gates. The higher error rate in 3-qubit case could be attributed to the much worse fidelity of the controlled-NOT gate. The circuit itself produces the absolutely correct result $|000\rangle$ in 84% of simulations. If the remaining errors were produced by uncorrelated bit flips, we would see outcomes with more than one wrong bit flip with total probability less than 1% but we found that such events have a much larger total probability $P_3 = 4.7\%$.

In defense of QEC, we note that probabilities of single bit flip errors were still several times larger than probabilities of multiple (two or three) wrong qubit flip errors. This means that at least partly QEC works, i.e., it corrects the state to $|000\rangle$ with 4.7% precision, versus the initially 16% in the wrong state. So, at least some part of the errors can be treated. However, an efficient error correction must show *exponential* suppression of errors, which was not observed in this test.

Summarizing, this brief test shows no improvements that would be required for efficient quantum error correction. The need to use more quantum gates and qubits to correct errors only leads to a larger probability of wrong interpretation of the final state. This problem will likely become increasingly much more important because without quantum error correction the whole idea of conventional quantum computing is not practically useful. Fortunately, IBM's quantum chips can be used for experimental studies of this problem. We also would like to note that quantum computers can provide computational advantages beyond standard quantum algorithms and using only classical error correction [115]. So, they must be developed even if problems with quantum error correction prove detrimental for conventional quantum computing schemes at achievable hardware quality.

ACKNOWLEDGMENTS

We would like to acknowledge the help from numerous readers who pointed out errors and misprints in the earlier version of the manuscript. The code and implementations accompanying the paper can be found at https://github.com/lanl/quantum_algorithms.

REFERENCES

- [1] [n.d.]. ibmq-device-information. <https://github.com/Qiskit/ibmq-device-information/tree/master/backends/tenerife/V1>. Accessed: 14-12-2019.
- [2] Scott Aaronson. 2015. Read the fine print. *Nature Physics* 11, 4 (2015), 291–293.
- [3] Scott Aaronson and Lijie Chen. 2017. Complexity-Theoretic Foundations of Quantum Supremacy Experiments. In *32nd Computational Complexity Conference, CCC 2017, July 6-9, 2017, Riga, Latvia (LIPIcs, Vol. 79)*, Ryan O'Donnell (Ed.). Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 22:1–22:67. <https://doi.org/10.4230/LIPIcs.CCC.2017.22>
- [4] Héctor Abraham, Ismail Yunus Akhalwaya, Gadi Aleksandrowicz, and yotamvakninibm. 2019. Qiskit: An Open-source Framework for Quantum Computing. <https://doi.org/10.5281/zenodo.2562110>
- [5] Andris Ambainis. 2007. Quantum walk algorithm for element distinctness. *SIAM J. Comput.* 37, 1 (2007), 210–239.
- [6] A. Ambainis, H. Buhrman, P. Høyer, M. Karpinski, and P. Kurur. 2002. Quantum matrix verification. (2002).
- [7] Andris Ambainis and R. Spalec. 2006. Quantum Algorithms for Matching and Network Flows. in *Lecture Notes in Computer Science: STACS 2006* 3884 (2006).
- [8] Itai Arad and Zeph Landau. 2010. Quantum computation and the evaluation of tensor networks. *SIAM J. Comput.* 39, 7 (2010), 3089–3121.
- [9] Frank Arute, Kunal Arya, Ryan Babbush, Dave Bacon, Joseph C Bardin, Rami Barends, Rupak Biswas, Sergio Boixo, Fernando GSL Brandao, David A Buell, et al. 2019. Quantum supremacy using a programmable superconducting processor. *Nature* 574, 7779 (2019), 505–510.
- [10] Dave Bacon, Isaac L Chuang, and Aram W Harrow. 2007. The quantum Schur and Clebsch-Gordan transforms: I. efficient qudit circuits. (2007), 1235–1244.
- [11] Dave Bacon and Wim Van Dam. 2010. Recent progress in quantum algorithms. *Commun. ACM* 53, 2 (2010), 84–93.
- [12] Stefanie Barz, Ivan Kassal, Martin Ringbauer, Yannick Ole Lipp, Borivoje Dakić, Alán Aspuru-Guzik, and Philip Walther. 2014. A two-qubit photonic quantum processor and its application to solving systems of linear equations. *Scientific reports* 4 (2014).
- [13] Robert Beals. 1997. Quantum computation of Fourier transforms over symmetric groups . In *Proceedings of STOC* (1997), 48–53.
- [14] Giuliano Benenti and Giuliano Strini. 2008. Quantum simulation of the single-particle Schrödinger equation. *American Journal of Physics* 76, 7 (2008), 657–662.
- [15] Charles H Bennett, Ethan Bernstein, Gilles Brassard, and Umesh Vazirani. 1997. Strengths and weaknesses of quantum computing. *SIAM journal on Computing* 26, 5 (1997), 1510–1523. <https://doi.org/10.1137/S0097539796300933>
- [16] E. Bernstein and U. Vazirani. 1993. Quantum complexity theory. In *Proc. of the Twenty-Fifth Annual ACM Symposium on Theory of Computing (STOC '93)*. 11–20. DOI:10.1145/167088.167097.
- [17] Dominic W Berry, Graeme Ahokas, Richard Cleve, and Barry C Sanders. 2007. Efficient quantum algorithms for simulating sparse Hamiltonians. *Communications in Mathematical Physics* 270, 2 (2007), 359–371.

- [18] Robin Blume-Kohout. 2010. Hedged maximum likelihood quantum state estimation. *Physical review letters* 105, 20 (2010), 200504.
- [19] Robin Blume-Kohout. 2010. Optimal, reliable estimation of quantum states. *New Journal of Physics* 12, 4 (2010), 043034.
- [20] Otakar Borůvka. 1926. O jistém problému minimálním. *Práce Mor. Přírodově d. spol. v Brně (Acta Societ. Scient. Natur. Moravicae)* 3 (1926), 37–58.
- [21] Michel Boyer, Gilles Brassard, Peter Høyer, and Alain Tapp. 1998. Tight bounds on quantum searching. *Fortschritte der Physik: Progress of Physics* 46, 4-5 (1998), 493–505.
- [22] Lucas T. Brady, Christopher L. Baldwin, Aniruddha Bapat, Yaroslav Kharkov, and Alexey V. Gorshkov. 2021. Optimal Protocols in Quantum Annealing and Quantum Approximate Optimization Algorithm Problems. *Physical Review Letters* 126 (2021), 070505. Issue 7. <https://doi.org/10.1103/PhysRevLett.126.070505>
- [23] G. Brassard et al. 2002. Quantum amplitude amplification and estimation. *Quantum Computation and Quantum Information* 9 (2002).
- [24] Carlos Bravo-Prieto, Ryan LaRose, Marco Cerezo, Yigit Subasi, Lukasz Cincio, and Patrick J Coles. 2019. Variational quantum linear solver: A hybrid algorithm for linear systems. *arXiv preprint arXiv:1909.05820* (2019).
- [25] Sergey Bravyi, Alexander Kliesch, Robert Koenig, and Eugene Tang. 2020. Obstacles to Variational Quantum Optimization from Symmetry Protection. *Physical Review Letters* 125 (2020), 260505. Issue 26. <https://doi.org/10.1103/PhysRevLett.125.260505>
- [26] H. Buhrman and R. Spalek. 2006. Quantum verification of matrix products. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete algorithm* (2006), 880–889.
- [27] X.-D. Cai, C. Weedbrook, Z.-E. Su, M.-C. Chen, M. Gu, M.-J. Zhu, L. Li, N.-L. Liu, C.-Y. Lu, and J.-W. Pan. 2013. Experimental Quantum Computing to Solve Systems of Linear Equations. *Physical Review Letters* 110, 23, Article 230501 (June 2013), 230501 pages. <https://doi.org/10.1103/PhysRevLett.110.230501> arXiv:1302.4310 [quant-ph]
- [28] Kevin K. H. Cheung and Michele Mosca. 2001. Decomposing Finite Abelian Groups. *Quantum Info. Comput.* 1, 3 (Oct. 2001), 26–32. <http://dl.acm.org/citation.cfm?id=2011339.2011341>
- [29] Nai-Hui Chia, András Gilyén, Han-Hsuan Lin, Seth Lloyd, Ewin Tang, and Chunhao Wang. 2020. Quantum-Inspired Algorithms for Solving Low-Rank Linear Equation Systems with Logarithmic Dependence on the Dimension. In *31st International Symposium on Algorithms and Computation (ISAAC 2020) (Leibniz International Proceedings in Informatics (LIPIcs), Vol. 181)*, Yixin Cao, Siu-Wing Cheng, and Minming Li (Eds.). Schloss Dagstuhl–Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 47:1–47:17. <https://doi.org/10.4230/LIPIcs.ISAAC.2020.47>
- [30] Andrew M Childs and Wim Van Dam. 2010. Quantum algorithms for algebraic problems. *Reviews of Modern Physics* 82, 1 (2010), 1.
- [31] Lukasz Cincio, Yiğit Subaşı, Andrew T Sornborger, and Patrick J Coles. 2018. Learning the quantum algorithm for state overlap. *New Journal of Physics* 20, 11 (2018), 113022.
- [32] Jill Cirasella. 2006. Classical and quantum algorithms for finding cycles. *MSc Thesis* (2006), 1–58.
- [33] C. Codsil and H. Zhan. 2011. Discrete-Time Quantum Walks and Graph Structures. (2011), 1–37. <https://arxiv.org/abs/1701.4474>
- [34] Rigetti Computing. 2017. Quantum Approximate Optimization Algorithm. Published online at <https://github.com/rigetticomputing/grove>. Accessed: 12/01/2017.
- [35] Jeremy Cook, Stephan Eidenbenz, and Andreas Bärtzchi. 2020. The Quantum Alternating Operator Ansatz on Maximum k-Vertex Cover. In *IEEE International Conference on Quantum Computing & Engineering QCE'20*. 83–92. <https://doi.org/10.1109/QCE49297.2020.00021> arXiv:1910.13483
- [36] Stephen A. Cook. 1971. The Complexity of Theorem-proving Procedures. In *Proceedings of the Third Annual ACM Symposium on Theory of Computing (Shaker Heights, Ohio, USA) (STOC '71)*. ACM, New York, NY, USA, 151–158. <https://doi.org/10.1145/800157.805047>
- [37] D. Coppersmith and S. Winograd. 1990. Matrix multiplication via arithmetic progressions. *Journal of symbolic computation* 9 (1990), 251–280.
- [38] Marcus Cramer, Martin B Plenio, Steven T Flammia, Rolando Somma, David Gross, Stephen D Bartlett, Olivier Landon-Cardinal, David Poulin, and Yi-Kai Liu. 2010. Efficient quantum state tomography. *Nature communications* 1 (2010), 149.
- [39] Sanjoy Dasgupta, Christos H. Papadimitriou, and Umesh Vazirani. 2008. *Algorithms*. McGraw-Hill, Inc., New York, NY, USA.
- [40] M. Dehn. 1911. Über unendliche diskontinuierliche Gruppen. *Math. Ann.* 71, 1 (01 Mar 1911), 116–144. <https://doi.org/10.1007/BF01456932>
- [41] D. Deutsch and R. Jozsa. 1992. Rapid solutions of problems by quantum computation. In *Proc. of the Royal Society of London A*. 439–553.
- [42] Simon J Devitt, William J Munro, and Kae Nemoto. 2013. Quantum error correction for beginners. *Reports on Progress in Physics* 76, 7 (2013), 076001.
- [43] B. L. Douglas and J. B. Wang. 2009. Efficient quantum circuit implementation of quantum walks. *Physical Review A* 79, 5 (2009), 052335.
- [44] Iain Dunning, Swati Gupta, and John Silberholz. 2018. What Works Best When? A Systematic Evaluation of Heuristics for Max-Cut and QUBO. *INFORMS Journal on Computing* 30, 3 (2018), 608–624. <https://doi.org/10.1287/ijoc.2017.0798>
- [45] Christoph Dürr, Mark Heiligman, Peter Høyer, and Mehdi Mhalla. 2006. Quantum query complexity of some graph problems. *SIAM J. Comput.* 35, 6 (2006), 1310–1328.
- [46] Christoph Durr and Peter Hoyer. 1996. A quantum algorithm for finding the minimum. *arXiv preprint quant-ph/9607014* (1996).

- [47] Jack Edmonds and Richard M. Karp. 1972. Theoretical improvements in algorithmic efficiency for network flow problems. *J. ACM* 19 (2) (1972), 248–264.
- [48] Nayak F. Magniez A, J. Roland, and M. Santha. 2011. Search via Quantum Walk. *SIAM J. Comput.* 40, 1 (2011), 142–164. <https://arxiv.org/abs/quant-ph/0608026>
- [49] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm Applied to a Bounded Occurrence Constraint Problem. *arXiv e-prints* (2014). arXiv:1412.6062.
- [50] Edward Farhi, Jeffrey Goldstone, and Sam Gutmann. 2014. A Quantum Approximate Optimization Algorithm.
- [51] Edward Farhi, Jeffrey Goldstone, Sam Gutmann, and Michael Sipser. 2000. Quantum Computation by Adiabatic Evolution. *arXiv e-prints* (2000). arXiv:quant-ph/0001106.
- [52] L. R. Ford and D. R. Fulkerson. 1956. Maximal flow through a network. *Canadian Journal of Mathematics* 8 (1956), 399–404.
- [53] R. Freivalds. 1979. Fast probabilistic algorithms. In *Proc. of 8th Symp. on Math. Foundations of Computer Science* (1979), 57–69.
- [54] Michael R. Garey and David S. Johnson. 1979. *Computers and Intractability; A Guide to the Theory of NP-Completeness*. W. H. Freeman & Co., USA.
- [55] Silvano Garnerone, Annalisa Marzuoli, and Mario Rasetti. 2007. Efficient quantum processing of 3-manifold topological invariants. *arXiv preprint quant-ph/0703037* (2007).
- [56] Julia M Georgescu, Sahel Ashhab, and Franco Nori. 2014. Quantum simulation. *Reviews of Modern Physics* 86, 1 (2014), 153.
- [57] Joseph Geraci. 2008. A new connection between quantum circuits, graphs and the Ising partition function. *Quantum Information Processing* 7, 5 (2008), 227–242.
- [58] Joseph Geraci and Daniel A Lidar. 2008. On the exact evaluation of certain instances of the Potts partition function by quantum computers. *Communications in Mathematical Physics* 279, 3 (2008), 735–768.
- [59] András Gilyén, Srinivasan Arunachalam, and Nathan Wiebe. [n.d.]. Optimizing quantum optimization algorithms via faster quantum gradient computation. In *ACM-SIAM Symposium on Discrete Algorithms, SODA'2019*. 1425–1444. <https://doi.org/10.1137/1.9781611975482.87>
- [60] Vittorio Giovannetti, Seth Lloyd, and Lorenzo Maccone. 2008. Quantum Random Access Memory. 100 (04, 2008), 160501.
- [61] Michel X. Goemans and David P. Williamson. 1995. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. *J. ACM* 42, 6 (1995), 1115–1145. <https://doi.org/10.1145/227683.227684>
- [62] David Gross, Yi-Kai Liu, Steven T Flammia, Stephen Becker, and Jens Eisert. 2010. Quantum state tomography via compressed sensing. *Physical review letters* 105, 15 (2010), 150401.
- [63] Lov K Grover. 1996. A fast quantum mechanical algorithm for database search. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*. ACM, 212–219.
- [64] Eran Halperin, Dror Livnat, and Uri Zwick. 2004. MAX CUT in cubic graphs. *Journal of Algorithms* 53, 2 (2004), 169–185. <https://doi.org/10.1016/j.jalgor.2004.06.001>
- [65] Matthew P. Harrigan et al. 2021. Quantum approximate optimization of non-planar graph problems on a planar superconducting processor. *Nature Physics* 17, 3 (2021), 332–336. <https://doi.org/10.1038/s41567-020-01105-y>
- [66] Aram W Harrow, Avinandan Hassidim, and Seth Lloyd. 2009. Quantum algorithm for linear systems of equations. *Physical review letters* 103, 15 (2009), 150502.
- [67] Zdenek Hradil. 1997. Quantum-state estimation. *Physical Review A* 55, 3 (1997), R1561.
- [68] Johan Hästad. 2001. Some Optimal Inapproximability Results. *J. ACM* 48, 4 (2001), 798–859. <https://doi.org/10.1145/502090.502098>
- [69] IBM Corporation. 2016. IBM Quantum Experience. Published online at <https://quantumexperience.ng.bluemix.net>. Accessed: 12/01/2017.
- [70] Daniel F. V. James, Paul G. Kwiat, William J. Munro, and Andrew G. White. 2001. Measurement of qubits. *Phys. Rev. A* 64 (2001), 052312. Issue 5.
- [71] Sonika Johri, Damian S Steiger, and Matthias Troyer. 2017. Entanglement spectroscopy on a quantum computer. *Physical Review B* 96, 19 (2017), 195136.
- [72] Stephan Jordan. 2011. Quantum Algorithm Zoo. Published online at <https://math.nist.gov/quantum/zoo/>. Accessed: 3/18/2018.
- [73] Stephen P. Jordan. 2009. Fast quantum algorithms for approximating some irreducible representations of groups . (2009), 1–21. arXiv:0811.0562 <https://arxiv.org/abs/0811.0562>
- [74] Petteri Kaski. 2002. Eigenvectors and spectra of Cayley graphs. http://www.tcs.hut.fi/Studies/T-79.300/2002S/esitelmat/kaski_paper_020506.pdf
- [75] J. Kempe. 2003. Quantum random walks - an introductory overview. *Contemporary Physics* 44, 4 (2003), 307–327. <http://www.tandfonline.com/doi/abs/10.1080/00107151031000110776>
- [76] V. Kendon. 2011. Where to quantum walk. (2011), 1–13. <https://arxiv.org/abs/1107.3795>
- [77] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. 2007. Optimal Inapproximability Results for MAXCUT and Other 2Variable CSPs? *SIAM J. Comput.* 37, 1 (2007), 319–357. <https://doi.org/10.1137/S0097539705447372>
- [78] Daphne Koller and Nir Friedman. 2009. *Probabilistic graphical models: principles and techniques*. MIT Press.

- [79] M W Krentel. 1986. The Complexity of Optimization Problems. In *Proceedings of the Eighteenth Annual ACM Symposium on Theory of Computing* (Berkeley, California, USA) (STOC '86). ACM, New York, NY, USA, 69–76. <https://doi.org/10.1145/12130.12138>
- [80] Thaddeus D Ladd, Fedor Jelezko, Raymond Laflamme, Yasunobu Nakamura, Christopher Monroe, and Jeremy Lloyd O'Brien. 2010. Quantum computers. *Nature* 464, 7285 (2010), 45.
- [81] R. LaRose, A. Tikku, É. O'Neel-Judy, L. Cincio, and P. J. Coles. 2019. Variational quantum state diagonalization. *npj Quantum Information* 5, 1 (2019), 57. <https://doi.org/10.1038/s41534-019-0167-6>
- [82] R. J. Lipton and K. W. Regan. 2014. Quantum Algorithms via Linear Algebra. (2014).
- [83] Seth Lloyd, Silvano Garnerone, and Paolo Zanardi. 2015. Quantum Algorithms for Topological and Geometric Analysis of Data. *Nature Communications* (2015). <https://doi.org/10.1038/ncomms10138>
- [84] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. 2013. Quantum algorithms for supervised and unsupervised machine learning. *arXiv preprint arXiv:1307.0411* (2013).
- [85] Seth Lloyd, Masoud Mohseni, and Patrick Rebentrost. 2014. Quantum principal component analysis. *Nature Physics* 10, 9 (2014), 631–633. <https://doi.org/10.1038/nphys3029> arXiv:arXiv:1307.0401v2
- [86] Neil B Lovett, Sally Cooper, Matthew Everitt, Matthew Trevers, and Viv Kendon. 2010. Universal quantum computation using the discrete-time quantum walk. *Physical Review A* 81, 4 (2010), 042330.
- [87] Frederic Magniez, Miklos Santha, and Mario Szegedy. 2007. Quantum algorithms for the triangle problem. *SIAM J. Comput.* (2007), 413–424.
- [88] Enrique Martin-Lopez, Anthony Laing, Thomas Lawson, Roberto Alvarez, Xiao-Qi Zhou, and Jeremy L O'brien. 2012. Experimental realization of Shor's quantum factoring algorithm using qubit recycling. *Nature photonics* 6, 11 (2012), 773–776.
- [89] Jarrod R McClean, Jonathan Romero, Ryan Babbush, and Alán Aspuru-Guzik. 2016. The theory of variational hybrid quantum-classical algorithms. *New Journal of Physics* 18, 2 (2016), 023023.
- [90] Ashley Montanaro. 2016. Quantum algorithms: an overview. *npj Quantum Information* 2 (2016), 15023.
- [91] Michele Mosca. 2012. Quantum algorithms. In *Computational Complexity*. Springer, 2303–2333.
- [92] Michael A. Nielsen and Isaac L. Chuang. 2016. *Quantum Computation and Quantum Information*. Cambridge University Press, Cambridge, United Kingdom. 10th Anniversary Edition.
- [93] Bryan O'Gorman, William J. Huggins, Eleanor G. Rieffel, and K. Birgitta Whaley. 2019. Generalized swap networks for near-term quantum computing. *arXiv e-prints* (2019). arXiv:1905.05118.
- [94] Brian Olson, Irina Hashmi, Kevin Molloy, and Amarda Shehu. 2012. Basin Hopping as a General and Versatile Optimization Framework for the Characterization of Biological Macromolecules. *Advances in Artificial Intelligence 2012* (2012), 674832. <https://doi.org/10.1155/2012/674832>
- [95] Karl Pearson. 1901. On lines and planes of closest fit to systems of points in space. *Philosophical Magazine Series 6* 2, 11 (1901), 559–572. <https://doi.org/10.1080/14786440109462720>
- [96] Alberto Peruzzo, Jarrod McClean, Peter Shadbolt, Man-Hong Yung, Xiao-Qi Zhou, Peter J. Love, Alán Aspuru-Guzik, and Jeremy L. O'Brien. 2014. A variational eigenvalue solver on a photonic quantum processor. *Nature Communications* 5 (July 2014), ncomms5213. <https://doi.org/10.1038/ncomms5213>
- [97] Martin Plesch and Časlav Brukner. 2011. Quantum-state preparation with universal gate decompositions. *Physical Review A* 83, 3 (2011), 032302.
- [98] Martin Plesch and Časlav Brukner. 2011. Quantum-state preparation with universal gate decompositions. *Phys. Rev. A* 83 (2011), 032302. Issue 3. <https://doi.org/10.1103/PhysRevA.83.032302>
- [99] Carl Pomerance. 1996. A tale of two sieves. *Notices Amer. Math. Soc* 43 (1996), 1473–1485.
- [100] John Preskill. [n.d.]. Quantum computing and the entanglement frontier. *Rapporteur talk at the 25th Solvay Conference on Physics, 19-22 October 2011* ([n.d.]).
- [101] Bo Qi, Zhibo Hou, Li Li, Daoyi Dong, Guoyong Xiang, and Guangcan Guo. 2013. Quantum state tomography via linear regression estimation. *Scientific reports* 3 (2013).
- [102] Patrick Rebentrost, Masoud Mohseni, and Seth Lloyd. 2014. Quantum support vector machine for big data classification. *Physical review letters* 113, 13 (2014), 130503.
- [103] E. Rieffel and W. Polak. 2011. Quantum Computing: A Gentle Introduction. (2011).
- [104] R. L. Rivest, A. Shamir, and L. Adleman. 1978. A Method for Obtaining Digital Signatures and Public-key Cryptosystems. *Commun. ACM* 21, 2 (Feb. 1978), 120–126. <https://doi.org/10.1145/359340.359342>
- [105] Mehdi Saeedi and Igor L Markov. 2013. Synthesis and optimization of reversible circuits - a survey. *ACM Computing Surveys (CSUR)* 45, 2 (2013), 21.
- [106] Miklos Santha. 2008. Quantum walk based search algorithms. In *International Conference on Theory and Applications of Models of Computation*. Springer, 31–46.
- [107] N. Santhi. 2017. Quantum Netlist Compiler (QNC) software repository. <http://gitlab.lanl.gov/QuantumProgramming2017/QNC> Applied for LANL LACC authorization for unlimited open-source release, December 2017.

- [108] Maria Schuld, Ilya Sinayskiy, and Francesco Petruccione. 2015. An introduction to quantum machine learning. *Contemporary Physics* 56, 2 (2015), 172–185.
- [109] Jiangwei Shang, Zhengyun Zhang, and Hui Khoon Ng. 2017. Superfast maximum-likelihood reconstruction for quantum tomography. *Physical Review A* 95, 6 (2017), 062336.
- [110] Vivek V. Shende and Igor L. Markov. 2009. On the CNOT-cost of TOFFOLI gates. *Quant. Inf. Comp.* 9, 5-6 (2009), 461–486. arXiv:0803.2316 <http://arxiv.org/abs/0803.2316>
- [111] Vivek V Shende, Igor L Markov, and Stephen S Bullock. 2004. Minimal universal two-qubit controlled-NOT-based circuits. *Physical Review A* 69, 6 (2004), 062321.
- [112] Neil Shenvi, Julia Kempe, and K Birgitta Whaley. 2003. Quantum random-walk search algorithm. *Physical Review A* 67, 5 (2003), 052307.
- [113] Peter W Shor. 1994. Algorithms for quantum computation: Discrete logarithms and factoring. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*. IEEE, 124–134. <https://doi.org/10.1109/SFCS.1994.365700>
- [114] Peter W. Shor. 1997. Polynomial-Time Algorithms for Prime Factorization and Discrete Logarithms on a Quantum Computer. *SIAM J. Comput.* 26, 5 (1997), 1484–1509. <https://doi.org/10.1137/S0097539795293172> arXiv:<https://doi.org/10.1137/S0097539795293172>
- [115] Nikolai A Sinitsyn. 2018. Computing with a single qubit faster than the computation quantum speed limit. *Physics Letters A* 382, 7 (2018), 477–481.
- [116] Robert S. Smith, Michael J. Curtis, and William J. Zeng. 2016. A Practical Quantum Instruction Set Architecture. arXiv:arXiv:1608.03355
- [117] John A Smolin, Jay M Gambetta, and Graeme Smith. 2012. Efficient method for computing the maximum-likelihood quantum state from measurements with additive gaussian noise. *Physical review letters* 108, 7 (2012), 070502.
- [118] Rolando D Somma. 2016. Quantum simulations of one dimensional quantum systems. *Quantum Information & Computation* 16, 13-14 (2016), 1125–1168.
- [119] Robert Spalek et al. 2006. *Quantum algorithms, lower bounds, and time-space tradeoffs*. ILLC, Amsterdam.
- [120] V. Strassen. 1969. Gaussian elimination is not optimal. *Numer. Math.* 13 (1969), 354–356.
- [121] Yiğit Subaşı, Rolando D Somma, and Davide Orsucci. 2019. Quantum algorithms for systems of linear equations inspired by adiabatic quantum computing. *Physical review letters* 122, 6 (2019), 060504.
- [122] J. A. K. Suykens and J. Vandewalle. 1999. Least Squares Support Vector Machine Classifiers. *Neural Process. Lett.* 9, 3 (June 1999), 293–300. <https://doi.org/10.1023/A:1018628609742>
- [123] Mario Szegedy. 2004. Quantum speed-up of Markov chain based algorithms. In *45th Annual IEEE symposium on foundations of computer science*. IEEE, 32–41.
- [124] Maika Takita, Antonio D Córcoles, Easwar Magesan, Baleegh Abdo, Markus Brink, Andrew Cross, Jerry M Chow, and Jay M Gambetta. 2016. Demonstration of weight-four parity measurements in the surface code architecture. *Physical review letters* 117, 21 (2016), 210505.
- [125] IBM QX Team. 2017. IBM Q experience backend information. <http://github.com/QISKit/ibmqx-backend-information>. Last accessed: 12 December, 2017.
- [126] Giacomo Torlai, Guglielmo Mazzola, Juan Carrasquilla, Matthias Troyer, Roger Melko, and Giuseppe Carleo. 2018. Neural-network quantum state tomography. *Nature Physics* 14, 5 (2018), 447.
- [127] L. M. K. Vandersypen, M. Steffen, G. Breyta, C. S. Yannoni, M. H. Sherwood, and I. L. Chuang. 2001. Experimental realization of Shor’s quantum factoring algorithm using nuclear magnetic resonance. *Nature* 414 (Dec. 2001), 883–887. <https://doi.org/10.1038/414883a> arXiv:quant-ph/0112176
- [128] Guifré Vidal. 2003. Efficient classical simulation of slightly entangled quantum computations. *Physical review letters* 91, 14 (2003), 147902.
- [129] Zhihui Wang, Stuart Hadfield, Zhang Jiang, and Eleanor G. Rieffel. 2018. Quantum approximate optimization algorithm for MaxCut: A fermionic view. *Physical Review A* 97 (2018), 022304. Issue 2. <https://doi.org/10.1103/PhysRevA.97.022304>
- [130] Zhihui Wang, Nicholas C. Rubin, Jason M. Dominy, and Eleanor G. Rieffel. 2020. XY mixers: Analytical and numerical results for the quantum alternating operator ansatz. *Physical Review A* 101, 1 (2020), 012320. <https://doi.org/10.1103/PhysRevA.101.012320>
- [131] Chu Ryang Wie. 2017. A quantum circuit to construct all maximal cliques using Grover’s search algorithm. (2017), 1–13. <https://arxiv.org/abs/1711.06146>
- [132] James R Wootton. 2017. Demonstrating non-Abelian braiding of surface code defects in a five qubit experiment. *Quantum Science and Technology* 2, 1 (2017), 015006.
- [133] James R Wootton. 2020. Benchmarking near-term devices with quantum error correction. *Quantum Sci. Technol.* 5, 4 (2020), 044004.
- [134] James R Wootton and Daniel Loss. 2018. Repetition code of 15 qubits. *Physical Review A* 97, 5 (2018), 052313.
- [135] Zhi-Cheng Yang, Armin Rahmani, Alireza Shabani, Hartmut Neven, and Claudio Chamon. 2017. Optimizing Variational Quantum Algorithms Using Pontryagin’s Minimum Principle. *Physical Review X* 7, 2 (2017), 021027. <https://doi.org/10.1103/physrevx.7.021027>
- [136] N. S. Yonofsky and M. A. Mannucci. 2008. *Quantum Computing for Computer Scientists*. (2008).
- [137] Yarui Zheng, Chao Song, Ming-Cheng Chen, Benxiang Xia, Wuxin Liu, Qiujiang Guo, Libo Zhang, Da Xu, Hui Deng, Keqiang Huang, et al. 2017. Solving Systems of Linear Equations with a Superconducting Quantum Processor. *Physical Review Letters* 118, 21 (2017),

- 210504.
- [138] Leo Zhou, Sheng-Tao Wang, Soonwon Choi, Hannes Pichler, and Mikhail D Lukin. 2020. Quantum Approximate Optimization Algorithm: Performance, Mechanism, and Implementation on Near-Term Devices. *Physical Review X* 10, 2 (2020), 021067. <https://doi.org/10.1103/PhysRevX.10.021067>

Just Accepted