Reference frame agreement in quantum networks

Tanvirul Islam,[1, 2, *] Loïck Magnin,[1, †] Brandon Sorg,[1] and Stephanie Wehner[1, 2, ‡]

[1] *Centre for Quantum Technologies, National University of Singapore, 3 Science Drive 2, 117543 Singapore*
[2] *School of Computing, National University of Singapore, 13 Computing Drive, 117417 Singapore*

In this work we design a multiparty protocol between m players to agree on a common direction without a prior reference frame shared between the players. Our protocol is tolerant to $t < m/3$ dishonest players with unbounded capabilities (the Byzantine problem). This is the first protocol to exchange non-fungible information in the Byzantine setting.

Quantum theory uses physical properties such as momentum, position, phase, or time to encode information. These properties are not absolute, but relative to a reference frame. When several parties need to communicate quantum information a shared reference frame between all the parties is very handy. Even though some tasks can be done without, sharing a common reference frame is an important tool to perform efficient shared computation, and cryptographic tasks such as QKD.

One can see, that to share a Cartesian reference frame one first have to be able to share a direction. But without a pre-shared reference frame, one cannot exchange directional information by exchanging only classical data. So, one must first exchange some physical object that points towards a certain direction. The information carried by such object is called *non-fungible information.* Sharing a direction can only be done approximately and with some probability of failure.

In this work, we examine the multiparty scenario in which m parties should align their reference frames in the Byzantine fault tolerance model, the players should succeed despite the faulty behavior of some of them. Byzantine fault tolerance does not consider only failures (in which players stop executing the protocol), but any arbitrary errors. In particular, faulty players can send incorrect messages, terminate the protocol, and even coordinate their actions to fool the honest players. This model also encompasses errors in the communication by considering the sender faulty. Faulty parties are also sometimes called dishonest players, and can be considered as a computationally unbounded adversary. This model allows the strongest possible errors, and has been introduced by Lamport *et al.* [1] in a distributed computation setting in which honest players reach a consensus on a bit.

The name "Byzantine" comes from a problem faced by the Byzantine army. Each of its divisions was led by a general, and these generals could communicate between them only by couriers. All the generals should agree on a strategy: either the all attack, or the all retreat. However some of the generals were disloyal to the Byzantine state. Thus the challenge for the Byzantine generals was to reach a consensus on bit despite the presence among them of dishonest generals. In this work, we show that quantum generals, can reach a consensus, not only on binary order such as attack or retreat, but on an arbitrary direction, without the use of a reference (like north, or the position of some stars).

Previous work. Byzantine consensus on a bit has been heavily studied in several variations of model of computation (deterministic, randomized, and quantum) with synchronous and asynchronous communication. The most celebrated results show that there is a fully polynomial protocol to Byzantine agreement on a bit tolerant to $t < m/3$ dishonest players, and that this value is optimal.

The use of quantum communication has been considered in [2] as a mean to decrease to total amount of communication needed, and in [3] in a fail-stop model where the authors show that 3 quantum players can reach a consensus even if one them is dishonest.

Our contributions

We introduce a quantum protocol to solve the Byzantine reference frame consensus problem that is tolerant to $t < m/3$ dishonest players. Our protocol succeeds with probability at least $q_{\text{succ}}^{m^2}$ and with approximation 30δ where δ is the approximation done by sending one direction between two players, and q_{succ} its success probability.

Model of communication. We assume that we have authenticated and synchronous classical and quantum channels for communications. This means that the players know when they are supposed to receive a message, and from whom. As a consequence, a player cannot be waiting an unlimited amount of time for a message, this ensures that our protocol terminates.

We only use quantum communication to send a direction between a sender and a receiver. We use one of the simplest protocol to do it. A sender creates many identical qubits with their Bloch vector pointing to the intended direction and the receiver measure them with Pauli measurements. From the statistics of the measurement outcomes, the receiver then estimates the Bloch vector's direction closely with high success probability. We do not require any quantum memory, or entangled states, which potentially simplifies any experimental implementation. But the downside of this choice is that our

[*] tanvir@locc.la
[†] loick@locc.la
[‡] steph@locc.la

protocol is not optimal in the number of qubits sent to achieve a certain accuracy. Optimal protocols can align frames in the so-called Heisenberg limit, that is they have a quadratic gain over the one we use here.

We also show that our protocol is tolerant to depolarizing noise in the quantum channels. This is an important fact since in the standard Byzantine setting, such noise would be handled by considering the sender dishonest, and thus all players would be dishonest!

Our protocol is independent of the implementation of this two-party estimate direction protocol, we only use it as a black box. Therefore, one can plug in his favorite two party direction estimating protocol. Throughout all this work, all our results are parametrized by the estimation error δ, and the success probability q_{succ} of this black box.

Overview. Our protocol is designed in several layers. The most important top layer protocol is called *King Consensus*. In this protocol, one player is a king: he chooses one arbitrary direction and sends it to all the other players. All the honest players should then decide to accept this direction, or they should all decide to reject it. More precisely, the protocol satisfies *Persistency*: if the king is honest, all the honest players will reach a consensus on directions which are very close to the direction initially chosen by the King. It also satisfies *Consistency*: no matter whether the king is honest or dishonest, the honest players either all agree on a common direction, or all of them declare the King a cheater.

Running the King Consensus protocol with $t+1$ different players as king ensures us then at least one of these king will be honest, and thus a consensus will be reached. The rest of this work is devoted to construct a King Consensus protocol.

Before introducing the inner working of the king consensus protocol, we will see two more protocols, namely Weak Consensus protocol and Graded Consensus protocol. As the name suggest, the Weak Consensus protocol satisfies a weaker version of persistency and consistency. This protocol is used by the Graded Consensus protocol, which satisfies stronger persistency and consistency requirements. Ultimately the King Consensus protocol uses the Graded Consensus to satisfy the previously mentioned properties.

Weak Consensus. We now see how Weak Consensus works. Here, every player starts with a direction w_i as input and outputs either a direction u_i or $\perp$. If all the inputs w_i of the honest players are δ-close to some common direction s, then the honest players should reach a consensus, i.e. they should output a direction u_i also close to s. However, we do not require that the honest players reach a consensus if all the w_i are not close to some s. In this case, they are allowed to output $\perp$. An honest player outputting $\perp$ can be interpreted as the honest player declaring that his input w_i is far from most of the inputs of the other honest players.

The protocol satisfies δ-Weak Persistency, which means, if all the honest players starts with a input direction which are at least δ-close to a certain direction s, which they are not aware of, then they will output directions which are also at least δ close to s. It also satisfies (8δ)-Weak Consistency, which says, any two honest players that output a non-$\perp$ direction, must be at least (8δ)-close to each other. This protocol succeeds with probability at least $q_{\text{succ}}^{m^{-}m}$. We recall that δ denotes the inaccuracy of approximation the two-party direction estimating protocol used, and q_{succ} its success probability.

Protocol 1: WEAK-CONSENSUS

```
Input  : Direction $w_i$
Output : Direction $u_i$ or $\perp$
Send $w_i$ to all other players
Receive $a_i[j] \leftarrow$ direction received from $P_j$
Create the set $S_i \leftarrow \{P_j : d(w_i, a_i[j]) \leq 3\delta\}$
if $|S_i| \geq m - t$ then
    Assign, $u_i \leftarrow w_i$
else
    Assign $u_i \leftarrow \perp$
Output $u_i$
```

Weak Consensus works by each player P_i making a set S_i of all the players who have sent a direction close to P_i's input direction w_i. If the set contains more than $m-t$ players, he outputs w_i. Otherwise, he outputs $\perp$.

Proving persistency is immediate. The proof of the (8δ)-consistency, that is $d(u_i, u_j) \leq 8\delta$ if the players P_i and P_j are honest, works as follows: First, we show that there exists an honest players P_k who is in both sets S_i and S_j. This is true since S_i and S_j contains more than $m/3$ honest players each. Secondly, we prove that $d(u_i, w_k) \leq 4\delta$ and $d(u_j, w_k) \leq 4\delta$. Indeed, by definition of the set S_i, the distance between u_i and the approximation $a_i[k]$ of w_k is less than 3δ, and that approximation is δ-close to w_k.

Graded Consensus. Using Weak Consensus we design a higher level protocol called *Graded Consensus*. Similarly to the Weak Consensus, the players have as input a direction w_i. However the outputs differ in two ways. First, the players are not allowed to output $\perp$ anymore, they have to output a direction; and secondly, the players output a grade $g_i \in \{0, 1\}$. The grade does not simply replace the output $\perp$. The $\perp$ output has a "local meaning" whereas the grade has a "global significance". Indeed, an honest player P_i outputting grade $g_i = 1$ means that *all* the honest players have reached an agreement—even if they are not all aware of it—, the $\perp$ output in a Weak Consensus protocol for an honest player P_i simply meant that his *own* input differs too much to most of the other honest players.

This protocol satisfies δ–Graded Persistency. That is, if all the honest players start with input which are at least δ close to some direction s, there all of them output grade 1 and their output directions are also at least δ-close to s. The protocol satisfies (30δ)–Graded Consistency, which states, if at least one honest player P_i outputs grade $g_i = 1$ then all the honest players are at least (30δ)-close to each other.

This protocol works by first running the Weak Consen-

Protocol 2: GRADED-CONSENSUS

```
Input  : A direction $w_i$
Output : A direction $v_i$ and a grade $g_i \in \{0, 1\}$
Run WEAK-CONSENSUS($w_i$)
if $u_i = \perp$ then
    Send flag $f_i = 0$ to all other players
else
    Send flag $f_i = 1$ to all other players
forall the Player $j$ do
    $f_i[j] \leftarrow$ Receive $f_j$
forall the Player $j$ do
    Create set $S_i[j] = \{P_k : f_i[k] = 1$, and
    $d(a_i[j], a_i[k]) \leq 10\delta\}$
Assign $l_i \leftarrow \arg\max\{|S_i[j]|\}$
if $f_i = 1$ then
    Assign $v_i \leftarrow w_i$
else
    Assign $v_i \leftarrow a_i[l_i]$
if $|S_i[l_i]| > m - t$ then
    Assign $g_i \leftarrow 1$
else
    Assign $g_i \leftarrow 0$
Output $(v_i, g_i)$
```

sus and keeping only outputs that are closed to each other. Those outputs are then clustered into sets, depending on the closeness of their Weak Consensus outcomes. The grade $g_i = 1$ is awarded only if there is a large set $S_i[j]$ that contains many honest players close to each other.

It is easy to see that the protocol is δ-persistent. The idea of the proof of the (30δ)-consistency is as follows: we show that if one honest player output grade 1, then the largest set of each honest players contains at least one honest player. Hence, each output is (10δ) close to a honest v, and the weak consensus assures us that all the honest v are (8δ)-close to each other.

King Consensus. The King Consensus protocols is build on top of the Graded Consensus protocol.

Protocol 3: KING-CONSENSUS

```
Input  : Id of the king, $P_k$.
Output : A direction $v_i$ or $\perp$
if I am the king then
    Fix an arbitrary direction $w_k$
    Send $w_k$ to all other players
else
    Receive $w_i$, approximation of $w_k$ from the king
Assign $(v_i, g_i) \leftarrow$ GRADED-CONSENSUS($w_i$)
Assign $y_i \leftarrow$ CLASSICAL-CONSENSUS($g_i$)
if $y_i = 1$ then
    Output $v_i$
else
    Output $\perp$
```

If the king is honest, all honest players will have grade $g_i = 1$, hence the classical consensus will be reached on $y_i = 1$ and the honest players will accept the direction shared by the king. If the king is dishonest, the only possibility for the honest players to reach a consensus on a direction, is to have $y_i = 1$. This is possible only if at least one of the grades of the honest players is $g_i = 1$. In this case the (30δ) Graded Consistency implies the (30δ) consistency of the King Consensus protocol, and thus of the complete protocol.

Discussion

We have presented the first protocol for reference frame agreement in a quantum network. Even in the classical setting, the algorithms to solve the Byzantine agreement problem are surprisingly complicated. We would be very keen to know if simpler and more efficient protocols could be designed for our setting, possibly by using entangled states. It is an interesting open question to construct protocols that also work in an asynchronous communication model. The latter is already challenging for the classical case [4–7], so we expect a similar behavior to hold here. Another interesting question is whether more dishonest players than $t < m/3$ can be tolerated. If our protocol were to succeed with probability 1 and δ sufficiently small, we can prove that it is optimal in that sense by adapting the classical proof [8] to our setting. However, for aligning reference frames, any protocol can only succeed with probability strictly less than 1. This problem has been partially studied in the classical case [9]. Even in the constant error scenario the optimal number of dishonest players that can be tolerated is not known for the classical Byzantine agreement problem [10]. This leaves hope to find protocols that can tolerate $t < m/2$ dishonest players when allowing constant success probability both for Byzantine and reference frame agreement.

Acknowledgments

We thank Esther Hänggi and Jürg Wullschleger for useful discussions. This work is funded by the Ministry of Education (MOE) and National Research Foundation Singapore, as well as MOE Tier 3 Grant MOE2012-T3-1-009.

[1] L. Lamport, R. Shostak, and M. Pease, ACM T. Prog. Lang. Sys. **4**, 382 (1982).

[2] M. Ben-Or and A. Hassidim, in *Proc. ACM STOC'05* (ACM, 2005), pp. 481–485.

[3] M. Fitzi, N. Gisin, and U. Maurer, Phys. Rev. Lett. **87**, 217901 (2001).

[4] G. Bracha, in *Proc. ACM PODC'84* (1984), pp. 154–162.

[5] R. Canetti and T. Rabin, in *Proc. ACM STOC'93* (ACM, 1993), pp. 42–51, ISBN 0-89791-591-7.

[6] I. Abraham, D. Dolev, and J. Y. Halpern, in *Proc. ACM PODC'08* (ACM, 2008), pp. 405–414.

[7] I. Abraham, M. K. Aguilera, and D. Malkhi, in *Proc. DISC'10* (2010), pp. 4–19.

[8] M. J. Fischer, N. A. Lynch, and M. Merritt, in *Proc. ACM PODC'85* (1985), pp. 59–70.

[9] R. L. Graham and A. C. Yao, in *Proc. ACM STOC'89* (1989), pp. 467–478.

[10] M. Fitzi, S. Wolf, and J. Wullschleger, in *Proc. IEEE ISIT'06* (2006), pp. 504–505.