Quantum Solution to the Byzantine Agreement Problem

Matthias Fitzi,[1] Nicolas Gisin,[2] and Ueli Maurer[1]

[1]*Department of Computer Science, Swiss Federal Institute of Technology (ETH), CH-8092 Zurich, Switzerland*
[2]*Group of Applied Physics, University of Geneva, CH-1211 Geneva 4, Switzerland*

(Received 25 July 2001; published 1 November 2001)

We present a solution to an old problem in distributed computing. In its simplest form, a sender has to broadcast some information to two receivers, but they have access only to pairwise communication channels. Unlike quantum key distribution, here the goal is not secrecy but agreement, and the adversary (one of the receivers or the sender himself) is not outside but inside the game. Using only classical channels this problem is provably impossible. The solution uses pairwise quantum channels and entangled qutrits.

DOI: 10.1103/PhysRevLett.87.217901 PACS numbers: 03.67.Lx, 02.50.Le

Entanglement is a resource that allows quantum physics to perform tasks that are classically impossible. This is the new leitmotif of quantum information processing. The best known examples are quantum cryptography [1,2] and Shor's algorithm to efficiently factorize large numbers [3]. In this Letter we consider an old information-theoretical problem in the field of fault-tolerant distributed computing, known as *the Byzantine agreement problem* [4] and present a solution which exploits entanglement between three qutrits (i.e., three 3-dimensional quantum systems).

Imagine that several divisions of the Byzantine army are camped outside an enemy city, each division commanded by its own general. The generals can communicate with one another only by messengers. One of the generals, the commanding general, after observing the enemy, must decide on a plan of action and communicate it to the other generals. However, some of the generals (especially the commanding general himself) might be traitors, trying to prevent the loyal generals from reaching agreement on the plan of action. The question hence is whether there is a protocol among the generals that, after its termination, satisfies the following conditions:

(1) All loyal generals agree on a common plan of action.

(2) If the commanding general is loyal, then all loyal generals agree on the commanding general's plan.

More precisely we define Byzantine agreement (shortly broadcast) as follows.

Definition 1: A protocol among n players such that one distinct player S (the sender) holds an input value $x_S \in \mathcal{D}$ (for some finite domain $\mathcal{D}$) and all other players (the receivers) eventually decide on an output value in $\mathcal{D}$ is said to achieve *broadcast* if the protocol guarantees that all honest players decide on the same output value $y \in \mathcal{D}$ and that $y = x_S$ whenever the sender is honest.

In modern terms, this problem concerns coordination in distributed computing (among several processors or computers) where some of the processors might fail. For example, a database can be replicated among several servers in order to guarantee access to the database even if some of the servers misbehave. Nevertheless, an inconsistent external update of the database must result in all honest servers having exactly the same views on the database. Consider, for instance, that the database contains the price of valuable goods, or currency exchange rates. It is then important that no adversary, not even an inside adversary, can affect the coordination in such a way that the prices would be low somewhere and high elsewhere.

The broadcast problem has been considered in a vast literature and has developed several variations [4]. Here we define the problem more precisely as follows. Three players are connected by pairwise authenticated classical and quantum channels; see Fig. 1. For simplicity, we assume the channels to be error-free—generally, errors would have to be additionally dealt with by means of error correction codes. The general purpose is that one of the players, namely, the sender (S for short), broadcasts a bit to his two partners, the receiving players R_0 and R_1. Both receivers should end with the same value. However,

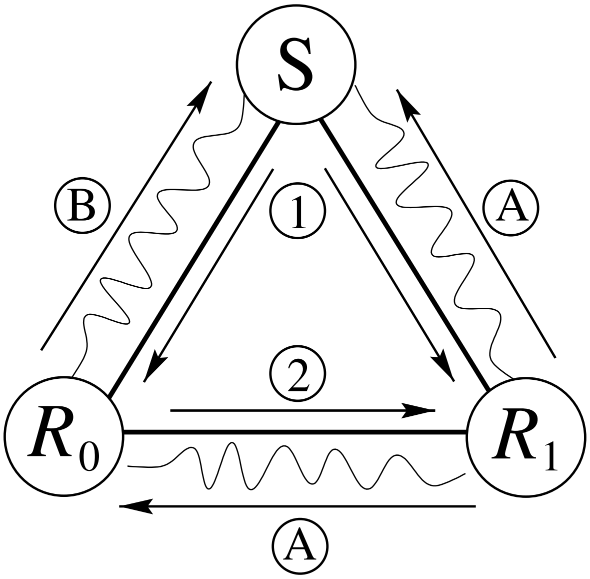

FIG. 1. Flow of classical (straight lines) and quantum (wavy lines) information. Note that both kinds of information flow exactly in opposite directions. This is needed to avoid the possibility that the adversary can bring in confusion in the last communication round.

one—and at most one—of the three players might actually be an active adversary. For instance, a dishonest sender could send different bit values to R_0 and R_1. The receivers may realize that there is a problem simply by exchanging their bits. But then, player R_1 cannot conclude whether the sender is dishonest and he should keep the bit received from R_0 or whether player R_0 is cheating and he should keep the bit received from the sender. It is not too difficult to convince oneself that because the players have access only to pairwise channels, this task is not obvious.

If the players have access only to classical pairwise authenticated channels, the broadcast problem is provably unsolvable [4,5]. This holds even for arbitrary pairwise communication, i.e., not even quantum channels can help to solve the problem [6]. However, we demonstrate that the additional resource of the quantum channel allows them to solve a slightly weaker problem, namely, detectable broadcast, which is powerful enough for a large range of applications of this problem.

Definition 2: A protocol among three players such that one distinct player S (the sender) holds an input value $x_S \in \mathcal{D}$ (for some finite domain $\mathcal{D}$) and the other two players (the receivers) eventually decide on an output value in $\mathcal{D}$ is said to achieve *detectable broadcast* if the protocol satisfies the following conditions: (1) If no player is corrupted then the protocol achieves broadcast. (2) If one or more players are corrupted then either the protocol achieves broadcast or all honest players abort the protocol.

Note that detectable broadcast cannot be solved only with pairwise authenticated classical channels. However, we demonstrate that pairwise authenticated classical *and quantum* channels are sufficient to solve the problem. Basically, we solve the problem by having the players (1) distribute entanglement, (2) check that the entangled states are not corrupted, and (3) use them to solve the problem.

At first sight, this sounds very similar to quantum cryptography. But, actually, it is very different. Indeed, here we do not require any secrecy: what counts is to avoid any discord. Also, here, in contrast to quantum cryptography, the adversary is not an outside player, but might be anyone among the three players.

The first point of the above program, i.e., distributing entanglement, is trivial (in theory), since we assume that quantum channels are available. The testing (i.e., the second point), however, is tricky. Indeed, the testing requires (classical) communication between the three players. But the adversary being inside the game could corrupt this communication phase. Especially at the last round of the communication phase, the adversary could send contradictory messages to the two honest players [7]. In other quantum information protocols involving more than two players, e.g., in quantum secret sharing [8], this problem is avoided by assuming that the players can broadcast their (classical) messages. But here broadcasting is not assumed among the primitives; on the contrary, it is the goal of the game. Below we show how to break this vicious circle. But first, we need to explain how the three players can use entangled qutrits to solve the problem.

Let us assume that the three players share many qutrits triplets Ψ_j, each in the Aharonov state $|\mathcal{A}\rangle$:

$$|\mathcal{A}\rangle = (|0,1,2\rangle_{\vec{m}} + |1,2,0\rangle_{\vec{m}} + |2,0,1\rangle_{\vec{m}} - |0,2,1\rangle_{\vec{m}} - |1,0,2\rangle_{\vec{m}} - |2,1,0\rangle_{\vec{m}})\frac{1}{\sqrt{6}}, \quad (1)$$

where $|0,1,2\rangle_{\vec{m}}$ denotes the tensor product state $|0\rangle_{\vec{m}} \otimes |1\rangle_{\vec{m}} \otimes |2\rangle_{\vec{m}}$. If one identifies qutrits with spin-1 and associates the state $|2\rangle_{\vec{m}}$ with the eigenvalue -1 of the spin operator $\vec{m}\vec{S}$, then the Aharonov state [9] is the unique three spin-1 state of total spin zero. Consequently—and analogously to the singlet state of qubit pairs—the Aharonov state is invariant under trilateral rotations: it keeps the same form (1) for all directions $\vec{m}$, where $|0\rangle_{\vec{m}}$, $|1\rangle_{\vec{m}}$ and $|2\rangle_{\vec{m}}$ are the three eigenvectors of the spin operator $\vec{m}\vec{S}$. We exploit the fact that whenever the three qutrits are all measured in the same basis, then all three results differ.

With the help of this additional resource, i.e., the Aharonov states, the protocol runs as described below. At each step we comment on the reasons why this is safe. Actually, all steps are rather trivial, except the last one which needs a careful analysis.

(1) First, the sender S sends the bit x to be broadcast to the two receivers R_0 and R_1, using the classical channels. Let us denote x_0 and x_1 the bits received by R_0 and R_1, respectively. Next, the sender S measures all his qutrits in the z basis. Whenever he gets the result x, S sends the index j to both receivers [10]. Accordingly, the players R_0 and R_1 receive each a set of indices, J_0 and J_1, respectively (label ① in Fig. 1).

(2) Both receivers test the consistency of their data. For this they measure their qutrits in the z basis. If all results with indices in J_p differ from x_p, then player R_p has consistent data and he sets a flag $y_p = x_p$. If a set of data is inconsistent, then the player sets his flag to $y_p = \perp$ (interpreted as *inconsistent*).

(3) The two receivers send their flags to each other. If both flags agree then the protocol terminates with all honest players agreeing on x.

(4) If $y_p = \perp$, player R_p knows that the sender is dishonest. He concludes that the other receiver is honest and he simply accepts the bit he receives from him (if $y_0 = y_1 = \perp$, then they both end with the "value" $\perp$).

(5) It remains only the interesting case that both receivers claim that they received consistent, but different, data. The strategy we propose then is that player R_1 will not change his bit y_1, unless player R_0 convinces him that he did indeed receive the bit y_0 from the sender in a consistent way. To convince his partner of his honesty, player R_0 sends him all the indices $k \in J_0$ for which he has the result $1 - y_0$ (label ② in Fig. 1).

(6) Receiver R_1 now checks that he gets "enough" indices k from R_0 such that (a) "almost all [11]" indices k from R_0 are not in R_1's index set J_1, and such that (b) these k indices correspond to qutrits for which R_1's results are "almost all" equal 2.

If R_0 indeed got an index set that is consistent with bit y_0 then S holds y_0, R_0 holds $1 - y_0$, and hence, R_0's result must be a 2. If the test succeeds, player R_1 changes his bit to y_0; otherwise he keeps y_1.

Let us examine why player R_0 cannot cheat (see Table I). Assume that R_0 receives the bit $x_0 = 0$, but pretends that he got 1. To convince the receiver R_1 to accept his "pretended bit," player R_0 must first announce that he received consistent data (which is true, but with bit value 0), and next send a sufficiently large set of indices $\{k\}$ with almost no intersection with J_1 and for which R_1 almost always has the result 2. Since player R_0 has no information on the indices outside J_0 for which he measured $1 - y_0$ (other than $1 - y_0$ itself), approximately half of the indices that player R_1 gets are different from 2—which is not accepted by player R_1.

Let us stress an important feature of this protocol: the last player, i.e., player R_1, almost never talks (he only sends his value y_1 to R_0). Moreover, if the sender is honest, then the last player never changes his mind. This is important for the distribution and test phase of the protocol described below.

So far we described a protocol assuming that the three players share a large collection of qutrit triplets in the Aharonov state (1). We now describe a protocol to distribute and test such states. This protocol uses only pairwise communication, in particular, no broadcasting is assumed. Nevertheless, the protocol has only two possible outcomes: global success or global failure. By global we mean that all honest players end with the same conclusion. In case of failure, the broadcasting protocol does not even start. In case of success, broadcasting can be realized reliably.

The distribution-and-testing protocol works as follows:

(1) Player R_1 prepares many qutrit triplets Ψ_j in the Aharonov state (1). For each index j he sends one qutrit to player S and one to R_0 (label Ⓐ in Fig. 1).

(2) Both S and R_0 check that their qutrits are in the maximally mixed state. In case of success, they set a flag f_p to 1, else to 0.

(3) Player R_0 sends a sample of his qutrits to S (label Ⓑ in Fig. 1). Player S tests that the sample of qutrit pairs he now holds is in the correct state [12]:

$$\rho_{s_{R_0}} = \mathrm{Tr}_{R_1}(|\mathcal{A}\rangle\langle\mathcal{A}|) = \tfrac{1}{3}\big(P_{|1,2\rangle-|2,1\rangle} + P_{|2,0\rangle-|0,2\rangle} + P_{|0,1\rangle-|1,0\rangle}\big). \quad (2)$$

If the test fails, then he sets his flag to 0: $f_S = 0$.

(4) Player R_1 sends a sample of his qutrits to R_0 and another sample to S. Both R_0 and S test their qutrit pairs as in the previous point (3). If the test fails they set their flag to 0.

(5) Player S and R_0 exchange their flag. If a player receives a 0, then he sets his flag to zero.

(6) Both players S and R_0 broadcast their flags using the protocol described previously.

(7) Any player with flag 1 who received a 0-flag changes his flag to 0.

At first, step (6) of the above protocol seems impossible, since the broadcast protocol requires reliable Aharonov states. Nevertheless, let us look closer at this step. If player R_1 does not produce the correct states, then, since by assumption there is no more than one dishonest player, players S and R_0 are honest and both will end with their flag on failure: $f_S = f_{R_0} = 0$. Let us thus assume that all states $\Psi_j = |\mathcal{A}\rangle$. Consequently, the broadcasting is reliable. All that a dishonest player S, R_0, or R_1 could do is to act in such a way that the flags are set to 0 [13]. But during the last step of the protocol, i.e., the broadcast sessions, both the one initiated by S and the one initiated by R_0 are reliable. Hence it is impossible that some players end this protocol thinking that a status of success has been reached, while another one thinks the opposite. Moreover, if all players agree on success, then they share Aharonov states and they can reliably run the broadcast protocol.

The field of quantum communication is still in its infancy. Only very few protocols concern more than two parties and almost all use qubits. In this Letter we presented a protocol among three players connected by pairwise quantum channels able to transmit qutrits and to preserve their entanglement. The protocol is a version of the well known Byzantine agreement problem, a very timely problem in today's information based society. Admittedly, the problem has been slightly adapted to fit into the quantum frame, a natural synergy between classical and quantum information theories. It is not too difficult to generalize our result to n players with $t < \frac{n}{2}$ cheaters, though this is beyond the present Letter [14].

One may question whether the use of qutrits is necessary or not for broadcasting. Clearly, the present protocol is

TABLE I. After measuring their qutrits, the sender's S and receivers' R_1 and R_0 results fall into six classes, labeled with Roman numbers. The index sets J_0 and J_1 associated to the bit values 0 and 1 correspond to the labels I, II and III, IV, respectively. If R_0 receives the bit 0 and the set J_0 he can announce to R_1 all cases where he has the bit 1: all cases labeled by I. For all these cases R_1 has the value 2. However, if R_0 tries to cheat and pretends to have received a bit 1, then he cannot differentiate between the cases labeled IV and V. For the latter, R_1 has a value 1; he can thus detect cheating R_0.

	I	II	III	IV	V	VI
S	0	0	1	1	2	2
R_0	1	2	2	0	0	1
R_1	2	1	0	2	1	0
	$\underbrace{\qquad}_{J_0}$		$\underbrace{\qquad}_{J_1}$			

intimately related to the Aharonov state, hence to qutrits. However, qutrits can be teleported with the help of several singlet states and classical communication. Hence, the protocol could be built on qubits. But this would be a waste of resources. Indeed, one can prove that the Aharonov state cannot be efficiently converted into singlet states [15]. Consequently, qutrits are not strictly necessary, but qutrits are definitively more efficient than qubits for broadcasting: the resource of entanglement will be used more efficiently if qutrits are entangled rather than qubits.

Two other features of our protocol should be mentioned. First, the quantum states are used "only" to distribute classical private random variables with specific correlation to the three players (a trit per player, each of a different value, all combination with equal probabilities). This is similar to quantum cryptography where quantum mechanics provides only key distribution. However, contrary to the "one-time-pad" algorithm used in conjunction with quantum key distribution, the present algorithm was itself inspired by the elegance of the Aharonov quantum state. Finally, experimental demonstration of the protocol can be realized with today's technology, using photons and 3-paths interferometers. Actually one would not need to prepare three entangled photons; two would suffice since the preparer R_1 could measure his qutrit immediately, similarly to the demonstration of quantum secret sharing using pairs of photons [16].

This work is supported by the Swiss Science Foundation and the IST-EQUIP European project.

[1] C. Bennett and G. Brassard, in *Proceedings of the International Conference on Computers, Systems, & Signal Processing, Bangalore, India, 1984* (Indian Institute of Science, Bangalore, 1984), pp. 175–179.

[2] A. Ekert, Phys. Rev. Lett. **67**, 661 (1991).

[3] P. W. Shor, in *Proceedings of the 35th Annual Symposium on Foundations of Computer Science, Santa Fe, NM, 1994,* edited by Shafi Goldwasser (IEEE Computer Society Press, Los Alamitos, CA, 1994), pp. 124–134.

[4] L. Lamport, R. Shostak, and M. Pease, ACM Trans. Programming Languages Syst. **4**, 382–401 (1982), and references therein.

[5] M. J. Fischer, N. A. Lynch, and M. Merritt, Distrib. Comput. **1**, 26–39 (1986).

[6] M. Fitzi, J. A. Garay, U. Maurer, and R. Ostrovsky, *Advances in Cryptology — CRYPTO 2001: Proceedings of the 21st Annual International Cryptology Conference, Santa Barbara, CA, 2001,* Lecture Notes in Computer Science (Springer, Berlin, New York, 2001).

[7] This is why no classical protocol can achieve detectable broadcast: in any protocol there must be a last communication round. An adversary could behave properly until then and only at the very last round send confusing data to the honest players. Since this was the last round, the honest players cannot detect that something went wrong.

[8] M. Hillery, V. Buzek, and A. Berthiaume, Phys. Rev. A **59**, 1829 (1999); A. Karlsson, M. Koashi, and N. Imoto, Phys. Rev. A **59**, 162 (1999); R. Cleve, D. Gottesman, and H.-K. Lo, Phys. Rev. Lett. **83**, 648 (1999).

[9] Yakir Aharonov and Sandu Popescu presented us with this state, arguing that it is so elegant that it must be useful for something (private communication).

[10] In practice, it is more realistic to assume that each player measures all the qutrits in randomly chosen bases. Then, each time they communicate a trit value, they need to add information about the measurement basis, and each time they receive a trit they ignore it unless it happens that they measured their corresponding qutrit in the same basis.

[11] This qualification is needed for statistical tests. In the limit of arbitrarily many qutrit-triplets, "almost all" translates into "with probability one."

[12] To see that this test is sufficient to guarantee that the player R_1 who prepared the states cannot cheat, consider a general purification Ψ of the mixed state $\rho_{s_{R_0}}$. Dividing the three parts into R_1 versus the two others, one can write Ψ in the Schmidt form. Using the fact that the eigenstates of $\rho_{s_{R_0}}$ are the three states $|n, n+1\rangle - |n+1, n\rangle$, $n = 0, 1, 2$, one obtains $\Psi = |\alpha_0\rangle \otimes (|1,2\rangle - |2,1\rangle) + |\alpha_1\rangle \otimes (|2,0\rangle - |0,2\rangle) + |\alpha_2\rangle \otimes (|0,2\rangle - |2,0\rangle)$. Since, by virtue of the Schmidt decomposition the three states $|\alpha_n\rangle$ are mutually orthogonal, this is precisely the Aharonov state (up to phases that can be changed locally and do not affect the correlation in the z basis). An equivalent test can be performed using only local measurements in randomly chosen bases and classical pairwise communication: S chooses the bases and R_0 announces his results.

[13] This is similar to quantum key distribution where Eve can block the key distribution, but not extract information without revealing herself.

[14] M. Fitzi, N. Gisin, U. Maurer, and O. von Rotz, "Unconditional Byzantine Agreement and Multiparty Computation Secure Against Dishonest Minorities from Scratch" (to be published).

[15] The proof is a direct application of Eqs. (3)–(6) of E. F. Galvão, M. B. Plenio, and S. Virmani, J. Phys. A **33**, 8809 (2000), using the relative entropy of entanglement computed in K. Audenaert *et al.,* quant-ph/0103096 (private communication by A. Acin).

[16] W. Tittel, N. Gisin, and H. Zbinden, Phys. Rev. A **63**, 042301 (2001).