ACM SIGACT News Distributed Computing Column 31

Quantum Computers Meet Distributed Computing

Idit Keidar
Dept. of Electrical Engineering, Technion
Haifa, 32000, Israel
idish@ee.technion.ac.il

After two columns on practical problems arising in current day technologies (multicores in Column 29; systems research in Column 30), this column takes a sharp turn towards the futuristic realm of quantum computations. More specifically, the column features two surveys of *distributed quantum computing*, which, unbeknownst to many distributed computing folks, is an active area of research.

First, Anne Broadbent and Alain Tapp provide a broad overview of distributed computations and multi-party protocols that can benefit from quantum mechanics, most notably from *entanglement*. Some of these are unsolvable with classical computing, for example, pseudo-telepathy. In other cases, like appointment scheduling, the problem's communication complexity can be reduced by quantum means.

Next, Vasil Denchev and Gopal Pandurangan critically examine the joint future of quantum computers and distributed computing, asking whether this is a new frontier . . . or science fiction. They give background to the lay reader on quantum mechanics concepts that provide added value over classical computing, (again, entanglement figures prominently). They also elaborate on the practical difficulties of implementing them. They then illustrate how these concepts can be exploited for two goals: (1) to distribute centralized quantum algorithms over multiple small quantum computers; and (2) to solve leader election in various distributed computing models. They conclude that the jury is still out on the cost-effectiveness of quantum distributed computing.

Both surveys outline open questions and directions for future research. Many thanks to Anne, Alain, Vasil and Gopal for their contributions!

Call for contributions: I welcome suggestions for material to include in this column, including news, reviews, open problems, tutorials and surveys, either exposing the community to new and interesting topics, or providing new insight on well-studied topics by organizing them in new ways.

Can Quantum Mechanics Help Distributed Computing?

Anne Broadbent and Alain Tapp
Département d'informatique et de recherche opérationnelle
Université de Montréal, C.P. 6128, Succ. Centre-Ville
Montréal (QC), H3C 3J7, Canada
{broadbea,tappa}@iro.umontreal.ca

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 in the field of theoretical computer science having only very limited knowledge of quantum computation. 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 [49, 14, 20, 16].

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 *qu*bit. 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 [44, 41, 33] 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 [32]. When the players share quantum correlations, this can be improved to $2n$ but not more [25]. 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. Note that the fact that quantum information cannot be copied does not imply that it cannot be teleported [12].

Quantum key distribution (QKD) [11] 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 [48], which is a polynomial-time algorithm to factor integers on a quantum computer. Note that the best known classical algorithm, the number field sieve, [36, 35] 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 [47].

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 [17] (although it does not appear in the paper). It involves the study of a physical phenomenon that was previously studied by physicists [30, 40]. 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 a 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 used 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 [27, 9, 22].

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 [16].

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 [52]. A good reference on *classical* communication complexity is [34]. 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 [43].

Yao is also responsible for pioneering work in the area of *quantum* communication complexity [53], 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 [23], 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 [21] 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 [12], 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 [19], and further improved to $O(\sqrt{n})$ in [2].

The first exponential separation between classical and quantum communication complexity was presented in [19] 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 [46]. See also related work [28]. Note that not all functions can be computed more efficiently using quantum communication or entanglement; this is the case of the binary inner product [25].

Other communication games

Fingerprinting

This interesting result was introduced in the context of communication complexity but is of general interest. It was shown in [18] 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 [29].

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 [26]. 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 [31] for a conceptually simple proof). This 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, quantum protocols exist that have arbitrarily small bias [42]. See also [13] 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 [51], 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 [39]. 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 [17], a simulation was presented using exactly 8 bits of communication and this was later improved to 1 bit [50].

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 [17] that n maximally-entangled qubits require an exponential amount of communication to be simulated perfectly. Some protocols actually exist that accomplish this almost tightly with an expected amount of communication for *general* measurements [38].

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 [12], 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 [15].

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 the underlying structure of the game.

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 [45] (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 [24].

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 [37], 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

Thanks to Idit Keidar for inviting us to write this survey which, we hope, will create awareness of quantum computation in the community. We also thank John Watrous for related insightful discussions.

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] 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.

[12] 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.

[13] G. Berlin, G. Brassard, F. Bussières, 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.

[14] G. Brassard. Quantum communication complexity. *Foundations of Physics*, 33:1593–1616, 2003.

[15] 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.

[16] G. Brassard, A. Broadbent, and A. Tapp. Quantum pseudo-telepathy. *Foundations of Physics*, 35:1877–1907, 2005.

[17] G. Brassard, R. Cleve, and A. Tapp. Cost of exactly simulating quantum entanglement with classical communication. *Physical Review Letters*, 83:1874–1878, 1999.

[18] H. Buhrman, R. Cleve, J. Watrous, and R. de Wolf. Quantum fingerprinting. *Physical Review Letters*, 87:167902 [4 pages], 2001.

[19] 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.

[20] 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.

[21] H. Buhrman, P. Høyer W. van Dam, and A. Tapp. Multiparty quantum communication complexity. *Physical Review A*, 60:2737–2741, 1999.

[22] 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.

[23] R. Cleve and H. Buhrman. Substituting quantum entanglement for communication. *Physical Review A*, 56:1201–1204, 1997.

[24] 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.

[25] 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 (QCQC'98)*, pages 61–74, 1998.

[26] C. Crépeau. *Encyclopedia of Cryptography and Security*, chapter Commitment, pages 83–86. Springer, 2005.

[27] A. Einstein, B. Podolsky, and N. Rosen. Can quantum-mechanical description of physical reality be considered complete? *Physical Review*, 47:777–780, 1935.

[28] 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*, pages 332–339, 2008.

[29] D. Gottesman and I. Chuang. Quantum digital signatures. Available as `http://arxiv.org/abs/quant-ph/0105032`, 2001.

[30] 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.

[31] 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.

[32] 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.

[33] P. Kaye, R. Laflamme, and M. Mosca. *An Introduction to Quantum Computing*. Oxford University Press, 2007.

[34] E. Kushilevitz and N. Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[35] 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.

[36] 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 [35], pp. 11–42.

[37] 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.

[38] 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.

[39] 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.

[40] N. D. Mermin. Quantum mysteries revisited. *American Journal of Physics*, 58(8):731–734, 1990.

[41] N. D. Mermin. *Quantum Computer Science: An Introduction*. Cambridge University Press, 2007.

[42] C. Mochon. Quantum weak coin flipping with arbitrarily small bias. Available as `http://arxiv.org/abs/0711.4114`, 2007.

[43] I. Newman. Private vs. common random bits in communication complexity. *Information Processing Letters*, pages 67–71, 1991.

[44] M. A. Nielsen and I. L. Chuang. *Quantum Computation and Quantum Information*. Cambridge University Press, 2000.

[45] R. Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27:763–803, 1998.

[46] 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.

[47] 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.

[48] 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.

[49] A. Steane and W. van Dam. Physicists triumph at guess my number. *Physics Today*, 53:35–39, 2000.

[50] B. F. Toner and D. Bacon. Communication cost of simulating Bell correlations. *Physical Review Letters*, 91:187904 [4 pages], 2003.

[51] 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.

[52] 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.

[53] 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.

Distributed Quantum Computing: A New Frontier in Distributed Systems or Science Fiction?

Vasil S. Denchev and Gopal Pandurangan[1]
Department of Computer Science
Purdue University
West Lafayette, IN 47907, USA
{vdenchev,gopal}@cs.purdue.edu

Abstract

Quantum computing and distributed systems may enter a mutually beneficial partnership in the future. On the one hand, it is much easier to build a number of small quantum computers rather than a single large one. On the other hand, the best results concerning some of the fundamental problems in distributed computing can potentially be dramatically improved upon by taking advantage of the superior resources and processing power that quantum mechanics offers. This survey has the purpose to highlight both of these benefits. We first review the current results regarding the implementation of arbitrary quantum algorithms on distributed hardware. We then discuss existing proposals for quantum solutions of leader election — a fundamental problem from distributed computing. Quantum mechanics allows leader election to be solved with no communication, provided that certain pre-shared entanglement is already in place. Further, an impossibility result from classical distributed computing is circumvented by the quantum solution of anonymous leader election — a unique leader is elected in finite time with certainty. Finally, we discuss the viability of these proposals from a practical perspective. Although, theoretically, distributed quantum computing looks promising, it is still unclear how to build quantum hardware and how to create and maintain robust large-scale entangled states. Moreover, it is not clear whether the costs of creating entangled states and working with them are smaller than the costs of existing classical solutions.

1 Introduction

In recent years, quantum computing has been widely advertised as the next ground-breaking technological innovation that holds the promise to fundamentally change the way we do computing. Futurists and lay people, as well as serious researchers from several diverse scientific areas, have been fascinated by the potential advantages that quantum computing shows.

But harnessing the counter-intuitive laws of quantum mechanics has proven to be a hard practical problem. Today there are just a few successful implementations of small quantum computers. Unfortunately, scalability is an insurmountable problem for all of them, which is why they are used only for trivial computations to illustrate the potential advantages of quantum algorithms. Consequently, some of the research that we discuss in this survey is motivated by the possibility to overcome the scalability problems of current technology. Such a goal can be achieved by using a collection of small quantum computers to solve problems in a distributed manner via casting known centralized algorithms into their distributed versions. We discuss this strategy in Section 4.

Distributed systems can benefit from quantum technology as well, assuming the ability to efficiently create and reliably use quantum entanglement. The phenomenon of entanglement counter-intuitively invalidates the notion of *local realism* by creating non-local relationships between quantum objects and blurring the physical state until a measurement is done. Considering two entangled particles, the state of each of them is a superposition of the possible values of some physical property and the joint state cannot be decomposed into a product of single-particle states. Non-locality is manifested by the fact that a measurement done on one of the particles not only collapses the superposition of the initial quantum state of the measured particle to a single definite value, but it also instantaneously collapses the state of the other particle to a corresponding definite value, regardless of the spatial separation of the two particles. It is this "spooky action at a distance" that strongly disturbed Einstein [13] but was nevertheless confirmed later [1]. Entanglement is briefly discussed in Section 2.3, but a far more detailed treatment can be found in [24].

The theoretical proposals that we overview in Section 5 offer impressive solutions for leader election — a fundamental problem that sometimes can be a performance bottleneck, because it has to be routinely solved in distributed computing. Section 5.1 presents quantum solutions for leader election without communication but with entanglement that has been shared among the participating processors. There is one main trick that makes these schemes work: choosing the specific form of entanglement in a way that ensures that measuring the entangled particles results in a collapsed global state that satisfies the requirements for a valid solution. Section 5.2 and Section 5.3 consider the anonymous version of leader election and show how an impossibility result [21] from classical distributed computing is violated in the quantum world — a unique leader in an anonymous network is elected in finite time with certainty. Classically, anonymous leader election in networks with arbitrary topology is solved with high probability by randomization [21, 33]. However, in the quantum world, entanglement can be used to break symmetry even in completely symmetric networks. Section 5.2 uses the same strategy of pre-shared entanglement as Section 5.1. Section 5.3 presents a more intricate algorithm that does not assume any pre-shared entanglement but uses quantum communication to create certain entangled states that do not necessarily guarantee that the leader is chosen in a single step. Nevertheless, they guarantee that a leader is chosen with certainty after a finite number of steps of gradual symmetry-breaking. The main trick here is quantum amplitude amplification [4], which is also the essential technique that is used in Grover's search algorithm [17]. We briefly introduce quantum amplitude amplification in Section 2.4.

The quantum solutions of distributed problems are quite impressive, but in practice there are very serious problems related to the implementation of useful quantum devices. Quantum entanglement appears as a basic requirement for the functioning of any quantum algorithm that claims any advantages over its classical counterpart, which motivates the conjecture that entanglement is the fundamental source of all quantum speedups [18, 19, 20, 15]. It appears that quantum entanglement is a new fundamental resource, the likes of which have never been known in classical computing. The difficulty here is that a complete understanding of entanglement has not been achieved yet. Simple cases of it have been explored by experimental physicists [23, 11, 34, 26], but nobody has attempted to build the large-scale entangled states that are assumed for the solutions of leader election. As a consequence, we currently do not know whether the costs of creating, using, and maintaining entanglement do not outweigh the advantages that it offers. For example, in the context of leader election, it might very well be the case that entangling a set of processors is harder than simply a-priori choosing a leader and equipping each processor with knowledge about the chosen leader. In the context of anonymous leader election, the additional cost of quantum processing brings up the question whether it is worth having a quantum algorithm that elects a unique leader in finite time with certainty when there are more efficient classical algorithms that solve the problem with high probability. Moreover, the anonymous leader election algorithm that we present in Section 5.3 suffers from two additional drawbacks: (i) The crucial underlying technique of quantum amplitude amplification may not be implementable at a constant cost as the algorithm assumes; (ii) Its physical implementation may not be able to preserve the theoretically promised certainty of always electing a unique leader in finite time.

We assume that the reader is familiar with the basic concepts of distributed computing. Comprehensive discussions on the existing classical algorithms for leader election can be found in [33, 21]. In the next section we give a basic and by no means complete introduction to the relevant background in quantum computing. We restrain ourselves just to defining the essential terms and concepts that are crucial for understanding our subsequent discussion.

2 Quantum Computing Background

2.1 Quantum States

A comprehensive introduction to quantum computing can be found in [24]. The qubit, the quantum version of the classical bit, is defined in two-dimensional complex vector space. Dirac notation is the accepted standard notation: two of the possible states for a qubit are $|0\rangle$ and $|1\rangle$, which are known as the computational basis states corresponding to the classical logical values of 0 and 1. Sometimes their vector representations are used: $|0\rangle = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ and $|1\rangle = \begin{pmatrix} 0 \\ 1 \end{pmatrix}$. The novelty here is that a qubit can be in many other states as well—different superpositions of the $|0\rangle$ and $|1\rangle$ basis states. In fact, there is an infinite number of possible quantum states, because a general quantum state is of the form $|\psi\rangle = \alpha|0\rangle + \beta|1\rangle$, where $\alpha, \beta \in \mathbb{C}$ are the probability amplitudes of the corresponding components of the superposition and they have to obey the normalization rule: $|\alpha|^2 + |\beta|^2 = 1$. Even though a qubit can be simultaneously storing two logical values, when a measurement in the $|0\rangle$, $|1\rangle$ basis is performed on it, the outcome is either one of them but not both. After a measurement is performed, and a classical bit is obtained from it, we say that the quantum state has been collapsed. The role of the probability amplitudes becomes apparent when we collect statistics about the measurement results for a multitude of identically prepared states. The square of the absolute value of a probability amplitude predicts the portion of the total number of measurements where the corresponding component is observed as the resulting collapsed state.

The joint state of multiple qubits is described by the tensor product ($\otimes$) of the single-qubit states of the individual qubits. For two vectors x and y of dimensions m and n, $x \otimes y$ is a vector of dimension mn. For example, $|\psi\rangle \otimes |\phi\rangle$, the two-qubit joint state of qubits $|\psi\rangle$ and $|\phi\rangle$, is a vector of dimension 4 because the single-qubit states are vectors of dimension 2. The $\otimes$ symbol can be omitted whenever the tensor product (multi-qubit state) is obvious, so the notation for this example can also be $|\psi, \phi\rangle$ or $|\psi\phi\rangle$. Also, $|\psi\rangle^{\otimes k} = \underbrace{|\psi\rangle \otimes |\psi\rangle \otimes \ldots \otimes |\psi\rangle}_{\text{k times}}$ denotes a k-qubit state in which all individual qubits are in the state $|\psi\rangle$ simultaneously.

2.2 Quantum Gates

A quantum gate transforms the state of a quantum system. A single-qubit input state of the general form $\alpha|0\rangle+\beta|1\rangle$ is transformed to $\gamma|0\rangle+\delta|1\rangle$, where γ and δ depend on α, β, and the definition of the transforming gate. By the laws of quantum mechanics, any such transformation must be unitary, i.e. any n-qubit quantum gate must be representable as a $2^n \times 2^n$ matrix U, where $U^\dagger U = I$ and $U^\dagger$ is the conjugate transpose of U. The action of a gate on an input state is described by matrix-vector multiplication: $Gx = y$, where G is some quantum gate, x is the input, and y is the output. Because unitarity is the only restriction on quantum gates, there is an infinite set of possible quantum gates, but only a small subset of them are used often and have become standardized. The standard quantum gates that are used in our subsequent discussion are the Z (corresponding to the Pauli Z matrix), Hadamard, rotation, controlled-NOT (CNOT), and swap gates. We refer the reader to [24] for descriptions of the other standard gates.

The Z and Hadamard gates are single-qubit gates: $Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}$ and $H = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$. The matrix-vector multiplication of the matrix for the Z gate with the vector representations of $|0\rangle$ and $|1\rangle$ shows that the Z gate does not change the $|0\rangle$ basis state and multiplies by -1 the $|1\rangle$ basis state. Similarly, the Hadamard gate outputs $(|0\rangle + |1\rangle)/\sqrt{2}$ when the input is $|0\rangle$ and $(|0\rangle - |1\rangle)/\sqrt{2}$ when the input is $|1\rangle$. Because of the linearity of quantum gates, when the input is not just one of the basis states but some superposition of them, the output consists of the superposed action of the gate on the basis states that compose the input state. For example, the output of the Z gate on the $(|0\rangle+|1\rangle)/\sqrt{2}$ input is $(|0\rangle-|1\rangle)/\sqrt{2}$. A rotation gate performs a phase rotation by an arbitrary angle in the Bloch sphere[1]. Rotation gates form a family of quantum gates, which serves as a generalization of all single-qubit gates.

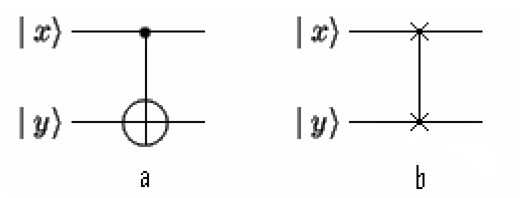

Figure 2: CNOT and swap gates.

The CNOT gate (Fig. 1a) has two inputs: the $|x\rangle$ input is called *control* and the $|y\rangle$ input is *target*. The gate performs addition modulo 2 of the control and target qubits, stores the result in the target qubit, and leaves the control qubit unchanged. In other words, the target is inverted when the control is $|1\rangle$ and is left unchanged when the control is $|0\rangle$, while the control always remains unchanged. Again, when the input is a superposition of the basis states, the output is a superposition of the actions of the gate on the corresponding basis-state components of the input state. For example, when the input to the CNOT gate is $(|00\rangle + |11\rangle)/\sqrt{2}$, the output is $(|00\rangle + |10\rangle)/\sqrt{2}$. This is so, because the $|00\rangle$ component of the input superposition is left unchanged by the gate (the control is $|0\rangle$) and the $|11\rangle$ component is transformed to $|10\rangle$ (the control is $|1\rangle$, so the target is inverted). The operation of the swap gate (Fig. 1b) is even more straightforward: the two inputs are simply exchanged.

Nielsen and Chuang provide in [24] a formal proof that the CNOT and general single-qubit gates form a universal set for quantum computation, i.e. any quantum algorithm can be expressed as a circuit consisting

[1] The Bloch sphere is the three-dimensional unit sphere. Any single-qubit quantum state can be represented as a point on the Bloch sphere, which is why this simple abstraction has traditionally served to describe single-qubit states and arbitrary transformations on them.

only of these gates. They also show that there exist unitary transformations that require compositions of exponential numbers of gates from the universal set. For example, the unitary transformation that solves an NP-C problem may require exponentially many such gates. However, the goal of quantum computation is exactly to find interesting transformations that can be performed more efficiently than what is possible in the classical world.

2.3 Quantum Entanglement

Quantum entanglement involves the joint state of two or more qubits. Informally, we say that when a collection of qubits is entangled, they are non-locally correlated in the sense that performing a measurement on one of them instantaneously affects the state of the other(s), even when there are arbitrarily large spatial separations between individual qubits.

However, there is an immediate concern that the presence of non-local correlations in entangled systems might be implying the possibility for super-luminal signaling, i.e. being able to communicate faster than the speed of light. Fortunately, this apparent conflict with Albert Einstein's Theory of Relativity [12] has been settled down by the No-Signaling Theorem [9, 27], which proves the impossibility to directly use the instantaneous effects of quantum entanglement to transmit useful information. As it turns out, entanglement alone cannot be used for communication. On the other hand, in the case of quantum teleportation [24], communication of quantum states is achieved with the help of classical signaling, which is clearly bounded by the speed of light.

There are two types of entanglement that we are interested in for the purposes of distributed computing: GHZ [16, 8] and W [8]. A collection of n qubits can be entangled in an n-partite GHZ state of the form:

$$|\psi\rangle = (|000\ldots000\rangle + |111\ldots111\rangle)/\sqrt{2} \tag{1}$$

The n-partite W state, on the other hand, looks like this:

$$|\gamma\rangle = (|00\ldots01\rangle + |00\ldots10\rangle + \cdots + |01\ldots00\rangle + |10\ldots00\rangle)/\sqrt{n} \tag{2}$$

It is not difficult to see that these two types of entanglement are quite different from each other. They exhibit different degrees of "persistency". Notice that all of the qubits in the GHZ state are collapsed to a definite state ($|0\rangle$ or $|1\rangle$) by measuring exactly one of them. In contrast, destroying the entanglement of a W state requires in general $n-1$ qubits to be measured, because for any fixed qubit, it has a single component that can give a measurement result of 1 and $n-1$ components that can give a measurement result of 0. Hence, if measuring one qubit gives 0, only a single component of the superposition is eliminated (the one that represented the possibility of measurement result of 1), but $n-1$ more components remain superposed.

Fig. 2 shows the circuit that creates a 2-partite GHZ state (also known as a Bell pair or EPR pair). The initial state of the two qubits is $|00\rangle$. The Hadamard gate puts the control qubit in an equal superposition: $(|0\rangle + |1\rangle)/\sqrt{2}$, so the joint state becomes $(|00\rangle + |10\rangle)/\sqrt{2}$. Now the CNOT gate is applied: the first component of the superposition does not change, because its control qubit is $|0\rangle$, but the target in the second component becomes $|1\rangle$ because the control is $|1\rangle$. As a result, the final state is $(|00\rangle + |11\rangle)/\sqrt{2}$. This entangled state can be augmented with more qubits by making each of them the target of a CNOT gate controlled by any of the already entangled qubits.

2.4 Quantum Amplitude Amplification

Quantum amplitude amplification appears as the essential technique in Grover's search algorithm [17]. An unstructured database is searched for an item that matches the search criteria according to some oracle. If the

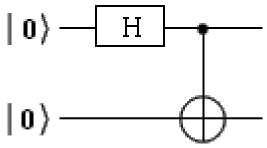

Figure 3: Locally entangling two qubits.

database is of size n and there are m items that match the search criteria, the necessary number of steps for finding one of the m items is $\Omega(n/m)$ according to the classical lower bound. However, Grover's algorithm succeeds with high probability in just $O(\sqrt{n/m})$ steps by using quantum amplitude amplification. The algorithm starts by preparing a uniform superposition of all items in the database. At this point, if the quantum state is measured, there is only an m/n chance of obtaining one of the desired items. Grover's algorithm does not do that, but instead, it proceeds in a number of successive steps of gradually increasing the amplitudes of the desired item(s) at the expense of the amplitudes of the rest of the items in the database. The novelty here is that just $\sqrt{n/m}$ such steps are enough to boost the amplitudes of the m good items, so that a measurement in the end yields one of them with high probability.

Grover's algorithm is just one application of quantum amplitude amplification, which is generalized in [4]. In general, this technique allows the amplitudes of chosen components from a superposition state to be amplified at the expense of others, whose amplitudes get attenuated. For example, consider the 2-qubit state consisting of qubits $|q_0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and $|q_1\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$: $|\phi\rangle = |q_0, q_1\rangle = (|0\rangle + |1\rangle)(|0\rangle + |1\rangle)/2 = (|00\rangle + |01\rangle + |10\rangle + |11\rangle)/2$. In this state, all of the four possible components are superposed with equal amplitudes. However, if the Hadamard gate is applied to $|q_0\rangle$, the 2-qubit state becomes $\phi' = |q_0', q_1\rangle = (|00\rangle + |01\rangle)/\sqrt{2}$, because the Hadamard gate transforms the state $|q_0\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ into $|q_0'\rangle = |0\rangle$. Clearly, the amplitudes of the $|00\rangle$ and $|01\rangle$ components of $|\phi\rangle$ are amplified at the expense of the $|10\rangle$ and $|11\rangle$ components, whose amplitudes get obliterated. Further, applying a Hadamard gate to $|q_1\rangle$ as well transforms the state $|\phi'\rangle$ into $|\phi''\rangle = |q_0', q_1'\rangle = |00\rangle$. Now, from the point of view of quantum amplitude amplification, the state $|\phi''\rangle$ is obtained from $|\phi\rangle$ by maximizing the amplitude of the $|00\rangle$ component at the expense of the rest of the components of $|\phi\rangle$.

In the anonymous leader election algorithm that is presented in Section 5.3, quantum amplitude amplification is used to achieve symmetry-breaking even in completely symmetric networks. The starting state is a superposition of symmetric components, whose amplitudes get obliterated to the benefit of resulting asymmetric components. When that state is measured and asymmetry is materialized, the processors can be divided in at least two non-overlapping non-empty groups, which allows a leader to be chosen with certainty within a finite number of such steps.

3 Models for Distributed Quantum Computing

Here we define the models for distributed quantum computing that are used by the research that we review later. A common assumption for all of them is that there are no faulty processors.

Section 4 and Section 5.1 use a standard distributed network with arbitrary topology and the added capability of individual processors to store, manipulate, and measure quantum states as well as classical states. The only means of communication between processors are classical channels that can transmit only classical information (bits). The restriction of only using classical communication while being able to do quantum processing locally is known as *Local Operations and Classical Communication (LOCC)*. The size of a message is assumed to be bounded by $O(\log n)$, and communication does not need to be synchronous.

An important requirement is that a sufficient amount of entanglement is created between processors at the time the network is set up. This is subsequently referred to as *pre-shared entanglement*. No a-priori knowledge is assumed. We refer to this model as *LOCC-ENTANGLE*.

Section 5.2 uses a model that has all of the features of *LOCC-ENTANGLE* with the only difference that here the network is anonymous. Anonymity means that individual processors are not guaranteed to have unique identities, so in a problem such as leader election, identity information cannot be used as a tie-breaker in order to guarantee correct solutions. This setting can be motivated by the situation in which a large number of generic sensors with no identities are parachuted from a plane and subsequently need to organize themselves for the purpose of conducting some computation. This model will be known as *LOCC-ENTANGLE-ANON*.

Section 5.3 is also in the anonymous setting, but here quantum channels are present in the sense that individual processors can send and receive quantum information (qubits) to and from other processors. Further, pre-shared entanglement is not assumed in Section 5.3. The algorithm presented there is formulated in a synchronous setting but can be modified to work in the asynchronously as well. The only a-priori knowledge that individual processors need is the total number of processors in the network or an upper bound of it. We refer to this model as *QCOMM-ANON*.

It is clear that the assumption of pre-shared entanglement is not a trivial one. Entanglement is the essential resource that allows quantum advantages to be materialized. Currently, building large-scale distributed entangled states has not been attempted, which is why it is not clear how efficiently they can be created and whether the costs of creating them are not larger than the advantages that they provide. Therefore, if quantum technology is ever to be used in a distributed manner in practice, future research must focus on providing pre-shared entanglement as a sufficiently efficient primitive. In subsequent sections we discuss different schemes for satisfying that assumption:

- Locally creating entangled pairs and exchanging qubits with neighbors if quantum channels are available (Section 4). Cost: one communicated qubit per shared entangled pair.
- Augmenting $n-1$ pre-shared entangled pairs to an n-partite GHZ state by using non-local CNOT gates in a binary-tree-like fashion (Section 4). Cost: $O(n)$ classical communication.
- Obtaining an n-partite GHZ state from $n-1$ pre-shared entangled pairs that are distributed in a spanning-tree-like fashion (Section 5.1). Cost: $O(n)$ classical communication.

The lack of complete knowledge about entanglement also causes its use in the anonymous settings of Section 5.2 to be questioned. In the motivating scenario of anonymity, the individual sensors do not have unique identities, because they have to be very cheap and consume very little power, which is why they cannot have any processing capabilities beyond what is absolutely necessary. Therefore, it is not clear whether it is fair to assume that such processors can be pre-entangled and given quantum-processing capabilities when they cannot even be afforded individual identities. Further, the total lack of failures cannot be practically guaranteed in such a scenario, which would then make the algorithms unusable.

4 Distributing Centralized Quantum Algorithms

Suppose there is an extremely useful centralized quantum algorithm, but only small quantum computers with just a few qubits each are available. If one is to do something useful with the algorithm, one has to find a way to distribute it over a collection of small quantum computers. The specific challenges related to that and how to overcome them are described in [14, 36].

To the best of our knowledge, Eisert et al. in [14] were the first to propose a solution to such a situation together with proving optimality of the resource and communication requirements of their solution. In general, the hurdles that one has to worry about in any porting of a centralized algorithm to a distributed setting are related to how to divide the problem into pieces and how to arrange for coordination between the individual computations. Local computations need to communicate with other parties at different times, which unavoidably incurs considerable communication overhead. Eisert et al. use the *LOCC-ENTANGLE* model.

The general strategy of distributing a centralized quantum algorithm is to take the quantum circuit that represents the centralized algorithm and draw horizontal lines that delineate the boundaries of each local computation. For example, one can distribute the CNOT gate over two different computers by having the control qubit at one computer and the target qubit at another (Fig. 3).

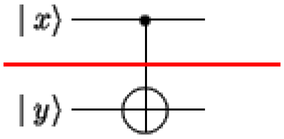

Figure 4: Distributed CNOT gate: The control and target are at two different computers.

Recall the universality of general single-qubit gates and the CNOT gate. It is clear that single-qubit gates do not induce any non-local interactions. Hence, the only gate that requires special treatment in the distributed context is the CNOT gate. Since all of the other multi-qubit gates that are of practical interest can be reduced to CNOT and single-qubit gates, the distributed CNOT gate is the necessary and sufficient primitive for building any distributed quantum circuit. Eisert et al. show a simple circuit (Fig. 4) for the distributed version of a CNOT gate and prove that one bit of classical communication in each direction and one previously shared entangled pair form a necessary and sufficient condition for a non-local implementation of the CNOT gate (assuming only LOCC). In fact, their circuit is a variation of quantum teleportation [24].

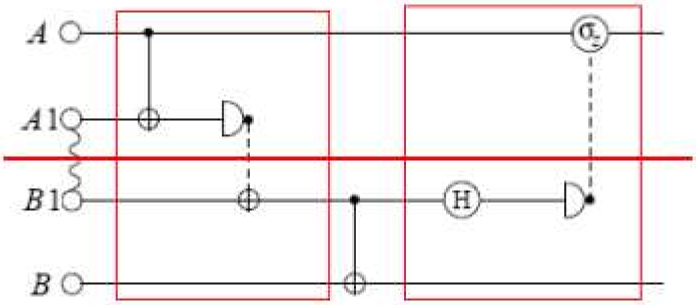

Figure 5: Circuit implementing the distributed CNOT gate. The rectangles delineate the two main parts of the circuit: "cat-entangler" and "cat-disentangler" as defined by Yimsiriwattana and Lomonaco Jr. [36, 35]. The horizontal line delineates the separation between the two computers.

In Fig. 4, qubits A and $A1$ are located at one party and qubits B and $B1$ are at another party. A is in some arbitrary quantum state and its purpose is to act as a control to the CNOT gate on B, where B is assumed to be in some other arbitrary quantum state. To achieve that, the two parties use a previously shared entangled pair ($A1$ and $B1$) to entangle A with $B1$, so that $B1$ can act as a local control qubit for the CNOT gate that is applied on B. This is done by applying a CNOT gate between A and $A1$, measuring $A1$, sending the measurement result as a classical bit to the other party (dashed line on the figure), and using it as a control to the CNOT gate at $B1$. At that point, A and $B1$ are entangled, from which it follows that $B1$ acting as control for the CNOT on B would be exactly as if A was the control. After that, a Hadamard gate and a measurement are applied on $B1$, after which the result is sent as a classical bit to A, where it is used to control the application of a Pauli Z gate. These last steps are performed in order to disentangle $B1$ and A, i.e. to completely restore A to the state in which it was at the beginning. In the process, the initially shared entanglement between the two parties is destroyed, and two classical bits are communicated in both directions. The paper by Eisert et al. gives more details about this circuit, including a step by step tracing of the intermediate states to show the desired result at the end. Since the CNOT gate is the only multi-qubit gate in the universal set, its distributed version is enough for implementing any quantum circuit in a distributed manner. Because the distribution concerns only the control qubit, the same technique works for any other controlled gate.

Eisert et al. in [14] also make the observation that generally, the required resources in terms of classical communication and entanglement are proportional to the number of distributed CNOT gates that are used, but as they point out, there may be remarkable exceptions. For example, when derived from the universal set, the swap gate requires three CNOT gates as shown in Fig. 5. This means that three entangled pairs and six classical bits of communication are the cost of implementing the swap gate by means of CNOT gates. On the other hand, it is rather intuitive that the swap gate's operation (simply exchanging the two input qubits) can be achieved by doing two teleportations, each of which requires only one entangled pair and the communication of two classical bits — a total of two entangled pairs and four classical bits, which is significantly cheaper than the first approach. The authors suggest that there may be other such cases that require fewer resources than what is required by the straightforward usage of the universal set.

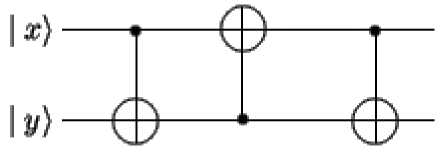

Figure 6: Swap gate implemented with three CNOT gates.

Yimsiriwattana and Lomonaco Jr. [36, 35] build on the work of Eisert et al. by distinguishing the two main parts of the distributed CNOT circuit, giving them the names "cat-entangler" (the first rectangle drawn in Fig. 4) and "cat-disentangler" (the second rectangle in Fig.4), and introducing the notion of "cat-like" state as the state that results from applying the cat-entangler on a general quantum state. The cat-like state is transformed back to the original quantum state after the cat-disentangler is applied on it. Their nomenclature is useful in terms of abstracting the basic parts of the circuit, so that they can be used as primitives in a simple manner later on, but the fundamental ideas were originated by Eisert et al.

Yimsiriwattana and Lomonaco Jr. also attempt to come to grips with the assumption of pre-shared entangled pairs between parties that share non-local gates. They propose several methods for creating the entangled pairs. One of them starts out by having each party locally create an entangled pair by using a Hadamard gate together with a CNOT gate (Fig. 2). After this, each party exchanges one of the qubits of its entangled pair with another party, and after a sufficient number of such exchanges, the global state is an n-partite GHZ state if n parties are involved. This approach requires the ability to physically transport the particles that carry the qubits — a quantum communication channel. However, the presence of quantum communication channels contradicts the assumption of *LOCC* that is used in Eisert et al.'s proof of necessity of entanglement.

The other possibility for satisfying the pre-shared entanglement assumption according to [36] is to use a Hadamard gate locally at one of the involved parties and a sequence of non-local CNOT gates that span the rest of the processors in a binary-tree-like fashion (Fig. 6). Note that this approach requires a pre-shared entangled pair for each of the non-local CNOT gates. Hence, even though the assumption of the pre-shared n-partite GHZ state is alleviated by this strategy, there is still the need to somehow prepare entangled pairs for the non-local CNOT gates that make this scheme work. In Fig. 6, we assume that the circuit starts in the state $|00000000\rangle$. After applying the Hadamard gate to the first qubit, the joint state at point 'a' is: $(|00000000\rangle + |10000000\rangle)/\sqrt{2}$. After applying a CNOT gate on qubit 5 with control qubit 1, the state at point 'b' is: $(|00000000\rangle + |10001000\rangle)/\sqrt{2}$, because qubit 5 gets inverted whenever qubit 1 (the control) is $|1\rangle$. Similarly, after applying CNOT gates on qubits 3 and 7 with controls 1 and 5, respectively, the state at point 'c' is: $(|00000000\rangle + |10101010\rangle)/\sqrt{2}$. Finally, CNOT gates are applied on qubits 2, 4, 6, and 8 with controls 1, 3, 5, and 7 respectively. The resulting final state is therefore $(|00000000\rangle + |11111111\rangle)/\sqrt{2}$. It can be easily seen that this scheme reduces the task of establishing an n-partite shared GHZ state to obtaining $n-1$ entangled pairs that are shared among parties in a binary-tree-like fashion. The time complexity is $\log n$ — the height of the binary tree — and the classical communication is $2(n-1)$ because each non-local CNOT gate communicates 2 bits.

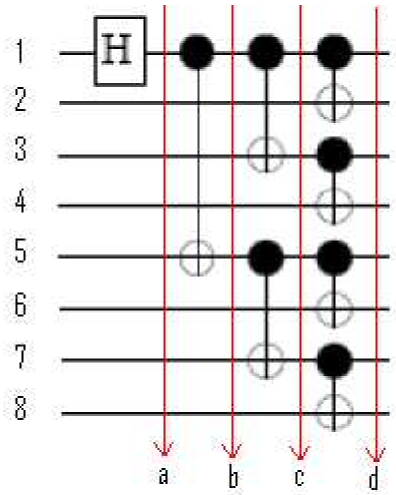

Figure 7: Creating an n-partite distributed GHZ state.

Finally, Yimsiriwattana and Lomonaco Jr. [36, 35] give two proof-of-concept examples (the quantum Fourier transform and Shor's algorithm [24, 28]) as direct applications of the distributing technique. They illustrate the straightforward observation that any centralized quantum algorithm with k gates can be distributed over m computers with a communication cost of $O(k/m)$.

5 Quantum Algorithms for Leader Election

5.1 Leader Election with Pre-shared Entanglement

In the leader election problem, each processor in a network participates in a computation that chooses one of the participating parties as the leader. It does not matter which party is chosen, as long as there is exactly one leader at the end, and all processors agree on the choice. The protocols with pre-shared entanglement circumvent the classical impossibility of leader election without communication by using entanglement instead of communication.

Pal et al. in [25] discuss the problem of leader election in the context of the *LOCC-ENTANGLE* model. Pal et al. assume $\log n$ n-partite pre-shared GHZ states, such that each of the n processors holds exactly one qubit from each of the $\log n$ entangled states. Consequently, each processor holds $\log n$ qubits, and the protocol consists of a single step — measure the local qubits. As a result, each processor holds a binary number that constitutes the address of the elected leader. In other words, each individual qubit of a processor is initially entangled in an n-partite GHZ state with the $n-1$ corresponding qubits at the other $n-1$ processors. Because of the form of the initial entanglement, quantum mechanics guarantees that the resulting binary number after measuring is the same at all processors, which is what is required for the protocol to be correct. Since the final outcome is determined by the party that is the quickest to measure its qubits, this scheme works in the asynchronous setting.

The measurements can be done asynchronously because whichever processor happens to measure its qubits first, the local measurement outcome instantaneously determines the measurement outcomes of the rest of the processors. This guarantees termination and uniqueness — the leader is elected with *no communication* at the cost of consuming $\log n$ n-partite GHZ states and performing $\log n$ measurements at each party. Additionally, since all measurement outcomes are completely random with no bias, all processors have equal chances. Thus, fairness is preserved as well.

This scheme would offer an extremely efficient way of solving the leader election problem. However, just as we noted in Section 4, the assumption of pre-shared entanglement is not a trivial one. Pal et al. point to [29], where they described a possible protocol to create the pre-shared n-partite GHZ states. According to the protocol, the creation of a single shared n-partite GHZ state requires $n-1$ EPR pairs of the form, $(|00\rangle + |11\rangle)/\sqrt{2}$ (the same as a 2-partite GHZ state). Additionally, the EPR pairs need to be a-priori distributed along the network in a specific way. If we interpret each EPR pair as supplying an "invisible" link between two processors, then the collection of the links supplied by the $n-1$ EPR pairs should form a spanning tree of the network. After that, Pal et al.'s protocol augments the entanglement provided by the $n-1$ EPR pairs to an n-partite GHZ state by using just *LOCC* at the cost of $O(n)$ communicated bits.

It is not clear whether the construction of $\log n$ n-partite GHZ states requires $\Omega(n \log n)$ communicated bits, given that a single n-partite GHZ state costs $O(n)$ bits of communication. Pal et al. do not raise this question, but in a way similar to the construction of the non-local swap gate that we discussed in Section 4, it may be possible to achieve some non-trivial savings when $\log n$ n-partite GHZ states are being constructed concurrently. This is an interesting question to be addressed in future research. Even so, we are left with another assumption — the presence of $n-1$ EPR pairs that form a spanning tree of the network. To the best of our knowledge there is no procedure to create the needed EPR pairs if we have the restriction of *LOCC* [2, 3]. On the other hand, if there is a quantum channel at hand, it is possible either to locally create an EPR pair (e.g. parametric down-conversion in photonic setups [23, 11, 34, 26]) and send one of the particles to a remote party or to entangle two spatially separated particles by making them interact with a third mediating particle [5, 22]. In short, the complexity of Pal et al.'s protocol is $O(n \log n)$ total classical communication in $O(n)$ rounds and $O(n \log n)$ quantum communication if the initial $n-1$ EPR pairs are created over quantum channels.

5.2 Anonymous Leader Election with Pre-shared Entanglement

Other research has focused on leader election in the *LOCC-ENTANGLE-ANON* model. Classically, the problem of anonymous leader election is known to be unsolvable because of the impossibility to simultaneously guarantee uniqueness and termination [21]. However, it appears that quantum mechanics can come to the rescue here. D'Hondt and Panangaden [7] prove that quantum entanglement is a necessary and sufficient resource for arriving at a correct solution, i.e., one that satisfies both uniqueness and termination. The specific kind of entanglement that is needed is an n-partite W state (see Section 2.3). Notice that each component of the superposition that makes up the W state has exactly one qubit as $|1\rangle$ and the rest of the qubits are $|0\rangle$. When the W state is destroyed after measuring all qubits, the resulting joint state is exactly one of the components of that superposition. If the W-state is initially prepared, so that each of the n qubits resides on a distinct processor, then after each processor measures its qubit, the one that gets 1 as a result becomes the leader. The W entanglement guarantees that the qubits of the rest of the processors are zero. Regardless of the specific moment when each processor does its measurement, as soon as one of them gets 1 as a measurement result, the superposition instantaneously collapses to the component where the qubit that is held by the lucky processor is $|1\rangle$ and the rest are $|0\rangle$.

Here again one faces the assumption that the entanglement needs to be taken care of before one starts solving the leader election problem. Even worse, each time an election is done, the entanglement is destroyed, so whatever efficient procedure there is to prepare it, that procedure must allow repeated usage in order to recreate/refresh the initial entangled condition of the network. There are no indications that the distributed n-partite W state is any easier to prepare than the corresponding GHZ state, so the research in this direction ends with the same problem as the previously considered cases — a practical implementation of this scheme needs to first have a way of preparing the n-partite W state. D'Hondt briefly considers this issue in [6]. She finds a quantum circuit to generate the 3-partite W state but finds it difficult to generalize to the n-partite case. She points to [23], where the experimental physics group of Mikami et al. offers a way to directly construct n-partite W states via a photonic setup.

5.3 Anonymous Leader Election Without Pre-shared Entanglement

Tani et al. in [31, 30, 32] assume the *QCOMM-ANON* model. They show that the general anonymous leader election problem has a correct quantum solution that can be achieved with certainty in polynomial time and communication without assuming any prior entanglement. Eliminating the dubious entanglement assumption that had to be made in Section 5.2 while still circumventing the classical impossibility for anonymous leader election is a very significant result. Tani et al. present several algorithms, whose common approach to solving the problem consists of gradual symmetry breaking by using quantum amplitude amplification, which is significantly different from the instant solution of D'Hondt and Panangaden. However, because no prior entanglement is assumed, these algorithms use quantum channels to create a number of n-partite shared entangled states that can be used to gradually break the symmetry in the network until a leader is chosen.

5.3.1 Tani et al.'s Algorithm

The complexities of the algorithm of Tani et al. [31, 30] that we describe here are $O(n^3)$ time and $O(n^4)$ quantum and classical communication. With n initially eligible parties, the algorithm proceeds in $n-1$ phases in each of which zero or more but not all parties become ineligible for election. We use l to denote the current number of eligible parties. Consequently, throughout the execution of a single phase, l decreases or stays the same but never increases or becomes zero. Each party i for $i = 0, \ldots, n-1$ has a number of quantum registers, initially in the state $|0\rangle$: $R0_i$, $R1_i$, S_i, $X0_i$, $X1_i$, $\ldots$, Xd_i, where d_i is the number of neighbors of i. Note that the network is anonymous and the identifier i is used only for notation purposes here. Also, each party i has classical registers k, z_i, and z_{max}. Register k is initialized to n and is decremented by 1 after a phase is completed. An important invariant to be clarified later is that $k \geq l$ for all phases. Registers z_i and z_{max} hold 2-bit numbers that are initialized to 0. The algorithm is divided in subroutines A, B, and C. The three subroutines, together with some local computations, form a single phase.

The execution starts by changing the state in $R0_i$ to $(|0\rangle + |1\rangle)/\sqrt{2}$ for all eligible parties and leaving it as initialized for all ineligible parties. Then Subroutine A is executed. This subroutine either creates a GHZ state over all eligible parties ("consistent state") or a state that guarantees the elimination of some of the eligible parties ("inconsistent state"). Consistent and inconsistent states are formally defined in [31, 30]. For the purposes of the algorithm, a consistent state, when measured, results in identical measurement results for all of the involved parties. On the other hand, with an inconsistent state, some parties have different measurement results from others. At the beginning of the subroutine, each party i locally entangles in a GHZ state (as shown in Section 2.3) the qubit from $R0_i$ with the qubits in $X0_i, \ldots, Xd_i$. Then i exchanges the contents of $X1_i, \ldots, Xd_i$ with its neighbors. Because all of the qubits in a GHZ state are equivalent to one another, it does not matter the qubit from which of the X registers is sent to which neighbor. Therefore, each party i is assumed to have chosen a random one-to-one mapping to map its $X1_i, \ldots, Xd_i$ registers to its ports before the algorithm begins. Then, the neighbor exchange is executed by sending each qubit from $X1_i, \ldots, Xd_i$ along the appropriate mapped port and placing the qubit that is received along that port in the appropriate register according to the mapping again. After the neighbor exchange, a simple local computation is done on the $X0_i, \ldots, Xd_i$ registers in order to determine consistency/inconsistency of the components of the state that is formed by them, and registers $X0_i$ and S_i are set to the outcome of this computation. It is assumed that $|0\rangle$ in register S_i means "consistent" and $|1\rangle$ means "inconsistent". Afterwards, by using $X0_i$ instead of $R0_i$ as the entangling register, Subroutine A repeats the described local entanglement, neighbor exchange, and local computation $n-1$ times in order for each i to obtain the consistency/inconsistency information about the components of the global state. The outcome of the execution of Subroutine A is an entangled state consisting of all registers S_i and $R0_i$:

$$|\psi\rangle = |S_0 \ldots S_{n-1}\rangle |R0_0 \ldots R0_{n-1}\rangle = \sum_{x \in \{0,1\}^n} |S_{0_x} \ldots S_{(n-1)_x}\rangle (|x\rangle + |\bar{x}\rangle)/\sqrt{2^{n+1}}, \tag{3}$$

where $\bar{x}$ represents the complement of the n-bit bit-string x. Each of the 2^n components of the global state defined by the $R0_i$ registers for all i is a superposition of a distinct bit-string of length n and its complement. With each such component is associated a consistency/inconsistency indication provided by the S_i registers. In each component of the superposition in the $|\psi\rangle$ state above (fixed x), the values of S_i are either all "consistent" or all "inconsistent" for all i. This determines the consistency/inconsistency of the associated component of the global state defined by the $R0_i$ registers. For example, if $n = 2$ and both parties are initially eligible, an execution of Subroutine A yields the state: $|\psi\rangle = |S_0, S_1\rangle |R0_0, R0_1\rangle = (|00\rangle(|00\rangle + |11\rangle) + |11\rangle(|01\rangle + |10\rangle) + |11\rangle(|10\rangle + |01\rangle) + |00\rangle(|11\rangle + |00\rangle))/2\sqrt{2}$.

The time complexity of Subroutine A is $O(n^2)$, because there are n rounds of communication — the neighbor exchanges done $n-1$ times. Each round takes $O(n)$ time, because Tani et al. assume that a message can only be sent to one neighbor at a time, and each party can have $O(n)$ neighbors. When the subroutine is executed in $n-1$ phases, the total time taken by it becomes $O(n^3)$. The quantum communication complexity of Subroutine A is $O(n^3)$, because each of the n parties does $n-1$ neighbor exchanges of $O(n)$ qubits, again because each party can have $O(n)$ neighbors. Execution in $n-1$ phases makes the total quantum communication $O(n^4)$.

After Subroutine A is executed, each party i measures its S_i register. This collapses the superposition in the $|\psi\rangle$ state above to one of its components. If the measurement outcome is "consistent", then the resulting global state defined by all $R0_i$'s has collapsed to a consistent state; otherwise, the global state is inconsistent. If the eligible parties find out in this way that they share a consistent state, they execute Subroutine B, which attempts to transform that state to a superposition of inconsistent states. Subroutine B does not always succeed, as will be explained below, but is guaranteed to have made possible the elimination of $n-1$ eligible parties by the end of the last phase. For technical reasons that we do not elaborate here, the transformation from a consistent to an inconsistent state is achieved by using an auxiliary qubit $R1_i$ at each processor i, so that each processor has two qubits in two quantum registers, $R0_i$ and $R1_i$. The purpose of Subroutine B is to transform the state consisting of $2l$ qubits (two qubits at each of the l eligible parties) to a state that has zero amplitudes for the superposition components that represent the possibilities of $|R0_i, R1_i\rangle$ simultaneously giving the same measurement results for all i. Observe that $|R0_i, R1_i\rangle$ can be $|00\rangle$, $|01\rangle$, $|10\rangle$, or $|11\rangle$ for any i. Therefore, in terms of the tensor product notation that was defined in Section 2.1, the components that need to be with zero amplitudes are exactly $|00\rangle^{\otimes k}$, $|01\rangle^{\otimes k}$, $|10\rangle^{\otimes k}$, and $|11\rangle^{\otimes k}$. These are the components that can cause identical measurement results everywhere. If they are not all with zero amplitudes, there is a non-zero probability that after the completion of a phase, no eligible party is excluded from the race.

To do its transformation, Subroutine B applies one of two possible gates — U and V in Fig. 7 — depending on whether the parameter $k = n - i$ in the i-th phase is odd or even. The reasons for using two different gates for odd and even phases as well as the definitions of these gates are entirely technical and are omitted from our discussion. The U and V gates are non-standard and are specified in [31, 30]. They can be derived from the more general concept of quantum amplitude amplification that was introduced in Section 2.4. Subroutine B is essentially an implementation of that technique in the sense that it obliterates the amplitudes of the undesirable components of the global state, i.e. the consistent components, and amplifies the amplitudes of the desirable ones, i.e. the inconsistent components. The circuits that simulate the cases $k = 3$ and $k = 2$ are shown in Fig. 7. Point 'a' of the circuit for $k = 3$ is the entry point of Subroutine B. Before that the 3-partite GHZ state consisting of the qubits in $R0_i$ for $i = 1, 2, 3$ is established, i.e. the three parties are sharing a consistent state. Between points 'a' and 'b', CNOT gates are applied on $R0_i$ and $R1_i$ at each of the three parties. As a result, at point 'b', the global state is $(|000000\rangle + |111111\rangle)/\sqrt{2}$. After that, each party applies the V gate on its qubits, which obliterates the amplitudes of the problematic consistent components, and the state is transformed into a large superposition of inconsistent states. The resulting superposition is too large to be given here, but the interested reader can easily implement the simulation using the first circuit from Fig. 7. The simulation software that we used can be obtained from [10]. For the case $k = 2$, point 'a' shows again the joint state at the subroutine entry: $(|0000\rangle + |1010\rangle)/\sqrt{2}$. Now each party applies the U gate and as a result, the consistent components are suppressed. The resulting state is a superposition of inconsistent states: $|R0_0, R1_0, R0_1, R1_1\rangle = -i(|0010\rangle + |1000\rangle)/\sqrt{2}$. It can be easily seen that both of the components of this state are inconsistent, i.e. when $R0_0$, $R1_0$, $R0_1$, and $R1_1$ are measured, the two parties are guaranteed to get different results.

A significant drawback of Subroutine B is that the U and V gates are parameterized over k, which is used as an upper bound for the number of eligible parties, l. Subroutine B successfully transforms a consistent state into an inconsistent superposition only when $k = l$. However, this algorithm does not operate with the exact value of l, because a significant amount of additional work would be required in order for each party to know the value of l for each phase. That work is circumvented here by just using the upper bound k, which gets gradually tightened in subsequent phases until it hits the actual value of l. At that point, Subroutine B is guaranteed to work, which makes it possible to decrease l by at least 1 and no more than $l - 1$. In the next phases, k continues to be an upper bound for l and the process of gradual tightening continues until $k = l$ again. At the conclusion of $n - 1$ phases, there is exactly one eligible party, which is the elected leader.

The elimination of eligible parties is attempted by Subroutine C, which succeeds whenever the global

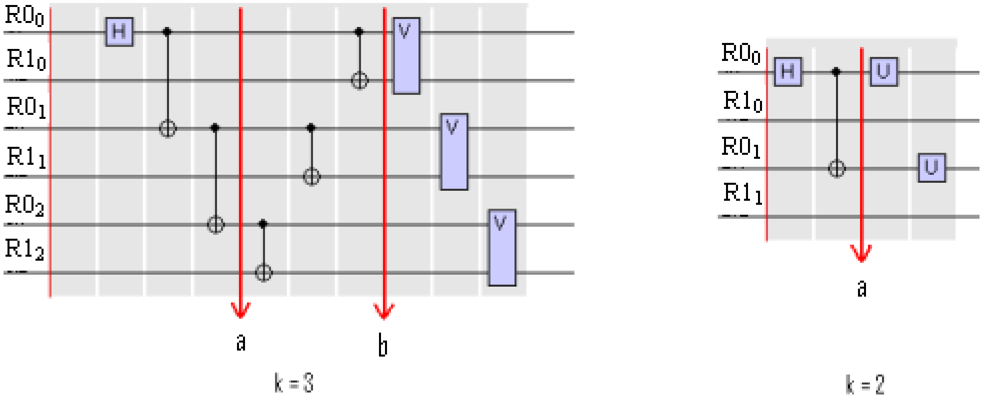

Figure 8: Quantum circuit simulations for the cases of $k = 3$ and $k = 2$ eligible parties.

state after Subroutine A has been collapsed to an inconsistent state, or whenever Subroutine B has succeeded in transforming a consistent state to an inconsistent one. Before Subroutine C, each of the eligible parties i measures its $R0_i$ and $R1_i$ to form a 2-bit number z_i. Only the parties with the highest resulting z_i remain eligible for the next phase. The elimination of the rest of the parties is done by having all eligible parties execute Subroutine C, which is just a simple classical computation of the maximum $z_{max} = max_{0 \leq i < n}(z_i)$ over the entire network. After Subroutine C is done, if party i's z_i is equal to z_{max}, i remains eligible; otherwise, it makes itself ineligible. Suppose that before the measurements of $R0_i$ and $R1_i$ for all eligible i, the global state formed by them was inconsistent. Then the measurements yield at least two distinct z_i values around the network, which forces some eligible parties to become ineligible after executing Subroutine C. However, in the case of a consistent state, all eligible processors i measure the same z_i values and nobody can be excluded after Subroutine C, because everybody's z_i equals z_{max}.

The time complexity of Subroutine C is $O(n^2)$ because z_{max} is computed in $n-1$ rounds, each of which takes $O(n)$ steps, because a message can be sent to only one neighbor at a time and each party can have $O(n)$ neighbors. When this is done in $n-1$ phases, the total time taken by this subroutine becomes $O(n^3)$. The classical communication complexity of Subroutine C is $O(n^3)$ because each of the n parties does $n-1$ neighbor exchanges of $O(n)$ bits, again because each party can have $O(n)$ neighbors. When this is done in $n-1$ phases, the total classical communication complexity adds up to $O(n^4)$.

It should be noted that it is not clear whether the U and V gates can be implemented with a constant number of gates from the universal set. Moreover, the matrix representations of the U and V gates contain irrational numbers, which makes it unclear whether they can be implemented exactly at all. Hence, when they are applied only with finite precision, the amplitudes of the consistent components are greatly suppressed by Subroutine B when $k = l$ but are still non-zero. For example, in the simulations that we did for the case $k = 3$, the amplitude for the state $|000000\rangle$ is $-2.7734^{-10} + 2.4813^{-10}i$ when the gate is applied with the precision of 9 significant figures. Apparently, some additional work is needed to eliminate such residues, because otherwise the correctness of the proposed algorithm is not guaranteed in practice. Further, Tani et al.'s algorithm has no tolerance for failures, because Subroutine B requires all eligible parties to execute it correctly in order for symmetry breaking to work with certainty. Another practical consideration concerns the complexities of Tani et al.'s algorithms — they are not competitive with the existing classical algorithms. Even though the classical algorithms in this setting are not guaranteed to be always correct, they are capable of electing a leader with high probabilities and much smaller time and communication costs. Tani et al. do not investigate the optimality of their algorithms, which leaves an open possibility for more efficient quantum algorithms.

5.3.2 Other Algorithms

The above algorithm is referred to as Algorithm I in Tani et al.'s original work. They also have Algorithm II [31, 30], which has the cost of $O(n^2 \log n)$ quantum communication and $O(n^6 \log^2 n)$ time and classical communication. The advantage over Algorithm I is that the quantum communication is lower, because the number of phases is $\log n$ instead of n. However, that is achieved at the price of much higher time and classical communication costs. Tani et al. in [32] give two more algorithms for solving anonymous leader election. The first has the same complexities as Algorithm I but explicitly uses the technique of exact quantum amplitude amplification that was introduced in Section 2.4. The second is for the special case when n is a power of 2. It has only a linear number of rounds, but its drawback is that it has $O(n^6 \log n)$ quantum communication.

6 Open Problems

Notably, quantum computing promises significant advantages over classical computing in some cases. Nevertheless, there are still important issues that have not been addressed fully in the research that we reviewed. Undoubtedly, the whole area of quantum computing can benefit greatly from new work that investigates them:

- Pre-shared multi-partite GHZ and W entanglement: These appear as assumed resources everywhere throughout Section 4, Section 5.1, and Section 5.2. To the best of our knowledge, there is no research work that gives a complete answer regarding the way in which such assumptions can be satisfied in a practical setting. The question whether this can be done efficiently must be answered if the distributed quantum computing schemes are ever to become practical. It is also possible that the cost of preparing multi-partite shared entanglement outweighs the advantage given by it. Preparation of the n-partite GHZ state seems to be easier than the corresponding W state. Indeed, with the presence of quantum communication links between all processors in the network, it is feasible to create the GHZ state by locally entangling and exchanging qubits. Locally entangled pairs can be produced on-demand by the parametric down-conversion techniques that were mentioned in Section 1. More experimental work on creating entanglement and manipulating it would be extremely beneficial.

- Resource savings in non-local circuits similar to the swap gate case: As discussed in Section 4, complex circuits that use large numbers of non-local gates may be able to do better than use one entanglement pair and two bits of classical communication per non-local gate.

- Generating $\log n$ n-partite GHZ states: The version of leader election that was discussed in Section 5.1 is solved when there are $\log n$ n-partite pre-shared GHZ states. Similarly to the possibility that was noted in Section 5.1, the cost per state of preparing the $\log n$ states together could be lower than the cost of preparing a single such state by itself. In particular, [29] describes how to prepare a single n-partite GHZ state with $O(n)$ communication cost. Is it possible to prepare $\log n$ such states at a cost lower than $O(n \log n)$?

- Fault-tolerant leader election: The algorithms with pre-shared entanglement do not work in the presence of faults, because even if a unique leader is agreed upon, that leader may be faulty. The algorithm by Tani et al. that was discussed in Section 5.3 also cannot tolerate faults, because the essential symmetry breaking via quantum amplitude amplification in Subroutine B works only if all eligible parties execute it correctly. The importance of fault-tolerance in practice motivates the search for fault-tolerant quantum leader election.

- Other applications: It would be beneficial to see how the discussed methods and results can be applied to versions of the leader election problem in settings other than the ones that have been considered thus far. Also, it is quite likely that other important distributed problems can benefit in similar ways.

7 Summary

We have reviewed the present research regarding the two aspects in which quantum computing can benefit from and contribute to distributed computing. First, since there is a perceived practical difficulty of scaling up existing quantum computing implementations, it could be possible to solve large problems by using a number of small quantum computers together. Second, important problems from classical distributed computing such as leader election can potentially benefit from using quantum resources. We presented trivial protocols that solve leader election with no communication but with the assumption that previously shared entanglement is in place. We also considered the anonymous leader election problem, where the best classical algorithms elect a unique leader in finite time with high probability, but the quantum algorithms solve the problem with certainty. Even though the research that we have reviewed has offered impressive solutions, it has done so with the non-trivial assumption of the ability to create and maintain shared entangled states. Not only does this assumption remain to be shown satisfiable, but also the additional cost incurred by satisfying it remains to be evaluated in order to see whether it does not outweigh the advantages of quantum processing. The only work that does not assume any prior entanglement is the work of Tani et al. that was discussed in Section 5.3. However, their algorithms have serious practical drawbacks as well. It is currently unknown whether the essential unitary transformations from Subroutine B that circumvent the classical impossibility result can be implemented exactly and how much overhead their implementations would incur. Also, the correct solution of the anonymous leader election problem that is achieved with certainty by the quantum algorithm comes at the price of higher time and communication complexities, when compared with the classical randomized algorithms that achieve a correct solution only with high probability.

References

[1] Bell, J. S. On the Einstein-Podolsky-Rosen paradox. *Physics 1* (1964), 195–200.

[2] Bennett, C. H., DiVincenzo, D. P., Smolin, J. A., and Wootters, W. K. Mixed-state entanglement and quantum error correction. *Physical Review A 54*, 5 (Nov 1996), 3824–3851.

[3] Bennett, C. H., Popescu, S., Rohrlich, D., Smolin, J. A., and Thapliyal, A. V. Exact and asymptotic measures of multipartite pure-state entanglement. *Physical Review A 63*, 1 (Dec 2000), 012307.

[4] Brassard, G., Hoyer, P., Mosca, M., and Tapp, A. Quantum amplitude amplification and estimation, 2000, quant-ph/0005055.

[5] Cubitt, T. S., Verstraete, F., Dür, W., and Cirac, J. I. Separable states can be used to distribute entanglement. *Physical Review Letters 91*, 3 (Jul 2003), 037902.

[6] D'Hondt, E. *Distributed quantum computation: a measurement-based approach*. PhD thesis, Vrije Universiteit Brussel, 2005.

[7] D'Hondt, E., and Panangaden, P. The Computational Power of the W and GHZ states. *Quantum Information and Computation 6*, 2 (2006), 173–183.

[8] Dur, W., Vidal, G., and Cirac, J. I. Three qubits can be entangled in two inequivalent ways. *Physical Review A 62* (2000), 062314.

[9] Eberhard, P. H. Bell's theorem and the different concepts of locality. *Nuovo Cimento 46B* (1978), 392–419. LBL-7537.

[10] Eck, M., Wocjan, P., and Beth, T. Design and implementation of a simulator for quantum circuits. Tech. rep., University of Karlsruhe, Department of Informatics, Institute for Algorithms and Cognitive Systems, 2000.

[11] Eibl, M., Gaertner, S., Bourennane, M., Kurtsiefer, C., Żukowski, M., and Weinfurter, H. Experimental observation of four-photon entanglement from parametric down-conversion. *Phys. Rev. Lett. 90*, 20 (May 2003), 200403.

[12] Einstein, A. *Relativity: the special and general theory*. Methuen and Co Ltd, 1916.

[13] Einstein, A., Podolsky, B., and Rosen, N. Can quantum-mechanical description of physical reality be considered complete? *Phys. Rev. 47*, 10 (May 1935), 777–780.

[14] Eisert, J., Jacobs, K., Papadopoulos, P., and Plenio, M. B. Optimal local implementation of non-local quantum gates. *Physical Review A 62* (2000), 052317.

[15] Ekert, A., and Jozsa, R. Quantum algorithms: Entanglement enhanced information processing. *ArXiv Quantum Physics e-prints* (1998), quant-ph/9803072.

[16] Greenberger, D. M., Horne, M. A., and Zeilinger, A. *Bell's theorem, Quantum Theory, and Conceptions of the Universe*. Kluwer Academic Press, 1989, ch. Going beyond Bell's theorem, pp. 73–76.

[17] Grover, L. K. A fast quantum mechanical algorithm for database search. In *STOC* (1996), pp. 212–219.

[18] Jozsa, R. Entanglement and quantum computation. In *The Geometric Universe*, S. Huggett, L. Mason, K. P. Tod, S. T. Tsou, and N. Woodhouse, Eds. Oxford University Press, January 1998.

[19] Jozsa, R., and Linden, N. On the role of entanglement in quantum-computational speed-up. *Royal Society of London Proceedings Series A 459* (Aug. 2003), 2011–2032.

[20] Linden, N., and Popescu, S. Good dynamics versus bad kinematics. Is entanglement needed for quantum computation? *Physical Review Letters 87* (2001), 047901.

[21] Lynch, N. *Distributed algorithms*. Morgan Kaufmann Publishers, 1997.

[22] Meter III, R. D. V. *Architecture of a Quantum Multicomputer Optimized for Shor's Factoring Algorithm*. PhD thesis, Keio University, 2006, quant-ph/0607065.

[23] Mikami, H., Li, Y., and Kobayashi, T. Generation of the four-photon W state and other multiphoton entangled states using parametric down-conversion. *Physical Review A 70*, 5 (Nov. 2004), 052308.

[24] Nielsen, M., and Chuang, I. *Quantum Computation and Quantum Informaiton*. Cambridge University Press, 2000.

[25] Pal, S. P., Singh, S. K., and Kumar, S. Multi-partite Quantum Entanglement versus Randomization: Fair and Unbiased Leader Election in Networks. *ArXiv Quantum Physics e-prints* (June 2003), quant-ph/0306195.

[26] Pan, J.-W., Bouwmeester, D., Weinfurter, H., and Zeilinger, A. Experimental entanglement swapping: Entangling photons that never interacted. *Phys. Rev. Lett. 80*, 18 (May 1998), 3891–3894.

[27] Shimony, A. Controllable and uncontrollable non-locality. In *Proceedings of the International Symposium on Foundations of Quantum Mechanics* (1984), S. Kamefuchi, Ed., The Physical Society of Japan.

[28] Shor, P. W. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *Society for Industrial and Applied Mathematics Journal on Computing 26*, 5 (1997), 1484–1509.

[29] Singh, S. K., Kumar, S., and Pal, S. P. Characterizing the combinatorics of distributed EPR pairs for multi-partite entanglement. *ArXiv Quantum Physics e-prints* (June 2003), quant-ph/0306049.

[30] Tani, S. *A study on algorithms for efficient multi-point communication in autonomous distributed networks*. PhD thesis, University of Tokyo, 2006.

[31] Tani, S., Kobayashi, H., and Matsumoto, K. Exact quantum algorithms for the leader election problem. In *Symposium on Theoretical Aspects of Computer Science Proceedings* (2005), V. Diekert and B. Durand, Eds., vol. 3404 of *Lecture Notes in Computer Science*, Springer, pp. 581–592.

[32] Tani, S., Kobayashi, H., and Matsumoto, K. Quantum leader election via exact amplitude amplification. In *Proceedings of ERATO conference on Quantum Information Science* (2005), pp. 11–12.

[33] Tel, G. *Introduction to distributed algorithms*. Cambridge University Press, 2000.

[34] Weinfurter, H., and Żukowski, M. Four-photon entanglement from down-conversion. *Phys. Rev. A 64*, 1 (Jun 2001), 010102.

[35] Yimsiriwattana, A., and Lomonaco, Jr., S. J. Distributed quantum computing: a distributed Shor algorithm. In *Quantum Information and Computation II* (Aug. 2004), E. Donkor, A. R. Pirich, and H. E. Brandt, Eds., vol. 5436 of *Proceedings of Society of Photographic Instrumentation Engineers Conference*, pp. 360–372.

[36] Yimsiriwattana, A., and Lomonaco, Jr., S. J. Generalized GHZ States and Distributed Quantum Computing. *ArXiv Quantum Physics e-prints* (Feb. 2004), quant-ph/0402148.