Enabling Classical Detectable Byzantine Broadcast and Consensus Through Quantum Error Correction Circuits

Farzam Nosrati
Nexus for Quantum Technologies,
University of Ottawa, Canada
IMDEA Networks Institute, Spain
fnosrati@uottawa.ca

Nicolás Borrajo
ETH Zurich, Switzerland
nico.borrajo@gmail.com

Antonio Fernández Anta
IMDEA Software Institute, Spain
IMDEA Networks Institute, Spain
antonio.fernandez@imdea.org

Vincenzo Mancuso
University of Palermo, Italy
IMDEA Networks Institute, Spain
vincenzo.mancuso@ieee.org

***Abstract*—We design and experimentally validate quantum circuits that prepare a quantum resource state which, when combined with authenticated classical communication, enables a protocol for detectable Byzantine broadcast and detectable Byzantine consensus. For the *broadcast* version of the problem, we propose an efficient leader–follower quantum circuit that prepares this resource state with depth $\mathcal{O}(\log n)$, using $\mathcal{O}(n)$ two-qubit gates and $\mathcal{O}(n)$ ancillary qubits. We also develop a majority-based leader-selection scheme, for the *consensus* version of the problem; the scheme is based on the same leader-follower quantum circuit and on coherent Hamming-weight computation, which prepares a fictitious leader qubit directly from any initial network state specified by the participant parties. The resulting fictitious leader state is then used as input to the leader–follower circuit to prepare a resource state consistent with majority voting on a common basis, thereby enabling detectable Byzantine consensus. In particular, we prove that, when supplemented with authenticated classical channels for the communication of quantum measurement outputs, the resulting protocols achieve detectable Byzantine broadcast and detectable Byzantine consensus, respectively. Under correct operation, all correct parties output the same value, namely the leader's value or the majority value. In the presence of Byzantine parties or quantum circuit deviations, all correct parties detect the inconsistency and abort. Furthermore, we demonstrate the feasibility of the approach through a proof-of-principle implementation on an IBM quantum processor.**

***Index Terms*—Quantum communications, Quantum error correction, Detectable Byzantine broadcast, Detectable Byzantine consensus, Quantum circuit design, Experiments.**

I. Introduction

Reaching agreement despite faulty components is one of the central challenges in distributed computing [1]–[3]. Classical formulations require that all correct processes agree on a common value despite the presence of malicious or Byzantine participants. These problems are central to the design of reliable distributed services, including replicated databases, blockchain systems, and fault-tolerant control infrastructures.

In distributed computing, the *Consensus Problem* concerns how multiple processes can decide on a common outcome for a computational task. It plays a pivotal role in distributed systems, ensuring that different components can agree on critical values even in the presence of noise, crashes, or adversarial (Byzantine) attacks [3]. Closely related to consensus, the *Byzantine Generals Problem* [1] captures the setting of Byzantine-tolerant broadcast, in which a designated leader sends an initial value; if that leader is correct, every correct process must decide on the same value.

Cheat-sensitive weaker variants of agreement problems are the *detectable* or *abortable* Byzantine consensus and, similarly, the detectable Byzantine broadcast has also been defined and studied as a weaker variant of the Byzantine generals problem [4]–[6]. For instance, in a detectable Byzantine broadcast, a designated sender (leader) distributes a value to all other processes (followers). Upon termination, the correct processes have only two possible outcomes: they decide the proposed value of the leader, or they all abort. Similarly, in detectable consensus either agreement is reached or all correct processes abort. These relaxations weaken classical guarantees on broadcast by allowing to abort the process. However, in return, they enable stronger fault tolerance under weaker assumptions and improve efficiency. This detectability of Byzantine broadcast or consensus captures an important operational requirement in the sense that, when the network cannot guarantee correct agreement due to adversarial behavior, it should fail explicitly rather than silently produce inconsistent outputs.

Considering the significance of the topic, it is not surprising that the Byzantine consensus and broadcast problems are also addressed using quantum genuine resources. Indeed, quantum information processing offers new tools for addressing different variants of consensus problems. Pre-shared entanglement among processes enables detectable Byzantine consensus through violations of quantum correlations using a classical channel [7]. It also reduces the number of communication rounds among processes to detect such inconsistencies, enabling so-called fast Byzantine agreement [8]. However, such approaches assume the existence of entanglement [7]–[12], rather than constructing a quantum state for consensus from an arbitrary initial quantum state of the network. Our goal is to design a quantum algorithm that, combined with

This work was supported by MADQuantum-CM, financed by EU NextGeneration-EU, CM and PRTR; and Grant CEX2024-001471-M funded by MICIU/AEI/10.13039/501100011033.

authenticated classical communication, realizes detectable Byzantine broadcast and consensus *without requiring a pre-shared entanglement*.

A. Related work

Quantum-enhanced technology offers unprecedented solutions to many problems, including computation, sensing, and secure communication. Quantum networks are expected to improve security and enable distributed quantum computing beyond classical capabilities by leveraging quantum resources and primitives such as quantum key distribution, entanglement swapping, and quantum repeaters [13]. Given the importance of consensus in classical networks, it is not surprising that the concept has been extended to quantum networks, motivated by the potential for enhanced performance through quantum computation and communication [14]. Beyond these primitives, there is a growing need to validate quantum computational tasks performed across the network, particularly when computation is distributed. Such validation tasks can be naturally interpreted as instances of consensus or broadcast, where participants must either agree on a common outcome or detect inconsistencies.

Quantum consensus can be defined based on several criteria [15]. Among these, entanglement-based approaches have played a central role in addressing quantum *Byzantine agreement* problems [7]–[12]. These protocols typically assume the availability of a *pre-shared* entangled resource, such as a Greenberger–Horne–Zeilinger (GHZ) state [16], distributed among the participants. The shared entanglement acts as a global correlation primitive, enabling agreement in constant time, often with a single round of classical communication. This constitutes a significant improvement over classical Byzantine protocols, which require multiple communication rounds to tolerate faults.

Moreover, entanglement-based protocols naturally support detection of Byzantine behavior. A deviation from the prescribed protocol leads to violations of the expected quantum correlations, which can be identified by correct parties [7]–[9], [12]. Experimental demonstrations further confirm the feasibility and advantages of these approaches [17], [18]. More broadly, quantum consensus techniques have been applied to distributed quantum computation [19], secure communication [20], and other tasks where quantum correlations provide computational or security benefits.

However, most existing approaches rely on the assumption that a suitable entangled resource is already available. Under this assumption, the protocol effectively *reveals* agreement encoded in the shared state, rather than *constructing* it. As a result, these approaches do not fully capture the operational aspect of consensus as a dynamic process, in which participant parties starting from arbitrary initial states must interact to reach agreement.

Alternative approaches adopt a more operational viewpoint. For example, quantum gossip and averaging protocols aim to drive the network toward consensus through repeated local interactions [21]–[23]. In these models, consensus is typically defined in terms of convergence to a statistical average. While these approaches provide insight into distributed quantum dynamics, they generally require asymptotically many iterations and repeated measurements to verify convergence. Furthermore, they may fail to satisfy classical *validity* requirements, in the sense that the final agreed value does not need to correspond to any of the initial input values, seen as the set containing the only *valid* output values [24].

In contrast to the above approaches, we consider an operational setting in which no pre-shared entanglement is assumed and the initial network state is arbitrary. We design a quantum circuit that steers the network toward a resource state enabling *detectable Byzantine broadcast/consensus*, ensuring that the correct processes either decide the same value or all abort.

B. Novel contributions of the paper

We design quantum circuits and experimentally validate them on an IBM quantum processor to establish a quantum resource state that, when combined with classical communication, is sufficient to realize detectable Byzantine broadcast and consensus. This work makes the following contributions:

- We propose a *leader-follower* quantum circuit for broadcast, in which *all qubit states collapse to the same value* upon measurement and follow the leader's measurement distribution in a pre-agreed basis. The leader-follower broadcast circuit is efficient, requiring only $\mathcal{O}(\log n)$ circuit depth and a linear number of two-qubit operations and ancillary qubits, where n is the number of parties of the system.
- We present a *leader selection* scheme compatible with the leader-follower circuit, designed to establish a single-qubit state based on the majority of the initial state of the network. This scheme uses a Hamming-weight quantum circuit to coherently prepare a single-qubit leader that accurately reflects the majority value in the initial quantum network, after which the leader-follower broadcast circuit is applied. The quantum circuit combining leader selection and leader-follower broadcast implements a consensus circuit in which the majority value is decided.
- We show that combining the leader-follower quantum circuit with synchronous authenticated secure classical channels yields protocols for detectable Byzantine broadcast and consensus for any $f < n$ Byzantine parties. Under correct operation, all correct parties output the same value. The decided value is the leader's value for broadcast and the majority value for consensus. If corrupted parties are present or the quantum circuit deviates from the prescribed behavior, then all correct parties detect the inconsistency and abort.
- We implement detectable broadcast on an IBM quantum computer as a proof of principle experiment.

II. Background and Technical Preliminaries

We now briefly revisit some of the classical and quantum primitives relevant to this work.

A. Detectable Byzantine Broadcast

The *Byzantine generals problem* [1] is a standard abstraction for studying agreement under Byzantine faults. It considers $n > 1$ generals, each commanding a distinct army division positioned around an enemy city, where communication is conducted solely by sending messengers. One distinguished general, the commanding general (leader), chooses a plan of action (attack or withdraw) and communicates it to the other generals. The model allows some generals, possibly even the commander, to be traitors who try to obstruct coordination among the loyal generals. The task is to design a protocol such that, upon termination: (i) every loyal general

reaches agreement on a single shared plan, and (ii) when the commander is loyal, all loyal generals must adopt the commander's proposed plan.

A weaker, yet still practically relevant, variant of this task is called *detectable byzantine agreement* [7]. It is defined as follows.

Definition 1 (Detectable Byzantine Broadcast)**.** *Consider a system of n processes, among which at most $f < n$ may be Byzantine. A distinguished process S, called the sender, has an input value $x_S \in \mathcal{D}$, where $\mathcal{D}$ is a finite domain of possible values. A protocol achieves* detectable Byzantine broadcast *if each correct process eventually produces an output in $\mathcal{D} \cup \{\perp\}$, where the $\perp$ symbol indicates that the correct process aborts. Moreover, the following properties are satisfied:*

- ***Agreement:*** *If two correct processes output values from $\mathcal{D}$, then those values are identical.*
- ***Validity:*** *If all processes are correct, then every process outputs x_S.*
- ***Detectability:*** *If the sender is Byzantine or the execution does not follow the prescribed protocol, then the correct processes either all output the same value $y \in \mathcal{D}$, or they all output $\perp$.*
- ***Termination:*** *Every correct process eventually outputs a value in $\mathcal{D} \cup \{\perp\}$.*

It has been shown that detectable Byzantine broadcast is achievable in three representative scenarios (cf. Theorem 1 in [4]): (1) in an authenticated-channel model with secure digital signatures and polynomial-time adversaries; (2) in a secure-channel model where messages are inaccessible to the adversary, with no computational bound assumed; and (3) in an authenticated-channel model augmented by possibly insecure quantum channels, again with computationally unbounded adversaries.

In the following, we show that an error-correction quantum circuit provides a quantum primitive within scenario (3) to achieve detectable Byzantine broadcast, i.e., it can be used to generate a quantum channel that is potentially insecure, still it allows to detect adversary party's actions in all cases.

B. Detectable Byzantine Consensus

We now consider a different variant of the problem, namely detectable Byzantine consensus. In contrast to a detectable Byzantine broadcast, there is no distinguished commanding general (leader). Instead, each process proposes an input value from a finite domain. The problem is defined formally as follows.

Definition 2 (Detectable Byzantine Consensus)**.** *Consider a system of n processes, among which at most $f < n$ may be Byzantine. Each process i proposes an input $x_i \in \mathcal{D}$, where $\mathcal{D}$ denotes the domain of possible values. A protocol achieves* detectable Byzantine consensus *if every correct process eventually outputs a value in $\mathcal{D} \cup \{\perp\}$. Accordingly, the following properties hold:*

- ***Agreement:*** *If any two correct processes output values in $\mathcal{D}$, then those values are equal.*
- ***Validity:*** *If all processes are correct and start with the same input value $v \in \mathcal{D}$, then every process outputs v.*
- ***Detectability:*** *If the prescribed protocol is not followed, the correct processes terminate either with a common value in $\mathcal{D}$ or with outputs $\perp$.*
- ***Termination:*** *Every correct process eventually outputs either a value in $\mathcal{D}$ or $\perp$.*

C. Classical Operations for Communications

In addition to quantum channels, we also need classical channels for communication to address the two problems defined earlier. Indeed, classical communication becomes indispensable because it allows correct parties to compare views, verify the outcome, and either commit to a common output or abort upon detecting deviations. Accordingly, throughout the paper, we assume access to *synchronous authenticated secure classical channels* between participants (cf. model (2) of Theorem 1 in [4]). The channel is a secure point-to-point link that provides message integrity and origin authentication (e.g., via MAC or authenticated encryption), ensuring that an adversary cannot impersonate a correct process, cannot read the messages sent, and cannot tamper with the traffic.

D. Quantum Operations

Qubits, Gates, and Circuits: A qubit is the fundamental unit of quantum information and is mathematically described by a unit vector in a two-dimensional complex Hilbert space $\mathcal{H}_2 \cong \mathbb{C}^2$ [25]. In the computational basis $\{|0\rangle, |1\rangle\}$, an arbitrary pure single-qubit state can be written as $|\varphi\rangle = \alpha_0 |0\rangle + \alpha_1 |1\rangle$, where α_0 and α_1 are complex probability amplitudes satisfying the normalization condition $\sum_{i=0}^{1} |\alpha_i|^2 = 1$. The squared modulus $\alpha_i \alpha_i^*$ gives the probability of measuring the corresponding basis state, where α_i^* denotes the complex conjugate of α_i. This superposition property enables non-classical phenomena such as entanglement and wavefunction collapse upon measurement. An n-qubit system is described in the tensor-product, denoted by $\otimes$, Hilbert space $\mathcal{H}^{(n)} = \mathcal{H}_2^{\otimes n} \cong \mathbb{C}^{2^n}$. For brevity, the tensor-product basis states are usually written as $|j_1, j_2, \ldots, j_n\rangle \equiv |j_1\rangle \otimes |j_2\rangle \otimes \cdots \otimes |j_n\rangle$. Qubits are realized through various possible physical systems, including superconducting circuits, trapped ions, photonic setups, and spin-based platforms [26].

Quantum gates [25] manipulate quantum states through unitary operations $\hat{U}$, which satisfy $\hat{U}\hat{U}^\dagger = \mathbb{I}$, where $\hat{U}^\dagger$ denotes the conjugate transpose and $\mathbb{I}$ is the identity operator. The action of a gate $\hat{U}$ on an initial state $|\psi_0\rangle$ produces a new state $|\psi'\rangle = \hat{U} |\psi_0\rangle$. These unitary gates serve as the fundamental building blocks of quantum computation. Single-qubit gates such as the Pauli X, Y, and Z gates apply specific bit- and phase-flip transformations. In the computational basis $\{|0\rangle, |1\rangle\}$, these operators are defined as

$$\hat{X} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \tag{1}$$

$$\hat{Y} = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \tag{2}$$

$$\hat{Z} = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}. \tag{3}$$

Two-qubit gates such as the controlled-NOT (CNOT) enable conditional logic. In the computational basis $\{|00\rangle, |01\rangle, |10\rangle, |11\rangle\}$, the CNOT gate is represented as

$$\mathrm{CNOT} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{4}$$

with the first qubit denoting the control and the second the target.

A quantum circuit is the quantum analogue of a classical Boolean circuit, consisting of a sequence of quantum gates applied to qubits over time. In this model, wires represent individual qubits and boxes denote gate operations. The circuit formalism captures the evolution of quantum states and supports complexity analyses in terms of gate count, circuit depth, and qubit number. Unlike classical circuits, quantum circuits exploit superposition and entanglement to perform computations, which may offer significant speedups for certain tasks, such as integer factorization [27].

Quantum Errors and Error Correction: Quantum systems are inherently fragile, as interactions with the environment introduce decoherence and errors that disrupt the intended evolution of quantum states [25], [28]. Common quantum errors include bit-flip and phase-flip operations, which change the computational state or its relative phase, respectively. Quantum error correction (QEC) provides a framework for protecting quantum information from such errors without violating the no-cloning theorem or collapsing the quantum state [25], [29]. This is accomplished by encoding a logical qubit into multiple physical qubits, enabling error detection and correction through carefully designed protocols. Ancillary qubits and unitary gates are used to extract error syndromes —signatures of specific error patterns— without measuring the encoded quantum information directly. While the space of possible errors grows exponentially with the system size, structured QEC schemes such as the Shor code [30], Steane code [31], and surface codes [32] allow correction of the most relevant single- and multi-qubit errors. These error-correcting codes are essential for achieving fault-tolerant quantum computation in large-scale quantum systems.

E. Quantum Network Model

A quantum network is modeled as an n-qubit system, with each qubit corresponding to a spatially distributed node (i.e., a party) that encodes local quantum information, such as sensed data, results from a distributed quantum computation, or data received from other sources. For brevity, we use the term node to refer to both the participant parties and their associated quantum systems (qubits). Collectively, the n-qubit state represents the global state of the network and captures its initial configuration. Nodes interact according to a specified network topology and may exchange information via classical communication channels after performing local single-qubit measurements.

Since nodes are initialized independently, the global network state may be arbitrary and the nodes must inevitably interact so that the global quantum state of the network evolves to a final state. However, since the nodes are spatially distributed, such interaction must be performed non-locally [33] through mediators, which we refer to as ancillary (auxiliary) qubits. In practical quantum network architectures, non-local operations such as CNOT gates between distant nodes are typically implemented using a combination of stationary physical qubits (e.g., trapped ions or superconducting circuits) and flying qubits—usually photons—that carry quantum information across spatial separation to mediate entangling operations. While our focus is algorithmic, similar physical implementations of non-local gates have been proposed and demonstrated in experimental settings (see, e.g., [34] for relevant examples). In our model, we assume that each pair of nodes shares access to an auxiliary qubit system that mediates such interaction. In Fig. 1, we illustrate this architecture between four nodes to reflect the network model and its key components.

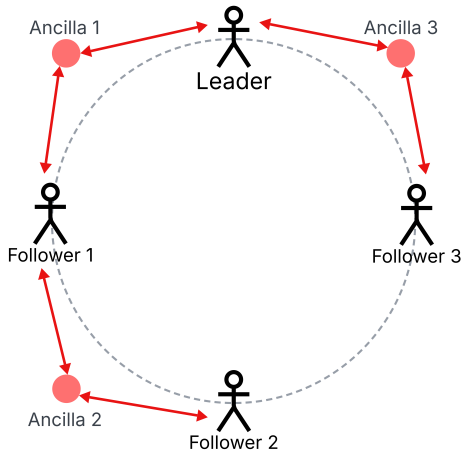

Fig. 1: **Leader-based quantum network for four nodes.** Each node (the leader and three followers) is equipped with a single qubit. Interactions between nodes (indicated by double arrows) are mediated by ancillary qubits.

Under this setting, we now formally define the problem.

F. Problem Formulation

Our objective is to design a protocol that builds, from an arbitrary n-qubit initial state, a quantum resource for an n-node quantum network which, together with authenticated classical channels, enables *detectable Byzantine broadcast* and *detectable Byzantine consensus*, as defined in Definition 1 and Definition 2. We assume that the network starts in an arbitrary pure state

$$|\psi_n\rangle = \sum_{j_1,j_2,\ldots,j_n \in \{0,1\}} \alpha_{j_1,j_2,\ldots,j_n} |j_1, j_2, \ldots, j_n\rangle , \qquad (5)$$

with $\sum_{j_1,\ldots,j_n} |\alpha_{j_1,\ldots,j_n}|^2 = 1.$. Here, $\alpha_{j_1,\ldots,j_n}$ are complex probability amplitudes over the computational-basis states in $\{0,1\}^n$, satisfying the normalization condition $\sum_{j_1,\ldots,j_n} |\alpha_{j_1,\ldots,j_n}|^2 = 1$.

For detectable Byzantine broadcast, a designated qubit serves as the leader, and a leader-follower quantum circuit prepares a resource state such that all correct processes reproduce the leader's measurement outcome in a pre-agreed basis. We then prove that, when combined with classical communication, the protocol achieves detectable Byzantine broadcast.

In the detectable Byzantine consensus, the leader is not predetermined or fixed beforehand. Instead, the quantum state of the leader is a majority rule extracted from the initial network state with a Hamming weight quantum circuit. The same leader-follower circuit then prepares a resource state such that all correct processes reproduce the majority outcome upon measurement. We further prove that, when combined with classical communication, the protocol achieves detectable Byzantine consensus.

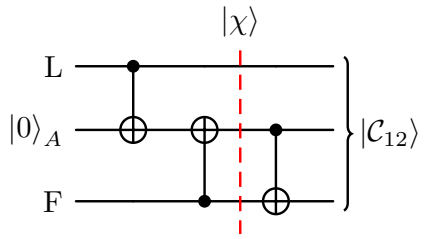

Fig. 2: **Simple QEC circuit.** Each quantum wire represents a single qubit, with the ancillary qubit labeled as $|0\rangle_{\mathrm{A}}$. The first two CNOT gates perform a parity check between the two nodes, while the final CNOT enforces agreement by driving the second qubit (bottom wire, labeled F) to follow the first qubit (top wire, labeled L), which acts as the leader.

G. Power of Byzantine Parties

We assume that up to $f < n$ parties of the system can be Byzantine, and can collude to prevent the detectable Byzantine broadcast or detectable Byzantine consensus to be achieved. However, it is important to explicitly describe what these parties can do.

As described above, we assume that each party has a qubit in an initial pure state. This is the input to a quantum circuit, and at the output of the circuit a measurement of the qubit is performed. Finally, a classical communication protocol is executed to complete the process. We assume that a Byzantine party can deviate from this process at any point. I.e, it can perturb the initial quantum state of its qubit, it can perturb the quantum computations in which its qubit is involved, and it can deviate from the final classical protocol arbitrarily.

III. Main Results

Before presenting the full quantum-circuit construction, we first consider the two-qubit case and show how a simple Quantum Error Correction (QEC) circuit, consisting of a sequence of three CNOT gates serves as the fundamental building block. In particular, we rely on bit-flip QEC, as described in what follows. Let the initial state of the two-qubit system be an arbitrary pure state in the computational basis, $|\varphi\rangle = \alpha_{00}|0,0\rangle + \alpha_{01}|0,1\rangle + \alpha_{10}|1,0\rangle + \alpha_{11}|1,1\rangle$ where α_{ij}, with $i,j \in \{0,1\}$, are complex probability amplitudes satisfying $\sum_{i,j}|\alpha_{ij}|^2 = 1$. In general, the state $|\varphi\rangle$ cannot be written as a product of two single-qubit states, i.e., $|\varphi\rangle \neq |\varphi_1\rangle \otimes |\varphi_2\rangle$, and may therefore be entangled.

We introduce an ancillary qubit initialized in the state $|0\rangle_{\mathrm{A}}$. As illustrated in Fig. 2, the first two CNOT gates are applied with the leader and follower qubits as control qubits and the ancillary qubit as the target. This configuration performs a parity check [25], namely, it determines whether the two data qubits are in the same computational-basis state, without collapsing their superposition. More explicitly, the action of the first two CNOT gates on the computational-basis states is

$$\begin{aligned} |0,0\rangle|0\rangle_{\mathrm{A}} &\mapsto |0,0\rangle|0\rangle_{\mathrm{A}}, & |0,1\rangle|0\rangle_{\mathrm{A}} &\mapsto |0,1\rangle|1\rangle_{\mathrm{A}}, \\ |1,0\rangle|0\rangle_{\mathrm{A}} &\mapsto |1,0\rangle|1\rangle_{\mathrm{A}}, & |1,1\rangle|0\rangle_{\mathrm{A}} &\mapsto |1,1\rangle|0\rangle_{\mathrm{A}}. \end{aligned} \tag{6}$$

By linearity, the initial state $|\varphi\rangle \otimes |0\rangle_{\mathrm{A}}$ is therefore transformed into

$$\begin{aligned} |\chi\rangle = &(\alpha_{00}|0,0\rangle + \alpha_{11}|1,1\rangle) \otimes |0\rangle_{\mathrm{A}} \\ &+ (\alpha_{01}|0,1\rangle + \alpha_{10}|1,0\rangle) \otimes |1\rangle_{\mathrm{A}}. \end{aligned} \tag{7}$$

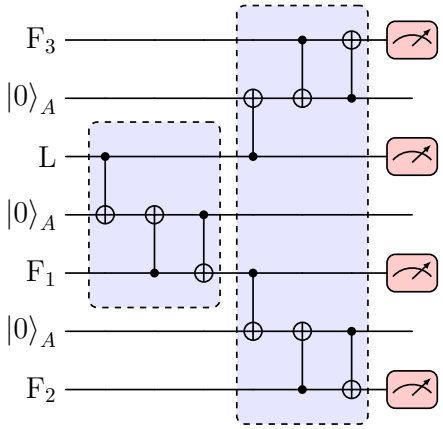

Fig. 3: **Leader-follower quantum circuit.** The leader qubit (labeled as the L wire) first drives the initial follower qubit (labeled as the F_1 to F_3 wire). In the second layer, the leader and first follower jointly drive two additional nodes toward agreement.

Eq. (7) shows that whenever the two data qubits agree, the ancillary qubit is found in the state $|0\rangle_{\mathrm{A}}$, whereas disagreement is encoded in the state $|1\rangle_{\mathrm{A}}$. In this sense, the ancillary qubit stores a syndrome analogous to a bit-flip error syndrome in quantum error correction [25], identifying whether the two qubits are consistent or not.

The third CNOT gate, shown in Fig. 2, uses the ancillary qubit as the control and the follower qubit as the target. Its action is to flip the follower qubit precisely when the ancillary qubit indicates disagreement. Applying this gate to Eq. (7) yields

$$\begin{aligned} |\mathcal{C}_{12}\rangle = &|0,0\rangle \otimes (\alpha_{00}|0\rangle_A + \alpha_{01}|1\rangle_A) \\ &+ |1,1\rangle \otimes (\alpha_{11}|0\rangle_A + \alpha_{10}|1\rangle_A). \end{aligned} \tag{8}$$

Hence, the two data qubits are always found either in the joint state $|0,0\rangle$ or in the joint state $|1,1\rangle$. In particular, if the leader qubit is measured in the computational basis and yields the outcome $|0\rangle$, then the follower qubit is deterministically found in the state $|0\rangle$ as well; similarly, if the leader yields $|1\rangle$, then the follower is also found in $|1\rangle$. Once this condition is met, the two parties can use classical communication to verify whether they obtained the same value.

It is also important to note that the probability of obtaining the outcome $|0,0\rangle$ is $P(|0\rangle_{\mathrm{L}}) = \sum_{i=0,1}|\alpha_{0i}|^2$, which is exactly the probability that the initial leader qubit is measured in the state $|0\rangle$. Therefore, the circuit preserves the leader's measurement distribution while forcing the follower to reproduce the leader's measurement outcome. This two-qubit construction provides the basic quantum resource used throughout the rest of the paper.

Below, we extend this building block to design a leader-follower quantum circuit for an n-qubit system with logarithmic depth.

A. Leader-follower quantum circuit

At a high level, the leader-follower quantum circuit drives the entire n-qubit network toward leader-aligned agreement, under the assumption of a designated leader. The circuit follows a layered architecture. In the first layer, the leader qubit interacts with a single follower qubit, establishing agreement between the two nodes. In the next layer, both

the leader and the first follower interact with two additional follower qubits, thereby propagating the leader's measurement outcome further across the network. After two layers, as illustrated in Fig. 3, four qubits have been aligned, allowing them to influence four more follower qubits in the next stage. This high-level mechanism is formally captured in the following lemma.

Lemma 1. *Assume that a leader node has been designated and, without loss of generality, is associated with the first qubit. Let the network initially be in the arbitrary pure state $|\psi_n\rangle$ of Eq.* (5)*, and append $n-1$ ancillary qubits prepared in the state $|0\rangle_A^{\otimes(n-1)}$. Then the leader–follower quantum circuit maps $|\psi_n\rangle \otimes |0\rangle_A^{\otimes(n-1)}$ to*

$$|\mathcal{C}_n\rangle = |0,0,\ldots,0\rangle \otimes |\xi_0\rangle_A + |1,1,\ldots,1\rangle \otimes |\xi_1\rangle_A\,, \quad (9)$$

where

$$\begin{aligned} |\xi_0\rangle_A &= \sum_{k_2,\ldots,k_n\in\{0,1\}} \alpha_{0,k_2,\ldots,k_n}\,|k_2,\ldots,k_n\rangle\,, \\ |\xi_1\rangle_A &= \sum_{k_2,\ldots,k_n\in\{0,1\}} \alpha_{1,k_2,\ldots,k_n}\,|1-k_2,\ldots,1-k_n\rangle\,, \end{aligned} \quad (10)$$

are unnormalized quantum states defined over the $n-1$ ancillary qubits. Consequently, a measurement of the network qubits in the computational basis yields a common bit value at all nodes, which means either all nodes obtain 0 or all nodes obtain 1. The probability of obtaining the all-0 outcome is exactly the probability that the leader qubit is in 0 initial state. Therefore, the circuit prepares the quantum resource state required for the detectable Byzantine broadcast. Moreover, the circuit has depth $\mathcal{O}(\log n)$, uses $3(n-1)$ CNOT gates, and requires $n-1$ ancillary qubits.

Proof. Since we have already shown how the minimal two-qubit instance circuit, verifying the correctness of the general quantum circuit is straightforward. In the leader-based construction, a qubit is designated as the leader, which in the first layer shares its state with a follower via CNOT gates that form a basic QEC block. In the second layer, both the leader and the first follower act as sources and share the state with two additional qubits. This pattern continues such that in each layer k, the circuit propagates agreement to up to 2^{k-1} new qubits.

The action of the first block (the first three CNOT gates in Fig. 3) evolves the initial state in Eq. (5) to $|\varphi_n'\rangle = |0,0\rangle \otimes \sum_{j_2,\ldots,j_n=\{0,1\}} \alpha_{0,\ldots,j_n}\,|j_3,\ldots,j_n\rangle \otimes |j_2,0,\ldots,0\rangle_A + |1,1\rangle \otimes \sum_{j_2,\ldots,j_n=\{0,1\}} \alpha_{1,\ldots,j_n}\,|j_3,\ldots,j_n\rangle \otimes |1-j_2,0,\ldots,0\rangle_A$. The first layer drives the first follower qubit into agreement with the leader qubit. Next, the second block (the following set of CNOT gates, as shown in Fig. 3), forces two additional follower qubits (one interacting with the leader and the other with the previous follower) into agreement, yielding $|\varphi_n''\rangle = |0,0,0,0\rangle \otimes \sum_{j_2,\ldots,j_n=\{0,1\}} \alpha_{0,\ldots,j_n}\,|j_5,\ldots,j_n\rangle \otimes |j_2,\ldots,j_4,0,\ldots,0\rangle_A + |1,1,1,1\rangle \otimes \sum_{j_2,\ldots,j_n=\{0,1\}} \alpha_{1,\ldots,j_n}\,|j_5,\ldots,j_n\rangle \otimes |1-j_2,\ldots,1-j_4,0,\ldots,0\rangle_A$. This construction naturally continues to n nodes, yielding the state expressions given in Eqs. (9) and (10). □

Remark. It is worth mentioning that although the leader-based circuit can theoretically achieve $\mathcal{O}(\log n)$ circuit depth, it remains adaptable to other computer architecture topology, achieving at worst $\mathcal{O}(n)$ circuit depth complexity in an alternative linear interaction topology, as illustrated in Fig 4. The linear construction of Fig 4 results in the same initial state but agreement is between qubits that are adjacent in the linear topology, one after the other, hence it takes linear circuit depth (i.e., no QEC block can be run in parallel with another). The probability of obtaining the all-zero outcome $|0,\ldots,0\rangle$ is

$$P_0 = \sum_{k_1,\ldots,k_{n-1}\in\{0,1\}} \left|\alpha_{0,k_1,\ldots,k_{n-1}}\right|^2, \quad (11)$$

which equals the probability that the leader qubit is initially in the state $|0\rangle$. The state in Eq. (9) serves as the quantum resource that, together with authenticated classical verification, enables detectable Byzantine broadcast. We formalize this high-level mechanism later.

However, let us first discuss the leader-selection procedure that allows establishing the detectable Byzantine consensus when no predefined leader is available.

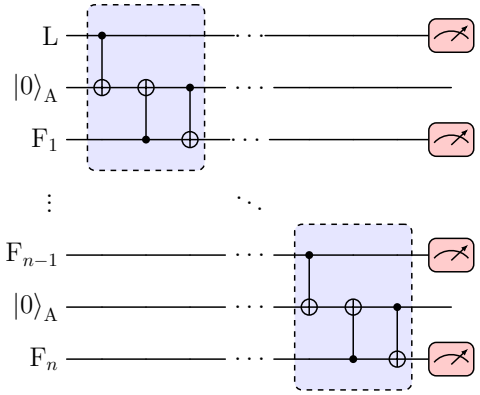

Fig. 4: **Leader–follower quantum circuit with linear topology.** The leader qubit (wire L) first drives the follower qubit F_1. Then F_1 becomes the new leader and drives the next follower. This process continues sequentially along the chain until the nth follower is reached.

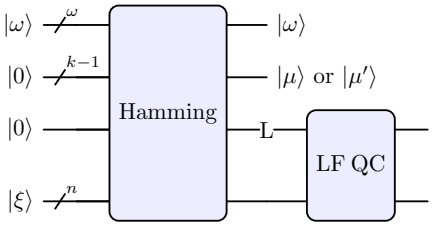

Fig. 5: **Leader selection via Hamming weight.** Bundle wires denote multi-qubit registers; the label above each bundle indicates its size. The Hamming circuit acts on the padding register $|\omega\rangle$ and a k-qubit Hamming-weight register initialized to $|0\rangle^{\otimes k}$. We take the kth weight qubit as the leader qubit, labeled L. The leader qubit L and the network input state $|\xi\rangle = \sum_{j_1,\ldots,j_n\in\{0,1\}} \alpha_{j_1,\ldots,j_n}\,|j_1,\ldots,j_n\rangle$ then enter the leader–follower circuit (LF QC).

B. Leader selection based on majority

We now describe how, algorithmically, a leader can be selected based on a majority rule. To prevent ties, we assume

that the number of parties in the network is odd[1]. This task involves preparing a single qubit—referred to as the *leader qubit*—whose state reflects the majority value (either $|0\rangle$ or $|1\rangle$) among the n-qubit initial state of the network. One possible way to do this is by calculating the *Hamming weight* of the input, defined as the number of qubits in the $|1\rangle$ state. Then, the highest order qubit from the Hamming-weight register can reflect the majority and serve as a leader, as formalized in the following lemma.

Lemma 2. *Let the initial state of an n-qubit quantum network be an arbitrary pure state, represented as:*

$$\sum_{j_1,\ldots,j_n\in\{0,1\}} \alpha_{j_1,\ldots,j_n} |j_1,\ldots,j_n\rangle \otimes |\omega\rangle \otimes |0\rangle^{\otimes k}, \quad (12)$$

where $k = \lceil \log_2(n+1) \rceil$ is the number of ancillary qubits used to store the binary representation of the Hamming weight, and $|\omega\rangle = |0\rangle^{\otimes w/2} \otimes |1\rangle^{\otimes w/2}$ is a padding register of size $w = 2^k - (n+1)$, used to preserve the original majority distribution during encoding. Then, there exists a quantum circuit that coherently computes the Hamming weight of the n-qubit state and encodes it into the ancillary register such that:

$$\sum_{j_1,\ldots,j_n\in\{0,1\}} \alpha_{j_1,\ldots,j_n} |j_1,\ldots,j_n\rangle \otimes |\omega\rangle \otimes |\mu,0\rangle, \text{ if } |J| \le \left\lfloor \frac{n}{2} \right\rfloor,$$
$$\sum_{j_1\ldots,j_n\in\{0,1\}} \alpha_{j_1,\ldots,j_n} |j_1,\ldots,j_n\rangle \otimes |\omega\rangle \otimes |\mu',1\rangle, \text{ if } |J| > \left\lfloor \frac{n}{2} \right\rfloor, \quad (13)$$

where $|J|$ denotes the Hamming weight of the n-qubit state vector $|j_1,\ldots,j_n\rangle$, and $|\mu\rangle$ and $|\mu'\rangle$ are the binary encodings of the Hamming weights on the first $k-1$ ancillary qubits. The highest-order qubit (i.e., the kth qubit) of the Hamming-weight register is designated as the leader qubit*, as its probability of being measured in state $|0\rangle$ or $|1\rangle$ corresponds directly to whether the majority class of the initial quantum state is composed of $|0\rangle$s or $|1\rangle$s, respectively.*

Proof. The correctness of Lemma 2 follows from existing quantum circuit constructions presented in Refs. [35], [36], which demonstrate how to coherently compute the Hamming weight of an n-qubit input state using either logarithmic or linear-depth quantum circuits.

Our contribution extends these constructions by interpreting the k-th qubit of the Hamming-weight register as the *leader qubit*, and by introducing a specific padding register of size $w = 2^k - (n+1)$. This padding ensures that the Hamming-weight register correctly encodes the majority class of the input distribution in its k-th qubit. Specifically, the probability that the leader qubit is in state $|0\rangle$—indicating that the majority of the input qubits are in state $|0\rangle$—is given by:

$$P_{|0\rangle_k} = \sum_{\substack{j_1,\ldots,j_n\in\{0,1\} \\ |J|\le\lfloor n/2\rfloor}} |\alpha_{j_1,\ldots,j_n}|^2. \quad (14)$$

This encoding enables direct use of the leader qubit in the leader-follower circuit, as described in Lemma 1, to generate resource state among the n qubits. □

[1]Generalization to an even number of parties is possible, but requires a pre-agreed method for tie-breaking. For instance, by adding an extra qubit to the input in the $|+\rangle$ state, so that ties are resolved by creating a superposition of the possible states.

Remark. The Hamming-weight circuit encodes basically counts qubits in state $|1\rangle$ in a binary system, and encodes the most significant qubit of the count as a qubit that is in state $|0\rangle$ (resp. $|1\rangle$) with the same probability that the majority of input qubits be $|0\rangle$ (resp. $|1\rangle$). This encoded qubit can be used as leader in the leader-follower circuit described before for the detectable Byzantine broadcast problem, where the followers are all the original qubits used as input of the Hamming-weight encoder, as shown in Fig. 5.

The state in Eq. (9), prepared either with a fixed leader or with a majority-based leader selected as in Lemma 2, together with authenticated classical verification, enables detectable Byzantine broadcast and detectable Byzantine consensus, respectively. This high-level concept is formalized in the following theorem.

Theorem 1. *For any pure state on an n-qubit network, the leader–follower quantum circuit, combined with classical communication over authenticated classical channels, solves detectable Byzantine broadcast when the leader is fixed in advance and solves detectable Byzantine consensus when the leader is chosen by majority vote.*

Proof. For detectable Byzantine broadcast, Lemma 1 shows that the circuit prepares the resource state in Eq. (9) such that all correct processes obtain the same bit upon measurement. To do so, a protocol over the synchronous authenticated secure classical channels can be used. With this protocol, parties compare their measurement outcomes, so that any inconsistency among the authenticated reports triggers the protocol's abort rule. Hence, if a Byzantine node sends inconsistent classical reports, or if a Byzantine deviation from the prescribed quantum circuit produces outcomes inconsistent with the required resource state Eq. (9), the protocol aborts. Otherwise, all correct processes accept a value that was proposed by the leader. Examples of such protocols are given in [4] that require $\mathcal{O}(f)$ rounds and polynomial bits exchanged. Since the protocol uses a fixed finite number of quantum gates in Lemma 1 and finite rounds of classical communications, every correct process terminates. Therefore, the protocol satisfies agreement, validity, detectability, and termination for detectable Byzantine broadcast in Definition 1.

For detectable Byzantine agreement, Lemma 2 shows that the Hamming-weight circuit prepares the quantum state associated with the leader selected by the prescribed majority rule. The quantum state associated with this majority-prepared leader is then used as input to the leader–follower circuit to prepare the resource state in Eq. (9). The parties then compare their measurement outcomes over synchronous authenticated secure classical channels as described above. Hence, if a Byzantine node sends inconsistent classical reports, or if a Byzantine deviation from the prescribed quantum circuit produces outcomes inconsistent with the required resource state Eq. (9), the protocol aborts. Likewise, if a Byzantine node deviates from the prescribed quantum circuit, the required resource state is not established. Consequently, the subsequent classical verification fails, and all correct processes abort. Otherwise, all correct processes accept the common value determined by the majority-selected leader. In particular, for validity, if all correct processes start with the same input, $|0\rangle$ or $|1\rangle$, and no abort occurs, then every correct process outputs

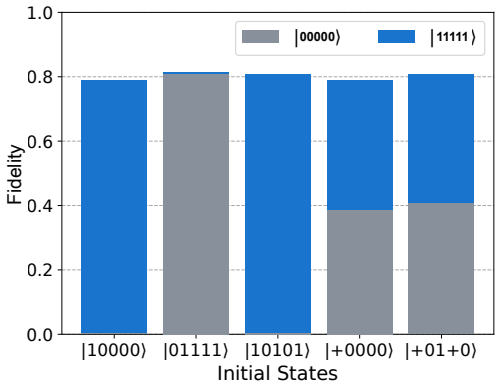

Fig. 6: Fidelity for the 5-qubit leader-follower circuit, tested with five different initial states, of which the leader is the left-most qubit. Each state is shown in a different color, where the sum of the two is the total fidelity.

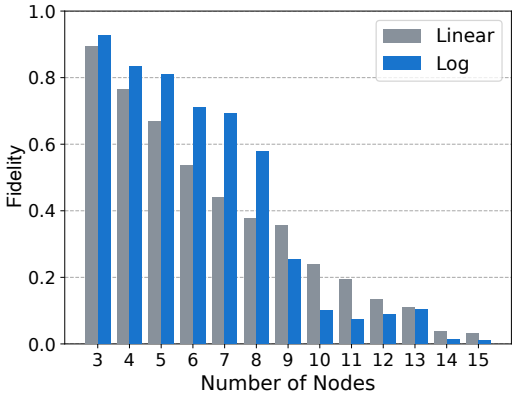

Fig. 7: Fidelities for the leader-follower circuits from 3 up to 15 nodes, for both the logarithmic and liner depth's versions in quantum circuit 5 and Fig 4.

the corresponding classical bit, 0 or 1. Since the protocol uses a fixed finite number of quantum gates in Lemmas 1 and 2 and finite rounds of classical communications, therefore, the protocol satisfies agreement, validity, detectability, and termination for detectable Byzantine agreement, as stated in Definition 2. □

IV. Experimental Results

We present a proof-of-principle experimental demonstration of both the leader–follower circuits of Fig. 3 ("Log" depth case) and Fig. 4 ("Linear" depth case) on IBM superconducting quantum hardware. We note that, although the proposed protocol is formulated for a distributed quantum-network setting, the experimental validation was carried out on a single quantum computer rather than on physically separated network nodes. The experiments should therefore be interpreted as a proof-of-principle demonstration of the circuit construction, rather than as a full distributed implementation. Even so, the logical structure of the circuit remains directly relevant to a genuine quantum-network setting, in which the corresponding non-local operations would be implemented through distributed quantum communication resources. As a minimal testbed, we evaluate systems with three qubits, using multiple initial states as shown in Fig. 6. As a figure of merit, we define the *fidelity* as the total probability of measuring either $|0\rangle^{\otimes n}$ or $|1\rangle^{\otimes n}$ at the circuit output. Each circuit instance is executed for 5,000 shots to estimate these probabilities reliably.

The leader-based circuit was firstly tested for the five-qubit system, achieving over 79% fidelity across various initial states, as shown in Fig. 6. Fidelity is below 100% even at this scale because, in a non-ideal open quantum system, the realized evolution deviates from the intended unitary due to hardware noise and imperfections [28]. On IBM devices, the dominant contributors include imperfect single- and two-qubit gate implementations, decoherence during gate execution and idle periods (finite T_1/T_2), and state-preparation-and-measurement (SPAM) errors (imperfect initialization and readout) [37]. These effects accumulate over the circuit duration and directly reduce the probability mass concentrated on the two possible outcomes.

To further investigate the scalability of the leader-based protocol, we extended our experiments to systems with up to 15 qubits, evaluating both the linear-depth and logarithmic-depth implementations of the leader-follower circuit. For consistency, we used a common initial state in which the leader qubit is initialized to $|1\rangle$ and all follower qubits to $|0\rangle$. The results, shown in Fig. 7, reveal a general decline in fidelity as the number of qubits increases. This degradation is expected on IBM superconducting hardware because non-idealities accumulate with circuit size. Indeed, additional qubits increase the number of gate operations (particularly two-qubit entangling gates, which typically dominate the error budget), and extend the total circuit duration.

Notably, the logarithmic-depth circuit outperforms the linear-depth variant at small system sizes because its shorter depth reduces the time exposed to decoherence and uses fewer sequential error opportunities. As n grows, however, the logarithmic-depth construction relies on long-range interactions. On IBM hardware, where two-qubit gates are available only along edges of a sparse native connectivity graph, these non-local interactions must be realized via transpiler-inserted routing (typically swap chains), which increases both the number of two-qubit gates and the effective depth. Since two-qubit operations dominate the error budget, this routing overhead grows quickly with n and can overwhelm the nominal depth advantage of the logarithmic circuit, leading to a faster fidelity decline at larger scales. In contrast, the linear-depth implementation primarily uses nearest-neighbor interactions and can be mapped more directly onto the native topology, requiring little routing. Consequently, its gate count and depth grow more predictably, and its fidelity degrades more gradually as the number of qubits increases.

Finally, to evaluate the resilience of the protocol under faults, we conducted additional experiments using the leader-follower circuit on both three- and five-qubit systems. For each experiment, a specific initial state was selected, and following the operation, a single qubit—the second qubit, representing the faults node—was artificially flipped (applying an X gate). The results indicate that, despite this adversarial intervention, the correct processes maintained a fidelity of 84% in the three-qubit circuit and 74% in the five-qubit circuit. In contrast, the probability of observing full detectable Byzantine broadcast across all parties dropped to less than 5% and 2%, respectively. These findings demonstrate that the effects of a Byzantine node can be distinguished from

the intrinsic noise of the quantum hardware. However, given the current limitations of quantum technologies—which we have clearly observed in our experiments—the compound operation of quantum and classical methods is key. Indeed, as we shown in this paper, it is not enough to prepare a quantum state on top of which to run detectable Byzantine broadcast or consensus, because one needs to be able to detect both Byzantine adversaries as well as deviations due to quantum noise.

V. Conclusions and perspectives

In conclusion, we introduced a quantum-circuit framework for realizing detectable Byzantine broadcast and detectable Byzantine consensus from arbitrary initial network states, using simple bit-flip QEC and Hamming-weight building blocks in the quantum domain, plus classical authenticated classical communications. In practice, we demonstrate that quantum resonance states enable all correct parties to either accept a common value (proposed by a leader or resulting from a majority) upon measurement and classical communication of the measured values, or abort upon detecting an inconsistency. The proposed QEC-based leader–follower construction is efficient and versatile, as it requires only logarithmic depth and linear ancillary and gate resources, and can be also used as consensus block in combination with a majority-based leader selection mechanism obtained through coherent computation of Hamming weights, achieving detectable Byzantine consensus.

We also validated the practicality of the proposed method through a proof-of-principle implementation on an IBM quantum processor. These results show that quantum circuits can provide the essential quantum resource for detectable Byzantine broadcast and detectable Byzantine consensus, opening a concrete path toward broadcast/consensus primitives in future distributed quantum networks.

As part of future work, it would be important to develop an explicit error model to assess whether the proposed protocol preserves its guarantees in the presence of failures, noise, and imperfect operations. Such an analysis would provide a clearer understanding of the operating regimes in which the protocol remains reliable and of the additional fault-tolerant techniques that may be required for practical deployment. Another promising direction is to extend the framework to *quantum conference key establishment* and the security analysis under various adversarial models.

Although in this paper we adopted simple bit-flip QEC primitives, in prospect, one may consider using more sophisticated error-correcting techniques, such as surface codes or other topological quantum codes, to further enhance the fault tolerance of circuits and enable scalability across larger quantum networks. Experimental realizations based on spatially separated quantum systems — particularly hybrid architectures combining photonic platforms with ion traps or superconducting qubits — offer promising avenues for implementing these protocols in practical settings.

References

[1] L. Lamport, R. Shostak, and M. Pease, "The byzantine generals problem," *ACM Transactions on Programming Languages and Systems*, vol. 4, no. 3, pp. 382–401, 1982.

[2] M. C. Pease, R. E. Shostak, and L. Lamport, "Reaching agreement in the presence of faults," *J. ACM*, vol. 27, no. 2, pp. 228–234, 1980. [Online]. Available: https://doi.org/10.1145/322186.322188

[3] M. J. Fischer, "The consensus problem in unreliable distributed systems (a brief survey)," in *International conference on fundamentals of computation theory*. Springer, 1983, pp. 127–140.

[4] M. Fitzi, D. Gottesman, M. Hirt, T. Holenstein, and A. Smith, "Detectable byzantine agreement secure against faulty majorities," in *Proceedings of the twenty-first annual symposium on Principles of distributed computing*, 2002, pp. 118–126.

[5] M. Fitzi, N. Gisin, U. M. Maurer, and O. von Rotz, "Unconditional byzantine agreement and multi-party computation secure against dishonest minorities from scratch," in *Advances in Cryptology - EUROCRYPT 2002, International Conference on the Theory and Applications of Cryptographic Techniques, Amsterdam, The Netherlands, April 28 - May 2, 2002, Proceedings*, ser. Lecture Notes in Computer Science, L. R. Knudsen, Ed. Springer, 2002, pp. 482–501. [Online]. Available: https://doi.org/10.1007/3-540-46035-7_32

[6] W. Chen, "Abortable consensus and its application to probabilistic atomic broadcast," Microsoft Research Asia, Beijing, China, Tech. Rep. MSR-TR-2006-135, Tech. Rep., 2007.

[7] M. Fitzi, N. Gisin, and U. Maurer, "Quantum solution to the byzantine agreement problem," *Physical Review Letters*, vol. 87, no. 21, p. 217901, 2001.

[8] M. Ben-Or and A. Hassidim, "Fast quantum byzantine agreement," in *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, 2005, pp. 481–485.

[9] R. Neigovzen, C. Rodó, G. Adesso, and A. Sanpera, "Multipartite continuous-variable solution for the byzantine agreement problem," *Phys. Rev. A*, vol. 77, p. 062307, Jun 2008.

[10] R. Rahaman, M. Wieśniak, and M. Żukowski, "Quantum byzantine agreement via hardy correlations and entanglement swapping," *Physical Review A*, vol. 92, no. 4, p. 042302, 2015.

[11] V. Cholvi, "Quantum byzantine agreement for any number of dishonest parties," *Quantum Information Processing*, vol. 21, no. 4, p. 151, 2022.

[12] L. Li, X. Sun, and J. Zhu, "Quantum byzantine agreement against full-information adversary," in *38th International Symposium on Distributed Computing (DISC 2024)*. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2024, pp. 32–1.

[13] R. Van Meter, *Quantum networking*. John Wiley & Sons, 2014.

[14] A. J. Daley, I. Bloch, C. Kokail, S. Flannigan, N. Pearson, M. Troyer, and P. Zoller, "Practical quantum advantage in quantum simulation," *Nature*, vol. 607, no. 7920, pp. 667–676, 2022.

[15] M. Marcozzi and L. Mostarda, "Quantum consensus: an overview," *arXiv preprint arXiv:2101.04192*, 2021.

[16] 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*. Springer, 1989, pp. 69–72.

[17] S. Gaertner, M. Bourennane, C. Kurtsiefer, A. Cabello, and H. Weinfurter, "Experimental demonstration of a quantum protocol for byzantine agreement and liar detection," *Phys. Rev. Lett.*, vol. 100, p. 070504, Feb 2008.

[18] X. Jing, C. Qian, C.-X. Weng, B.-H. Li, Z. Chen, C.-Q. Wang, J. Tang, X.-W. Gu, Y.-C. Kong, T.-S. Chen *et al.*, "Experimental quantum byzantine agreement on a three-user quantum network with integrated photonics," *Science Advances*, vol. 10, no. 34, p. eadp2877, 2024.

[19] L. Jiang, J. M. Taylor, A. S. Sørensen, and M. D. Lukin, "Distributed quantum computation based on small quantum registers," *Phys. Rev. A*, vol. 76, p. 062323, Dec 2007.

[20] N. Gisin, G. Ribordy, W. Tittel, and H. Zbinden, "Quantum cryptography," *Rev. Mod. Phys.*, vol. 74, pp. 145–195, Mar 2002.

[21] L. Mazzarella, A. Sarlette, and F. Ticozzi, "Consensus for quantum networks: Symmetry from gossip interactions," *IEEE Transactions on Automatic Control*, vol. 60, no. 1, pp. 158–172, 2014.

[22] G. Shi, D. Dong, I. R. Petersen, and K. H. Johansson, "Reaching a quantum consensus: Master equations that generate symmetrization and synchronization," *IEEE Transactions on Automatic Control*, vol. 61, no. 2, pp. 374–387, 2015.

[23] X. Jin, Z. Cao, Y. Tang, and J. Kurths, "Distributed partial quantum consensus of qubit networks with connected topologies," *IEEE Transactions on Cybernetics*, vol. 54, no. 9, pp. 4986–4997, 2024.

[24] W. Golab and H. Tan, "A closer look at quantum distributed consensus," in *Proceedings of the 32nd ACM Symposium on Parallelism in Algorithms and Architectures*, 2020, pp. 539–541.

[25] M. A. Nielsen and I. L. Chuang, *Quantum computation and quantum information*. Cambridge university press, 2010.

[26] T. D. Ladd, F. Jelezko, R. Laflamme, Y. Nakamura, C. Monroe, and J. L. O'Brien, "Quantum computers," *nature*, vol. 464, no. 7285, pp. 45–53, 2010.

[27] P. W. Shor, "Algorithms for quantum computation: discrete logarithms and factoring," in *Proceedings 35th annual symposium on foundations of computer science*. Ieee, 1994, pp. 124–134.

[28] H.-P. Breuer and F. Petruccione, *The theory of open quantum systems*. OUP Oxford, 2002.

[29] D. Gottesman, "An introduction to quantum error correction," in *Proceedings of Symposia in Applied Mathematics*, vol. 58, 2002, pp. 221–236.
[30] P. W. Shor, "Scheme for reducing decoherence in quantum computer memory," *Physical review A*, vol. 52, no. 4, p. R2493, 1995.
[31] A. Steane, "Multiple-particle interference and quantum error correction," *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, vol. 452, no. 1954, pp. 2551–2577, 1996.
[32] E. Dennis, A. Kitaev, A. Landahl, and J. Preskill, "Topological quantum memory," *Journal of Mathematical Physics*, vol. 43, no. 9, pp. 4452–4505, 2002.
[33] D. Gottesman and I. L. Chuang, "Demonstrating the viability of universal quantum computation using teleportation and single-qubit operations," *Nature*, vol. 402, no. 6760, pp. 390–393, 1999.
[34] S. Daiss, S. Langenfeld, S. Welte, E. Distante, P. Thomas, L. Hartung, O. Morin, and G. Rempe, "A quantum-logic gate between distant quantum-network modules," *Science*, vol. 371, no. 6529, pp. 614–617, 2021.
[35] S. Rethinasamy, M. L. LaBorde, and M. M. Wilde, "Logarithmic-depth quantum circuits for hamming weight projections," *Physical Review A*, vol. 110, no. 5, p. 052401, 2024.
[36] M. Plesch and V. Bužek, "Efficient compression of quantum information," *Physical Review A—Atomic, Molecular, and Optical Physics*, vol. 81, no. 3, p. 032317, 2010.
[37] IBM Quantum, "Qpu information," IBM Quantum Documentation, 2024, accessed 2026-01-13.