Quantum Leader Election via Exact Amplitude Amplification

Seiichiro Tani[1 2 *] Hirotada Kobayashi[3 †] Keiji Matsumoto[3 2 ‡]

[1] *NTT Communication Science Laboratories, NTT Corporation*
3-1, Morinosato-Wakamiya, Atsugi, Kanagawa 243-0198, Japan
[2] *ERATO Quantum Computation and Information Project, JST*
Hongo White Building, 5-28-3 Hongo, Bunkyo-ku, Tokyo 113-0033, Japan
[3] *Foundations of Information Research Division, National Institute of Informatics*
2-1-2 Hitotsubashi, Chiyoda-ku, Tokyo 101-8430, Japan

Abstract. It is well-known that no classical algorithm can solve exactly (i.e., in bounded time without error) the leader election problem in anonymous networks. Recently, Tani, Kobayashi and Matsumoto proved that the problem can be exactly solved when the parties are connected by quantum communication links. This paper gives another two quantum algorithms that exactly solve the problem. The algorithms take a different approach to the problem from the existing ones. The first algorithm uses the quantum amplitude amplification in a distributed manner under the anonymous condition, and runs in $O(n^2)$ rounds and $O(n^4)$-qubit communication for any network topology, where n is the number of parties. The second one is restricted to the case where the number of parties is a power of two, and requires $O(n^6 \log n)$-qubit communication, but takes only linear rounds in n.

Keywords: leader election, quantum amplitude amplification, anonymous networks

1 Introduction

The leader election problem is a core problem in traditional distributed computing, which has been studied for decades (see, e.g., [4]). The goal of the leader election problem is to elect a unique leader from among distributed parties. Obviously, it is possible to deterministically elect a unique leader if each party has a unique identifier. On the other hand, a lot of papers examined the case wherein the network is anonymous, i.e., no party has an identifier [1, 3, 6, 7]. In this setting, no classical exact algorithm (i.e., an algorithm that runs in bounded time and solves the problem with zero error) exists for a broad class of network topologies including regular graphs, even if the network topology (and thus the number of parties) is known to each party prior to algorithm invocation [6]. In the quantum setting (i.e., every party can perform quantum computation and communication and each adjacent pair of parties has a bidirectional quantum communication link between them, but does not share any prior entanglement), the situation is quite different. It was recently proved that the problem can be exactly solved even when the network is anonymous [5].

This paper gives two quantum algorithms that, given the number n of parties, exactly solve the leader election problem in an anonymous network using approaches that are quite different from those of the algorithms in [5]. Our first algorithm elects a unique leader from among n parties by applying exact amplitude amplification [2] in a distributed manner under the anonymous condition. It attains the same performance as that of the algorithm in [5], i.e., it takes $O(n^2)$ rounds and $O(n^4)$ communication complexity for synchronous network of *any* topology. Our second algorithm is restricted to the case wherein the number of parties is a power of two. It takes at most $6n$ rounds for any topology, while the communication complexity is $O(n^6 \log n)$. Both algorithms are easily modified to support their use in asynchronous networks.

*tani@theory.brl.ntt.co.jp
†hirotada@nii.ac.jp
‡keiji@nii.ac.jp

2 $O(n^2)$-round Quantum Algorithm for General Case

First we introduce the concept of *consistent* and *inconsistent* strings. Suppose that each party l has a c-bit string x_l. That is, the n parties share cn-bit string $x = x_1 x_2 \cdots x_n$. For convenience, we may consider that each x_l expresses an integer, and identify string x_l with the integer it expresses. Given a set $E \subseteq \{1, \ldots, n\}$, string x is said to be *consistent* over E if x_l has the same value for all l in E. Otherwise x is said to be *inconsistent* over E. We also say that a cn-qubit pure state $|\psi\rangle = \sum_x \alpha_x |x\rangle$ shared by the n parties is *consistent (inconsistent)* over E if $\alpha_x \neq 0$ only for x's that are consistent (inconsistent) over E.

Next, we quote the exact quantum amplitude amplification theorem, which our algorithm is based on.

Proposition 1 ([2]) *Let $\mathcal{A}$ be any quantum algorithm that uses no measurements to search a truth assignment for any Boolean function χ. Given the initial success probability $a \geq 0.25$ of $\mathcal{A}$, $Q(\mathcal{A}, \chi, \phi, \psi)\mathcal{A}|0\rangle$ gives a correct assignment with certainty by setting ψ and ϕ $(0 \leq \psi, \phi < 2\pi)$ to some appropriate values depending on a, where $Q(\mathcal{A}, \chi, \phi, \psi) = -\mathcal{A}F_0(\phi)\mathcal{A}^{-1}F_\chi(\psi)$ such that: $F_\chi(\psi)$ transforms $|x\rangle$ into $e^{i\psi}|x\rangle$ if $\chi(x) = 1$ and $|x\rangle$ if $\chi(x) = 0$, and $F_0(\phi)$ transforms $|x\rangle$ into $e^{i\phi}|x\rangle$ if $x = 0 \cdots 0$ and $|x\rangle$ otherwise.*

Initially all parties are eligible to become the unique leader. The algorithm repeats one procedure exactly $(n-1)$ times, each of which is called a *phase*. In each phase, the number of parties eligible to be the leader either decreases or remains the same, but never increases or becomes zero. After $(n-1)$ phases the number of eligible parties becomes one with certainty.

Each phase has a parameter denoted by k, whose value is $(n-i+1)$ in the ith phase. In each phase i, let $E_i \subseteq \{1,\ldots,n\}$ be the set of all ls such that party l is still eligible. Each phase prepares the uniform superposition of $2^{|E_i|}$ of x's where $x = x_1 \cdots x_{|E_i|}$ and x_l is a bit possessed by the lth party in E_i. The phase then amplifies the amplitude of any inconsistent state over E_i so that we can get an inconsistent string over E_i with certainty by measurement. This is possible if every party knows the number $|E_i|$ of eligible parties by using Proposition 1, since he/she can compute the exact initial success probability of getting an inconsistent string, which is clearly $1 - \frac{2}{2^{|E_i|}}$. Actually, every party does not know $|E_i|$, however, we can detour to avoid this issue as described later. When amplifying the amplitude, $F_\chi(\psi)$ and $F_0(\phi)$ need to be realized in a distributed manner, i.e., in a way that every party performs the same operation. This can be done as follows. For $F_\chi(\psi)$, every party multiplies the amplitude of any inconsistent state by the factor of $e^{i\frac{1}{n}\psi}$, which multiplies it as a whole by the factor of $e^{i\psi}$. For $F_0(\phi)$, every party multiplies the amplitude of the all-zero state $|00\cdots0\rangle$ by the factor of $e^{i\frac{1}{n}\phi}$, which multiplies it as a whole by the factor of $e^{i\phi}$. Notice that every party can distinguish inconsistent states or the all-zero state by applying a flooding algorithm to the superposition as described in [5].

Once the parties in E_i share an inconsistent string, the number of eligible parties can be reduced with certainty by excluding $l \in E_i$ from E_i such that party l does not have the maximum one-bit value among all one-bit values in the string.

In each phase i, every party uses k instead of $|E_i|$ to set ψ and ϕ to some appropriate values. Hence, the eligible parties may share a consistent state since k does not necessarily represent $|E_i|$. In this case, the above operations that attempt to reduce the eligible parties does not change E_i. In the case of $k = |E_i|$ by chance, $|E_i|$ is reduced by at least one (but not to zero), although they cannot recognize the case. It is clear from this observation that k is always at least $|E_i|$ in each phase i since k is $n = |E_1|$ in the first phase and is decreased by 1 after each phase. It follows that exactly one leader is elected after the last phase.

Theorem 2 *Let m be the number of edges of the underlying graph for the network topology. Given the number n of parties, the algorithm exactly elects a unique leader in $\Theta(n^2)$ rounds. The total communication complexity over all parties is $\Theta(mn^2)$.*

If each party initially knows only the upper bound N of the number of parties, each party has only to perform the above algorithm with N instead of n. The complexity in this case is described simply by replacing every n by N in Theorem 2.

3 Linear-Round Quantum Algorithm for n a Power of Two

First, every party prepares $\frac{|0\rangle+|1\rangle}{\sqrt{2}}$ and $|0\rangle$ in one-qubit registers $\mathbf{R}_0$ and $\mathbf{S}$, respectively. They then set the content of $\mathbf{S}$ to the Hamming weight $(\mathrm{mod}\ 2)$ of the contents in all $\mathbf{R}_0$s by computing the view [6] of depth $2(n-1)$ in a superposition, regarding the contents of $\mathbf{R}_0$'s as node labels. It takes at most $2n$ rounds to construct the view, and after computing the Hamming weight $(\mathrm{mod}\ 2)$ from the view, it takes another $2n$ rounds to invert every computation and communication that was performed to construct the view. Next every party measures the qubit in $\mathbf{S}$ in the $\{|0\rangle,|1\rangle\}$ basis and stores the result into variable y. If $y = 0$ (1), the resulting state $|\psi\rangle$ is the uniform superposition of the n-bit strings that have the Hamming weights of even (resp. odd) values. In the case where $y = 0$, every party applies two unitary operators V and U_n (in this order) to the qubit in $\mathbf{R}_0$ to share a superposition of only the strings that have the Hamming weights of odd values, where

$$U_n = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & e^{-i\frac{\pi}{n}} \\ -e^{i\frac{\pi}{n}} & 1 \end{pmatrix}, \quad V = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & -1 \\ 1 & 1 \end{pmatrix}.$$

By measuring the qubit in $\mathbf{R}_0$, every party gets classical value z. Finally, every party calls another subroutine, which elects a leader by constructing the view of depth $2(n-1)$ whose node labels are z values in $2n$ rounds. From the property of view, we can prove that, when n is a power of two, if the number of parties having 1 is odd, no two parties have an isomorphic view. By regarding the view of every party as its identifier, the parties can elect a unique leader from among them. Therefore the total number of rounds required is at most $6n$ rounds. If we use the technique called f-view [5], the communication complexity is $O(mn^4 \log n)$, where m is the set of edges of the underlying network topology.

If each party initially knows only the upper bound N of the number of parties, each party performs the above algorithm for $k = 2, 3, \ldots, N$ instead of n in parallel, and with a bit elaboration, the parties can elect a unique leader in $6N$ rounds and $O(mN^5 \log N)$.

References

[1] D. Angluin. Local and global properties in networks of processors (extended abstract). In *Proc. of 20th ACM STOC*, pages 82–93, 1980.

[2] G. Brassard, P. Høyer, M. Mosca, and A. Tapp. Quantum amplitude amplification and estimation. In *Quantum Computation and Quantum Information: A Millennium Volume*, volume 305 of *AMS Contemporary Mathematics Series*, pages 53–74. 2002.

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

[4] N. A. Lynch. *Distributed Algorithms*. Morgan Kaufman Publishers, 1996.

[5] S. Tani, H. Kobayashi, and K. Matsumoto. Exact quantum algorithms for the leader election problem. In *Proc. of 22nd STACS*, volume 3404 of *Lecture Notes in Computer Science*, pages 581–592, 2005.

[6] M. Yamashita and T. Kameda. Computing on anonymous networks: Part I – characterizing the solvable cases. *IEEE Trans. Parallel Distrib. Syst.*, 7(1):69–89, 1996.

[7] M. Yamashita and T. Kameda. Computing on anonymous networks: Part II – decision and membership problems. *IEEE Trans. Parallel Distrib. Syst.*, 7(1):90–96, 1996.