THE INDIA FORUM
A JOURNAL-MAGAZINE ON CONTEMPORARY ISSUES

July 1, 2026

The Quantum Frontier: Its Promise and Challenges

By: Jaikumar Radhakrishnan,Sandip Trivedi

Quantum technologies hold out the promise of reshaping computation, and improving secure communication and precision measurement. But in spite of decades of research, the realisation of this promise with practical applications remains elusive.

1. Introduction

The potential of quantum mechanics for computation has been recognised and discussed since the 1970s. For two decades it remained in the realm of intellectual speculation with no clear consequences in sight. Then, in 1993, everybody was jolted out of their indifference by a singular event: Peter Shor's announcement of the quantum factoring algorithm.

> There is now great excitement among many scientists and engineers, as this article will hopefully bring out, that the coming few decades hold much promise, as diverse researchers put their minds together and push forward exploring this new and wondrous frontier.

Shor's algorithm would efficiently decompose a large number as a product of smaller numbers, provided we could build a computer that ran on the principles of quantum mechanics. Such a computer would wreck most widely used cryptosystems, which are based on the assumption that such a decomposition must take millions of years on regular 'classical' computers. With the realization that commonly used cryptography-the bedrock on which today's online commercial infrastructure rests-was vulnerable to 'quantum attacks', the discussion of quantum computing became more fervent on the blackboards and laboratories of academia, and suddenly moved into the boardrooms of technological and commercial organisations.

Around the world, computer science and physics departments started offering courses; government agencies, software companies and banks set aside substantial sums for quantum technologies, hoping to exploit its power, or at least stanch the damage it was likely to cause. It has now been 30 years, and a practical quantum computer capable of factoring sufficiently large numbers is still distant. It appears clear that practical quantum computing at a reasonable scale still requires several more years of technological progress and development. Yet, the intense attention this has focused on our understanding of the world around us, as well as our abilities to probe and manipulate it, has undeniably opened up a new intellectual and technological frontier, the *Quantum Frontier*.

Progress pushing forward in this direction is promising to change our understanding of many areas of science, including physics, material science, chemistry and computer science, and have significant impact on important new technologies, including the building of computers which are much more powerful than hitherto imagined, new methods of communication well protected from eavesdropping, and exquisitely more precise methods of sensing and measurement.

> The reason why it has taken so long, and progress, in many respects, has been slow so far, is that the quantum world is truly bizarre; far removed from our day-to-day experiences and the intuition we develop and carry with us based on these experiences.

This new frontier has to do with our increasingly precise understanding of Quantum Mechanics, and the growing realisation that the strange properties of the quantum realm can in fact be a resource for humankind to exploit for the building of powerful new technologies. It might seem strange at first that such a frontier exists in science at all. After all, almost a century has passed since the advent of quantum mechanics, roughly at the turn of the 20th century. Shouldn't the lessons of quantum mechanics have been well absorbed by now and its potential for technology be well developed already? The reason why it has taken so long, and progress, in many respects, has been slow so far, is that the quantum world is truly bizarre; far removed from our day-to-day experiences and the intuition we develop and carry with us based on these experiences. As a result, it has taken a century or so to fully understand its meaning, significance and potential. Of late though different strands of thought, from physics and computer science, for example, have come together, and our pace of understanding is growing and accelerating. There is now great excitement among many scientists and engineers, as this article will hopefully bring out, that the coming few decades hold much promise, as diverse researchers put their minds together and push forward exploring this new and wondrous frontier.

1.1 What defines the quantum world?

Two aspects are key to the mysteries and strangeness of the quantum world: *superposition* and *entanglement*. While these are technical terms, they draw their meaning from familiar experiences in the world around us. Water waves, for example, coming from different directions in the sea, merge, getting *superimposed*, to result in a net disturbance at a given location. The total power in the disturbance of the sea, at that location, is given by taking the sum of the superimposed amplitudes of the different waves to obtain a resultant amplitude and then squaring it (energy is always a square of the displacement or wave amplitude).

In quantum mechanics something similar happens, except it is much more bizarre. The underlying description of a quantum system is intrinsically probabilistic-this by itself is a radical departure, scientifically and philosophically, from a deterministic classical world view. Unlike a clock or a planet whose position you can determine into the future, one cannot, typically, predict with certainty the outcome of quantum events. But it is even stranger than that! The fundamental notion underlying the quantum description is that of a probability amplitude, or wave function, which places the object being studied in a superposition of states. Just like with a water wave, you add the different parts of a quantum phenomenon together to get a total amplitude; squaring that amplitude gives you the probability of the outcome. Even the most precise description of a quantum state can only tell us the probability with which an event might happen.

There is no analogue to this bizarreness in the classical world. We are familiar with probabilities in day-to-day life, but they arise because we do not keep track of all the information, not because key parts of the information are missing, or not-yet-there, in the system. Even more importantly, there is no analogue of the probability amplitudes, which allows in effect for probabilities to cancel each other. Several experiments provide compelling demonstrations of this cancellation, e.g., the double-slit experiment and the Mach-Zehnder interferometer.

> These two distinct properties of superposition and entanglement govern much of the strangeness in the quantum world. Attempts to understand them more deeply and exploit them in new and powerful ways, lie at the heart of the recent progress and excitement in this new frontier.

The concept of entanglement, used to model quantum correlations between different parts of a joint system, is quite subtle and it too has no classical analogue. For example, in the classical world a particle is described by its position and velocity. To describe two particles, we need to simply keep track of each particle's position and velocity. There can be interactions between the two particles which will affect their positions and velocities. But the data or information content for the two particles is obtained by simply keeping track of each particle's information. This is very different from the quantum world, where the relevant data for two quantum particles is more intertwined, and cannot be obtained from each particle's information alone. These two distinct properties of superposition and entanglement govern much of the strangeness in the quantum world. Attempts to understand them more deeply and exploit them in new and powerful ways, lie at the heart of the recent progress and excitement in this new frontier.

Despite their profound differences, classical and quantum system can be placed, compared and understood in a common framework. A physical system can be thought of as possessing basic states: there are two states in which we can find a coin, six states for a die, a door can be open or closed, a car in a city could be moving east, west, north or south. As the system evolves over time, its state changes. For our purposes, it will be enough to consider systems that change their states at discrete points in time. The evolution can be deterministic or probabilistic. If the evolution is probabilistic, the resulting state of the system is not determined to be one of the basic states; the best we can do is assign probabilities to the basic states, which measures the tendency of the system for acquiring those states.

2. Classical Computing versus Quantum Computing

A computer is a physical device that evolves over time based on the laws of physics. It is a physical apparatus; every run of the computer is an experiment. The initial configuration of the apparatus is set up so that it reflects the input we wish to process: for example, in traditional electronic computing the voltages are set to low and high, to reflect the values 0 and 1. The device itself is so rigged up that its evolution inevitably takes the initial state to a final state, from which the desired computational result can be read out. The laws can be deterministic: the state the device will be in the next moment is determined by its state at the current moment. Or they can be stochastic: for symmetry we cannot predict with certainty what the next state would be, for we only have probabilities that describe its tendencies. A quantum computer is similar - the laws that govern the device's evolution involve quantum mechanics, and the outcomes are generally not deterministic.

Classical algorithms have been studied since antiquity. In elementary school, we are taught methods for addition, subtraction, multiplication and division of numbers. Even before computers became available, sophisticated methods for solving a system of equations and practical methods for determining the optimum solutions for problems arising in economics, operations research and manufacturing had been devised. These methods involved deterministic well-defined steps, and were readily implemented once classical computers became available.

> One now devised algorithms that relied on random actions, no longer deterministic. Instead of demanding certainty in the results, we were willing to settle for the result being correct with high probability.

While the foundations of deterministic computing were being developed based on ideas from formal logic, a revolution was sweeping through our understanding of the physical world. The rules of probability, it was realized, explained a range of physical phenomena: the behaviour of gases, the formation of galaxies, the unique patterns on the furs of animals. The understanding of the underlying random processes was crucial to obtaining an explanation of these phenomena, not the precise determination of the initial conditions or the details of the mechanics underlying the subsequent evolution. Inevitably, probability theory influenced computation too, this time, as a computational paradigm, to arrive at solutions to problems, not just as a mathematical framework to explain natural phenomena. One now devised algorithms that relied on random actions, no longer deterministic. Instead of demanding certainty in the results, we were willing to settle for the result being correct with high probability.

In our discussion above, we imagined classical deterministic computation as consisting of a series of operations that nudged the system to a desired output state. Classical computation can also be randomized. Classical randomized computation too involves nudges, except that outcomes of these nudges have probabilities associated with them. For a particular input, some of these outcomes may lead to the desired state, some others may not. The designer of a good randomized algorithm places these nudges in such way that for every input a great majority of the outcomes lead to the desired states; for different inputs, the outcomes that are favourable might be different, but the design of the algorithm ensures that the favourable outcomes always outnumber the unfavourable ones. The randomized algorithm by its design ensures that most of the probability accumulates on the desired final configuration, and only an insignificant sprinkling reaches the undesired states. Classical randomized computation shares some of these features with quantum computation, in particular, the probabilistic nature of the evolution of states, but we will see that quantum computing is more subtle.

When one commonly encounters randomness in everyday life, it appears to be more of a nuisance: the unpredictability in weather or traffic, the variation in the quality of industrially manufactured products, the noise in communication, etc. It is not immediately obvious how randomness can help in efficient computation. Here are two appealing examples that illustrate how randomness can be leveraged to provide efficient solutions.

Example 1: Fingerprinting

Suppose two parties A and B, each have a copy of a long movie on their respective computers, each containing say exactly $N = 10^{12}$ bits. They would like to know if their files are identical (not a bit differs). Any deterministic method for ascertaining this would require a total communication of N bits: one or the other party must send the relevant part of the movie over for the other to verify and report.

There is a remarkable randomized method, however, where the communication can be significantly reduced. Party A picks a short random string and computes a short fingerprint of her copy of the movie and sends the random string and the fingerprint to party B, who then applies the same procedure to his copy of the movie and computes the fingerprint. Party B then tells party A if the fingerprints matched. The fingerprinting procedure is so designed that even if the movies differed in one bit, the fingerprints are likely to differ with probability at least 0.9999. We do not get certainty, but in return get a huge savings in communication.

This remarkable idea of fingerprinting has other applications besides saving on communication. If a government agency confiscates your devices, they are supposed to allow you to compute the fingerprint of the contents, which they must certify as correctly computed. If at a later time, they claim that they found something incriminating on your computer; you could demand the contents of the computer and check if the fingerprint matches the one you kept with yourself. The procedure for computing fingerprints is usually so robust that there is negligible chance that the agency can come up with contents for which the fingerprint just happens to match yours and yet the contents are different and incriminating.

Example 2: Checking Flow Charts

Imagine two medical diagnosis flowcharts, System A and System B, that need to be deployed at a hospital. Each flowchart asks the same 100 questions, but not in the same sequence. Will the two flowcharts give the same diagnosis for every patient they are administered to? In order to be 100% confident, one might test the two flowcharts on every patient and compare the results. For a possible 2^{100} patient profiles, there would be roughly 1.26 nonillion possibilities, clearly too computationally intensive for it to be attempted.

A randomised efficient fingerprinting procedure can be devised to verify if two flowcharts give the same conclusion for all patient profiles. Even if the two flowcharts differ on even one profile, the fingerprints will not match with high probability. Once again, we need to accept a small probability of error, but in return the cost of computing and matching the fingerprints is minuscule when compared to the exhaustive method of comparing the flowcharts for each potential patient profile.

These two examples a illustrate the power of randomness for computational tasks; the power emanates from the algorithms ability to make choices that cannot be predicted in advance. This makes designing hard inputs for them difficult; in many cases, one can argue that no matter what input is provided most choices the algorithm makes will lead to an acceptable solution. A quantum algorithm enjoys similar advantages-its choices are also unpredictable, but in an even more subtle way.

> It, however, appears unlikely that quantum technologies will revolutionize computing in the way the advent of semiconductors did. We do not know how to use quantum computers to speedily solve several important hard problems.

Instead of deterministic bits, a quantum algorithm works with qubits. A qubit can be put in a state where we cannot predict with certainty if its value is 0 or 1, until we observe it. Again, all we have is a measure of its tendency of turning up 0 or 1; these tendencies are described by amplitudes. Just as the operations in a randomized algorithm nudge the states so that the probability accumulates on desirable states, the operations of the quantum algorithm nudge amplitudes so that they accumulate on desirable states. The quantum algorithm, however, has an additional trick up its sleeve: *the power of destructive interference*. In addition to accumulating tendencies towards desirable states, it is also allowed to let a system evolve so that it accumulates a tendency to reach undesirable states, while simultaneously sending an antidote to the system to reduce this tendency, so that eventually the system does not reach such unwanted states. How a combination of accumulation and destruction can lead to more efficient algorithms is an involved analysis, and is still being understood.

2.1 Can quantum algorithms do better...

We do not at present know an efficient method (deterministic or randomized) for simulating a general quantum system using a classical computer. In the context of quantum circuits this means that even if we are given the circuit diagram with registers, gates and wires, we do not know how to efficiently (in a number of steps that grows only moderately with the size of the circuit) estimate the probability that we will find a register in state 1 at the end. This presents us with an opportunity! After all, if a physics experiment

does things that we cannot easily simulate, then why not rig up the experiment in such a way that it performs computations for us that we cannot at present perform on classical computers?

> …[I]t is too early to dismiss the potential of quantum computers based on the evidence of the past two decades: once viable quantum devices become available, the attention and effort will be focused towards effectively exploiting them, and unforeseen important practical applications might well emerge.

Do we have examples where this does result in faster algorithms? Yes, in principle, but only a few. The most significant and celebrated is the problem for factoring numbers. The computational task is the following: we are given an n-digit number N , and asked to tell if N is prime, and if it is not prime we must present a factorization of N as a product of two numbers N1 and N2, both at least two. The number of steps needed by the best classical algorithm we know today, approximately scales exponentially with $n^{1/3}$. It is practically impossible, e.g., to factor numbers that are say 600 digits long.[1] In 1994, Peter Shor showed that if quantum devices can be constructed with sufficient precision and can be operated so that they hold on to their state for sufficiently long, then we can factor n-digit numbers in time that scales only as n^3. This is absolutely stunning.[2]

The factoring problem, while not lacking in appeal for everybody who knows numbers, might appear specialized, and far removed from problems we hope computers will solve for us: such as predicting the weather or the stock market, or problems in manufacturing, say determining the optimal way to cut a sheet of metal into pieces of prescribed sizes. However, much rides on this problem: it is our assumption that factoring is hard for classical computers that underlies our confidence in the cryptographic protocols employed in our communication today. A quantum computer capable of factoring a thousand-digit number would leave much of today's communication vulnerable.

It, however, appears unlikely that quantum technologies will revolutionize computing in the way the advent of semiconductors did. We do not know how to use quantum computers to speedily solve several important hard problems: it is widely believed that the notoriously hard NP-complete problems (e.g., *the travelling salesperson problem*) are not amenable to quantum computing methods.[3] Yet, it is too early to dismiss the potential of quantum computers based on the evidence of the past three decades: once viable quantum devices become available, the attention and effort will be focused towards effectively exploiting them, and unforeseen important practical applications might well emerge.

2.2 Reliable cryptography

As we mentioned above, Shor's algorithm and ideas that emanate from would make the most commonly used methods of encryption insecure if quantum computers become available. Interestingly, however, there are cyrptographic methods whose security is based on the inherently quantum aspects of nature; these methods provide, perhaps, the most compelling and practically credible applications of quantum computing for the near-term.

A standard method for one party to securely transmit information to another when a third party tries to eavesdrop is to scramble the message. The scrambling and unscrambling methods are known only to the two parties but not the eavesdropper. Here is a textbook method, popularly known as *one-time pad*. Suppose two parties, Alice and Bob, are at a common location; they determine that at a later time they will be separated and will only have an insecure method for communication. However, at that later point Alice would like to send Bob a number in the range 0000 to 9999 (say, advice on which stock to buy, or a one-time-password for a bank transaction). How can Alice send this information to Bob without the eavesdropper learning of it?

> Once quantum devices capable of faithfully transmitting and maintaining quantum information become available, cryptography whose security is guaranteed by the very laws of quantum mechanics is likely to become ubiquitous.

The classical idea goes something like this: when they are together, Alice and Bob pick a random number r uniformly in the range 0000 to 9999, which they keep securely with themselves. Let the four digits of r be r_1, r_2, r_3, r_4 (so if r is 4562, then r_1 is 4, r_2 is 5, r_3 is 6 and r_4 is 2). Eventually, when Alice needs to send Bob the four digit number n, with digits n_1, n_2, n_3, n_4, over the insecure communication medium, she scrambles n using r before sending the number over. This is how the scrambling works. Imagine that n_1, n_2, n_3, n_4 appear on a number lock that we commonly see on a safe or a suitcase. Now, Alice moves the first digit r_1 units up, the second digit r_2 units up, the third digits r_3 units up, and the fourth digits r_4 units up. Alice sends Bob the four digits s_1, s_2, s_3, s_4 that now show up. On receiving s_1, s_2, s_3, s_4, Bob imagines that they are on a number lock, and unscrambles them using r_1, r_2, r_3, r_4, by

moving the first digit down r_1 units, the second r_2 units, and so on. (Figure 1).

Figure 1: The One-Time Pad

- Alice and Bob agreed on a number

  $$r = \boxed{4}\boxed{5}\boxed{6}\boxed{2}.$$

  They ensure that no eavesdropper gets to know r.
- At a later point Alice wants to send Bob the number $n = \boxed{7}\boxed{8}\boxed{6}\boxed{6}$.
- She scrambles the numbers using the digits of r to obtain

  $$s = \boxed{1}\boxed{3}\boxed{2}\boxed{8}.$$

- On receiving s, Bob (who knows r) reverses Alice's actions and recovers n from s. For the eavesdropper, who does not know r, s provides no clue about n.

For this scheme to work, Alice and Bob need to be at the same place initially (or possess some other secure means of communicating the digits of r). Bennett and Brassard in 1984 pointed out that Alice and Bob could use quantum mechanics to generate the random number r without ever being at the same place at some earlier point. Alice and Bob are now assumed to possess a channel that can carry qubits. The idea is for Alice to generate qubits that are entangled and send some of the qubits to Bob over the quantum channel (which is perhaps being monitored by an eavesdropper). What saves the day is the extraordinary quantum phenomenon that the state of a quantum system changes when it is observed. So, if the eavesdropper actually attempts to copy some of the information being sent, the nature of the qubits that Bob eventually receives will have changed. Alice and Bob, with some additional classical communication (this time in full view of the eavesdropper), can eliminate the tainted qubits, and arrive at the random number r of which they can be reasonably sure the eavesdropper has no idea. This proposal of Bennett and Brassard constitutes a quantum key (the number r in our example) distribution protocol. Once quantum devices capable of faithfully transmitting and maintaining quantum information become available, cryptography whose security is guaranteed by the very laws of quantum mechanics is likely to become ubiquitous. In fact, several companies already supply products for cryptography based on quantum devices.

3. Future Prospects and Conclusions

To summarise, despite the great promise of quantum technologies, there are many challenges which still need to be overcome for building workable quantum computers. And the problems that quantum computers can solve, in a much more efficient manner compared to their classical counterparts, still remain just a few.[4] One may be tempted to look for parallels in the spectacular development of classical computers. Are we today where classical computers were in the late 1940s or early 1950s? The comparison is perhaps weak, even untenable.

> How significant the impact of ideas drawn from quantum mechanics will be in reshaping computing in the future… remains to be seen.

The technological world had learnt to control electrical and electromagnetic signals decades before it embarked on making a computer. Millenia old algorithms of various kinds already existed that could potentially run on a computer. A couple of decades after the transistor was invented, the computer moved into people's homes and eventually onto their palms. In comparison, our ability to control quantum signals is still rudimentary. We do not yet know what form a quantum computer will take. When the technological challenges are surmounted, the quantum computer will likely be deployed for rather specialized tasks whose impact on day-to-day life would probably be rather indirect. How significant the impact of ideas drawn from quantum mechanics will be in reshaping computing in the future, therefore, remains to be seen.

There are of course still considerable grounds for optimism. Some of these stem from the fact that entanglement, one of the quintessential properties tied to quantum mechanics that we mentioned at the outset, itself allows for errors that would accumulate during the running of a quantum computers, to be controlled and corrected. In fact, a race is on, among many research and industry

groups now, to see if scalable computing systems can be built using entanglement to tame the errors and make reliable quantum hardware.

The prospects of quantum-based cryptography, which would be inherently more safe seem even brighter, and we can be sure that these will find important practical applications in the foreseeable future. There are a few other exciting directions that also hold considerable promise in this new frontier.

One idea has to do with the building of *quantum simulators*. It is well known that quantum systems, with even a modest degree of complexity, are notoriously difficult to understand using classical computers that we have today. This is because the computational resources needed for their study seem to grow exponentially with the number of degrees of freedom the system possesses. The idea behind a simulator is to instead build a *designer* quantum system, whose behaviour will mimic, or simulate, that of the actual system of interest. The simulator may well involve very different moving parts, for example, it may have atoms, or ions, instead of electron in the actual system. The quantum interactions between the atoms or ions are then designed such that the simulator behaves like the system of interest, thereby allowing us, in effect, to efficiently carry out the computations needed to understand the system of interest.

> …[O]ne can hope to do much better measurements of magnetic fields, which can be used for more building more advanced MRI machines, and help in the study of biological and chemical systems. It is quite certain that developments in quantum technologies will play an important role in the areas of communication and sensing.

Another idea is to use quantum mechanics for building exquisitely more precise measuring instruments. This, it is hoped, would allow us to control and exploit the world around in hitherto unimagined ways. For example, while there has already been remarkable progress in the building of precise clocks, one can hope to do even better in the future using the additional power of quantum mechanics-and this will help with building better navigational tools, GPS systems, etc. Similarly, one can hope to do much better measurements of magnetic fields, which can be used for more building advanced MRI machines, and help in the study of biological and chemical systems. It is quite certain that developments in quantum technologies will play an important role in the areas of communication and sensing.

Finally, as our understanding of the nature of quantum information grows, it should help us understand the physics of quantum systems itself in a deeper manner. This has already begun to happen in a significant way, both in the study of new materials which involve highly entangled states of matter, and the study of black holes in quantum gravity.

A last comment for why the readers might care, from an Indian perspective, is that the Government of India, recognising the potential of this new frontier, has recently announced a National Quantum Mission, with a significant outlay of about Rs 6,000 crore, to spur progress along many of the directions we discussed above. This makes the future prospects for the *quantum frontier* all the more exciting.

Jaikumar Radhakrishnan is a theoretical computer scientist with research interest in algorithms, combinatorics, computation complexity, information theory, and quantum information processing; he works at the International Centre for Theoretical Sciences, Tata Institute of Fundamental Research, Bengaluru.

Sandip Trivedi is a theoretical physicist with research interests in string theory, cosmology, particle physics, and quantum information; he works at the Tata Institute of Fundamental Research, Mumbai.

Footnotes:

1 See https://en.wikipedia.org/wiki/RSA_numbers for an account of factoring challenges posed in 1991, some of which have not yet been successfully met.

2 Even faster methods have been found recently.

3 In the travelling salesperson problem we are given a set of n cities and the costs of travelling between them. A salesperson wishes to visit each of the cities and return to the city they started at. In what order must they visit the city so that the total cost of travel is minimum. We do not know an efficient method to solve this problem.

4 The most noteworthy among them are Grover's quantum algorithm for database search (1996), and the algorithm of Harrow, Hassidim and Lloyd for solution of a system of linear equations (2008).