Leader Election and Distributed Consensus with Quantum Resources

Ellie D'Hondt and Prakash Panangaden

January 13, 2005

Abstract

In this paper we study paradigmatic tasks from classical distributed computing – leader election and distributed consensus – in the presence of quantum resources. Our main contribution is a demonstration of the special computational power of the W-state, and also of the GHZ-state. We find that totally correct leader election *is* possible in anonymous quantum networks, which is in stark contrast with the classical situation. Next, we prove that the specific entanglement provided by the W- and GHZ-states, and their generalizations, is the *only* kind that exactly solves leader election and distributed consensus respectively. At the heart of the proofs of these impossibility results lie symmetry arguments.

1 Introduction

The use of quantum resources in computational tasks has led to a revolution in algorithms [NC00]. Models of computation have been developed which serve as the basis for the design of quantum algorithms. In the present paper we study paradigmatic tasks – leader election and distributed consensus – from classical distributed computing [Lyn96, Tel94] in the presence of quantum resources. Traditional algorithms are typically intended to establish an input-output correspondence; the main interest in quantum algorithms is in the use of quantum resources to reduce the time complexity of such algorithms. By contrast, our problems are about *joint decision making* by a group of autonomous agents. It is much closer in spirit to communication protocols.

As with classical distributed algorithms a whole new arena for establishing impossibility results becomes available. The notion of "universality" commonly used in algorithms is not relevant here because that notion assumes that one can entangle any two (or more) qubits. In a distributed system one has resources in separated locations and one's actions are limited to what one can do in a particular region. One cannot just demand that a particular global sequence of operations be carried out; it is necessary to arrange coordination – usually involving communication – between the agents acting at the separate locations (regions). Instead one has to ask what can be achieved within the limits of a particular computational model with communication and other limitations built in. In our case we study the quantum resources needed to carry out leader election and distributed consensus in *anonymous* networks. We focus on achieving these tasks exactly (to be defined precisely below) rather than probabilistically. The ultimate goal is understanding how information flows between different agents in a quantum setting.

Leader election algorithms have been studied extensively in numerous papers and for many different network settings. Distributed consensus is usually studied in a fault-tolerant context [FLP85], a matter upon which we do not touch in this paper. In networks where processors have unique identifiers the symmetry is inherently broken via processor names [LeL77]. By contrast, in an anonymous situation purely deterministic leader election is impossible: there is nothing that can break the symmetry if all the processors do the same thing [Ang80]. If each process has a coin then they can elect a leader: for example they can each toss a coin and if they get a head they are the leader. Of course this is not guaranteed to work, there may be more than one leader or no leaders and the process will have to be repeated in the next round. This idea was first put forward in [Ang80] and later generalized in [IR90]. With probability one this will terminate eventually but termination is not guaranteed. There is no bound on how long this process will take, though the expected number of rounds is just 2. When two processors share a singlet state they can just measure it, such that the one who gets, say, 1 is the leader: this terminates in one step and always succeeds. How does this generalize to more than two processors? Is such a protocol possible within our framework?

Our main result is that only very special shared quantum resources can be used to achieve the tasks at hand. More precisely, we show that in an anonymous network if the processors share the so-called W-state then a trivial protocol allows them to solve the leader election problem and, more importantly, *this is the only possible shared resource that allows this problem to be solved at all*. Similarly the GHZ-state is the only shared quantum resource that allows solution of distributed consensus. These are then essentially impossibility results: without the W-state leader election is not solvable in the sense that we make precise below. The proof uses the concept of *symmetric configurations*, and as such is similar to the work of Angluin. She writes that for anonymous networks " [...] it seems intuitively that the behaviour of the network can only depend upon the 'local appearance' of the underlying graph". Our results are in the same vein, but with the symmetry breaking being dictated by the (non-local) properties of the shared entangled quantum state. The details of the proof are quite different, however.

The structure of this paper is as follows. In Sec. 2 we define the particular framework in which our results are situated. Sec. 3 formalises symmetry concepts within this setting, which are then used in the impossibility proofs of Secs. 4 and 5 concerning leader election and distributed consensus respectively. We conclude in Sec. 6.

2 The model

A distributed system is a system in which several inter-communicating parties carry out computations concurrently. While this definition is quite general, there are a whole series of specific implementations, depending on the communication structure, the degree of synchronization and the computational capabilities of individual processors. The setting of this paper is that of synchronous, anonymous, distributed quantum networks with broadcasts, where the network size is known. Let us clarify these terms. First of all, while each processor is allowed a quantum state as well as a classical state, the communication between processors remains classical[1]. Classical and quantum states are assumed to be finite – in fact, throughout this paper we work with qubit states. Communication occurs via *broadcasts*, which means that all messages are sent out publicly to all processors along classical channels. Processors as well as communication are assumed to be non-faulty. So far this framework is very similar to what is known in as LOCC, for *local operations and classical communication*.

In the anonymous setting, all processors are completely identical, that is they do not carry individual names with which they can be identified. As such the initial network specification must be invariant under permutations of processors. One of the implications thereof is that processors start out in identical local classical states. However, one has to carefully restate this when quantum states are allowed. Indeed, due to the phenomenon of quantum entanglement, the network quantum state is in general not a product state of individual (i.e. local) processor states. The hallmark example is that of a network of two processors, call them A and B, such that each own one qubit of the following *Bell state* $|0\rangle_A|0\rangle_B + |1\rangle_A|1\rangle_B$, where the indices were added to clarify which qubit belongs to which processor. A simple calculation shows that it is not possible to write this state as a product of a quantum state owned by A and one owned by B. In this situation, the most sensible definition of anonymity is to demand that the network quantum state is invariant under permutations of processor subspaces. This has as an immediate consequence that each processor has the same local view on their quantum state, which in quantum mechanics is formally expressed by the reduced density matrix. However, we use the slightly stronger assumption of invariance with respect to subspace swaps, thus avoiding initial network states as for example $|0\rangle_A|0\rangle_B + e^{i\theta}|1\rangle_A|1\rangle_B$. Hence, we propose the following definition.

Definition 2.1 *An anonymous distributed quantum network has the property that each processor* i*, where* $i \in \{1, \cdots, n\}$ *executes the same local protocol and has identical initial classical state, and furthermore that the initial network quantum state* $|\psi\rangle \in \mathcal{H}_1 \otimes \ldots \otimes \mathcal{H}_n$ *is invariant under any permutation of the processor subspaces* $\mathcal{H}_1, \ldots, \mathcal{H}_n$*.*

Note that anonymity implies that all $\mathcal{H}_i$ are identical. We denote the closure over all permutations of a state $|\psi\rangle$ by $\mathrm{Perm}|\psi\rangle$; for example $\mathrm{Perm}|001\rangle = |001\rangle + |010\rangle + |100\rangle$.

Finally, the system is synchronous, which means that a protocol proceeds in a sequence of *rounds*. During each round a processor receives messages that were sent to it in the previous round, performs a local operation, and then broadcasts messages. The local operation, which in general depends on received messages, consists of quantum operations followed by a classical post-processing stage, such that, because of anonymity, each processor has identical decision trees.

We investigate distributed protocols within this framework, where we make one further distinction.

Definition 2.2 *A* totally correct *distributed protocol is a protocol that is* terminating, *i.e. it reaches a terminal configuration in each computation, and* partially correct, *i.e. for each of the reachable terminal configurations the goal of the protocol is achieved.*

[1]Initial qubits are distributed at the time the system is set up.

Note that the above definition does not exclude the possibility of non-deterministic processes. In this paper we study totally correct leader election and distributed consensus protocols. For a *leader election* protocol, a correct terminating configuration is one in which there is exactly one leader in the state *leader*, while all other processes are in the state *follower*. In the case of *distributed consensus*, all processors should terminate with an identical bit value, which can be 0 or 1 with equal probability. Probabilistic algorithms violate either termination or partial correctness. Probabilistic leader election algorithms are generally Las Vegas algorithms, that is, they terminate with positive probability and are partially correct, i.e. there are no reachable forbidden configurations. In the current literature, probabilistic algorithms come about by equipping each processor with a randomization tool, i.e. an random number generator or electronic coin. In our framework however, we allow only quantum operations and deterministic classical operations. Any qubit can of course be used to implement a coin toss.

These definitions lead to the following propositions.

Proposition 2.3 *No totally correct leader election protocol exists without prior shared entanglement.*

Proof *Outline.* Without entanglement the network's quantum state is in a product state, and no entanglement can be created through LOCC operations [BPR$^+$01]. Therefore each processor is essentially equipped with a coin throughout the protocol. As a result, either termination or partial correctness is compromised, since there is in a nonzero probability that local measurement results are identical. ∎

Proposition 2.4 *Totally correct leader election algorithms for anonymous quantum networks are* fair, *i.e. each processor has equal probability of being elected leader.*

Proof Recall that the processor state is finite, i.e. each processor can be in one of finitely many classical states and has only a finite number of qubits. Moreover, with qubits branching due to quantum measurements occurs in a finite way as well. Hence there are only finitely many computations, so it suffices to count these. Suppose then that by a terminating and correct computation configuration C is reached, wherein, say, process A is elected leader. Then there exists a corresponding configuration C' in which process B is elected leader, by defining a permutation of processors σ such that $\sigma(A) = B$ and running C' on the permuted set of processors. C' is necessarily correct because of the anonymity of the network. This reasoning holds for arbitrary processors A and B in the network, hence there are equally many terminal configurations that elect each of the processors as leader respectively, and our result follows. ∎

The protocols described below are presented with respect to the computational basis. Specifically, quantum measurements occur within this same basis – this is in fact quite general since measurements in any other basis can be brought back to these by first applying the appropriate unitary transform in an anonymous way, i.e. locally and identical for all processors. We frequently denote basis states in their integer representation if the number of qubits is clear, for example $|2\rangle = |010\rangle$ for a 3-qubit state. Normalization factors are suppressed throughout this paper because the crux of the argument will depend upon the symmetry rather than on the actual probability amplitudes. Furthermore, we only address symmetry-breaking capabilities of the quantum parts of our protocol, since it is known from previous results that classical protocols, and thus classical post-processing, cannot break the symmetry in anonymous networks.

3 Symmetric moves

We first need to set up some machinery. Specifically, we prove that certain types of superposition terms are possibly present throughout a computation. This corresponds to a branch of the computation in which a group of processors evolve symmetrically.

Definition 3.1 *Suppose an n-partite state $|\psi\rangle \in \mathcal{H}^{\otimes n}$, where $\mathcal{H}$ is a 2^m-dimensional Hilbert space, is distributed over n processors. We say that there exists a k-*symmetric move *for the processors $i_1, \ldots, i_k$ with respect to $|\psi\rangle$, where $0 < k \leq n$, if for all observables $M = \sum_{j=1}^{J} \lambda_j P_j$, with $J \leq 2^m$ and all P_j projectors, we have that*

$$\exists l \in \{1, \ldots, J\} : (P_l)^{\otimes k}_{i_1,\ldots,i_k} (P_{j_{k+1} \neq l})_{i_{k+1}} \ldots (P_{j_n \neq l})_{i_n} |\psi\rangle \neq 0 \tag{1}$$

The idea is that all measurements potentially give identical measurement results for k out of the n processors. Because anonymous networks are invariant under permutations we need not specify any particular subset of processors.

Lemma 3.2 *There exists a k-symmetric move for an anonymous quantum state $|\psi\rangle \in \mathcal{H}^{\otimes n}$ if and only if $|\psi\rangle = \mathrm{Perm}|\omega\rangle + |\omega'\rangle$, where $|\omega\rangle = |\eta\rangle^{\otimes k} \bigotimes_{j=k+1}^{n} |\eta_j\rangle$, $|\eta\rangle$ and $|\eta_j\rangle \in \mathcal{H}$, $|\omega\rangle$ and $|\omega'\rangle \in \mathcal{H}^{\otimes n}$, and $\langle \eta_j | \eta \rangle = 0$ for all j.*

This is trivial in one direction and an easy calculation in the other. Only when $\mathcal{H} = \mathbf{C}^2$ k-symmetric moves imply $(n-k)$-symmetric moves; for example, 1- and $(n-1)$-symmetric moves exist for all network states containing a term of the form $|1\rangle|0\rangle^{\otimes(n-1)}$. Note that for any state $|\psi\rangle$ some incomplete measurements, for which $J < 2^m$, always result in more than k processors obtaining the same results (the trivial example being the identity projector), while for complete measurements $J = 2^m$ there is always a branch leading to exactly k of the processors measuring identical values. The point of Def. 3.1 is of course that it is a statement about *all* possible measurements. A particular execution of a distributed protocol follows a *k-symmetric branch* if it occurs by a sequence of k-symmetric moves. We have the following result.

Proposition 3.3 *Any anonymous distributed quantum protocol that can follow a k-symmetric move initially has a k-symmetric branch.*

Proof Without loss of generality, we assume that local quantum operations during one round consist of an isometric operation $\mathbf{U}$, i.e. a unitary operation along with creation of ancillae, followed by a measurement M. If Def. 3.1 holds for the initial network state $|\psi\rangle$, then it must hold for $\mathbf{U}^{\otimes n}|\psi\rangle$; indeed, after such an operation we again obtain a network state as in Lemma 3.2. Suppose that for the subsequent measurement M the protocol follows the existing k-symmetric move, corresponding to identical measurement results j and projections on $|\phi_j\rangle$. Knowing that classical post-processing cannot break symmetry in anonymous networks, in this case identical measurement results are broadcast, such that the local operations applied in the next round, depending on these results, are identical. Moreover, at this point the network state still allows a k-symmetric move, since it contains a term of the form $|\phi_j\rangle^{\otimes k}|\phi_j\rangle^{\perp}$. Then by induction we can construct a k-symmetric branch for the entire protocol. ∎

4 Quantum leader election

Theorem 4.1 *A necessary and sufficient condition for a totally correct anonymous quantum leader election (QLE) protocol, where each processor owns 1 qubit initially, is that processor qubits are entangled in a W-state.*

We prove this theorem in both directions in the following two subsections.

4.1 W is sufficient

The idea is to share a specific entangled state between all parties, which allows to break the symmetry in one step. The state used is known as the W-state, where

$$W_n = \sum_{j=1}^{n} |2^j\rangle \tag{2}$$

For example $W_3 = |001\rangle + |010\rangle + |100\rangle$. This state can be used as a symmetry-breaking quantum resource. Each processor i carries out the protocol below.

1. $q \leftarrow i$th qubit of W_n
   b=0
   result=wait
2. **b**:= measure q
3. if **b** $= 1$ then **result**:= leader, else **result**:=follower.

Protocol 1: The QLE protocol.

This is a totally correct protocol with time complexity $\mathcal{O}(1)$; no message passing is required[2]. Note that Protocol 4.1 works also for different communication graphs or in asynchronous networks.

4.2 W is necessary

For this part of the proof we use the tools from Sec.3. Specifically, we prove that for certain types of superposition terms in the initial quantum state total correctness is compromised.

Proof Any protocol for which k-symmetric branches exist with k different from 1 or $n-1$ is not totally correct. Indeed, for such a branch the protocol either does not terminate or it terminates in a forbidden configuration. Hence by Prop. 3.3, the initial network state $|\psi\rangle$ cannot allow k-symmetric moves for $k \notin \{1, n-1\}$. By Lemma 3.2 with $|\psi\rangle \in (\mathbf{C}^2)^{\otimes n}$ anonymous in the sense of Def. 2.1, this leaves us with W_n, or unitary transforms thereof, as the only possibility. ∎

4.3 More qubits per processor

One can repeat the same symmetry argument in the case where each processor has m qubits, i.e. $\mathcal{H} = \mathbf{C}^{2m}$. Taking $\{|\phi_i\rangle\}$ to be a basis for $\mathcal{H}$, this leads to initial network states of the form

$$|\psi\rangle = \sum_{i=1}^{2^l} \alpha_i \mathrm{Perm}|\phi_i\rangle \bigotimes_{j=2}^{n} |\phi_{i_j}\rangle \tag{3}$$

Here we need one extra ingredient: as before, each processor knows beforehand which measurement results lead to them becoming a leader. However, since up to m different results are possible the situation is slightly more complicated. So suppose $\mathcal{L}$ is spanned by the leader labels $\{|\phi_i\rangle, i = 1, \ldots, 2^l\}$ and $\mathcal{F}$ is spanned by the follower labels $\{|\phi_{i_j}\rangle\}$. Then $|\psi\rangle$ cannot allow $(k > 1)$-symmetrical moves with respect to $\mathcal{L}$, which means concretely that none of the $|\phi_i\rangle$ can appear in the tensor product in Eq. (3), or in other words we take $\mathcal{H} = \mathcal{L} \oplus \mathcal{F}$. Note that Eq. (3) includes the more stringent dual situation where $n-1$ processors are symmetric w.r.t. $\mathcal{F}$. A QLE protocol would then succeed by measuring $|\psi\rangle$, such that the processor obtaining a result in $\mathcal{L}$ appoints itself leader, while those obtaining results in $\mathcal{F}$ become followers. As a result, we obtain a family of W-like states as the only possible entanglement resources for totally correct anonymous QLE protocols.

4.4 An instructive example

The following example shows the drastic impact of anonymity on the success of a protocol. Suppose we have an odd number of processors n, and each processor carries out Protocol 4.1 given below[3]. Because n is odd one of the candidates always gets more votes. However, in order to be

[2] We could also have used the state $\overline{W}_n$, the complement of W_n, to carry out the protocol; in this case the processor measuring $|0\rangle$ becomes the leader.

[3] H is the Hadamard transform, defined by $H|0\rangle = |0\rangle + |1\rangle$ and $H|1\rangle = |0\rangle - |1\rangle$.

able to appoint a leader, either the voters have to be able to name the candidate they voted for, or the candidates must differ in that they know which votes are intended for them. Both possibilities violate anonymity. Note however, that Protocol 4.4 would work when adapted for a network where the communication graph is a ring and the processors have a sense of direction. Indeed, suppose each candidate sends a message in say, the clockwise direction, such that the first candidate receiving this message proclaims itself the leader. This works because when n is odd both messages are ensured to arrive in different rounds. Both time and message complexity are in this case $\mathcal{O}(n)$.

1. $q \leftarrow i$th qubit of $W_{2,n-2} = \mathrm{Perm}(|1\rangle^{\otimes 2}|0\rangle^{\otimes(n-2)})$
   b=0
   result=wait
2. **b**:=measure q
3. if **b** $= 1$ then **result**:= candidate, else **result**:=voter
4. if **result**:=voter then $\{$**b**=measure H(q), broadcast **b**$\}$

Protocol 2: Attempting odd-party leader election.

5 Quantum distributed consensus

The results here are similar in spirit to the results for leader election: they also depend on a symmetry property, this time on symmetry preservation rather than symmetry breaking.

Theorem 5.1 *A necessary and sufficient condition for a totally correct anonymous quantum distributed consensus (QDC) protocol, where each processor owns 1 qubit initially, is that processor qubits are entangled in a GHZ-state.*

We prove this theorem in both directions in the following two sections.

5.1 GHZ is sufficient

The trick is to share a specific entangled state between all parties, which allows to create symmetrical knowledge in one step. The state used is known in the quantum computation community as the GHZ-state, where

$$GHZ_n = |0\rangle^{\otimes n} + |1\rangle^{\otimes n} \tag{4}$$

This state can be used as a symmetry-creating quantum resource. Each processor i carries out the following protocol below.

1. $q \leftarrow i$th qubit of GHZ_n
   result=wait
2. **result**:= measure q

Protocol 3: The QDC protocol.

This is a totally correct protocol with time complexity $\mathcal{O}(1)$; no message passing is required. Note that Protocol 5.1 works as well for different communication graphs or in asynchronous networks.

5.2 GHZ is necessary

The proof is analogous to that of the previous section.

Proof For a QDC protocol to be totally correct it can have only n-symmetric branches. Indeed, any k-symmetric branch with $k < n$ results in a non-zero probability that only k processors obtain symmetrical knowledge. Thus neither partial correctness nor termination can be guaranteed, since either k processors terminate with different knowledge as the $n-k$ others, or they do not terminate precisely because of this. Hence by Prop. 3.3, the initial network state $|\psi\rangle$ should allow only n-symmetric moves. With $|\psi\rangle \in (\mathbf{C}^2)^{\otimes n}$ anonymous in the sense of Def. 2.1, this leaves us with GHZ_n, or unitary transforms thereof, as the only possibility. ∎

5.3 More qubits per processor

Again one can repeat the same symmetry argument in the case where each processor has m qubits. Taking $\{|\phi_i\rangle\}$ to be a basis for $\mathcal{H}$, this leads to initial network states of the form

$$|\psi\rangle = \sum_{i=1}^{2^m} \alpha_i |\phi_i\rangle^{\otimes n} \tag{5}$$

An 2^m-valued QDC protocol would then consist of measuring this state. Again, requiring total correctness means that one cannot but use states of this type as a resource.

6 Conclusion

The main contribution of this paper is a demonstration of the special computational power of the W-state, and also of the GHZ-state, and generalizations thereof. A number of new results are established. First, totally correct leader election *is* possible in anonymous quantum networks, which is in stark contrast with the classical situation. Next, we prove that the specific entanglement provided by the W- and GHZ-states, and their generalizations, is the *only* kind that exactly solves leader election and distributed consensus respectively. The W-state has been thought about less in quantum information theory than many other entangled states though it does possess remarkable properties [DVC00]. It is highly persistent for example, unlike the graph states[RBB03] it requires many more measurements on average to destroy the entanglement.

In the programming languages community the relative power of synchronous vs. asynchronous process calculi were compared using symmetry breaking arguments: in fact on the ability to implement leader election [Pal03]. The results of the present paper would have similar consequences on the expressive power of quantum process calculi. In joint work with others we are developing such calculi.

One can and should study the role of the W-state more thoroughly. For example what can be done with it in a variant of the one-way model based on W-states? We are actively investigating these and related questions.

Acknowledgements

We would like to thank Elham Kashefi, Vincent Danos and especially Patrick Hayden for helpful discussions. This work was begun at the Oxford University Computing Laboratory; we would like to thank that institution for its hospitality. Prakash Panangaden was supported by EPSRC (U.K.) and NSERC (Canada) and Ellie D'Hondt by the Vrije Universiteit Brussel.

References

[Ang80] Dana Angluin. Local and global properties in networks of processors. In *Proceedings of the 12th annual ACM symposium on Theory of computing*, pages 82–93. ACM Press, 1980.

[BPR+01] C. H. Bennett, S. Popescu, D. Rohrlich, J. A. Smolin, and A. V. Thapliyal. Exact and asymptotic measures of multipartite pure-state entanglement. *Phys. Rev. A*, 63(1), January 2001.

[DVC00] Wolfgang Dür, Guifre Vidal, and Ignacio Cirac. Three qubits can be entangled in two inequivalent ways. *Phys. Rev. A*, 62(6), December 2000.

[FLP85] Michael J. Fischer, Nancy A. Lynch, and Michael S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, 1985.

[IR90] Alon Itai and Michael Rodeh. Symmetry breaking in distributed networks. *Inf. Comput.*, 88(1):60–87, 1990.

[LeL77] G. LeLann. Distributed systems, towards a formal approach. *Inf. Proc. Lett.*, (77):155–160, 1977.

[Lyn96] Nancy Lynch. *Distributed algorithms*. Morgan Kaufman Publishers, 1996.

[NC00] Michael Nielsen and Isaac Chuang. *Quantum computation and quantum information*. Cambridge university press, 2000.

[Pal03] Catuscia Palamidessi. Comparing the expressive power of the synchronous and the asynchronous pi-calculus. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.

[RBB03] Robert Raussendorf, Daniel E. Browne, and Hans J. Briegel. Measurement-based quantum computation on cluster states. *Phys. Rev. A*, 68(2):022312, 2003.

[Tel94] Gerard Tel. *Introduction to Distributed Algorithms*. Cambridge University Press, first edition, 1994.