



Non-Interactive and Non-Destructive Zero-Knowledge Proofs on Quantum States and Multi-Party Generation of Authorized Hidden GHZ States

LÉO COLISSON, LIP6, Paris, France, Centrum voor Wiskunde en Informatica, Amsterdam, Netherlands, and LJK, Grenoble, France

FRÉDÉRIC GROSSHANS, LIP6, Paris, France

ELHAM KASHEFI, School of Informatics, The University of Edinburgh, Edinburgh, United Kingdom of Great Britain and Northern Ireland and LIP6, Sorbonne Université Campus Pierre et Marie Curie, Paris, France

We propose the first generalization of the famous Non-Interactive Zero-Knowledge (NIZK) proofs to *quantum* languages (NIZKoQS) and we provide a protocol to prove advanced properties on a received quantum state non-destructively and non-interactively (a single message being sent from the prover to the verifier).

In our second orthogonal contribution, we improve the costly Remote State Preparation protocols [Cojocaru et al. 2019; Gheorghiu and Vidick 2019] that can classically fake a quantum channel (this is at the heart of our NIZKoQS protocol) by showing how to create a multi-qubit state from a single superposition.

Finally, we generalize these results to a multi-party setting and prove that multiple parties can anonymously distribute a GHZ state in such a way that only participants knowing a secret credential can share this state, which could have applications to quantum anonymous transmission, quantum secret sharing, quantum onion routing and more.

CCS Concepts: • **Security and privacy** → **Public key (asymmetric) techniques**;

Additional Key Words and Phrases: Quantum cryptography, remote state preparation, zero-knowledge, learning with errors

ACM Reference Format:

Léo Colisson, Frédéric Grosshans, and Elham Kashefi. 2025. Non-Interactive and Non-Destructive Zero-Knowledge Proofs on Quantum States and Multi-Party Generation of Authorized Hidden GHZ States. *ACM Trans. Quantum Comput.* 6, 2, Article 16 (April 2025), 56 pages. <https://doi.org/10.1145/3721487>

This work has been supported in part by the French ANR Project ANR-18-CE39-0015 (CryptiQ), the French ANR project ANR-18-CE47-0010 (QUDATA), the EU Flagship Quantum Internet Alliance (QIA) project and the European Union as H2020 Programme under grant agreement number ERC-669891. EK acknowledges support from the EPSRC Verification of Quantum Technology grant (EP/N003829/1), the EPSRC Hub in Quantum Computing and Simulation (EP/T001062/1), and the UK Quantum Technology Hub: NQIT grant (EP/M013243/1).

Authors' Contact Information: Léo Colisson, LIP6, Paris, Île-de-France, France, Centrum voor Wiskunde en Informatica, Amsterdam, Netherlands, and LJK, Grenoble, Auvergne-Rhône-Alpes, France; e-mail: leo.colisson-palais@univ-grenoble-alpes.fr; Frédéric Grosshans, LIP6, Paris, Île-de-France, France; e-mail: frederic.grosshans@lip6.fr; Elham Kashefi, School of Informatics, The University of Edinburgh, Edinburgh, United Kingdom of Great Britain and Northern Ireland and LIP6, Sorbonne Université Campus Pierre et Marie Curie, Paris, Île-de-France, France; e-mail: ekashefi@gmail.com.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM 2643-6817/2025/04-ART16

<https://doi.org/10.1145/3721487>

1 Introduction

Due to the special no-cloning principle, quantum states appear to be very useful in cryptography. But this very same property also has drawbacks: when receiving a quantum state, it is nearly impossible for the receiver to efficiently check non-trivial properties on that state without destroying it. This allows a sender to send maliciously crafted states (potentially entangled with a larger system) without being detected.

To illustrate this, let us imagine the following simple goal. A receiver would like to obtain a quantum state $|\psi\rangle$ and verify, without destroying that state, that this state belongs to some “*quantum language*”, say the language composed of BB84 states (so $|\psi\rangle \in \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$). Since any direct measurement would destroy that state, a first solution could be to use a generic **quantum secure multiparty computing (QSMPC)** protocol [Dulek et al. 2020; Dupuis et al. 2012; Kapourniotis et al. 2025] between the sender and the receiver in order to generate that state. However, these protocols are interactive and require at least 2 messages (depending on the number of users and on the complexity of the prepared state, the number of rounds can increase significantly). Therefore, the following question was left open:

Is it possible to receive via a single message a quantum state and test non-trivial properties on it non-destructively?

At the heart of our method are *classical-client Remote State Preparation (RSP)* protocols [Cojocar et al. 2019, 2021; Gheorghiu and Vidick 2019] being itself based on the groundbreaking work of [Mahadev 2018b]. These protocols can be used to fake a quantum channel using a purely classical channel. However, they are particularly heavy to run: a single qubit can be sent at a time at the expense of doing a very large and expensive quantum superposition. To prepare more complex states, the solution is typically to repeat the above protocol n times before combining appropriately the different runs. Unfortunately, this method is very costly since the superposition needs to be recreated n times. Therefore, a second question naturally arises, orthogonal to the first one:

Is it possible to classically send multi-qubits states while paying the cost of a single superposition?

1.1 Our Results

In this work, we answer the two above interrogations in the affirmative, and provide various use cases together with a generalization of RSP to multiparty:

- (1) We first provide a method to classically prepare large states on n qubits (where a GHZ state is hidden in between $|0\rangle$ and $|1\rangle$ qubits) using a single superposition. It allows us to improve the existing RSP protocols, spreading the cost to create one state among the multiple qubits. We are therefore able to asymptotically obtain a quadratic improvement: the number of operations required to prepare this n qubits state drops from $O(nMN)$ to $O((M+n)N) = O(nN)$, $M > N$ being very large constants (for a fixed security parameter λ but variable number of qubits n , note that M and N do depend on the security parameter λ).
- (2) We lift the above construction to a multiparty setting in which each sender can control one part of the hidden GHZ state. This construction is still non-interactive in the sense that a single message is sent from each sender to the receiver.
- (3) We initiate the study of **Non-Interactive and Non-Destructive Zero-Knowledge Proofs on Quantum States (NIZKoQS)**. We obtain these guarantees by sending a classical NIZK proof on *classical* instructions used to produce a quantum state. This approach has numerous advantages: the verification is *non-destructive*, *non-interactive*, can be done over a purely *classical channel* (the sender does not need any quantum capabilities), and it is now *possible to check much more involved properties* on the quantum state, potentially linked with secret classical information or involving *multiple parties*. For instance, we can force the sender to

send different categories of states depending on whether or not they know a classical password or signature, without revealing to the receiver to which category the state belongs to.

We emphasize that this method is very generic and can be applied to any (non-interactive) RSP protocol, but also to a non-interactive computation performed using the protocol from [Mahadev 2018b].¹ However, we focus in this article on a generalization of [Cojocararu et al. 2019] which allows more efficient GHZ-like state generation and can be extended to a multiparty protocol.

While it is not the first time that LWE-based methods are used in quantum cryptography to protect the sender, to the best of our knowledge, it is the first time that this kind of cryptography is used to also protect the receiver in the context of RSP protocols. This allows us to *bypass fundamental quantum limitations*: we expect this result to have further independent interests.

- (4) We present a concrete use case in the context of the *generation of multi-party certified hidden GHZ states*. We derive a protocol between a server and n “applicants”. At the end of the protocol, the applicants end-up sharing a GHZ state in such a way that, if the server is honest, only the applicants knowing a classical secret (password/secret key/signature/...) can be part of the GHZ (they are then said to be *supported*). Other applicants will not be entangled in any way with the supported applicants, and therefore, cannot disturb a potential future protocol. Our protocol has the following properties:
 - The protocol is *blind*, in the sense that no adversary (corrupting potentially the server and some applicants) can learn whether the honest applicants are supported or not.
 - As soon as the server is honest, it is also impossible for any malicious or noisy applicant to significantly change the GHZ state obtained by supported applicants. This is particularly interesting to avoid “Denial Of Service” kind of attacks where some noisy applicants would always force the protocol to abort.

One direct application is the following: a server can do a quantum secret sharing between parties that have been authorized, say, by some Certification Authority. The identity of these authorized parties stay hidden, even against a malicious coalition involving the server and applicants.

- (5) We provide a series of simpler protocols that only provide blindness. In these protocols a classical trusted third party is in charge of choosing who is supported and who is not. These protocols could prove useful in the context of Anonymous Transmission, *Quantum Onion Routing*, ...
- (6) We give an *instantiation of a cryptographic family* required in our protocol whose security relies on the hardness of LWE (with superpolynomial modulus-to-noise ratio).

1.2 Related Works and Further Applications

Classical **Zero-Knowledge (ZK)** proofs and Interactive Proofs systems have been introduced thirty years ago [Babai and Moran 1988; Goldwasser et al. 1989], and allow a prover to prove a statement to a verifier without revealing anything beyond the fact that this statement is true. ZK proofs have been proposed for any language in NP [Goldreich et al. 1991]. While **Non-Interactive Zero-Knowledge (NIZK)** proofs are known to be impossible in the plain model [Fiat and Shamir 1987], NIZK can be obtained in the Common Reference String model [Blum et al. 1988] or in the Random Oracle model by using the famous Fiat–Shamir transformation [Goldreich and Oren 1994]. The security of classical ZK proofs were also extended to protect against malicious quantum provers [Bitansky and Shmueli 2020; Unruh 2012; Watrous 2009], and much work has been done

¹In that later case, we can prove any property on the output state before the application of the Pauli keys. Note also that when dropping the non-interactive requirements, any property on the final state should be verifiable (and it is also possible to use more involved classical protocols like classical Multi-Party computation [Goldreich et al. 1987] to classically manipulate quantum states), but this is out of the scope of this article.

to extend these protocols to remove interactivity [Don et al. 2019; Liu and Zhandry 2019] or deal with multiple (potentially quantum) provers, targeting larger classes like IP, QMA, MIP, or MIP*, sometimes considering “dequantized” verifiers ([Alagic et al. 2020; Bartusek et al. 2021; Broadbent and Grilo 2020; Broadbent et al. 2016; Coladangelo et al. 2020; Grilo et al. 2019; Impagliazzo and Yung 1988; Morimae and Yamakawa 2022; Shmueli 2021; Vidick and Zhang 2020], see also the review [Vidick and Watrous 2016]).

However, these works focus on *classical* languages (even in QMA, the language is still classical, even if the witness can be quantum). To prove similar properties on *quantum* languages, a first solution is to use a generic QSMPC protocol [Dulek et al. 2020; Dupuis et al. 2012; Kapourniotis et al. 2025]. However, these protocols are interactive, and to the best of our knowledge no work has been done to provide one-shot ZK proofs on quantum states.

In this article, we use internally classical cryptography to leverage quantum cryptography. [Mahadev 2018b] introduced for the first time this notion, and this idea led to the development of many quantum protocols based on classical cryptography [Brakerski 2018b; Brakerski et al. 2018; Mahadev 2018c; Metger et al. 2021].

The first classical-client RSP protocol was proposed in [Cojocaru et al. 2021]. Informally, in this protocol, a purely classical client can create a state on a quantum server, in such a way that only the client knows the description of that qubit. This protocol named QFactory was later improved in [Cojocaru et al. 2019], and proven secure against an arbitrary malicious adversary. Another protocol [Gheorghiu and Vidick 2019] can also be used to obtain a verifiable classical-client RSP. While this later article provides stronger guarantees in terms of verification, it cannot be directly extended to produce multi-party GHZ states, requires interaction and offers a security scaling polynomially with the security parameter for both verifiability (which is certainly hard to avoid) and blindness. For these reasons, we will build our approach upon the QFactory protocol. Some general impossibility results on classical RSP protocols are also given in [Badertscher et al. 2020; Morimae and Takeuchi 2020]: therefore, we will not consider the Constructive Cryptography framework in this article.

In this work, we focus on the distributed generation of **Greenberger–Horne–Zeilinger** (GHZ) state [Greenberger et al. 1989]. This state (and its special case, the Bell state) is very popular, and appears to be useful in many protocols, such as in Quantum Secret Sharing [Hillery et al. 1999], Quantum Teleportation [Bennett et al. 1993], Entanglement Distillation [Bennett et al. 1996; Bennett et al. 1996, Bennett et al. 1996], Device-Independent Quantum-Key-Distribution [Mayers and Yao 1998], Anonymous Transmission [Christandl and Wehner 2005], and Quantum Routing [Meignant et al. 2019; Priker et al. 2018]. Our approach could be used as an initial step in order to run this kind of protocol between an unknown subset of authorized parties. We note that even if our protocols are computationally secure, one may still be able to obtain unconditional guarantees, for example when the source is not corrupted.

Our protocol could also be used to achieve new functionalities, such as a Quantum Onion-like Routing protocol in order to route a quantum message through an untrusted quantum network, hiding the exact taken route. The idea would be to ask each intermediate server to blindly generate a large state in which a Bell pair is hidden: this Bell pair could then be used to teleport the qubit to the next server, without revealing its identity to the previous one.

Technical comparison. While the conceptual notion of ZK proofs on quantum states has never been considered in the past (at least to the best of our knowledge), and neither the notion of multipartite RSP, our method to obtain an efficient multi-qubit RSP protocol do generalizes some constructions introduced in previous work. For instance, our approach is based on an encryption similar to the multi-bit encryption of the dual-Regev encryption scheme [Gentry et al. 2008a], which were used in some RSP-related works [Brakerski 2018a; Cojocaru et al. 2019; Mahadev 2018a],

but only in the single-bit encryption regime. The generalization from using single-bit dual-Regev to multi-bit dual-Regev is an important step to obtain an efficient multi-qubit RSP protocol. Other related protocols like [Brakerski et al. 2018; Gheorghiu and Vidick 2019] use a different method to obtain claw-free functions based on a adaptive hardcore bit property, which is, as far as we understand, unrelated to the dual-Regev encryption and quite fundamentally relies on a series of tests that is not needed in our approach. Another noticeable difference between [Brakerski 2018a; Brakerski et al. 2018; Gheorghiu and Vidick 2019; Mahadev 2018a] versus this work and [Cojocaru et al. 2019] is that in the former the superposition is done by using different amplitudes, while in the later the superposition is done uniformly over the input space of the function (which we find conceptually simpler to interpret). Related to our contribution on ZK on quantum states, we are not aware of any other work combining classical ZK proofs with RSP protocols to provide guarantees against a malicious sender.

Historical update: OT from ZKoQS. After the first online appearance of this work, [Colisson et al. 2023] built upon our ZKoQS concept in order to obtain the first (quantum) Oblivious Transfer protocol—without structure as it is using only hash functions—with optimal round-complexity (2 messages). Note however that their ZKoQS protocol is very different: while they do not need the structure offered by LWE and can get statistical security (answering two open questions we raised here), they guarantee weaker assumptions on the received quantum state (which is however enough to obtain OT). In particular, using our construction, we could also obtain for free a 2-message Oblivious Transfer protocol, but the interest is less obvious as we rely on LWE, which is already known to imply 2-message classical OT [Brakerski and Döttling 2018]. Nevertheless, this opens interesting questions to characterize the set of quantum languages that fundamentally need additional structure to be prepared in a NIZK way.

1.3 Article Organization

In the Section 2, we give a generic overview of our result: Section 2.1 presents the role of the different parties involved in the target protocols that will benefit from NIZQoS and multiparty large state preparation, Section 2.2 gives the list of the different protocols we realize, and Section 2.3 gives an insight on the technical methods used in this article. The preliminaries can be found in Section 3. Then, we give in Section 4 the cryptographic requirements of our protocols. In Section 6, the reader will find the definition of our different protocols—including one desirable but impossible protocol—with our main protocol AUTH – BLIND_{can}^{dist} in Section 6.5. Section 5 defines NIZKoQS and proves that our protocols exhibit NIZKoQS. Finally, Section 7 explains how to construct the required cryptographic primitives: Section 7.1 shows how to obtain an δ -GHZ^H capable family, and Section 7.2 describes a generic way to construct a distributable δ' -GHZ^{can} capable family.

2 Technical Overview and Presentation of the Setup

2.1 Initial Setup

Our two first orthogonal results (efficient large state preparation and NIZKoQS) will ultimately be combined to obtain multiple protocols allowing the (secret) distribution of a GHZ state between n clients. The protocols involve several parties:

- There are n (weakly²) quantum *applicants* a_i . At the end of the protocols, a (secret) subset \mathcal{S} of these applicants—the *supported*³ applicants—will share a GHZ state.

²The applicants need only basic quantum skills, depending on the protocol, they may have nothing to do except receiving a state, or may need to apply a few gates.

³We also call them *supported* because they are part of the support of the hidden GHZ state. We may also refer to this as being the *support status* of an applicant.

- A quantum⁴ server, *Bob*. Bob is quantumly connected to all applicants, and it is the only quantum communication required in the upcoming protocols.
- *Cupid*⁵ is an honest classical party (the only party that will always be considered as honest when present. In our last protocol, we won't need Cupid anymore). Cupid will be in charge of choosing the supported applicants that will share a GHZ at the end of the protocol. We assume Cupid can communicate classically with the server, as well as with all applicants—using secure (i.e., confidential⁶ and authenticated) channels.

Note that if the supported applicants want to perform at the end some quantum protocols in order to use the previously shared GHZ state, *Cupid* may also be used as a proxy forwarding the messages to the good applicants in case classical communication is required between participants.⁷

Note also that these parties may not always be different entities. For example, when a user wants to send a qubit to a secret recipient, this user could be both considered as an applicant and as Cupid. Similarly, the server may want to be part of the applicants. Moreover, in our final protocol, the role of Cupid is distributed among all applicants: each applicant will be able to choose whether they want to be part of the GHZ or not, and in that setting all parties may potentially be malicious. However, since the simpler protocols could be of independent interest, we also study them separately.

2.2 List of the Protocols Proposed in This Work

In the following, a canonical GHZ state will be a state of the form $\frac{1}{\sqrt{2}}(|0 \dots 0\rangle + |1 \dots 1\rangle)$ (we will drop all normalization factors later), a generalized GHZ state is a GHZ state in which we applied some local X or Z gates, and a hidden GHZ is a permutation of a generalized GHZ state tensored with some $|0\rangle$ and $|1\rangle$ qubits.

We propose 4 different protocols in Section 6 (BLIND , $\text{BLIND}^{\text{sup}}$, $\text{BLIND}_{\text{can}}^{\text{sup}}$, and $\text{AUTH} - \text{BLIND}_{\text{can}}^{\text{dist}}$) to distribute a GHZ state. Note that, we show in Section 5 how BLIND and $\text{AUTH} - \text{BLIND}_{\text{can}}^{\text{dist}}$ can exhibit NIZKoQS.

$\text{AUTH} - \text{BLIND}_{\text{can}}^{\text{dist}}$ is the most involved protocol, and we also show Section 6.3 the impossibility of a desirable variant of these protocols, $\text{BLIND}_{\text{can}}$. All these protocols derive from the BLIND protocol, in which Cupid chooses the support status of each applicant, and at the end of the protocol each supported applicant is supposed to obtain a generalized GHZ state, while non-supported applicant obtain random not entangled qubits in the computational basis (at that step no applicants know if they are supported or not). The server is simply used as an intermediate party in charge of creating and distributing states. The other protocols differ slightly:

- The subscript $_{\text{can}}$ denotes the fact that at the end of the protocol each supported applicant ends up with a canonical GHZ state instead of a generalized GHZ state.

⁴Bob needs to perform quite heavy quantum computations.

⁵Besides having a name starting with a 'C', Cupid, the roman god of love, is famous for sending arrows at the heart of Humans to designate the beloved among the applicants.

⁶Since secure channels may still leak the length of the exchanged message, in this article, we assume that for a given round, all messages have the same size.

⁷In that case, it is also important to make sure that no adversary can detect the source and the destination of a message, otherwise it would be trivial for the adversary to obtain the set of supported applicants. This could be achieved by either asking to all non-supported applicants to also send/receive dummy encrypted messages to Cupid when a message is expected to be exchanged (it is the simplest solution but may not be the most efficient solution), by using a classical channel that is not controlled by the adversary to exchange these messages, by using a noisy cover channel in which communications can be made undetectable, or by relying on an anonymous onion-like routing protocol if the adversary controls only parts of the network and if enough unrelated messages are exchanged this way to make sure that the adversary cannot apply timing attacks/flow correlation attacks...

- The superscript \cdot^{sup} denotes the fact that at the end of the protocol each applicant knows their own support status, chosen by Cupid.
- In the protocol AUTH – BLIND $_{\text{can}}^{\text{dist}}$ each applicant chooses their own support status (Cupid is not required anymore), and the server can check that each participant is “authorized”, meaning that if they are supported, then they know a secret (note that this does *not* reveal to the server if a given applicant is supported). To obtain this last property, we use NIZKoQS.

In term of security, we typically expect that no malicious group of applicants, potentially colluding with the server, should learn the support status of honest applicants.⁸ In BLIND, we can even prove that applicants can’t even learn their own support status. Moreover, in AUTH – BLIND $_{\text{can}}^{\text{dist}}$ we show that an honest server is guaranteed that any supported applicant “knows” the secret, and that no malicious or noisy applicant can significantly alter the state obtained by honest applicants.

A simple application of the AUTH – BLIND $_{\text{can}}^{\text{dist}}$ protocol would be to allow Bob to teleport a quantum state $|\psi\rangle$ to an unknown applicant who knows a secret password, or who has access to a signature provided by a Certification Authority. The applicant would be allowed to hide its identity to Bob, and Bob can be sure that only applicants knowing the password/signature could obtain the state $|\psi\rangle$. In addition, if several applicants know the secret, then Bob is in fact secret sharing its qubit $|\psi\rangle$ among all applicants knowing this secret [Hillery et al. 1999]. Other applications are discussed at the end of Section 1.2.

2.3 Quick Overview of Our Methods

Cryptographic assumptions. We need to use a classical cryptographic family of functions $\{f_k: \mathcal{X} \rightarrow \mathcal{Y}\}_{k \in \mathcal{K}}$, together with a function $h: \mathcal{X} \rightarrow \{0, 1\}^n$ having several properties. The exact list of requirements is given in Section 4, but here are the important properties. For any $\mathbf{d}_0 \in \{0, 1\}^n$ (intuitively representing the support status of all applicants, with $\mathbf{d}_0[i] = 1$ iff the i th applicant is supported), we can generate using a function $\text{Gen}(1^\lambda, \mathbf{d}_0)$ an index k and a trapdoor t_k such that

- f_k is 2-to-1 (i.e., for all x , $|f_k^{-1}(f_k(x))| = 2$, we will also generalize this definition to approximate δ -2-to-1 functions).
- f_k can be efficiently computed given k , but should be hard to invert without t_k . Moreover, it should be hard to obtain any information on \mathbf{d}_0 given k .
- Given the trapdoor t_k , f_k can be efficiently inverted.
- For any $x \neq x'$ such that $f(x) = f(x')$, $h(x) \oplus h(x') = \mathbf{d}_0$.

We will say that such a family is δ -GHZ^H capable (more details in Definition 4.1). While these properties will be enough to understand our “binding” approach, in practice we will also require other properties in order to allow each applicant to choose their own support status. More precisely, we will require the existence of a local generation procedure $(k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$ such that $k = (k^{(1)}, \dots, k^{(n)})$ (since $k^{(i)}$ is enough to fix $\mathbf{d}_0[i]$, we may write $\mathbf{d}_0[i] := \mathbf{d}_0(k^{(i)})$). Moreover $t_k^{(i)}$ can be used to obtain partial information about the preimages of f_k : one can, for example, obtain the i th bit of $h(x)$ given $f_k(x)$ and $t_k^{(i)}$. These additional properties will be important to allow the creation of a multi-party state, and we will say that such family is a distributable δ -GHZ^{can} capable family.

Efficient large RSP. The first key point of our method is to bind a classical message to a quantum state. Instead of receiving directly a quantum state, the receiver receives classical instructions

⁸In BLIND $_{\text{can}}^{\text{sup}}$ and AUTH – BLIND $_{\text{can}}^{\text{dist}}$, we expect at least one supported applicant to be honest when the adversary is allowed to corrupt supported applicants (the identity of this honest applicant may be unknown to the adversary), otherwise there is a trivial attack against any such protocol.

encrypting a class of quantum states and derives from it the final quantum state. More precisely, since in our case we are interested in the preparation of hidden GHZ states (i.e., a state which is equal, up to some local operations, to a permutation of $(|0 \dots 0\rangle + |1 \dots 1\rangle) \otimes |0 \dots 0\rangle$), we proceed as follows: If denote by \mathbf{d}_0 the support of the GHZ state (i.e., qubit i is part of the GHZ iff $\mathbf{d}_0 = 1$), then the encryption of the quantum state will be k , where k is generated using $\text{Gen}(1^\lambda, \mathbf{d}_0)$. Then, in order to produce the quantum state, the receiver will run the unitary corresponding to $x \mapsto (h(x), f_k(x))$ in superposition, obtaining the state

$$\sum_x |x\rangle |h(x)\rangle |f_k(x)\rangle = \sum_y (|x\rangle |h(x)\rangle + |x'\rangle |h(x')\rangle) |y\rangle, \quad (1)$$

where in the right hand side, we sum over the elements y in the image of f_k , and (x, x') are the two preimages of y (reminder: the function is 2-to-1). After measuring the last register, we obtain a y , and the following quantum state (where x and x' are the two preimages of y):

$$|x\rangle |h(x)\rangle + |x'\rangle |h(x')\rangle. \quad (2)$$

Now, we measure the first register in the Hadamard basis, and we obtain

$$|h(x)\rangle + (-1)^\alpha |h(x')\rangle \quad (3)$$

for some $\alpha \in \{0, 1\}$, depending on the outcome of the measurement. This state is now a hidden GHZ state, whose support is equal to \mathbf{d}_0 due to the property $h(x) \oplus h(x') = \mathbf{d}_0$. Indeed, when $\mathbf{d}_0[i] = 0$, $h(x)[i] = h(x')[i]$, so we can factor out this qubit. More precisely, in order to come back to a canonical GHZ on the other qubits, one just needs to apply an X gate on the i th qubit for all i such that $1 = h(x)[i] \neq h(x')[i]$ (the trapdoor would be required to learn $h(x)$), followed by a Z operation on one supported qubit if $\alpha = 1$.

Note that here we managed to prepare a n -qubit state using a single superposition, which improves the efficiency of existing RSP protocols. With the actual implementation of f_k , it drops the required number of steps from $O(nMN)$ to $O((M+n)N)$ where $M > N$, N and M being very large parameters of f_k .

NIZQoQS. It is now possible to turn our Non-Interactive RSP protocol into what we call a NIZKoQS in order to prove properties on the generated state. We formalize this notion in Section 5, but roughly speaking, we say that a protocol is a ZKoQS protocol for a *quantum* language $\mathcal{L} \subseteq \mathcal{H}$ (representing the set of allowed quantum states) if it is:

- **Complete** (or correct): When the prover and verifier act honestly, the verifier ends up with a state $|\psi\rangle \in \mathcal{L}$.
- **Sound**: If the prover is malicious, then an honest verifier will obtain, when accepting the transcript, a state $|\psi\rangle \in \mathcal{L}$.
- **ZK**: A malicious verifier should be unable to learn (non-trivial things about) the classical description of the generated state $|\psi\rangle$. For instance, in our case we can hide the basis of the final states (in the one-qubit setting), or more generally the set of qubits that are part of the final GHZ state.

In practice, to obtain a NIZQoQS, the key element to note is that in our above protocol, a single *classical* message k is sent to the receiver. This message can be interpreted as the instructions to follow to produce a given hidden GHZ state whose support is \mathbf{d}_0 , and it is therefore possible to send a classical NIZK proof—using your favorite post-quantum NIZK construction—about these instructions to indirectly (and, therefore, non-destructively) obtain guarantees on the produced quantum state. More precisely, the sender simply needs to send, together with k , a NIZK proof (where t_k is part of the secret witness) proving that (i) the message k is indeed a δ -GHZ^H capable

function⁹ (ii) that $\text{Auth}(\mathbf{d}_0, w) = 1$, where Auth can be any efficiently computable function, and w any secret witness, depending on the wanted property on $\mathbf{d}_0(t_k)$. This last function and witness could be virtually anything, like ensuring that *There exists only two indices $i \neq j$ such that $\mathbf{d}_0(i) = \mathbf{d}_0(j) = 1$* (i.e., it proves that the final state contains only two entangled qubits forming a Bell state, w is not needed here), or *Either $\mathbf{d}_0[42] = 0$ or $(\mathbf{d}_0[42] = 1 \text{ and } I \text{ know the private key corresponding to the bitcoin wallet } 12c6DSiU4Rq3P4ZxziKxzl5LmMBrzjrjX)$* (i.e., it proves that the 42th qubit is entangled to the rest of the GHZ only if the sender is Satoshi Nakamoto... of course without revealing to the receiver if the sender is Satoshi Nakamoto; w being here the private key of the wallet). This kind of property will be particularly interesting in the AUTH – BLIND_{can}^{dist} protocol.

Usage in protocols. We can use the above ideas in order to achieve the above protocols, notably AUTH – BLIND_{can}^{dist}. Each applicant will be asked to sample $(k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0)$, then $k^{(i)}$ will be sent to the server. In order to prove that they are authorized, each applicant will also include a NIZK proof confirming, as explained in the last paragraph, that they know a classical witness w_i such that $\text{Auth}(\mathbf{d}_0[i], w_i) = 1$ (the NIZK proof also checks that $\mathbf{d}_0[i] = \mathbf{d}_0(k^{(i)})$ and that $k^{(i)}$ is honestly prepared). Since this protocol is a ZK protocol, the server will not be able to learn any information about \mathbf{d}_0 . In return, the server will have the guarantee that it will indeed produce a hidden GHZ state distributed among authorized applicants. Therefore, the server can run the quantum circuit described above, and distribute each qubit to the corresponding applicant. In order to come back to a canonical GHZ, each applicant will use their local trapdoor $t_k^{(i)}$ to compute $h(x)$ in order to apply the corresponding X correction. Moreover, in order to compute the α needed to apply the Z correction, all parties will need to run a **Multi-Party Computation (MPC)**. Since α could leak some information about the state, each supported applicant will in fact obtain from the MPC a linear secret sharing of α , i.e., a random $\hat{\alpha}_i$ such that $\bigoplus_{i \in \mathcal{S}} \hat{\alpha}_i = \alpha$ (\mathcal{S} being the set of supported applicants). That way, each supported applicant will be able to locally apply a $Z^{\hat{\alpha}_i}$ correction, and it will be equivalent to applying a single Z^α gate on the overall state.

How to construct the cryptographic assumptions. In order to implement the family $\{f_k\}_{k \in \mathcal{K}}$, we proceed in two steps: we first build a family which is δ -GHZ^H capable (see intuitive definition above, or Definition 4.1). The δ is needed because in practice the functions are only 2-to-1 with a probability $1 - \delta$, where δ can be made negligible by relying on the hardness of GapSVP _{γ} with a superpolynomial ratio γ . We then provide a quite generic method to turn a δ -GHZ^H capable family into a δ' -GHZ^{can} capable distributable family.

The security of our construction reduces to the hardness of the LWE assumption (more details in Section 3.5), and builds on ideas introduced in [Cojocar et al. 2019]. The starting point is the trapdoor construction provided by [Micciancio and Peikert 2012]. They provide an algorithm to generate a matrix $\mathbf{A} \in \mathbb{Z}_q^{M \times N}$ (q will be a power of two) that looks random, together with a trapdoor matrix \mathbf{R} . If the noise $\mathbf{e} \in \mathbb{Z}_q^M$ is sufficiently small,¹⁰ the function $g_{\mathbf{A}}(\mathbf{s}, \mathbf{e}) := \mathbf{A}\mathbf{s} + \mathbf{e}$ is injective. Moreover, given the trapdoor \mathbf{R} , one can easily invert the function $g_{\mathbf{A}}$, otherwise inverting this function is hard: $g_{\mathbf{A}}(\mathbf{s}, \mathbf{e})$ is indistinguishable from a random vector given \mathbf{A} .

From that, we can first see how to get a δ -2-to-1 family of functions. Note that the larger the noise \mathbf{e} is, the larger δ is. So a perfect (but insecure) 2-to-1 family would use $\mathbf{e} = \mathbf{0}$: therefore, to better understand this construction, it may be practical to imagine that $\mathbf{e} = \mathbf{0}$ during a first reading. The idea of the construction is to sample first an image vector $\mathbf{y}_0 := \mathbf{A}_u \mathbf{s}_0 + \mathbf{e}_0 \in \mathbb{Z}_q^M$ (by

⁹In our construction, it boils down to proving that the trapdoor t_k has small enough singular values.

¹⁰In the actual construction, we also require \mathbf{s} to be small because we rely on the equivalent but more efficient normal-form of LWE, but for simplicity we use the classic LWE problem in this overview.

sampling first \mathbf{s}_0 uniformly at random over \mathbb{Z}_q^N and $\mathbf{e}_0 \in \mathbb{Z}_q^m$ according to a discrete Gaussian $\mathcal{D}_{\mathbb{Z}, \alpha q}^m$, and then to define for all $\mathbf{s} \in \mathbb{Z}_q^N$, $c \in \{0, 1\}$ and errors \mathbf{e} belonging to some set $E \subseteq \mathbb{Z}_q^M$ (to be defined), the function $f_{\mathbf{A}_u, \mathbf{y}_0}(\mathbf{s}, \mathbf{e}, c \in \{0, 1\}) := \mathbf{A}_u \mathbf{s} + \mathbf{e} + c \times \mathbf{y}_0$. That way, if all vectors in E are small enough, $f_{\mathbf{A}_u, \mathbf{y}_0}$ has at most two preimages, one for $c = 0$ and one for $c = 1$: $f(\mathbf{s}, \mathbf{e}, 0) = \mathbf{A}_u \mathbf{s} + \mathbf{e} = \mathbf{A}_u(\mathbf{s} - \mathbf{s}_0) + (\mathbf{e} - \mathbf{e}_0) + \mathbf{y}_0 = f(\mathbf{s} - \mathbf{s}_0, \mathbf{e} - \mathbf{e}_0, 1)$. We remark that in order to have two preimages, we want to make sure that both $\mathbf{e} \in E$ and $\mathbf{e} - \mathbf{e}_0 \in E$: we will rely on LWE with superpolynomial modulus to noise ratio to ensure this is true for a superpolynomially large fraction of inputs. Then, in order to obtain the bit string $\mathbf{d}_0 \in \{0, 1\}^n$ and the function h such that for any two preimages x, x' , $h(x) \oplus h(x') = \mathbf{d}_0$, we will update the previous construction and now sample \mathbf{y}_0 as follows: We will first sample additional lines $\mathbf{A}_l \xleftarrow{\$} \mathbb{Z}_q^{n \times N}$ to add to the matrix \mathbf{A}_u . Then, we will sample \mathbf{s}_0 uniformly at random over \mathbb{Z}_q^N (as before) and $\mathbf{e}_0 \in \mathbb{Z}_q^{M+n}$ will be sampled according to the discrete Gaussian $\mathcal{D}_{\mathbb{Z}, \alpha q}^{M+n}$. Finally, if we also denote by \mathbf{d}_0 the binary vector in \mathbb{Z}_q^n composed of elements of \mathbf{d}_0 , we compute $\mathbf{y}_0 := \begin{bmatrix} \mathbf{A}_u \\ \mathbf{A}_l \end{bmatrix} \mathbf{s}_0 + \mathbf{e}_0 + \frac{q}{2} \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d}_0 \end{bmatrix}$. We update similarly our function f by adding a parameter $\mathbf{d} \in \{0, 1\}^n$:

$$f_{\mathbf{A}_u, \mathbf{A}_l, \mathbf{y}_0}(\mathbf{s}, \mathbf{e}, c, \mathbf{d}) := \begin{bmatrix} \mathbf{A}_u \\ \mathbf{A}_l \end{bmatrix} \mathbf{s} + \mathbf{e} + \frac{q}{2} \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d} \end{bmatrix} + c \times \mathbf{y}_0. \quad (4)$$

Then, we can remark that, because q is even:

$$f_{\mathbf{A}_u, \mathbf{A}_l, \mathbf{y}_0}(\mathbf{s}, \mathbf{e}, 0, \mathbf{d}) = \begin{bmatrix} \mathbf{A}_u \\ \mathbf{A}_l \end{bmatrix} \mathbf{s} + \mathbf{e} + \frac{q}{2} \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d} \end{bmatrix} \quad (5)$$

$$= \begin{bmatrix} \mathbf{A}_u \\ \mathbf{A}_l \end{bmatrix} (\mathbf{s} - \mathbf{s}_0) + (\mathbf{e} - \mathbf{e}_0) + \frac{q}{2} \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d} \oplus \mathbf{d}_0 \end{bmatrix} + \mathbf{y}_0 \quad (6)$$

$$= f_{\mathbf{A}_u, \mathbf{A}_l, \mathbf{y}_0}(\mathbf{s} - \mathbf{s}_0, \mathbf{e} - \mathbf{e}_0, 1, \mathbf{d} \oplus \mathbf{d}_0) \quad (7)$$

and that, therefore, (skipping a small technicality) for any two preimages $(\mathbf{s}, \mathbf{e}, 0, \mathbf{d})$ and $(\mathbf{s}', \mathbf{e}', 1, \mathbf{d}')$, we have $\mathbf{d} \oplus \mathbf{d}' = \mathbf{d}_0$. So we just need to define $h(\mathbf{s}, \mathbf{e}, c, \mathbf{d}) = \mathbf{d}$ to get the XOR property. Finally, the indistinguishability property is simple to obtain: since \mathbf{A}_u is indistinguishable from a random matrix, and \mathbf{A}_l is actually a random matrix, thus $\mathbf{A} := \begin{bmatrix} \mathbf{A}_u \\ \mathbf{A}_l \end{bmatrix}$ is indistinguishable from a random matrix. Therefore, under the decision-LWE $_{q, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ assumption, $\mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$ is indistinguishable from a random vector. Therefore, since adding a constant vector to a uniformly sampled vector does not change its distribution, one cannot distinguish $\mathbf{A}\mathbf{s}_0 + \mathbf{e}_0$ from $\mathbf{A}\mathbf{s}_0 + \mathbf{e}_0 + \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d}_0 \end{bmatrix}$, or from any vector of the form $\mathbf{A}\mathbf{s}_0 + \mathbf{e}_0 + \begin{bmatrix} \mathbf{0}^M \\ \mathbf{d} \end{bmatrix}$. We provide in Section 7.1 a more in-depth analysis in order to properly handle the noise, we find an explicit set of parameters allowing f to be $\text{negl}(\lambda)$ -2-to-1, and we give a method to prove that a maliciously sampled f is indeed $\text{negl}(\lambda)$ -2-to-1 and has the XOR property. This is required to turn it into a distributable δ' -GHZ^{can} capable family.

Finally, one can obtain a distributable δ' -GHZ^{can} capable family from a δ -GHZ^H capable family (which has an additional property that the two preimages have the form $x = (0, \bar{x})$ and $x' = (1, \bar{x}')$). The idea is to sample one δ -GHZ^H capable function for each $\mathbf{d}_0[i]$. Then, we can define our new function to be $f_{k(1), \dots, k(n)}(c, \bar{x}^{(1)}, \dots, \bar{x}^{(n)}) = f_{k(1)}(c, \bar{x}^{(1)}) \dots |f_{k(n)}(c, \bar{x}^{(n)})$. More details can be found in Section 7.

2.4 Open Questions

Since this work introduces NIZKoQS, there are many questions left open:

- **Characterization of the quantum languages verifiable non-interactively:** Using classical-client **Universal Blind Quantum Computing (UBQC)** and ZK proofs, classical-client **Quantum Fully Homomorphic Encryption (QFHE)** and ZK proofs, or directly

quantum MPC (QMPC), one could possibly prove ZKoQS statements on arbitrary languages (even using three messages with the QFHE+ZK solution). However, this must be written formally, in particular by proving that revealing additional information regarding the quantum state like the one-time pad does not weaken the security. Moreover, characterizing completely \equiv with either possibility or impossibility results \equiv the set of quantum languages that can be verified *non-interactively* is an open question. This article solves this problem for a specific quantum language, and using QFHE+ZK we could possibly characterize a broader set of quantum languages (basically any property which is true up to a quantum one-time pad, characterizing notably entanglement). However, the question of the feasibility of NIZKoQS for broader quantum languages is left open.

- **Minimal assumptions:** In the present article, we rely heavily on the hardness of the LWE problem. However, studying the properties obtainable with weaker assumptions (like one-way functions) could be a thrilling exploration path.
- **Statistical vs. computational soundness:** In the protocol presented in this article, an unbounded verifier could learn completely the produced quantum state. Finding NIZKoQS protocols secure against an unbounded verifier could be another interesting question, with application to statistically secure protocols.
- **Applications:** We have shown how NIZKoQS can help boosting the round-efficiency of protocols sharing hidden GHZ states. However, we expect NIZKoQS to be of paramount importance to achieve optimal round complexity in a much wider variety of protocols as they basically avoid cut-and-choose approaches. Similarly, applications of our protocols sharing GHZ states should be further explored.
- **RSP efficiency and applications:** We showed how RSP protocols can produce multi-qubits states with an asymptotically quadratic improvement. However, further improving this bound (potentially using LWE with a polynomial noise ratio) and finding multi-qubit states that are both efficient to prepare and compatible with blind quantum computing is an exciting challenge.

3 Preliminaries

3.1 Notation

The applicants and Cupid are designated with the gender neutral pronoun *they*, while the server—a machine—is designated by *it*.

If $\mathbf{b} \in \{0, 1\}^n$ is a bit string or a vector, and i is an integer, we may use \mathbf{b}_i or ¹¹ $\mathbf{b}[i]$ to denote the i th bit of \mathbf{b} . The symbol $|$ will be used to denote string concatenation, and we define $\bar{\mathbf{b}} := (b_1 \oplus 1) | \dots | (b_n \oplus 1)$. We define also $X^b := X^{b[1]} \otimes X^{b[2]} \dots \otimes X^{b[n]}$ and Z_i is the operator that applies the Z gate on the i th qubit, i.e., $Z_i = \underbrace{I \otimes \dots \otimes I}_{i-1} \otimes Z \otimes I \otimes \dots \otimes I$. For any permutation σ ,

we denote by PERM_σ the quantum unitary that performs a permutation of the qubits, such that the i th qubit is mapped on $\sigma(i)$. For any bit strings x, x' and b of the same length, we also define $\langle b, x \rangle := \bigoplus_i b_i x_i$, $x \oplus x' := (x_1 \oplus x'_1) | \dots | (x_n \oplus x'_n)$.

\mathbb{R} is the set of reals, \mathbb{Z} the set of integers, \mathbb{Z}_q is the set of integers modulo q and $[n] := \{1, \dots, n\}$. For any element $x \in \mathbb{Z}_q$, if \hat{x} is the representant of x in $[-\frac{q}{2}, \frac{q}{2})$ we will define the modular

¹¹Note however that we often refer in this article to a bit string $\mathbf{d}_0 \in \{0, 1\}^n$ (we use this notation in order to keep a consistent notation with previous works), so to avoid confusion between \mathbf{d}_0 (the bit string) and d_0 (the first bit of d), we will always use the bracket notation $\mathbf{d}[i]$ or $\mathbf{d}_0[i]$ when the d letter is involved. Since \mathbf{d}_0 will also be interpreted as a vector in the second part of the article, and to highlight the difference with a previous work in which d_0 was a bit string, we will use a bold font for this kind of bit string.

rounding $\text{RoundMod}_q(x) = 0$ if $\hat{x} \in [-\frac{q}{4}, \frac{q}{4}]$ and $\text{RoundMod}_q(x) = 1$ otherwise. We extend this notation when \mathbf{x} is a vector by doing it component wise. Similarly, for any element $\mathbf{x} \in \mathbb{Z}_q^n$, $\|\mathbf{x}\|_\infty$ will be understood in a modular way: $\|\mathbf{x}\|_\infty := \max_i \hat{x}[i]$. For any vector $\mathbf{x} \in \mathbb{R}^n$, \mathbf{x}^T is the transpose operation, $\|\mathbf{x}\|_2 := \sqrt{\mathbf{x}^T \mathbf{x}}$ is the Euclidean norm of \mathbf{x} . For any matrix $\mathbf{R} \in \mathbb{R}^{n \times m}$, $\|\mathbf{R}\|_2 := \sup_{\mathbf{x}} \|\mathbf{R}\mathbf{x}\|_2 / \|\mathbf{x}\|_2$ is the spectral norm of \mathbf{R} , and $\sigma_{\max}(\mathbf{R})$ is the highest singular value of \mathbf{R} . Note that $\sigma_{\max}(\mathbf{R}) = \|\mathbf{R}\|_2$. For an Hilbert space \mathcal{H} , $\mathcal{L}_o(\mathcal{H})$ is the set of density operators on \mathcal{H} (positive linear operators on \mathcal{H} with trace 1).

PPT stands for Probabilistic Polynomial Time and QPT stands for Quantum Polynomial Time. We denote by $\text{negl}(\lambda)$ any function that goes asymptotically to 0 faster than any inverse polynomial in λ : we say that this function is *negligible*. We say that a probability is *overwhelming* when it is $1 - \text{negl}(\lambda)$. We use the Landau notation $O(\cdot)$, and we say that $f(x) = \tilde{O}(g(x))$ if there exists k such that $f(x) = O(g(x) \log^k g(x))$. We use $\omega(\sqrt{\log N})$ to denote a function, fixed across all the article, such that $\lim_{N \rightarrow \infty} \sqrt{\log N} / \omega(\sqrt{\log N}) = 0$. For instance, we can take $\omega(\sqrt{\log N}) = \log N$.

For any finite set X , we denote $x \xleftarrow{\$} X$ when $x \in X$ is sampled uniformly at random. The uniform distribution on X is also denoted by $\mathcal{U}(X)$, and the statistical distance between the distributions \mathcal{D}_1 and \mathcal{D}_2 is $\Delta(\mathcal{D}_1, \mathcal{D}_2) := \frac{1}{2} \sum_x |\mathcal{D}_1(x) - \mathcal{D}_2(x)|$. For any distribution \mathcal{D} , we write $x \leftarrow \mathcal{D}$ if x is sampled according to \mathcal{D} , and $y := f(x)$ when f is deterministic or if it is the definition of y . If an adversary \mathcal{A} interacts n times with its environment, without loss of generality we may decompose $\mathcal{A} = (\mathcal{A}_1, \dots, \mathcal{A}_n)$ and denote $y_1 \leftarrow \mathcal{A}_1(x_1); y_2 \leftarrow \mathcal{A}_2(x_2); \dots; y_n \leftarrow \mathcal{A}_n(x_n)$ where we implicitly assume that \mathcal{A}_i gives its internal state to \mathcal{A}_{i+1} . In the hybrid games, we will also strike out some lines to denote differences between two games: it only removes the content of the current line, without changing the line numbering of the other lines.

3.2 Generalized Hidden GHZ State

In this article, we assume basic familiarity with quantum computing [Nielsen and Chuang 2000]. The GHZ quantum state [Greenberger et al. 1989] is a generalization of the Bell state on multiple qubits. This state turns out to be useful in many applications, going from Quantum Secret Sharing [Hillery et al. 1999] to Anonymous Transmission [Christandl and Wehner 2005].

Definition 3.1 ((Canonical) GHZ State). We denote by $|\text{GHZ}_n\rangle$ a n qubits (canonical) GHZ state, i.e., a state of the form

$$|\text{GHZ}_n\rangle := \frac{1}{\sqrt{2}}(|00 \dots 0\rangle + |11 \dots 1\rangle).$$

Definition 3.2 (GHZ^G : Generalized GHZ State). A quantum state on n qubits is said to be a *generalized GHZ state* (sometimes abbreviated as GHZ^G), and usually denoted $|\text{GHZ}_{\alpha, \mathbf{d}}^G\rangle$ if it is a GHZ state up to an eventual π phase, and up to local X gates, i.e., if there exist a bit string $\mathbf{d} \in \{0, 1\}^n$ and $\alpha \in \{0, 1\}$ such that

$$|\text{GHZ}_{\alpha, \mathbf{d}}^G\rangle := X^{\mathbf{d}} Z_1^\alpha |\text{GHZ}_n\rangle = \frac{1}{\sqrt{2}}(|\mathbf{d}\rangle + (-1)^\alpha |\bar{\mathbf{d}}\rangle).$$

We call (α, \mathbf{d}) the GHZ key, since with the help of (α, \mathbf{d}) it is possible to turn a generalized GHZ state into a canonical GHZ state using only local X and Z gates. The term “key” is used since the generalized GHZ state can be seen as a kind of **one-time padded (OTP)** canonical GHZ state.

Definition 3.3 (GHZ^H : Hidden Generalized GHZ State). A quantum state $|\phi\rangle$ on n qubits is said to be a *hidden (generalized) GHZ state* (sometimes abbreviated as GHZ^H) if it is a permutation of a

state composed of one generalized GHZ state and the tensor product of qubits in the computational basis, i.e., if there exist a permutation σ , a bit $\alpha \in \{0, 1\}$, a bit string \mathbf{d} , and two integers $(l, m) \in \mathbb{N}^2$ such that

$$|\phi\rangle = \text{PERM}_\sigma(|\text{GHZ}^G_{\alpha, \mathbf{d}}\rangle \otimes |0\rangle^l \otimes |1\rangle^m).$$

Equivalently, it means that there exist two bit strings $(\mathbf{d}, \mathbf{d}') \in (\{0, 1\}^n)^2$ and a bit $\alpha \in \{0, 1\}$ such that

$$|\phi\rangle = \frac{1}{\sqrt{2}}(|\mathbf{d}\rangle + (-1)^\alpha |\mathbf{d}'\rangle) = |\text{GHZ}^H_{\alpha, \mathbf{d}, \mathbf{d}'}\rangle.$$

For such a state $|\text{GHZ}^H_{\alpha, \mathbf{d}, \mathbf{d}'}\rangle$, the qubits that are not in the computational basis (and therefore that are part of the original GHZ^G state) are said to be in the *support* of the GHZ^H state, and it is easy to see that the support is

$$\text{supp}(|\text{GHZ}^H_{\alpha, \mathbf{d}, \mathbf{d}'}\rangle) := \{i \mid \mathbf{d}[i] \neq \mathbf{d}'[i]\}.$$

We will assimilate the string $\mathbf{d}_0 := \mathbf{d} \oplus \mathbf{d}'$ (bitwise XOR) to the support of the GHZ^H since $\mathbf{d}_0[i] = 1$ iff $i \in \text{supp}(|\text{GHZ}^H_{\alpha, \mathbf{d}, \mathbf{d}'}\rangle)$. Moreover, we will call $(\alpha, \mathbf{d}, \mathbf{d}')$ the *GHZ key* since it is possible to come back to a canonical GHZ state tensored with some $|0\rangle$ qubits by using $(\alpha, \mathbf{d}, \mathbf{d}')$.

3.3 Classical ZK Proofs and Arguments for NP

In the first (NIZKoQS) and last protocol presented in this article, we need to use a classical ZK protocol for NP (to obtain NIZKoQS, we also expect the protocol to be non-interactive, but interactivity does not change security or correctness). Intuitively, in a ZK protocol for a language $\mathcal{L} \in \text{NP}$, a prover must convince a verifier that a word x belongs to \mathcal{L} , in such a way that the verifier should not learn anything more about x beyond the fact that x belongs to \mathcal{L} . Because \mathcal{L} is in NP, \mathcal{L} is described by a relation $\mathcal{R}_\mathcal{L}$, in such a way that a word x belongs to \mathcal{L} iff there exists a witness w such that $w \in \mathcal{R}_\mathcal{L}(x)$. Moreover, deciding if a witness w belongs to $\mathcal{R}_\mathcal{L}(x)$ must be doable in polynomial time.

We will now formalize the above security statements, taking definition from [Bitansky and Shmueli 2020]. But first, we introduce a few notations: if (P, V) is a protocol between two parties P and V , we denote by $\text{OUT}_V(P, V)(x)$ the output of the party V after having followed the protocol with P , x being a common input of P and V . An honest verifier outputs a single bit (1 if they accept and 0 if they reject), but a malicious verifier can output an arbitrary quantum state.

Definition 3.4 (Computational Indistinguishability [Bitansky and Shmueli 2020]). Two maps of quantum random variables $X := \{X_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$ and $Y := \{Y_i\}_{\lambda \in \mathbb{N}, i \in I_\lambda}$ over the same set of indices $I = \cup_{\lambda \in \mathbb{N}} I_\lambda$ are said to be *computationally indistinguishable*, denoted by $X \approx_c Y$, if for any non-uniform quantum polynomial-time distinguisher $D := \{(D_\lambda, \rho_\lambda)\}_{\lambda \in \mathbb{N}}$, there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, $i \in I_\lambda$,

$$|\Pr[D_\lambda(X_i, \rho_\lambda) = 1] - \Pr[D_\lambda(Y_i, \rho_\lambda) = 1]| \leq \mu(\lambda). \quad (8)$$

Definition 3.5 (Post-Quantum ZK Classical Protocol [Bitansky and Shmueli 2020, Definitions 2.1 and 2.6]). Let (P, V) be a protocol between an honest PPT prover P and an honest PPT verifier V . Then (P, V) is said to be a Post-Quantum ZK Classical Protocol for a language $\mathcal{L} \in \text{NP}$ if the following properties are respected:

- (1) **Perfect Completeness:** For any $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cap \{0, 1\}^\lambda$, $w \in \mathcal{R}_\mathcal{L}(x)$,

$$\Pr[\text{OUT}_V(P(w), V)(x) = 1] = 1. \quad (9)$$

(2) **Soundness:** The protocol satisfies one of the following:

- **Computational Soundness:** For any non-uniform QPT malicious prover $P^* = \{(P_\lambda^*, \rho_\lambda)\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for any security parameter $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,

$$\Pr [\text{OUT}_V(P_\lambda^*(\rho_\lambda), V)(x) = 1] \leq \mu(\lambda). \quad (10)$$

A protocol with computational soundness is called an argument.

- **Statistical Soundness:** There exists a negligible function $\mu(\cdot)$ such that for any unbounded prover P^* , any security parameter $\lambda \in \mathbb{N}$ and any $x \in \{0, 1\}^\lambda \setminus \mathcal{L}$,

$$\Pr [\text{OUT}_V(P^*, V)(x) = 1] \leq \mu(\lambda). \quad (11)$$

A protocol with statistical soundness is called a proof.

(3) **Quantum ZK:** There exists a QPT simulator Sim such that for any QPT verifier $V^* = \{(V_\lambda^*, \rho_\lambda)\}_{\lambda \in \mathbb{N}}$,

$$\{\text{OUT}_{V_\lambda^*}(P(w), V_\lambda^*(\rho_\lambda))(x)\}_{\lambda, x, w} \approx_c \{\text{Sim}(x, V_\lambda^*, \rho_\lambda)\}_{\lambda, x, w}, \quad (12)$$

where $\lambda \in \mathbb{N}$, $x \in \mathcal{L} \cup \{0, 1\}^\lambda$, $w \in \mathcal{R}_{\mathcal{L}}(x)$, and V^* is given to Sim by sending the circuit description of V^* .

A Non-Interactive ZK protocol will be denoted NIZK.

In our last protocol, in order to get stronger guarantees, we may also want to ensure that the prover “knows” the secret. This is known as Proof of Knowledge, but since we do not really rely on it in the proofs, we only describe it in more detail in Definition A.1.

Several Post-Quantum ZK protocols have been proposed in the literature [Bitansky and Shmueli 2020; Unruh 2012; Watrous 2009] and have been shown to obey properties similar to both Definitions A.1 and 3.5. Moreover, [Don et al. 2019; Liu and Zhandry 2019] explain how to obtain quantum-secure NIZK (which are also Proof of Knowledge) using the Fiat–Shamir transformation and the hardness of the LWE problem in a Quantum Random Oracle model. In the following, we are agnostic of the used NIZK protocol and we assume the existence of a NIZK protocol obeying Definitions A.1 and 3.5.

3.4 Classical MPCs

In the last protocol presented in this article, we also need to use a classical¹² MPC protocol Π . An MPC protocol works as follows: given n (public and deterministic) functions (f_1, \dots, f_n) , at the end of the protocol involving n parties P_1, \dots, P_n , we expect party P_i to get $f_i(x_1, \dots, x_n)$, where x_j is the (secret) input of the party P_j . Moreover, we expect that no party can learn anything more than what they can already learn from $f_i(x_1, \dots, x_n)$. For simplicity, we define $f(x_1, \dots, x_n) = (f_1(x_1, \dots, x_n), \dots, f_n(x_1, \dots, x_n))$.

We can formalize the above security statements using the usual (quantum) real/ideal world paradigm. Informally, the protocol Π will be said to be secure if it is impossible to distinguish two “worlds”. On the one hand, we have a *real world* in which an adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ can corrupt a subset $\mathcal{M} \subseteq [n]$ of parties and interact in an arbitrary way with the other honest parties. On the other hand, we have an *ideal world*, in which a simulator Sim interacts with a functionality, this functionality behaving as a trivially-secure trusted third-party. If these two worlds are indistinguishable, it means that the protocol is “secure” because any secret obtained from the real world would also be obtainable from the ideal world... which is impossible because the ideal world is trivially secure.

¹²But of course post-quantum secure.

More precisely, we define, following [Agarwal et al. 2021], the real and ideal world as follows, where $\vec{x} := (x_1, \dots, x_n)$ is the inputs of the parties:

Definition 3.6 ($\text{REAL}_{\Pi, \mathcal{A}}(\lambda, \vec{x}, \rho_\lambda)$). \mathcal{A}_λ is given ρ_λ , and gives a subset $\mathcal{M} \subseteq [n]$ of corrupted (malicious) parties. Then \mathcal{A}_λ receives the inputs x_i of all corrupted parties P_i ($i \in \mathcal{M}$), sends and receives all the messages on the behalf of these corrupted parties, and communicates in an arbitrary quantum polynomial time way with the honest parties that follow the protocol Π . At the end of the interaction, \mathcal{A}_λ outputs an arbitrary state ρ , and we define as \vec{y} the output of the honest parties P_j , $j \notin \mathcal{M}$. Finally, we define $\text{REAL}_{\Pi, \mathcal{A}}(\lambda, \vec{x}, \rho_\lambda)$ as the random variable corresponding to (ρ, \vec{y}) .

Definition 3.7 ($\text{IDEAL}_{f, \text{Sim}}(\lambda, \vec{x}, \rho_\lambda)$). Sim (playing the role of the adversary) receives ρ_λ , outputs a set $\mathcal{M} \subseteq [n]$ of corrupted parties, interacts with a trusted party (called the ideal functionality) defined below, and outputs at the end a state ρ . The ideal functionality also outputs at the end a message \vec{y} corresponding to the output of the trusted party. We then define $\text{IDEAL}_{f, \text{Sim}}(\lambda, \vec{x}, \rho_\lambda)$ as the random variable corresponding to (ρ, \vec{y}) . Now we define the ideal functionality:

- The ideal functionality receives the set $\mathcal{M} \subseteq [n]$ a subset of corrupted parties, and for each party P_i , it receives an input x'_i : if P_i is honest ($i \notin \mathcal{M}$), we have $x'_i = x_i$, otherwise x'_i can be arbitrary.
- Then, it computes $(y_1, \dots, y_n) := f(x'_1, \dots, x'_n)$, and sends $\{(i, y_i)_{i \in \mathcal{M}}\}$ to the simulator.
- The simulator can choose to abort by sending a message \perp to the ideal functionality. Otherwise, it sends a “continue” message \top . If the message received by the ideal functionality is \perp , then it outputs \perp to each honest party, which is formalized by outputting $\vec{y} := \{(i, \perp)\}_{i \notin \mathcal{M}}$. Otherwise, it outputs $\vec{y} := \{(i, y_i)\}_{i \notin \mathcal{M}}$.

Definition 3.8 (Secure MPC [Agarwal et al. 2021]). Let f be a deterministic function with n inputs and n outputs, and Π be an n -party protocol. Protocol Π securely computes f if for every non-uniform quantum polynomial-time adversary $\mathcal{A} = \{\mathcal{A}_\lambda\}_{\lambda \in \mathbb{N}}$ corrupting a set of at most $n - 1$ players, there exists a non-uniform quantum polynomial-time ideal-world adversary Sim such that for any combination of inputs $\vec{x} \in (\{0, 1\}^*)^n$ and any non-uniform quantum advice $\rho = \{\rho_\lambda\}_{\lambda \in \mathbb{N}}$,

$$\{\text{REAL}_{\Pi, \mathcal{A}}(\lambda, \vec{x}, \rho_\lambda)\}_{\lambda \in \mathbb{N}} \approx_c \{\text{IDEAL}_{f, \text{Sim}}(\lambda, \vec{x}, \rho_\lambda)\}_{\lambda \in \mathbb{N}}. \quad (13)$$

3.5 Introduction to the Learning With Errors (LWE) Problem

The LWE problem. The LWE problem was introduced in [Regev 2005]. We briefly recall the definitions here, more details and useful lemmas can be found in Section B.

Definition 3.9 (Continuous and Discrete Gaussian). For any $s \in \mathbb{R}_{>0}$ and any vector $\mathbf{x} \in \mathbb{R}^N$, we define $\rho_s(\mathbf{x}) := \exp(-\pi \mathbf{x}^T \mathbf{x} / s^2)$. We define the continuous Gaussian distribution on \mathbb{R}^N as $\mathcal{D}_s^N(\mathbf{x}) := \rho_s(\mathbf{x}) / s^N$ and the discrete Gaussian distribution on \mathbb{Z}_q^N as $\mathcal{D}_{\mathbb{Z}, s}^N(\mathbf{x}) := \rho_s(\mathbf{x}) / (\sum_{\mathbf{y} \in \mathbb{Z}_q^N} \rho_s(\mathbf{y}))$. Note that sampling from \mathcal{D}_s^N is equivalent to sampling each coordinate from $\mathcal{D}_s^1 := \mathcal{D}_s^1$.

Definition 3.10 (LWE Problem). Let (N, M) be two integers, $\alpha \in (0, 1)$, $q = q(N) \in \mathbb{N}_{\geq 2}$ be a modulus. For a distribution χ on \mathbb{R} (respectively, on \mathbb{Z}_q), the decision-LWE $_{q, \chi}$ problem (matrix version) is to distinguish a sample of the uniform distribution $\mathcal{U}(\mathbb{Z}_q^{M \times N} \times [0, q]^M)$ (respectively, $\mathcal{U}(\mathbb{Z}_q^{M \times N} \times \mathbb{Z}_q^M)$) from $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$, where $\mathbf{A} \in \mathbb{Z}_q^{M \times N}$ and $\mathbf{s} \in \mathbb{Z}_q^N$ are sampled uniformly at random, and \mathbf{e} is sampled from χ^M .

The decision-LWE $_{q, \mathcal{D}_{\alpha q}}$ problem is widely supposed to be hard to solve even for quantum computers and is the basic building block of many post-quantum cryptographic protocols [Peikert 2016]. In particular, it is on average as hard as worst-case problems on lattices.

LEMMA 3.11 (HARDNESS OF LWE [PEIKERT ET AL. 2017]). *Let N, q be integers and $\alpha \in (0, 1)$ be such that $\alpha q > 2\sqrt{N}$. If there exists an algorithm that solves decision-LWE $_{q, \mathcal{D}_{\alpha q}}$, then there exists an efficient quantum algorithm that approximates the decision version of the shortest vector problem (GapSVP $_{\gamma}$) and the shortest independent vectors problem (SIVP $_{\gamma}$) to within $\gamma := \tilde{O}(N/\alpha)$.*

As there is no known algorithm to solve GapSVP $_{\gamma}$ for $\gamma = \tilde{O}(2^{N^\epsilon})$ and $\epsilon \in (0, \frac{1}{2})$, we will assume (like [Boneh et al. 2014]) that GapSVP $_{\gamma}$ is secure for this kind of superpolynomial γ (see Section B.2 for more details).

While the hardness assumption given in Lemma 3.11 targets a continuous noise distribution, it can also be adapted to discrete Gaussians (note that the result is trivial for less used “rounded Gaussians” [Regev 2005, Lemma 4.3]).

In particular, [Peikert 2010, Theorem 3.1] can be used to show that if decision-LWE $_{q, \mathcal{D}_{\alpha q}}$ is hard, then the discrete version decision-LWE $_{q, \mathcal{D}_{\mathbb{Z}_s}}$ is hard for $s := \sqrt{(\alpha q)^2 + \omega(\sqrt{\log \lambda})^2}$ (details in Corollary B.4). Moreover, we can also choose to sample \mathbf{s} according to $\mathcal{D}_{\mathbb{Z}_s}$ without weakening the security: this is known as the **Hermite Normal Form (HNF)** (details in Lemma B.5).

The [Micciancio and Peikert 2012] construction. In order to realize the primitives described in Definitions 4.1 and 4.3, we will use the trapdoor system presented in [Micciancio and Peikert 2012]. This work introduced an algorithm MP.Gen(1^λ) that samples a matrix \mathbf{A} and a trapdoor \mathbf{R} . In addition, \mathbf{A} is indistinguishable from a random matrix (without \mathbf{R}), and $g_{\mathbf{A}}(\mathbf{s}, \mathbf{e}) := \mathbf{A}\mathbf{s} + \mathbf{e}$ is injective and can be inverted given \mathbf{R} for any $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}$ (\mathcal{X} will be defined later as sets of elements having a small norm). The details of the construction and parameters required by the [Micciancio and Peikert 2012] construction are provided in Section B.

4 Cryptographic Requirements

All the protocols are based on the existence of a (post-quantum secure) cryptographic family of functions, which is said to be δ -GHZ^H capable. Intuitively, in the definition below, the string \mathbf{d}_0 is such that $\mathbf{d}_0[i] = 1$ iff applicant a_i is part of the support of the final GHZ):

Definition 4.1. Let $\lambda \in \mathbb{N}$ be a security parameter, and $n \in \mathbb{N}$. We say that a family of functions $\{f_k : \mathcal{X}_\lambda \rightarrow \mathcal{Y}_\lambda\}_{k \in \mathcal{K}_\lambda}$ with $\mathcal{X}_\lambda \subseteq \{0, 1\}^l$ is δ -GHZ^H capable if there exists a function $h : \mathcal{X}_\lambda \rightarrow \{0, 1\}^n$ (h could be extended to depend on k) such that the following properties are respected:

- **efficient generation**: for all $\mathbf{d}_0 \in \{0, 1\}^n$ a PPT machine can efficiently sample $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0)$ to generate (with overwhelming probability) an index $k \in \mathcal{K}_\lambda$ and a trapdoor $t_k \in \mathcal{T}_\lambda$.
- **efficient computation**: for any index k , the function f_k is efficiently computable by a PPT algorithm $\text{Eval}(k, x)$.
- **trapdoor**: for any trapdoor t_k and any y , there exists a procedure Invert that efficiently inverts f_k when y has two preimages. More precisely, if y has exactly two distinct preimages, we have $\text{Invert}(t_k, y) = f^{-1}(y)$. If the number of preimages is not 2, we expect $\text{Invert}(t_k, y) = \perp$.
- **quantum input superposition**: there exists a QPT algorithm that, on input 1^λ generates a uniform superposition $\sum_{x \in \mathcal{X}_\lambda} |x\rangle$.

For instance, if $\mathcal{X}_\lambda = \{0, 1\}^l$, it can be trivially done by computing $H^{\otimes l} |0\rangle^{\otimes l}$.

- **δ -2-to-1**¹³: for all $k \in \mathcal{K}$, when sampling an input x uniformly at random in \mathcal{X}_λ , the probability that $y := f_k(x)$ has exactly two distinct preimages (denoted by x_y and x'_y or simply x and x') is larger than $1 - \delta$. When $\delta = 0$, we just say that the function is 2-to-1.

¹³This kind of function is sometimes said to be δ -2-regular.

- **XOR of h :** for all k , there exists $\mathbf{d}_0 \in \{0, 1\}^n$ such that for all y , if y has exactly 2 distinct preimages x and x' (i.e., $f_k^{-1}(y) = \{x, x'\}$ with $x \neq x'$), then

$$h(x) \oplus h(x') = \mathbf{d}_0$$

Moreover, if k was obtained from $\text{Gen}(1^\lambda, \mathbf{d}_0^*)$, then $\mathbf{d}_0 = \mathbf{d}_0^*$. We will always assume that \mathbf{d}_0 is easy to obtain from t_k (it is always possible to append \mathbf{d}_0 to t_k). Since, fixing k fixes also \mathbf{d}_0 , in the following we may use interchangeably $\mathbf{d}_0(k)$, $\mathbf{d}_0(t_k)$ or simply \mathbf{d}_0 .

- **indistinguishability:** If $\text{Gen}(1^\lambda, \cdot)$ is seen as an encryption function, then, similarly to IND-CPA security (see for example [Goldreich 2004]), a quantum adversary cannot learn \mathbf{d}_0 from the index k , seen as an “encryption” of \mathbf{d}_0 . More formally, if we formulate it as an indistinguishability game $\text{IND} - \text{D}\mathcal{O}_{\text{Gen}}^{\mathcal{A}}$, where $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ is a non-uniform QPT adversary (as explained in Section 3.1, \mathcal{A}_1 gives implicitly its internal state to \mathcal{A}_2), then any non-uniform QPT adversary \mathcal{A} has only a negligible advantage of winning this game:

$\text{IND} - \text{D}\mathcal{O}_{\text{Gen}}^{\mathcal{A}}(\lambda)$	
1 :	$(\mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 :	$c \xleftarrow{\$} \{0, 1\}$
3 :	$(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0^{(c)})$
4 :	$\tilde{c} \leftarrow \mathcal{A}_2(k)$
5 :	return $\tilde{c} = c$

$$\Pr [\text{IND} - \text{D}\mathcal{O}_{\text{Gen}}^{\mathcal{A}}(\lambda)] \leq \frac{1}{2} + \text{negl}(\lambda). \quad (14)$$

Note that we provide in Section 7 an explicit implementation of a δ -GHZ^H capable function.

The above properties are enough to prove the security of BLIND and BLIND^{sup} against an arbitrary corruption of parties, and can be used to prove the security of BLIND_{can} and BLIND_{can}^{sup} when the adversary corrupts only the server and the non-supported applicants. However, if we want to prove the security of these last two protocols in a stronger attack model, namely, when the adversary can also corrupt some supported applicants, we also require our function to have a stronger property. Intuitively, the PartInfo function will list the messages to send to all applicants: if it contains a \mathbf{X} for applicant i , it means that applicant i is not part of the support of the GHZ, if it is a 0 or a 1, it means that the applicant gets a GHZ canonical state—up to a local X correction if it is a 1—and, if it is a \perp , it means that the protocol aborts “locally”.

Remark 4.2. This local abort is interesting since it triggers when y has only one preimage, and this means that the server is malicious with overwhelming probability.¹⁴ Note that one may want to send this abort bit to all applicants, however it is not yet known if leaking this bit to other corrupted applicants could reduce the security of the protocol (for example, it is not clear if a malicious server could force the protocol to abort when one specific honest applicant is not part of the GHZ^C). To avoid that issue, we introduce this notion of local abort, that tells locally to participants if the server was behaving honestly. Note that it is important to make sure that this abort bit doesn’t leak to the adversary later, otherwise the security is not guaranteed anymore.

Definition 4.3. A δ -GHZ^H capable family of function $\{f_k\}$ is said to be δ -GHZ^{can} capable if this additional property is respected:

- **indistinguishability with partial knowledge:** We want to show that, by leaking some information about the “key” of the GHZ state owned by malicious applicants, we don’t reveal additional information about the support status of applicants. More precisely, there exists a PPT algorithm PartInfo: $\mathcal{T}_\lambda \times \mathcal{Y}_\lambda \rightarrow \{0, 1, \mathbf{X}, \perp\}^n$ with the following properties:
 - **correctness:** $\forall k \in \mathcal{K}_\lambda, y \in \mathcal{Y}_\lambda$, and $v \leftarrow \text{PartInfo}(t_k, y)$:

¹⁴This is the case when $\delta = \text{negl}(\lambda)$, which is possible to obtain using LWE with superpolynomial noise ratio.

- * y has exactly two preimages iff there is no partial abort (see discussion above): $|f_k^{-1}(y)| = 2$ iff $\perp \notin v$
- * for all i , if $v[i] \in \{0, 1\}$ then $\mathbf{d}_0[i] = 1$, and if $v[i] = \mathbf{X}$, then $\mathbf{d}_0[i] = 0$ (required to make sure \mathbf{X} is sent only to non-supported applicants and that 0/1 is sent only to supported applicants).
- * if $|f_k^{-1}(y)| = 2$, $\exists u \in \{h(x), h(x')\}$, such that $\forall i$, if $\mathbf{d}_0[i] = 1$ then $v[i] = u[i]$ (required to make sure that all corrections are correct).

IND – PARTIAL _{Gen,PartInfo} ^{\mathcal{A}} (λ)
1 : $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 : if $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$: return false fi
3 : $c \xleftarrow{\$} \{0, 1\}$
4 : $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0^{(c)})$
5 : $y \leftarrow \mathcal{A}_2(k)$
6 : $v \leftarrow \text{PartInfo}(t_k, y)$
7 : $\tilde{c} \leftarrow \mathcal{A}_3(\{(i, v[i])\}_{i \in \mathcal{M}})$
8 : return $\tilde{c} = c$

– **security**: The game on the left is impossible to win with non negligible advantage for any QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ (note that \mathcal{M} is intuitively the set of malicious corrupted applicants, and the condition line 2 is added because otherwise there is a trivial uninteresting way to distinguish).

For our protocol AUTH – BLIND_{can}^{dist}, we also need to make sure that this family of functions can be computed in a distributed manner among users:

Definition 4.4. A δ -GHZ^{can} capable family of function $\{f_k\}$ is said to be distributable if the above procedures can be computed after gathering partial results from the parties. More precisely:

- There exists Gen_{Loc} , a “local” generation procedure such that sampling $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0)$ can be done by first sampling for all i : $(k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$ and defining $k := (k^{(1)}, \dots, k^{(n)})$ and $t_k := (t_k^{(1)}, \dots, t_k^{(n)})$. We will denote as $\mathcal{K}_{\lambda, \text{Loc}}$ the set of such $k^{(i)}$, and we assume that $\mathcal{K} = \mathcal{K}_{\lambda, \text{Loc}}^n$.
- Similarly, there exists $\text{PartInfo}_{\text{Loc}}$, a “local” version of PartInfo such that sampling $v \leftarrow \text{PartInfo}(t_k, y)$ can be done by sampling for all i : $v[i] \leftarrow \text{PartInfo}_{\text{Loc}}(t_k^{(i)}, y)$.
- Finally, there exists a method $\text{PartAlpha}_{\text{Loc}}$ such that for any bit string b and for any y such that $f_k^{-1}(y) = \{x, x'\}$ with $x \neq x'$ we have $\langle b, x \oplus x' \rangle = \oplus_i \text{PartAlpha}_{\text{Loc}}(i, t_k^{(i)}, y, b)$.

Moreover, we cannot assume anymore that people running these functions will be honest. Therefore, if we want to make sure that a non-supported malicious applicant cannot alter the state obtained by supported applicants (for example, by providing a function which is not δ -2-to-1), we also require the existence of a circuit $\text{CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)})$ that returns **true** iff $k^{(i)} \in \mathcal{K}_{\lambda, \text{Loc}}$ ¹⁵ and if $k^{(i)}$ is the public key corresponding to the trapdoor $t_k^{(i)}$, embedding the bit $\mathbf{d}_0[i]$. This circuit can in particular be combined with a ZK protocol to prove in a ZK way that $k^{(i)} \in \mathcal{K}_{\lambda, \text{Loc}}$.

We also provide Section 7.2, a generic construction that turns a δ -GHZ^H capable family of functions into a δ' -GHZ^{can} capable distributable family of functions, with $\delta' = 1 - (1 - \delta)^n \leq \delta n$. In particular, if δ is a negligible function of λ as in Section 7.1 and $n = O(\lambda)$, δ' is negligible.

¹⁵This is particularly important when $\mathcal{K}_{\lambda, \text{Loc}} \subseteq \{0, 1\}^*$ and when there exists no efficient algorithm to decide if a bit string $k^{(i)}$ is indeed an element of $\mathcal{K}_{\lambda, \text{Loc}}$. For example, with our construction, it is easy to produce a key k' that is indistinguishable from a key $k \in \mathcal{K}$, and such that the function $f_{k'}$ is injective instead of δ -2-to-1.

5 NIZKoQS

In this section, we first formally define our new concept of NIZKoQS, and define a protocol achieving NIZKoQS. The more involved protocol AUTH – BLIND_{can}^{dist} defined in the next section will also exploits NIZKoQS (but this protocol will have more than one message as it is also consuming the state produced by the NIZKoQS), while the other simpler protocols will only rely on the correctness (or completeness) of the following NIZKoQS protocol. Before giving the formal definition, let us motivate and describe informally the definition.

Quantum language. In classical NIZK, a language \mathcal{L} is a set of strings, so similarly we will define a quantum language \mathcal{L}_Q as a set of quantum states. For instance, we could consider:

- the quantum language made of BB84 states $\mathcal{L}_Q^{\text{BB84}} = \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$,
- the quantum language $\mathcal{L}_Q^{\text{ex}} = \{|\mathbf{d}\rangle \pm |\mathbf{d} \oplus \mathbf{d}_0\rangle \mid (\mathbf{d}, \mathbf{d}_0) \in (\{0, 1\}^n)^2, \mathbf{d}_0[1] = 1, w_H(\mathbf{d}_0) = 2\}$ (w_H denotes the Hamming weight), corresponding to all hidden GHZ states whose first qubit is supported and where the support has size 2 (i.e., only two qubits are entangled forming a Bell state, where one of these qubits is at the first position),
- but we can also consider quantum languages referring to classical secrets, for instance if p_k is a public key (say of a bitcoin wallet), and if $\text{Ver}_{p_k}(s_k) = 1$ iff s_k is the private key of p_k , we can define $\mathcal{L}_Q^{p_k} = \{|\mathbf{d}\rangle \pm |\mathbf{d} \oplus \mathbf{d}_0\rangle \mid (\mathbf{d}, \mathbf{d}_0) \in (\{0, 1\}^n)^2, \mathbf{d}_0[1] = 0 \vee (\mathbf{d}_0[1] = 1 \wedge \exists s_k, \text{Ver}_{p_k}(s_k) = 1)\}$ that informally allows the prover to “send” a hidden GHZ state where the first qubit is supported *only* if the prover knows the private key of p_k .

Classically, both the prover and the verifier typically have a copy of the word x , and since information can be copied classically, the verification process cannot “destroy” x . Quantumly, this is not true anymore, therefore, instead of saying that all parties agree on ρ before the protocol, what matters is that *at the end* of the protocol, the verifier should end up with a $\rho \in \mathcal{L}_Q$.

Relation, witness, and quantum ZK. Classically, to check if a word x belongs to a language \mathcal{L} , we usually define a relation \mathcal{R} between a *witness* w and the word, saying that $x \in L \Leftrightarrow \exists w, w\mathcal{R}x$. The prover typically knows the witness w and the ZK property ensures that the verifier has no way to learn the witness w , formalizing the fact that the verifier learns nothing beyond the fact that the statement is true. Quantumly, we mimic this definition by defining a relation between classical *witnesses* (or *classes*¹⁶) ω and quantum states \mathcal{R} , saying that a quantum state ρ belongs to a quantum language \mathcal{L}_Q if and only if there exists ω such that $\omega\mathcal{R}\rho$. Similarly, we want to ensure that the verifier has no way to learn ω .

However, even if our definition does not say anything more about witnesses, we need to choose them appropriately to obtain a meaningful and secure protocol. Moreover, at that stage it may not even be clear what could be used as a witness. For instance, in the quantum language $\mathcal{L}_Q^{\text{BB84}}$ defined above, what would be the witness of $|0\rangle$? Because the ZK property ensures that no information leaks about the witness, while we typically want to ensure that no information is leaked about the received state, one could naively say that the witness is the classical description of the state. Unfortunately if each witness ω_ρ is in a 1-to-1 correspondence to its corresponding state $\rho \in \mathcal{L}_Q$, then it would be impossible to obtain the ZK property: for instance, given a random state in $\mathcal{L}_Q^{\text{BB84}} = \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$, it is possible to rule out one of the four states: for instance by measuring the state in the computational basis, if we measure b we know that the state can’t be $|1 - b\rangle$ and, therefore, we know that the witness can’t be $\omega_{|1-b\rangle}$, contradicting the ZK property.

¹⁶Unlike in classical ZK, the witness ω cannot be used to verify that a quantum state belongs to the quantum language, for this reason the term *class* may be more appropriate. This also justifies the usage of a different notation ω instead of w .

To overcome this issue, a single witness ω must characterize a *class* of states. For instance, for the language $\mathcal{L}_Q^{\text{BB84}}$ we will define two witnesses 0 and 1 characterizing the basis of the state and we, therefore, define the relation $0\mathcal{R}|0\rangle$, $0\mathcal{R}|1\rangle$, $1\mathcal{R}|+\rangle$ and $1\mathcal{R}|-\rangle$. These classes will be used to characterize two of the three wanted properties:

- **Completeness** (or correctness): An honest prover should be able to choose ω and generate on the side of the verifier a state in $\mathcal{L}_\omega := \{\rho \mid \omega\mathcal{R}\rho\}$.
- **ZK**: A malicious verifier should be unable to learn the witness ω with significant advantage over a random guess. Because of the completeness property, the verifier should therefore be unable to learn the class \mathcal{L}_ω chosen by the prover that was supposed to contain the target state.
- **Soundness**: Finally, we also expect that if the prover is malicious, then an honest verifier will obtain a state in \mathcal{L} whenever it accepts. (Note that this property does not depend on \mathcal{R})

From the above properties, it clearly appears that when the relation \mathcal{R} is thinner (i.e., when $|\mathcal{L}_\omega|$'s are smaller and the number of witnesses increases), we get a stronger result: indeed, an honest prover can choose more precisely the sent state and a malicious verifier is more confused as there are more classes to which a state could belong to. In particular, it is always possible to define a trivial NIZKoQS protocol for any language \mathcal{L}_Q if there is a single witness ω_0 such that $\forall \rho \in \mathcal{L}_Q, \omega_0\mathcal{R}\rho$: the prover would not do anything and the verifier would simply generating an arbitrary state in \mathcal{L}_Q . However, the guarantees are quite poor in that case as the verifier can fully describe the state... For this reason, we will focus on non-trivial relations, and we will always specify the relation associated with a quantum language.

Note that in some cases, it may be cumbersome to write separately the language and the relation, especially when the witness is an arbitrary label and when only the equivalence class formed by the relation matters. In that case, we may abuse notations and write directly $\mathcal{L} = \{\mathcal{L}_\omega\}_\omega$, like $\mathcal{L}_Q^{\text{BB84}} = \{|0\rangle, |1\rangle, |+\rangle, |-\rangle\}$. This more succinct notation makes it clearer that the prover can choose in advance the basis (computation or Hadamard) of the state obtained by the verifier, that an honest verifier would always obtain a BB84 state, and that a malicious verifier would be unable to learn the basis of the output state. For the other above examples of languages $\mathcal{L}_Q^{\text{ex}}$ and $\mathcal{L}_Q^{\text{pk}}$, the witness that we will consider is the support \mathbf{d}_0 of the GHZ state (therefore, no malicious verifier will be able to learn any information about \mathbf{d}_0 beyond the fact that it respects the constraints that are specified in the language). We can now provide the formal definition of (NI)ZKoQS.

Definition 5.1 (Zero-Knowledge Proof on Quantum State (ZKoQS)). Let (P, V) be a protocol between an honest QPT prover¹⁷ P and an honest QPT verifier V (that also outputs a final quantum state). Let $E = \cup_{n \in \mathbb{N}} \mathcal{L}_o(\mathcal{H}_n)$ be the set of finite dimensional quantum states (where \mathcal{H}_n is the Hilbert space of dimension n), $\mathcal{R} \subseteq \{0, 1\}^* \times E$ be a relation between bit strings (called *witnesses* or *classes*) and quantum states, and $\mathcal{L}_Q = \{\rho \in E \mid \exists \omega, \omega\mathcal{R}\rho\}$ be a *quantum language* defined by \mathcal{R} . Then (P_λ, V_λ) is said to be a ZKoQS for \mathcal{L}_Q if the following properties are respected:

- (1) **Completeness**: There exists a negligible function $\mu(\cdot)$ such that for any $\lambda \in \mathbb{N}$ and ω such that $\exists \rho' \in \mathcal{L}, \omega\mathcal{R}\rho'$,

$$\Pr [a = 1 \text{ and } \omega\mathcal{R}\rho \mid (a, \rho) \leftarrow \text{OUT}_V(P_\lambda(\omega), V_\lambda)] = 1 - \mu(\lambda). \quad (15)$$

- (2) **Soundness**: For any non-uniform QPT malicious prover $P^* = \{(P_\lambda^*, \sigma_\lambda)\}_{\lambda \in \mathbb{N}}$, there exists a negligible function $\mu(\cdot)$ such that for any security parameter $\lambda \in \mathbb{N}$,

$$\Pr [a = 1 \text{ and } \rho \notin \mathcal{L} \mid (a, \rho) \leftarrow \text{OUT}_V(P_\lambda^*(\sigma_\lambda), V_\lambda)] \leq \mu(\lambda). \quad (16)$$

When P^* is unbounded, it is called a *proof* otherwise an *argument*.

¹⁷In our case the prover is actually PPT.

- (3) **Quantum ZK**: There exists a QPT simulator Sim_λ such that for any QPT verifier $V^* = \{(V_\lambda^*, \sigma_\lambda)\}_{\lambda \in \mathbb{N}}$,

$$\{\text{OUT}_{V_\lambda^*}(P_\lambda(\omega), V_\lambda^*(\sigma_\lambda))\}_{\lambda, \omega} \approx_c \{\text{Sim}_\lambda(V_\lambda^*, \sigma_\lambda)\}_{\lambda, \omega}, \quad (17)$$

where $\lambda \in \mathbb{N}$, $\omega \in \{\omega \mid \exists \rho, \omega \mathcal{R} \rho\}$, and V^* is given to Sim_λ by sending the circuit description of V^* .

A Non-Interactive ZKoQS protocol—in which a single message is sent, from the prover to the verifier—will be denoted NIZKoQS.

Remark 5.2 (Relaxation of the Guarantees). This definition is fairly strict, in the sense that it is binary: either a state belongs to \mathcal{L} or not, without notion of distance to \mathcal{L} . As a consequence, if a protocol is a ZKoQS according to this definition, it should be hard for a malicious prover to be accepted while sending a state not in \mathcal{L} , but it might be easy to send a state really close to a state in \mathcal{L} . While in practice, our protocol is also secure against this stronger definition, we do not generalize this definition in this article for simplicity, but readers interested by a more generic definition can look into our follow up (partially) related article [Colisson et al. 2023] (see the historical note in the introduction for comparisons with this work). Note also that we leave open the question of the potential implications of ZKoQS, but our follow-up work [Colisson et al. 2023] shows interesting applications regarding Oblivious Transfer.

Now, we define a protocol where we can prove any property on the set of entangled qubits of a hidden GHZ state in a NIZKoQS fashion. Note that we formalise the notion of “any property” by relying on a function $\text{Auth}: \{0, 1\}^n \times \mathcal{W}_c \rightarrow \{0, 1\}$ specifying if a given support $\mathbf{d}_0 \in \{0, 1\}^n$ (which corresponds to the witness/class of the NIZKoQS) is allowed or not, where the decision can optionally depend on another classical secret witness $w \in \mathcal{W}_c$. This additional optional set of witnesses \mathcal{W}_c is needed to allow NIZKoQS relying on classical secrets like in the $\mathcal{L}_Q^{P_k}$ example provided above, where the \mathcal{W}_c would correspond to the set of private keys.

Remark 5.3. Note that we also require here the existence of $\text{CheckTrapdoor}(\mathbf{d}_0, t_k, k)$ that will check if $k \in \mathcal{K}$ and if \mathbf{d}_0 corresponds to the constant involved in the XOR property of k . The reason is that we cannot anymore be sure that Alice (the prover) is honest: therefore we need to check that the k sent by Alice was honestly prepared. This is particularly important when $\mathcal{K}_\lambda \subseteq \{0, 1\}^*$ and when there exists no efficient algorithm to decide if a bit string k is indeed an element of \mathcal{K}_λ . For example, with our construction, it is easy to produce a key k' that is indistinguishable from a key $k \in \mathcal{K}$, and such that the function $f_{k'}$ is injective instead of δ -2-to-1. Our construction based on LWE does guarantee that there exists a function CheckTrapdoor , which internally checks if the singular values of the trapdoor \mathbf{R} are small enough and if the norm of $(\mathbf{s}_0, \mathbf{e}_0)$ is small enough (see Lemmas 7.2 and 7.5 for more details).

THEOREM 5.4 (NIZKoQS). *Let $n \in \mathbb{N}$ be the size of the quantum state outputted by Bob in BLIND – ZK and $\delta = \text{negl}(\lambda)$. The protocol BLIND – ZK (where Alice plays the role of the prover P and Bob the verifier V) is a NIZKoQS for the quantum language, parameterized by the witnesses/classes $\mathbf{d}_0 \in \{0, 1\}^n$, defined by all hidden GHZ states ρ on n qubits whose support \mathbf{d}_0 is such that there exists w_c such that $\text{Auth}(\mathbf{d}_0, w_c)$ (defining the relation $\mathbf{d}_0 \mathcal{R} \rho$). More precisely, when the verifier Bob accepts, with overwhelming probability the state outputted by Bob is the hidden GHZ state $|\mathbf{d}\rangle + (-1)^\alpha |\mathbf{d}'\rangle$, with*

$$\mathbf{d} = h(x) \quad \mathbf{d}' = h(x') \quad \alpha = \bigoplus_i b_i(x_i \oplus x'_i) = \langle b, x \oplus x' \rangle. \quad (18)$$

In particular, since by definition of f_k we have $\mathbf{d} \oplus \mathbf{d}' = h(x) \oplus h(x') = \mathbf{d}_0$, the support of the hidden GHZ is \mathbf{d}_0 as explained in Section 3.2, where $\exists w_c, \text{Auth}(\mathbf{d}_0, w_c)$.

PROTOCOL 1: BLIND – ZK

Assumptions: There exists a $\text{negl}(\lambda)$ -GHZ^H capable family of functions (Definition 4.1), together with an efficient function $\text{CheckTrapdoor}_\lambda(\mathbf{d}_0, t_k, k)$ outputting **true** if $k \in \mathcal{K}$ and if $\mathbf{d}_0 = \mathbf{d}_0(t_k)$.

Parties: A classical sender/prover (Alice) and a quantum receiver/verifier (Bob).

Common inputs: The size n of the hidden GHZ and an efficiently computable function $\text{Auth}: \{0, 1\}^n \times \mathcal{W}_c \rightarrow \{0, 1\}$ where $\{0, 1\}^n$ corresponds to the witness defined in Definition 5.1 and \mathcal{W}_c is a set of other optional witnesses.

Alice's input: The support $\mathbf{d}_0 \in \{0, 1\}^n$ of the hidden GHZ state and a witness $w_c \in \mathcal{W}_c$ such that $\text{Auth}(\mathbf{d}_0, w_c) = 1$.

Bob's output: Bob can reject or accept and output a quantum state if he thinks that there exist \mathbf{d}_0 and w such that $\text{Auth}(\mathbf{d}_0, w_c) = 1$ and such that the quantum state is a hidden GHZ state of support \mathbf{d}_0 .

Protocol:

- Alice: Generate $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0)$ and a NIZK proof π proving that $\text{CheckTrapdoor}(\mathbf{d}_0, t_k, k) \wedge \text{Auth}(\mathbf{d}_0, w_c) = 1$ (the witness being (\mathbf{d}_0, t_k, w_c) and the word k) and send (k, π) to Bob.
 - Bob: Check that π is correct (if not, reject), then perform the following operations (circuit in Figure 1):
 - create the state $\sum_{x \in \mathcal{X}_\lambda} |x\rangle |h(x)\rangle |f_k(x)\rangle$ by applying in superposition the unitary mapping $|x\rangle |0\rangle |0\rangle \mapsto |x\rangle |h(x)\rangle |f_k(x)\rangle$ on the uniform superposition of all inputs,
 - measure the third register in the computational basis, which gives a y ,
 - measure the first register in the Hadamard basis, which gives a bit string b
- Then, output the resulting quantum state on the second register.

PROOF. The protocol is non interactive since a single message (k, π) is sent (the rest of the proof will also work for interactive protocols). First, for completeness we can see that when both parties are honest, the verifier accepts ($a = 1$) with probability 1 thanks to the perfect completeness of the classical NIZK protocol. Then, one can see that after the measurement of the third register, Bob gets on the first two registers the state $\sum_{x \in f^{-1}(y)} |x\rangle |h(x)\rangle$. But because the first register was created as a uniform superposition, y has exactly two preimages x, x' with probability $1 - \delta$. So with probability $1 - \delta$, the state obtained by an honest Bob is $(|x\rangle |h(x)\rangle + |x'\rangle |h(x')\rangle)$. Then, after applying the Hadamard gates (preparing the measurement in the Hadamard basis), we obtain the state $\sum_b (-1)^{\langle b, x \rangle} |b\rangle |h(x)\rangle + (-1)^{\langle b, x' \rangle} |b\rangle |h(x')\rangle$. So after measuring a b , the second register contains the state

$$(-1)^{\langle b, x \rangle} |h(x)\rangle + (-1)^{\langle b, x' \rangle} |h(x')\rangle = (-1)^{\langle b, x \rangle} (|h(x)\rangle + (-1)^{\langle b, x \oplus x' \rangle} |h(x')\rangle) \quad (19)$$

which is equal to the state $\text{GHZ}_{\alpha, \mathbf{d}, \mathbf{d}'}^H = |\mathbf{d}\rangle + (-1)^\alpha |\mathbf{d}'\rangle$, once we get rid of the global phases and define \mathbf{d}, \mathbf{d}' and α as above. Since we have $\mathbf{d} \oplus \mathbf{d}' = h(x) \oplus h(x') = \mathbf{d}_0$, the support of this hidden GHZ state is \mathbf{d}_0 as explained in Section 3.2, which concludes the proof of completeness.

The soundness property relies on the soundness property of the classical NIZK protocol, and again on the correctness of the circuit performed by the server: the probability of accepting a k which is not in \mathcal{K}_λ or such that there exists no w_c such that $\text{Auth}(\mathbf{d}_0, w_c)$ is negligible; the correctness of the protocol described above is enough to conclude that the hidden GHZ has the expected properties.




Fig. 1. Circuit performed by the server Bob.

For the ZK property, we define the following simulator, where Sim_{ZK} is the simulator of the classical (NI)ZK protocol, and $k \rightsquigarrow V_\lambda^*$ is the machine obtained by running V_λ^* , and sending k as first message:

$\text{Sim}_\lambda(V_\lambda^*, \sigma_\lambda)$	
1 :	$\mathbf{d}'_0 \leftarrow \mathbb{S} \{0, 1\}^n$
2 :	$(k', t_{k'}) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}'_0)$
3 :	return $\text{Sim}_{\text{ZK}}(k', k' \rightsquigarrow V_\lambda^*, \sigma_\lambda)$

To prove that the output of $\text{Sim}_\lambda(V_\lambda^*, \sigma_\lambda)$ is indistinguishable from the real world, we define an hybrid distribution:

$\text{Game1}(\mathbf{d}_0, V_\lambda^*, \sigma_\lambda)$	
1 :	$(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0)$
2 :	return $\text{Sim}_{\text{ZK}}(k, k \rightsquigarrow V_\lambda^*, \sigma_\lambda)$

First, one can see that $\{\text{Game1}(\mathbf{d}_0, V_\lambda^*, \sigma_\lambda)\}_{\lambda, \mathbf{d}_0} \approx_c \{\text{Sim}_\lambda(V_\lambda^*, \sigma_\lambda)\}_{\lambda, \mathbf{d}_0}$. Indeed, if a non-uniform distinguisher D can distinguish between these two distributions, then we can use D to break the game $\text{IND} - \text{D0}_{\text{Gen}}^{\omega}(\lambda)$ by simply sending for any λ a random \mathbf{d}_0 and the \mathbf{d}_0 which maximizes the distinguishing probability (anyway, D is already non-uniform). Then, $\{\text{Game1}(\mathbf{d}_0, V_\lambda^*, \sigma_\lambda)\}_{\lambda, \mathbf{d}_0} \approx_c \{\text{OUT}_{V_\lambda^*}(\langle P_\lambda(\mathbf{d}_0), V_\lambda^*(\sigma_\lambda) \rangle)\}_{\lambda, \mathbf{d}_0}$ since Game1 is exactly the same as the RHS, except that we replaced the actual ZK protocol with its simulator, which is an indistinguishable process by definition of ZK. \square

Remark 5.5. In particular, note that the above protocol provides a NIZKoQS for the examples given at the beginning of the section. The quantum language $\mathcal{L}_Q^{\text{BB84}}$ can be certified by remarking that a BB84 state is a hidden GHZ state of size one, where $\mathbf{d}_0 \in \{0, 1\}$ represents the basis, and therefore, we just need to take $\text{Auth}(\mathbf{d}_0, _) = 1$. The language $\mathcal{L}_Q^{\text{ex}}$ can be certified using

$\text{Auth}(\mathbf{d}_0, _) = (\mathbf{d}_0[1] = 1 \wedge w_H(\mathbf{d}_0) = 2)$ and finally \mathcal{L}_Q^{pk} can be obtained using $\text{Auth}(\mathbf{d}_0, s_k) = (\text{Ver}_{p_k}(s_k) = 1 \implies \mathbf{d}_0[1] = 1)$.

6 The Different Protocols

In this section, we define the protocols BLIND (Protocol 2), BLIND^{sup} (Protocol 3), the impossible BLIND_{can} (Section 6.3), BLIND_{can}^{sup} (Protocol 4), and finally our main protocol AUTH – BLIND_{can}^{dist} (Protocol 5).

6.1 The Protocol BLIND

We now define the protocols BLIND (Protocol 2), which is the basic building block of all the other protocols. Note that this protocol can already easily be extended to exhibit NIZKoQS as shown in Section 5. This protocol allows a trusted classical party, Cupid, to create a hidden GHZ state on a server. This hidden GHZ state is then distributed to all applicants (one qubit per applicant), who just need to store the received qubit.

LEMMA 6.1 (CORRECTNESS OF BLIND AND BLIND^{sup}). *At the end of an honest run of protocol BLIND, when y has exactly two distinct preimages x, x' (which occurs with probability $1 - \delta$ according to Definition 4.1, which is overwhelming when we use the construction defined in Section 7.1), the state shared between all applicants is a hidden generalized GHZ^H _{$\alpha, \mathbf{d}, \mathbf{d}'$} state, with*

$$\mathbf{d} = h(x) \quad \mathbf{d}' = h(x') \quad \alpha = \bigoplus_i b_i(x_i \oplus x'_i) = \langle b, x \oplus x' \rangle. \quad (20)$$

In particular, since by definition of f_k we have $\mathbf{d}_0 = h(x) \oplus h(x') = \mathbf{d} \oplus \mathbf{d}'$, the support of the hidden GHZ is \mathbf{d}_0 .

The proof of this lemma is a direct consequence of the completeness of the protocol BLIND – ZK proven in the first part of the proof of Theorem 5.4.

LEMMA 6.2 (SECURITY OF BLIND). *At the end of a fully malicious interaction during the protocol BLIND, where all applicants and the server can be fully malicious and can all collude together, the set of supported applicants is completely hidden. More precisely, if we define a game following the spirit of IND-CPA security, no QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2)$ can win the game IND – BLIND with probability better than $\frac{1}{2} + \text{negl}(\lambda)$.*

IND – BLIND _{Gen} ^{sc} (λ)	
1 :	$(\mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 :	$c \xleftarrow{\$} \{0, 1\}$
3 :	$(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0^{(c)})$
4 :	$(y, b, \tilde{c}) \leftarrow \mathcal{A}_2(k)$
5 :	/ No more interaction
6 :	return $\tilde{c} = c$

PROOF. This is a trivial reduction to the indistinguishability property of the family $\{f_k\}$: since y and b are not used, we can remove them without changing the probability of winning, and we get exactly the game IND – D0 introduced Definition 4.1. This game is impossible to win for probability better than $\frac{1}{2} + \text{negl}(\lambda)$ by assumption on the family $\{f_k\}$, which ends the proof. \square

PROTOCOL 2: BLIND

Inputs: Cupid gets as input $\mathbf{d}_0 \in \{0, 1\}^n$, a bit string describing the final supported applicants: applicant a_i will be supported iff $\mathbf{d}_0[i] = 1$. $\lambda \in \mathbb{N}$ is a public, fixed, security parameter.

Assumptions: There exists a δ -GHZ^H capable family of functions with $\delta = \text{negl}(\lambda)$.

Protocol:

- Cupid: Generate $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0)$, and send k to the server.
- Server: Perform the quantum operations described in BLIND – ZK (circuit in Figure 1). Then, send (y, b) to Cupid, and for all i , send the i th qubit of the (unmeasured) second register to applicant a_i .
- All applicants: just receive and store the qubit sent by the server.

PROTOCOL 3: BLIND^{sup}

Inputs: Same as BLIND: Cupid gets \mathbf{d}_0 and λ .

Assumptions: Same as BLIND (δ -GHZ^H capable family with $\delta = \text{negl}(\lambda)$).

Protocol:

- Run the protocol BLIND, so that Cupid gets (b, y) and each applicant a_i the i -th qubit
- Cupid: For all i , send $\mathbf{d}_0[i]$ to applicant a_i , so that each applicant knows whether they are supported or not.

6.2 The Protocol BLIND^{sup}

We now describe the protocol BLIND^{sup} (Protocol 3). In this protocol, all the applicants will obtain a qubit part of a hidden GHZ state, and they will learn their own support status. However, they will not know the “key” of the hidden GHZ state.

Now, in order to prove the security of the BLIND^{sup} protocol, we first need to define what we mean by security. Since in this protocol Cupid reveal to all applicants their respective support status, we can’t use the previous definition of security.

LEMMA 6.3 (SECURITY OF BLIND^{sup}). *If we allow in the protocol BLIND^{sup} the fully malicious server Bob to corrupt an arbitrary subset of applicants, then the support status of the remaining honest applicants is completely hidden. More precisely, no QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ can win the game IND – BLIND^{sup} with probability better than $\frac{1}{2} + \text{negl}(\lambda)$. In the following, \mathcal{M} is the set of malicious applicants corrupted by Bob, and the condition $\forall i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] = \mathbf{d}_0^{(1)}[i]$ is required to avoid a trivial uninteresting distinguishing strategy.*

IND – BLIND ^{sup, \mathcal{A}} _{$\{f_k\}$} (λ)
1 : $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 : if $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$:
3 : return false fi
4 : $c \xleftarrow{\$} \{0, 1\}$
5 : $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0^{(c)})$
6 : $(y, b) \leftarrow \mathcal{A}_2(k)$
7 : / The adversary has only access
8 : / to the messages sent by Cupid
9 : / to corrupted applicants:
10 : $\tilde{c} \leftarrow \mathcal{A}_3(\{(i, \mathbf{d}_0^{(c)}[i])\}_{i \in \mathcal{M}})$
11 : return $\tilde{c} = c$

PROOF. To prove the security of this scheme, we will assume by contradiction that there exists an adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ that can win the game IND – BLIND^{sup} with probability $p_{\mathcal{A}} := \frac{1}{2} + \frac{1}{\text{poly}(\lambda)}$, and we will construct an adversary $\mathcal{A}' = (\mathcal{A}'_1, \mathcal{A}'_2)$ that can win the game IND – BLIND

with a non negligible advantage (which is impossible by assumption). So we define $\mathcal{A}'_1(\lambda)$ as follows: \mathcal{A}'_1 runs in a blackbox way $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1$, returns $(\mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)})$ and keeps $(\mathcal{M}, \mathbf{d}_0^{(0)})$ in its internal state. We then define

$$\mathcal{A}'_2(k, \text{state}_1 := (\mathcal{M}, \mathbf{d}_0^{(0)})) := (y, b) \leftarrow \mathcal{A}_2(k); \tilde{c} \leftarrow \mathcal{A}_3(\{(i, \mathbf{d}_0^{(0)}[i])\}_{i \in \mathcal{M}}); \text{return } \tilde{c}. \quad (21)$$

It is then easy to see that \mathcal{A}' wins the game $\text{IND} - \text{D0}$ with probability greater than $p_{\mathcal{A}}$. Indeed, when \mathcal{A}_1 outputs a $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)})$ that does not respect the condition $\forall i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] = \mathbf{d}_0^{(1)}[i]$, then \mathcal{A} always lose (while \mathcal{A}' may win the game). Moreover, when the condition is respected, since $\{(i, \mathbf{d}_0^{(0)}[i])\}_{i \in \mathcal{M}} = \{(i, \mathbf{d}_0^{(1)}[i])\}_{i \in \mathcal{M}} = \{(i, \mathbf{d}_0^{(c)}[i])\}_{i \in \mathcal{M}}$, we can replace the input of \mathcal{A}'_3 with $\{(i, \mathbf{d}_0^{(c)}[i])\}_{i \in \mathcal{M}}$: the game is now exactly equivalent to $\text{IND} - \text{BLIND}^{\text{sup}}$, so in that case \mathcal{A}' win with the exact same probability as \mathcal{A} . Therefore, \mathcal{A}' wins the game $\text{IND} - \text{D0}$ with probability greater than $p_{\mathcal{A}} = \frac{1}{2} + \frac{1}{\text{poly}(n)}$: contradiction. \square

6.3 The Impossible Protocol $\text{BLIND}_{\text{can}}$

One may be interested by a protocol $\text{BLIND}_{\text{can}}$, that would make sure that all supported applicants share a canonical GHZ, but that at the same time none of them know if they are part of the GHZ or not. We state here that such a security guarantee is impossible, and why it is therefore meaningful to use $\text{BLIND}_{\text{can}}^{\text{sup}}$ protocols instead.

LEMMA 6.4 (IMPOSSIBILITY OF A SECURE $\text{BLIND}_{\text{can}}$ PROTOCOL). *There exists no protocol $\text{BLIND}_{\text{can}}$ such that, at the end of an honest interaction, all supported applicants share a canonical GHZ, and such that none of them know their own support status. More precisely, there exists always an adversary \mathcal{A} that can win the game [ImpossibleGame1](#).*

ImpossibleGame1

```

1 :  $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$ 
   / Avoid trivial attack: check if at least one honest applicant is in the GHZ
2 : if  $(\forall i \notin \mathcal{M}, \mathbf{d}_0^{(0)}[i] = 0)$  or  $(\forall i \notin \mathcal{M}, \mathbf{d}_0^{(1)}[i] = 0)$  then return false fi
3 : / Run  $\text{BLIND}_{\text{can}}$  with adversary.
4 :  $\tilde{c} \leftarrow \mathcal{A}_3()$ 
5 : return  $\tilde{c} = c$ 
```

PROOF. So for simplicity, we do here a sketch of the proof, in order to give the general ideas. So when it comes to proving the security of the protocol, we realize that at least one of the supported applicants needs to be honest, otherwise it is trivial to distinguish any correct protocol: The attacker can always send $\mathbf{d}_0^{(0)} = (1 \ 1 \ 0 \dots 0)$ and $\mathbf{d}_0^{(1)} = (0 \ 0 \ 0 \dots 0)$, and at the end of any (correct) protocol run honestly, the attacker will get either a Bell pair on the first two qubits or two qubits not entangled. It is therefore easy to distinguish, so the condition in Game 1 is indeed required. But it is not enough: even if we assume that one applicant is honest (let's say the first one), it is still impossible to prove the security of the protocol.

Indeed, let's consider an adversary that sends $\mathbf{d}_0^{(0)} = (1 \ \dots 1)$ and $\mathbf{d}_0^{(1)} = (1 \ \dots 1 \ 0 \dots 0)$ (first half of the qubits is 1 second half is 0). Then, a first remark is that at the end of an honest protocol, all the qubits that are not entangled must be all equal, i.e., if $c = 1$, the state obtained is $(|0 \dots 0\rangle + |1 \dots 1\rangle) \otimes |0 \dots 0\rangle$ or $(|0 \dots 0\rangle + |1 \dots 1\rangle) \otimes |1 \dots 1\rangle$. Indeed, if some qubits in the second half are different, then a measurement in the computational basis will reveal some different outcomes with high probability (while when $c = 0$ all measurements are equal since the state is a

PROTOCOL 4: $\text{BLIND}_{\text{can}}^{\text{sup}}$

Inputs: Same as BLIND : Cupid gets \mathbf{d}_0 and λ .

Assumptions: There exists a δ -GHZ^{can} capable family of functions with $\delta = \text{negl}(\lambda)$.

Protocol:

- Run the protocol BLIND , so that Cupid gets (b, y) and each applicant a_i the i -th qubit
- Cupid: Compute $v \leftarrow \text{PartInfo}(t_k, y)$, and if $f_k^{-1}(y) = \{x, x'\}$ with $x \neq x'$, compute $\alpha := \langle b, x \oplus x' \rangle$ (otherwise, sample α randomly). Computes the supported set $\mathcal{S} = \{i \mid \mathbf{d}_0[i] = 1\}$. Sample uniformly at random $\hat{\alpha} \leftarrow \{0, 1\}^n$ such that $\alpha = \oplus_{i \in \mathcal{S}} \hat{\alpha}_i$. For all i , send $(\hat{\alpha}_i, v[i])$ to applicant a_i .
- All applicants: When receiving the message $(\hat{\alpha}_i, v[i])$:
 - If $v[i] = \mathbf{X}$, then it means that the applicant is not part of the support of the final GHZ. The end.
 - If $v[i] = \perp$, it is a local abort. It's likely that the server was malicious. Do not reveal this information to the server. The end.
 - If $v[i] \in \{0, 1\}$, it means that this applicant is part of the final GHZ state. Apply $Z^{\hat{\alpha}_i} X^{v[i]}$ on the qubit sent by the server.

canonical GHZ state by the correctness property). But even in that case, it is still easy to distinguish: when we do the measurement, in the first case, we either get $1 \dots 1$ or $0 \dots 0$. In the second case, however, the first part may be different compared to the second part, i.e., we can measure either a $0 \dots 0$ or a $1 \dots 10 \dots 0$ with probability $\frac{1}{2}$. This last measurement is enough to distinguish, we can just ask to \mathcal{A} to measure the state in the computational basis: if all measurements are equal, \mathcal{A} picks \tilde{c} uniformly at random, otherwise \mathcal{A} outputs $\tilde{c} = 1$. \mathcal{A} will succeed with non-negligible advantage. \square

Therefore, it is not possible to hide to an adversary its support status, so the best we can get is to prove that no adversary can learn the support status of the honest applicants, which is the goal of the protocol $\text{BLIND}_{\text{can}}^{\text{sup}}$.

6.4 The Protocol $\text{BLIND}_{\text{can}}^{\text{sup}}$

We present now the protocol $\text{BLIND}_{\text{can}}^{\text{sup}}$ (Protocol 4): at the end of the protocol, the supported applicants share a canonical GHZ state, and each applicant knows their own support status.

LEMMA 6.5 (CORRECTNESS OF $\text{BLIND}_{\text{can}}^{\text{sup}}$). *If all parties are honestly running the $\text{BLIND}_{\text{can}}^{\text{sup}}$ protocol, then at the end of the protocol, with probability $1 - \delta$ (so with overwhelming probability when δ is negligible), all supported applicants share a canonical GHZ, and all applicants know whether or not they are supported.*

PROOF. With probability $1 - \delta$, the y obtained by Cupid has exactly two preimages. In that case, due to the correctness property of PartInfo given Definition 4.3 (part 1) we get $\forall i, v[i] \neq \perp$, so $v[i] \in \{0, 1, \mathbf{X}\}$. Then, using Lemma 6.1, we know that the state shared by all participants after the BLIND part is $|h(x)\rangle + (-1)^\alpha |h(x')\rangle$, with $h(x) \oplus h(x') = \mathbf{d}_0$. We can combine this using part 2 of the correctness property given in Definition 4.3, (that states that $v[i] \in \{0, 1\}$ iff $\mathbf{d}_0[i] = 1$): because the set of supported participants is $\mathcal{S} := \{i \mid \mathbf{d}_0[i] = 1\}$, we have for all $i \notin \mathcal{S}$: $h(x)[i] = h(x')[i]$. Thus the register of each applicant $i \notin \mathcal{S}$ is in a tensor product with all the other qubits, so we can factor them out, and consider only the state shared by applicants $i \in \mathcal{S}$ (in that

case $h(x')[i] = 1 \oplus h(x)[i]$:

$$\bigotimes_{i \in \mathcal{S}} |h(x)[i]\rangle + (-1)^\alpha \bigotimes_{i \in \mathcal{S}} |h(x')[i]\rangle = \bigotimes_{i \in \mathcal{S}} |h(x)[i]\rangle + (-1)^\alpha \bigotimes_{i \in \mathcal{S}} |1 \oplus h(x)[i]\rangle. \quad (22)$$

After the corrections, the state becomes:

$$\bigotimes_{i \in \mathcal{S}} X^{v[i]} |h(x)[i]\rangle + (-1)^{\alpha \oplus \bigoplus_{i \in \mathcal{S}} \hat{\alpha}_i} \bigotimes_{i \in \mathcal{S}} X^{v[i]} |1 \oplus h(x)[i]\rangle. \quad (23)$$

And due to the fact that $\alpha = \bigoplus_{i \in \mathcal{S}} \hat{\alpha}_i$, we can get rid of the phase. Moreover, we can now use the part 3 of Definition 4.3 which states that there exists $u \in \{h(x), h(x')\}$ such that if $\mathbf{d}_0[i] = 1$, then $v[i] = u[i]$. So if $\forall i \in \mathcal{S}$, we have $u[i] = v[i] = h(x)[i]$, then after the correction we get the state $|0 \dots 0\rangle + |1 \dots 1\rangle$, which is a canonical GHZ, and if $u[i] = v[i] = h(x')[i] = 1 - h(x)[i]$, then we get $|1 \dots 1\rangle + |0 \dots 0\rangle$, which is the same canonical GHZ state. \square

Similarly, we define now the security of $\text{BLIND}_{\text{can}}^{\text{sup}}$:

LEMMA 6.6 (SECURITY OF $\text{BLIND}_{\text{can}}^{\text{sup}}$). *If we allow in the protocol $\text{BLIND}_{\text{can}}^{\text{sup}}$ the fully malicious server to corrupt a subset of applicants in such a way that either at least one supported applicant is not corrupted or no supported applicant is corrupted,¹⁸ then the support status of honest applicants is completely hidden. More precisely, no QPT adversary $\mathcal{A} = (\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3)$ can win the game $\text{IND} - \text{BLIND}_{\text{can}}^{\text{sup}}$ (Figure 2) with probability better than $\frac{1}{2} + \text{negl}(\lambda)$. In the following, \mathcal{M} is the set of malicious applicants corrupted by Bob, and the condition $\forall i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] = \mathbf{d}_0^{(1)}[i]$ is required to avoid a trivial uninteresting distinguishing strategy.*

PROOF. The first step in the proof is to note that, due to the condition line 3, we have either:

- $\forall i \in \mathcal{M}, \mathbf{d}_0^{(0)} = 0$, i.e., $\mathcal{M} \cap \mathcal{S} = \emptyset$. But since we send $\hat{\alpha}_i$ to the adversary only if $i \in \mathcal{M}$, and since the line 10 does not put any restriction on the sampling $\hat{\alpha}_i$ when $i \notin \mathcal{S}$, the line 9 and 10 can be replaced with a single line $\hat{\alpha} \leftarrow \{0, 1\}^n$.
- or there exists $j \notin \mathcal{M}$ such that $\mathbf{d}_0^{(c)}[j] = 1$, i.e., such that $j \in \mathcal{S}$. Therefore, $\hat{\alpha}$ can be sampled by choosing for all $i \neq j$, $\hat{\alpha}_i$ randomly, and finally by setting $\hat{\alpha}_j = \alpha \oplus \bigoplus_{i \in \mathcal{S} \setminus \{j\}} \hat{\alpha}_i$ (this is statistically indistinguishable). But since $\hat{\alpha}_j$ is never sent to \mathcal{A} because $j \notin \mathcal{M}$, we can also remove lines 9 and 10 and replace them with $\hat{\alpha} \leftarrow \{0, 1\}^n$.

This gives us a new game GAME1:

GAME1 ^{\mathcal{A}}
/ ... First 7 lines like $\text{IND} - \text{BLIND}_{\text{can}}^{\text{sup}}$
9 : $\hat{\alpha} \leftarrow \{0, 1\}^n$
10 : $\tilde{c} \leftarrow \mathcal{A}_3(\{(i, \hat{\alpha}_i, v[i])\}_{i \in \mathcal{M}})$
11 : return $\tilde{c} = c$

Since the two games are exactly equivalent, we have $\Pr \left[\text{IND} - \text{BLIND}_{\text{can}}^{\text{sup}, \mathcal{A}} \right] = \Pr \left[\text{GAME1}^{\mathcal{A}} \right]$. Then, define a new game GAME2 by removing the condition lines 3 and 4:

¹⁸Otherwise, there is a trivial, fundamental, attack to any protocol which consists in setting $\mathbf{d}_0^{(0)} = (01)$, $\mathbf{d}_0^{(1)} = (11)$, $\mathcal{M} = \{2\}$ and then testing if the quantum state obtained by the party 2 is a $|+\rangle$ or not. However, this attack is not possible anymore if the adversary is not in possession of one part of the GHZ (for example, if we replace $\mathbf{d}_0^{(0)} = (00)$ and $\mathbf{d}_0^{(1)} = (10)$ in the above example), that is the reason why we can provide a stronger security guarantee when no supported applicant is supported.

$\text{IND} - \text{BLIND}_{\text{can}\{f_k\}}^{\text{sup}, \mathcal{A}}(\lambda)$ <hr/> <pre> 1: $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$ 2: if $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$ then return false fi 3: if $(\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)} = 1)$ and $((\forall i \notin \mathcal{M}, \mathbf{d}_0^{(0)}[i] = 0) \text{ or } (\forall i \notin \mathcal{M}, \mathbf{d}_0^{(1)}[i] = 0))$ 4: then return false fi 5: $c \xleftarrow{\\$} \{0, 1\}$ 6: $(k, t_k) \leftarrow \text{Gen}(1^\lambda, \mathbf{d}_0^{(c)})$ 7: $(y, b) \leftarrow \mathcal{A}_2(k)$ 8: $v \leftarrow \text{PartInfo}(t_k, y)$ 9: if $\perp \notin v$ then $\alpha := \langle b, x \oplus x' \rangle$ else $\alpha \xleftarrow{\\$} \{0, 1\}$ fi 10: $\hat{\alpha} \leftarrow \{\hat{\alpha} \mid \hat{\alpha} \in \{0, 1\}^n, \bigoplus_{i \in \mathcal{S}} \hat{\alpha}_i = \alpha \text{ or } \mathcal{S} = \emptyset\}$ 11: / The adversary has only access to the messages sent by Cupid to corrupted applicants: 12: $\tilde{c} \leftarrow \mathcal{A}_3(\{(i, \hat{\alpha}_i, v[i])\}_{i \in \mathcal{M}})$ 13: return $\tilde{c} = c$ </pre>
--

Fig. 2. Game $\text{IND} - \text{BLIND}_{\text{can}}^{\text{sup}}$ required in Lemma 6.6.

$\text{GAME2}^{\mathcal{A}}$ <hr/> <pre> / ... Remove line 3 and 4 of GAME1 3: if $(\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)} = 1)$ and $((\forall i \notin \mathcal{M}, \mathbf{d}_0^{(0)}[i] = 0) \text{ or } (\forall i \notin \mathcal{M}, \mathbf{d}_0^{(1)}[i] = 0))$ 4: then return false fi 5: / ... Rest is like GAME1 </pre>

We can remark that this condition cannot help the adversary to win since entering inside this condition always returns “false”, therefore, $\Pr[\text{GAME1}^{\mathcal{A}}] \leq \Pr[\text{GAME2}^{\mathcal{A}}]$. But now, we remark that GAME2 is very similar to the game $\text{IND} - \text{PARTIAL}$ defined Definition 4.3, except that we provide an additional random string $\hat{\alpha}$ to \mathcal{A} . But since this string is random, it is easy to see that we can turn any adversary \mathcal{A} winning GAME2 with probability p into an adversary \mathcal{A}' winning the game $\text{IND} - \text{PARTIAL}$ with the same probability p by defining $\mathcal{A}'(\{(i, v[i])\}_{i \in \mathcal{M}})$ as an adversary sampling a uniformly random bit string $\hat{\alpha}$ and calling $\mathcal{A}(\{(i, \hat{\alpha}_i, v[i])\}_{i \in \mathcal{M}})$ (and reciprocally, any adversary that can win $\text{IND} - \text{PARTIAL}$ without access to $\hat{\alpha}$ can win GAME2 with access to $\hat{\alpha}$ by simply forgetting this value). So we get

$$\max_{\text{QPT } \mathcal{A}} \Pr[\text{GAME2}^{\mathcal{A}}] = \max_{\text{QPT } \mathcal{A}} \Pr[\text{IND} - \text{PARTIAL}^{\mathcal{A}}]. \quad (24)$$

But by assumption, for any $\text{QPT } \mathcal{A}$, $\Pr[\text{IND} - \text{PARTIAL}^{\mathcal{A}}] \leq \frac{1}{2} + \text{negl}(\lambda)$. So by combining all the inequations we showed on games, we also get

$$\Pr[\text{IND} - \text{BLIND}_{\text{can}\{f_k\}}^{\text{sup}, \mathcal{A}}(\lambda)] \leq \frac{1}{2} + \text{negl}(\lambda) \quad (25)$$

which ends the proof. \square

6.5 The Protocol AUTH – BLIND_{can}^{dist}

We can now define our main protocol AUTH – BLIND_{can}^{dist}. Similarly to the BLIND_{can}^{sup} protocol, each supported applicant is supposed to end up with a canonical GHZ state, and the support status of each applicant should be unknown to the other applicants and to the server. However, in this protocol the trusted party Cupid is not needed anymore: each applicant is supposed to choose themselves their own support status, and they will be assured that no malicious party (including the server) can learn it.

Moreover, the server can have some guarantees on the support status of the applicants: for example, the server can ensure that *if* some applicants are supported, then they all know a classical secret (but the server has no way to know whether or not a given applicant is supported). This secret can be any witness of a NP relation: it could be a password, a private key linked with some known public key, a signature from a third party Certification Authority, the proof of any famous theorem... We formalize it by defining n deterministic functions $\text{Auth}_i: \{0, 1\} \times \{0, 1\}^* \rightarrow \{0, 1\}$ responsible of the “authorization” of the applicants: the server will allow applicant i to be part of the protocol iff they can prove in a NIZK way that they know w_c such that $\text{Auth}_i(\mathbf{d}_0[i], w_c) = 1$. For instance, we can use the Auth_i function to ensure that an applicant is part of the GHZ iff they know a password whose hash by h is x by defining $\text{Auth}_i(\mathbf{d}_0[i], \tilde{s}_i) := (\mathbf{d}_0[i] = 0 \vee h(\tilde{s}_i) = x)$. Again, we emphasize that Auth_i does *not* reveal the value of $\mathbf{d}_0[i]$: it just reveals that *if* the user is supported, then they know the password. This verification only requires a single message from the client(s) and is therefore achieving NIZKoQS, as formalized in Section 5.

Remark 6.7. Note that one of the motivation of our protocol compared to running directly QMPC is to have a more efficient protocol, e.g., in the number of rounds of communication. So one might be worried to see that we use inside our protocol a MPC protocol. Note however that this is a *classical* MPC protocol, which is significantly less costly than any QMPC protocol: QMPC not only typically uses MPC protocols internally as a black box, but also require lots of extra quantum rounds of communication (which are particularly hard to implement in practice since each quantum communication will add some extra noise, latency etc). For instance, in [Dulek et al. 2020], on each applied gate, a quantum state must travel through all parties, i.e., a single gate costs n quantum rounds (plus extra classical rounds due to the use of a classical MPC protocol with a complexity close to the one we are computing in our own protocol) which results in a way larger number of rounds.

We will now prove the correctness and security of the AUTH – BLIND_{can}^{dist} protocol. Note that an honest server can obtain guarantees on the distributed state even in the presence of malicious or noisy applicants. Assuming here an honest server is not absurd, notably when the server wants to use this GHZ, for example to share a quantum state or if a verification is done afterward. This centralization is also useful in the presence of many noisy clients (a single hardware needs to be noiseless, while in a decentralized MPC computation the protocol is likely to always abort if a single client is noisy). We provide here an informal version of the correctness property, more details can be found in Lemma 6.9.

LEMMA 6.8 (CORRECTNESS OF AUTH – BLIND_{can}^{dist} IN THE PRESENCE OF MALICIOUS APPLICANTS, INFORMAL). *In the presence of malicious applicants, an honest server is guaranteed that with overwhelming probability, the protocol will either abort, or the applicants will obtain a GHZ state, up to some unavoidable local deviation performed by supported malicious applicants on their own parts of the GHZ.*

LEMMA 6.9 (CORRECTNESS OF AUTH – BLIND_{can}^{dist} IN THE PRESENCE OF MALICIOUS APPLICANTS). *If the server is honest, and if we allow an attacker to corrupt an arbitrary subset \mathcal{M} of applicants, then with overwhelming probabilities, the protocol either aborts, or the k received by the server*

PROTOCOL 5: AUTH – BLIND_{can}^{dist}

Inputs: Each applicant i gets λ . They also get $\mathbf{d}_0[i] \in \{0, 1\}$ and $w_i \in \{0, 1\}^*$ such that $\text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$ ($\mathbf{d}_0[i] = 1$ iff applicant i wants to be supported). The authentication functions $\{\text{Auth}_i\}_{i \in [n]}$ are also public.

Assumptions: There exists a δ -GHZ^{can} capable distributable family of functions with $\delta = \text{negl}(\lambda)$. Cupid is not required anymore, and instead we require the existence of a classical (but quantum-secure) MPC protocol.

Protocol:

- Each applicant a_i : Run $(k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$, send $k^{(i)}$ to the server, and continue the protocol.
- Server: Run (as a verifier) a ZK protocol with each applicant (the prover) to check that $k^{(i)}$ is well prepared, and that the applicant can authenticate the quantum state. More precisely, each applicant i proves to the server that they know $(\mathbf{d}_0[i], t_k^{(i)}, w^{(i)})$ such that $\text{CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)}) \wedge \text{Auth}_i(\mathbf{d}_0[i], w^{(i)}) = 1$. If the protocol fails with at least one applicant, abort after sending \perp to all applicants (the server can also output if needed the identity of the applicant who were malicious in case other actions should be performed with respect to them, e.g., in further runs). Otherwise, the protocol continues.
- Server: Compute $k := (k^{(1)}, \dots, k^{(n)})$, run the quantum circuit already described in protocol BLIND, and for all i , send the i th qubit of the second register to applicant a_i together with (y, b) .
- For each applicant i : Compute $v[i] \leftarrow \text{PartInfo}_{\text{Loc}}(t_k^{(i)}, y)$, and compute via a MPC protocol the function CombineAlpha which returns a secret share of α between supported applicants. More precisely, it returns to each applicant i a random bit $\hat{\alpha}_i$ such that $\bigoplus_{i \in \mathcal{S}} \hat{\alpha}_i = \alpha = \bigoplus_i \text{PartAlpha}_{\text{Loc}}(i, t_k^{(i)}, y, b)$, where \mathcal{S} is the set of supported applicants (details in Figure 3). The reason we use a MPC protocol is that $\text{PartAlpha}_{\text{Loc}}(i, t_k^{(i)}, y, b)$ can leak^a information about the bit $\mathbf{d}_0[i]$, so this bit should not be revealed directly.
- For each applicant i : If the outcome $\hat{\alpha}_i$ of the MPC is \perp , abort. Otherwise, similarly to the last step of the BLIND_{can}^{sup} protocol: if $v[i] \in \{0, 1\}$, apply the correction $X^{v[i]} Z^{\hat{\alpha}_i}$ on the qubit, else discard the qubit.

^aThe attack would be as follows: the malicious server Bob can run the QFactory [Cojocaru et al. 2019] protocol with $k^{(i)}$ which gives him a BB84 state in the basis $\mathbf{d}_0[i]$: if $\mathbf{d}_0[i] = 0$ then it gets either $|0\rangle$ or $|1\rangle$, but if $\mathbf{d}_0[i] = 1$ then it gets the state $|+\rangle$ if $\alpha_p^{(i)} \alpha_i = 0$ and a $|-\rangle$ otherwise. So the trick is to measure the state in the Hadamard basis: if the measurement is different from $\alpha_p^{(i)}$ then we know that $\mathbf{d}_0[i] = 0$, and otherwise, the server will randomly guess the value of $\mathbf{d}_0[i]$. It is easy to see that if $\mathbf{d}_0[i]$ is chosen uniformly at random, then the server has a non-negligible advantage in guessing $\mathbf{d}_0[i]$.

belongs to \mathcal{K} , and for all applicant i , there exists w_i such that $\text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$. If the ZK protocol is also a Proof of Knowledge protocol, then the applicant “knows” w_i , in the sense that if the adversary can pass the test with non-negligible probability, there exists an extractor that can extract w_i given the applicant’s circuit with non-negligible probability (this is a direct application of Definition A.1).¹⁹

Moreover, if the protocol did not abort before, at the end of the protocol, with probability $1 - \delta - \text{negl}(\lambda)$ (i.e., overwhelming if δ is negligible), the protocol will either abort, or a state will be obtained by

¹⁹In particular, if it is impossible to forge with non-negligible probability a w_i such that $\text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$ (for instance because this w_i is a signature coming from an unforgeable signature scheme), then it means that with overwhelming probability the applicant is indeed in possession of w_i .

```

CombineAlpha((k, y, b), (tk(1), d0[1], k(1), y(1), b(1)), ..., (tk(n), d0[n], k(n), y(n), b(n)))
1 : / Check if the input are honestly prepared
2 : if k ≠ (k(1), ..., k(n)) or
3 :   ∃i, y(i) ≠ y or b(i) ≠ b or ¬CheckTrapdoorλ(d0[i], tk(i), k(i))
4 :   then return ⊥n+1 fi
5 : / Compute the correction α, the set of supported applicants, sample a first version of α̂
6 : α = ⊕i PartAlphaLoc(i, tk(i), y, b); S = {i | d0[i] = 1}; ∀i, α̂i = ⊕l r(l)[i]
7 : if S ≠ ∅ then / If at least one person is supported, ensure ⊕i∈S α̂i = α
8 :   j = maxi∈S i / Pick an arbitrary j ∈ S to change
9 :   α̂j = α ⊕ ⊕i∈S\{j} α̂i
10 : fi / Return α̂i to applicant i and ⊤ to the server to indicate no problem occurred.
11 : return (⊤, α̂1, ..., α̂n)

```

Fig. 3. The function to compute in the AUTH – BLIND_{can}^{dist} protocol in a MPC way. The first input is the input of the server, and the other inputs are from the applicants (the $y^{(i)}$ and $b^{(i)}$ are supposed to be equal to y and b and are just used to ensure that the server provided coherent inputs in the MPC, and $r^{(i)} \in \{0, 1\}^n$ is a string supposed to be sampled uniformly at random).

applicants. In this later case, if we denote by $\rho_{A,\mathcal{M}}$ the joint state of the honest applicants (register A) and of the adversary (register \mathcal{M}) obtained at the end of the protocol, then $\rho_{A,\mathcal{M}}$ can be written as a **Completely Positive Trace Preserving (CPTP)** map²⁰ applied on a GHZ state shared among all parties i such that $d_0[i] = 1$, in such a way that the CPTP map leaves untouched the qubits of the GHZ state owned by honest applicants i . In particular, if all supported parties are honest, they all share a GHZ state.

PROOF. The first action of the server (which is assumed to be honest here) is to run a ZK protocol to check that $\forall i, k^{(i)} \in \mathcal{K}_{\lambda, \text{Loc}}$ and $\text{Auth}_i(d_0[i], w_i) = 1$. Therefore, we can use the soundness property of the ZK protocol to claim that with overwhelming probability $\forall i, k^{(i)} \in \mathcal{K}_{\lambda, \text{Loc}}$ and there exist w_i such that $\text{Auth}_i(d_0[i], w_i) = 1$ (since the provers are the applicants, they are bounded so we can rely on both computational or statistical soundness). The fact that w_i is actually “known” to the applicant comes directly from the fact that the ZK protocol is a Proof of Knowledge and is extractable. So with overwhelming probability, $k := (k^{(1)}, \dots, k^{(n)})$ belongs to $\mathcal{K} := \mathcal{K}_{\lambda, \text{Loc}}^n$. Therefore, since the server is honest, with probability $1 - \delta$ it will measure a y such that $|f_k^{-1}(y)| = 2$, and due to the properties of the family f_k , the two preimages x and x' are such that $h(x) \oplus h(x') = d_0$. So the state sent by the server is $|h(x)\rangle + (-1)^a |h(x')\rangle$ with $a = \langle b, x \oplus x' \rangle$. Then, the MPC protocol will be performed. If the MPC abort, then we are already in the setting of the theorem. If the MPC does not aborts, then, due to the fact that in quantum mechanics, operations performed by two non-communicating parties commute, without any loss of generality we can assume that the honest applicants will apply the correction before the deviation of the malicious party. Moreover, since the honest corrections are unitary, we can also assume without any loss of generality that the first

²⁰Note that for simplicity, we just require the existence of this CPTP map, and therefore, it can depend on any quantity, including k, d_0, \dots . Therefore, we don't require this map to be efficiently computable since it will not be useful for us. However, a similar “efficient” version should be derivable if we make sure our ZK protocol is extractable, i.e., that the trapdoor of each $k^{(i)}$ can be extracted by a simulator. This is however out of the scope of this article.

step of the malicious party is to apply the honest correction on the state received from the server and then deviate (eventually by starting to undo the correction). Note that, we do not even ask this correction to be efficiently computable by the adversary (see Lemma 6.9) since we just claim that such deviation exists. Due to the definition of PartInfo, after applying the X correction, the parties i for which $\mathbf{d}_0[i] = 1$ share a state $|0 \dots 0\rangle + (-1)^\alpha |1 \dots 1\rangle$. Now, we have two cases:

- (1) If there exists at least one supported applicant which is malicious, then the proof is done; no matter what are the values of $\hat{\alpha}_i$ which will be used by the honest applicant to correct the state, we can always include in the CPTP map a first step that applies $Z^{\alpha \oplus \bigoplus_{i \in S, i \notin \mathcal{M}} \hat{\alpha}_i}$ on the qubit of the malicious applicant to map the step back to a GHZ state. Again, this is possible since we just require the existence of the CPTP map. Then, any CPTP deviation can be applied on the state owned by the malicious adversary, including undoing the previous Z and X corrections.
- (2) If there exists no malicious supported applicant, and if the probability of having no abort and no malicious supported applicant is non-negligible,²¹ then with overwhelming probability we must have $\bigoplus_{i \in S} \hat{\alpha}_i = \alpha$. Indeed, if it is not the case, then it is possible to distinguish the real world from the ideal world of the MPC computation. Therefore, after the Z correction the honest applicants having $\mathbf{d}_0[i] = 1$ will share a canonical GHZ, which ends the proof.

□

LEMMA 6.10 (BLINDNESS OF AUTH – BLIND_{can}^{dist} IN THE PRESENCE OF MALICIOUS APPLICANTS). *If the server corrupts a set of applicants, in such a way that at least one supported applicant is not corrupted, or that no supported applicant is corrupted, then the support status of the honest applicants is hidden in the AUTH – BLIND_{can}^{dist} protocol, beyond the fact that server knows whether or not they can pass the authorization step. More formally, no adversary can win the game IND – AUTH – BLIND_{can}^{dist}.*

IND – AUTH – BLIND_{can}^{dist}_{f_k}(λ)

```

1 : ( $\mathcal{M}, \mathbf{d}_0^{(0)}, \{(i, w_i^{(0)})\}_{i \in [n] \setminus \mathcal{M}}, \mathbf{d}_0^{(1)}, \{(i, w_i^{(1)})\}_{i \in [n] \setminus \mathcal{M}} \leftarrow \mathcal{A}_1(1^\lambda)$ 
2 : if  $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$  then return false fi
3 : if  $(\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)} = 1)$  and  $((\forall i \notin \mathcal{M}, \mathbf{d}_0^{(0)}[i] = 0) \text{ or } (\forall i \notin \mathcal{M}, \mathbf{d}_0^{(1)}[i] = 0))$ 
4 : then return false fi
   / Check that the adversary did not gave wrong witnesses  $w_i$ :
5 : if  $\exists i \in [n], \exists c \in \{0, 1\}, \text{Auth}_i(\mathbf{d}_0^{(c)}[i], w_i^{(c)}) \neq 1$  then return false fi
6 :  $c \xleftarrow{\$} \{0, 1\}; \mathbf{d}_0 := \mathbf{d}_0^{(c)}; w_i := w_i^{(c)}$ 
7 : Run with  $\mathcal{A}_2$  the protocol AUTH – BLINDcandist.
8 :  $\tilde{c} \leftarrow \mathcal{A}_3$ 
9 : return  $\tilde{c} = c$ 

```

PROOF. The above game is more formally defined in Figure 4. We will prove the above theorem by using a hybrid argument. First, we can easily see that if the adversary corrupts all applicants ($\mathcal{M} = [n]$), then it cannot win the game with probability better than $\frac{1}{2}$. Indeed, the line 2 forces $\mathbf{d}_0^{(0)} = \mathbf{d}_0^{(1)}$ (and both maps w_i are empty), therefore, the view of the adversary is exactly the

²¹If on the other hand this quantity is negligible, then this second case occurs with negligible probability so it is absorbed in the $\text{negl}(\lambda)$ of the theorem.

GAME1 ^{\mathcal{A}} (λ)	
1 :	$(\mathcal{M}, \mathbf{d}_0^{(0)}, \{(i, w_i^{(0)})\}_{i \in [n] \setminus \mathcal{M}}, \mathbf{d}_0^{(1)}, \{(i, w_i^{(1)})\}_{i \in [n] \setminus \mathcal{M}}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 :	if $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$ then return false fi
3 :	if $(\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)} = 1)$ and $((\forall i \notin \mathcal{M}, \mathbf{d}_0^{(0)}[i] = 0) \text{ or } (\forall i \notin \mathcal{M}, \mathbf{d}_0^{(1)}[i] = 0))$
4 :	then return false fi
	/ Check that the adversary did not gave wrong witnesses w_i :
5 :	if $\exists i \in [n], \exists c \in \{0, 1\}, \text{Auth}_i(\mathbf{d}_0^{(c)}[i], w_i^{(c)}) \neq 1$ then return false fi
6 :	$c \xleftarrow{\$} \{0, 1\}; \mathbf{d}_0 := \mathbf{d}_0^{(c)}; w_i := w_i^{(c)}$
7 :	$\forall i \notin \mathcal{M}, (k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$
8 :	$\mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$
9 :	for $i \notin \mathcal{M}$ do
10 :	Prove in ZK with $\mathcal{A}_{2,i}$ that $\text{CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)}) \wedge \text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$
11 :	endfor
12 :	if \mathcal{A}_2 aborts then wait for \tilde{c} from \mathcal{A}_2 . return $\tilde{c} = c$ fi
13 :	$(y, b) \leftarrow \mathcal{A}_3$
14 :	Compute in a MPC way the CombineAlpha function, where \mathcal{A}_4 controls adversaries in \mathcal{M} .
	/ All others operations are not sent to the adversary, and operations
	/ applied on the quantum state do not change anything due to non-signaling.
15 :	$\tilde{c} \leftarrow \mathcal{A}_5$
16 :	return $\tilde{c} = c$

Fig. 4. Definition of GAME1.

same for $c = 0$ and $c = 1$. So we can define a new hybrid game GAME2 in which we return false if $\mathcal{M} = [n]$:

GAME2 ^{\mathcal{A}}	
	/ Just update the line 2 of GAME1 as follows:
2 :	if $\mathcal{M} = [n]$ or $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$ then return false fi
	/ Rest is like GAME1...

Then, we can turn any adversary \mathcal{A} winning GAME1 with probability p into another adversary \mathcal{A}' winning GAME2 with probability p . To do so, \mathcal{A}' runs first \mathcal{A}_1 : if \mathcal{A}_1 returns $\mathcal{M} \neq [n]$, then \mathcal{A}' continues normally with \mathcal{A} , otherwise if $\mathcal{M} = [n]$ then \mathcal{A}' removes one element of \mathcal{M} (of course, $[n]$ is assumed to be non empty...), and \mathcal{A}' aborts when \mathcal{A}_2 is supposed to run, and output a random \tilde{c} . Therefore, we have

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME1}^{\mathcal{A}}](\lambda) = \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME2}^{\mathcal{A}'}](\lambda). \quad (26)$$

The second hybrid game that we define is the same as the game GAME2, except that we replace lines 14 to 15 with one line $(\tilde{c}, \tilde{y}) \leftarrow \text{REAL}_{\Pi, \mathcal{A}'}(\lambda, \vec{x}, \rho_3)$, where ρ_3 is the final internal state of the adversary \mathcal{A}_3 , \vec{x} contains the honest inputs of the MPC computation for the non-corrupted adversaries and dummy inputs for the corrupted adversaries (we will ignore them anyway), and

\mathcal{A}' is the adversary that outputs the corrupted set \mathcal{M} , that runs $\mathcal{A}_4(\rho_3)$ followed by $\tilde{c} \leftarrow \mathcal{A}_5(\rho_4)$, where ρ_4 is the final internal state of \mathcal{A}_4 , and that finally returns \tilde{c} . This defines a new game GAME3:

GAME3 ^{\mathcal{A}}

14 : ~~Compute in a MPC way the CombineAlpha function, where \mathcal{A}_4 controls adversaries in \mathcal{M} .~~

15 : ~~$\tilde{c} \leftarrow \mathcal{A}_5$~~ $(\tilde{c}, \vec{y}) \leftarrow \text{REAL}_{\Pi, \mathcal{A}'}(\lambda, \vec{x}, \rho_3)$

Since this is perfectly equivalent from the point of view of the adversary (due to the definition of REAL), the probability of winning these two games are exactly the same: $\Pr[\text{GAME2}^{\mathcal{A}}] = \Pr[\text{GAME3}^{\mathcal{A}}]$. Now, because the MPC protocol is secure, there exists a simulator Sim fitting Definition 3.8. Therefore, we can now define a new game GAME4, in which we replace the real world with the ideal world:

GAME4 ^{\mathcal{A}}

15 : ~~$(\tilde{c}, \vec{y}) \leftarrow \text{REAL}_{\Pi, \mathcal{A}'}(\lambda, \vec{x}, \rho_3)$~~ $(\tilde{c}, \vec{y}) \leftarrow \text{IDEAL}_{\text{CombineAlpha}, \text{Sim}}(\lambda, \vec{x}, \rho_3)$

Then, we have $\Pr[\text{REAL}_{\Pi, \mathcal{A}'}(\lambda, \vec{x}, \rho_3)[0] = c] \leq \Pr[\text{IDEAL}_{\text{CombineAlpha}, \text{Sim}}(\lambda, \vec{x}, \rho_3)[0] = c] + \text{negl}(\lambda)$ (otherwise, we could distinguish between the real and ideal worlds), and therefore, $\Pr[\text{GAME3}^{\mathcal{A}}] \leq \Pr[\text{GAME4}^{\mathcal{A}}] + \text{negl}(\lambda)$. Now, we can define CombineRandom, which is an adaptation of CombineAlpha that does not depend anymore on the secret values of the honest parties:

CombineRandom($((k, y, b), (t_k^{(1)}, \mathbf{d}_0[1], r^{(1)}, k^{(1)}, y^{(1)}, b^{(1)}), \dots, (t_k^{(n)}, \mathbf{d}_0[n], r^{(n)}, k^{(n)}, y^{(n)}, b^{(n)}))$)

1 : / Check if the input are honestly prepared

2 : **if** $k \neq (k^{(1)}, \dots, k^{(n)})$ or $\exists i, y^{(i)} \neq y$ or $b^{(i)} \neq b$ or $\exists i \in \mathcal{M}, \neg \text{CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)})$

3 : **then return** \perp^{n+1} **fi**

4 : $\forall i, \hat{\alpha}_i = \bigoplus_l r^{(l)}[i]$

5 : **return** $(\top, \hat{\alpha}_1, \dots, \hat{\alpha}_n)$

We can now define a new game GAME5 in which we substitute the CombineAlpha function with the CombineRandom function:

GAME5 ^{\mathcal{A}}

15 : ~~$(\tilde{c}, \vec{y}) \leftarrow \text{IDEAL}_{\text{CombineAlpha}, \text{Sim}}(\lambda, \vec{x}, \rho_3)$~~ $(\tilde{c}, \vec{y}) \leftarrow \text{IDEAL}_{\text{CombineRandom}, \text{Sim}}(\lambda, \vec{x}, \rho_3)$

Then, we have $\Pr[\text{GAME4}^{\mathcal{A}}] = \Pr[\text{GAME5}^{\mathcal{A}}]$. Indeed, by construction, the inputs of the honest parties always pass the CheckTrapdoor test, so removing this test for the honest parties cannot help the adversary to distinguish the two games. Moreover, since at least one applicant is honest, the string $\bigoplus_l r^{(l)}$ is indistinguishable from a random string. Therefore, we can use the same trick used already in the proof of Lemma 6.6: the condition line 3 gives us two cases.

- If all malicious applicants are not supported, then, since the output of honest applicants are never given back the adversary, we don't need to update $\hat{\alpha}_j$.
- Similarly, if at least one honest applicant is supported, then instead of updating $\hat{\alpha}_j$, we can update the $\hat{\alpha}_j$ of this applicant... But since the output of honest applicants are never given to the adversary, we don't even need to update it.

Therefore, $\Pr [\text{GAME4}^A] = \Pr [\text{GAME5}^A]$. In the next hybrid, we are going to remove completely the MPC computation, and the previous line that can now be merged in a single one:

GAME6 ^A	
12 :	if \mathcal{A}_2 aborts then wait for \tilde{c} from \mathcal{A}_2. return $\tilde{c} = c$ fi
13 :	$(y, b) \leftarrow \mathcal{A}_3$
12 :	$(\tilde{c}, \vec{y}) \leftarrow \text{IDEAL}_{\text{CombineRandom, Sim}}(\lambda, \vec{x}, \rho_3)$ $\tilde{c} \leftarrow \mathcal{A}_3$

The reason is that now, since the CombineRandom function does not depend on any secret own by honest parties, the input of honest parties can be replaced with wrong trapdoors $t_k^{(i)}$ and $\mathbf{d}_0[i]$. Therefore, since the adversary knows already $k^{(i)}$, it can simulate locally the ideal world. More precisely, from an adversary \mathcal{A} winning the game GAME5^A with probability p , we can create another adversary \mathcal{A}' winning the game GAME6^A with the same probability p : \mathcal{A}' will run \mathcal{A}_1 and \mathcal{A}_2 against the challenger, keeping locally the $k^{(i)}$. If \mathcal{A}_2 aborts and sends \tilde{c} , then it returns \tilde{c} directly. Otherwise \mathcal{A}' runs as a blackbox \mathcal{A}_3 to obtain (y, b) , and then it locally runs Sim to compute $\text{IDEAL}_{\text{CombineRandom, Sim}}(\lambda, \vec{x}, \rho_3)$, by feeding the input of honest parties input \vec{x} with the $k^{(i)}$ that it got before, $y^{(i)} := y$, $b^{(i)} := b$, $r^{(i)} \leftarrow_{\$} \{0, 1\}^n$ and since the values of $t_k^{(i)}$ do not matter anymore, it can put any value here. Finally, \mathcal{A}' outputs the \tilde{c} obtained from the simulation IDEAL. Therefore, since we do not change what is done, but who's doing what, we get

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME5}^A] (\lambda) = \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME6}^{\mathcal{A}'}] (\lambda). \quad (27)$$

So now, the game GAME6 is exactly like GAME1 except that the lines 12–15 are replaced with a single line $\tilde{c} \leftarrow \mathcal{A}_3$. We will now do a similar strategy to remove the ZK protocol (line 9). The first step is to formalize this line (we will also merge it with the next line). We define $\mathcal{R}_{\mathcal{L}}$ as the relation $(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)}) \in \mathcal{R}_{\mathcal{L}}(k^{(i)})$ iff $\text{CheckTrapdoor}_{\lambda}(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)}) \wedge \text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$, and P the honest ZK prover associated with $\mathcal{R}_{\mathcal{L}}$. If we define ρ_1 as the internal state at the end of \mathcal{A}_1 , and $V^*(\{k^{(i)}\}_{i \notin \mathcal{M}}, \rho_1) := (\rho_2 \leftarrow \mathcal{A}_2(\rho_1); \tilde{c} \leftarrow \mathcal{A}_3(\rho_2))$ then we can merge the line 12 of GAME6 with the line 9 as follows:

GAME7 ^A	
8 :	$\rho \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$ / We just explicit the internal state after \mathcal{A}_2
9 :	for $i \notin \mathcal{M}$ do
10 :	Prove in ZK with $\mathcal{A}_{2,i}$ that $\text{CheckTrapdoor}_{\lambda}(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)}) \wedge \text{Auth}_i(\mathbf{d}_0[i], w_i) = 1$
	$\rho \leftarrow \text{OUT}_{\mathcal{A}_{2,i}}(P(\mathbf{d}_0[i], t_k^{(i)}, w_i), \mathcal{A}_{2,i}(\rho))(k^{(i)})$
11 :	endfor

Since both games are exactly identical (up to the notation), we get $\Pr [\text{GAME6}^A] = \Pr [\text{GAME7}^A]$. Now, due to the fact that the MPC protocol respects the property *Quantum ZK* defined in Definition 3.5, there exist for all $i \notin \mathcal{M}$ a simulator Sim_i fitting Definition 3.5. To be completely formal, one should define a series of games in which we replace in the loop only one OUT at a time by the simulated version, and we can then claim that the probability of having $\tilde{c} = c$ in each hybrid game is negligibly close to the probability of having $\tilde{c} = c$ in the first game, otherwise we could distinguish between the real world and the ideal world. This gives us at the end a new game GAME8:

GAME8 ^A	
10 :	$\rho \leftarrow \text{OUT}_{\mathcal{A}_{2,i}}(P(\mathbf{d}_0[i], t_k^{(i)}, w_i), \mathcal{A}_{2,i}(\rho))(k^{(i)})$ $\rho \leftarrow \text{Sim}_i(k^{(i)}, V^*, \rho)$

And using the above argument, $\Pr [\text{GAME7}^{\mathcal{A}}] \leq \Pr [\text{GAME8}^{\mathcal{A}}] + \text{negl}(\lambda)$. Now, the simulators Sim_i can be fully simulated by the adversary since there is no more secret information. So, exactly like we did for the MPC computation, we can move the loop into the adversary:

GAME7 ^A	
8 :	$p \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$ $\tilde{c} \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$
9 :	for $i \notin \mathcal{M}$ do
10 :	$p \leftarrow \text{Sim}(k^{(i)}, V^*, p)$
11 :	endfor
12 :	$\tilde{c} \leftarrow \mathcal{A}_3$

and we get

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME8}^{\mathcal{A}}] (\lambda) = \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME9}^{\mathcal{A}'}] (\lambda). \quad (28)$$

So now, GAME9 is like GAME1 except that all lines starting from line 8 are replaced with a single line $\tilde{c} \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$. We see now that the conditions line 5 and line 3 can only decrease the probability of winning the game. Therefore, we can remove them (as well as w_i 's which are not used anymore). This gives us this new game (after removing empty lines):

GAME10 ^A (λ)	
1 :	$(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 :	if $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$ then return false fi
3 :	$c \xleftarrow{\$} \{0, 1\}; \mathbf{d}_0 := \mathbf{d}_0^{(c)}$
4 :	$\forall i \notin \mathcal{M}, (k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$
5 :	$\tilde{c} \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M}})$
6 :	return $\tilde{c} = c$

Since we only increase the probability of success, we have

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME9}^{\mathcal{A}}] (\lambda) \leq \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME10}^{\mathcal{A}'}] (\lambda). \quad (29)$$

Similarly, we can decide to give more advices to the adversary, by running Gen_{Loc} on all $i \in [n]$ instead of only on the $i \notin \mathcal{M}$:

GAME11 ^A	
4 :	$\forall i \notin \mathcal{M}$ $\forall i \in [n], (k^{(i)}, t_k^{(i)}) \leftarrow \text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$
5 :	$\tilde{c} \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \notin \mathcal{M} \cup [n]})$

Since we give more advice to the adversary, its probability of winning the game can only increase (it can always decide to drop this additional information). Therefore,

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME10}^{\mathcal{A}}] (\lambda) \leq \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME11}^{\mathcal{A}'}] (\lambda). \quad (30)$$

However, since Gen is defined as the concatenation of Gen_{Loc} , then this game is actually exactly the $\text{IND} - \text{D}\mathcal{O}_{\text{Gen}}^{\mathcal{A}}(\lambda)$ game defined Definition 4.1, and by assumption, the probability of winning this

$\text{Gen}_{\mathcal{P}}(1^\lambda, \mathbf{d}_0)$	$\text{Invert}_{\mathcal{P}}(t_k := (\mathbf{R}, \mathbf{d}_0, \mathbf{s}_0, \mathbf{e}_0, \mathbf{A}), \mathbf{y})$
1 : $(\mathbf{A}_u, \mathbf{R}) \leftarrow \text{MP.Gen}(1^\lambda)$	1 : $\left[\frac{\mathbf{y}_u \in \mathbb{Z}_q^M}{\mathbf{y}_l \in \mathbb{Z}_q^n} \right] := \mathbf{y}; \quad \left[\frac{\mathbf{A}_u}{\mathbf{A}_l} \right] := \mathbf{A}$
2 : $\mathbf{A}_l \xleftarrow{\$} \mathbb{Z}_q^n$	2 : $(\mathbf{s}, \mathbf{e}_u) \leftarrow \text{MP.Invert}(\mathbf{R}, \mathbf{A}, \mathbf{y}_u)$
3 : $\mathbf{A} := \left[\frac{\mathbf{A}_u}{\mathbf{A}_l} \right]$	3 : if $\mathbf{s} = \perp$ then return \perp fi
4 : $\mathbf{s}_0 \leftarrow \mathcal{D}_{\mathbb{Z}, \alpha q}^N$	4 : $\mathbf{d} := \text{RoundMod}_q(\mathbf{y}_l - \mathbf{A}_l \mathbf{s})$
5 : $\mathbf{e}_0 \leftarrow \mathcal{D}_{\mathbb{Z}, \alpha q}^{M+n}$	5 : $\mathbf{e} := \left[\frac{\mathbf{e}_u}{\mathbf{y}_l - \mathbf{A}_l \mathbf{s} - \mathbf{d}} \right]$
6 : $\mathbf{y}_0 := \mathbf{A} \mathbf{s}_0 + \mathbf{e}_0 + \frac{q}{2} \left\lceil \frac{\mathbf{0}^M}{\mathbf{d}_0} \right\rceil$	6 : $\mathbf{s}' := \mathbf{s} - \mathbf{s}_0; \quad \mathbf{e}' := \mathbf{e} - \mathbf{s}_0$
7 : $k := (\mathbf{A}, \mathbf{y}_0)$	7 : if $(\mathbf{s}, \mathbf{e}) \notin \mathcal{X}$ or $(\mathbf{s}', \mathbf{e}') \notin \mathcal{X}$ then
8 : $t_k := (\mathbf{R}, \mathbf{d}_0, \mathbf{s}_0, \mathbf{e}_0, \mathbf{A})$	8 : return \perp fi
9 : return (k, t_k)	9 : return $((\mathbf{s}, \mathbf{e}, 0, \mathbf{d}), (\mathbf{s}', \mathbf{e}', 1, \mathbf{d}' \oplus \mathbf{d}_0))$
<hr/>	
$\text{Eval}_{\mathcal{P}}(k := (\mathbf{A}, \mathbf{y}_0), x := (\mathbf{s}, \mathbf{e}, c, \mathbf{d})) = f_k(x)$	$h(x := (\mathbf{s}, \mathbf{e}, c, \mathbf{d}))$
1 : return $\mathbf{A} \mathbf{s} + \mathbf{e} + \left\lceil \frac{\mathbf{0}^M}{\mathbf{d}} \right\rceil + c \times \mathbf{y}_0$	1 : return \mathbf{d}

Fig. 5. Definition of the $\delta\text{-GHZ}^H$ capable family.

game is smaller than $\frac{1}{2} + \text{negl}(\lambda)$. Therefore, we also have

$$\max_{\text{QPT}, \mathcal{A}} \Pr \left[\text{GAME1}^{\mathcal{A}} \right] (\lambda) \leq \frac{1}{2} + \text{negl}(\lambda) \quad (31)$$

which ends the proof. \square

7 Function Construction

7.1 Construction of a $\delta\text{-GHZ}^H$ capable Family

In this section, we will explain how to derive a $\delta\text{-GHZ}^H$ capable family. See Section 2.3 to get an intuitive explanation of our method, which extends the construction introduced in [Cojocaru et al. 2019] (itself based on [Cojocaru et al. 2021]). In the following, MP.Gen and MP.Invert are the functions defined in [Micciancio and Peikert 2012] (to, respectively, generate a couple public key/trapdoor (\mathbf{A}, \mathbf{R}) and to invert the function $g_{\mathbf{A}}(\mathbf{s}, \mathbf{e}) := \mathbf{A} \mathbf{s} + \mathbf{e}$). Details are in Section B.3.

Definition 7.1. For the parameters $\mathcal{P} := (k, N, \alpha, r_{\max}, n, \mathcal{X})$ with $(k, N, n, r_{\max}) \in \mathbb{N}$, $0 < \alpha < 1$, $\mathcal{X} \subseteq \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n}$ where $q := 2^k$ and $M := N(1 + k)$, we define Figure 5 the algorithms $\text{Gen}_{\mathcal{P}}$, $\text{Invert}_{\mathcal{P}}$, $\text{Eval}_{\mathcal{P}}$ (to compute f_k) and h . We use $\mathbf{d}_0 \in \{0, 1\}^n \subseteq \mathbb{Z}_q^n$ (same for \mathbf{d}), $\mathbf{s}_0 \in \mathbb{Z}_q^N$, $\mathbf{e}_0 \in \mathbb{Z}_q^M$, $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}$, $c \in \{0, 1\}$, $\mathbf{A} \in \mathbb{Z}_q^{(M+n) \times N}$ and \mathbf{R} is the trapdoor obtained via the [Micciancio and Peikert 2012] algorithm.

While most of the correctness and security proofs are similar to [Cojocaru et al. 2019], we also need a worst-case analysis (for malicious applicants) and ensure that the functions are $\text{negl}(\lambda)$ -2-to-1 for appropriate parameters.

We obtain conditions to check that our function is δ -2-to-1 and have the XOR property explained in Definition 4.1, which is exactly what the H.CheckTrapdoor function (required later in Theorem 7.8) must do (i.e., check that the trapdoor is well formed, and that it has the properties given

in Lemma 7.2). Note that the δ given here is a kind of worst-case analysis (which is required if the key is maliciously sampled): on average, if we sample random functions we expect to have a smaller δ .

LEMMA 7.2 (CONDITIONS FOR f_k TO BE δ -2-TO-1). *Let \mathcal{P} be like in Definition 7.1. We define $r_{\text{safe}} = r_{\text{max}} - \alpha q \sqrt{N + M + n}$, $\mathcal{X} + (\hat{s}_0, \hat{e}_0) := \{(s + \hat{s}_0, e + \hat{e}_0) \mid (s, e) \in \mathcal{X}\}$ and*

$$\delta := 1 - \min \left\{ \frac{|\mathcal{X} \cap (\mathcal{X} + (\hat{s}_0, \hat{e}_0))|}{|\mathcal{X}|} \mid (\hat{s}_0, \hat{e}_0) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n}, \left\| \begin{bmatrix} \hat{s}_0 \\ \hat{e}_0 \end{bmatrix} \right\|_2 \leq \alpha q \sqrt{N + M + n} \right\}. \quad (32)$$

Let $s_0 \in \mathbb{Z}_q^N$, $e_0 \in \mathbb{Z}_q^{M+n}$, $d_0 \in \{0, 1\}$, $\hat{A} \in \mathbb{Z}_q^{N \times N}$, $A_l \in \mathbb{Z}_q^{n \times n}$ and $R = [R_1 \mid R_2] \in \mathbb{Z}_q^{Nk \times 2N}$. We define as before:

$$A_u := \left[\frac{\hat{A}}{G - R_2 \hat{A} - R_1} \right] \quad A := \left[\frac{A_u}{A_l} \right] \in \mathbb{Z}_q^{(M+n) \times N} \quad (33)$$

together with $y_0 := As_0 + e_0 + \frac{q}{2} \lceil \frac{q}{d_0} \rceil$ and $k := (A, y_0)$.

If $\sqrt{\sigma_{\text{max}}(R)^2 + 1} \leq \frac{q}{4r_{\text{max}}}$, $\| \lceil \frac{s_0}{e_0} \rceil \|_2 \leq \alpha q \sqrt{N + M + n}$, and

$$X \subseteq \left\{ (s, e) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n} \mid \left\| \begin{bmatrix} s \\ e \end{bmatrix} \right\|_2 \leq r_{\text{safe}} \right\} \quad (34)$$

then the function $f_k(x)$ described in Definition 7.1 is δ -2-to-1, trapdoor, and for any y having exactly two preimages x and x' , we have $x \neq x'$, $h(x) \oplus h(x') = d_0$.

On the other hand, if (k, t_k) is sampled according to $\text{Gen}(1^\lambda, d_0)$, if Equation (34) is true, and if

$$\sqrt{\left(C \times \alpha q \times \sqrt{N}(\sqrt{k} + \sqrt{2} + 1) \right)^2 + 1} \leq \frac{q}{4r_{\text{max}}} \quad (35)$$

then with overwhelming probability on N , the function f_k is δ -2-to-1, trapdoor, and for any y having exactly two preimages x and x' , we have $x \neq x'$, $h(x) \oplus h(x') = d_0$.

PROOF. Let us first prove that for all $c \in \{0, 1\}$, the function $f_k(\cdot, \cdot, c, \cdot)$ is injective. Let $c \in \{0, 1\}$, and s, e, d, s', e', d' be such that $f_k(s, e, c, d) = f_k(s', e', c, d')$. Then, if we consider only the upper part of this equation (and denote e_u the upper part of the vector e), we get $A_u s + e_u + c \times y_{0,l} = A_u s' + e'_u + c \times y_{0,l}$, i.e., $A_u s + e_u = A_u s' + e'_u$. But because $\sqrt{\sigma_{\text{max}}(R)^2 + 1} \leq \frac{q}{4r_{\text{max}}}$ and due to the condition on \mathcal{X} given in Equation (34) and the fact that $r_{\text{safe}} \leq r_{\text{max}}$, according to Lemma B.9 the function $(s, e) \rightarrow A_u s + e$ is injective. So $s = s'$ and $e_u = e'_u$. Now, we focus on the upper part of the above equation: we have $A_l s + e_l + \frac{q}{2} d + c \times y_{0,l} = A_l s' + e'_l + \frac{q}{2} d' + c \times y_{0,l}$. Because $s = s'$, we obtain $e_l + \frac{q}{2} d = e'_l + \frac{q}{2} d'$. Because $1 \leq \sqrt{\sigma_{\text{max}}(R)} + 1 < \frac{q}{4r_{\text{max}}}$, we have $r_{\text{max}} < \frac{q}{4}$. Therefore, we get for all i , $|e_l[i]| < \frac{q}{4}$, so $\text{RoundMod}_q(e_l[i] + \frac{q}{2} d[i]) = \text{RoundMod}_q(e'_l[i] + \frac{q}{2} d'[i])$, i.e., $d[i] = d'[i]$. So $d = d'$ and, therefore, we also get $e_l = e'_l$: the function $f_k(\cdot, \cdot, c, \cdot)$ is injective.

Therefore, f_k has at most two preimages, one for $c = 0$ and one for $c = 1$. We now prove that $f_k(s, e, 0, d) = f_k(s', e', 1, d')$, iff $(s', e', d') = (s - s_0, e - e_0, d \oplus d_0)$. One implication is trivial: if $(s', e', d') = (s - s_0, e - e_0, d \oplus d_0)$ then because q is even, $f_k(s, e, 0, d) = f_k(s', e', 1, d')$. We prove now the second implication. By definition of f_k , if we consider again the upper part of the equation and replace y_0 by its definition, we have $A_u s + e_u = A_u(s' + s_0) + (e'_u + e_{0,u})$. But the triangle inequality gives

$$\left\| \begin{bmatrix} s' + s_0 \\ e'_u + e_{0,u} \end{bmatrix} \right\|_2 \leq \left\| \begin{bmatrix} s' \\ e'_u \end{bmatrix} \right\|_2 + \left\| \begin{bmatrix} s_0 \\ e_{0,u} \end{bmatrix} \right\|_2 \leq r_{\text{safe}} + \alpha q \sqrt{N + M + n} = r_{\text{max}}.$$

Therefore, we can use again the injectivity property given in Lemma B.9, which gives $(\mathbf{s}, \mathbf{e}_u) = (\mathbf{s}' + \mathbf{s}_0, \mathbf{e}' + \mathbf{e}_{0,u})$. We can now analyse the lower part of the equation: $\mathbf{A}_l \mathbf{s} + \mathbf{e}_l + \frac{q}{2} \mathbf{d} = \mathbf{A}_l (\mathbf{s}' + \mathbf{s}_0) + (\mathbf{e}'_l + \mathbf{e}_{0,l}) + \frac{q}{2} (\mathbf{d}' + \mathbf{d}_0)$. Because $\mathbf{s} = \mathbf{s}' + \mathbf{s}_0$ and q is even, we have $\mathbf{e}_l + \frac{q}{2} \mathbf{d} = (\mathbf{e}'_l + \mathbf{e}_{0,l}) + \frac{q}{2} (\mathbf{d}' + \mathbf{d}_0)$. Using again the triangle inequality, we prove the same way that $\|\mathbf{e}'_l + \mathbf{e}_{0,l}\|_2 < \frac{q}{2}$. As before, by rounding the previous equation using RoundMod_q , we obtain $\mathbf{d} = \mathbf{d}' \oplus \mathbf{d}_0$ and $\mathbf{e}_l = \mathbf{e}'_l$ which concludes the proof. In particular, if x and x' are the two preimage, we have $h(x) \oplus h(x') = \mathbf{d} \oplus \mathbf{d}' = \mathbf{d}_0$. We remark that this proof follows exactly the algorithm *Invert*, therefore the correctness of *Invert* follows quite directly and thus the function is trapdoor.

Now, we prove that the function f_k is δ -2-to-1. The total number of elements in the domain of f_k is $2|\mathcal{X}| \times 2^n$. Let $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}$ and $\mathbf{d} \in \{0, 1\}^n$. Then, using the result proven above, $f_k(\mathbf{s}, \mathbf{e}, 0, \mathbf{d})$ has a second preimages $(\mathbf{s} - \mathbf{s}_0, \mathbf{e} - \mathbf{e}_0, 1, \mathbf{d} \oplus \mathbf{d}_0)$ iff $(\mathbf{s} - \mathbf{s}_0, \mathbf{e} - \mathbf{e}_0) \in \mathcal{X}$, i.e., iff $(\mathbf{s}, \mathbf{e}) \in \mathcal{X} + (\mathbf{s}_0, \mathbf{e}_0)$. So the number of elements x such that $|f_k^{-1}(f_k(x))| = 2$ is equal to $2|\mathcal{X} \cap (\mathcal{X} + (\mathbf{s}_0, \mathbf{e}_0))|2^n$, therefore, if we define

$$\delta_k := 1 - \frac{2|\mathcal{X} \cap (\mathcal{X} + (\mathbf{s}_0, \mathbf{e}_0))|2^n}{2|\mathcal{X}|2^n} = 1 - \frac{|\mathcal{X} \cap (\mathcal{X} + (\mathbf{s}_0, \mathbf{e}_0))|}{|\mathcal{X}|} \quad (36)$$

this function is δ_k -2-to-1. But $\|[\frac{\mathbf{s}_0}{\mathbf{e}_0}]\|_2 \leq \alpha q \sqrt{N + M + n}$, so by definition of δ , $\delta_k \leq \delta$. So f_k is also δ -2-to-1.

To prove the last part of the theorem, we use Equation (35) and Lemma B.9: with overwhelming probability (on N), $\sqrt{\sigma_{\max}(\mathbf{R})^2 + 1} < \frac{q}{4r_{\max}}$. Moreover, because $[\frac{\mathbf{s}_0}{\mathbf{e}_0}]$ is sampled according to $\mathcal{D}_{\mathbb{Z}, \alpha q}^{N+(M+n)}$, we get, using²² Lemma B.3, that with overwhelming probability (on $N + M + n$): $\|[\frac{\mathbf{s}_0}{\mathbf{e}_0}]\|_2 \leq \alpha q \sqrt{N + M + n}$. We can now end the proof using the first (already proven) part of the theorem. \square

Note that we did not yet give an explicit definition of \mathcal{X} . The most natural way to define \mathcal{X} may be to define it following Equation (34) as

$$\mathbf{X}_\bullet := \left\{ (\mathbf{s}, \mathbf{e}) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n} \mid \left\| \left[\frac{\mathbf{s}}{\mathbf{e}} \right] \right\|_2 \leq r_{\text{safe}} \right\}. \quad (37)$$

However, for our protocol to work, one needs to be able to create quantumly a uniform superposition over all elements in \mathcal{X} . A first naive method would be a *rejection sampling* method: we create a uniform superposition $\sum_x |x\rangle$ over the hypercube of length $2r_{\text{safe}}$ (see later how to do), add an auxiliary qubit, set (in superposition) this qubit to 1 if $\|x\|_2 \leq r_{\text{safe}}$, and to 0 otherwise, and we finally measure it. If the output is 0, we discard the state and start again from scratch, otherwise we have the wanted state. Unfortunately, this method is inefficient as the probability of not rejecting is negligible when N tends to the infinity. While some more efficient methods may exist, it will be easier to focus rather on \mathcal{X} being a hypercube.

Definition 7.3. For any $(\mu, N, M, n) \in \mathbb{N}^4$, we define the hypercube $\mathcal{X}_{\bullet\mu}$ as

$$\mathbf{X}_{\bullet\mu} := \left\{ (\mathbf{s}, \mathbf{e}) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n} \mid \left\| \left[\frac{\mathbf{s}}{\mathbf{e}} \right] \right\|_{\infty} \leq \mu \right\}. \quad (38)$$

Remark 7.4. It is now easy to sample from $\mathcal{X}_{\bullet\mu}$: we can do a rejection sampling as explained above, except that we proceed coordinate per coordinate (this is much more efficient than doing a rejection sampling on the final high dimensional state): if we use the binary two's complement notation, we apply Hadamard gates on $\lceil \log_2(2\mu + 1) \rceil$ qubits, and use an auxiliary qubit to check

²²Note that the original lemma applies to Gaussian distributions that are not reduced modulo q , but reducing the Gaussian distribution modulo q can only decrease the length of the vector.

if the state is projected on the superposition of elements having size $\leq 2\mu + 1$ (this should happen with probability $\frac{2\mu+1}{2^{\lceil \log_2(2\mu+1) \rceil}} \geq 1/2$). If the test passes, we add $|0\rangle$ “significant qubits” until having $k = \log_2(q)$ qubits, and run the quantum unitary that subtracts μ modulo q . We repeat until having $N + M + n$ successful projections (this requires, therefore, $O(N + M + n)$ samplings). Moreover, we can even get completely rid of the rejection sampling if we slightly change the definition of \mathcal{X}_{μ} and if we make it less symmetric by asking that there exists $k' \in \mathbb{N}$ such that for all i , $[\frac{s}{e}][i] \in [-2^{k'}, 2^{k'} - 1]$. The superposition procedure is the same except that we work on $k' + 1$ qubits, and we don’t need the rejection sampling. However this notation slightly complicates the computations with no clear benefit (if simplifies slightly the sampling part, but it may decrease the value of δ since the term $2^{k'}$ must be a power of 2), so for simplicity we will keep our initial notation. Note also that in the following, for the sake of simplicity we won’t try to give tight bounds.

LEMMA 7.5. *Let $(N, M, n, \mu) \in \mathbb{N}^4$, $\mathcal{X} = \mathcal{X}_{\mu}$, $\alpha \in (0, 1)$, δ be as in Equation (32) and $\mu' := \lfloor \mu - \alpha q \sqrt{N + M + n} \rfloor$. Then if $\mu' \geq 0$:*

$$\delta \leq 1 - \left(\frac{2\mu' + 1}{2\mu + 1} \right)^{N+M+n} \leq \frac{(\alpha q + 1)(N + M + n)^{3/2}}{\mu + 1/2}. \quad (39)$$

PROOF. Let $(s, e) \in \mathcal{X}_{\mu}$ and $(\hat{s}_0, \hat{e}_0) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^{M+n}$ such that $\|[\frac{\hat{s}_0}{\hat{e}_0}]\|_2 \leq \alpha q \sqrt{N + M + n}$. Then, $(s, e) \in \mathcal{X}_{\mu} + (\hat{s}_0, \hat{e}_0)$ iff $(s - \hat{s}_0, e - \hat{e}_0) \in \mathcal{X}$, i.e., iff $\|[\frac{s - \hat{s}_0}{e - \hat{e}_0}]\|_{\infty} \leq \mu$. If we assume that $\|[\frac{s}{e}]\|_{\infty} \leq \mu'$ (this will not be super tight, but it is good enough for our analysis), then

$$\left\| \left[\frac{s - \hat{s}_0}{e - \hat{e}_0} \right] \right\|_{\infty} \leq \left\| \left[\frac{s}{e} \right] \right\|_{\infty} + \left\| \left[\frac{\hat{s}_0}{\hat{e}_0} \right] \right\|_{\infty} \leq \left\| \left[\frac{s}{e} \right] \right\|_{\infty} + \left\| \left[\frac{\hat{s}_0}{\hat{e}_0} \right] \right\|_2 \leq \mu. \quad (40)$$

Therefore, $\mathbf{X}_{\mu'} \subseteq \mathbf{X}_{\mu} \cap (\mathbf{X}_{\mu} + (\hat{s}_0, \hat{e}_0))$. So $|\mathbf{X}_{\mu} \cap (\mathbf{X}_{\mu} + (\hat{s}_0, \hat{e}_0))| \geq |\mathbf{X}_{\mu'}| = (2\mu' + 1)^{N+M+n}$, and because $|\mathcal{X}_{\mu}| = (2\mu + 1)^{N+M+n}$, we get

$$\begin{aligned} \delta &\leq 1 - \left(\frac{2\mu' + 1}{2\mu + 1} \right)^{N+M+n} = 1 - \left(1 + \left(-1 + \frac{2\mu' + 1}{2\mu + 1} \right) \right)^{N+M+n} \\ &\leq 1 - \left(1 + (N + M + n) \left(-1 + \frac{2\mu' + 1}{2\mu + 1} \right) \right) \quad (\text{Bernouilli's inequality}) \\ &= (N + M + n) \left(1 - \frac{2\mu' + 1}{2\mu + 1} \right) \\ &\leq (N + M + n) \left(1 - \frac{2\mu - 2\alpha q \sqrt{N + M + n} + 3}{2\mu + 1} \right) \quad (\text{Remove } \lfloor \cdot \rfloor) \\ &= (N + M + n) \left(\frac{\alpha q \sqrt{N + M + n} + 1}{\mu + 1/2} \right) \\ &\leq \frac{(\alpha q + 1)(N + M + n)^{3/2}}{\mu + 1/2}. \end{aligned} \quad (41)$$

□

LEMMA 7.6 (CONDITIONS ON PARAMETERS). *Let $\lambda \in \mathbb{N}$ be a security parameter and let $(k, n) \in \mathbb{N}$ and $\alpha \in (0, 1)$ be parameters that depend on λ , and $C \approx \frac{1}{\sqrt{2\pi}}$ (see Lemma B.9). We define $N := \lambda$,*

$q = 2^k, M := N(1 + k),$

$$r_{\max} := \frac{q}{4\sqrt{(C \times \alpha q \times \sqrt{N}(\sqrt{k} + \sqrt{2} + 1))^2 + 1}}, \quad (42)$$

$$r_{\text{safe}} := r_{\max} - \alpha q \sqrt{N + M + n}, \quad (43)$$

$$\mu := \left\lfloor \frac{r_{\text{safe}}}{\sqrt{N + M + n}} \right\rfloor, \quad (44)$$

$$\mathcal{X} := \mathcal{X}_{\bullet\mu}, \quad (45)$$

$$\delta_m := \frac{(\alpha q + 1)(N + M + n)^{3/2}}{\mu + 1/2}. \quad (46)$$

Then, if $\left| \mu - \alpha q \sqrt{N + M + n} \right| \geq 0$, the construction given in Definition 7.1 is δ_m -GHZ^H capable (see Definition 4.1) assuming the security of decision-LWE _{$q, \mathcal{D}_{\mathbb{Z}, \alpha q}$} . Moreover, if we define

$$\alpha_0 := \frac{1}{q} \sqrt{(\alpha q)^2 - \omega(\sqrt{\log N})^2} \quad (47)$$

$$\gamma := \tilde{O}\left(\frac{N}{\alpha_0}\right) \quad (\text{See constants in [Peikert et al. 2017]})$$

and if $\alpha_0 q > 2\sqrt{N}$, then the construction is secure if GapSVP _{γ} is hard. In particular, we are interesting in the regime in which δ_m is negligible (correctness) and in which $\gamma = \tilde{O}(2^{N^\epsilon})$ for some $\epsilon \in (0, 1/2)$ (security).

PROOF. For the first part of the theorem, the efficient generation and computation properties are trivial to check. The fact that the function is trapdoor and the property on the XOR is a direct consequence of Lemma 7.2. The δ_m -2-to-1 property comes from Lemma 7.2 and Lemma 7.5. The method to efficiently create a uniform superposition of elements in \mathcal{X} is given in Remark 7.4.

To prove the indistinguishability property on game IND - D0, we assume that there exists an adversary \mathcal{A} that can win this game with a non-negligible advantage. Because \mathcal{A} has only access to $(\mathbf{A}, \mathbf{y}_0 = \mathbf{A}\mathbf{s}_0 + \mathbf{e}_0 + \frac{q}{2} \lfloor \frac{\mathbf{0}^M}{\mathbf{d}_0^{(c)}} \rfloor)$, we can use \mathcal{A} to break the decision-LWE problem: given a challenge $(\mathbf{A}', \mathbf{y})$, we run the adversary and send to \mathcal{A} the couple $(\mathbf{A}', \mathbf{y} + \frac{q}{2} \lfloor \frac{\mathbf{0}^M}{\mathbf{d}_0^{(c)}} \rfloor)$. If the guess \tilde{c} of \mathcal{A} equals c , we guess that we get the non-uniform distribution normal- $A_{\mathbf{s}, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ (i.e., the distribution where \mathbf{s} is also sampled according to $\mathcal{D}_{\mathbb{Z}, \alpha q}$), otherwise we guess that we get the uniform distribution U . We remark that if \mathbf{y} is a vector chosen uniformly at random, then the distribution of $\mathbf{y} + \frac{q}{2} \lfloor \frac{\mathbf{0}^M}{\mathbf{d}_0^{(c)}} \rfloor$ is statistically uniform. So in that case, \mathcal{A} cannot guess c with probability better than 1/2. Now, if \mathbf{y} is sampled from normal- $A_{\mathbf{s}, \mathcal{D}_{\mathbb{Z}, \alpha q}}$, then \mathcal{A} must guess correctly the value of c with non-negligible advantage, otherwise it means that we can distinguish a matrix obtained by MP.Gen from a uniform matrix (we already know it is not possible, see Lemma B.7). Therefore, the probability p of guessing the correct distribution is

$$p \geq \frac{1}{2} \times \frac{1}{2} + \frac{1}{2} \left(\frac{1}{2} + \text{negl}(\lambda) \right) = \frac{1}{2} + \text{negl}(\lambda) \quad (48)$$

which is absurd since we assumed the hardness of LWE. Note that we don't exactly have an instance of decision-LWE, because it is an instance of the normal version of decision-LWE (\mathbf{s}_0 is sampled according to the same distribution as \mathbf{e}_0). However, Lemma B.5 shows that the normal problem is harder, and keeping only K samples can also only make the problem harder. Therefore, no adversary can win IND - D0 with non-negligible advantage.

For the second part, if we assume GapSVP_γ to be hard, then using Lemma 3.11 we get that decision- $\text{LWE}_{q, \mathcal{D}_{\alpha_0 q}}$ is also hard ($\alpha_0 q > 2\sqrt{N}$, and $0 < \alpha_0 \leq \alpha < 1$). We can now discretize the distribution using Corollary B.4 ($\lambda = N$) to obtain that decision- $\text{LWE}_{q, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ is hard. Indeed, if we assume the existence of an adversary \mathcal{A} that can distinguish with non-negligible advantage the distribution U from $\mathcal{A}_{s, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ for a vector s chosen uniformly at random, then \mathcal{A} has also a non-negligible advantage in distinguishing $\mathcal{A}_{s, \chi}$ where χ is the marginal distribution of e in Corollary B.4 (otherwise we could use \mathcal{A} to distinguish χ from $\mathcal{D}_{\mathbb{Z}, \alpha q}$, which is impossible because they are statistically negligibly close). But it also means that \mathcal{A} can also be used to break $\text{LWE}_{q, \mathcal{D}_{\alpha_0 q}}$ by first discretizing $\mathcal{D}_{\alpha_0 q}$ (it works because the transformation given in Corollary B.4 also maps the uniform distribution on itself). Since we already prove the security when we assume the hardness of decision- $\text{LWE}_{q, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ above, the proof is finished. \square

We prove now that there exists an instantiation that fulfills the requirements of Lemma 7.6. Note that for simplicity, we only verify the properties asymptotically. Moreover, we do not attempt to give any particularly optimized construction (note that there is a tradeoff between security, correctness, and simplicity of the quantum superposition preparation circuit).

THEOREM 7.7 (EXISTENCE OF A $\text{negl}(\lambda)$ -GHZ^H CAPABLE FAMILY). *Let $\varepsilon \in (0, \frac{1}{2})$ be a constant, and $\lambda \in \mathbb{N}$ be a security parameter. Let $n = \text{poly}(\lambda) \in \mathbb{N}$ and $N := \lambda$. If we assume the hardness of the GapSVP_γ problem for any $\gamma = \tilde{O}(2^{N^\varepsilon})$, then there exists a $\text{negl}(\lambda)$ -GHZ^H capable family of functions. More precisely, if we define the fixed function $\omega(\sqrt{\log N}) := \log N$, $k := \lfloor N^\varepsilon \rfloor$, $q := 2^k$,*

$$\alpha := \frac{\sqrt{4N + \omega(\sqrt{\log N})^2 + 1}}{q} \quad (49)$$

and $M, r_{\max}, r_{\text{safe}}, \mu, \mathcal{X}, \delta_m$ as in Lemma 7.6, the construction given in Definition 7.1 is δ_m -GHZ^H capable (for sufficiently large²³ λ), with $\delta_m = \text{negl}(\lambda)$.

PROOF. We just need to check that for sufficiently large λ the properties of Lemma 7.6 are respected. Because $1/q = \text{negl}(N)$, it is easy to see that for sufficiently large λ , $a \in (0, 1)$ since $\alpha = \text{poly}(N)/q$. Then, $\alpha q > \sqrt{4N + \omega(\sqrt{\log N})^2}$, so using notation from Lemma 7.6, we directly get $\alpha_0 q > 2\sqrt{N}$. Moreover, multiplying Equation (49) by q , we get $\alpha q = \text{poly}(N)$, so it means that $\alpha_0 = \text{poly}(N)/q = \text{negl}(\lambda)$. Therefore, $\gamma = \tilde{O}(N/\alpha_0) = \tilde{O}(\sqrt{N}q) = \tilde{O}(\sqrt{N}2^{N^\varepsilon}) = \tilde{O}(2^{N^\varepsilon})$ which is assumed to be hard. Next, let us study μ and δ_m . Because $\alpha q = \text{poly}(n)$ and $q = 1/\text{negl}(\lambda)$ is superpolynomial, that r_{\max} is also superpolynomial, and same for r_{safe}, μ and $\lfloor \mu - \alpha q \sqrt{N + M + N} \rfloor$ (we only subtract or divide by terms that are $\text{poly}(N)$). Therefore, for a large enough $\lambda (= N)$, we have $\lfloor \mu - \alpha q \sqrt{N + M + N} \rfloor \geq 0$. Finally, $\delta_m = \text{poly}(N)/(\mu + 1/2)$. But we showed that μ is superpolynomial, so $\frac{1}{\mu}$ is negligible, and therefore, δ_m is also negligible, which ends the proof. \square

7.2 Generic Construction to Create Distributable δ' -GHZ^{can} capable Primitives from δ -GHZ^H capable Primitives

We prove in this section that we can create a distributable δ' -GHZ^{can} capable family of functions from a δ -GHZ^H capable family having a small assumption.

²³The function may not be well defined for too small λ because the input set may be empty. For example, with this instantiation, if we take $\varepsilon = \frac{1}{3}$, the function is well defined for $N \geq 7 \times 10^5$. Moreover, when $N = 6 \times 10^6$, we get $k = 181$ and $\delta_m < 2^{-80}$. There is surely place for optimisation, but only existence matters here.

$\text{Gen}_{\text{Loc}}(1^\lambda, \mathbf{d}_0[i])$	$h((x^{(1)}, \dots, x^{(n)}))$
1 : $(k^{(i)}, t_k^{(i)}) \leftarrow \text{H.Gen}(1^\lambda, \mathbf{d}_0[i])$ return $(k^{(i)}, t_k^{(i)})$	1 : return $h(x^{(1)}) \dots h(x^{(n)})$
$\text{Eval}_{\mathcal{P}}((k^{(1)}, \dots, k^{(n)}), (c, \bar{x}^{(1)}, \dots, \bar{x}^{(n)}))$	$\text{PartAlpha}_{\text{Loc}}(i, t_k^{(i)}, y, b)$
1 : return $(\text{H.Eval}(k^{(1)}, (c, \bar{x}^{(1)})),$	1 : $(y^{(1)}, \dots, y^{(n)}) := y; \quad (b_c, b^{(1)}, \dots, b^{(n)}) := b$
2 : $\dots, \text{H.Eval}(k^{(n)}, (c, \bar{x}^{(n)})))$	2 : $\{(0, \bar{x}), (1, \bar{x}')\} \leftarrow \text{H.Invert}(t_k^{(i)}, y)$
$\text{PartInfo}_{\text{Loc}}(t_k^{(i)}, y)$	3 : if $\bar{x} = \perp$ or $\bar{x}' = \perp$ then return \perp fi
1 : if $\mathbf{d}_0(t_k^{(i)}) = 0$ then return \perp fi	4 : if $i = 1$ then return $b_c \oplus \langle b^{(i)}, \bar{x} \oplus \bar{x}' \rangle$
2 : $\{(0, \bar{x}), (1, \bar{x}')\} \leftarrow \text{H.Invert}(t_k^{(i)}, y)$	5 : else return $\langle b^{(i)}, \bar{x} \oplus \bar{x}' \rangle$ fi
3 : if $\bar{x} = \perp$ or $\bar{x}' = \perp$ then return \perp	$\text{CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)})$
4 : return $\text{H.h}((0, \bar{x}))$	1 : return $\text{H.CheckTrapdoor}_\lambda(\mathbf{d}_0[i], t_k^{(i)}, k^{(i)})$

Fig. 6. Construction distributable δ' -GHZ^{can} capable family.

THEOREM 7.8. *Let $\delta \in [0, 1]$, and $\{f_k\}_{k \in \mathcal{K}}$ be a δ -GHZ^H capable family²⁴ of functions, such that $f_k(x)$ can be written as $f_k((c, \bar{x}))$ with $c \in \{0, 1\}$ a bit labelling the preimage,²⁵ i.e., such that when a given y has exactly two preimages, one preimage has the form $(0, \bar{x})$ and the other $(1, \bar{x}')$. Then there exists a family $\{f'_k\}_{k \in \mathcal{K}'}$ which is a distributable δ' -GHZ^{can} capable family with $\delta' = 1 - (1 - \delta)^n \leq \delta n$. In particular, if δ is negligible (and n polynomial) then δ' is negligible. Moreover, if the family $\{f_k\}$ admits a circuit $\text{H.CheckTrapdoor}_\lambda(\mathbf{d}_0, t_k, k)$ that returns 1 iff t_k is the trapdoor of k , $k \in \mathcal{K}$ and $\mathbf{d}_0 = \mathbf{d}_0(t_k)$, then there exists a function CheckTrapdoor for $\{f'_k\}$ having the properties from Definition 4.4.*

The family $\{f'_k\}_{k \in \mathcal{K}'}$ can be obtained by generating n independent functions in $\{f_k\}$ (one for each $\mathbf{d}_0[i]$). More precisely, we define in Figure 6 the precise construction (where H.Gen , H.Enc , H.Invert and H.Eval are coming from the family $\{f_k\}$).

PROOF. Most of the properties are simple to check. We just precise the δ' -2-to-1 proof and blindness. First, we show that the function f'_k are δ' -2-to-1 with $\delta = 1 - (1 - \delta)^n$. Let $\#_2(f)$ be the number of images having exactly 2 preimages by f (we will call this kind of preimages “twin”), and $|\mathcal{K}|$ the number of elements in \mathcal{K} . Then by definition, for all k , $1 - \delta \leq \#_2(f_k)/|\mathcal{K}|$. Let $k' = (k^{(1)}, \dots, k^{(n)}) \in \mathcal{K}^n = \mathcal{K}'$, we want to show that $1 - \delta' := (1 - \delta)^n \leq \#_2(f'_{k'})/(|\mathcal{K}'|)$. First, we compute $\#_2(f'_{k'})$. Because of the assumption of the shape $(0, \bar{x})$ and $(1, \bar{x}')$ of all the couples of preimages, we can define for any $k^{(i)}$ the sets

$$A_0^{(i)} := \{\bar{x} \mid f_{k^{(i)}}^{-1}(f_{k^{(i)}}(x)) = \{(0, \bar{x}), (1, \bar{x}')\}\} \quad (50)$$

and

$$A_1^{(i)} := \{\bar{x}' \mid f_{k^{(i)}}^{-1}(f_{k^{(i)}}(x)) = \{(0, \bar{x}), (1, \bar{x}')\}\}. \quad (51)$$

Moreover, due to this same condition, we have $|A_0^{(i)}| = |A_1^{(i)}| = \frac{1}{2}\#_2(f_{k^{(i)}})$. Now, we compute a lower bound on the number of twin preimages of $f'_{k'}$. Let $(\bar{x}^{(1)}, \dots, \bar{x}^{(n)}) \in A_0^{(1)} \times \dots \times A_0^{(n)}$: then

²⁴In fact we only require this function to work when \mathbf{d}_0 is a single bit.

²⁵This is quite similar to the concept of claw-free functions used in [Mahadev 2018b].

for all i there exists a unique $\bar{x}^{(1)'} \in A_1^{(i)}$ such that $f_{k^{(i)}}((0, \bar{x}^{(i)})) = f_{k^{(i)}}((1, \bar{x}^{(i)}))$. So

$$f_{k'}(0, \bar{x}^{(1)}, \dots, \bar{x}^{(n)}) = f_{k^{(1)}}(0, \bar{x}^{(1)}) \dots |f_{k^{(n)}}(0, \bar{x}^{(n)}) \quad (52)$$

$$= f_{k^{(1)}}(1, \bar{x}^{(1)'}) \dots |f_{k^{(n)}}(1, \bar{x}^{(n)'}) \quad (53)$$

$$= f_{k'}(1, \bar{x}^{(1)'}, \dots, \bar{x}^{(n)'}) \quad (54)$$

So we found at least one different preimage, and due to the uniqueness of the above $\bar{x}^{(1)'}$ it is the only second preimage. Therefore,

$$\frac{\#_2(f_{k'})}{|\mathcal{X}'|} = \frac{2 \times |A_0^{(1)}| \times \dots \times |A_0^{(n)}|}{2 \left(\frac{|\mathcal{X}|}{2}\right)^n} \quad (55)$$

$$= \frac{\#_2(f_{k^{(1)}}) \times \dots \times \#_2(f_{k^{(n)}})}{(|\mathcal{X}|)^n} \quad (56)$$

$$\geq (1 - \delta)^n. \quad (57)$$

Which concludes the proof.

To prove the inequality, we use the Bernoulli's inequality: since $\delta \in [0, 1]$ and n is a non-negative integer, we get: $(1 - \delta)^n \geq 1 - \delta n$, so

$$\delta' = 1 - (1 - \delta)^n \leq 1 - (1 - \delta n) = \delta n. \quad (58)$$

Since the keys of $\{f'_k\}$ are keys of $\{f_k\}$ (except that \mathbf{d}_0 is a single bit), the properties of CheckTrapdoor come directly from the properties of H.CheckTrapdoor. All the other correctness properties are true by construction.

The security is quite intuitive: since all trapdoors are independently sampled, if one can learn information about the d_0 sampled by another party, then it can break the $\text{IND} - \text{D}\emptyset_{\text{Gen}}^{\mathcal{A}}(\lambda)$ game of f_k . More formally, because of the properties on Gen_{Loc} , the game $\text{IND} - \text{PARTIAL}_{\text{Gen, PartInfo}}^{\mathcal{A}}(\lambda)$ can equivalently be rewritten as follows:

GAME1

```

1 :  $(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$ 
2 : if  $\exists i \in \mathcal{M}, \mathbf{d}_0^{(0)}[i] \neq \mathbf{d}_0^{(1)}[i]$  : return false fi
3 :  $c \xleftarrow{\$} \{0, 1\}$ 
4 :  $\forall i, (k^{(i)}, t_k^{(i)}) \leftarrow \text{H.Gen}(1^\lambda, \mathbf{d}_0^{(c)}[i])$ 
5 :  $y \leftarrow \mathcal{A}_2(k^{(1)}, \dots, k^{(n)})$ 
6 :  $\forall i, v[i] \leftarrow \text{PartInfo}_{\text{Loc}}(i, t_k^{(i)}, y)$ 
7 :  $\tilde{c} \leftarrow \mathcal{A}_3(\{(i, v[i])\}_{i \in \mathcal{M}})$ 
8 : return  $\tilde{c} = c$ 

```

Then, we can define the following game in which the sampling of malicious $t_k^{(i)}$, the computing of $\text{PartInfo}_{\text{Loc}}$ and the initial condition are removed:

GAME2	
1 :	$(\mathcal{M}, \mathbf{d}_0^{(0)}, \mathbf{d}_0^{(1)}) \leftarrow \mathcal{A}_1(1^\lambda)$
2 :	$c \xleftarrow{\$} \{0, 1\}$
3 :	$\forall i \in \mathcal{M}, (k^{(i)}, t_k^{(i)}) \leftarrow \text{H.Gen}(1^\lambda, \mathbf{d}_0^{(c)}[i])$
4 :	$\tilde{c} \leftarrow \mathcal{A}_2(\{k^{(i)}\}_{i \in \mathcal{M}})$
5 :	return $\tilde{c} = c$

Then, because \mathcal{A}_2 in GAME2 can do itself the sampling and computing done in GAME1, and because the removing the condition line 2 can only increase the probability of winning the game, we have

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME1}^{\mathcal{A}}] (\lambda) \leq \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME2}^{\mathcal{A}'}] (\lambda). \quad (59)$$

Then, we define a series a hybrid games in which we gradually replace the $\text{H.Gen}(1^\lambda, \mathbf{d}_0^{(c)}[i])$'s with $\text{H.Gen}(1^\lambda, 0)$, which \mathcal{Z} starting from \emptyset until $\mathcal{Z} = \mathcal{M}$.

GAME3 $_{\mathcal{Z}}$	
3 :	$\forall i \in \mathcal{M}, (k^{(i)}, t_k^{(i)}) \leftarrow \text{H.Gen}(1^\lambda, \mathbf{d}_0^{(c)}[i])$
	$\forall i \in \mathcal{M}, \text{if } i \in \mathcal{Z} \text{ then } x := 0 \text{ else } x := \mathbf{d}_0^{(c)}[i] \text{ fi}; (k^{(i)}, t_k^{(i)}) \leftarrow \text{H.Gen}(1^\lambda, x)$

This is possible because the game IND – D0 of H.Gen can be seen as a CPA secure encryption (where Gen is the concatenation of the key generation and encryption of \mathbf{d}_0), itself equivalent to semantic security: i.e., the ciphertext does not give any advice on the clear text. Therefore, we can replace the clear text with 0 for example, which gives for all j :

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{GAME3}_{\mathcal{Z}}^{\mathcal{A}}] (\lambda) \leq \max_{\text{QPT}, \mathcal{A}'} \Pr [\text{GAME3}_{\mathcal{Z} \cup \{j\}}^{\mathcal{A}'}] (\lambda). \quad (60)$$

At the end of the hybrid series, when $\mathcal{Z} = \mathcal{M}$, we get a final game where no information about c is given to the adversary: it is, therefore, impossible to win this game $\text{GAME3}_{\mathcal{M}}$ with probability better than $\frac{1}{2}$. Therefore, we get

$$\max_{\text{QPT}, \mathcal{A}} \Pr [\text{IND} - \text{PARTIAL}_{\text{Gen, PartInfo}}^{\mathcal{A}}] (\lambda) \leq \frac{1}{2} + \text{negl}(\lambda). \quad (61)$$

Which ends the proof.

Finally, the security of the game IND – PARTIAL implies the security of IND – D0 since when $\mathcal{M} = \emptyset$, both games are equivalent. \square

Acknowledgement

We thank AirFrance for the unscheduled one night delay of a flight which was instrumental in starting this line of research. We are also very grateful to Alex Grilo, Chris Peikert, Florian Bourse, Geoffroy Couteau, the StackExchange user [Mark](#), Michael Reichle, Thomas Vidick, and Romain Gay for very valuable discussions.

References

Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. 2021. Post-quantum multi-party computation. In *Proceedings of the Advances in Cryptology – EUROCRYPT 2021*, Anne Canteaut and François-Xavier Standaert (Eds.). Lecture Notes in Computer Science, Springer International Publishing, Cham, 435–464. DOI : https://doi.org/10.1007/978-3-030-77870-5_16

- Gorjan Alagic, Andrew M. Childs, Alex B. Grilo, and Shih-Han Hung. 2020. Non-interactive classical verification of quantum computation. In *Proceedings of the Theory of Cryptography*, Rafael Pass and Krzysztof Pietrzak (Eds.). Lecture Notes in Computer Science, Springer International Publishing, Cham, 153–180. DOI : https://doi.org/10.1007/978-3-030-64381-2_6
- Martin Albrecht, Rachel Player, and Sam Scott. 2015. On the concrete hardness of learning with errors. *Journal of Mathematical Cryptology* 9 (2015). DOI : <https://doi.org/10.1515/jmc-2015-0016>
- Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. 2009a. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proceedings of the Advances in Cryptology - CRYPTO 2009, 29th Annual International Cryptology Conference (Lecture Notes in Computer Science)*, Vol. 5677. Springer, 595–618. DOI : https://doi.org/10.1007/978-3-642-03356-8_35
- Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. 2009b. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In *Proceedings of the Advances in Cryptology - CRYPTO 2009*, Shai Halevi (Ed.). Springer, Berlin, 595–618.
- László Babai and Shlomo Moran. 1988. Arthur-Merlin games: A randomized proof system, and a hierarchy of complexity classes. *Journal of Computer and System Sciences* 36, 2 (1988), 254–276. DOI : [https://doi.org/10.1016/0022-0000\(88\)90028-1](https://doi.org/10.1016/0022-0000(88)90028-1)
- Christian Badertscher, Alexandru Cojocaru, Léo Colisson, Elham Kashefi, Dominik Leichte, Atul Mantri, and Petros Wallden. 2020. Security limitations of classical-client delegated quantum computing. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 667–696.
- Wojciech Banaszczyk. 1993. New bounds in some transference theorems in the geometry of numbers. *Mathematische Annalen* 296, 1 (1993), 625–635.
- James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. 2021. On the round complexity of two-party quantum computation. In *Advances in Cryptology CRYPTO 2021*, T. Malkin and C. Peikert (Eds.). Springer International Publishing, Cham, 406–435.
- Charles H. Bennett, Herbert J. Bernstein, Sandu Popescu, and Benjamin Schumacher. 1996. Concentrating partial entanglement by local operations. *Physical Review A* 53, 4 (1996), 2046.
- Charles H. Bennett, Gilles Brassard, Claude Crépeau, Richard Jozsa, Asher Peres, and William K. Wootters. 1993. Teleporting an unknown quantum state via dual classical and einstein-podolsky-rosen channels. *Physical Review Letters* 70, 13 (1993), 1895.
- Charles H. Bennett, Gilles Brassard, Sandu Popescu, Benjamin Schumacher, John A. Smolin, and William K. Wootters. 1996. Purification of noisy entanglement and faithful teleportation via noisy channels. *Physical Review Letters* 76, 5 (1996), 722.
- Charles H. Bennett, David P. DiVincenzo, John A. Smolin, and William K. Wootters. 1996. Mixed-state entanglement and quantum error correction. *Physical Review A* 54, 5 (1996), 3824.
- Nir Bitansky and Omri Shmueli. 2020. Post-quantum zero knowledge in constant rounds. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, 269–279.
- Manuel Blum, Paul Feldman, and Silvio Micali. 1988. Non-interactive zero-knowledge and its applications. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC'88)*. Association for Computing Machinery, New York, NY, USA, 103–112. DOI : <https://doi.org/10.1145/62212.62222>
- Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. 2014. Fully key-homomorphic encryption, arithmetic circuit ABE and compact garbled circuits. In *Proceedings of the Advances in Cryptology – EUROCRYPT 2014*, Phong Q. Nguyen and Elisabeth Oswald (Eds.). Springer, Berlin, 533–556.
- Zvika Brakerski. 2018a. Quantum FHE (almost) as secure as classical. In *Proceedings of the Advances in Cryptology – CRYPTO 2018*, Hovav Shacham and Alexandra Boldyreva (Eds.). Lecture Notes in Computer Science, Springer International Publishing, 67–95. DOI : https://doi.org/10.1007/978-3-319-96878-0_3
- Zvika Brakerski. 2018b. Quantum FHE (almost) as secure as classical. In *Proceedings of the Annual International Cryptology Conference*. Springer, 67–95.
- Zvika Brakerski, Paul Christiano, Urmila Mahadev, Umesh Vazirani, and Thomas Vidick. 2018. A cryptographic test of quantumness and certifiable randomness from a single quantum device. In *Proceedings of the 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS'18)*. IEEE, 320–331.
- Zvika Brakerski and Nico Döttling. 2018. Two-message statistically sender-private OT from LWE. In *Proceedings of the Theory of Cryptography*, Amos Beimel and Stefan Dziembowski (Eds.). Lecture Notes in Computer Science, Vol. 11240, Springer International Publishing, 370–390. DOI : https://doi.org/10.1007/978-3-030-03810-6_14
- A. Broadbent and A. B. Grilo. 2020. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. In *Proceedings of the 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS'20)*, 196–205. DOI : <https://doi.org/10.1109/FOCS46700.2020.00027>
- Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. 2016. Zero-knowledge proof systems for QMA. In *Proceedings of the 2016 IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS'16)*. IEEE, 31–40.

- Matthias Christandl and Stephanie Wehner. 2005. Quantum anonymous transmissions. In *Proceedings of the International Conference on the Theory and Application of Cryptology and Information Security*. Springer, 217–235.
- Alexandru Cjocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. 2019. QFactory: Classically-instructed remote secret qubits preparation. In *Proceedings of the Advances in Cryptology – ASIACRYPT 2019*, Steven D. Galbraith and Shihō Moriai (Eds.). Springer International Publishing, 615–645.
- Alexandru Cjocaru, Léo Colisson, Elham Kashefi, and Petros Wallden. 2021. On the possibility of classical client blind quantum computing. *Cryptography* 5, 1 (2021), 3.
- Andrea Coladangelo, Thomas Vidick, and Tina Zhang. 2020. Non-interactive zero-knowledge arguments for QMA, with preprocessing. In *Proceedings of the Annual International Cryptology Conference*. Springer, 799–828.
- Léo Colisson, Garazi Muguruza, and Florian Speelman. 2023. Oblivious Transfer from Zero-Knowledge Proofs, or How to Achieve Round-Optimal Quantum Oblivious Transfer and Zero-Knowledge Proofs on Quantum States. In *Advances in Cryptology ASIACRYPT 2023*, J. Guo and R. Steinfeld (Eds.). Springer Nature, Singapore, 338. arXiv:2303.01476. Retrieved from <https://arxiv.org/abs/2303.01476>
- Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. 2019. Security of the fiat-shamir transformation in the quantum random-oracle model. In *Proceedings of the Advances in Cryptology – CRYPTO 2019*, Alexandra Boldyreva and Daniele Micciancio (Eds.). Lecture Notes in Computer Science, Springer International Publishing, Cham, 356–383. DOI: https://doi.org/10.1007/978-3-030-26951-7_13
- Yfke Dulek, Alex B. Grilo, Stacey Jeffery, Christian Majenz, and Christian Schaffner. 2020. Secure multi-party quantum computation with a dishonest majority. In *Proceedings of the Advances in Cryptology – EUROCRYPT 2020*, Anne Canteaut and Yuval Ishai (Eds.). Lecture Notes in Computer Science, Springer International Publishing, 729–758. DOI: https://doi.org/10.1007/978-3-030-45727-3_25
- Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. 2012. Actively secure two-party evaluation of any quantum operation. In *Proceedings of the Advances in Cryptology – CRYPTO 2012*, Reihaneh Safavi-Naini and Ran Canetti (Eds.). Lecture Notes in Computer Science, Springer, Berlin, 794–811. DOI: https://doi.org/10.1007/978-3-642-32009-5_46
- Amos Fiat and Adi Shamir. 1987. How to prove yourself: Practical solutions to identification and signature problems. In *Proceedings of the Advances in Cryptology – CRYPTO’86*, Andrew M. Odlyzko (Ed.). Lecture Notes in Computer Science, Springer, Berlin, 186–194. DOI: https://doi.org/10.1007/3-540-47721-7_12
- Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008a. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC’08)*. Association for Computing Machinery, 197–206. DOI: <https://doi.org/10.1145/1374376.1374407>
- Craig Gentry, Chris Peikert, and Vinod Vaikuntanathan. 2008b. Trapdoors for hard lattices and new cryptographic constructions. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing (STOC’08)*. Association for Computing Machinery, New York, NY, USA, 197–206. DOI: <https://doi.org/10.1145/1374376.1374407>
- A. Gheorghiu and T. Vidick. 2019. Computationally-secure and composable remote state preparation. In *Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*. 1024–1033. DOI: <https://doi.org/10.1109/FOCS.2019.00066>
- Oded Goldreich. 2004. *Foundations of Cryptography. II: Basic Applications*. DOI: <https://doi.org/10.1017/CBO9780511721656>
- O. Goldreich, S. Micali, and A. Wigderson. 1987. How to play ANY mental game. In *Proceedings of the 19th Annual ACM Symposium on Theory of Computing (STOC’87)*. Association for Computing Machinery, New York, NY, USA, 218–229. DOI: <https://doi.org/10.1145/28395.28420>
- Oded Goldreich, Silvio Micali, and Avi Wigderson. 1991. Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems. *Journal of the ACM* 38, 3 (1991), 690–728. DOI: <https://doi.org/10.1145/116825.116852>
- Oded Goldreich and Yair Oren. 1994. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology* 7, 1 (1994), 1–32. DOI: <https://doi.org/10.1007/BF00195207>
- Shafi Goldwasser, Silvio Micali, and Charles Rackoff. 1989. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing* 18, 1 (1989), 186–208. DOI: <https://doi.org/10.1137/0218012>
- Daniel M. Greenberger, Michael A. Horne, and Anton Zeilinger. 1989. Going beyond Bell’s theorem. In *Proceedings of the Bell’s Theorem, Quantum Theory and Conceptions of the Universe*. Springer, 69–72.
- Alex B. Grilo, William Slofstra, and Henry Yuen. 2019. Perfect zero knowledge for quantum multiprover interactive proofs. In *Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS’19)*. IEEE, 611–635.
- Mark Hillery, Vladimír Bužek, and André Berthiaume. 1999. Quantum secret sharing. *Physical Review A* 59, 3 (1999), 1829.
- R. Impagliazzo and M. Yung. 1988. Direct minimum-knowledge computations (Extended Abstract). In *Advances in Cryptology – CRYPTO’87*, C. Pomerance (Ed.). Springer, Berlin, Heidelberg, 40–51.
- T. Kapourniotis, E. Kashefi, D. Leichtle, L. Music, and H. Ollivier. 2025. Asymmetric secure multi-party quantum computation with weak clients against Dishonest Majority. *Quantum Science and Technology* 10, 2, (Feb. 2025), 025015.
- Qipeng Liu and Mark Zhandry. 2019. Revisiting post-quantum fiat-shamir. In *Proceedings of the Advances in Cryptology – CRYPTO 2019*, Alexandra Boldyreva and Daniele Micciancio (Eds.). Lecture Notes in Computer Science, Springer International Publishing, Cham, 326–355. DOI: https://doi.org/10.1007/978-3-030-26951-7_12

- Urmila Mahadev. 2018a. Classical homomorphic encryption for quantum circuits. In *Proceedings of the 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 332–338. DOI : <https://doi.org/10.1109/FOCS.2018.00039>
- Urmila Mahadev. 2018b. Classical homomorphic encryption for quantum circuits. In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7–9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 332–338.
- Urmila Mahadev. 2018c. Classical verification of quantum computations. In *Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7–9, 2018*, Mikkel Thorup (Ed.). IEEE Computer Society, 259–267.
- D. Mayers and A. Yao. 1998. Quantum cryptography with imperfect apparatus. In *Proceedings 39th Annual Symposium on Foundations of Computer Science (Cat. No.98CB36280)*. 503–509. DOI : <https://doi.org/10.1109/SFCS.1998.743501>
- Clément Meignant, Damian Markham, and Frédéric Grosshans. 2019. Distributing graph states over arbitrary quantum networks. *Physical Review A* 100, 5 (2019), 052333. DOI : <https://doi.org/10.1103/PhysRevA.100.052333>
- Tony Metger, Yfke Dulek, Andrea Coladangelo, and Rotem Arnon-Friedman. 2021. Device-independent quantum key distribution from computational assumptions. *New Journal of Physics*.
- Daniele Micciancio and Chris Peikert. 2012. Trapdoors for lattices: Simpler, tighter, faster, smaller. *Lecture Notes in Computer Science* (2012), 700–718. DOI : https://doi.org/10.1007/978-3-642-29011-4_41
- Daniele Micciancio and Oded Regev. 2009. Lattice-based cryptography. In *Post-Quantum Cryptography*, D. J. Bernstein, J. Buchmann, and E. Dahmen (Eds.). Springer, Berlin, Heidelberg, 147–191.
- Tomoyuki Morimae and Yuki Takeuchi. 2020. Trusted center verification model and classical channel remote state preparation. arXiv:2008.05033. Retrieved from <https://arxiv.org/abs/2008.05033>
- Tomoyuki Morimae and Takashi Yamakawa. 2022. Classically verifiable (dual-mode) NIZK for QMA with preprocessing. In *Advances in Cryptology ASIACRYPT 2022*, S. Agrawal and D. Lin (Eds.). Springer Nature Switzerland, Cham, 599–627.
- Michael A. Nielsen and Isaac Chuang. 2000. *Quantum Computation and Quantum Information*. Cambridge University Press.
- Chris Peikert. 2010. An efficient and parallel gaussian sampler for lattices. In *Proceedings of the Advances in Cryptology – CRYPTO 2010*, David Hutchison, Takeo Kanade, Josef Kittler, Jon M. Kleinberg, Friedemann Mattern, John C. Mitchell, Moni Naor, Oscar Nierstrasz, C. Pandu Rangan, Bernhard Steffen, Madhu Sudan, Demetri Terzopoulos, Doug Tygar, Moshe Y. Vardi, Gerhard Weikum, and Tal Rabin (Eds.). Vol. 6223, Springer, Berlin, 80–97. DOI : https://doi.org/10.1007/978-3-642-14623-7_5
- Chris Peikert. 2016. A decade of lattice cryptography. *Foundations and Trends in Theoretical Computer Science* 10, 4 (2016), 283–424. DOI : <https://doi.org/10.1561/04000000074>
- Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. 2017. Pseudorandomness of ring-LWE for any ring and modulus. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC’17)*. Association for Computing Machinery, New York, NY, USA, 461–473. DOI : <https://doi.org/10.1145/3055399.3055489>
- Chris Peikert and Brent Waters. 2011. Lossy trapdoor functions and their applications. *SIAM Journal on Computing* 40, 6 (2011), 1803–1844.
- A. Priker, J. Wallnöfer, and W. Dür. 2018. Modular architectures for quantum networks. *New Journal of Physics* 20, 5 (2018), 053054. DOI : <https://doi.org/10.1088/1367-2630/aac2aa>
- Oded Regev. 2005. On lattices, learning with errors, random linear codes, and cryptography. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing – STOC ’05*. ACM Press. DOI : <https://doi.org/10.1145/1060590.1060603>
- C. P. Schnorr. 1987. A hierarchy of polynomial time lattice basis reduction algorithms. *Theoretical Computer Science* 53, 2 (1987), 201–224.
- Omri Shmueli. 2021. Multi-theorem (malicious) designated-verifier NIZK for QMA. In *Advances in Cryptology CRYPTO 2021*, T. Malkin and C. Peikert (Eds.). Springer International Publishing, Cham, 375–405.
- Dominique Unruh. 2012. Quantum proofs of knowledge. In *Proceedings of the Advances in Cryptology – EUROCRYPT 2012*, David Pointcheval and Thomas Johansson (Eds.). Springer, Berlin, 135–152.
- Thomas Vidick and John Watrous. 2016. Quantum proofs. *Foundations and Trends in Theoretical Computer Science* 11, 1-2 (2016), 1–215. DOI : <https://doi.org/10.1561/04000000068>
- Thomas Vidick and Tina Zhang. 2020. Classical zero-knowledge arguments for quantum computations. *Quantum* 4 (2020), 266.
- John Watrous. 2009. Zero-knowledge against quantum attacks. *SIAM Journal on Computing* 39, 1 (2009), 25–58.

SUPPLEMENTARY MATERIAL

A Details on ZK

In our last protocol, in order to get stronger guarantees, we may also want to ensure that the prover “knows” the secret. Therefore, to get stronger guarantees, we will require the ZK protocol to also be

a Proof of Knowledge protocol. Intuitively, we would like to check that any malicious prover P^* that can convince a verifier with a non-negligible probability has the witness w “encoded in its source code or memory”. We formalize this notion by saying that there exist a QPT circuit K , the extractor, which can recover w with non-negligible probability from a full description of P^* and its input.

This is usually enough, since being able to obtain a witness with non-negligible probability is usually enough to break the security: for example it could be used to forge a signature and break the unforgeability property as explained in [Unruh 2012]. More precisely:

Definition A.1 (Post-Quantum ZK Proof of Knowledge [Unruh 2012]). We say that a Post-Quantum ZK protocol (P, V) for a relation $\mathcal{R}_{\mathcal{L}}$ is a *Proof of Knowledge* protocol, if it is *quantum extractable* with knowledge error $\kappa = \text{negl}(\lambda)$, i.e., if there exists a constant $d > 0$, a polynomially-bounded function $p > 0$, and a QPT K such that for any interactive QPT malicious prover P^* , any polynomial l , any security parameter $\lambda \in \mathbb{N}$, any state ρ , and any $x \in \{0, 1\}^\lambda$, we have

$$\Pr [\text{OUT}_V(P^*(\rho), V)(x) = 1] \geq \kappa(\lambda) \\ \implies \Pr [w \in \mathcal{R}_{\mathcal{L}}(x) \mid w \leftarrow K(P^*, \rho, x)] \geq \frac{1}{p(\lambda)} (\Pr [\text{OUT}_V(P^*(\rho), V)(x) = 1] - \kappa(\lambda))^d$$

B Details on LWE

B.1 Definitions

We provide here a more in-depth introduction to LWE.

Definition B.1 (LWE [Regev 2005]).

Let $N \in \mathbb{N}$,²⁶ $q = q(N) \in \mathbb{N}_{\geq 2}$ be a modulus and χ a distribution on \mathbb{R} (χ may be continuous or $\chi \subseteq \mathbb{Z}$ (we will then say that χ is discrete), and will always be reduced modulo q). For any $s \in \mathbb{Z}_q^N$ we define $A_{s, \chi}$ as the distribution on $\mathbb{Z}_q^N \times [0, q)$ obtained by sampling $a \in \mathbb{Z}_q^N$ uniformly at random, $e \leftarrow \chi$ and outputting $(a, a^T s + e \bmod q)$.

We say that an algorithm solves the *search-LWE* $_{q, \chi}$ problem (in the worst case, with overwhelming probability) if for any $s \in \mathbb{Z}_q^N$, given an arbitrary number of samples from $A_{s, \chi}$, it outputs s with overwhelming probability. We say that an algorithm solves the *decision-LWE* $_{q, \chi}$ problem (on the average, with non-negligible advantage) if it can distinguish with non-negligible advantage between the distribution $A_{s, \chi}$ where $s \leftarrow \mathbb{Z}_q^N$, and the uniform distribution $U := \mathcal{U}(\mathbb{Z}_q^N \times [0, q))$ (when χ is discrete, we consider instead the uniform distribution $U := \mathcal{U}(\mathbb{Z}_q^N \times \mathbb{Z}_q)$).

We can also formulate this problem using matrices by grouping a fixed number $M \in \mathbb{N}$ of samples: For any $s \in \mathbb{Z}_q^N$, we can sample $A \leftarrow \mathbb{Z}_q^{M \times N}$ and $e \leftarrow \chi^M$ a (typically small) vector where each of its component is sampled according to χ . Then let $b := As + e$. The *search* problem consists in finding s given (A, b) . The *decision* problem consists in deciding, when receiving a couple (A, b) , if b has been sampled uniformly at random (over $[0, q)^M$ if χ is continuous, or over \mathbb{Z}_q^M if χ is discrete) or if b has been sampled according to the procedure described above (and therefore $b = As + e$).

Note that having access to less samples can only make the problem harder. On the other hand, one can show that having access to polynomially many samples is enough to generate arbitrary many further samples, with only a minor degradation in the error [Albrecht et al. 2015; Applebaum et al. 2009a; Gentry et al. 2008b]. The worst-case and average-case search problems are in fact equivalent: [Regev 2005, Lemma 4.1] shows how it is possible to turn a distinguisher that can solve the decision-LWE problem with non-negligible advantage into a better distinguisher that can solve the decision-LWE problem with overwhelming probability.

²⁶While usually the parameter N is written in lowercase, we will use this notation here since n already represents the number of applicants.

The distribution χ can be instantiated in many different ways: for example when χ is always equal to 0, these problems are trivial, and when χ is uniform, they are impossible. In practice, χ is usually a (discrete or continuous) Gaussian (or more rarely a rounded Gaussian [Regev 2005]):

Definition B.2 (Continuous and Discrete Gaussian). For any $s \in \mathbb{R}_{>0}$ and any vector $\mathbf{x} \in \mathbb{R}^N$, we define $\rho_s(\mathbf{x}) := \exp(-\pi(\frac{\|\mathbf{x}\|_2}{s})^2) = \exp(-\pi\mathbf{x}^T\mathbf{x}/s^2)$. By applying a linear transformation on \mathbf{x} , we can generalize this notion: for any positive-definite matrix $\Sigma > \mathbf{0}$, we define

$$\rho_{\sqrt{\Sigma}}(\mathbf{x}) := \exp(-\pi \cdot \mathbf{x}^T \Sigma^{-1} \mathbf{x}). \quad (62)$$

In particular, if $\Sigma = s^2\mathbf{I}$, we have $\rho_s = \rho_{\sqrt{\Sigma}}$: therefore, s will be used as a shortcut for $\sqrt{\Sigma} = \sqrt{s^2\mathbf{I}}$. The normalization of the expression gives $\int_{\mathbb{R}^n} \rho_{\sqrt{\Sigma}}(\mathbf{x}) = \sqrt{\det \Sigma}$ and $\int_{\mathbb{R}^n} \rho_s(\mathbf{x}) = s^N$. We can now define the continuous Gaussian distribution:

$$\mathcal{D}_s^N(\mathbf{x}) := \rho_s(\mathbf{x})/s^N \quad \mathcal{D}_{\sqrt{\Sigma}}(\mathbf{x}) := \rho_{\sqrt{\Sigma}}(\mathbf{x})/\sqrt{\det \Sigma}. \quad (63)$$

Note that sampling from \mathcal{D}_s^N is equivalent to sampling each component from $\mathcal{D}_s := \mathcal{D}_s^1$. Moreover, due to our choice of normalization, s and Σ are not exactly equal to the standard deviation and to the covariance matrix: \mathcal{D}_s has standard deviation $\sigma := s/(\sqrt{2\pi})$ and the actual covariance of $\mathcal{D}_{\sqrt{\Sigma}}$ is $\Sigma' := \Sigma/(2\pi)$.

If $\Lambda \subseteq \mathbb{R}^N$ is a lattice (i.e., a discrete additive subgroup of \mathbb{R}^N), we define for any $\mathbf{c} \in \mathbb{R}^N$ the coset $\Lambda + \mathbf{c} = \{\mathbf{x} + \mathbf{c} \mid \mathbf{x} \in \Lambda\}$ and $\rho_{\sqrt{\Sigma}}(\Lambda + \mathbf{c}) := \sum_{\mathbf{x} \in \Lambda + \mathbf{c}} \rho_{\sqrt{\Sigma}}(\mathbf{x})$. We can define now the discrete Gaussian on $\Lambda + \mathbf{c}$ by simply normalizing $\rho_{\sqrt{\Sigma}}(\mathbf{x})$. For any $\mathbf{x} \in \mathbb{R}^N$, if $\mathbf{x} \notin \Lambda + \mathbf{c}$, $\mathcal{D}_{\Lambda + \mathbf{c}, \sqrt{\Sigma}}(\mathbf{x}) = 0$ and if $\mathbf{x} \in \Lambda + \mathbf{c}$:

$$\mathcal{D}_{\Lambda + \mathbf{c}, \sqrt{\Sigma}}(\mathbf{x}) := \frac{\rho_{\sqrt{\Sigma}}(\mathbf{x})}{\rho_{\sqrt{\Sigma}}(\Lambda + \mathbf{c})}. \quad (64)$$

In the following, for $\alpha \in (0, 1)$, the discrete Gaussian on \mathbb{Z} , $\mathcal{D}_{\mathbb{Z}, \alpha q} = \mathcal{D}_{\mathbb{Z}, \sqrt{(\alpha q)^2\mathbf{I}}}$, will be particularly important to sample the noise.

The next lemma is useful to bound the length of a vector sampled according to a discrete Gaussian.

LEMMA B.3 (PARTICULAR CASE OF [BANASZCZYK 1993, LEMMA 1.5][MICCIANCIO AND PEIKERT 2012, LEMMA 2.6]). *For any $s > 0$, we have*

$$\Pr \left[\|\mathbf{x}\|_2 \geq s\sqrt{n} \mid \mathbf{x} \leftarrow \mathcal{D}_{\mathbb{Z}, s}^n \right] \leq 2^{-n}. \quad (65)$$

B.2 Best Algorithms to Solve LWE

To have strong security guarantees, we usually want γ to be polynomial in N , but this is sometimes impossible. According to [Peikert 2016] the best algorithm for solving these problems in polynomial time works for only slightly sub-exponential approximation factor $\gamma = 2^{\Theta(N \log \log N / \log N)}$. They also mention the algorithm [Schnorr 1987] that provides a tradeoff between the approximation γ and the running time: an approximation of $\gamma = 2^k$ can be obtained in time $2^{\tilde{\Theta}(N/k)}$. This suggest that there is no efficient algorithm for $\gamma = 2^{N^\epsilon}$ for $\epsilon \in (0, \frac{1}{2})$ (the algorithm provided by [Schnorr 1987] would indeed run in subexponential time $2^{N^{1-\epsilon}}$). In practice, many works rely on the security of LWE when γ is superpolynomial ([Boneh et al. 2014] uses for example the above assumption that $\gamma = 2^{N^\epsilon}$ for some $\epsilon \in (0, \frac{1}{2})$), and we will follow this same path here. The following theorem (which is a corollary of [Peikert 2010, Theorem 3.1]), allows us to turn a continuous Gaussian sampling into a discrete Gaussian.

COROLLARY B.4 (FROM CONTINUOUS GAUSSIAN TO DISCRETE GAUSSIAN). *Let $\lambda, q \in \mathbb{N}$, $\alpha \in (0, 1)$. If e_c is sampled according to $\mathcal{D}_{\alpha q}$ and $e \leftarrow e_c + \mathcal{D}_{\mathbb{Z}-e_c, \omega(\sqrt{\log \lambda})}$, then the marginal distribution of e is within negligible statistical distance $\Delta = 1/(\exp(\pi\omega(\sqrt{\log \lambda})^2 - \ln(2\lambda)) - 1) = \text{negl}(\lambda)$ of $\mathcal{D}_{\mathbb{Z}, s}$ with $s := \sqrt{(\alpha q)^2 + \omega(\sqrt{\log \lambda})^2}$. Moreover, if there exists $x \in \mathbb{Z}_q$ (for example $x = \mathbf{a}^T \mathbf{s}$ for some $\mathbf{s} \in \mathbb{Z}_q^N$ and $\mathbf{a} \in \mathbb{Z}_q^N$) such that e_c is distributed according to $x + \mathcal{D}_{\alpha q} \bmod q$, then the statistical distance between the distribution $e \bmod q$ and $x + \mathcal{D}_{\mathbb{Z}, s} \bmod q$ is $\Delta = \text{negl}(\lambda)$. Finally, if e_c is uniformly sampled over $[0, q)$, the marginal distribution of e is uniform over \mathbb{Z}_q .*

PROOF. The first part of this corollary is a direct application of [Peikert 2010, Theorem 3.1]: We define $\mathbf{c}_1 = \mathbf{0}$, $\Lambda_1 = \mathbb{Z}$, $\Sigma_2 = (\alpha q)^2 \mathbf{I}_1$, $\Sigma_1 = \omega(\sqrt{\log \lambda})^2 \mathbf{I}_1$ with $\sqrt{\Sigma_1} = \omega(\sqrt{\log \lambda}) \geq \eta_\varepsilon(\mathbb{Z})$, where ε is a negligible function of N (the last inequality comes from [Peikert 2010, Lemma 2.5]). When $\mathbf{x}_2 (= e_c)$ is chosen according to a continuous Gaussian, the marginal distribution of e is within statistical distance $8\varepsilon = \text{negl}(\lambda)$ of $\mathcal{D}_{\mathbb{Z}, \sqrt{\Sigma}}$ (the constant can actually be improved in this specific case), with $\Sigma = \Sigma_1 + \Sigma_2 = (\alpha q)^2 \mathbf{I}_1 + \omega(\sqrt{\log \lambda})^2 \mathbf{I}_1 = s^2 \mathbf{I}_1$, which concludes the first part of the proof.

To see that the equality also holds when $e_c \leftarrow x + \mathcal{D}_{\alpha q} \bmod q$ for some $x \in \mathbb{Z}$, we remark that for any $\bar{e} \in \mathbb{Z}_q$ and for any $e_c = x + e'_c \in \mathbb{R}$:

$$p := \Pr \left[\bar{e} = e \bmod q \mid e \leftarrow e_c + \mathcal{D}_{\mathbb{Z}-e_c, \omega(\sqrt{\log \lambda})} \right] \quad (66)$$

$$= \Pr \left[\bar{e} = e \mid e \leftarrow e_c + \mathcal{D}_{\mathbb{Z}-e_c, \omega(\sqrt{\log \lambda})} \bmod q \right] \quad (67)$$

$$= \Pr \left[\bar{e} = e \mid e \leftarrow x + e'_c + \mathcal{D}_{\mathbb{Z}-e'_c, \omega(\sqrt{\log \lambda})} \bmod q \right], \quad (68)$$

where the last equality comes from $\mathbb{Z} - e_c = \mathbb{Z} - (e'_c \bmod q)$. But we already know that $e'_c + \mathcal{D}_{\mathbb{Z}-e'_c, \omega(\sqrt{\log \lambda})}$ is statistically close to $\mathcal{D}_{\mathbb{Z}, s}$ from the first part of the corollary, which concludes this part of the proof.

Now, let us assume that e_c is sampled uniformly at random over $[0, q)$. Because e_c can be uniquely decomposed into $e_c = e_{c,1} + e_{c,2}$ where $e_{c,1} \in \{0, \dots, q-1\}$ and $e_{c,2} \in [0, 1)$, and because $\mathbb{Z} - e_c = \mathbb{Z}^n - e_{c,2}$ we have

$$\begin{aligned} & \Pr \left[\bar{e} = e \bmod q \mid e_c \leftarrow [0, q), e \leftarrow e_c + \mathcal{D}_{\mathbb{Z}-e_c, \omega(\sqrt{\log \lambda})} \right] \\ &= \frac{1}{q} \int_0^1 de_{c,2} \sum_{e_{c,1}=0}^{q-1} \sum_{\substack{\hat{e} \in \mathbb{Z} \\ \hat{e} \equiv e_{c,1} \pmod{q}}} \frac{\rho_{\omega(\sqrt{\log \lambda})}(\hat{e} - e_{c,1} - e_{c,2})}{\rho_{\omega(\sqrt{\log \lambda})}(\mathbb{Z} - e_{c,2})}. \end{aligned}$$

Similarly, because any integer $\hat{e} \in \mathbb{Z}$ can be decomposed uniquely into $\hat{e} = e - e_{c,1}$ where $e_{c,1} \in \{0, \dots, q-1\}$, and $e = \bar{e} \bmod q$, we can merge the two sums into a single sum over $\hat{e} \in \mathbb{Z}$, replace the $e - e_{c,1}$ with \hat{e} , and use the fact that $\sum_{\hat{e} \in \mathbb{Z}} \rho_{\omega}(\hat{e} - e_{c,2}) = \rho_{\omega}(\mathbb{Z} - e_{c,2})$ to conclude that the probability is equal to $1/q$: it corresponds to a uniform sampling over \mathbb{Z}_q . \square

While in the usual LWE problem, \mathbf{s} is sampled uniformly at random over \mathbb{Z}_q^N , it is also possible to sample \mathbf{s} according to a small Gaussian. Since this can be seen as a reformulation of the problem in its HNF, this new sampling method is actually at least as secure as the initial uniform sampling, and it appears to be more efficient. Moreover, the construction given in [Micciancio and Peikert 2012] will naturally be formulated in this form, so in this article, we will also sample \mathbf{s} according to a small Gaussian (however, one can easily come back to the initial formulation as we will see later).

LEMMA B.5 (NORMAL LWE PROBLEM [APPLEBAUM ET AL. 2009b, LEMMA 2]). *Let $q = p^k$ be a prime power. There is a deterministic polynomial-time transformation T that, for arbitrary $\mathbf{s} \in \mathbb{Z}_q^N$ and error distribution χ , maps $A_{\mathbf{s}, \chi}$ to $A_{\tilde{\mathbf{s}}, \chi}$ where $\tilde{\mathbf{s}} \leftarrow \chi^N$, and maps $U(\mathbb{Z}_q^N \times \mathbb{Z}_q)$ to itself.*

The idea of the proof given in [Applebaum et al. 2009b, Lemma 2] (following [Micciancio and Regev 2009], see also [Peikert 2016, p. 23]) is to first obtain and select enough samples $(\tilde{\mathbf{A}}, \tilde{\mathbf{y}} := \tilde{\mathbf{A}}\mathbf{s} + \tilde{\mathbf{s}})$, $\tilde{\mathbf{s}}$ being sampled according to χ , to ensure that $\tilde{\mathbf{A}} \in \mathbb{Z}_q^{N \times N}$ is invertible. Then, any new sample $(\mathbf{a}, y := \langle \mathbf{a}, \mathbf{s} \rangle + e)$ can be updated into (\mathbf{a}', b') where $\mathbf{a}' := -(\tilde{\mathbf{A}}^T)^{-1}\mathbf{a}$ and $b' := y + \langle \mathbf{a}', \tilde{\mathbf{y}} \rangle = \langle \mathbf{a}', \tilde{\mathbf{s}} \rangle + e$.

B.3 The [Micciancio and Peikert 2012] Construction

We give here more details on the construction introduced in [Micciancio and Peikert 2012]. Note that [Micciancio and Peikert 2012] uses a left-style multiplication $\mathbf{s}^T \mathbf{A}^T$, here we will prefer a right-style multiplication $\mathbf{A}\mathbf{s}$ for consistency. In this article, we will also focus on the (more efficient) computationally-secure construction presented in [Micciancio and Peikert 2012] (we also require the modulus $q := 2^k$ to be a power of 2), but the same method should extend to other constructions with even q . We give in Definition B.6 the construction we will use, and we explain after the intuition behind it.

Definition B.6 ([Micciancio and Peikert 2012]). Let $\lambda \in \mathbb{N}$ be a security parameter, and $\mathcal{P}_0 = (k, N, \alpha, r_{\max})$ with $(k, N) \in \mathbb{N}^2$, $\alpha \in (0, 1)$ and $r_{\max} \in \mathbb{R}$, be some parameters that can depend on λ . We define $M := N(1 + k)$, $q := 2^k$, $\mathbf{X}_g := \{(\mathbf{s}, \mathbf{e}) \in \mathbb{Z}_q^N \times \mathbb{Z}_q^M \mid \|\tilde{\mathbf{e}}\|_2 \leq r_{\max}\}$, $\mathbf{g} := [1 \ 2 \ 4 \ \dots \ 2^{k-1}]^T \in \mathbb{Z}_q^k$, and the gadget matrix \mathbf{G} as

$$\mathbf{G} := \mathbf{I}_n \otimes \mathbf{g} = \begin{bmatrix} \vdots & & & & \\ & \mathbf{g} & & & \\ \vdots & & & & \\ & & \ddots & & \\ & & & \mathbf{g} & \\ & & & & \vdots \end{bmatrix} \in \mathbb{Z}_q^{Nk \times N}. \quad (69)$$

We define now in Figure 7 a procedure to sample a public matrix and its trapdoor $(\mathbf{A}, \mathbf{R}) \leftarrow \text{MP.Gen}_{\mathcal{P}_0}(1^\lambda)$, and for any $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}_g$, we define $\mathbf{y} := \mathbf{g}_\mathbf{A}(\mathbf{s}, \mathbf{e})$ and its inversion procedure $(\tilde{\mathbf{s}}, \tilde{\mathbf{e}}) := \text{MP.Invert}(\mathbf{R}, \mathbf{A}, \mathbf{y})$.

The idea of the construction given in Definition B.6 is to use a gadget matrix \mathbf{G} which is easy to invert even in the presence of noise, and then to hide this matrix inside a random looking matrix \mathbf{A} . \mathbf{G} is easy to invert because \mathbf{G} basically encodes all bits of the binary representation of each component of \mathbf{s} in a different component (where a 1 is encoded by $q/2 + \text{noise}$, and 0 is encoded by $0 + \text{noise}$): the inversion of \mathbf{G} is doable by a rounding operation, starting from the least significant bits of the components of \mathbf{s} . Then, as we will see, \mathbf{R} can be used to invert $\mathbf{A}\mathbf{s} + \mathbf{e}$: with \mathbf{R} we can obtain a vector of the form $\mathbf{G}\mathbf{s} + \mathbf{e}'$ (\mathbf{e}' is small if \mathbf{R} has sufficiently small singular values), and then since \mathbf{G} is easy to invert we can obtain \mathbf{s} easily. We formalize these statements now.

LEMMA B.7 ([MICCIANCIO AND PEIKERT 2012]). *If $\text{LWE}_{q, D_{\mathbb{Z}, \alpha q}}$ is hard and if $Nk = \text{poly}(\lambda)$, then the matrix \mathbf{A} obtained via MP.Gen is indistinguishable from a uniform random matrix.*

PROOF. For completeness, we sketch the proof given in [Micciancio and Peikert 2012]. Since \mathbf{G} is a fixed matrix, it is easy to subtract \mathbf{G} from \mathbf{A} and transpose the matrices: \mathbf{A} looks random iff $(\hat{\mathbf{A}}^T, \hat{\mathbf{A}}^T \mathbf{R}_2^T + \mathbf{R}_1^T)$ looks random. But this is nearly an exact LWE instance in its normal form

$\text{MP.Gen}_{\mathcal{P}_0}(1^\lambda)$ <hr/> 1: $\hat{\mathbf{A}} \xleftarrow{\$} \mathbb{Z}_q^{N \times N}$ 2: $\mathcal{R} = [\mathbf{R}_1 \mid \mathbf{R}_2] \leftarrow \mathcal{D}_{\mathbb{Z}, \alpha q}^{Nk \times 2N}$ 3: $\mathbf{A} := \left[\frac{\hat{\mathbf{A}}}{\mathbf{G} - \mathbf{R}_2 \hat{\mathbf{A}} - \mathbf{R}_1} \right] \in \mathbb{Z}_q^{M \times N}$ 4: return (\mathbf{A}, \mathbf{R})	$\text{InvertSmallGadget}_{\mathcal{P}_0}(\mathbf{y} = [y_0 \ \dots y_{k-1}]^T)$ <hr/> 1: / Return $s \in \mathbb{Z}_q$ such that $\mathbf{y} = \mathbf{g}s + \mathbf{e}$ 2: $s := 0$ 3: for $i = k-1, \dots, 0$ do 4: if $y_i - 2^i s \notin \left[-\frac{q}{4}, \frac{q}{4}\right) \bmod q$ 5: $s := s + 2^{k-1-i}$ fi endfor 6: return s
$g_{\mathbf{A}}(\mathbf{s}, \mathbf{e})$ <hr/> 1: return $\mathbf{A}\mathbf{s} + \mathbf{e}$	$\text{InvertGadget}_{\mathcal{P}_0}(\mathbf{y} = [\mathbf{y}_1^T \ \dots \mathbf{y}_N^T]^T)$ <hr/> 1: / Return $\mathbf{s} \in \mathbb{Z}_q^n$ such that $\mathbf{y} = \mathbf{G}\mathbf{s} + \mathbf{e}$ 2: for $i = 1, \dots, N$ do 3: $s_i := \text{InvertSmallGadget}_{\mathcal{P}_0}(\mathbf{y}_i)$ 4: endfor 5: $\mathbf{s} := [s_1 \ \dots \ s_n]^T$ 6: return \mathbf{s}
$\text{MP.Invert}_{\mathcal{P}_0}(\mathbf{R} := [\mathbf{R}_1 \mid \mathbf{R}_2], \mathbf{A}, \mathbf{y})$ <hr/> 1: / Return $(\tilde{\mathbf{s}}, \tilde{\mathbf{e}}) \in \mathcal{X}_g$ s.t. $\mathbf{y} = g_{\mathbf{A}}(\mathbf{s}, \mathbf{e})$ 2: $\tilde{\mathbf{s}} := \text{InvertGadget}_{\mathcal{P}_0}([\mathbf{R}_2 \mid \mathbf{I}_{Nk}]\mathbf{y})$ 3: $\tilde{\mathbf{e}} := \mathbf{y} - \mathbf{A}\tilde{\mathbf{s}}$ 4: if $\left\ \begin{bmatrix} \tilde{\mathbf{s}} \\ \tilde{\mathbf{e}} \end{bmatrix} \right\ _2 > r_{\max}$ return \perp fi 5: return $(\tilde{\mathbf{s}}, \tilde{\mathbf{e}})$	

Fig. 7. Construction from [Micciancio and Peikert 2012].

(\mathbf{R}_2 is indeed sampled according to a small Gaussian). The only difference is that the $\hat{\mathbf{A}}$ samples are “reused” multiple times since \mathbf{R}_i^T are matrices and not vectors. However, as shown in [Peikert and Waters 2011, Lemma 6.2], an hybrid argument can be made (by gradually replacing each column with random elements) to prove that it is still hard to distinguish it from a random matrix if $\text{LWE}_{q, \mathcal{D}_{\mathbb{Z}, \alpha q}}$ is hard (since we obtain Nk hybrid games, we need $Nk = \text{poly}(\lambda)$). \square

Remark B.8. It is possible to easily translate the normal form into a more usual form in which \mathbf{s} is sampled uniformly at random: one can sample a random invertible matrix $\mathbf{A}_r \in \mathbb{Z}_q^{N \times N}$, and define $\mathbf{A}' := \begin{bmatrix} \mathbf{I} \\ \mathbf{A} \end{bmatrix} \mathbf{A}_r$.

LEMMA B.9. Let $\hat{\mathbf{A}} \in \mathbb{Z}_q^{N \times N}$, $(\mathbf{R}_1, \mathbf{R}_2) \in (\mathbb{Z}^{Nk \times N})^2$, $\mathbf{A} := \left[\frac{\hat{\mathbf{A}}}{\mathbf{G} - \mathbf{R}_2 \hat{\mathbf{A}} - \mathbf{R}_1} \right] \in \mathbb{Z}_q^{M \times N}$. If the highest singular value $\sigma_{\max}(\mathbf{R})$ of \mathbf{R} is such that $\sqrt{\sigma_{\max}(\mathbf{R})^2 + 1} < \frac{q}{4r_{\max}}$, then $g_{\mathbf{A}} : \mathcal{X}_g \rightarrow \mathbb{Z}_q^M$ is injective and for all $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}_g$, $\text{MP.Invert}([\mathbf{R}_1 \mid \mathbf{R}_2], \mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e}) = (\mathbf{s}, \mathbf{e})$.

Moreover, if we denote by $C \approx \frac{1}{\sqrt{2\pi}}$ the universal constant defined in [Micciancio and Peikert 2012, Lemma 1.9], and if we have for the parameters \mathcal{P}_0 defined in Definition B.6:

$$\sqrt{\left(C \times \alpha q \times \sqrt{N}(\sqrt{k} + \sqrt{2} + 1)\right)^2 + 1} \leq \frac{q}{4r_{\max}} \quad (70)$$

then with overwhelming probability (on N) $\geq 1 - 2e^{-N}$, we have $\sqrt{\sigma_{\max}(\mathbf{R})^2 + 1} < \frac{q}{4r_{\max}}$ and, therefore, $g_{\mathbf{A}}$ is injective.

PROOF. This proof can be obtained by combining theorems from [Micciancio and Peikert 2012]. For completeness, we put the full proof of this lemma here. For the injectivity, we assume that there exist $(\mathbf{s}, \mathbf{s}') \in S$ and $(\mathbf{e}, \mathbf{e}') \in E$ such that $g_{\mathbf{A}}(\mathbf{s}, \mathbf{e}) = g_{\mathbf{A}}(\mathbf{s}', \mathbf{e}')$. So $\mathbf{A}(\mathbf{s} - \mathbf{s}') + (\mathbf{e} - \mathbf{e}') = \mathbf{0}$,

and it is enough now to prove that $\mathbf{s} = \mathbf{s}'$ to obtain $\mathbf{e} = \mathbf{e}'$. We multiply (on the left) this equation by $\begin{bmatrix} \mathbf{R}_2 & \mathbf{I}_{Nk} \end{bmatrix}$: we obtain $\mathbf{G}(\mathbf{s} - \mathbf{s}') - \mathbf{R}_1(\mathbf{s} - \mathbf{s}') + \begin{bmatrix} \mathbf{R}_2 & \mathbf{I}_{Nk} \end{bmatrix}(\mathbf{e} - \mathbf{e}') = \mathbf{0}$, i.e., if we define $\tilde{\mathbf{R}} := \begin{bmatrix} -\mathbf{R}_1 & \mathbf{R}_2 & \mathbf{I}_{Nk} \end{bmatrix}$, $\tilde{\mathbf{e}} := \tilde{\mathbf{R}} \begin{bmatrix} \mathbf{s} - \mathbf{s}' \\ \mathbf{e} - \mathbf{e}' \end{bmatrix}$ and $\tilde{\mathbf{s}} = \mathbf{s} - \mathbf{s}'$, we have $\mathbf{G}\tilde{\mathbf{s}} + \tilde{\mathbf{e}} = \mathbf{0}$. First, we remark that

$$\begin{aligned}
 \sigma_{\max}(\tilde{\mathbf{R}}^T) &= \|\tilde{\mathbf{R}}^T\|_2 \\
 &= \max_{\mathbf{x}, \|\mathbf{x}\|_2=1} \left\| \begin{bmatrix} -\mathbf{R}_1^T \\ \mathbf{R}_2^T \\ \mathbf{I}_{Nk} \end{bmatrix} \mathbf{x} \right\|_2 \\
 &= \max_{\mathbf{x}, \|\mathbf{x}\|_2=1} \sqrt{\left\| \begin{bmatrix} -\mathbf{R}_1^T \\ \mathbf{R}_2^T \end{bmatrix} \mathbf{x} \right\|_2^2 + \|\mathbf{x}\|_2^2} \\
 &= \sqrt{\max_{\mathbf{x}, \|\mathbf{x}\|_2=1} \left\| \begin{bmatrix} \mathbf{R}_1^T \\ \mathbf{R}_2^T \end{bmatrix} \mathbf{x} \right\|_2^2 + 1} \\
 &= \sqrt{\max_{\mathbf{x}, \|\mathbf{x}\|_2=1} \|\mathbf{R}^T \mathbf{x}\|_2^2 + 1} \\
 &= \sqrt{\sigma_{\max}(\mathbf{R}^T)^2 + 1} \\
 &= \sqrt{\sigma_{\max}(\mathbf{R})^2 + 1}.
 \end{aligned} \tag{71}$$

We prove now that for all i , $\tilde{\mathbf{e}}[i] \in (-\frac{q}{2}, \frac{q}{2})$: If we denote by \mathbf{u}_i the vector such that $\mathbf{u}_i[i] = 1$ and for all $j \neq i$, $\mathbf{u}_i[j] = 0$, we have

$$|\tilde{\mathbf{e}}[i]| = |\mathbf{u}_i^T \tilde{\mathbf{e}}| = \left| \mathbf{u}_i^T \tilde{\mathbf{R}} \begin{bmatrix} \mathbf{s} - \mathbf{s}' \\ \mathbf{e} - \mathbf{e}' \end{bmatrix} \right| = \left\langle \tilde{\mathbf{R}}^T \mathbf{u}_i, \begin{bmatrix} \mathbf{s} - \mathbf{s}' \\ \mathbf{e} - \mathbf{e}' \end{bmatrix} \right\rangle. \tag{72}$$

Using the Cauchy-Schwarz inequality we get

$$|\tilde{\mathbf{e}}[i]| \leq \|\tilde{\mathbf{R}}^T \mathbf{u}_i\|_2 \left\| \begin{bmatrix} \mathbf{s} - \mathbf{s}' \\ \mathbf{e} - \mathbf{e}' \end{bmatrix} \right\|_2. \tag{73}$$

Using $\|\mathbf{A}\|_2 = \sigma_{\max}(\tilde{\mathbf{R}}^T)$, and the definition of \mathcal{X}_g (Definition B.6), then Equation (71) and the assumption on $\sigma_{\max}(\mathbf{R})$ we obtain

$$|\tilde{\mathbf{e}}[i]| \leq \sigma_{\max}(\tilde{\mathbf{R}}^T) \times (2r_{\max}) < \frac{q}{2}. \tag{74}$$

Now, if we write for all $i \in [N]$, $\tilde{\mathbf{s}}[i] = \sum_{j=0}^{k-1} 2^j \tilde{\mathbf{s}}[i]_j$, with for all j , $\tilde{\mathbf{s}}[i]_j \in \{0, 1\}$ (we also use the same notation for $\tilde{\mathbf{e}}$), we can prove, following the same path as the InvertSmallGadget algorithm that $\forall i, j$, $\tilde{\mathbf{s}}[i]_j = 0$. Let $i \in [n]$. If we consider only the line $l := (i-1)k + k$ of $\mathbf{G}\tilde{\mathbf{s}} + \tilde{\mathbf{e}} = \mathbf{0}$, we obtain $2^{k-1}\tilde{\mathbf{s}} + \tilde{\mathbf{e}}[l] = \frac{q}{2}\tilde{\mathbf{s}}[i]_0 + \tilde{\mathbf{e}}[i] \bmod q = \mathbf{0}$. Since $\tilde{\mathbf{e}}[l] \in (-\frac{q}{2}, \frac{q}{2})$, we can't have $\tilde{\mathbf{s}}[i]_0 = 1$, so $\tilde{\mathbf{s}}[i]_0 = 0$. We can then iterate the same process for $m = 1 \dots k-1$ with the line $l := (i-1)k + (k-m)$ to show that $\tilde{\mathbf{s}}[i]_m = 0$, i.e., $\tilde{\mathbf{s}} = \mathbf{0}$, which concludes the proof of the injectivity of g_A . Because this proof follows exactly the algorithm MP.Invert, it is easy to see using the same argument that this algorithm correctly inverts any $\mathbf{y} = \mathbf{A}\mathbf{s} + \mathbf{e}$ with $(\mathbf{s}, \mathbf{e}) \in \mathcal{X}_g$.

We now prove the second part of the theorem. Because \mathbf{R} is sampled according to a discrete Gaussian of parameter αq , so according to [Micciancio and Peikert 2012, Lemma 2.8], this distribution is 0-subgaussian, and therefore, we can apply [Micciancio and Peikert 2012, Lemma 2.9] with,

for example, $t = \sqrt{N/\pi}$ (we divide by π just to simplify the probability, and to transform the \leq into a $<$). So with probability $1 - 2 \exp(-\pi t^2) = 1 - 2e^{-N}$,

$$\sigma_{\max}(\mathbf{R}) \leq C \times \alpha q \times (\sqrt{Nk} + \sqrt{2N} + \sqrt{\frac{N}{\pi}}) < C \times \alpha q \times \sqrt{N}(\sqrt{k} + \sqrt{2} + 1).$$

To conclude the proof, we inject this equation inside Equation (71) and use Equation (70). □

Received 13 November 2023; revised 30 November 2024; accepted 15 February 2025