

TOWARD A WORLD WITH QUANTUM COMPUTERS

*Surveying the recent past and projecting future developments
and applications involving quantum information science.*

When one naively follows the trajectory of Moore's Law to its endpoint one finds computers whose constituents are only a few atoms in size. In a world where computer gates are made of such microscopic entities, the physical laws of quantum theory dominate, and as such, the rules governing information processing change drastically. This realization, that at the end of Moore's Law lie machines dominated by quantum effects, has given birth to an entirely new and highly active subdiscipline of computer science and physics: quantum information science. Quantum information science, and the new quantum technology it seeks to create, are not meant as a solution for how to go beyond Moore's Law, but are better viewed as the realization that an entirely new form of information processing is possible. This new form of information processing, among whose many forms are quantum

MOTIVATED BY THEORETICAL INDICATIONS OF THE SURPRISING POWER OF QUANTUM INFORMATION PROCESSING DEVICES, A WORLDWIDE EFFORT TO CONSTRUCT LARGE-SCALE QUANTUM COMPUTERS IS UNDER WAY.

computing, quantum information theory, and quantum cryptography, has turned out to be loaded with important solutions to computational and cryptographic problems and new fundamental insights into the physical nature of computation. Motivated by theoretical indications of the surprising power of quantum information processing devices, a worldwide effort to construct large-scale quantum computers is currently under way.

The basic idea of quantum information science is deceptively simple—a brief tutorial is provided in the sidebar “Basic Elements of Quantum Information Science”—yet the implications this has are profound. Here, we survey several of the themes dominating modern research in quantum information science (these being representative and in no way meant to be exhaustive; a near-complete record of papers on quantum computing is available at the online archive <http://arxiv.org/archive/quant-ph>).

THE QUEST TO BUILD A QUANTUM COMPUTER

Quantum theory and computer science are both disciplines that have dominated the 20th century. However, merging these two disciplines took an incredibly long period of time. Starting in the mid-1960s, an intrepid band of researchers began to consider the idea of quantum information science (see the sidebar “Timeline for Quantum Informa-

tion Science”). Yet in spite of the dominance of the quantum theory of nature and the rising success of computation, the field of quantum information science remained effectively a small interest group on the fringe of legitimacy. In 1994, that all changed and quantum computing moved into the limelight when Peter Shor discovered that quantum computers could efficiently factor and compute the discrete logarithm [8]. While this sounds like esoterica only a theoretical computer scientist could appreciate, it is actually of huge practical significance. This is because the most widely used public key cryptosystems, RSA and Diffie-Hellman, are guaranteed to be secure only if these two problems, factoring and discrete log, are computationally intractable. Thus a quantum computer, it seemed, would render insecure a central component of modern cryptography, but only if a large-scale quantum computer could indeed be physically constructed.

So can a large-scale quantum computer be built? The answer to this question was the second great triumph of quantum computation after Shor’s algorithm: the development of quantum error correction and fault-tolerant quantum computation. The concern about quantum computers is that they are essentially analog computers. If we take a model of analog computation where our information is stored and manipulated as infinite precision

BASIC ELEMENTS OF QUANTUM INFORMATION SCIENCE

A classical bit is used to describe two different configurations of an information processing machine, say, labeled by 0 and 1. If we only know probabilities about which configuration the system is in, then it is represented by a two-component vector of the probabilities. A gate is then a 2×2 transition matrix that preserves total probability. In the quantum world, something similar happens, with the probability for a configuration replaced by a complex amplitude (the magnitude squared of which represents the probability of observing the corresponding configuration), and a quantum gate is given by a 2×2 unitary matrix. Quantum gates seem to form a continuous set. Luckily, there exist finite universal sets of one- and two-qubit gates that can be used to compose approximate versions of any unitary gate on all n qubits. Furthermore, accuracy is easily addressed by the theory of fault-tolerant quantum computation (see [1]). A quantum algorithm is a classical prescription for: preparing a quantum system in an initial configuration; performing a sequence of gates from a universal gate set; and making a measurement of the configuration afterward. The resulting configuration is then the output of the computation. It is important to realize that a quantum computer is no more an analog computer than a classical digital computer that uses randomness. That the transition from probabilities to amplitudes leads to speedups is not, intrinsically, a result of moving to an analog model with complex numbers, but to a model where real probabilities are replaced by complex amplitudes. Why this simple change leads to speedups in computation is one of the deep mysteries of quantum computing. **C**

numbers, then this model of computation has extraordinary computational power, being able to solve NP-complete problems in polynomial time, for example. This power, however, is (believed to be) destroyed if we move from infinite precision continuous variables for our information to variables that are affected by imprecision and noise. Quantum computers, at first glance, appear to be similar. Indeed, quantum systems are allowed to exist in superpositions of different configurations, the amplitudes of these different configurations being complex infinite precision numbers.

However, looks can be deceiving. Consider, for example, a probabilistic classical computer. Allowing probabilities, which are real numbers, in our computation, might, in analogy with analog computers give us unwarranted computational power. This belief, however, does not bear itself out. A key observation (due to Claude Shannon) is that faithful communication in the presence of noise, the simplest model of probabilistic information transfer, is readily achieved to any desired degree of success by using error-correcting codes—by judicious use of redundancy in our representation of information. So one can ask, is it possible to construct quantum error-correcting codes to protect quantum information from quantum noise? Earlier, William Wootters and Wojciech Zurek showed that quantum information cannot be cloned, creating widespread belief that quantum error correcting is impossible.

The answer, it turns out, is that quantum error correction is possible. In the years from 1995 to 1997, a large number of researchers developed a theory of quantum error-correcting codes. Despite the impossibility of cloning and the continuous representation of quantum information, it was proven that quantum error-correcting codes exist and are effective in protecting quantum data from the effects of quantum noise. The main insight is that it is possible to represent quantum noise in a discrete manner, which can be detected and corrected. But simply being able to protect quantum information in a communication setting is not enough. In addition, if we want to build a quantum computer, we must also show that quantum computing itself can be done when every component of the computer can possibly fail due to noise. For classical computers, this problem was essentially solved by John von Neumann's NAND multiplexing approach. His basic idea was to compute on information encoded into classical error-correction codes, in a manner that would not propagate errors. Could a similar approach be used for quan-

TIMELINE FOR QUANTUM INFORMATION SCIENCE

- 1960s** Stephen Weisner proposes using quantum information for a secure money scheme.
- 1973** Alexander Holevo studies the transmission of classical information over quantum channels.
- Early 1980s** Paul Benioff, Charles Bennett, David Deutsch, Richard Feynman, and Yuri Manin conceive of the idea of a quantum computer. Feynman notes that naive simulations of quantum systems is classically difficult. David Deutsch suggests that quantum computers challenge the strong Church-Turing thesis.
- 1984** Charles Bennett and Gilles Brassard invent a secure quantum key distribution scheme.
- 1992** David Deutsch and Richard Jozsa show the first quantum speedup over classical computers in a black-box setting.
- 1993–1994** Ethan Bernstein and Umesh Vazirani show the first superpolynomial speedup for quantum computers followed closely by Dan Simon showing the first exponential speedup for quantum computers.
- 1994** Peter Shor discovers a quantum algorithm for efficiently factoring and computing the discrete log.
- 1995** Peter Shor and Andrew Steane invent quantum error-correcting codes.
- 1995** David Wineland's group at NIST implements the first two-qubit gate in ion traps using a scheme invented by Ignacio Cirac and Peter Zoller the previous year.
- 1996** Lov Grover shows that quantum computers offer quadratic speedups for unstructured search problems.
- 1996–1997** The four groups of Dorit Aharonov and Michael Ben-Or, Alexei Kitaev, and Emanuel Knill, Raymond Laflamme and Wojciech Zurek and, in unpublished work, Daniel Gottesman and John Preskill, produce the first threshold theorems for fault-tolerant quantum computing.
- 1999** Ran Raz shows that quantum entanglement allows an exponential decrease in the number of bits of communication that is necessary to solve certain distributed problems.
- 2001** Isaac Chuang and his group at IBM use an NMR quantum computer to implement Shor's algorithm and factor the number 15.
- 2002–2003** A Swiss company, id Quantique, and an American company, MagicQ, begin selling commercial quantum key distribution systems.
- 2005** David Wineland's group at NIST builds a five-qubit ion trap quantum computer and Rainer Blatt's group in Innsbruck entangles six qubits in an ion trap.
- 2006** A group at the University of Waterloo and MIT benchmarks a 12-qubit NMR quantum computer.

tum computers? The answer to this is one of the most profound discoveries on the border between physics and computer science—the threshold theorem for fault-tolerant quantum computation. The threshold theorem states that if the rate of noise on a quantum computer is small enough, and the control over the quantum computer is precise enough, then one can simulate a quantum algorithm with a desired failure rate using an overhead that scales only polylogarithmically with the inverse of the failure rate. (For a comprehensive introduction and some latest results in this fascinating area see [1].) This, remarkably, is exactly the scaling achieved for similar statements in classical probabilistic computation. Therefore, the model of quantum computation is much closer to classical computation than the unrealistic analog model of computation. Thus, if we are to accept classical computation as a valid, scalable, model of computation, it seems that we must do the same for quantum computers.

The threshold theorem was truly a triumphant discovery of a new form of computation, but is it true that our universe is so generous as to provide physical systems that satisfy the conditions of this theorem out of which we can build a large-scale quantum computer? The answer to this question is what experimental physicists have been working on in the decade since the discovery of Shor's algorithm. Numerous experimental implementations for quantum computing have been proposed and are currently being implemented, including, but not limited to, trapped ions, trapped neutral atoms, superconducting circuits, photons, and quantum dots. Research in these fields is seeking to create the basic constituents of a quantum computer, the so-called quantum bit (qubit) and to show that these units can be manipulated in a regime below the threshold for quantum computation.

At this point it is still too early to tell which implementation will be the ultimate technology for quantum computation, but current successes include ion-trap quantum computers with up to eight qubits [4] and NMR quantum computers with up to 12 qubits [7]. Large-scale efforts to scale ion-trap quantum computers into the tens to hundreds of qubits are currently being pursued worldwide. While the task of building a scalable ion-trap quantum computer is daunting, the various pieces of technology needed to do so have all been experimentally demonstrated. The challenge now is to get these pieces to work together and in a manner whose economics and engineering complexity scales in a reasonable manner. Solid-state implementations of quantum computers, unlike their

ion-trap brethren, are just now achieving demonstrations of the basic one- and two-qubit manipulations necessary for quantum computation [9]. However, due to the vast fabrication infrastructure already developed for solid-state implementations, it is entirely conceivable that, having demonstrated the most basic manipulations, they will easily scale to larger numbers of qubits and quickly catch up with ion-trap quantum computers.

Interestingly, the quest to build a quantum computer has now engendered the first generation of quantum computer architects. Quantum computing architecture offers a different set of challenges over classical architecture due to the unique properties of quantum theory. A second renaissance of quantum error correction is currently under way, motivated, in large part by the quest to establish reasonable micro-architectures to achieve fault-tolerant quantum computation within the confines of the currently investigated physical systems.

The quest to build a large-scale quantum computer is still a field in its infancy, comparable in a classical setting to the world before the invention of the first transistor. There appear, however, no theoretical obstacles to building a large-scale quantum computer. Whether this will be a multidecade slog toward bigger computers or a revolution with the invention of, for want of a better name, one could call a quantum transistor, is the fundamental question.

QUANTUM CRYPTOGRAPHY: FROM THEORY TO PRODUCT

While quantum information processing plays a role in code breaking, the same quantum effects also bring in new methods for cryptography. In fact, cryptographic applications of quantum systems predated the idea of quantum computation by two decades. In the 1960s, Stephen Wiesner proposed a simple (theoretical) quantum scheme for money that resists counterfeiting by the laws of physics—when information is properly encoded in quantum states, it is impossible for some other illegitimate party to access that information without disturbing the state, thus providing valuable means to detect various dishonest behavior. The same principle was turned into spectacular use in 1984 by Charles Bennett and Gilles Brassard, in their quantum key distribution scheme [2]. The security of the scheme is based on the possibility to detect eavesdropping, thereby achieving a type of security impossible by classical means. Unconditional security in the presence of channel noise (allowing the most powerful, joint, attack permissible within quantum mechanics) was proved by Dominic

Mayers [6]. These results were improved by many others so that protocols are simpler to implement and provably secure for higher error rates and in the presence of other device imperfections.¹ It is also remarkable that quantum key distribution requires little quantum coherent manipulation—besides the preparation, transmission, and measurement of a small number of simple quantum states, the rest of the processing is classical. It is expected to become fully realized even if a large-scale quantum computer cannot be built.

State-of-the-art experiments have demonstrated unconditionally secure key distribution up to a distance of approximately 100km. Several attacks exploiting implementation imperfections have been patched. Two companies, MagiQ and id Quantique, are marketing the first generalization quantum key distribution devices, paving the way to full commercialization. These are not the only companies currently investing in quantum technologies, however: Hewlett-Packard, Microsoft, IBM, Lucent, Toshiba, and NEC all have active research programs in quantum information technology.

BEYOND KILLER APPLICATIONS

Besides Shor's algorithm and quantum key distribution, what else can we gain from quantum information science? Having more quantum algorithms certainly provides an answer. An entire family of algorithms has been developed along the direction given by Lov Grover's quantum search algorithm [3] for an unsorted database with runtime square root of the size of the database. Meanwhile, Sean Hallgren found an algorithm for solving Pell's equation [5], which is exponentially faster than the best known classical algorithm. Other recent quantum algorithms include those estimating Gauss sums, solving certain hidden shift problems, efficient gradient calculation, and counting points on certain curves.

Quantum information science has also inspired many new ideas in physics and computer science. A partial list of results include efficient simulation of certain quantum systems, better formulation of the blackhole information loss paradox, quantum proofs for exponential lower bounds on classical locally decodable codes, elegant quantum proofs for properties of classical complexity classes, and

security proofs of public key cryptographic systems based on quantum hardness.

CONCLUSION

Quantum information science has evolved from its nascency only a decade ago into a mature research discipline with an extraordinary breadth of concerns and achievements. The quest to build a large-scale quantum computer is among the great technological races of the 21st century, a race whose results may profoundly alter the manner in which we compute. Today, with the advent of commercial quantum cryptography, the first effects of our quantum future are beginning to reverberate. What the future holds we cannot say, but the last decade has taught us that we should not be surprised by the depth and abilities of our future quantum information processing machines. **C**

REFERENCES

1. Aliferis, A., Gottesman, D., and Preskill, J. Quantum accuracy threshold for concatenated distance-3 codes. *Quant. Inf. Comput.* 6, (2006), 97–165; quant-ph/0504218.
2. Bennett, C. and Brassard, G. Quantum cryptography: Public-key distribution and coin-tossing. In *Proceedings of IEEE Conference on Computers, Systems, and Signal Processing* (1984), 175–179.
3. Grover, L. A fast quantum mechanical algorithm for database search. In *Proceedings of the 28th Annual ACM Symposium on the Theory of Computation*. ACM Press, New York, 1996, 212–219; quant-ph/9605043.
4. Haffner, H. et al. Scalable multiparticle entanglement of trapped ions. *Nature* 438 (2005), 643–646; quant-ph/0603217.
5. Hallgren, S. Polynomial-time quantum algorithms for Pell's equation and the principal ideal problem. In *Proceedings of the 34th Annual ACM Symposium on the Theory of Computation*. ACM Press, New York, 2002, 653–658.
6. Mayers, D. Quantum key distribution and string oblivious transfer in noisy channels. In *Advances in Cryptography—Proceedings of Crypto'96*. Springer-Verlag, New York, 1996, 343–357; quant-ph/9606003.
7. Negrevergne, C. et al. Benchmarking quantum control methods on a 12-qubit system. *Phys. Rev. Lett.* 96 (2006); quant-ph/0603248.
8. Shor, P.W. Algorithms for quantum computation: Discrete log and factoring. In S. Goldwasser, Ed., *Proceedings of the 35th Annual Symposium on the Foundations of Computer Science*, IEEE Computer Society, Los Alamitos, CA, 1994, 124–134; quant-ph/9508027.
9. Steffen, M. et al. Measurement of the entanglement of two superconducting qubits via state tomography. *Science* 33 (2006), 1423–1425.

DAVE BACON (dabacon@cs.washington.edu) is an assistant research professor in the Department of Computer Science and Engineering with an adjunct appointment in the Department of Physics at the University of Washington in Seattle, WA.

DEBBIE LEUNG (wcleung@math.uwaterloo.ca) is an assistant professor in the Department of Combinatorics and Optimization and the Institute for Quantum Computing at the University of Waterloo in Ontario, Canada.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

¹A partial list of other major results in quantum key distribution includes Ekert's protocol proposed in 1991, a simple and versatile unconditional security proof for Ekert's protocol by Lo and Chau in 1999, and a prescription by Shor and Preskill in 2000 for converting a Lo-Chau type of security proof to one for the much simpler prepare-measure schemes.