



HAL
open science

Anonymous and private parameter estimation in networks of quantum sensors

Jarn de Jong, Santiago Scheiner, Naomi R Solomons, Ziad Chaoui, Damian Markham, Anna Pappa

► To cite this version:

Jarn de Jong, Santiago Scheiner, Naomi R Solomons, Ziad Chaoui, Damian Markham, et al.. Anonymous and private parameter estimation in networks of quantum sensors. *Physical Review Applied*, 2025, 24 (5), pp.054053. <10.1103/lbfk-cykl>. <hal-05458726>

HAL Id: hal-05458726

<https://hal.science/hal-05458726v1>

Submitted on 14 Jan 2026

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY-SA 4.0 - Attribution - ShareAlike - International License

Anonymous and private parameter estimation in networks of quantum sensors

Jarn de Jong,¹ Santiago Scheiner,² Naomi R. Solomons,² Ziad Chaoui,¹ Damian Markham,² and Anna Pappa¹

¹*Electrical Engineering and Computer Science Department, Technische Universität Berlin, Einsteinufer 17, 10587 Berlin, Germany*

²*LIP6, CNRS, Sorbonne Université, 4 place Jussieu, F-75005 Paris, France*

Anonymity and privacy are two key properties of modern communication networks. In quantum networks, distributed quantum sensing has emerged as a powerful use case, with applications to clock synchronisation, detecting gravitational effects and more. In this work, we develop a new protocol that, for the first time, combines the different cryptographic properties of anonymity and privacy for the task of distributed parameter estimation. That is, we present a protocol that allows a selected subset of network participants to anonymously collaborate in estimating the average of their private parameters. Crucially, this is achieved without disclosing either the individual parameter values or the identities of the participants, neither to each other nor to the broader network.

I. INTRODUCTION

The rapid advancement of quantum technologies promises a new era of secure communication, computing, data processing and sensing. In particular, quantum sensing allows for exceedingly precise and exact measurements, and has applications in civil engineering [1] as well as for clock synchronisation [2–4] and phase estimation [5–7]. Another cornerstone in modern quantum technologies is the development of quantum networks, i.e., infrastructures that leverage the unique properties of quantum mechanics to enable transmission of information with higher security or efficiency than their classical counterparts [8–10]. These quantum networks can be seen as part of the effort to design a future quantum Internet that promises unparalleled levels of robustness and security [11].

An interesting combination of the above leads to a network of spatially distributed agents, each holding a quantum sensor that measures some local parameter. In this setting, the agents aim to estimate a linear function over their parameters, commonly referred to as distributed parameter estimation. As in other networking tasks, malicious adversaries form an unavoidable threat that needs to be counteracted. In many applications (particularly in military, medical or even agricultural settings [12]) the network should be safeguarded against eavesdroppers who wish to learn the outcome of the estimation or the individual parameters themselves. This can be ensured by combining quantum sensing with quantum cryptographic methods [13–21].

In [13], the authors present a protocol for private parameter estimation, which considers a quantum network of n agents, each hosting an individual parameter (e.g. as a property of a locally held sample). The protocol allows the agents to compute the desired linear function (such as the mean) of all parameters, without, importantly, any agent having to reveal their individual parameter to anyone else inside or outside the network. In other words, every member of the network can obtain information only about their own parameter and the global average. This notion of *privacy* was further expanded and formalised in detail in [18, 19].

Anonymity is another crucial aspect of modern communication. Here anonymity refers to the goal that the *identities* of communicating parties need to remain hidden rather than the information itself. Anonymity was first considered in the quantum setting in [22], where it was shown how to

anonymously send quantum and classical messages, and has since been developed to more robust protocols [23, 24], and guaranteeing anonymity in other scenarios such as conference key agreement [25, 26]. Like covert sensing [27, 28], we obscure participation in the scheme. However instead of considering security against external adversaries we focus on security internal to a network, hiding information about other users' behaviour. Our scheme is also device-agnostic as it is implemented on the communication level, as opposed to depending on the choice of physical infrastructure.

In this work, we combine the notions of privacy and anonymity to present the first protocol for *anonymous private parameter estimation*. Specifically, we study the setting of a quantum network of n agents, who each receive one qubit of a GHZ state (from a potentially untrusted source). One special node, referred to as Alice, chooses the set of *participants*, a subset of $m \leq n$ agents, from whose parameters she wishes to estimate the average. Apart from Alice, every agent in the network is aware of only their own role, but not which other agents are contained in the subset, or which agent is acting as co-ordinator (Alice). All nodes in the network together coordinate a scheme that allows Alice to learn the average, while maintaining the privacy requirement that individual parameters are kept secret, by exploiting the non-local correlations arising from the GHZ state.

This work builds on the idea of private parameter estimation from [13] which we extend towards an anonymous setting using insights from [25, 26, 29]. Care must be taken, however, when combining cryptographic properties; any part of the protocol involving parameter values or participant identities has the potential to compromise either privacy or anonymity respectively. Therefore, anonymous and private schemes cannot be straightforwardly combined in sequence, and one must ensure all desired properties are safeguarded at all times. Our efforts culminate in the first ANONYMOUS PRIVATE PARAMETER ESTIMATION protocol and mark the first application of anonymity beyond the scope of secure communication within quantum technologies. The significance of our contribution lies in adding anonymity to private parameter estimation while maintaining the same levels of accuracy and privacy. Further, we highlight that our protocol can easily be modified according to the required level of privacy or anonymity.

We structure our work as follows. We first give some back-

ground information on the task at hand, presenting the basic parameter estimation scheme in Section II A, and then explain the setting considered here in more detail in Section II B. We then present our protocol in Section III detailing every step and sub-protocol. In Section IV, we prove that our protocol is integrous (the parameter estimation can be trusted), private and anonymous. Finally, in Section V, we discuss adaptations to our security requirements as well as our protocol and conclude our work.

II. BACKGROUND

In this section, we briefly describe the key mechanism for private parameter estimation in a quantum network, we introduce the anonymous setting we will consider in this work, and we describe the notions of privacy and anonymity that we will use.

A. Private parameter estimation

Studies in quantum metrology have long suggested the use of large entangled states for the joint estimation of a linear function of locally held parameters [2, 30–35], with early proof-of-concept experiments now being implemented [36, 37]. The quantum advantage in this scheme comes from an asymptotically quadratic improvement in the number of probe-sample interactions (that is, the number of interactions that the local qubit, e.g. a single photon, must make with the material or field to be measured) [7, 38, 39].

Recent work has considered this task in a cryptographic setting, where these parameters should remain unknown to any other node of the network [13, 18, 19]. In other words, the different parties that form the network are able to collectively estimate a function of the local parameters, without the need to communicate this parameter to any other party. This principle is the basis of *private parameter estimation*.

Analogous to [13], in this work, we focus on the estimation of a particular linear function of the local parameters: the mean. Such a functionality is in line with use cases such as clock synchronisation [2], although it is possible to calculate other linear functions using different states as a resource [19, 33]. In [13], the process for the private estimation of the average value of the local parameters proceeds using a verified GHZ state across an n -user network. Each agent implements a local rotation $\Lambda(\theta_i) = |0\rangle\langle 0| + e^{i\theta_i/n} |1\rangle\langle 1|$ to the received qubit, where θ_i is their private parameter, then measures it in the computational basis and announces the outcome.

The key idea of the protocol is that these individual announcements reveal no information about the private parameters, but allow the nodes to estimate $\bar{\theta} = \frac{1}{n} \sum_i \theta_i$ with high precision. More specifically the probability that the announcements have even parity in each round is $\frac{1}{2} (1 + \cos(\bar{\theta}))$. Repeated rounds therefore allow the network agents to gain a precise estimate of $\bar{\theta}$.

B. Anonymous setting

There are many reasons why anonymity may be required. Even in networks where every member is associated with a public identity, it may be necessary to obscure the relationships between users. For example, in medical use cases, a subgroup may be chosen according to particular characteristics, which participants may wish to remain private, or in political scenarios it may be preferred to keep secret who is involved in certain operations.

The goal of this work is to allow one agent, *Alice* (\mathcal{A}), to act anonymously as an orchestrator of a scheme where she can choose a set of $m \leq n$ participants, \mathcal{P} , and where she is able to estimate $\bar{\theta}_{\mathcal{P}} = (1/m) \sum_{i \in \mathcal{P}} \theta_i$, the average of the local parameters of the agents in \mathcal{P} , while preserving both the anonymity of all agents involved and the privacy of the individual parameters.

Alice has full information regarding the identities of the network members, that is, she knows who is a participant and who is not. Regarding the parameters however, she is as constrained as any participant by the privacy conditions, meaning she will only have access to her own private parameter while estimating $\bar{\theta}_{\mathcal{P}}$.

On the other hand, the property of *anonymity* ensures that any other agent in the network, be it a participant in \mathcal{P} or a non-participant in $\bar{\mathcal{P}}$, will only know their personal role in the protocol (i.e. whether they are a participant or not) and the number of participants. They will however not know who Alice or the other (non-)participants are. Regarding the parameters, they will only know their own private parameter.

We utilise the definition of full anonymity from [26] and adapt it to our protocol and setting: we define anonymity in terms of closeness to an *ideal* output state σ , which contains all of the quantum and classical information of the different parties. Ultimately, the definition ensures that the reduced state of σ on any relevant subset \mathcal{G} of agents in the network is independent of the choice of Alice and the participants.

Table I summarises the information available to the agents of the network in this protocol.

Agent a_i	$a_i = \mathcal{A}$	$\{a_j \in \mathcal{P}\}_{j \neq i}$	$ \mathcal{P} = m$	$\bar{\theta}_{\mathcal{P}}$	θ_i	$\theta_{j \neq i}$
$a_i = \mathcal{A}$	✓	✓	✓	✓	✓	✗
$a_i \neq \mathcal{A}$	✗	✗	✓	✗	✓	✗

Table I. Information available to the different agents of the network. **Identities:** Alice (\mathcal{A}), as the orchestrator, is the only one with information about who are the participants. All agents learn m (the size of \mathcal{P}). **Parameters:** all agents have access to their own parameter θ_i , but none of them are able to get any information about other agents' parameters $\theta_{j \neq i}$. Only Alice can estimate the average of the participants' parameters, $\bar{\theta}_{\mathcal{P}}$.

III. PROTOCOL

We now describe our protocol with all its sub-protocols in detail. The protocol starts by running the NOTIFICATION sub-protocol, anonymously informing each participant

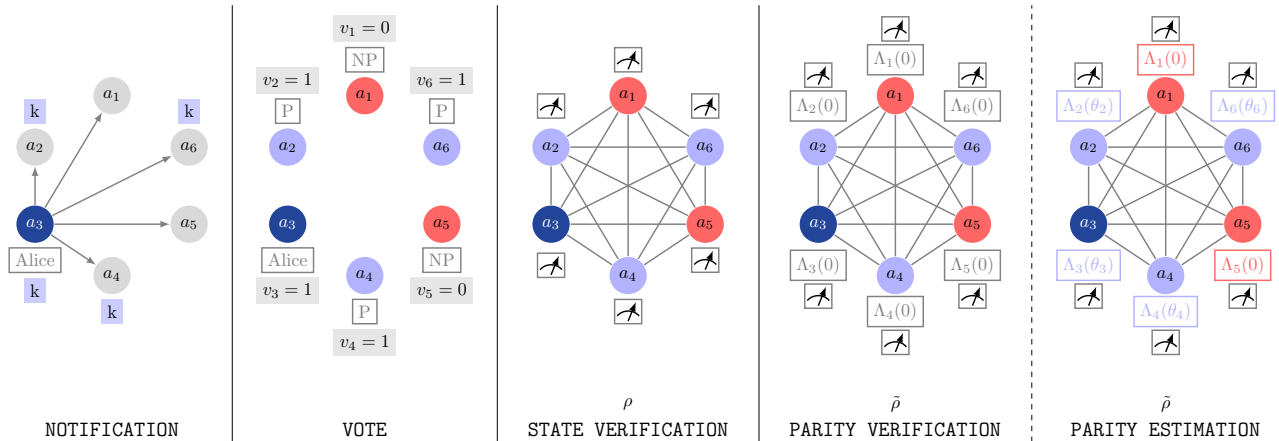


Figure 1. Illustration of ANONYMOUS PRIVATE PARAMETER ESTIMATION on an example of a network of $n = 6$ agents, where a_3 is the co-ordinator Alice, \mathcal{A} , and starts the protocol. In the first step, the nodes run NOTIFICATION to allow Alice to anonymously notify the set of $m = 4$ participants $\{a_2, a_3, a_4, a_6\}$. In the second step the agents run VOTE for the participants to verify m . During this round, participants vote anonymously $v_2 = v_3 = v_4 = v_6 = 1$, and, in an honest scenario, non-participants vote $v_1 = v_5 = 0$. Next the network runs STATE VERIFICATION to ensure that all agents share a state sufficiently close to a GHZ state. Using a secret key Alice can then communicate to the participants whether the state is used for PARITY VERIFICATION (PV) or PARITY ESTIMATION (PE). In PV, everyone measures their qubit in the X -basis and it is used as TRAP round to verify that non-participants do not tamper with the desired estimation. In PE, the participants first apply a rotation on their qubit using their private parameter and then measure in the X -basis, while non-participants are expected to just measure. By repeating PE enough times, Alice can estimate $\theta_{\mathcal{P}}$.

whether or not they are included in \mathcal{P} . Crucially, each agent learns only their own status and gains no information regarding the inclusion of others. This notification step can be implemented using classical protocols, such as those described in [25, 40]; see also Section E.

If only a single agent a_i is included in \mathcal{P} , a malicious Alice can retrieve the private parameter of said agent. Specifically, when $\mathcal{P} = \{\mathcal{A}, a_i\}$ the average value becomes $\theta_{\mathcal{P}} = \frac{1}{2}(\theta_{\mathcal{A}} + \theta_i)$. Since Alice knows $\theta_{\mathcal{A}}$ and $\theta_{\mathcal{P}}$ (see Table I) she can compute the private parameter θ_i . The participants therefore need a guarantee that \mathcal{P} is a large enough set. To this end, the network runs a self tallying majority voting protocol, VOTE, e.g. from [40]. In the VOTE protocol, agents in \mathcal{P} input 1 as their vote, and non-participants input 0. The protocol counts the number of 1 votes and reveals the number of agents in \mathcal{P} , m , to the network, allowing participants to abort if m is not sufficiently large. Without loss of generality, we assume that non-participants act honestly in VOTE, because the only scenario where cheating would be of advantage is if they collaborate with Alice, in which case it would be equally advantageous to simply include the dishonest agents in \mathcal{P} (otherwise, Alice would notice that the total number of participants is incorrect, which constitutes a denial-of-service attack).

In order to be able to run the parameter estimation, the participants need to guarantee that the network shares a GHZ state. The agents therefore run STATE VERIFICATION (SV), a protocol that queries a source to distribute a GHZ which the network verifies. Suitable SV protocols are presented in [41, 42].

This is repeated L times, and each of these shared GHZ states is then (sequentially) used either for estimat-

ing $\bar{\theta}_{\mathcal{P}}$ (a ‘PARITY ESTIMATION (PE) round’), or to detect dishonest behaviour of non-participants (a ‘PARITY VERIFICATION (PV) round’). To coordinate these choices between the participants, they make use of a secret key κ : a bit string of length L with k ‘1’s at random positions indicating the PV rounds, and $\nu = L - k$ ‘0’s indicating the PE rounds, as illustrated in Figure 2. It is vital that the non-participants do not learn what rounds are PV rounds, so in order to establish κ the participants make use of a suitable ANONYMOUS CONFERENCE KEY AGREEMENT (ACKA) protocol (see section V for more details).

Protocol 1 - PARITY ESTIMATION (PE)

Input : Parameters $\{\theta_i\}_{i \in \{1, \dots, n\}}$.
Goal : Alice obtains parity estimation bit χ .

- 1: Each agent a_i applies the unitary $\Lambda_i(\theta_i) = |0\rangle\langle 0| + e^{i\frac{\theta_i}{m}} |1\rangle\langle 1|$ to their qubit.
- 2: Each agent a_i measures their qubit of the GHZ state in the X -basis and gets outcome o_i .
- 3: Each agent except for Alice announces o_i . Alice announces a random bit.
- 4: Alice computes and stores $\chi = \bigoplus_{i=1}^n o_i$.

In PE (Protocol 1) every agent a_i applies the unitary $\Lambda_i(\theta_i) = |0\rangle\langle 0| + e^{i\frac{\theta_i}{m}} |1\rangle\langle 1|$ to their qubit, with $\theta_i = 0$ if $a_i \notin \mathcal{P}$. All agents measure their qubit in the X -basis and

announce their outcome, except for Alice who announces a random bit. Over many runs, Alice can then use the parity of all measurement outcomes (including hers) to estimate $\bar{\theta}_{\mathcal{P}}$, as the probability that the parity is even is $\frac{1}{2}(1 + \cos(\bar{\theta}_{\mathcal{P}}))$ (see Section II A).

Protocol 2 - PARITY VERIFICATION (PV)

Goal: Alice obtains parity verification bit γ .

- 1: Each agent a_i measures their qubit of the GHZ state in the X -basis and gets outcome o_i .
- 2: Each agent except for Alice announces o_i . Alice announces a random bit.
- 3: Alice computes and stores $\gamma = \bigoplus_{i=1}^n o_i$.

With PV (Protocol 2) Alice implicitly verifies that non-participants act honestly. All agents measure their qubit in the X -basis and share their outcomes, except for Alice who again announces a random bit. Here, the total parity of all outcomes must always be 0. Unintended behaviour of any non-participant, such as applying a unitary that would influence the estimation process, would disrupt this condition, allowing Alice to detect dishonest behaviour. Alice computes δ , the relative number of incorrect PV rounds, which she can use to obtain a bound on the accuracy of the parameter estimation.

Finally, Alice estimates $\bar{\theta}_{\mathcal{P}}$ using the measurement outcomes of the PE rounds and the fact that $\Pr(\chi_j = 0) = (1 + \cos(\bar{\theta}_{\mathcal{P}}))/2$ for every j . She bounds its accuracy in terms of δ , using Equation (1).

The full protocol, which we denote by ANONYMOUS PRIVATE PARAMETER ESTIMATION (APPE), is given in Protocol 3 and a small example with 6 agents is illustrated in Figure 1.

As well as the aforementioned GHZ states, the protocol takes the secret identity of Alice and her selection \mathcal{P} consisting of m participants (potentially including herself) as an input.

To ensure the integrity, privacy and anonymity of the protocol, there are certain expectations imposed on the end users, phrased as (*resource*) *requirements*, which is common in cryptography. Some of these requirements arise from the APPE protocol itself, and some are inherited from the sub-protocols. For both the main protocol and the sub-protocols, we assume that all classic channels are authenticated, a standard assumption. Furthermore, we assume that the network has access to an n -partite, high fidelity GHZ state source.

The sub-protocols NOTIFICATION and VOTE rely on pairwise private channels, and VOTE in particular relies on a simultaneous broadcasting channel. We note that no simultaneous broadcasting is necessary to announce the measurement outcomes in the main protocol, because Alice encrypts her outcome. Different implementations of SV can rely on a public source of randomness, and the ACKA protocol (step 4 of Protocol 3) will rely on assumptions as well (in particular, those discussed in section V are either private pairwise clas-

Protocol 3 - ANONYMOUS PRIVATE PARAMETER ESTIMATION (APPE)

Input: GHZ state source, parameters $\{\theta_i\}$, designated co-ordinator Alice and her choice of participants. Parameters ν and k .

Goal: Alice obtains accurate estimate of $\bar{\theta}_{\mathcal{P}}$. Privacy and anonymity are maintained.

- 1: Run NOTIFICATION with \mathcal{A} as coordinator.
- 2: Run VOTE.
- 3: Each agent $a_i \notin \mathcal{P}$ sets $\theta_i = 0$.
- 4: Establish a biased secret key κ of length L between \mathcal{P} , with k '1's and $L - \nu$ '0's.
For round $1 \leq j \leq L$:
 - 5: Distribute N $|\text{GHZ}_n\rangle$ states.
 - 6: Establish a verified GHZ state through an SV protocol.
 - 7-a: If $\kappa_j = 0$: run PE; Alice obtains and records outcome as γ_j .
 - 7-b: If $\kappa_j = 1$: run PV; Alice obtains and records outcome as β_j .
- 8: Alice computes the relative number of incorrect PV rounds δ using $\{\gamma_j\}$.
- 9: Alice estimates $\bar{\theta}_{\mathcal{P}}$ using $\{\chi_j\}$, using the fact that $\Pr(\chi_j = 0) = (1 + \cos(\bar{\theta}_{\mathcal{P}}))/2$ for every j .

sical communication, or sharing additional n -partite, high fidelity GHZ states in the bounded storage model).

IV. PROPERTIES

In this section, we define and discuss integrity, privacy, and anonymity. We provide proof outlines demonstrating how our APPE protocol satisfies these properties and refer the reader to the appendix for complete and detailed proofs.

These security properties do not prevent denial-of-service attacks, which are possible at many stages. They also cannot prevent participants from inputting an incorrect value for their parameter (as this is still a legitimate use of the protocol). However, these security guarantees are not compromised (in the sense defined) under dishonest behaviour, or honest-but-curious behaviour, by any member of the network, eavesdropper, the source, or any collaboration of these.

A. Integrity

The *integrity* of the protocol represents Alice's confidence that the output of the protocol (i.e., her estimate of $\bar{\theta}_{\mathcal{P}}$, represented by the $\hat{\cdot}$ symbol) is both *accurate* and *precise*, even in the presence of malicious adversaries (represented

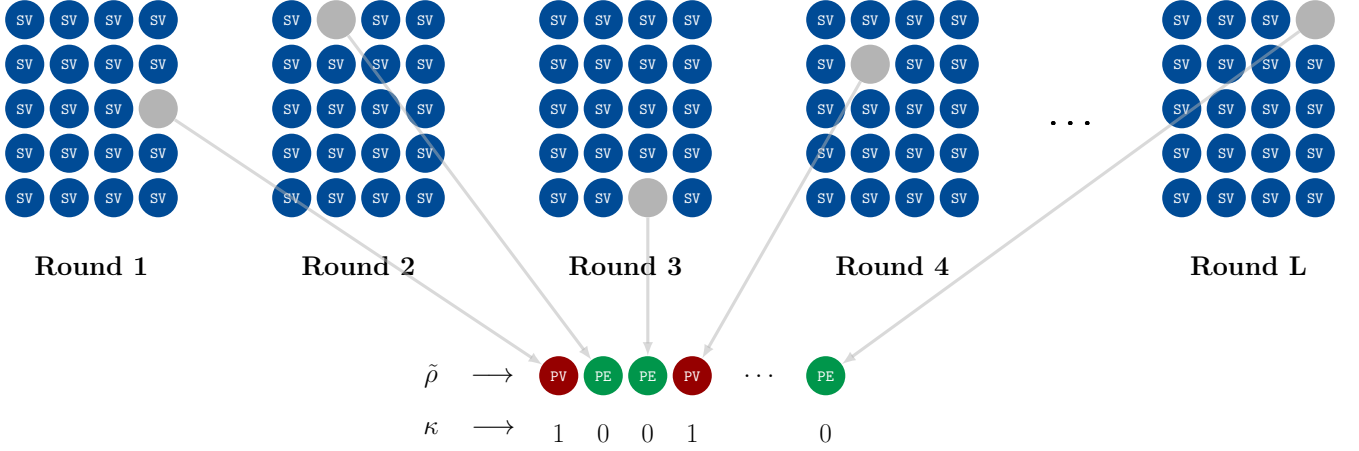


Figure 2. Illustrative example of the selection and use of target states $\tilde{\rho}$ from each round. In this example, five of the total L rounds are shown. In each round a certain number of states are measured and tested to be GHZ states. Of the remaining states one is chosen to be the target state $\tilde{\rho}$ highlighted in grey in the figure. The target state is then used for PE or PV, according to the key κ generated during ACKA.

by the $'$ symbol). These adversaries could include non-participants who should not be participating in parameter estimation, and yet wish to influence Alice's outcome. Accuracy is represented through constraining the bias, i.e. the difference between the expected and actual estimations produced by the protocol: $|\mathbb{E}(\hat{\theta}'_{\mathcal{P}}) - \bar{\theta}_{\mathcal{P}}|$. The uncertainty is represented through any effective change to the variance, $|\Delta^2 \hat{\theta}'_{\mathcal{P}} - \Delta^2 \hat{\theta}_{\mathcal{P}}|$.

Central to this guarantee is the use of PV, which is enabled by the secret key shared by the participants, and hence the security of ACKA is also a vital consideration. We use these verification rounds to ensure that the only individuals able to influence PE rounds are the participants, i.e., an adversary cannot impact a round of PE without this being recorded by a wrong outcome in PV. In particular, note that the outcome of PE is a bit string, which is then used to estimate $\theta_{\mathcal{P}}$; we show that there is an exponentially small probability of adversaries remaining undetected in PV rounds, relative to the number of PE bits that they are able to flip.

Given that we can limit any perturbation that an adversary can enact on the bitstring output by PE, we then show that we can similarly constrain the bias of the estimation of $\bar{\theta}_{\mathcal{P}}$. More specifically, when ν rounds are used for PE and k rounds are used for PV, we limit the bias by:

$$\begin{aligned} & \Pr\left(|\mathbb{E}(\hat{\theta}'_{\mathcal{P}}) - \bar{\theta}_{\mathcal{P}}| \geq \eta\right) \\ & \leq \exp\left(-2(f(\eta, \theta) - \delta)^2 \frac{\nu k^2}{(k + \nu)(k + 1)}\right) \end{aligned} \quad (1)$$

where $f(\eta, \theta)$ is a polynomial in η and θ , and δ is the proportion of PV rounds which give the outcome 1. This shows that there is an exponentially low probability of an attack causing a bias without causing a similarly significant response in the PV rounds.

Furthermore, we show that this attack in fact has no effect on the expected variance of the estimation of $\bar{\theta}_{\mathcal{P}}$, that is:

$$|\Delta^2 \hat{\theta}'_{\mathcal{P}} - \Delta^2 \hat{\theta}_{\mathcal{P}}| = 0. \quad (2)$$

Further details are given in Section B.

B. Privacy

Privacy is understood as the ability of any agent to contribute to the estimation of a global linear function of local parameters $\{\theta_i\}$, such that any subset of dishonest agents \mathcal{D} in the network obtains no more information about any parameter θ_i than the information they already have from knowing the global parameter, the local parameters of the agents in the dishonest set D , and any function of these values. The notion of privacy in this work is relevant only during PE, as this is the only stage in which the parameters are used to estimate the average of the participants' parameters.

The proof for this property follows from [13, 18, 19]. The core idea is that a set of agents can *privately* estimate a global function of local parameters if the quantum Fisher information matrix that depends on the state, after parameter encoding, is a rank-1 matrix. This means that the target linear function can be estimated with arbitrary precision, while no information about any other function of the local parameters can be extracted from the system.

In particular, this work is concerned with the mean of the local parameters. This function can be represented by $\mathbf{w} = (1/m, \dots, 1/m)^T$, a vector of weights [43], so that:

$$\bar{\theta}_{\mathcal{P}} = \mathbf{w}^T \vec{\theta} = \left(\frac{1}{m}, \dots, \frac{1}{m}\right) \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_m \end{pmatrix}, \quad (3)$$

which means that $\mathbf{w}\mathbf{w}^T$ is a constant matrix. This, together with the continuity relation of the Fisher information, implies that any state $\rho_{\vec{\theta}}$, the state after the local encoding of the parameters, that satisfies

$$\partial_i \rho_{\vec{\theta}} = \partial_j \rho_{\vec{\theta}}, \quad \forall i, j \quad (4)$$

can be used for the private estimation of linear functions of the local parameters.

Furthermore, for the case of the mean of the local parameters $\bar{\theta}_{\mathcal{P}}$ it is shown that for an n -partite GHZ state and encoding unitaries $\Lambda(\vec{\theta}) = \bigotimes_{i=1}^n (|0\rangle\langle 0| + e^{i\theta_i/m} |1\rangle\langle 1|)$, $\bar{\theta}_{\mathcal{P}}$ is the only information that can be retrieved from the resulting state:

$$|\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle^{\otimes n} + e^{i\bar{\theta}_{\mathcal{P}}} |1\rangle^{\otimes n} \right). \quad (5)$$

More generally, for an ideal state σ , a state verification scheme produces an output ρ with the guarantee that $\Pr(\|\sigma - \rho\|_{\text{tr}} \geq \varepsilon_{\text{SV}}) \leq \alpha$, for some security parameter ε_{SV} and some confidence bound α ($\|\cdot\|_{\text{tr}}$ denotes the trace norm). In [18], ‘ ε -privacy’ is used to quantify privacy, introducing $\varepsilon_{\text{priv}}$ (defined in terms of the coefficients of the Quantum Fisher Information matrix) which is associated with the maximum amount of information that can be extracted about each local parameter θ . From this definition, and using the outcome of state verification, we can bound the leakage of information as:

$$\varepsilon_{\text{priv}} \leq 2\varepsilon_{\text{SV}} \quad (6)$$

with probability $1 - \alpha$.

The detailed proof schemed in this subsection can be found in Section C.

C. Anonymity

We prove the anonymity of our protocol under an adapted version of *full anonymity* as defined in [26], which requires that an ideal output state of the APPE protocol over all relevant registers (containing both classical and quantum information) needs to be *fully ideally anonymous* - a definition that captures the notion that the state is independent of the choice of sender and participants, for every subset of network agents $\mathcal{G} \subset \{a_i\}_{i=1}^n$ that does not include the sender. Anonymity is then quantified by the ε_a -closeness of the actual output state to such an ideal output state, which we call ρ^{out} and σ^{out} , respectively.

Our protocol consists of several sub-protocols that all contribute differently to the final output state of the APPE protocol and play different roles for anonymity. We can implicitly model the relevant steps of our protocol as a CPTP map Γ , so that $\rho^{\text{out}} = \Gamma(\rho^{\text{in}})$ for some input state ρ^{in} , and $\sigma^{\text{out}} = \Gamma(\sigma^{\text{in}})$, where σ^{in} is the ideal input state defined as:

$$\sigma^{\text{in}} = |\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle\langle\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})|_R^{\otimes L} \otimes \sigma_C \otimes \sigma_E, \quad (7)$$

with R the register that contains the quantum states shared in the network after step 4, C a classical register containing a transcript of all public communication and E the quantum register of Eve.

To prove the anonymity of our protocol we first show that this ideal output state σ^{out} satisfies our definition of a fully

ideally anonymous state (Definition 2 in the appendix). Subsequently, the ε_a -anonymity defined in Definition 3 in the appendix then follows from the fact that the states after step 4 are verified GHZ states.

Indeed, any state verification scheme ensures that $\Pr(\|\rho^{\text{in}} - \sigma^{\text{in}}\|_{\text{tr}} \geq \varepsilon_{\text{SV}}) \leq \alpha$, for some security parameter ε_{SV} and some confidence bound α that decreases with growing N .

By the data processing inequality, we also have $\|\Gamma(\rho^{\text{in}}) - \Gamma(\sigma^{\text{in}})\|_{\text{tr}} \leq \|\rho^{\text{in}} - \sigma^{\text{in}}\|_{\text{tr}}$ which immediately implies that

$$\|\rho^{\text{out}} - \sigma^{\text{out}}\|_{\text{tr}} \leq \varepsilon_a \quad (8)$$

holds for $\varepsilon_a = \varepsilon_{\text{SV}}$ with probability $1 - \alpha$. Our protocol is therefore ε_a -anonymous by Definition 3.

We refer to Section D for more details, but note that we have omitted the confidence window there, essentially taking $\alpha = 0$.

Remark. *While privacy and anonymity are two cryptographic properties of a protocol, neither is a functionality in itself. There are thus no protocols that simply achieve anonymity or privacy but rather we design protocols for a certain task - in our case parameter estimation - that safeguard these properties. In this work these properties coexist independently of one another, while both directly depend on the accuracy of the GHZ state verification protocol.*

V. DISCUSSION

In this work we contribute to the growing number of cryptographic protocols for quantum sensing applications. We specifically manage to improve the private parameter estimation protocol presented in [15] in order to also provide anonymity, while at the same time maintaining integrity and privacy. Indeed, we have presented the first protocol that allows an agent in a network to estimate the average of the parameters of some chosen subset of agents, while their identities as well as the parameter values, remain hidden.

Nevertheless, several challenges remain. Most notably, the protocol heavily relies on verified GHZ states obtained through state verification, which implies the utilisation of a quantum memory to prevent loss of anonymity, or leaking private information. Moreover, many methods for creating and distributing GHZ states suffer both in success rate and robustness to noise [44, 45] when the number of agents increases; our anonymous setting requires GHZ states shared between the entire network rather than only sharing them between the participants as in [13]. The possible anonymity and privacy guarantees that can be achieved directly depend on the quality of GHZ states that can be assumed through state verification. These considerations pose practical limitations that must be addressed with improvements in the actual protocol and in future implementations.

Towards such improvements, consider the fact that the protocol and its proofs (as presented in sections B to D) rely on the fact that the GHZ state is verified; the guarantees on both privacy and anonymity are derived from this. However,

improved proof techniques could alleviate this strong requirement, instead ensuring the privacy and anonymity from (the announcements during) `PARITY VERIFICATION` instead¹. Note that such proof techniques could drastically improve the efficiency of our protocol in terms of both the number of necessary GHZ states that need to be distributed, and the size of the necessary quantum memory. Additionally, `VOTE` relying on simultaneous broadcasting is another challenge, because this may not always be feasible in realistic network settings. Future work should explore alternative approaches to mitigate these constraints, ensuring that the protocol remains both scalable and practical.

There are also various methods for the participants to establish the secret key κ during step 4 of Protocol 3. Essentially, to establish κ the participants need to run a fully anonymous `ANONYMOUS CONFERENCE KEY AGREEMENT (ACKA)` scheme, e.g. in [25, 26]. Note that the participants do not need to run `ACKA` to create a shared secret key of length L : because κ contains only a fraction k/L of ‘1’s, it can be compressed to a bit string of length $h_2(k/L) \cdot L$ (where $h_2(\cdot)$ denotes the binary entropy). The participants only need to establish a secret key of that length, after which they can individually ‘decompress’ it to the desired length L .

As presented, the protocol allows \mathcal{A} to estimate the average of the parameters $\{\theta_i\}_{i \in \mathcal{P}}$. By having all agents change their local rotation $\Lambda(\theta_i)$ to $\Lambda(\theta_i, a_i) = |0\rangle\langle 0| + e^{ia_i\theta_i/n} |1\rangle\langle 1|$ (with $a_i \in \mathbb{R}$), \mathcal{A} can estimate any linear function $\sum_{i \in \mathcal{P}} a_i \theta_i$ instead. However, this approach only works when the agents are aware of their weights a_i , and when the agents are able to change their local rotation. In contrast, a fixed interaction might prove more relevant in a quantum sensing setting, so that the local rotation cannot be adapted. This was addressed in [19], which departed from utilising GHZ states to other multi-partite entangled states: the state that is distributed is carefully adapted by the source to reflect the weights a_i , while maintaining the privacy of all the agents. Future research has to determine whether these states can also be used in our scheme, additionally safeguarding the anonymity of the involved parties.

Beyond adaptations for generic linear functions, adjustments to the protocol can be made in terms of the anonymity. Our definition of anonymity, as adapted from [26], is called *full*, and can be understood as the most stringent form of anonymity where (except for \mathcal{A}) no agent is aware of any-

one else’s role; `NOTIFICATION` and `VOTE` are necessary in our protocol because of this requirement. If the identity of \mathcal{A} is known within the network and she is trusted, `VOTE` can be omitted because the participants do not need an extra guarantee for the privacy of their parameters. Moreover, in such a setting `NOTIFICATION` can be replaced by simple private pairwise communication between \mathcal{A} and all other agents. Alleviating the anonymity even further, one arrives at *partial* anonymity, where the set of participants is completely aware of each other. Note that in this setting it is considerably easier to obtain the shared secret key because *partially anonymous* protocols are easier to implement [26, 29]; it is even customary to assume that a pre-shared key is already in place.

Given the commercial and societal promise of quantum sensors, as well as the growing interest in utilising large scale quantum networks, function estimation is particularly appealing for combining the cryptographic and metrological advantages of quantum correlations, with recent proof-of-concept experimental demonstrations of some of the protocols considered in this work [37, 44, 45] or their adaptations [46]. Despite the remaining challenges in its implementation, the protocol described in this paper extends the functionality of the scheme, thereby increasing its practicality and security and bringing us closer to profiting from quantum networks at scale.

ACKNOWLEDGMENTS

We thank Sean W. Moore for useful discussions.

JdJ, SS, NRS, DM & AP acknowledge the Quantum Internet Alliance (QIA), which has received funding from the European Union’s Horizon 2020 research and innovation programme under grant agreement No 820445 and from the Horizon Europe grant agreements 101080128 and 101102140. SS, NRS & DM acknowledge the PEPR integrated project EPiQ ANR-22-PETQ-0007 part of Plan France 2030. NRS & DM acknowledge support from the ANR project QNS ANR-24-CE97-0005-01. AP acknowledges support from the Emmy Noether DFG grant No. 41829458. ZC and AP acknowledge funding from the Hector Fellow Academy.

[1] N. Shettell, K. S. Lee, F. E. Oon, E. Maksimova, C. Hufnagel, S. Wei, and R. Dumke, Geophysical survey based on hybrid gravimetry using relative measurements and an atomic

¹ Indeed, compare with modern entanglement-based QKD protocols. There, instead of obtaining security by verifying that Alice and Bob share an EPR pair $|00\rangle + |11\rangle$, they merely cross-compare some of their measurement results to verify the security of their key through e.g. entropic uncertainty relations, thereby improving their key rates with orders of magnitude.

gravimeter as an absolute reference, *Scientific Reports* **14**, 6511 (2024).

- [2] P. Kómár, E. M. Kessler, M. Bishof, L. Jiang, A. S. Sørensen, J. Ye, and M. D. Lukin, A quantum network of clocks, *Nature Physics* **10**, 582 (2014).
- [3] P. Kómár, T. Topcu, E. M. Kessler, A. Derevianko, V. Vuletić, J. Ye, and M. D. Lukin, Quantum network of atom clocks: A possible implementation with neutral atoms, *Physical Review Letters* **117**, 060506 (2016).
- [4] H. Dai, Q. Shen, C.-Z. Wang, S.-L. Li, W.-Y. Liu, W.-Q. Cai, S.-K. Liao, J.-G. Ren, J. Yin, Y.-A. Chen, Q. Zhang, F. Xu, C.-Z. Peng, and J.-W. Pan, Towards satellite-based quantum-

- secure time transfer, *Nature Physics* **16**, 848 (2020).
- [5] P. A. Knott, T. J. Proctor, A. J. Hayes, J. F. Ralph, P. Kok, and J. A. Dunningham, Local versus global strategies in multiparameter estimation, *Phys. Rev. A* **94**, 062312 (2016).
- [6] P. C. Humphreys, M. Barbieri, A. Datta, and I. A. Walmsley, Quantum enhanced multiple phase estimation, *Phys. Rev. Lett.* **111**, 070403 (2013).
- [7] C. N. Gagatsos, D. Branford, and A. Datta, Gaussian systems for quantum-enhanced multiple phase estimation, *Phys. Rev. A* **94**, 042342 (2016).
- [8] G. Chiribella, G. M. D’Ariano, and P. Perinotti, Theoretical framework for quantum networks, *Phys. Rev. A* **80**, 022339 (2009).
- [9] R. V. Meter, Quantum networking and internetworking, *IEEE Network* **26**, 59 (2012).
- [10] W. Kozłowski and S. Wehner, Towards large-scale quantum networks, in *Proceedings of the Sixth Annual ACM International Conference on Nanoscale Computing and Communication*, NANOCOM ’19 (Association for Computing Machinery, New York, NY, USA, 2019).
- [11] S. Wehner, D. Elkouss, and R. Hanson, Quantum internet: A vision for the road ahead, *Science* **362** (2018).
- [12] T. Ederer, M. Ivancsits, and I. Ivkić, Temperature monitoring of agricultural areas in a secure data room (2023).
- [13] N. Shettell, M. Hassani, and D. Markham, Private network parameter estimation with quantum sensors (2022).
- [14] N. Shettell and D. Markham, Quantum metrology with delegated tasks, *Phys. Rev. A* **106**, 052427 (2022).
- [15] N. Shettell, E. Kashefi, and D. Markham, Cryptographic approach to quantum metrology, *Physical Review A* **105**, L010401 (2022).
- [16] Y. Takeuchi, Y. Matsuzaki, K. Miyayoshi, T. Sugiyama, and W. J. Munro, Quantum remote sensing with asymmetric information gain, *Phys. Rev. A* **99**, 022325 (2019).
- [17] H. Okane, H. Hakoshima, Y. Takeuchi, Y. Seki, and Y. Matsuzaki, Quantum remote sensing under the effect of dephasing, *Phys. Rev. A* **104**, 062610 (2021).
- [18] M. Hassani, S. Scheiner, M. G. A. Paris, and D. Markham, Privacy in networks of quantum sensors, *Phys. Rev. Lett.* **134** (2025).
- [19] L. Bugalho, M. Hassani, Y. Omar, and D. Markham, Private and Robust States for Distributed Quantum Sensing, *Quantum* **9**, 1596 (2025).
- [20] S. W. Moore and J. A. Dunningham, Secure quantum-enhanced measurements on a network of sensors, *Physical Review A* **111**, 012616 (2025).
- [21] Z. Huang, C. Macchiavello, and L. Maccone, Cryptographic quantum metrology, *Physical Review A* **99**, 022314 (2019).
- [22] M. Christandl and S. Wehner, Quantum anonymous transmissions, in *Advances in Cryptology – ASIACRYPT 2005*, Lecture Notes in Computer Science, Vol. 3788, edited by B. Roy (Springer, Berlin, Heidelberg, 2005) pp. 217–235.
- [23] G. Brassard, A. Broadbent, J. Fitzsimons, S. Gambs, and A. Tapp, Anonymous quantum communication, in *Advances in Cryptology – ASIACRYPT 2007*, Lecture Notes in Computer Science, Vol. 4833, edited by K. Kurosawa (Springer, Berlin, Heidelberg, 2007) pp. 460–473.
- [24] A. Unnikrishnan, I. J. MacFarlane, R. Yi, E. Diamanti, D. Markham, and I. Kerenidis, Anonymity for practical quantum networks, *Phys. Rev. Lett.* **122**, 240501 (2019).
- [25] F. Hahn, J. de Jong, and A. Pappa, Anonymous Quantum Conference Key Agreement, *PRX Quantum* **1**, 020325 (2020).
- [26] F. Grasselli, G. Murta, J. de Jong, F. Hahn, D. Bruß, H. Kampmann, and A. Pappa, Secure Anonymous Conferencing in Quantum Networks, *PRX Quantum* **3**, 040306 (2022).
- [27] M. Tahmasbi, B. A. Bash, S. Guha, and M. Bloch, Signaling for Covert Quantum Sensing, in *2021 IEEE International Symposium on Information Theory (ISIT)*, pp. 1041–1045.
- [28] S. Hao, H. Shi, C. N. Gagatsos, M. Mishra, B. Bash, I. Djordjevic, S. Guha, Q. Zhuang, and Z. Zhang, Demonstration of entanglement-enhanced covert sensing, *Phys. Rev. Lett.* **129**, 010501 (2022).
- [29] J. de Jong, F. Hahn, J. Eisert, N. Walk, and A. Pappa, Anonymous conference key agreement in linear quantum networks, *Quantum* **7**, 1117 (2023).
- [30] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum-enhanced measurements: beating the standard quantum limit, *Science* **306**, 1330 (2004).
- [31] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum Metrology, *Physical Review Letters* **96**, 010401 (2006).
- [32] V. Giovannetti, S. Lloyd, and L. Maccone, Advances in Quantum Metrology, *Nature Photonics* **5**, 222 (2011).
- [33] Z. Eldredge, M. Foss-Feig, J. A. Gross, S. L. Rolston, and A. V. Gorshkov, Optimal and secure measurement protocols for quantum sensor networks, *Physical Review A* **97**, 042337 (2018).
- [34] C. L. Degen, F. Reinhard, and P. Cappellaro, Quantum sensing, *Reviews of Modern Physics* **89**, 035002 (2017).
- [35] T. J. Proctor, P. A. Knott, and J. A. Dunningham, Multiparameter estimation in networked quantum sensors, *Phys. Rev. Lett.* **120**, 080501 (2018).
- [36] L.-Z. Liu, Y.-Z. Zhang, Z.-D. Li, R. Zhang, X.-F. Yin, Y.-Y. Fei, L. Li, N.-L. Liu, F. Xu, Y.-A. Chen, and J.-W. Pan, Distributed quantum phase estimation with entangled photons, *Nature Photonics* **15**, 137 (2021).
- [37] J. Ho, J. W. Webb, R. M. J. Brooks, F. Grasselli, E. Gauger, and A. Fedrizzi, Quantum-private distributed sensing (2024).
- [38] V. Giovannetti, S. Lloyd, and L. Maccone, Quantum metrology, *Phys. Rev. Lett.* **96**, 010401 (2006).
- [39] G. Toth and I. Apellaniz, Quantum metrology from a quantum information science perspective, *Journal of Physics A: Mathematical and Theoretical* **47**, 424006 (2014).
- [40] A. Broadbent and A. Tapp, Information-Theoretic Security Without an Honest Majority, in *Advances in Cryptology – ASIACRYPT 2007*, Vol. 4833, edited by K. Kurosawa (Springer Berlin Heidelberg, Berlin, Heidelberg, 2007) pp. 410–426.
- [41] A. Unnikrishnan and D. Markham, Verification of graph states in an untrusted network, *Physical Review A* **105**, 052420 (2022).
- [42] A. Pappa, A. Chailloux, S. Wehner, E. Diamanti, and I. Kerenidis, Multipartite entanglement verification resistant against dishonest parties, *Physical Review Letters* **108** (2012).
- [43] T. J. Proctor, P. A. Knott, and J. A. Dunningham, Networked quantum sensing (2017).
- [44] J. W. Webb, J. Ho, F. Grasselli, G. Murta, A. Pickston, A. Ulibarrena, and A. Fedrizzi, Experimental anonymous quantum conferencing, *Optica* **11**, 872 (2024).
- [45] C. Thalacker, F. Hahn, J. de Jong, A. Pappa, and S. Barz, Anonymous and secret communication in quantum networks, *New Journal of Physics* **23**, 083026 (2021).
- [46] L. Rückle, J. Budde, J. de Jong, F. Hahn, A. Pappa, and S. Barz, Experimental anonymous conference key agreement using linear cluster states, *Phys. Rev. Res.* **5**, 033222 (2023).
- [47] M. Tomamichel and A. Leverrier, A largely self-contained and complete security proof for quantum key distribution, *Quantum* **1**, 14 (2017).
- [48] S. W. Moore, *Secure quantum-enhanced networks of remote sensors*, Ph.D. thesis, University of Sussex (2024).

- [49] J. J. . Bollinger, W. M. Itano, D. J. Wineland, and D. J. Heinzen, Optimal frequency measurements with maximally correlated states, *Physical Review A* **54**, R4649 (1996).
- [50] C. W. Helstrom, Quantum detection and estimation theory, *Journal of Statistical Physics* **1**, 231 (1969).
- [51] J. Liu, H. Yuan, *et al.*, Quantum fisher information matrix and multiparameter estimation, *Journal of Physics A: Mathematical and Theoretical* **53**, 023001 (2020).
- [52] C. Portmann and R. Renner, Cryptographic security of quantum key distribution (2014).

Appendix A: Notation

\mathcal{A}	Alice/co-ordinator
\mathcal{P}	Set of participants
$\overline{\mathcal{P}}$	Set of non participants
n	Size of network, $n = \mathcal{P} \cup \overline{\mathcal{P}} $
m	Size of set of participants, $ \mathcal{P} $
a_i	i -th agent (node) in the network
θ_i	Parameter of a_i
$\vec{\theta}$	Vector of all parameters, $\vec{\theta} = (\theta_1, \dots, \theta_n)$
$\vec{\theta}_{\mathcal{P}}$	Vector of participants' parameters
$\bar{\theta}_{\mathcal{P}}$	Average of participants' parameters, $\bar{\theta}_{\mathcal{P}} = \frac{1}{m} \sum_{\{i:a_i \in \mathcal{P}\}} \theta_i$
$\hat{\theta}_{\mathcal{P}}$	Estimation of $\bar{\theta}_{\mathcal{P}}$; for readability we also use notation $\hat{\theta} := \hat{\theta}_{\mathcal{P}}$
$\hat{\theta}'_{\mathcal{P}}$	Estimation of $\bar{\theta}_{\mathcal{P}}$, after it has been perturbed by adversarial behaviour; we also use notation $\hat{\theta}' := \hat{\theta}'_{\mathcal{P}}$
L	The total number of states used in APPE (corresponding to the number of states to be produced by SV)
ν	The number of rounds of PE which are carried out (expected to be $L - k$, unless some states are discarded)
k	The number of rounds of L used for PV
δ	The proportion of incorrect PV rounds

Appendix B: Integrity

Integrity refers to the ability to accurately calculate the objective function – that is, to retain the desired functionality of the protocol – even in the presence of malicious adversaries. This is closely related to the *soundness* of the protocol, i.e. the ability to detect any malicious activity. In general, we follow the notation and definitions from [15].

More precisely, the actual mean to be calculated is $\bar{\theta}_{\mathcal{P}}$, and the correct functioning of the protocol produces an estimate to this, $\hat{\theta}_{\mathcal{P}}$ (for readability, for the remainder of this section, we will instead use the notation $\hat{\theta}$ to represent the estimate to the mean). The protocol is unbiased if the expected estimation of the mean matches the true mean, $\mathbb{E}(\hat{\theta}) = \bar{\theta}_{\mathcal{P}}$. In general, a realistic implementation of the protocol with some adversaries will produce a new estimate $\hat{\theta}'_{\mathcal{P}}$ (similarly, we will use the notation $\hat{\theta}'$ for the remainder of this section). We are hence interested in the bias:

$$|\mathbb{E}(\hat{\theta}') - \bar{\theta}_{\mathcal{P}}|. \quad (\text{B1})$$

We are also concerned with a measure of the uncertainty introduced by malicious parties [15]:

$$|\Delta^2 \hat{\theta}' - \Delta^2 \hat{\theta}|. \quad (\text{B2})$$

By the expected behaviour of the protocol, members of \mathcal{P} have complete freedom to input any angle from 0 to 2π at each parameter estimation round (although the desired behaviour is to enter a value in the range 0 to $2\pi/m$, this is not technically enforced), and hence in this analysis we do not consider the dishonest behaviour of participants, who could freely produce any bias or uncertainty. Instead, we consider dishonest behaviour from any subset of the non-participants, or another adversary, given that any of these may have control over the source, or quantum channels between the source and members of the network.

1. Constraining the affected rounds of PARITY ESTIMATION

The estimate $\hat{\theta}$ is produced within the final stage of APPE. During this step, a bitstring is created of length L , of which we expect that k are used for PV, which produces error rate δ . The remaining $\nu = L - k$ bits are used for PE. The subset of rounds, V , which are used for PV are decided by a preshared key between the participants. We use an ACKA protocol that has an exponentially low probability of failure, and therefore we will at first assume that the adversary has no knowledge of which bits of this block are used for PE and which are used for PV. Therefore, any attack is permutation invariant with regards to the L bits of APPE – it has equal chance of landing on a PE or PV round.

We also assume that we start the APPE protocol with a GHZ state distributed across the whole network. This is based on the correct functioning of SV, with the assumption that the source of randomness to decide which rounds are used for verification is called after the state distribution.

The outcome of all rounds of PV can be expressed as a test function, $\{0, 1\}^k \rightarrow \{\checkmark, \emptyset\}$, where the outcome is \emptyset if $\sum_{i=1}^k v_i > \delta k$ and \checkmark otherwise, where v_i is the parity of an individual verification round (computed by Alice) and δ represents some accepted level of error (that is, the proportion of verification rounds that fail, where the overall test still passes). If more than the proportion δ of these rounds give parity 1 ($v_i = 1$), the protocol is abandoned. Alternatively, Alice can place no requirements on δ , but take this as an outcome of the PV rounds and use it in calculating the expected bias of her estimate.

Consider the situation where, unknown to the adversaries, all of the participants input $\theta_i = 0$, but otherwise behave honestly (that is, they simply make X measurements and announce the outcome). We can use the following result from [47] (proof omitted):

Lemma 1. *Consider a set of binary random variables $Z = (Z_1, Z_2, \dots, Z_L)$ where $L = \nu + k$. Let V be an independent, uniformly distributed random subset of size k . Then:*

$$\Pr\left(\sum_{i \in V} Z_i \leq k\delta \wedge \sum_{i \in \bar{V}} Z_i \geq (L - k)(\delta + \omega)\right) \leq \exp\left(-2\omega^2 \frac{(L - k)k^2}{L(k + 1)}\right). \quad (\text{B3})$$

This is applied such that the random variables Z are the outcomes of each round of APPE, where the subset V are used for verification and \bar{V} are used for parameter estimation. The correct outcome for each round in both verification and estimation is 0, but any bit may be flipped to a 1 by the action of adversaries. However, as the adversaries do not know which rounds are in V , then each Z_i has the same distribution.

The term $\sum_{i \in V} Z_i \leq k\delta$ then specifies the case where the protocol passes. It is to be expected that up to $\delta\nu$ of the parameter estimation rounds come out as 1 (this is the accepted error). However, we are interested in a further bias that affects a further $\alpha\nu$ of the remaining outcomes. The probability of this occurring when the protocol passes is exponentially small (as shown in Lemma 1).

This result is still useful in the case that $\theta_i \neq 0$. Note that, using Lemma 1 we can limit (by α) the proportion of rounds in which an adversary behaves in such a way that, should the compromised quantum state or classical information be used for verification, the output bit would be 1. Thus, we need to confirm that any state which would pass a verification round can be used effectively for parameter estimation (that is, giving 0 or 1 with the correct probability).

Now we consider the sorts of attacks that can be carried out by dishonest participants. Regarding classical communication, the honest participants publicly announce their measurement result, and Alice keeps her measurement result private and obscures it by announcing a random bit, and hence the only attack that can be carried out using only the classical communication is to flip the outcome bit. Alternatively, simultaneous broadcast could be used, in which case the adversaries again have no knowledge of the parity of the rest of the bit string before announcing their own bit.

Now we consider any attacks on the quantum state, starting with the assumption of a distributed GHZ state. As the local operations performed by the parties commute, we can assume that Alice and the participants, as well as honest non-participants, have all made their X measurements and now have the appropriate outcome. The state across the l remaining members of the network is therefore:

$$\frac{1}{\sqrt{2}} \left(|0\rangle^{\otimes l} + (-1)^h e^{i\bar{\theta}_{\mathcal{P}}} |1\rangle^{\otimes l} \right) \quad (\text{B4})$$

where h is the parity of all of the X measurement outcomes of honest members of the network, \mathcal{H} – which is unknown to the dishonest users, as \mathcal{H} contains Alice, who obscures her outcome. We must assume that the dishonest users have knowledge of $\bar{\theta}_{\mathcal{P}}$, as many rounds may have already occurred, or the protocol may have been run previously.

Given that $h = 0$ or 1 with equal probability, we can see that the density matrix of this state is independent of $\bar{\theta}_{\mathcal{P}}$ (indeed it is $\frac{1}{2}\mathbb{1}$, the maximally mixed state, in the basis spanned by $|0\rangle^{\otimes l}, |1\rangle^{\otimes l}$). Therefore any rotations or measurements made locally among the adversaries has the same impact on the classical information (and in particular the only relevant information, the parity of their shared l -bit string) for any value of $\bar{\theta}_{\mathcal{P}}$ – that is, any behaviour that causes a bit flip of the parity in a round of parameter estimation does the same for parameter verification.

This situation is different if Alice does not encode her outcome. In this case, the dishonest non-participants could potentially force a particular outcome (e.g. parity 0 for a round), which would not be detected by the verification scheme, which would damage integrity (or at least present a denial-of-service attack). Hence, in this case simultaneous broadcast should be enforced.

2. Bias

We would like to use Lemma 1 to bound the potential bias of $\hat{\theta}'$ (see, for example, Fig. 7.5 of [48]). The perturbation of the outcome by an adversary is given by $\alpha = \delta + \omega$, the proportion of the ν rounds which have bits flipped. Let β be the correct proportion of the $L - k$ bits used for parameter estimation which have the value 0:

$$\beta = \frac{1}{2} \left(1 + \cos(\hat{\theta}) \right) \quad (\text{B5})$$

(recall that we are using the shorthand $\hat{\theta} := \hat{\theta}_P$). The order of the bits (i.e. which are 0 and which are 1) is random, and hence the new expected proportion is:

$$\begin{aligned} \beta' &= \beta(1 - \alpha) + (1 - \beta)\alpha = \beta + \alpha - 2\alpha\beta \\ &= \frac{1}{2} \left(1 + \cos(\hat{\theta}') \right). \end{aligned} \quad (\text{B6})$$

We set $\eta := \hat{\theta}' - \hat{\theta}$, and then we can find:

$$|\alpha| = \left| \sin(\hat{\theta} + \eta/2) \sin(\eta/2) / \cos(\hat{\theta}) \right|. \quad (\text{B7})$$

We now aim to show that if $|\alpha|$ is sufficiently small, then $|\eta|$ is also small.

We can lower bound the size of α according to η and $\hat{\theta}$ (using truncated Taylor expansions):

$$\begin{aligned} |\alpha| &= \left| \sin(\eta/2) \left(\tan(\hat{\theta}) \cos(\eta/2) + \sin(\eta/2) \right) \right| \\ &\geq \frac{1}{2} \left| \tan(\hat{\theta}) \sin(\eta) \right| \\ &\geq \frac{1}{2} \left| \hat{\theta} + \frac{\hat{\theta}^3}{3} + \frac{2\hat{\theta}^5}{15} \right| \left| \frac{\eta}{2} \right|. \end{aligned} \quad (\text{B8})$$

That is, α is at least polynomial in η . Hence, it is not possible to achieve arbitrarily large bias without incurring a polynomial cost in the number of parameter estimation rounds affected (which we can bound).

3. Uncertainty

We now consider the uncertainty potentially introduced by adversaries. The variance of the estimation typically scales as $1/\nu$ [49] (or $1/\nu n^2$ if the proportion is instead expressed as $\beta = \frac{1}{2}(1 + \cos(n\theta))$, using slightly different notation).

Recall that we have defined α to be some proportion of the bitstring produced by PE that are flipped by malicious behaviour. For the binomial distribution we can use the result that the variance for a single trial is $\beta(1 - \beta)$ in the original case, and:

$$\begin{aligned} \Delta^2 \hat{\beta}' &= (\beta + \alpha - 2\alpha\beta)(1 - \beta - \alpha + 2\alpha\beta) \\ &= \sin^2(\hat{\theta})/4 - \alpha(1 - \alpha) \cos^2(\hat{\theta}) \end{aligned} \quad (\text{B9})$$

in the corrupted case.

Hence, we can use the error propagation:

$$\begin{aligned} \Delta \theta' &= \frac{d\hat{\theta}'}{d\beta'} \Delta \beta' = \frac{2}{\sin(\hat{\theta}')} \Delta \beta' \\ &= \sqrt{\frac{\sin^2(\hat{\theta})}{\sin^2(\hat{\theta}')} - 4\alpha(1 - \alpha) \frac{\cos^2(\hat{\theta})}{\sin^2(\hat{\theta}')}}. \end{aligned} \quad (\text{B10})$$

Using:

$$\begin{aligned}\sin^2(\hat{\theta}') &= 1 - (1 - 2\beta')^2 \\ &= 1 - \cos^2(\hat{\theta})(1 - 2\alpha)^2,\end{aligned}\tag{B11}$$

we arrive at:

$$\Delta\hat{\theta}' = 1.\tag{B12}$$

Note that this is for a single round, but over the ν rounds, we reintroduce the $1/\nu$ factor. Therefore, we have rederived the unperturbed variance, and we can see that:

$$|\Delta^2\hat{\theta}' - \Delta^2\hat{\theta}| = 0.\tag{B13}$$

The fact that the proportion of perturbed rounds, α , does not influence the variance, may be understood by considering that for a fixed α , the impact of the malicious behaviour on the parameter estimation is well-behaved. There is no prior assignment of which rounds are 0 and which are 1, and therefore being able to reverse the output of random rounds has a fixed effect on the expectation value without necessarily introducing any additional noise, and simply adding a bias. Alternatively, this can be understood through an attack strategy: if an adversarial non-participant was able to add a $\pi/4$ rotation, and remain undetected by PV, this would disturb a high proportion of the PE bits but could be reformulated as the estimation of a new parameter with the same variance.

4. Summary of integrity

The condition on the integrity of parameter estimation is the bias:

$$\Pr\left(|\mathbb{E}(\hat{\theta}') - \bar{\theta}_{\mathcal{P}}| \geq \eta\right) \leq \exp\left(-2(f(\eta, \hat{\theta}) - \delta)^2 \frac{(L-k)k^2}{L(k+1)}\right)\tag{B14}$$

where:

$$f(\eta, \hat{\theta}) = \frac{1}{2} \left| \hat{\theta} + \frac{\hat{\theta}^3}{3} + \frac{2\hat{\theta}^5}{15} \right| \left| \frac{\eta}{2} \right|.\tag{B15}$$

This rests on several assumptions, in particular that non-participants have no information about which rounds are used for PV and that a true GHZ state is distributed to the network. The first assumption depends on the ACKA protocols used. If the ACKA protocol allows a bits of key to be leaked, this corresponds directly to αL bits which can be flipped without detection (as in Equation (B3)).

The quality of the GHZ state is assured through the verification protocol. Much like the analysis in Section B 1, we note that a state which is able to pass the SV rounds would also produce the correct outcome at PE, and hence, if there is a failure rate of a rounds, this can once again be interpreted as αL affected bits of PE, although we note that this may have an outsized change to β' , as adversaries can in this case force an outcome, e.g. make sure that the outcome bit is 1 – hence, we can follow the same argument but instead with $\beta' = \beta + \alpha$.

It is important to highlight the role of δ here: this can be seen as a known bias introduced in the parameter estimation. On the one hand, as δ increases in size, Alice's confidence in the bias bound expressed in Eq. B14 decreases. This can alternatively be expressed that there is already a bias η' given by:

$$|\delta| = \left| \sin(\eta'/2) \left(\tan(\hat{\theta}) \cos(\eta'/2) + \sin(\eta'/2) \right) \right|\tag{B16}$$

and the total bias can now be expressed as:

$$\Pr\left(|\mathbb{E}(\hat{\theta}') - \bar{\theta}_{\mathcal{P}}| \geq \eta + \eta'\right) \leq \exp\left(-2f(\eta, \hat{\theta})^2 \frac{(L-k)k^2}{L(k+1)}\right).\tag{B17}$$

On the other hand, assuming δ is sufficiently small, using the measured PV proportion β' , and δ , Alice can correct the error by calculating:

$$\beta = \frac{\beta' - \delta}{1 - 2\delta}.\tag{B18}$$

The bias is then simply Equation (B14), with $\delta = 0$. This correction procedure assumes that the noise that creates the error δ in PV is equally likely to cause a bit flip in both PV and PE rounds; as shown previously, this is the case for attacks by potential adversaries, but may be true for other sources of noise (such as in Alice's experimental apparatus). Therefore, this relies on the experimental noise being sufficiently low, or well-characterised, as to only consider noise from adversaries or dishonest non-participants.

Appendix C: Privacy

In this work, we define the notion of *privacy* as the property of a given scheme to be executed while ensuring that sensitive information that belongs to individual agents remains inaccessible to unauthorised parties.

1. Definition of privacy

Formally, *privacy* in the context of networked sensing can be defined as follows (see the diagram below for an example).


Definition 1. Let $\mathcal{D} = \{d_1, \dots, d_{m_d}\}$ be a subset of m_d dishonest agents selected from $\{a_1, \dots, a_n\}$, the complete set of agents in the network; and let $\mathcal{T} = \mathcal{D}^c = \{t_1, \dots, t_{n-m_d}\}$ be the target subset of size $n - m_d$, given by the complement of set \mathcal{D} .

Let $\{\theta_i^{(d)}\}$ be the local parameter of agent $d_i \in \mathcal{D}$, and $\{\theta_j^{(t)}\}$ the parameter of a target agent $t_j \in \mathcal{T}$.

A protocol is **private** if the quantum Fisher information \mathcal{Q} that any subset of agents \mathcal{D} can extract about any function $f = f(\theta_1^{(t)}, \dots, \theta_{n-m_d}^{(t)})$ throughout the protocol is:

$$\mathcal{Q}(f|\rho_{\mathcal{D}}) \leq \mathcal{Q}(f|\bar{\theta}_{\mathcal{P}}, \theta_1^{(d)}, \theta_2^{(d)}, \dots, \theta_{m_d}^{(d)}). \quad (\text{C1})$$

Here ρ represents the global classical-quantum state produced by the protocol, encompassing the final distributed quantum state and any exchanged classical information, and $\rho_{\mathcal{D}}$ is the partial state that is shared by the dishonest agents. In other words, agents in \mathcal{D} do not get more information about $f(\theta_1^{(t)}, \dots, \theta_{n-m_d}^{(t)})$ than they already have from sharing the state ρ and from knowing the local parameters $\{\theta_i^{(d)}\}$ of agents in \mathcal{D} , the public value $\bar{\theta}_{\mathcal{P}}$, and any function $g(\bar{\theta}_{\mathcal{P}}, \theta_1^{(d)}, \theta_2^{(d)}, \dots, \theta_{m_d}^{(d)})$.



2. Proof overview

A quantifier of privacy in this scenario needs to capture the idea that, once the information of each local parameter θ_i is locally encoded in the state:

$$\rho = \frac{|0\rangle^{\otimes n} + |1\rangle^{\otimes n}}{\sqrt{2}} \quad \longrightarrow \quad \rho_{\bar{\theta}} = \frac{|0\rangle^{\otimes n} + e^{i\bar{\theta}_{\mathcal{P}}} |1\rangle^{\otimes n}}{\sqrt{2}}, \quad (\text{C2})$$

only information about $\bar{\theta}_{\mathcal{P}} = (1/m) \sum_{i:a_i \in \mathcal{P}} \theta_i$ can be extracted from the system, but each secret parameter θ_i remains hidden. In our protocol, this guarantee essentially follows from the work of [13, 18, 19], as any step that involves parameter input uses the same states and operations as described in previous iterations of the protocol.

With this definition of privacy in mind, we can introduce L_i , the symmetric logarithmic derivative (SLD) for the parameter θ_i . L_i is a Hermitian operator given by the relation

$$L_i \rho_{\bar{\theta}} + \rho_{\bar{\theta}} L_i = 2 \frac{\partial \rho_{\bar{\theta}}}{\partial \theta_i} \quad (\text{C3})$$

and it allows us to compute the coefficients of the quantum Fisher information matrix (QFI, \mathcal{Q}) [50, 51].

The QFI is a symmetric matrix with real elements that quantifies the amount of extractable information about different unknown parameters over all possible measurements. Using the SLD, the elements of this matrix can be calculated as

$$\mathcal{Q}_{ij}[\bar{\theta}] = \frac{1}{2} \text{Tr} [\rho_{\bar{\theta}} \{L_i, L_j\}]. \quad (\text{C4})$$

This means that the QFI for parameter θ_i is

$$\mathcal{Q}_{ii}[\bar{\theta}] = \text{Tr} [\rho_{\bar{\theta}} L_i^2] \quad (\text{C5})$$

from which we can see that the presence of non-zero off-diagonal entries of the QFI implies statistical correlation between the local parameters, as there is a way of extracting information of θ_i and θ_j simultaneously. Equivalently, if the different SLDs do not commute, the different parameters cannot be estimated independently. On the other hand, if the QFI is block diagonal, parameters in different blocks are *information orthogonal*, in the sense that their maximum likelihood estimates are asymptotically uncorrelated.

As the aim is to estimate one function of the unknown parameters, $\bar{\theta}_{\mathcal{P}} = f(\vec{\theta})$, the corresponding QFI may be obtained by reparametrisation, that is:

$$\mathcal{Q}'[\vec{\theta}] = B^T \mathcal{Q}[\vec{\theta}] B$$

where B is a transformation matrix into an orthogonal basis such that the first element is the desired linear combination. Privacy (as defined in this work) is ensured if \mathcal{Q}' is a rank-1 matrix, as this implies that only information about one linear combination of the parameters can be retrieved: the mean of all the local parameters.

Following the derivation from [18], the linear function of interest can be encoded in a vector $\mathbf{w} = (\omega_1, \dots, \omega_m)^T$ (cf. [43]), so that:

$$\bar{\theta}_{\mathcal{P}} = \mathbf{w}^T \vec{\theta} = (\omega_1, \dots, \omega_m) \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_m \end{pmatrix}. \quad (\text{C6})$$

As we are only interested in the function given by \mathbf{w} , we can choose the matrix:

$$W \equiv \mathbf{w}\mathbf{w}^T = \begin{pmatrix} \omega_1\omega_1 & \omega_1\omega_2 & \cdots & \omega_1\omega_m \\ \omega_2\omega_1 & \omega_2\omega_2 & \cdots & \omega_2\omega_m \\ \vdots & \vdots & \ddots & \vdots \\ \omega_m\omega_1 & \omega_m\omega_2 & \cdots & \omega_m\omega_m \end{pmatrix}, \quad (\text{C7})$$

and then $\mathcal{Q}' \propto W$ can be used for the desired purpose of this work, as it is a rank-1 matrix that carries the information of the average value. This implies that:

$$\mathcal{Q}_{ij} \propto \omega_i \omega_j \quad \forall i, j. \quad (\text{C8})$$

Using the continuity relation of the quantum Fisher information, as shown in detail in [18], the following condition must be satisfied:

$$\|\partial_i \rho_{\vec{\theta}} - \partial_j \rho_{\vec{\theta}}\|_{\text{tr}} \propto |\omega_i - \omega_j|, \quad \forall i \neq j. \quad (\text{C9})$$

In this particular case, the aim is to compute the average value of the local parameters:

$$\bar{\theta}_{\mathcal{P}} = \mathbf{w}^T \vec{\theta} = \left(\frac{1}{m}, \dots, \frac{1}{m} \right) \begin{pmatrix} \theta_1 \\ \vdots \\ \theta_m \end{pmatrix}, \quad (\text{C10})$$

which means that Eq. (C7) becomes:

$$\mathbf{w}\mathbf{w}^T = \frac{1}{m^2} \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{pmatrix} \quad (\text{C11})$$

and, since $\omega_i = \omega_j = \frac{1}{m}$ for all i, j , Eq. (C9) becomes simply:

$$\partial_i \rho_{\vec{\theta}} = \partial_j \rho_{\vec{\theta}}, \quad \forall i, j. \quad (\text{C12})$$

Therefore, any state that satisfies Eq. (C12) can be used for the private estimation of the average value.

Using this, [18] shows that for an n -partite GHZ state and encoding unitaries $\Lambda(\vec{\theta}) = \bigotimes_{i=1}^n (|0\rangle\langle 0| + e^{i\frac{\theta_i}{m}} |1\rangle\langle 1|)$, the only information that can be retrieved from the resulting state:

$$|\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle = \frac{1}{\sqrt{2}} \left(|0\rangle^{\otimes n} + e^{i(\theta_1 + \dots + \theta_n)/m} |1\rangle^{\otimes n} \right) \quad (\text{C13})$$

corresponds to the sum of all the local parameters, and therefore the local parameters of users remain private.

More generally, for an ideal state σ , any state ρ that originates from the output of any state verification scheme that ensures

$$\Pr(\|\sigma - \rho\|_{\text{tr}} \geq \varepsilon_{\text{SV}}) \leq \alpha, \quad (\text{C14})$$

for some security parameter ε_{SV} and some confidence bound α , which can be made arbitrarily small. From [18], we know that the ‘ ε -privacy’ (a variable that quantifies the leakage of information about the local parameters θ_i) can be bounded with probability $1 - \alpha$ as:

$$\varepsilon_{\text{priv}} \leq 4 \|M\|_{\infty} \|\sigma - \rho\|_{\text{tr}} \leq 4 \|M\|_{\infty} \varepsilon_{\text{SV}} \quad (\text{C15})$$

where M is an operator that contains the information about the local encoding of the parameters. For the case of the average value of the local parameters, $\|M\|_{\infty} = \|\sigma_z/2\|_{\infty} = \frac{1}{2}$.

A similar definition of ε -privacy is defined in [19].

Appendix D: Anonymity

We make use of the definition of *full anonymity* from [26] and adapt it to our setting. This definition compares the output state of the protocol against an *ideally anonymous* state, which is any state that is perfectly anonymous under a specific requirement, introduced below. Note that there is not one unique ideally anonymous state, but that any state which adheres to the requirement will be ideally anonymous. We first introduce the structure of the output state, and state the requirement for the output state to be ideally anonymous. Utilising the concept of ideally anonymous states, we can define anonymity of an APPE protocol, which is in terms of closeness to any such ideally anonymous state. Subsequently, we prove our protocol’s anonymity under this definition.

The anonymity statement in Section IV C presents our results with respect to a confidence window inherited from the preceding SV protocol. In the following however we omit the confidence window, essentially taking $\alpha = 0$.

Ideally anonymous states and definition of full anonymity

The output state $\sigma_{PTKCE}^{\text{out}}$ is defined on the registers P, T, K, C and E , all containing classical information. Each of these registers can be indexed by the agents, so that e.g. $T_t = 0.01$ indicates that agent a_t holds the value 0.01 in their part of the register T . Note that an agent can only access the entries of the registers at their own index. Within this section, let i be the specific index such that $a_i = \mathcal{A}$, and let \vec{j} be a vector with 1 at indices corresponding to participants and 0 otherwise. Using this notation, the registers are defined as:

P : This register contains the information regarding the roles in the network for every agent. It holds that:

$$P_t = \begin{cases} \vec{j} & \text{if } t = i \\ 1 & \text{if } a_t \in \mathcal{P} \setminus \mathcal{A}, \\ 0 & \text{otherwise.} \end{cases}$$

T : This register holds the parameters of the agents: $T_t = \theta_t$.

K : This register contains the secret key κ for the participants:

$$K_t = \begin{cases} \kappa & \text{if } a_t \in \mathcal{P}, \\ \emptyset & \text{otherwise.} \end{cases}$$

C : This register is divided into three sub-registers, defined as follows:

- C_{NV} contains all (classical) information from NOTIFICATION and VOTE,
- C_{SV} contains all information from SV, and
- C_{PP} contains all public classical communication from PV and PE.

E : This register contains all classical and quantum information held by the adversary Eve. This may include the information of one or more dishonest nodes.

Note that we have omitted Alice's measurement outcome and $\bar{\theta}_P$ from being included in any of the registers to ease our analysis. While these are crucial for integrity and privacy, they are irrelevant for anonymity, because Alice never announces her measurement outcome.

The output state σ^{out} depends on Alice and her choice of participants, as well as on the set of dishonest parties, $\mathcal{D} \subsetneq \{1, \dots, n\}$. We capture this by using the following notation:

$$\sigma_{PTKCE|i, \vec{j}}^{\text{out}, \mathcal{D}}$$

This notation allows us to precisely state our requirement for ideally anonymous states. Even though the output state σ^{out} is dependent on Alice and her choice of participants, we define a state to be anonymous if the reduced state of σ^{out} on any relevant subset is independent of Alice and the participants. More specifically, let $\mathcal{G} \subseteq \{1, \dots, n\}$ be a subset of agents, and $\mathcal{G}^c = \{1, \dots, n\} \setminus \mathcal{G}$ be its complement. For a given register S , we define $S_{\mathcal{G}} = \{S_i \in S | i \in \mathcal{G}\}$, and for a state ρ_S , we have $\rho_{S_{\mathcal{G}}} = \text{Tr}_{S_{\mathcal{G}^c}}(\rho_S)$. Using this notation, we can now precisely state the requirement for $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ to be ideally anonymous.

Definition 2. Let \mathcal{D} be the set of dishonest agents in an APPE protocol with an output state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$. The state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is fully ideally anonymous with respect to a subset $\mathcal{G} \subseteq \{1, \dots, n\}$ if it holds that:

$$\sigma_{P_{\mathcal{G}}T_{\mathcal{G}}K_{\mathcal{G}}C_{\mathcal{G}}E_{\mathcal{G}}|i, \vec{j}}^{\text{out}, \mathcal{D}} = \sigma_{P_{\mathcal{G}}T_{\mathcal{G}}K_{\mathcal{G}}C_{\mathcal{G}}E_{\mathcal{G}}|i', \vec{j}'}^{\text{out}, \mathcal{D}}, \quad (\text{D1})$$

for any i, i', \vec{j}, \vec{j}' such that $i, i' \notin \mathcal{G} \cup \mathcal{D}$, $\vec{j} \cap \mathcal{D} = \vec{j}' \cap \mathcal{D}$ and $\vec{j} \cap \mathcal{G} = \vec{j}' \cap \mathcal{G}$. Moreover, because we specifically fix the size of the set of participants, it should hold that $H(\vec{j}) = H(\vec{j}')$, where $H(\vec{j}) = \sum_k j_k$ denotes the Hamming weight of \vec{j} , i.e. its total number of 1's.

This state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is fully ideally anonymous with respect to the set of dishonest agents \mathcal{D} if it is fully ideally anonymous with respect to every subset $\mathcal{G} \subseteq \{1, \dots, n\}$.

With the definition of a fully ideally anonymous state, we are now equipped to define ε_a -full anonymity of our APPE protocol.

Definition 3. An APPE protocol with an output state $\rho_{PTKCE}^{\text{out}, \mathcal{D}}$ is ε_a -fully anonymous with respect to a set of dishonest agents \mathcal{D} if there exists a fully ideally anonymous state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ such that:

$$\|\rho_{PTKCE}^{\text{out}, \mathcal{D}} - \sigma_{PTKCE}^{\text{out}, \mathcal{D}}\|_{\text{tr}} \leq \varepsilon_a, \quad (\text{D2})$$

where $\|\cdot\|_{\text{tr}}$ denotes the trace distance.

Note that, similar to security in quantum key distribution (QKD) protocols [47, 52], we have defined anonymity with respect to an anonymity parameter ε_a . This captures the notion of *quantifiable* anonymity, which allows for the protocol to be fully anonymous in non-perfect scenarios, e.g. due to noise. One can then obtain anonymity *approximately*, meaning that the anonymity statements hold except for increasingly small probabilities. In the following subsections, we prove that our protocol satisfies Definition 3.

Proof overview

In order to show that our APPE protocol is ε_a -fully anonymous, we show that its output state $\rho_{PTKCE}^{\text{out}, \mathcal{D}}$ is ε_a -close to a particular fully ideally anonymous output state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$.

This specific output state is defined as the output of APPE in a completely ideal scenario; we can implicitly model our protocol as a CPTP map Γ , so that $\sigma_{PTKCE}^{\text{out}, \mathcal{D}} = \Gamma(\sigma^{\text{in}})$, where σ^{in} is an idealised input state to our protocol.

More specifically, $\sigma^{\text{in}} = \sigma_{R,I,E}^{\text{in}}$ is defined over three registers. R contains the quantum states used in the protocol, I the other inputs to the protocol (i.e. the index of Alice, her choice of participants, and the parameters of all the agents), and E Eve's quantum and classical side information.

We first observe that SV is completely independent of the roles of the agents, and therefore we can ignore this step in our proof. This means that the register R can be regarded to only hold the verified GHZ states, so that we can write $\tau_{R,I,E}$ as (recalling that $L = \nu + k$ is the total number of verified GHZ states used for PE and PV):

$$\sigma^{\text{in}} = |\text{GHZ}_n\rangle \langle \text{GHZ}_n|_R^{\otimes L} \otimes \sigma_I \otimes \sigma_E. \quad (\text{D3})$$

Furthermore, SV guarantees that for some $\varepsilon_{\text{ver}} > 0$ it holds that $\|\rho_R^{\text{in}} - \sigma_R^{\text{in}}\|_{\text{tr}} \leq \varepsilon_{\text{ver}}$, where ρ^{in} is the actual input state of the protocol - note that, because σ_R^{in} is pure, it can be concluded that $\rho^{\text{in}} = \rho_R^{\text{in}} \otimes \rho_{I,E}^{\text{in}}$. As the register I does not change

between the ideal and real case, it holds that $\rho_T^{\text{in}} = \sigma_T^{\text{in}}$. Finally, all communication after SV is only correlated with the register R , so there is no loss of generality in assuming that $\rho_E^{\text{in}} = \sigma_E^{\text{in}}$, so that it holds that $\|\rho^{\text{in}} - \sigma^{\text{in}}\|_{\text{tr}} \leq \varepsilon_{\text{ver}}$.

We then have $\rho_{PTKCE}^{\text{out}, \mathcal{D}} = \Gamma(\rho^{\text{in}})$ and can leverage the contractivity of the trace distance under quantum channels (data processing inequality) to obtain:

$$\|\rho^{\text{out}} - \sigma^{\mathcal{D}}\|_{\text{tr}} = \|\Gamma(\rho^{\text{in}}) - \Gamma(\sigma^{\text{in}})\|_{\text{tr}} \leq \|\rho^{\text{in}} - \sigma^{\text{in}}\|_{\text{tr}} \leq \varepsilon_{\text{ver}}. \quad (\text{D4})$$

Proving anonymity of our APPE protocol therefore reduces to showing that $\sigma_{PTKCE}^{\text{out}, \mathcal{D}} = \Gamma(\sigma^{\text{in}})$ is a fully ideally anonymous state under Definition 2. Indeed, from (D4) it then follows that our APPE protocol is ε_{ver} -fully anonymous under Definition 3.

In the remainder of this section we prove that $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is a fully ideally anonymous state, first in the honest ($\mathcal{D} = \emptyset$), and later in the dishonest setting (i.e. for arbitrary \mathcal{D} , as long as $\mathcal{A} \notin \mathcal{D}$).

$\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is fully ideally anonymous in the honest setting

Our goal is now to show that $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is fully ideally anonymous under Definition 2 in the honest setting, i.e. where $\mathcal{D} = \emptyset$, allowing us to write $\sigma_{PTKCE}^{\text{out}, \mathcal{D}} = \sigma_{PTKCE}^{\text{out}}$.

As a starting point, we note that although there exists correlations between various registers P , T , K , C and E the ideal state is separable over a certain partitioning of the registers, which allows us to write $\sigma_{PTKCE}^{\text{out}}$ in tensor product form:

$$\sigma_{PTKCE}^{\text{out}} = \sigma_{PTKC_{NV}} \otimes \sigma_{C_{PP}} \otimes \sigma_{C_{SV}} \otimes \sigma_E. \quad (\text{D5})$$

The registers P , T , K and C_{NV} cannot be written in tensor product form. Indeed, only the participants obtain the secret key so that P and K are correlated, similarly only the participants keep their parameter in their register T , and NOTIFICATION and VOTE correlate with the contents of P as well. An important part of our later analysis shows that C_{PP} is not correlated with any other register. The tensor product structure of (D5) then follows from the fact SV is independent of the rest of the protocol, so that $\sigma_{C_{SV}}$ is separate from the rest of the output state. Finally, there is no loss of generality in assuming that σ_E is uncorrelated, because all relevant (side)-information (e.g. the public communication) can be understood to be explicitly considered already within the other registers.

To prove that $\sigma_{PTKCE}^{\text{out}}$ is fully ideally anonymous under Definition 2, it suffices to show that the tensor factors $\sigma_{PTKC_{NV}}$, $\sigma_{C_{PP}}$, $\sigma_{C_{SV}}$ and σ_E individually satisfy eq. (D1).

σ_E NOTIFICATION and VOTE rely on private pairwise communication, therefore Eve cannot learn anything from these protocols that was not shared publicly. Moreover, the input state $\tau_{R,I,E}$ is separable between all three registers (by (D3)), so that the σ_E remains completely separable from the networks' registers, and the identities of the agents. Hence, σ_E trivially satisfies (D1).

$\sigma_{C_{SV}}$ As noted before, there is no distinction between participants and non-participants in SV , i.e. every agent in the network performs the exact same steps, regardless of their role. This means that $\sigma_{C_{SV}}$ trivially satisfies (D1).

$\sigma_{PTKC_{NV}}$ By construction, the registers P , T and K together obey (D1) for any subset \mathcal{G} and any choice of i, i' and \vec{j}, \vec{j}' . From [40] we know that both NOTIFICATION and VOTE are anonymous, so that the classical communication involved with the protocols are independent of Alice and her choice of participants. However, VOTE outputs m , which makes C_{NV} dependent on P . Nevertheless, Definition 2 specifically only compares states with the same number of participants (i.e. $H(\vec{j}) = H(\vec{j}')$), so that the requirement (D1) is still met. It therefore holds that $\sigma_{PTKC_{NV}}$ is fully ideally anonymous under Definition 2.

$\sigma_{C_{PP}}$ To show that $\sigma_{C_{PP}}$ satisfies (D1), we prove that its contents are uniformly random by virtue of Alice announcing a random bit instead of her actual measurement outcome.

Apart from Alice's random bit, $\sigma_{C_{PP}}$ contains the announced outcomes of the X -basis measurement on $|\text{GHZ}_n\rangle = (1/\sqrt{2}) (|0\rangle^{\otimes n} + |1\rangle^{\otimes n})$ in the PV rounds and $|\text{GHZ}_n(\bar{\theta}_P)\rangle = (1/\sqrt{2}) (|0\rangle^{\otimes n} + e^{i\bar{\theta}_P} |1\rangle^{\otimes n})$ in the PE rounds. While the outcomes of the X -basis measurements obey parity statistics for both PE and PV rounds, this is only the case when Alice's outcome is also considered. However Alice announces a random bit, ensuring that the *announced* measurement outcomes stored in C_{PP} are uniformly random and uncorrelated, and thus independent of the choice of \mathcal{P} .

To this end we show that for any strict subset it holds that all measurement outcomes are equally likely. Indeed, consider $A \subsetneq \{1, 2, \dots, n\}$, we can then write the reduced states on A as:

$$\text{Tr}_{A^c}(|\text{GHZ}_n\rangle \langle \text{GHZ}_n|) = \text{Tr}_{A^c}(|\text{GHZ}_n(\bar{\theta}_P)\rangle \langle \text{GHZ}_n(\bar{\theta}_P)|) = \frac{|0\rangle\langle 0|^{\otimes |A|} + |1\rangle\langle 1|^{\otimes |A|}}{2}. \quad (\text{D6})$$

Consider now the probability $\Pr(\{o_j\}_{j \in A})$ that the agents in A obtain a specific set of outcomes $\{o_j\}_{j \in A} \in \{0, 1\}^{|A|}$ for their X -basis measurements. A straightforward calculation reveals that the probability of obtaining this specific outcome $\{o_j\}_{j \in A}$ is $\Pr(\{o_j\}_{j \in A}) = 1/2^{|A|}$. That is, each outcome is equally likely. As this holds for any proper subset A , it follows that the partial measurement outcomes are uniformly random and uncorrelated. It follows that the contents of C_{PP} are independent of the choice of \mathcal{P} , and therefore that $\sigma_{C_{PP}}$ obeys (D1).

Finally, we note that due to the fact that the correlations in the measurement outcomes are only between the complete set of outcomes, and that any subset of measurement outcomes is uniformly random and uncorrelated, it is actually not necessary for Alice to announce a random bit instead of her actual measurement outcome to safeguard anonymity. However, the tensor product structure of (D5) is then lost, as C_{PP} will depend on T and P , making the notation less clear. Furthermore, as briefly discussed in Section V, this would introduce the need of a simultaneous broadcast channel, or an adaptation to the protocol so that the independence of all measurement outcome announcements can be guaranteed.

We proved that σ_{PTKCNV} , $\sigma_{C_{PP}}$, σ_{CSV} and σ_E all satisfy (D1) and therefore by (D5) we know that $\sigma_{PTKCE}^{\text{out}}$ obeys Definition 2 and thus our protocol is fully anonymous in the honest setting, where $\mathcal{D} = \emptyset$.

$\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ is fully ideally anonymous in the dishonest setting

We now turn our attention to the case where \mathcal{D} is non-empty, i.e. there are nodes that are deviating from the protocol to try and learn the roles of other agents of the network. Importantly, the dishonest agents may base their strategy on the announcements of the honest agents. Indeed, while the honest agents announce their measurement results for both $\mathbb{P}V$ and $\mathbb{P}E$, the dishonest members can delay their measurement and announce anything else instead. The intermediate output state relevant to our analysis, which is the state *before* the dishonest agents perform their measurement, is therefore:

$$\sigma_{PTKCNV} \otimes \sigma_{SV} \otimes \rho_{\tilde{C}_{PP}R_{\mathcal{D}}} \otimes \sigma_E, \quad (\text{D7})$$

where, $\rho_{\tilde{C}_{PP}R_{\mathcal{D}}}$ is the classical-quantum state on the classical register \tilde{C}_{PP} and the quantum register $R_{\mathcal{D}}$. We write \tilde{C}_{PP} to emphasize that the contents of the register differ from the honest case, because the dishonest agents have not measured their qubits and announced something else. $R_{\mathcal{D}}$ is the part of the quantum register that has not been measured, i.e. the qubits of the agents in \mathcal{D} . Ultimately, the output state $\sigma_{PTKCE}^{\text{out}, \mathcal{D}}$ will be obtained from the state in (D7).

To show that the output state is fully ideally anonymous it therefore suffices to show that the state in (D7) is itself fully ideally anonymous (similar to Definition 2 but adapted to reflect the change of registers). We can again exploit the tensor product structure and analyse the tensor factors individually; neither the states σ_{PTKCNV} nor σ_{SV} have changed from the honest setting, which means that they are still fully ideally anonymous.

Although the output registers C_{PP} and E may become correlated by any deviation of the protocol by the dishonest agents, the intermediate state of (D7) indeed does allow the tensor product structure, so that we may analyse $\rho_{\tilde{C}_{PP}R_{\mathcal{D}}}$ separately from σ_E . In general, the state $\rho_{\tilde{C}_{PP}R_{\mathcal{D}}}$ is a CQ-state that correlates any possible set of measurement outcomes $\{o_j\}_{\mathcal{D}^c}$ of the honest agents with a state $\rho_{R_{\mathcal{D}}|\{o_j\}_{\mathcal{D}^c}}^{\mathcal{D}}$ on $R_{\mathcal{D}}$, which, in turn, can be computed as:

$$\rho_{R_{\mathcal{D}}|\{o_j\}_{\mathcal{D}^c}}^{\mathcal{D}} = (\langle \{o_j\}_{\mathcal{D}^c} | \otimes I_{\mathcal{D}}) |\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle \langle \text{GHZ}_n(\bar{\theta}_{\mathcal{P}}) | (\{o_j\}_{\mathcal{D}^c} \otimes I_{\mathcal{D}}). \quad (\text{D8})$$

In order to prove that $\rho_{R_{\mathcal{D}}|\{o_j\}_{\mathcal{D}^c}}^{\mathcal{D}}$ is fully ideally anonymous, we now show that it is independent of the choice of participants. To this effect we first rewrite the GHZ state as:

$$|\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle = \frac{1}{\sqrt{2^{n+1}}} \left((1 + e^{i\bar{\theta}_{\mathcal{P}}}) \sum_{\Delta(\{o_j\})=0} |\{o_j\}\rangle + (1 - e^{i\bar{\theta}_{\mathcal{P}}}) \sum_{\Delta(\{o_j\})=1} |\{o_j\}\rangle \right), \quad (\text{D9})$$

where $\Delta_{\{o_i\}}$ denotes the parity of $\{o_i\}$, i.e. its Hamming weight modulo two. From this it follows:

$$\begin{aligned} (\langle \{o_j\}_{\mathcal{D}^c} | \otimes I_{\mathcal{D}}) |\text{GHZ}_n(\bar{\theta}_{\mathcal{P}})\rangle &= \frac{1}{\sqrt{2^{|\mathcal{D}|+1}}} \left((1 + e^{i\bar{\theta}_{\mathcal{P}}}) \sum_{\Delta_{\{o_j\}_{\mathcal{D}}} = \Delta_{\{o_j\}_{\mathcal{D}^c}} |\{o_j\}_{\mathcal{D}}\rangle \right. \\ &\quad \left. + (1 - e^{i\bar{\theta}_{\mathcal{P}}}) \sum_{\Delta_{\{o_j\}_{\mathcal{D}}} = 1 - \Delta_{\{o_j\}_{\mathcal{D}^c}} |\{o_j\}_{\mathcal{D}}\rangle \right). \end{aligned} \quad (\text{D10})$$

It is directly apparent that this state is independent of the actual choice of participants but rather depends only on $\bar{\theta}_{\mathcal{P}}$ which is irrelevant for anonymity. There is therefore no information on the roles of the agents that dishonest members can gain by deviating from the protocol. Finally we can conclude that the protocol is fully anonymous under Definition 3 in the dishonest setting as well.

Appendix E: Sub-protocols

In this section, we present classical sub-protocols used throughout our work, with adapted notation. The APPE protocol presented in this work makes use of two classical subprotocols: NOTIFICATION and VOTE. Although these steps can be implemented in any way that preserves the anonymity of the participants (see Section D), a particular choice of algorithms for these initial steps is presented by Broadbent & Tapp [40].

In the NOTIFICATION protocol, a coordinating agent is able to notify all participating agents while preserving the anonymity of every party involved. The output, which every agent computes locally and in private, reveals to each agents whether they participate or not in the overall APPE protocol, and no other information can be extracted from it. This subprotocol, with adapted notation, is presented as follows.

Protocol 4 - NOTIFICATION

(Broadbent & Tapp 2007)

Input : A coordinator Alice (\mathcal{A}), a set of m participants $\mathcal{P} \subseteq \{a_1, \dots, a_n\}$.

Goal : \mathcal{A} notifies the m participants.

1: Each agent $\{a_i\}_{i=1}^n$ sends every other agent $\{a_j\}_{j=1}^n$, n random bits $\{r_{ijk}\}_{k=1}^n$, such that:

- if $a_i = \mathcal{A}$, the random bits satisfy $\bigoplus_{j=1}^n r_{ijk} = \begin{cases} 1, & \text{if } a_k \in \mathcal{P} \\ 0, & \text{if } a_k \in \overline{\mathcal{P}} \end{cases}$
- if $a_i \neq \mathcal{A}$, the random bits satisfy $\bigoplus_{j=1}^n r_{ijk} = 0$

2: Each agent $\{a_j\}_{j=1}^n$ receives the bits $\{r_{ijk}\}_{i=1}^n$, computes $t_{jk} = \bigoplus_{i=1}^n r_{ijk}$ and sends $\{t_{jk}\}_{k=1}^n$ to agent a_k .

3: Each agent $\{a_k\}_{k=1}^n$ receives $\{t_{jk}\}_{j=1}^n$, and computes:

$$z_k = \bigoplus_{j=1}^n t_{jk} = \begin{cases} 1, & \text{if } a_k \in \mathcal{P} \\ 0, & \text{if } a_k \in \overline{\mathcal{P}} \end{cases} \quad (\text{E1})$$

For the second stage of the APPE protocol, two subprotocols are needed: PARITY and VOTE. In the first one, all agents are able to input a binary value, and all agents can compute the global parity of the network. This is used as a subroutine in the VOTE protocol. In this subprotocol, agents anonymously input their *choice* (in this case, 0 or 1, depending on whether they are participants or non-participants). At the end of this stage, each agent is able to compute the number of participants in the scheme, and decide whether to continue with the protocol or not. The PARITY and VOTE protocols are presented below, with adapted notation. The probability of an incorrect outcome is exponentially small in the number of rounds, s , and simultaneous broadcast is required to ensure that corrupt participants do not receive partial information which could be used in future rounds, potentially compromising the anonymity of the scheme.

Protocol 5 - PARITY

(Broadbent & Tapp 2007)

Input : Parities $\{x_i\}$, with each $x_i \in \{0, 1\}$.

Goal : Each agent computes the global parity $y = x_1 \oplus x_2 \oplus \dots \oplus x_n$.

Each agent a_i does the following:

- 1: Select uniformly at random an n -bit string $r_i = r_i^{(1)} r_i^{(2)} \dots r_i^{(n)}$ with Hamming weight of parity x_i .
- 2: Send $r_i^{(j)}$ to agent a_j using the private channel; keep bit $r_i^{(i)}$ to themselves.
- 3: Compute z_i , the parity of all the bits received, including $r_i^{(i)}$.
- 4: Use the simultaneous broadcast channel to announce z_i .
- 5: After the simultaneous broadcast is finished, compute $y = \bigoplus_{j=1}^n z_j$, the outcome of the protocol. If the simultaneous broadcast fails, the protocol aborts.

Protocol 6 - VOTE

(Broadbent & Tapp 2007)

Input : Choice $x_i \in \{0, 1\}$ voted by each agent ($x_i = 1$ if $a_i \in \mathcal{P}$, $x_i = 0$ if $a_i \in \overline{\mathcal{P}}$), the number of rounds s .

Goal : Each agent computes the tally $y = (y[0], y[1]) = (n - m, m)$ with the number of votes for each candidate (i.e. the number of non-participants, and the number of participants, respectively).

For each choice $b \in \{0, 1\}$:

Phase A:

For each round $j \in \{1, \dots, s\}$:

- 1: each agent a_i sets the value of p_i in the following way: if $x_i \neq b$, $p_i = 0$; otherwise, $p_i = 1$ with probability $\frac{1}{n}$ and $p_i = 0$ with complementary probability.
- 2: the participants execute the PARITY protocol to compute the parity of p_1, p_2, \dots, p_n , but instead of broadcasting their output bit z_i , they store it as $z[b]_i^j$.

Phase B:

All agents $\{a_i\}$ simultaneously broadcast $z[b]_i^j$ ($j \in \{1, \dots, s\}$). If the simultaneous broadcast is not successful, the protocol aborts.

Phase C:

To compute the tally $y[b]$, each participant sets: $p[b]_j = \bigoplus_{i=1}^n z[b]_i^j$, $\sigma[b]_i = \sum_{j=1}^s p[b]_j / s$ and if there exists an integer v such that $|\sigma[b]_i - p_v| < \frac{1}{2e^{2n}}$, where:

$$p_v - \frac{1}{2} \left(\frac{n-2}{n} \right)^v \left(\left(\frac{n}{n-2} \right)^v - 1 \right),$$

then $y[b] = v$.

If, for any b , no value v exists or $y[0] + y[1] \neq n$, the protocol aborts.