

The Quantum Computer Puzzle

Gil Kalai

Communicated by Joel Hass

DILBERT



Quantum computers are hypothetical devices, based on quantum physics, which would enable us to perform certain computations hundreds of orders of magnitude faster than digital computers. This feature is coined “quantum supremacy”, and one aspect or another of such quantum computational supremacy might be seen by experiments in the near future: by implementing quantum error-correction or by systems of noninteracting bosons or by exotic new phases of matter called anyons or by quantum annealing, or in various other ways. We concentrate in this paper on the model of a universal quantum computer that allows the full computational potential for quantum systems, and on the restricted model, called “BosonSampling”, based on noninteracting bosons.

A main reason for concern regarding the feasibility of quantum computers is that quantum systems are inherently noisy. We will describe an optimistic hypothesis regarding quantum noise that will allow quantum computing and a pessimistic hypothesis that won't. The

quantum computer puzzle is to decide between these two hypotheses. We list some remarkable consequences of the optimistic hypothesis, giving strong reasons for the intensive efforts to build quantum computers, as well as good reasons for suspecting that this might not be possible. For systems of noninteracting bosons, we explain how quantum supremacy achieved without noise is replaced, in the presence of noise, by a very low yet fascinating computational power.¹ Finally, we describe eight predictions about quantum physics and computation from the pessimistic hypothesis.²

Are quantum computers feasible? Is quantum supremacy possible? My expectation is that the pessimistic hypothesis will prevail, leading to a negative answer. Rather than regarding this possibility as an unfortunate failure that impedes the progress of humanity, I believe that the failure of quantum supremacy itself leads to important consequences for quantum physics, the theory of computing, and mathematics. Some of these will be explored here.

Gil Kalai is Henry and Manya Noskwith Professor of Mathematics at the Hebrew University of Jerusalem and is also affiliated with Yale University. His email address is gil.kalai@gmail.com.

This work was supported in part by ERC advanced grant 320924, BSF grant 2006066, and NSF grant DMS-1300120. The author is thankful to an anonymous referee, Bill Casselman, Irit Dinur, Oded Goldreich, Joel Hass, and Abby Thompson for helpful comments, and to Neta Kalai for drawing Figures 2 and 4.

¹Based on G. Kalai and G. Kindler, “Gaussian noise sensitivity and BosonSampling”, arXiv:1409.3093.

²Based on G. Kalai, “How quantum computers fail: quantum codes, correlations in physical systems, and noise accumulation”, arXiv:1106.0485, and a subsequent Internet debate with Aram Harrow and others.

For permission to reprint this article, please contact: reprint-permission@ams.org.

DOI: <http://dx.doi.org/10.1090/noti1380>

A Brief Summary

Here is a brief summary of the author's pessimistic point of view as explained in the paper: understanding quantum computers in the presence of noise requires consideration of behavior at different scales. In the small scale, standard models of noise from the mid-90s are suitable, and quantum evolutions and states described by them manifest a very low-level computational power. This small-scale behavior has far-reaching consequences for the behavior of noisy quantum systems at larger scales. On the one hand, it does not allow reaching the starting points for quantum fault tolerance and quantum supremacy, making them both impossible at all scales. On the other hand, it leads to novel implicit ways for modeling noise at larger scales and to various predictions on the behavior of noisy quantum systems.

The Vision of Quantum Computers and Quantum Supremacy

Circuits and Quantum Circuits

The basic memory component in classical computing is a bit, which can be in two states, “0” or “1”. A computer (or circuit) has n bits, and it can perform certain logical operations on them. The NOT gate, acting on a single bit, and the AND gate, acting on two bits, suffice for *universal* classical computing. This means that a computation based on another collection of logical gates, each acting on a bounded number of bits, can be replaced by a computation based only on NOT and AND. Classical circuits equipped with random bits lead to *randomized algorithms*, which are both practically useful and theoretically important.

Quantum computers (or circuits) allow the creation of probability distributions that are well beyond the reach of classical computers with access to random bits. A qubit is a piece of quantum memory. The state of a qubit can be described by a unit vector in a two-dimensional complex Hilbert space H . For example, a basis for H can correspond to two energy levels of the hydrogen atom or to horizontal and vertical polarizations of a photon. Quantum mechanics allows the qubit to be in a *superposition* of the basis vectors, described by an arbitrary unit vector in H . The memory of a quantum computer consists of n qubits. Let H_k be the two-dimensional Hilbert space associated with the k th qubit. The state of the entire memory of n qubits is described by a unit vector in the tensor product $H_1 \otimes H_2 \otimes \cdots \otimes H_n$. We can put one or two qubits through *gates* representing unitary transformations acting on the corresponding two- or four-dimensional Hilbert spaces, and as for classical computers, there is a small list of gates sufficient for universal quantum computing. Each step in the computation process consists of applying a unitary transformation on the large 2^n -dimensional Hilbert space, namely, applying a gate on one or two qubits, tensored with the identity transformation on all other qubits. At the end of the computation process, the state of the entire computer can be *measured*, giving a probability distribution on 0-1 vectors of length n .

A few words on the connection between the mathematical model of quantum circuits and quantum physics: In quantum physics, states and their evolutions (the way they change in time) are governed by the Schrödinger equation. A solution of the Schrödinger equation can be described as a unitary process on a Hilbert space, and quantum computing processes as we just described form a large class of such quantum evolutions.

A Very Brief Tour of Computational Complexity

Computational complexity is the theory of *efficient computations*, where “efficient” is an asymptotic notion referring to situations where the number of computation steps (“time”) is at most a polynomial in the number of input bits. The complexity class \mathbf{P} is the class of algorithms that can be performed using a polynomial number of steps in the size of the input. The complexity class \mathbf{NP} refers to nondeterministic polynomial time. Roughly speaking, it




Figure 1. The (conjectured) view of some main computational complexity classes. The red ellipse represents efficient quantum algorithms.

refers to questions where we can *provably* perform the task in a polynomial number of operations in the input size, provided we are given a certain polynomial-size “hint” of the solution. An algorithmic task A is \mathbf{NP} -hard if a subroutine for solving A allows solving any problem in \mathbf{NP} in a polynomial number of steps. An \mathbf{NP} -complete problem is an \mathbf{NP} -hard problem in \mathbf{NP} . A useful analog is to think about the gap between \mathbf{NP} and \mathbf{P} as similar to the gap between finding a proof of a theorem and verifying that a given proof of the theorem is correct. \mathbf{P} and \mathbf{NP} are two of the lowest computational complexity classes in the *polynomial hierarchy* \mathbf{PH} , which is a countable sequence of such classes, and there is a rich theory of complexity classes beyond \mathbf{PH} .

There are intermediate problems between \mathbf{P} and \mathbf{NP} . Factoring an n -digit integer is not known to be in \mathbf{P} , as the best algorithms are exponential in the cube root of the number of digits. Factoring is in \mathbf{NP} , but it is unlikely that factoring is \mathbf{NP} -complete. Shor’s famous algorithm shows that quantum computers can factor n -digit integers efficiently—in $\sim n^2$ steps! Quantum computers are not known to be able to solve efficiently \mathbf{NP} -complete problems, and there are good reasons to think that they cannot. Yet, quantum computers can efficiently perform certain computational tasks beyond \mathbf{NP} .

Two comments: First, our understanding of the computational complexity world depends on a whole array of conjectures: $\mathbf{NP} \neq \mathbf{P}$ is the most famous one, and a stronger conjecture asserts that \mathbf{PH} does not collapse, namely, that there is a strict inclusion between the computational complexity classes defining the polynomial hierarchy. Second, computational complexity insights, while asymptotic, strongly apply to finite and small algorithmic tasks. Paul Erdős famously claimed that finding the value of the Ramsey function $R(n, n)$ for $n = 6$ is well beyond mankind’s ability. This statement is supported by computational complexity insights that consider the difficulty of computations as $n \rightarrow \infty$, while not directly implied by them.

Noise

Noise and Fault-Tolerant Computation

The main concern regarding the feasibility of quantum computers has always been that quantum systems are inherently noisy: we cannot accurately control them, and we cannot accurately describe them. To overcome this difficulty, a theory of quantum fault-tolerant computation based on quantum error-correction codes was developed.³ Fault-tolerant computation refers to computation in the presence of errors. The basic idea is to represent (or “encode”) a single piece of information (a bit in the classical case or a qubit in the quantum case) by a large number of physical components so as to ensure that the computation is robust even if some of these physical components are faulty.

The main concern regarding the feasibility of quantum computers has always been that quantum systems are inherently noisy.

freedom, namely, approximating the process on a large Hilbert space by a process on a small Hilbert space. For controlled quantum systems and, in particular, quantum computers, H represents the controlled part of the system, and the large unitary process on H' represents, in addition to an “intended” controlled evolution on H , also the uncontrolled effects of the *environment*. The study of noise is relevant not only to controlled quantum systems but also to many other aspects of quantum physics.

A second, mathematically equivalent way to view noisy states and noisy evolutions is to stay with the original Hilbert space H but to consider a mathematically larger class of states and operations. In this view, the state of a noisy qubit is described as a classical probability distribution on unit vectors of the associated Hilbert spaces. Such states are referred to as *mixed states*. It is convenient to think about the following form of noise, called *depolarizing noise*: in every computer cycle a qubit is not affected with probability $1 - p$, and, with probability p , it turns into the *maximal entropy mixed state*, i.e., the average of all unit vectors in the associated Hilbert space. In this example, p is the error rate, and, more generally, the error rate can be defined as the probability that a

What is noise? Solutions of the Schrödinger equation (quantum evolutions) can be regarded as unitary processes on Hilbert spaces. Mathematically speaking, the study of noisy quantum systems is the study of *pairs* of Hilbert spaces (H, H') , $H \subset H'$, and a unitary process on the larger Hilbert space H' . Noise refers to the general effect of neglecting degrees of

qubit is corrupted at a computation step conditioned on it surviving up to this step.

Two Alternatives for Noisy Quantum Systems

The quantum computer puzzle is, in a nutshell, deciding between two hypotheses regarding properties of noisy quantum circuits: the *optimistic hypothesis* and the *pessimistic hypothesis*.

Optimistic Hypothesis: It is possible to realize universal quantum circuits with a small bounded error level regardless of the number of qubits. The effort required to obtain a bounded error level for universal quantum circuits increases moderately with the number of qubits. Therefore, large-scale fault-tolerant quantum computers are possible.

Pessimistic Hypothesis: The error rate in every realization of a universal quantum circuit scales up (at least) linearly with the number of qubits. The effort required to obtain a bounded error level for any implementation of universal quantum circuits increases (at least) exponentially with the number of qubits. Thus, quantum computers are not possible.

Some explanations: For the optimistic hypothesis, we note that the main theorem of quantum fault tolerance asserts that (under some natural conditions on the noise) if we can realize universal quantum circuits with a sufficiently small error rate (where the threshold is roughly between 0.001 and 0.01), then quantum fault tolerance and hence universal quantum computing are possible. For the pessimistic hypothesis, when we say that the rate of noise per qubit scales up linearly with the number of qubits, we mean that when we double the number of qubits in the circuit, the probability for a single qubit to be corrupted in a small time interval doubles. The pessimistic hypothesis does not require new modeling for the noise for universal quantum circuits, and it is just based on a different assumption on the rate of noise. However, it leads to interesting predictions and modeling and may lead to useful computational tools, for more general noisy quantum systems. We emphasize that both hypotheses are assertions about physics (or physical reality), not about mathematics, and both of the hypotheses represent scenarios that are compatible with quantum mechanics.

The constants are important, and the pessimistic view regarding quantum supremacy holds that every realization of universal quantum circuits will fail for a handful of qubits long before any quantum supremacy effect is witnessed and long before quantum fault tolerance is possible. The failure to reach universal quantum circuits for a small number of qubits and to manifest quantum supremacy for small quantum systems is crucial for the pessimistic hypothesis, and Erdős’s statement about $R(6, 6)$ is a good analogy for this expected behavior.

Both on the technical and conceptual levels we see here what we call a “wide-gap dichotomy”. On the technical level, we have a gap between small constant error rate per qubit for the optimistic view and linear increase of rate

³M. A. Nielsen and I. L. Chuang, *Quantum Computation and Quantum Information*, Cambridge University Press, Cambridge, 2000, Ch. 10.




Figure 2. The optimistic hypothesis: Classical fault-tolerance mechanisms can be extended, via quantum error-correction, allowing robust quantum information and computationally superior quantum computation. Drawing by Neta Kalai.

per qubit (in terms of the number of qubits in the circuit) on the pessimistic side. We also have a gap between the ability to achieve large-scale quantum computers on the optimistic side and the failure of universal quantum circuits already for a handful of qubits on the pessimistic side. On the conceptual level, the optimistic hypothesis asserts that quantum mechanics allows superior computational powers, while the pessimistic hypothesis asserts that quantum systems without specific mechanisms for robust classical information that leads only to classical computing are actually computationally inferior. We will come back to both aspects of this wide-gap dichotomy.

Potential Experimental Support for Quantum Supremacy

A definite demonstration of quantum supremacy of controlled quantum systems—namely, building quantum systems that outperform, even for specific computational tasks, classical computers—or a definite demonstration of quantum error correction will falsify the pessimistic hypothesis and will give strong support for the optimistic hypothesis. (The optimistic hypothesis will be completely verified with full-fledged universal quantum computers.) There are several ways people plan, in the next few years, to demonstrate quantum supremacy or the feasibility of quantum fault tolerance.

- (1) Attempts to create small universal quantum circuits with up to “a few tens of qubits.”
- (2) Attempts to create stable logical qubits based on surface codes.
- (3) Attempts to have BosonSampling for 10–50 bosons.
- (4) Attempts to create stable qubits based on anyonic states.
- (5) Attempts to demonstrate quantum speed up based on quantum annealing.

Each of attempts (1)–(4) represents many different experimental directions carried out mainly in academic institutions, while (5) represents an attempt by a commercial company, D-wave.⁴ There are many different avenues for realizing qubits, of which ion-trapped qubits and superconducting qubits are perhaps the leading ones. Quantum supremacy via nonabelian anyons stands out as a very different direction based on exotic new phases of matter and very deep mathematical and physical issues. BosonSampling (see the next section) stands out in the quest to demonstrate quantum supremacy for narrow physical systems without offering further practical fruits.

The pessimistic hypothesis predicts a decisive failure for *all* of these attempts to demonstrate quantum supremacy or very stable logical qubits and that this failure will be witnessed for small systems. A reader may ask how the optimistic hypothesis can be falsified beyond repeated failures to demonstrate universal quantum computers or partial steps toward them as those listed above. My view is that the optimistic hypothesis can be largely falsified if we can understand the absence of quantum supremacy and quantum error correction as a physical principle with predictive power that goes beyond these repeated failures, both in providing more detailed predictions about these failures themselves (such as scaling-up of errors, correlations between errors, etc.) and in providing predictions for other natural quantum systems. Mathematical modeling of noisy quantum systems based on the pessimistic hypothesis is valuable, not only if it represents a general physical principle, but also if it represents temporary technological difficulties or if it applies to limited classes of quantum systems.

BosonSampling

Quantum computers allow the creation of probability distributions that are beyond the reach of classical computers with access to random bits. This is manifested by BosonSampling, a class of probability distributions representing a collection of noninteracting bosons that quantum computers can efficiently create. It is a restricted subset of distributions compared to the class of distributions that a universal quantum computer can produce, and it is not known if BosonSampling distributions can be used for efficient integer factoring or for other “useful” algorithms. BosonSampling was introduced by Troyansky and Tishby in 1996 and was intensively studied by Aaronson and Arkhipov,⁵ who offered it as a quick path for experimentally showing that quantum supremacy is a real phenomenon.

Given an n by n matrix A , let $\det(A)$ denote the determinant of A , and let $\text{per}(A)$ denote the permanent of A . Thus $\det(A) = \sum_{\pi \in S_n} \text{sgn}(\pi) \prod_{i=1}^n a_{i\pi(i)}$, and $\text{per}(A) = \sum_{\pi \in S_n} \prod_{i=1}^n a_{i\pi(i)}$. Let M be a complex $n \times m$

⁴D-wave is attempting to demonstrate quantum speedup for NP-hard optimization problems and even to compute Ramsey numbers.

⁵S. Aaronson and A. Arkhipov, “The computational complexity of linear optics”, *Theory of Computing* 4 (2013), 143–252; arXiv:1011.3245.

matrix, $m \geq n$. Consider all $\binom{m}{n}$ subsets S of n columns, and for every subset consider the corresponding $n \times n$ submatrix A . The algorithmic task of sampling subsets S of columns according to $|\det(M')|^2$ is called FermionSampling. Next consider all $\binom{m+n-1}{n}$ submultisets S of n columns (namely, allow columns to repeat), and for every submultiset S consider the corresponding $n \times n$ submatrix A (with column i repeating r_i times). BosonSampling is the algorithmic task of sampling those multisets S according to $|\text{per}(A)|^2 / (r_1! r_2! \cdots r_n!)$. Note that the algorithmic task for BosonSampling and FermionSampling is to sample according to a specified probability distribution. They are not decision problems, where the algorithmic task is to provide a yes/no answer.

Let us demonstrate these notions by an example for $n = 2$ and $m = 3$. The input is a 2×3 matrix:

$$\begin{pmatrix} 1/\sqrt{3} & i/\sqrt{3} & 1/\sqrt{3} \\ 0 & 1/\sqrt{2} & i/\sqrt{2} \end{pmatrix}.$$

The output for FermionSampling is a probability distribution on subsets of two columns, with probabilities given according to absolute values of the square of determinants. Here we have probability $1/6$ for columns $\{1, 2\}$, probability $1/6$ for columns $\{1, 3\}$, and probability $4/6$ for columns $\{2, 3\}$. The output for BosonSampling is a probability distribution according to absolute values of the square of permanents of submultisets of two columns. Here, the probabilities are: $\{1, 1\} \rightarrow 0$, $\{1, 2\} \rightarrow 1/6$, $\{1, 3\} \rightarrow 1/6$, $\{2, 2\} \rightarrow 2/6$, $\{2, 3\} \rightarrow 0$, $\{3, 3\} \rightarrow 2/6$.

FermionSampling describes the state of n noninteracting fermions, where each individual fermion is described as a superposition of m “modes”. BosonSampling describes the state of n noninteracting bosons, where each individual boson is described by m modes. A few words about the physics: Fermions and bosons are the main building blocks of nature. Fermions, such as electrons, quarks, protons, and neutrons, are particles characterized by Fermi–Dirac statistics. Bosons, such as photons, gluons, and the Higgs boson, are particles characterized by Bose–Einstein statistics.

Moving to computational complexity, we note that Gaussian elimination gives an efficient algorithm for computing determinants, but computing permanents is very hard: it represents a computational complexity class called $\#\mathbf{P}$ (in words, “number \mathbf{P} ” or “sharp \mathbf{P} ”) that extends beyond the entire polynomial hierarchy. It is commonly believed that even quantum computers cannot efficiently compute permanents. However, a quantum computer can efficiently create a bosonic (and a fermionic) state based on a matrix M and therefore perform efficiently both BosonSampling and FermionSampling. A classical computer with access to random bits can sample FermionSampling efficiently, but, as proved by Aaronson and Arkhipov, a classical computer with access to random bits cannot perform BosonSampling unless the polynomial hierarchy collapses!

Predictions from the Optimistic Hypothesis

Barriers Crossed. Quantum computers would dramatically change our reality.

- (1) A universal machine for creating quantum states and evolutions will be built.
- (2) Complicated evolutions and states with global interactions, markedly different from anything witnessed so far, will be created.
- (3) It will be possible to experimentally time-reverse every quantum evolution.
- (4) The noise will not respect symmetries of the state.
- (5) There will be fantastic computational complexity consequences.
- (6) Quantum computers will efficiently break most current public-key cryptosystems.

Items (1)–(4) represent a vastly different experimental reality than that of today, and items (5) and (6) represent a vastly different computational reality.

Magnitude of Improvements. It is often claimed that quantum computers can perform certain computations that even a classical computer of the size of the entire universe cannot perform! Indeed it is useful to examine not only things that were previously impossible and that are now made possible by a new technology but also the improvement in terms of orders of magnitude for tasks that could have been achieved by the old technology. Quantum computers represent enormous, unprecedented order-of-magnitude improvement of controlled physical phenomena as well as of algorithms. Nuclear weapons represent an improvement of 6–7 orders of magnitude over conventional ordnance: the first atomic bomb was a million times stronger than the most powerful (single) conventional bomb at the time. The telegraph could deliver a transatlantic message in a few seconds compared to the previous three-month period. This represents an (immense) improvement of 4–5 orders of magnitude. Memory and speed of computers were improved by 10–12 orders of magnitude over several decades. Breakthrough algorithms at the time of their discovery also represented practical improvements of no more than a few orders




Photo courtesy of Yifan Wang.

Aram Harrow and Gil Kalai shake hands at MIT after their internet debate.

of magnitude. Yet implementing BosonSampling with a hundred bosons represents more than a hundred orders of magnitude improvement compared to digital computers, and a similar story can be told about a large-scale quantum computer applying Shor's algorithm.

Computations in Quantum Field Theory. Quantum electrodynamics (QED) computations allow one to describe various physical quantities in terms of a power series

$$\sum c_k \alpha^k,$$

where c_k is the contribution of Feynman's diagrams with k loops and α is the fine structure constant (around $1/137$). Quantum computers will (likely)⁶ allow one to compute these terms and sums for large values of k with hundreds of digits of accuracy, similar to computations of the digits of e and π on today's computers, even in regimes where they have no physical meaning!

My Interpretation. I regard the incredible consequences from the optimistic hypothesis as solid indications that quantum supremacy is "too good to be true" and that the pessimistic hypothesis will prevail. Quantum computers would change reality in unprecedented ways, both qualitatively and quantitatively, and it is easier to believe that we will witness substantial theoretical changes in modeling quantum noise than that we will witness such dramatic changes in reality itself.

BosonSampling Meets Reality

How Does Noisy BosonSampling Behave?

BosonSampling and Noisy BosonSampling (i.e., BosonSampling in the presence of noise) exhibit radically different behavior. BosonSampling is based on n noninteracting, indistinguishable bosons with m modes. For noisy BosonSamplers these bosons will not be perfectly noninteracting (accounting for one form of noise) and will not be perfectly indistinguishable (accounting for another form of noise). The same is true if we replace bosons by fermions everywhere. The state of n bosons with m modes is represented by an algebraic variety of decomposable symmetric tensors of real dimension $2mn$ in a huge relevant Hilbert space of dimension $2m^n$. For the fermion case this manifold is simply the Grassmannian.

We have already discussed the rich theory of computational complexity classes beyond **P**, and there is also a rich theory below **P**. One very low-level complexity class consists of computational tasks that can be carried out by bounded-depth polynomial-size circuits. In this model the number of gates is, as before, at most polynomial in the input side, but an additional severe restriction is that the entire computation is carried out in a bounded number of rounds. Bounded-depth polynomial-size circuits cannot even compute or approximate the parity of n bits, but they can approximate real functions described

⁶This plausible conjecture, which motivated quantum computers to start with, is supported by the recent work of Jordan, Lee, and Preskill and is often taken for granted. A mathematical proof is still beyond reach.




Figure 3. The huge computational gap (left) between BosonSampling (purple) and FermionSampling (green) vanishes in the noisy versions (right).

by bounded-degree polynomials and can sample approximately according to probability distributions described by real polynomials of bounded degree.

Theorem 1 (Kalai and Kindler). *When the noise level is constant, BosonSampling distributions are well approximated by their low-degree Fourier-Hermite expansion. Consequently, noisy BosonSampling can be approximated by bounded-depth polynomial-size circuits.*

It is reasonable to assume that for all proposed implementations of BosonSampling, the noise level is at least a constant, and therefore an experimental realization of BosonSampling represents, asymptotically, bounded-depth computation. The next theorem shows that implementation of BosonSampling will actually require pushing down the noise level below $1/n$.

Theorem 2 (Kalai and Kindler). *When the noise level is $\omega(1/n)$ and $m \gg n^2$, BosonSampling is very sensitive to noise, with a vanishing correlation between the noisy distribution and the ideal distribution.*⁷

Theorems 1 and 2 give evidence against expectations of demonstrating "quantum supremacy" via BosonSampling: experimental BosonSampling represents an extremely low-level computation, and there is no precedence for a "bounded-depth machine" or a "bounded-depth algorithm" that gives a practical advantage, even for small input size, over the full power of classical computers, not to mention some superior powers.

Bounded-Degree Polynomials

The class of probability distributions that can be approximated by low-degree polynomials represents a severe restriction below bounded-depth computation. The description of noisy BosonSampling with low bounded-degree polynomials is likely to extend to small noisy quantum circuits and other similar quantum systems, and this would support the pessimistic hypothesis. This

⁷The condition $m \gg n^2$ can probably be removed by a more detailed analysis.

description is relevant to important general computational aspects of quantum systems in nature, as we now discuss.

Why Is Robust Classical Information Possible? The ability to approximate low-degree polynomials still supports robust *classical* information. The (“Majority”) Boolean function⁸ $f(x_1, x_2, \dots, x_n) = \text{sgn}(x_1 + x_2 + \dots + x_n)$ allows for very robust bits based on a large number of noisy bits and admits excellent low-degree approximations. Quantum error correction is also based on encoding a single qubit as a function $f(q_1, q_2, \dots, q_n)$ of many qubits, and also for quantum codes the quality of the encoded qubit grows with the number of qubits used for the encoding. But for quantum error-correction codes, implementation with bounded-degree polynomial approximations is not available, and I conjecture that no such implementation exists. This would support the insight that quantum mechanics is limiting the information one can extract from a physical system in the absence of mechanisms leading to robust classical information.

Why Can We Learn the Laws of Physics from Experiments? Learning the parameters of a process from examples can be computationally intractable, even if the process belongs to a low-level computational task. (Learning even a function described by a depth-two Boolean circuit of polynomial size does not admit an efficient algorithm.) However, the approximate value of a low-degree polynomial can efficiently be learned from examples. This offers an explanation for our ability to understand natural processes and the parameters defining them.

Predictions from the Pessimistic Hypothesis

Under the pessimistic hypothesis, universal quantum devices are unavailable, and we need to devise a specific device in order to implement a specific quantum evolution. A sufficiently detailed modeling of the device will lead to a familiar detailed Hamiltonian modeling of the quantum process that also takes into account the environment and various forms of noise. Our goal is different: we want to draw from the pessimistic hypothesis predictions on noisy quantum circuits (and, at a later stage, on more general noisy quantum processes) that are common to *all* devices implementing the circuit (process).

The basic premises for studying noisy quantum evolutions when the specific quantum devices are not specified are as follows: First, modeling is implicit; namely, it is given in terms of conditions that the noisy process must satisfy. Second, there are systematic relations between the noise and the entire quantum evolution and also between the target state and the noise.

In this section we assume the pessimistic hypothesis, but we note that the previous section proposes the following picture in support of the pessimistic hypothesis: evolutions and states of quantum devices in the small scale are described by low-degree polynomials. This allows, for a larger scale, the creation of robust classical information and computation but does not provide the necessary

⁸A Boolean function is a function from $\{-1, 1\}^n$ to $\{-1, 1\}$.

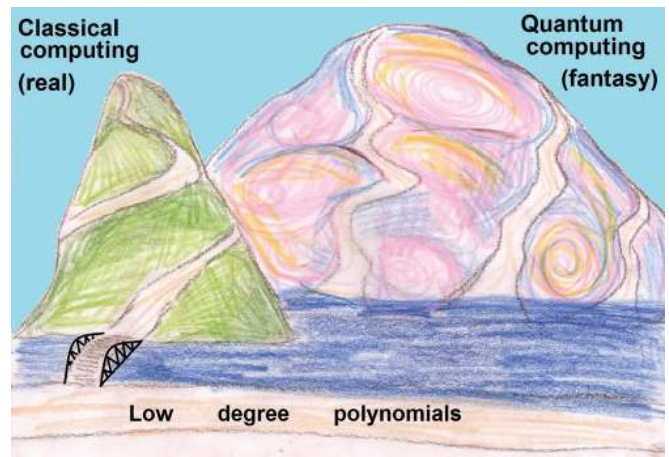


Figure 4. The pessimistic hypothesis: Noisy quantum evolutions, described by low-degree polynomials, allow via the mechanisms of averaging/repetition robust *classical* information and computation but do not allow reaching the starting points for quantum supremacy and quantum fault tolerance. Drawing by Neta Kalai.

starting point for quantum fault tolerance or for any manifestation of quantum supremacy.

No Quantum Fault Tolerance: Its Simplest Manifestation

Entanglement and Cat States. Entanglement is a name for quantum correlation, and it is an important feature of quantum physics and a crucial ingredient of quantum computation. A *cat state* of the form $\frac{1}{\sqrt{2}}|00\rangle + \frac{1}{\sqrt{2}}|11\rangle$ represents the simplest form of entanglement between two qubits. Let me elaborate: the Hilbert space H representing the states of a single qubit is two-dimensional. We denote by $|0\rangle$ and $|1\rangle$ the two vectors of a basis for H . A pure state of a qubit is a *superposition* of basis vectors of the form $a|0\rangle + b|1\rangle$, where a, b are complex and $|a|^2 + |b|^2 = 1$. Two qubits are represented by a tensor product $H \otimes H$, and we denote it by $|00\rangle = |0\rangle \otimes |0\rangle$. Now, a superposition of two vectors can be thought of as a quantum analog of a coin toss in classical probability—a superposition of $|00\rangle$ and $|11\rangle$ is a quantum analog of correlated coin tosses: two heads with probability $1/2$, and two tails with probability $1/2$. The name “cat state” refers, of course, to Schrödinger’s cat.

Noisy Cats. The following prediction regarding noisy entangled pairs of qubits (or “noisy cats”) is perhaps the simplest prediction on noisy quantum circuits under the pessimistic hypothesis.

Prediction 1: Two-qubits behavior. Any implementation of quantum circuits is subject to noise, for which errors for a pair of entangled qubits will have substantial positive correlation.

Prediction 1, which we will refer to as the “noisy cat prediction”, gives a very basic difference between the optimistic and pessimistic hypotheses. Under the optimistic hypothesis gated qubits will manifest correlated noise,

but when quantum fault tolerance is in place, such correlations will be diminished for most pairs of qubits. Under the pessimistic hypothesis quantum fault-tolerance is not possible, and without it there is no mechanism to remove correlated noise for entangled qubits. Note that the condition on noise for a pair of entangled qubits is implicit, as it depends on the unknown process and unknown device leading to the entanglement.

Further Simple Manifestations of the Failure of Quantum Fault Tolerance.

Prediction 2: Error synchronization. For complicated (very entangled) target states, highly synchronized errors will occur.

Error synchronization refers to a substantial probability that a large number of qubits, much beyond the average rate of noise, are corrupted. Under the optimistic hypothesis error synchronization is an extremely rare event.

Prediction 3: Error rate. For complicated evolutions, and for evolutions approximating complicated states, the error rate, in terms of qubit-errors, scales up linearly with the number of qubits.

The three predictions 1–3 are related. Under natural assumptions, the noisy cat prediction implies error synchronization for quantum states of the kind involved in quantum error correction and quantum algorithms. Roughly speaking, the noisy cat prediction implies positive correlation between errors for every pair of qubits, and this implies a substantial probability for the event that a large fraction of qubits (well above the average rate of errors) will be corrupted at the same computer cycle. Error synchronization also implies, again under some natural assumptions, that error rate in terms of qubit errors is at least linear in the number of qubits. Thus, the pessimistic hypothesis itself can be justified from the noisy cat prediction, together with natural assumptions on the rate of noise. Moreover, this also explains the wide-gap dichotomy in terms of qubit errors.

The optimistic hypothesis allows creating via quantum error correction very stable “logical” qubits based on stable raw physical qubits.

Prediction 4: No logical qubits. Logical qubits cannot be substantially more stable than the raw qubits used to construct them.

No Quantum Fault-Tolerance: Its Most General Manifestation⁹

We can go to the other extreme and try to examine consequences of the pessimistic hypothesis for the most general quantum evolutions. We start with a prediction related to the discussion in the section “BosonSampling Meets Reality”.

Prediction 5: Bounded-depth and bounded-degree approximations. Quantum states achievable by any implementation of quantum circuits are limited by bounded-depth polynomial-size quantum computation.

⁹This section is more technical and assumes more background on quantum information.

Even stronger: low-entropy quantum states in nature admit approximations by bounded-degree polynomials.

The next items go beyond the quantum circuit model and do not assume that the Hilbert space for our quantum evolution has a tensor product structure.

Prediction 6: Time smoothing. Quantum evolutions are subject to noise, with a substantial correlation with time-smoothed evolutions.

Time-smoothed evolutions form an interesting restricted class of noisy quantum evolutions aimed to model evolutions under the pessimistic hypothesis when fault tolerance is unavailable to suppress noise propagation. The basic example for time-smoothing is the following: Start with an ideal quantum evolution given by a sequence of T unitary operators, where U_t denotes the unitary operator for the t th step, $t = 1, 2, \dots, T$. For $s < t$ we denote $U_{s,t} = \prod_{i=s}^{t-1} U_i$ and let $U_{s,s} = I$ and $U_{t,s} = U_{s,t}^{-1}$. The next step is to add noise in a completely standard way: consider a noise operation E_t for the t th step. We can think about the case where the unitary evolution is a quantum computing process and E_t represents a depolarizing noise with a fixed rate acting independently on the qubits. And finally, replace E_t with a new noise operation E'_t defined as the average

$$(1) \quad E'_t = \frac{1}{T} \cdot \sum_{s=1}^T U_{s,t} E_s U_{s,t}^{-1}.$$

Prediction 7: Rate. For a noisy quantum system a lower bound for the rate of noise in a time interval is a measure of noncommutativity for the projections in the algebra of unitary operators in that interval.

Predictions 6 and 7 are implicit and describe systematic relations between the noise and the evolution. We expect that time-smoothing will suppress high terms for some Fourier-like expansion, thus relating Predictions 5 and 6. We also note that Prediction 7 resembles the picture about the “unsharpness principle” from symplectic geometry and quantization.¹⁰

Locality, Space and Time

The decision between the optimistic and pessimistic hypotheses is, to a large extent, a question about modeling locality in quantum physics. Modeling natural quantum evolutions by quantum computers represents the important physical principle of “locality”: quantum interactions are limited to a few particles. The quantum circuit model enforces local rules on quantum evolutions and still allows the creation of very nonlocal quantum states. This remains true for noisy quantum circuits under the optimistic hypothesis. The pessimistic hypothesis suggests that quantum supremacy is an artifact of incorrect modeling of locality. We expect modeling based on the pessimistic hypothesis, which relates the laws of the “noise” to the laws of the “signal”, to imply a strong form of locality for both.

We can even propose that spacetime itself emerges from the absence of quantum fault tolerance. It is a

¹⁰L. Polterovich, “Symplectic geometry of quantum noise”, *Comm. Math. Phys.* 327 (2014), 481–519; arXiv:1206.3707.

familiar idea that since (noiseless) quantum systems are time reversible, time emerges from quantum noise (decoherence). However, also in the presence of noise, with quantum fault tolerance, every quantum evolution that can experimentally be created can be time-reversed, and, in fact, we can time-permute the sequence of unitary operators describing the evolution in an arbitrary way. It is therefore both quantum noise and the absence of quantum fault tolerance that enable an arrow of time.

Next, we note that with quantum computers one can emulate a quantum evolution on an arbitrary geometry. For example, a complicated quantum evolution representing the dynamics of a four-dimensional lattice model could be emulated on a one-dimensional chain of qubits. This would be vastly different from today's experimental quantum physics, and it is also in tension with insights from physics, where witnessing different geometries supporting the same physics is rare and important. Since a universal quantum computer allows the breaking of the connection between physics and geometry, it is noise and the absence of quantum fault tolerance that distinguish physical processes based on different geometries and enable geometry to emerge from the physics.

Classical Simulations of Quantum Systems

Prediction 8: Classical simulations of quantum processes. Computations in quantum physics can, in principle, be simulated efficiently on a digital computer.

This bold prediction from the pessimistic hypothesis could lead to specific models and computational tools. There are some caveats: heavy computations may be required for quantum processes that are not realistic to start with, for a model in quantum physics representing a physical process that depends on many more parameters than those represented by the input size, for simulating processes that require knowing internal parameters of the process that are not available to us (but are available to nature), and when we simply do not know the correct model or relevant computational tool.



Gil Kalai lecturing at Adam Mickiewicz University in Poland.

Quid est Noster Computationis Mundus?¹¹

Deciding between the optimistic and pessimistic hypotheses reflects a far-reaching difference in the view of our computational world. Is the wealth of computations we witness in reality only the tip of the iceberg of a supreme computational power used by nature and available to us, or is it the case that the wealth of classical computations we witness represents the full computational power that can be extracted from natural quantum physics processes?

I expect that the pessimistic hypothesis will prevail, yielding important outcomes for physics, the theory of computing, and mathematics. Our journey through probability distributions described by low-degree polynomials, implicit modeling for noise, and error synchronization may provide some of the ingredients needed for solving the quantum computer puzzle.

¹¹What is our computational world?