Survey on classical quantum computer simulators

Miguel Murça
Instituto Superior Técnico, University of Lisbon, Portugal
(Dated: July 12, 2023)

I. INTRODUCTION

A. Quantum computation fundamentals

1. State vector

Quantum computation is the field of study concerned with computation performed within the postulates imposed by the theory of quantum mechanics. The notion of "computation", as here used, is the same as in the field of (classical) computability and complexity theory, and so we refer to, for example, Ref. [1] for a formal overview of the topic. However, for simplicity, computation can be taken to mean the evaluation of some function over some discrete input range. In any case, this function (or, more generally, the computational task at hand) will be made concrete in context.

The postulates of quantum mechanics, on the other hand, are well defined and may succinctly be stated as follows [2]:

State space: An isolated quantum system is represented by a ray of unit two-norm vectors in a complex vector space equipped with an inner product (i.e., a Hilbert space). This space is referred to as the *state space*.

Evolution: The evolution of an isolated quantum system, as represented by its quantum state vector $\vec{\psi}$, is given by a unitary operator acting on that quantum state vector. I.e., let $\mathbb{H}$ be the Hilbert space such that the quantum state vector is, at some moment, $\vec{\psi} \in \mathbb{H}$. Then, there exists an operator $U \in \mathbb{H} \times \mathbb{H}$, satisfying $UU^\dagger = U^\dagger U = I$ [1] such that, after some time, the state of the system is given by the quantum state vector $U\vec{\psi}$.

Measurement: A set of operators $\{M_m\}_m$, acting on $\mathbb{H}$ and satisfying $\sum_m M_m^\dagger M_m = I$, define a "measurement". Each operator M_m has an associated outcome m, such that the measurement operation yields outcome m with probability $\|M_m\vec{\psi}\|_2^2$. After the measurement, the quantum system is described by the quantum state vector $M_m\vec{\psi}/\|M_m\vec{\psi}\|_2$.

Composite Systems: If a quantum system is composed of multiple quantum subsystems, then the corresponding state vector space is given by the tensor product of the quantum state vector spaces of the subsystems.

$\|\psi\rangle$	Vector ψ, or "ket" ψ. Corresponds to $\vec{\psi}$.
$\langle\psi\|$	Dual of $\|\psi\rangle$, or "bra" ψ. Corresponds to $\vec{\psi}^\dagger$.
$\langle a\|b\rangle$	Inner product between vectors $\|a\rangle$ and $\|b\rangle$.
$\|a\rangle \otimes \|b\rangle$	Tensor product between vectors $\|a\rangle$ and $\|b\rangle$. If both $\|a\rangle, \|b\rangle \in \mathbb{H}_n$, $\|a\rangle \otimes \|b\rangle \in \mathbb{H}_{n^2}$.
$\|a\rangle\|b\rangle$	Shortened notation for $\|a\rangle \otimes \|b\rangle$.
$\langle a\| A \|b\rangle$	Inner product between $\|a\rangle$ and $A\|b\rangle$, or, equivalently, $A^\dagger\|a\rangle$ and $\|b\rangle$. If $\|a\rangle = \|b\rangle$, may be referred to as the *expectation value* of A under $\|a\rangle$.
$\|a\rangle\langle a\|$	Projector onto the span of $\|a\rangle$.

Table I. Summary of "Dirac notation", or "bra-ket notation".

A quantum computation is, thus, a computation carried out within these postulates. Taking the computational task to be a decision problem (i.e., a question to which the computer should output a "yes" or "no" answer), we may, without loss of generality, consider that the final answer is produced by a final measurement (as defined in the measurement postulate), with outcomes "yes" ($m = 1$) or "no" ($m = 0$).

In the classical case, it is well known that a binary alphabet – the "bit" – is sufficient to perform computation (in the sense that it requires only a logarithmic overhead in comparison to a larger alphabet; see [3, Claim 1.5]). To perform quantum computation, we will likewise work with a quantum analogue of the bit, the "qubit". In particular, a qubit is a quantum state of a Hilbert space of dimension 2, $\mathbb{H}_2$. For concreteness, we choose two quantum state vectors of $\mathbb{H}_2$ that are orthogonal,

$$|0\rangle, |1\rangle \tag{1}$$

and that thus form a basis of $\mathbb{H}_2$, the "computational basis". We've also here introduced the so-called "Dirac notation", or "bra-ket notation", common in quantum mechanics, and by extension in quantum computing works. We present a summary of Dirac notation in Table I, and we will henceforth use this notation.

In analogy to how, in the classical case, a binary alphabet can represent larger alphabets, multiple qubits can

[1] The dagger symbol ($\dagger$) is used to denote conjugate transposition.

be used to span a larger Hilbert space, by the *composite systems* postulate. Indeed, n qubits span 2^n orthogonal states in their collective state space, which we may label by the binary string given by each of the qubits, or the corresponding number:

$$\begin{aligned}
|0\rangle|0\rangle\ldots|0\rangle|0\rangle &\equiv |0_b\rangle, \\
|0\rangle|0\rangle\ldots|0\rangle|1\rangle &\equiv |1_b\rangle, \\
|0\rangle|0\rangle\ldots|1\rangle|0\rangle &\equiv |2_b\rangle, \\
&\ldots \\
|1\rangle|1\rangle\ldots|1\rangle|1\rangle &\equiv |2^n-1_b\rangle.
\end{aligned} \tag{2}$$

Per the *state space* postulate, a state of n qubits may be given by a linear combination of these basis states:

$$|\psi\rangle = \sum_{j=0}^{2^n-1} \alpha_j \, |j_b\rangle \tag{3}$$

$$\alpha_j \in \mathbb{C}, \quad \sum_{j=0}^{2^n-1} |\alpha_j|^2 = 1$$

If more than one α_j is non-zero, we say the state is in "superposition". Measuring a state in superposition produces different outcomes, depending on the state's overlap with the outcome state. To see this, consider the measurement

$$\{M_m = |m_b\rangle\langle m_b| \qquad m = 0,\ldots,2^n-1\}. \tag{4}$$

Since every $|j_b\rangle$ has unit norm, and $\{|j_b\rangle\}_{j=0,\ldots,2^n-1}$ span the Hilbert space being considered, this set satisfies the conditions outlined in the *measurement* postulate. We conclude also from the postulate that the probability of observing outcome j is given by

$$\Pr_{|\psi\rangle}[j] = |\alpha_j|^2. \tag{5}$$

2. Density matrix

Consider the measurement (4) on state (3). After performing such a measurement, one holds state $|j_b\rangle$ with probability $\Pr_{|\psi\rangle}[j]$. This is not correctly described by a superposition. To see this, suppose a single qubit, and a Hadamard gate:

$$H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \qquad \begin{aligned} H\,|0\rangle &= \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle) \\ H\,|1\rangle &= \tfrac{1}{\sqrt{2}}(|0\rangle - |1\rangle) \end{aligned} \tag{6}$$

From this definition and the previous discussion,

$$\begin{aligned} \Pr_{H|0\rangle}[0] &= \Pr_{H|1\rangle}[0] = 1/2 \\ \Pr_{H|0\rangle}[1] &= \Pr_{H|1\rangle}[1] = 1/2 \end{aligned} \tag{7}$$

and we note that this is also true of the states

$$|+\rangle = \frac{|0\rangle + |1\rangle}{\sqrt{2}} \qquad |-\rangle = \frac{|0\rangle - |1\rangle}{\sqrt{2}}. \tag{8}$$

We conclude from Eq. (7) that by taking a $|0\rangle$ state, applying a Hadamard gate, measuring, and again applying a Hadamard gate and measuring, we should observe 0 with probability 1/2, and likewise 1 with probability 1/2. But, following the postulates and attempting to describe the intermediate situation by either $|+\rangle$ or $|-\rangle$, we find that

$$\begin{aligned} H\,|+\rangle &= |0\rangle \\ H\,|-\rangle &= |1\rangle \end{aligned} \tag{9}$$

which would indicate either $\Pr[0] = 1$ or $\Pr[1] = 1$.

Indeed, the system may be described by one of several states not in superposition, while we are uncertain about which which state describes it. The density matrix formalism (or *density operator* formalism) gives a formal tool for describing this situation [2]. If a quantum state is in state $|\psi_j\rangle$ with probability p_j, the associated density operator is

$$\rho = \sum_j p_j \, |\psi_j\rangle\langle\psi_j| \, . \tag{10}$$

In general, to be a valid density operator, ρ must have unit trace and be a positive operator. In particular, the density operator associated to a state vector $|\psi\rangle$ is

$$\rho = |\psi\rangle\langle\psi| \, , \tag{11}$$

and a statistical mixture between multiple ρ_j with corresponding probability p_j is given by the density operator

$$\rho = \sum_j p_j \rho_j. \tag{12}$$

Density matrices are a "complete" formalism, in the sense that we may rephrase the postulates of quantum mechanics only in terms of the density matrix operator; in fact one may check that the two sets of postulates are equivalent.

State space *(Density matrix)***:**
An isolated physical system is completely described by a *density operator*, which is a positive operator of unit trace in a Hilbert space. If a system is in state ρ_j with probability p_j, its density operator is $\rho = \sum_j p_j \rho_j$.

Evolution *(Density matrix)***:**
Let a quantum system be described, at some moment, by the density operator ρ. Then, there exists a unitary operator U such that, after some time, the quantum system is now described by the density operator $\rho' = U\rho U^\dagger$.

---

[2] Alternatively, the density matrix formalism allows one to express both "quantum randomness", as resulting from the measurement of a state in superposition, and "classical randomness", in the sense of operations conditioned on a random/unknown bit string.

Measurement *(Density matrix)*:
A set of operators $\{M_m\}_m$, acting on the space of ρ, and satisfying $\sum_m M_m^\dagger M_m$, define a "measurement". Each operator M_m has an outcome m associated to it. If a quantum state is described by a density operator ρ, the outcome of a measurement is m with probability $\mathrm{Tr}\big(M_m^\dagger M_m \rho\big)$, and, after the measurement, the system is now described by the density matrix $\rho' = M_m \rho M_m^\dagger / \mathrm{Tr}\big(M_m^\dagger M_m \rho\big)$.

Composite Systems *(Density matrix)*:
The density operator describing a quantum system composed of multiple quantum subsystems is given by the tensor product of the density operators of each of the subsystems.

A quantum state with density matrix ρ for which there exists

$$|\psi\rangle \quad \text{such that} \quad \rho = |\psi\rangle\langle\psi| \tag{13}$$

is said to be a *pure* state, while a state that cannot satisfy this is said to be a *mixed* state.

A density operator may be used to describe a quantum subsystem: if a system is composed of subsystems A and B, jointly described by the density operator ρ, then subsystem A is described by density operator

$$\rho_A = \mathrm{Tr}_B(\rho) = \sum_j (I_A \otimes \langle j|_B)\rho(I_A \otimes |j\rangle_B) \tag{14}$$

where Tr_B is the newly defined *partial trace* operation, I_A is the identity in the state space of A, and we take $\{|j\rangle\}_j$ to be a basis over the state space of B.

3. *Quantum circuits*

To conclude this section, we introduce a common notation to denote unitary transformations, and by extension quantum algorithms: quantum circuits.

Recall that a quantum computation may be described by a sequence of unitary evolutions and measurements, and a final measurement. While the measurements correspond (by the *measurement* postulate) to a physical procedure, it is not necessarily clear how to implement a given unitary transformation[3]. Instead, we assume the ability to physically realize a number of elementary operations, and compose these operations to build more sophisticated unitary evolutions. Table II lists some common basic operations. Critically, a limited set of these elementary operations can be sufficient to express any unitary evolution. I.e., the notion of a *universal quantum gate set* is well defined. The set of gates specified in Table II form a universal quantum gate set. One could even reduce the size of this set while maintaining universality; for example, the set of $\{R_X, R_Y, R_Z, {}_CX\}$ gates is also universal [2]. A common choice of universal gate set is the "Clifford+T" set, where the Clifford gate set, $\{S, H, {}_CX\}$, is augmented with the T gate. Note that the Clifford gate set generates all the Pauli gates. The term "Clifford group" is used to denote the set of all operations generated by the Clifford gate set.

Symbol	Definition	Description
X	$X\|0\rangle = \|1\rangle$ $X\|1\rangle = \|0\rangle$	X-Pauli or "not" gate
Y	$Y\|0\rangle = -i\|1\rangle$ $Y\|1\rangle = i\|0\rangle$	Y-Pauli gate
Z	$Z\|0\rangle = \|0\rangle$ $Z\|1\rangle = -\|1\rangle$	Z-Pauli gate
H	$H\|0\rangle = \sqrt{1/2}(\|0\rangle + \|1\rangle)$ $H\|1\rangle = \sqrt{1/2}(\|0\rangle - \|1\rangle)$	Hadamard gate
S	$S\|0\rangle = \|0\rangle$ $S\|1\rangle = i\|1\rangle$	Phase gate
T	$T\|0\rangle = \|0\rangle$ $T\|1\rangle = e^{i\pi/4}\|1\rangle$	$\pi/8$ gate
$R_X(\theta)$	$\exp\{-i\theta X/2\}$	X-rotation gate
$R_Y(\theta)$	$\exp\{-i\theta Y/2\}$	Y-rotation gate
$R_Z(\theta)$	$\exp\{-i\theta Z/2\}$	Z-rotation gate
${}_CX$	$\|0\rangle\langle 0\| \otimes I + \|1\rangle\langle 1\| \otimes X$	Controlled-not gate

Table II. Common "native" operations, often assumed to be physically realizable, to be composed into other operations.

Quantum circuits provide a visual notation to denote composition of elementary gates into larger unitary evolutions and measurements. The main elements of a quantum circuit are given in Table III. Assuming every operation in the elementary gate set can be performed in a time step, it follows that the number of vertical slices in a quantum circuit correspond to the running time of the circuit. This is referred to as the circuit's *depth*. The circuit's *width* is the number of qubits acted upon non-trivially by the circuit, and relates to the space requirements of the circuit.

Because a quantum algorithm corresponds to known unitary evolutions and measurements, it is expressible in quantum circuit form. We say, then, that a quantum algorithm is efficient if the corresponding quantum circuit's width and depth scale at most polynomially with input size.

[3] Rigorously, this depends on how the unitary is "given". Here, consider that a unitary is given by specifying its action over each element $|j\rangle$ of a set spanning the state space. By linearity (cf. the *evolution* postulate), this determines the action of the unitary over any vector in the state space.

Symbol	Definition	Description
	Wire	Reperesents the state space associated to a qubit, i.e., $\mathbb{H}_2$. Multiple wires, vertically aligned, represent the collective Hilbert space, given by the tensor product of the individual $\mathbb{H}_2$ spaces (cf. the *composite systems* postulate).
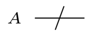	Register	Represents multiple wires, i.e., a subspace of dimension 2^n, where n is the number of wires in the register. As shown, a register may be labelled.
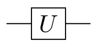	Gate	Unitary U acting on the state space associated to the incoming wire (left-hand side). The outgoing wire (right-hand side) corresponds to the state space after the action of the gate.
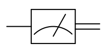	Measurement	Measurement of the state subspace associated to the incoming wire (left-hand side), according to the measurement $\{(\|0\rangle\langle 0\| , 0), (\|1\rangle\langle 1\| , 1)\}$ ("computational basis measurement"). The double wire represents the resulting classical bit (i.e., a $\mathbb{Z}_2$ space).
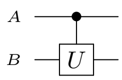	Controlled operation	Denotes the operation $\|0\rangle\langle 0\|_A \otimes I_B + \|1\rangle\langle 1\|_A \otimes U_B$ acting on the collective state space of A and B.
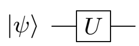	Application	The quantum state resulting from application of the unitary represented by the quantum circuit on the quantum state specified on the left-hand side. The resulting state may be denoted on the right-hand side. Here, represents the state $U\|\psi\rangle$.
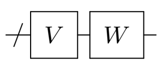	Composition	Circuits are "read" left-to-right. Thus, the concatenation of two circuits denoting the actions of operators, respectively, V and W, results in a circuit denoting the operator WV.

Table III. Quantum circuit notation.

B. Simulation theory

The quantum computational model, as introduced in the previous section, is believed to be more powerful than the classical computing model [4]. One canonical example of evidence for this separation is the existence of Shor's Algorithm for efficient factoring [5, 6]. Thus, the task of classically simulating a quantum computation becomes doubly significant: on the one hand, an efficient classical algorithm to simulate an arbitrary quantum computation would have a fundamental impact in the current understanding of computational complexity (but, for this reason, is not expected to exist). On the other hand, it is necessary to ensure that a particular instance of a quantum circuit cannot be classically simulated in practice in order to claim that a quantum computation has been carried out "with advantage" (i.e., beyond a classical regime, or what is often referred to as "quantum supremacy") [7–10]. Finally, even in the quantum advantage regime, small-scale classical simulations remain relevant as a source of reference data for validation [11].

As discussed in section I A 1, a quantum state of n qubits is determined by a vector of 2^n complex values, up to a global phase factor, and subject to a normalization constraint. Therefore, on first approach, one may believe that a quantum circuit of more than 40–45 qubits cannot be classically simulated, simply due to the memory requirements of maintaining a quantum state vector. However, this consideration ignores the details of any particular problem instance, such as the existence of structure or constraints in the quantum circuit to be ran, or specifications on the desired output. In this section, we review some key theoretical results regarding classical simulation of quantum computations, as well as some general techniques.

1. *Strong simulation vs. weak simulation*

As defined in section I A, the final output of a quantum computation results from a measurement. Due to the nature of measurements, the outcome is a random variable, and its distribution depends on the underlying state vector before the measurement. Thus, should a classical simulation of a quantum algorithm:

i. generate a random outcome observing the same output distribution as the quantum counterpart, or;

ii. explicitly specify the distribution of the generated output?

These two problem specifications correspond, respectively, to the notions of *weak simulation* and *strong simulation.* Requiring either strong or weak simulation may significantly affect the computational hardness of the task; indeed there exist circuits for which classical weak simulation is easy, but classical strong simulation is hard [12]. Stating these two notions more formally:

Definition 1 (Strong simulation [12])**.** *Given a description of a quantum circuit of n qubits, which corresponds to unitary operator U, terminating with a measurement of the first qubit in the computational basis, output* $\mathrm{Tr}\big(|0\rangle\langle 0|\ U|0_b\rangle\langle 0_b|\,U^\dagger\big)$.

Definition 2 (Weak simulation [12])**.** *Given a description of a quantum circuit of n qubits, which corresponds to unitary operator U, terminating with a measurement of the first qubit in the comptuational basis, output* 0 *with probability* $\mathrm{Tr}\big(|0\rangle\langle 0|\ U|0_b\rangle\langle 0_b|\,U^\dagger\big)$, *or* 1 *otherwise.*

In practice, one may wish to simulate the circuit only up to some point, or to inspect the intermediate state of a small-scale computation [11]. This motivates a different definition of strong simulation by some authors. Namely, recalling Eq. (3), note that knowledge of α_j is enough to determine the probability of observing any measurement in the computational basis. However, the converse is not true: because $\Pr[j] = |\alpha_j|^2$, knowledge of $\Pr[j]$ fails to inform about the complex phase of α_j[4]. So, strong simulation may be taken to mean:

Definition 3 (Strong simulation, wave-function version [13])**.** *Given a quantum circuit of n qubits, corresponding to a unitary operator U, and an n-bit string j, output* $\langle j_b|\,U\,|0_b\rangle$.

Note that, in all of the definitions above, we took the quantum circuits to be described by a unitary operator, which may not be trivially true if measurements are performed half-way in the circuit (cf. section I A 2), or if classical post-processing is employed. However, a well-known result, which we review in appendix A, allows us to defer all measurements to the end of the circuit, such that the whole of the computation is carried out unitarily.

2. Clifford circuits and the Gottesman-Knill theorem

Recall that Clifford gates are gates in the Clifford set, i.e., any quantum circuit that can be written in terms of phase, Hadamard, and Controlled-Not gates (cf. Table II). Then, the Gottesman-Knill theorem states the following:

[4] By the *state space* postulate, a global phase factor is physically irrelevant. However, relative phase differences should not be disregarded. As a simple example, consider the action of the Hadamard gate (eqs. (7),(8)): the only difference between the $|+\rangle$ and $|-\rangle$ states is the relative phase difference between the $|0\rangle$ and $|1\rangle$ components – however, the result from acting with the Hadamard gate is completely different.

Theorem 1 (Gottesman-Knill [14])**.** *Every (uniform family of) Clifford circuit(s), when applied to the input state* $|0_b\rangle \equiv |0\rangle|0\rangle\ldots|0\rangle$, *and when followed by a computational basis measurement of the first qubit, can be efficiently simulated classically in the strong sense.*

The theorem is constructive, and in Ref. [15] Gottesman and Aaronson provide a high-performance (weak) simulator of Clifford circuits that can scale up to tens of thousands of qubits. Van den Nest [12] gives an alternative derivation of the theorem that allows for direct strong simulation as well (both the regular and wave-function version).

Despite this result, Clifford circuits very easily extend to the universality regime: not only by augmentation of the gate set – the Clifford+T gate set is already universal – but also by choice of the input state. Indeed, there exist "magic states", such that a supply of these (pre-prepared) quantum states and Clifford operations are enough to perform universal quantum computation [16].

This move from Clifford-based computation to quantum universality entails a "jump", since Clifford circuits are not as powerful as classical circuits. In Ref. [12], this computational gap is discussed and eliminated, by giving a superclass of Clifford circuits, "HT circuits", that is equivalent to classical computation and can be weakly simulated.

3. Schrödinger simulation

Schrödinger simulation refers to the straightforward approach of maintaining the global state-vector, updating it as new unitary operations are encountered [17]. As such, it is inherently a form of the wave-function version of strong simulation (definition 3). This method also requires, by definition, that 2^n complex values are maintained for a state of n qubits, such that it cannot physically scale beyond a certain number of qubits (about 45-50 qubits, corresponding to a petabyte or more of memory, if each amplitude is represented within 8 bytes; barring adaptative models admitting error, such as in Ref. [17]).

Despite this constraint, a significant amount of research has been devoted to Schrödinger simulation in the memory-tractable regime ($n \lesssim 50$), specifically in pushing this limit and speeding up the running time [9, 10, 17–23].

A key observation is that, if each gate is considered at a time, the matrix-vector products being calculated are of very sparse and strongly structured matrices. Namely, the gates are *local*, in the sense that they involve non-trivially at most a small number k of qubits. For example, in Ref. [9], this is used to ensure that compute resources are maximally utilized via parallelization. Following a different strategy, in Ref. [21], advance knowledge of the action of common blocks of operations in quantum computing is used to speed up over the simulation of each gate individually. Gate clustering, different encoding techniques and cache-related considerations, as well as employment of compiler intrinsics, also allow for speed improvements [9, 10, 19, 21].

Otherwise, the problem may be regarded as a classical large-sparse-matrix and vector product, a well-researched problem (see, e.g., Ref. [24]).

4. Feynman simulation

The Feynman simulation method [13, 25, 26] trades the 2^n memory requirement by an exponential time computation, but in linear space. Being also a form of the wave-function version of strong simulation (definition 3), the Feynman simulation method follows from noticing the following:

$$\begin{aligned} &\langle x_b| U_L U_{L-1} U_{L-2} \cdots U_2 U_1 |0_b\rangle = \\ &= \langle x_b| U_L (\Sigma_{j=0}^{2^n-1} |j_b\rangle\langle j_b|) U_{L-1} (\Sigma_{j'=0}^{2^n-1} |j'_b\rangle\langle j'_b|) \\ &\qquad U_{L-2} \cdots U_2 (\Sigma_{j''=0}^{2^n-1} |j''_b\rangle\langle j''_b|) U_1 |0_b\rangle = \\ &= \sum_{\{y_{(t)}\} \in \{0,1,\ldots,2^n-1\}^{L-1}} \prod_{t=0}^{L-1} \langle y_{(t+1)}| U_t |y_{(t)}\rangle \end{aligned} \tag{15}$$

since $\{|j_b\rangle\}_{j=0,\ldots,2^n-1}$ form an orthonormal basis of the state vector space, and letting $|y_{(L)}\rangle \equiv |x_b\rangle$. Now, take each of the U_t to be the (local, sparse) unitary corresponding to a quantum gate in a quantum circuit, such that L is the depth of the quantum circuit. One may conclude this scheme requires $\mathcal{O}(n \cdot (2d)^{n+1})$ time and $\mathcal{O}(n \log d)$ space [13].

This approach may be interpreted as a discrete version of the Feynman path integral formulation, where, simply put, every possible "computation path" (corresponding to a choice of $\{y_{(t)}\}$) is considered separately, in order to determine the resulting constructive or destructive interference between the paths; each path requires a linear amount of memory to compute, but there are exponentially many paths to consider, which interfere among themselves. This is illustrated in figure 1.

5. Schrödinger-Feynman simulation

It is possible to establish an intermediate scheme between Schrödinger simulation (section I B 3) and Feynman simulation (section I B 4) [13, 25, 27]. Thus, this scheme, which allows for a controllable trade-off between space and time complexity, is referred to by some authors as Schrödinger-Feynman simulation [19, 25, 28].

The main idea of the technique is to divide the quantum circuit into disjoint registers, performing Schrödinger simulation for operations that involve only

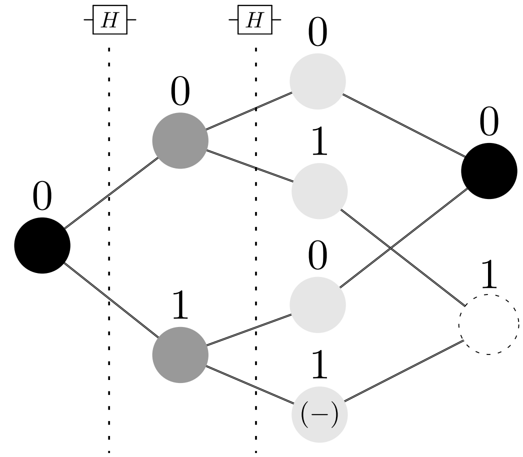

Figure 1. The action of two Hadamard gates (eq. 6) acting on a single qubit initialized to $|0\rangle$, as a trivial example of interference. Each node's tone reflects the absolute value of the associated amplitude: darker corresponds to greater amplitude. The node with a negative amplitude is marked with a $(-)$. In a Feynman path integral interpretation (see section I B 4), each of the drawn paths is considered separately, and then summed, resulting in the destructive interference of the $|1\rangle$ state.

qubits in the same register, but summing over "paths" resulting from operations across registers. This allows the computation to be distributed (across different choices of computation paths for cross-register operations), while maximally exploiting the memory available to each process.

The choice of (maximum) size of the registers determines the memory consumption, with bigger sized registers requiring more memory but less computational paths to consider. Thus, for at-most k-qubit sized registers in an n-qubit circuit of depth d, one requires $\mathcal{O}(n2^{n-k} \cdot (2d)^{k+1})$ time, and $\mathcal{O}(2^{n-k} \log d)$ space [13].

This method is well suited for simulating circuits with a grid-like connectivity graph between circuits, as is common in practical implementations [7, 29, 30], and so is used in multiple works pushing the boundary of quantum advantage [9, 18, 27].

6. Tensor network simulation

TODO

Appendix A: Delayed measurement

The delayed measurement lemma states:

Lemma 1 (Delayed measurement [2])**.**

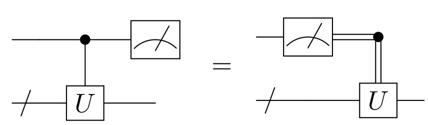

(A1)

i.e., measurements can always be moved from an intermediate stage of a quantum circuit to the end of the circuit, replacing conditional classical operations by controlled quantum operations.

This statement may be proven by explicitly calculating the density matrix resulting from the action of each circuit, for an arbitrary input, and checking that the result is the same. It follows that, when speaking of a quantum algorithm, one may always take the procedure to be described by a unitary operation followed by measurements.

---

[1] S. Homer and A. L. Selman, *Computability and Complexity Theory* (Springer New York, 2001).

[2] M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, 2012).

[3] S. Arora and B. Barak, *Computational Complexity: A Modern Approach* (Cambridge University Press, 2009).

[4] E. Bernstein and U. Vazirani, Quantum complexity theory, in *Proceedings of the twenty-fifth annual ACM symposium on Theory of computing - STOC '93* (ACM Press, 1993).

[5] P. Shor, Algorithms for quantum computation: discrete logarithms and factoring, in *Proceedings 35th Annual Symposium on Foundations of Computer Science* (IEEE Comput. Soc. Press, 1994).

[6] A. Y. Kitaev, Quantum measurements and the abelian stabilizer problem, arXiv:9511026 [quant-ph] (1995), unpublished.

[7] F. Arute, K. Arya, R. Babbush, D. Bacon, J. C. Bardin, R. Barends, R. Biswas, S. Boixo, F. G. S. L. Brandao, D. A. Buell, B. Burkett, Y. Chen, Z. Chen, B. Chiaro, R. Collins, W. Courtney, A. Dunsworth, E. Farhi, B. Foxen, A. Fowler, C. Gidney, M. Giustina, R. Graff, K. Guerin, S. Habegger, M. P. Harrigan, M. J. Hartmann, A. Ho, M. Hoffmann, T. Huang, T. S. Humble, S. V. Isakov, E. Jeffrey, Z. Jiang, D. Kafri, K. Kechedzhi, J. Kelly, P. V. Klimov, S. Knysh, A. Korotkov, F. Kostritsa, D. Landhuis, M. Lindmark, E. Lucero, D. Lyakh, S. Mandrà, J. R. McClean, M. McEwen, A. Megrant, X. Mi, K. Michielsen, M. Mohseni, J. Mutus, O. Naaman, M. Neeley, C. Neill, M. Y. Niu, E. Ostby, A. Petukhov, J. C. Platt, C. Quintana, E. G. Rieffel, P. Roushan, N. C. Rubin, D. Sank, K. J. Satzinger, V. Smelyanskiy, K. J. Sung, M. D. Trevithick, A. Vainsencher, B. Villalonga, T. White, Z. J. Yao, P. Yeh, A. Zalcman, H. Neven, and J. M. Martinis, Quantum supremacy using a programmable superconducting processor, Nature **574**, 505 (2019).

[8] B. M. Terhal, Quantum supremacy, here we come, Nature Physics **14**, 530 (2018).

[9] T. Häner and D. S. Steiger, 0.5 petabyte simulation of a 45-qubit quantum circuit, in *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (ACM, 2017).

[10] I. Markov, A. Fatima, S. Isakov, and S. Boixo, Quantum supremacy is both closer and farther than it appears, arXiv:1807.10749 [quant-ph] (2018), unpublished.

[11] S. Bravyi and D. Gosset, Improved classical simulation of quantum circuits dominated by clifford gates, Physical Review Letters **116**, 10.1103/physrevlett.116.250501 (2016).

[12] M. V. den Nest, Classical simulation of quantum computation, the gottesman-knill theorem, and slightly beyond, Quantum Information and Computation **10**, 258 (2010).

[13] S. Aaronson and L. Chen, Complexity-theoretic foundations of quantum supremacy experiments, in *Proceedings of the 32nd Computational Complexity Conference*, CCC '17 (Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, DEU, 2017).

[14] S. Corney, R. Delbourgo, and P. Jarvis, *Group22: Proceedings of the XXII International Colloquium on Group Theoretical Methods in Physics, Hobart, July 13-17, 1998*, International Press lectures and conference proceedings in physics (International Press, 1999) pp. 32–43.

[15] S. Aaronson and D. Gottesman, Improved simulation of stabilizer circuits, Phys. Rev. A **70**, 052328 (2004).

[16] S. Bravyi and A. Kitaev, Universal quantum computation with ideal clifford gates and noisy ancillas, Physical Review A **71**, 10.1103/physreva.71.022316 (2005).

[17] H. D. Raedt, F. Jin, D. Willsch, M. Willsch, N. Yoshioka, N. Ito, S. Yuan, and K. Michielsen, Massively parallel quantum computer simulator, eleven years later, Computer Physics Communications **237**, 47 (2019).

[18] M. Smelyanskiy, N. P. D. Sawaya, and A. Aspuru-Guzik, qhipster: The quantum high performance software testing environment, arXiv:1601.07195 [quant-ph] (2016), unpublished.

[19] A. Fatima and I. L. Markov, Faster schrödinger-style simulation of quantum circuits, in *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)* (IEEE, 2021).

[20] T. Jones, A. Brown, I. Bush, and S. C. Benjamin, QuEST and high performance simulation of quantum computers, Scientific Reports **9**, 10.1038/s41598-019-47174-9 (2019).

[21] T. Haner, D. S. Steiger, M. Smelyanskiy, and M. Troyer, High performance emulation of quantum circuits, in *SC16: International Conference for High Performance Computing, Networking, Storage and Analysis* (IEEE, 2016).

[22] J. Niwa, K. Matsumoto, and H. Imai, General-purpose parallel simulator for quantum computing, Physical Review A **66**, 10.1103/physreva.66.062317 (2002).

[23] H. D. Raedt and K. Michielsen, *Handbook of Theoretical and Computational Nanotechnology*, edited by M. Rieth and W. Schommers (American Scientific Publishers, 2006) pp. 2–48.

[24] G. Xiao, C. Yin, T. Zhou, X. Li, Y. Chen, and K. Li, A survey of accelerating parallel sparse linear algebra, ACM Computing Surveys 10.1145/3604606 (2023).

[25] S. Boixo, S. V. Isakov, V. N. Smelyanskiy, and H. Neven, Simulation of low-depth quantum circuits as complex undirected graphical models, arXiv:1712.05384 [quant-ph] (2017), unpublished.

[26] E. Bernstein and U. Vazirani, Quantum complexity theory, SIAM Journal on Computing **26**, 1411 (1997).

[27] Z.-Y. Chen, Q. Zhou, C. Xue, X. Yang, G.-C. Guo, and

G.-P. Guo, 64-qubit quantum circuit simulation, Science Bulletin **63**, 964 (2018).
[28] L. Burgholzer, H. Bauer, and R. Wille, Hybrid schrödinger-feynman simulation of quantum circuits with decision diagrams, in *2021 IEEE International Conference on Quantum Computing and Engineering (QCE)* (IEEE, 2021).
[29] Y. Kim, A. Eddins, S. Anand, K. X. Wei, E. van den Berg, S. Rosenblatt, H. Nayfeh, Y. Wu, M. Zaletel, K. Temme, and A. Kandala, Evidence for the utility of quantum computing before fault tolerance, Nature **618**, 500 (2023).
[30] M. Saffman, T. G. Walker, and K. Mølmer, Quantum information with rydberg atoms, Reviews of Modern Physics **82**, 2313 (2010).