

Can quantum mechanics help distributed computing?

Anne Broadbent

Alain Tapp

Abstract

We present a brief survey of results where quantum information processing is useful to solve distributed computation tasks. We describe problems that are impossible to solve using classical resources but that become feasible with the help of quantum mechanics. We also give examples where the use of quantum information significantly reduces the need for communication. The main focus of the survey is on communication complexity but we also address other distributed tasks.

Keywords: pseudo-telepathy, communication complexity, quantum games, simulation of entanglement

Quantum computation and entanglement

This survey is aimed at researchers having very limited knowledge of quantum computation, but that have a basic understanding of complexity from the theoretical computer science perspective. We address the topics of communication complexity and pseudo-telepathy, as well as other problems of interest in the field of distributed computation. The goal of this survey is not to be exhaustive but rather to cover many different aspects and give the highlights and intuition into the power of distributed quantum computation. Other relevant surveys are available [53, 15, 21, 17].

In classical computation, the basic unit of information is the bit. In quantum computation, which is based on quantum mechanics, the basic unit of information is the *qubit*. A string of bits can be described by a string of zeroes and ones; quantum information can also be in a classical state represented by a binary string, but in general it can be in *superposition* of all possible strings with different *amplitudes*. Amplitudes are complex numbers and thus the complete description of a string of n qubits requires 2^n complex numbers. The fact that quantum information uses a continuous notation does not mean that qubits are somewhat equivalent to analog information: although the description of a quantum state is continuous, quantum measurement, the method of extracting classical information from a quantum state, is discrete. Only n bits of information can be extracted from an n -qubit state. Depending of the choice of measurement, different properties of the state can be extracted but the rest is lost for ever. Another way to see this is that measurement disturbs a quantum state irreversibly. In quantum algorithms, it is possible to compute a

function on all inputs at the same time by only one use of a quantum circuit. The difficult part is to perform the appropriate measurement to extract useful information about the function. We refer the reader to [48, 45, 36] for introductory textbooks to quantum information processing.

One of the most mysterious manifestations of quantum information is *entanglement*, according to which distant parties can share correlations that go beyond what is feasible with classical information alone: *quantum correlations*! Entanglement is strange, useful and not completely understood. Some of the results described in this survey will shed light on this facet of quantum mechanics. In the absence of quantum correlations (if two players do not share entanglement), it is necessary to transmit n qubits to convey n bits of information [35]. When the players share quantum correlations, this can be improved to $2n$ but not more [27]. One would therefore think that quantum mechanics cannot reduce the amount of communication required in distributed tasks (by more than a constant). Surprisingly, this intuition is wrong!

We are beginning to get the idea that classical information and quantum information are quite different. As further evidence, note that classical information can trivially be copied, but quantum information is disturbed by observation and therefore cannot be faithfully copied in general. However, the fact that quantum information cannot be copied does not imply that it cannot be teleported [13].

Quantum key distribution (QKD) [12] is one of the founding results of quantum information processing. This amazing breakthrough is an amplification protocol for private shared keys. Another result that propelled quantum computation into the attractive area of research that it is today is Peter Shor's factoring algorithm [52], which is a polynomial-time algorithm to factor integers on a quantum computer. Note that the best known classical algorithm, the number field sieve, [40, 39] takes time in $O(2^{cn^{1/3}(\log n)^{2/3}})$ where n is the number of bits of the number to be factored. The importance of this result is evidenced by the fact that the security of most sensitive transactions on the Internet is based on the assumption that factoring is difficult [51].

Since quantum information cannot, even in theory, be copied, and since it is very fragile in its physical implementations, it was initially believed by some that errors would be an unsurmountable barrier to building a quantum computer. Actually, this was the first and only serious theoretical threat to quantum computers. Fortunately, quantum error correction and fault tolerant computation were shown to be possible with realistic assumptions if the rate of errors is not too big. This implies that a noisy quantum computer can perform an arbitrary long quantum computation efficiently as soon as some threshold of gate quality is attained [4]. We will not discuss quantum computer implementations but let us mention that experiments are only in their infancy. Quantum communication is the most successful present-day implementation, with QKD being implemented by dozens of research groups and being commercially available [1].

We now begin a survey of the main results in distributed computation. We will not give the quantum algorithms or protocols that solve the presented problems; they are usually

quite simple. Most of the time, the difficulty is to provide a proof of their correctness or to show that a classical computer cannot be as efficient.

Pseudo-telepathy

The term *pseudo-telepathy* originates from the authors of [18] (although it does not appear in the paper). It involves the study of a physical phenomenon that was previously studied by physicists [33, 44]. We introduce this strange behaviour of quantum mechanics with a story.

Alice and Bob claim that they have mysterious powers that enable them to perform telepathy. However surprising that this may seem, they are willing to prove their claim to a pair of physicists that do not know about quantum mechanics. Imagine that they are willing to bet a substantial amount of money. To be more precise, Alice and Bob do not claim that they can send emails by thought alone, but they claim that they can win the following game with certainty without talking to each other. As you will see, their claim is very surprising because it appears that it is impossible to satisfy!

A magic square (see Figure 1) is a 3 by 3 matrix of binary digits such that the sum of each row is even and the sum of each column is odd. A simple parity argument is sufficient to convince oneself that a magic square cannot exist: since the sum of each row is even, the sum of the whole square has to be even. But since the sum of each column is odd, the sum of the whole square has to be odd. This is a contradiction and therefore such a square cannot exist.

0	1	1
1	1	0
0	1	?

Figure 1: A partial magic square. In a magic square, the sum of each row is even and the sum of each column is odd.

In the game that Alice and Bob agree to play, they will behave exactly as if they actually agreed on a collection of such squares (at least, in a context where they cannot talk to each other). The physicists will prevent Alice and Bob from communicating during the game; an *easy* solution is to place Alice and Bob several light years away. According to relativity, any message they would exchange would take several years to arrive.

To test the purported telepathic abilities, each physicist is paired with a participant. They then ask simultaneously questions: Alice is asked to give a row of the square (either row 1, 2 or 3) and Bob is asked to give a column (either column 1, 2 or 3). Each time the experiment is performed, Alice and Bob claim to use a different magic square. After a certain number of repetitions, the physicists get together and verify that the sum of each

row is even and the sum of each column is odd, and moreover that the bit at the intersection of the row and column is the same. It is not so difficult to see that if Alice and Bob do not communicate after the onset of the game, there is no strategy that wins this game with probability more than $8/9$. This is the outcome that the pair of physicists would expect. Instead, they are astounded to see that Alice and Bob always win, no matter how many times they repeat the game! Alice and Bob have managed to win their bet and accomplish a task that provably requires communication, but without communicating! Hence the name *pseudo-telepathy*. How is this possible? Thanks to quantum mechanics, Alice and Bob can win with probability 1. In addition to agreeing on a strategy before the experiment, Alice and Bob share enough entangled particles. If you think winning such a game is amazing, then now you understand a bit more why we consider entanglement to be such a wonderful and strange resource. This simple thought experiment has very important consequences on our understanding of the world in which we live, both in the physical and philosophical perspectives [30, 9, 24].

More formally, a pseudo-telepathy game is a distributed k -player game where the players can agree on a strategy and can share entanglement. While the players are not allowed to communicate, each player is asked a question and should provide an answer. The game must be such that quantum players can win with probability 1 but classical players cannot. The example we presented comes from [7]. We refer the reader to a survey specifically on this subject [17].

Communication complexity

Communication complexity is the study of the amount of communication required to compute functions on distributed data in a context of honest cooperating players. It was first introduced by Harold Abelson [3] and given in its current form by Andrew Yao [57]. A good reference on *classical* communication complexity is [37]. There are several variations of the basic model; here, we concentrate on the most natural one. Let F be a k -input binary function. We are in a context where the k players each have one of the inputs to the function. The *probabilistic* communication complexity is the amount of bits that have to be broadcast by the players in order for player number one to be able to compute F with probability at least $2/3$ (in the worst case). We assume that the players share some random bits and that they cooperate. The value $2/3$ is arbitrary and can be very efficiently improved by parallel repetition. Note that in this model, we do not care about the computational complexity for every player, but in general the computation required by the players is polynomial. The trivial solution that works for all functions is for each player (except the first one) to broadcast his input. We will see that sometimes, but not always, the players can do much better.

Let us illustrate the concept with a simple example. Suppose we have two players, Alice and Bob, who each have a huge electronic file and they want to test if these are

identical. More formally, they want to compute the equality function. If one insists that the probability of success be 1, then Bob has to transmit his entire file to Alice: any solution would require an amount of communication equal to Bob's file size. Obviously, if we are willing to tolerate some errors, there is a more efficient solution. Let x be Alice's input and y be Bob's, and assume Alice and Bob share z , a random string of the same length as x and y . If $x = y$, obviously $x \cdot z = y \cdot z$ but it is not too hard to see that if $x \neq y$, the probability that $x \cdot z = y \cdot z$ is exactly $1/2$ (here, $x \cdot z$ is taken to be the *binary* inner product: the inner product of x and z , modulo 2). In order for Alice to learn this probabilistic information, Bob only has to send one bit. By executing this twice, we have that the function can be computed correctly with probability $3/4$.

One might argue that we are cheating by allowing Alice and Bob to share random bits and not counting this in the communication cost. We have decided to concentrate on this model since it is natural to compare it to the quantum case. Also, in general, if Alice and Bob do not share randomness, they can obtain the same result only with an additional $\log n$ bits of communication [47].

Yao is also responsible for pioneering work in the area of *quantum* communication complexity [58], in which he asked the question: what if the players are allowed to communicate qubits (quantum information) instead of classical bits. No answer to this question was initially advanced. In [25], Richard Cleve and Harry Buhrman introduced a variation on the model, for which they showed a separation between the classical and quantum models: the players communicate classically but they share entanglement instead of classical random strings. This time, the goal is to compute the function with certainty. They exhibited a function (more specifically, a *relation*, also called a *promise problem*) for three players such that in the broadcast model, any protocol that computes the function requires 3 bits of communication. In contrast, if the players share entanglement, it can be computed exactly with only 2 bits of classical communication. The function they studied is not very interesting by itself but the result is revolutionary: we knew that entanglement cannot replace communication, and what this result shows is that entanglement can be used to reduce communication in a context of communication complexity.

Harry Buhrman, Wim van Dam, Peter Høyer and Alain Tapp [22] improved the above result by exhibiting a k -player function (again with a promise) such that the communication required for computation with probability 1 is in $\Theta(k \log k)$, but if the players share quantum entanglement, it is in $\Theta(k)$. They also showed that it is possible to substitute the quantum entanglement for quantum communication, resulting in a protocol still with $O(k)$ communication. This was the first non-constant gap between quantum and classical communication complexity. Once more, the function that was studied is not natural.

Quantum teleportation [13], shows that two classical bits of communication, coupled with entanglement, enable the transfer of a qubit. Applying this, we get that any two-player protocol using quantum communication can be simulated by a protocol using entanglement and classical communication, at the cost of only doubling the communication.

The first problem of practical interest where quantum information was shown to be very

useful is the appointment scheduling problem. For this problem, Alice has an appointment calendar that, for each day, indicates whether or not she is free for lunch. Bob also has his own calendar, indicating whether or not he is free. The players wish to know if there is a day where they are both free for a lunch meeting. In the classical model, the amount of communication required to solve the problem is in $\Theta(n)$. In the quantum model, this was reduced to $O(\sqrt{n} \log n)$ in [20], and further improved to $O(\sqrt{n})$ in [2].

The first exponential separation between classical and quantum communication complexity was presented in [20] but it was in the case where the function must be computed exactly. Later, Ran Raz gave an exponential separation in the more natural probabilistic model that we have presented, but for a contrived problem [50]. See also related work [31]. Note that not all functions can be computed more efficiently using quantum communication or entanglement; this is the case of the binary inner product [27].

Other communication games

Fingerprinting

This interesting result was introduced in the context of communication complexity but is of general interest. It was shown in [19] that to any bitstring or message, a unique and very short (logarithmic) quantum fingerprint can be associated. Although the fingerprint is very small and generated deterministically, when two such fingerprints are compared, it is possible to determine with high probability if they are equal. The concepts of quantum fingerprinting were used in the context of quantum digital signatures [32].

Coin tossing

Moving to a more cryptographic context, one of the simplest and most useful primitives is the ability to flip coins fairly in an adversarial scenario. *Strong coin tossing* encompasses the intuitive features of such a protocol: it allows k players to generate a random bit with no bias (or an exponentially small one), where *bias* is the notion of a player being able to choose the outcome. The trivial method of allowing a single player to flip a coin and announce the result is biased: the player could choose the outcome to his advantage.

It is possible to base the fairness of a coin toss on computational assumptions: this is due to the fact that bit commitment can be used to implement coin toss and that bit commitment itself can be implemented with computational assumptions [28]. However, we know that when quantum computers become available, some of the assumptions on which these protocols are based will unfortunately become insecure. Is there a way to implement a coin toss using quantum information? It was shown by Andris Ambainis [5] that if two players can use quantum communication, this task can be approximated to some extent without computational assumptions. If both Alice and Bob are honest, the coin flip will be fair, otherwise one player can bias the coin toss by 25% but no more. This is

almost tight since it was proven that in this context, the bias cannot be reduced lower than approximately 21% (this result is due to unpublished work of Alexei Kitaev; see [34] for a conceptually simple proof). Recent work by André Chailloux and Iordanis Kerenidis [23] establishes an optimal coin tossing protocol. Kitaev's lower bound discouraged quantum cryptographers but it was misleading. In a context where the coin toss is used to choose a winner (a very natural application), then we know in which direction each player is trying to bias the coin toss. Surprisingly, in this context, called *weak* coin tossing, quantum protocols exist that have arbitrarily small bias [46]. See also [14] for a loss-tolerant protocol.

Quantum proofs

An area of theoretical computer science that is very important and related to complexity is the field of *proofs*. The concept of short classical proofs for a statement is captured by the complexity class NP and is the most natural. We know that many difficult problems actually have short witnesses or proofs. Can we generalize this concept in a useful and meaningful way to the quantum world? What would be a quantum proof? Would it be useful?

In a seminal paper by John Watrous [56], a specific problem, group non-membership, was shown to have short quantum proofs. It is not known (and believed to be impossible) to come up in general with a short classical proof that an element is not part of a group when the description of the group is given as a list of generators. What is amazing is that there exist quantum states that can be associated to such a problem that are short quantum proofs. More specifically, if the verifier has a quantum computer, there is a quantum algorithm that will efficiently verify the witness: if the element is in the group, no quantum state will make the verifier's algorithm accept with non-negligible probability, whereas if the element is not in the group, there is a quantum state that will make the algorithm accept with probability 1.

Classical simulation of entanglement

In previous sections, we presented several examples where entanglement can be used to solve distributed computing problems more efficiently. In physics and computer science, an active area of research is dealing with the opposite problem, the simulation of entanglement using classical communication. The objective is to exactly reproduce the distribution of measurement outcomes, as if they were performed on entangled qubits. The distant players are assumed to share continuous random variables; otherwise it is known to be impossible. The first protocol to simulate a maximally entangled pair of qubits using classical communication was presented in [43]. The protocol uses an expected 1.74 bits of communication but to be able to simulate a maximally entangled pair of qubits perfectly, the amount of communication is not bounded. In [18], a simulation was presented using exactly 8 bits of communication and this was later improved to 1 bit [55].

In general, looking at the classical communication complexity (with shared randomness) for pseudo-telepathy games tells us how difficult it is to simulate entanglement. Using this idea, it is proved in [18] that n maximally-entangled qubits require an exponential amount of communication for perfect simulation. Some protocols actually exist that accomplish this almost tightly with an expected amount of communication for *general* measurements [42].

Leader election

An important problem in distributed computation is *leader election*. The task is the following: within a group of distributed parties that communicate according to an underlying graph structure, create the situation where a single participant is singled out as the *leader*. In [54], Seiichiro Tani, Hirotsada Kobayashi, and Keiji Matsumoto give an error-free protocol for the leader election problem in the anonymous networks setting. This contrasts with the fact that any classical bounded-time protocol must have a non-zero error probability. See also [29].

Byzantine agreement

In *Byzantine agreement* [38] (also called *consensus*), a group of participants must decide on a common output bit value, despite the intervention of an adversary. The participants each have an input bit. The protocol must be such that all honest participants agree on the same output bit, and if the inputs were all $c \in \{0, 1\}$, the common output bit is c . The complexity of the protocol is measured by the maximal expected number of rounds it takes for the protocol to terminate. In [11], Michael Ben-Or and Avinatan Hassidim present a fast quantum Byzantine agreement protocol that can reach agreement in $O(1)$ expected communication rounds. This beats the known classical lower bound of $O(\sqrt{n/\log n})$ and the $O(n)$ complexity of best known classical solutions.

Protocols for quantum information

If we choose to deal with tasks involving quantum information instead of classical information, there are a lot of results and possibilities. Quantum teleportation is the most famous [13], but all sorts of channels have been studied for quantum communication. On the cryptography side, we know protocols to encrypt [6] and authenticate quantum messages [8]. It is possible to perform multi-party computation with quantum inputs and outputs in a secure way [10]. It is also possible to anonymously transmit quantum messages [16].

Conclusion

We have given the reader a glimpse of distributed computing in the quantum world. Following the main lines of our survey, we now present a partial list of open questions.

Characterization of games that exhibit pseudo-telepathy. One way to recognize a pseudo-telepathy game is to find a perfect quantum strategy and then show that there is no such classical strategy. We would like a more natural way to recognize such a game, relying more on its underlying structure.

Quantum parallel repetition. What is the best probability of success for Alice and Bob who are involved in many *parallel* instances of the same game, using entanglement? For purely classical games, the probability of success decreases at an exponential rate [49] (as surprising as it sounds, the probability does not decrease at the same rate as one might expect and this result is far from being trivial). This question asks whether or not there is a similar theorem for the case that the players use shared entanglement. A special case was answered in the affirmative by [26].

Quantum communication complexity: qubits versus entanglement. As mentioned, we know that teleportation can be used to transform any two-player protocol using quantum communication into a protocol using entanglement, at a cost of only two classical bits per qubit in the original protocol. This question asks whether or not we can do the same thing, up to a constant factor, in the *other* direction. Related work in this direction includes [41], where it is shown that in a slightly different scenario, there exist tasks for which no finite amount of entanglement yields an optimal strategy.

Simulation of multi-party entanglement. In contrast to the two-party case, very little is known about the simulation of multi-party entangled states. In particular, it is not even known if this general task is possible with bounded communication.

Acknowledgements

We thank John Watrous for related insightful discussions and Michael Ben-Or for pointing out references on quantum leader election and Byzantine agreement.

References

- [1] <http://www.idquantique.com/>.
- [2] S. Aaronson and A. Ambainis. Quantum search of spatial regions. *Theory of Computing*, 1:47–79, 2005.
- [3] H. Abelson. Lower bounds on information transfer in distributed computations. *Journal of the ACM*, 27:384–392, 1980. First published in 1978.
- [4] P. Aliferis, D. Gottesman, and J. Preskill. Quantum accuracy threshold for concatenated distance-3 codes. *Quantum Information and Computation*, 6:97–165, 2006.

- [5] A. Ambainis. A new protocol and lower bounds for quantum coin flipping. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing (STOC 2001)*, pages 398–416, 2001.
- [6] A. Ambainis, M. Mosca, A. Tapp, and R. de Wolf. Private quantum channels. In *Proceedings of the 41st Annual Symposium on Foundations of Computer Science (FOCS 2000)*, pages 547–553, 2000.
- [7] P.K. Aravind. Bell’s theorem without inequalities and only two distant observers. *Foundations of Physics Letters*, 15:397–405, 2002.
- [8] H. Barnum, C. Crépeau, D. Gottesman, and A. Tapp. Authentication of quantum messages. In *Proceedings of the 43rd Annual IEEE Symposium on Foundations of Computer Science (FOCS 2002)*, pages 449–458, 2002.
- [9] J.S. Bell. On the Einstein-Podolsky-Rosen paradox. *Physics*, 1:195–200, 1964.
- [10] M. Ben-Or, C. Crépeau, D. Gottesman, A. Hassidim, and A. Smith. Secure multiparty quantum computation with (only) a strict honest majority. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2006)*, pages 249–260, 2006.
- [11] M. Ben-Or and A. Hassidim. Fast quantum Byzantine agreement. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC 2005)*, pages 481–485, 2005.
- [12] C.H. Bennett and G. Brassard. Quantum cryptography: Public key distribution and coin tossing. In *Proceedings of IEEE International Conference on Computers, Systems and Signal Processing*, pages 175–179, 1984.
- [13] C.H. Bennett, G. Brassard, C. Crépeau, R. Jozsa, A. Peres, and W.K. Wootters. Teleporting an unknown quantum state via dual classical and Einstein-Podolsky-Rosen channels. *Physical Review Letters*, 70:1895–1899, 1993.
- [14] G. Berlin, G. Brassard, F. Bussi eres, and N. Godbout. Loss-tolerant quantum coin flipping. In *Proceedings of the Second International Conference on Quantum, Nano, and Micro Technologies (ICQNM08)*, pages 1–9, 2008.
- [15] G. Brassard. Quantum communication complexity. *Foundations of Physics*, 33:1593–1616, 2003.
- [16] G. Brassard, A. Broadbent, J. Fitzsimons, S. Gambs, and A. Tapp. Anonymous quantum communication. In *Proceedings of 13th International Conference on the Theory and Application of Cryptology and Information Security (ASIACRYPT 2007)*, pages 460–473, 2007.

- [17] G. Brassard, A. Broadbent, and A. Tapp. Quantum pseudo-telepathy. *Foundations of Physics*, 35:1877–1907, 2005.
- [18] G. Brassard, R. Cleve, and A. Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83:1874–1878, 1999.
- [19] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87:167902 [4 pages], 2001.
- [20] H. Buhrman, R. Cleve, and A. Wigderson. Quantum vs. classical communication and computation. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing (STOC 1998)*, pages 63–68, 1998.
- [21] H. Buhrman and H. Röhrig. Distributed quantum computing. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science 2003 (MFCS 2003)*, pages 1–20, 2003.
- [22] H. Buhrman, P. Høyer W. van Dam, and A. Tapp. Multiparty quantum communication complexity. *Physical Review A*, 60:2737–2741, 1999.
- [23] A. Chailloux and I. Kerenidis. Optimal quantum strong coin flipping. Available as <http://arxiv.org/abs/0904.1511>, 2009.
- [24] F. Clauser, M. A. Horne, A. Shimony, and R. A. Holt. Proposed experiment to test local hidden-variable theories. *Physical Review Letters*, 23:880–884, 1969.
- [25] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56:1201–1204, 1997.
- [26] R. Cleve, W. Slofstra, F. Unger, and S. Upadhyay. Perfect parallel repetition theorem for quantum XOR proof systems. In *Proceedings of the 22nd Annual IEEE Computational Complexity Conference (CCC '07)*, pages 109–114, 2007.
- [27] R. Cleve, W. van Dam, M. Nielsen, and A. Tapp. Quantum entanglement and the communication complexity of the inner product function. In *Proceedings of Quantum computing and quantum communication: First NASA International Conference (QCC'98)*, pages 61–74, 1998.
- [28] C. Crépeau. *Encyclopedia of Cryptography and Security*, chapter Commitment, pages 83–86. Springer, 2005.
- [29] E. D’Hondt and P. Panangaden. The computational power of the W and GHZ states. *Quantum Information and Computation*, 6:173–183, 2006.
- [30] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.

- [31] D. Gavinsky and P. Pudlák. Exponential separation of quantum and classical non-interactive multi-party communication complexity. In *Proceedings of the 23rd Annual IEEE Conference on Computational Complexity (CCC '08)*, pages 332–339, 2008.
- [32] D. Gottesman and I. Chuang. Quantum digital signatures. Available as <http://arxiv.org/abs/quant-ph/0105032>, 2001.
- [33] D.M. Greenberger, M. A. Horne, and A. Zeilinger. Going beyond Bell’s theorem. In *Bell’s Theorem, Quantum Theory and Conceptions of the Universe*, pages 69–72, 1988.
- [34] G. Gutoski and J. Watrous. Toward a general theory of quantum games. In *Proceedings of the 39th ACM Symposium on Theory of Computing (STOC 2007)*, pages 565–574, 2007.
- [35] A. Holevo. Bounds for the quantity of information transmitted by a quantum communication channel. *Problemy Peredachi Informatsii*, 9:3–11, 1973. English translation in *Problems of Information Transmission*, 9:177–183, 1973.
- [36] P. Kaye, R. Laflamme, and M. Mosca. *An Introduction to Quantum Computing*. Oxford University Press, 2007.
- [37] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.
- [38] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4:382–401, 1982.
- [39] A. K. Lenstra and H. W. Lenstra, Jr., editors. *The Development of the Number Field Sieve*, volume 1554 of *Lecture Notes in Mathematics*. Springer, 1993.
- [40] A. K. Lenstra, H. W. Lenstra, Jr., M. S. Manasse, and J. M. Pollard. The number field sieve. In *Proceedings of the 22nd Annual ACM Symposium on Theory of Computing (STOC 1990)*, pages 564–572, 1990. An expanded version appears in [39], pp. 11–42.
- [41] D. Leung, B. Toner, and J. Watrous. Coherent state exchange in multi-prover quantum interactive proof systems. Available as <http://arxiv.org/abs/0804.4118>, 2008.
- [42] S. Massar, D. Bacon, N. J. Cerf, and R. Cleve. Classical simulation of quantum entanglement without local hidden variables. *Physical Review A*, 63:052305 [8 pages], 2001.
- [43] T. Maudlin. Bell’s inequality, information transmission, and prism models. In *Proceedings of the Biennial Meeting of the Philosophy of Science Association*, volume one: contributed papers, pages 404–417, 1992.

- [44] N. D. Mermin. Quantum mysteries revisited. *American Journal of Physics*, 58(8):731–734, 1990.
- [45] N. D. Mermin. *Quantum Computer Science: An Introduction*. Cambridge University Press, 2007.
- [46] C. Mochon. Quantum weak coin flipping with arbitrarily small bias. Available as <http://arxiv.org/abs/0711.4114>, 2007.
- [47] I. Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, pages 67–71, 1991.
- [48] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.
- [49] R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27:763–803, 1998.
- [50] R. Raz. Exponential separation of quantum and classical communication complexity. *Proceedings of the 31st Annual ACM Symposium on Theory of Computing (STOC 1999)*, pages 358–367, 1999.
- [51] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Communications of the ACM*, 21:120–126, 1978.
- [52] P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM Journal on Computing*, 26:1484–1509, 1997. First published in 1995.
- [53] A. Steane and W. van Dam. Physicists triumph at guess my number. *Physics Today*, 53:35–39, 2000.
- [54] S. Tani, H. Kobayashi, and K. Matsumoto. Exact quantum algorithms for the leader election problem. In *Proceedings of the 22nd International Symposium on Theoretical Aspects of Computer Science (STACS 2005)*, pages 581–592, 2005.
- [55] B. F. Toner and D. Bacon. Communication cost of simulating Bell correlations. *Physical Review Letters*, 91:187904 [4 pages], 2003.
- [56] J. Watrous. Succinct quantum proofs for properties of finite groups. In *Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2000)*, pages 537–546, 2000.
- [57] A. C.-C. Yao. Some complexity questions related to distributive computing. In *Proceedings of the 11th Annual ACM Symposium on Theory of Computing (STOC 1979)*, pages 209–213, 1979.

- [58] A. C.-C. Yao. Quantum circuit complexity. In *Proceedings of the 34th Annual IEEE Symposium on Foundations of Computer Science (STOC 1993)*, pages 222–227, 1993.