ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Βελτίωση Πολυπλοκότητας Επικοινωνίας Κβαντικών Κρυπτογραφικών Πρωτοκόλλων

Improving Communication Complexity of Quantum Cryptographic Protocols

Γεώργιος Ορέστης Χαρδούβελης
A.M. : 03115100

Επιβλέπων : Αριστείδης Παγουρτζής
Καθηγητής ΕΜΠ

Συνεπιβλέπων : Giulio Malavolta
Assistant Professor MPI-SP

Αθήνα
Ιούλιος 2021

ΕΘΝΙΚΟ ΜΕΤΣΟΒΙΟ ΠΟΛΥΤΕΧΝΕΙΟ
ΣΧΟΛΗ ΗΛΕΚΤΡΟΛΟΓΩΝ ΜΗΧΑΝΙΚΩΝ ΚΑΙ ΜΗΧΑΝΙΚΩΝ ΥΠΟΛΟΓΙΣΤΩΝ
ΤΟΜΕΑΣ ΤΕΧΝΟΛΟΓΙΑΣ ΠΛΗΡΟΦΟΡΙΚΗΣ ΚΑΙ ΥΠΟΛΟΓΙΣΤΩΝ

ΔΙΠΛΩΜΑΤΙΚΗ ΕΡΓΑΣΙΑ

Βελτίωση Πολυπλοκότητας Επικοινωνίας Κβαντικών Κρυπτογραφικών Πρωτοκόλλων

Improving Communication Complexity of Quantum Cryptographic Protocols

Γεώργιος Ορέστης Χαρδούβελης
Α.Μ. : 03115100

Επιβλέπων : Αριστείδης Παγουρτζής
Καθηγητής ΕΜΠ

Συνεπιβλέπων : Giulio Malavolta
Assistant Professor MPI-SP

Τριμελής Επιτροπή Εξέτασης

(Υπογραφή)	(Υπογραφή)	(Υπογραφή)
.....	.....	.....
Αριστείδης Παγουρτζής	Δημήτριος Φωτάκης	Νικόλαος Παπασπύρου
Καθηγητής	Αν. Καθηγητής	Καθηγητής
ΕΜΠ	ΕΜΠ	ΕΜΠ

Ημερομηνία Εξέτασης:
14 Ιουλίου 2021

(Υπογραφή)

.....

Γεώργιος Ορέστης Χαρδούβελης
Διπλωματούχος Ηλεκτρολόγος Μηχανικός και Μηχανικός Υπολογιστών Ε.Μ.Π.

Περίληψη

Στην παρούσα διπλωματική εργασία, μελετάμε την πολυπλοκότητα επικοινωνίας των κβαντικών σχημάτων πλήρως ομομορφικής κρυπτογράφησης (QFHE) και των συστημάτων μηδενικής γνώσης (ZK) για την κλάση QMA (το κβαντικό ανάλογο της NP). Θεωρώντας την κβαντική δυσκολία του learning with errors προβλήματος (LWE) για τα πρώτα, και την οιονεί-πολυωνυμική (quasi-polynomial) κβαντική δυσκολία του LWE για τα δεύτερα, έχουμε τα παρακάτω αποτελέσματα:

- Κβαντικό Σχήμα Πλήρως Ομομορφικής Κρυπτογράφησης Ρυθμού-1, που επιτρέπει στην Alice να στείλει στον Bob το κβαντικό της μήνυμα $|\psi\rangle$ κρυπτογραφημένο, ώστε ο Bob, έχοντας ένα κύκλωμα C, να μπορεί να υπολογίσει το $C(|\psi\rangle)$ χωρίς να χρειάζεται να αποκτυπτογραφήσει το μήνυμα. Επιτυγχάνουμε πολυπλοκότητα

  $$(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$$

  που είναι σχεδόν βέλτιστη.

- Πρωτόκολλό στατιστικής Μηδενικής Γνώσης 4 γύρων για την κλάση QMA στο απλό μοντέλο. Αυτό είναι το πρώτο πρωτόκολλο που επιτυγχάνει *στατιστική* μηδενική γνώση σε σταθερό αριθμό γύρων για την κλάση QMA.

- Πρωτόκολλο υπολογιστικής (αντ. στατιστικής) Μηδενικής Γνώσης 2 γύρων στο χρονικό μοντέλο, θεωρώντας επιπλέον την ύπαρξη μετακβαντικών μη-παραλληλοποιήσιμων συναρτήσεων (αντ. time-lock puzzles).

Όλα τα παραπάνω πρωτόκολλα επιτυγχάνουν την βέλτιστη πολυπλοκότητα επικοινωνίας των αντίστοιχων NP πρωτοκόλλων με ασφάλεια έναντι σε κλασσικούς αντιπάλους.

Λέξεις Κλειδιά— κβαντική κρυπτογραφία, πλήρως ομομορφική κρυπτογράφηση, μηδενική γνώση, πολυπλοκότητα επικοινωνίας, LWE

Abstract

In this diploma dissertation, we study the communication complexity of Quantum Fully Homomorphic Encryption (QFHE) schemes and Zero-Knowledge (ZK) for QMA (the quantum analogue of NP). Assuming the quantum hardness of the learning with errors problem (LWE) for the first and the quantum quasi-polynomial hardness for the latter, we obtain the following results:

- Rate-1 Quantum Fully Homomorphic Encryption Scheme, which allows Alice to send an encrypted version of her quantum input $|\psi\rangle$ to Bob, such that he (holding a circuit C) can compute $C(|\psi\rangle)$ without first decrypting it. We achieve communication complexity
$$(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$$
which is nearly optimal.
- 4-Round statistical Zero-Knowledge Arguments for QMA in the plain model, additionally assuming the existence of quantum fully homomorphic encryption. This is the first protocol for constant-round *statistical* zero-knowledge arguments for QMA.
- 2-Round computational (statistical, resp.) Zero-Knowledge arguments for QMA in the timing model, additionally assuming the existence of post-quantum non-parallelizing functions (time-lock puzzles, resp.).

All of these protocols match the best communication complexity known for the corresponding protocols for NP with security against classical adversaries.

Keywords— quantum cryptography, fully homomorphic encryption, zero-knowledge, LWE, communication complexity

Ευχαριστίες

Καταρχάς θα ήθελα να ευχαριστήσω τον επιβλέποντα καθηγητή της παρούσας διπλωματικής εργασίας, Αριστείδη Παγουρτζή, για την άψογη συνεργασία μας και τις πολύτιμες συμβουλές του.

Η εργασία αυτή πραγματοποιήθηκε στα πλαίσια υποτροφίας στο Max Planck Institute for Security and Privacy. Οφείλω ένα μεγάλο ευχαριστώ στον υπεύθυνό μου Dr. Giulio Malavolta για την ανεκτίμητη καθοδήγηση του και τις γνώσεις που με βοήθησε να αποκτήσω.

Τέλος, θα ήθελα να ευχαριστήσω τους φίλους και την οικογένεια μου για την στήριξη τους όλα αυτά τα χρόνια.

Contents

List of Figures

Εκτεταμένη Ελληνική Περίληψη

1 Εισαγωγή

Η κρυπτογραφία διαδραματίζει πλέον μείζονα ρόλο στην καθημερινή μας ζωή, εφοδιάζοντας μας με εργαλεία που καθιστούν δυνατή την ασφαλή επικοινωνία μεταξύ δύο ή παραπάνω ατόμων. Πέρα από τα ευρύτερα γνωστά κρυπτογραφικά πρωτόκολλα, όπως η κρυπτογράφηση δημοσίου κλειδιού που επιτρέπει την ασφαλή επικοινωνία μέσω ενός δημοσίου καναλιού, υπάρχουν και περισσότερο περίπλοκα πρωτόκολλα που επιλύουν δυσκολότερα προβλήματα όπως:

- Η Πλήρως Ομομορφική Κρυπτογράφηση (FHE), που επιτρέπει τον υπολογισμό συναρτήσεων με είσοδο κρυπτογραφημένα δεδομένα.
- Οι Αποδείξεις Μηδενικής Γνώσεις (ZK), που αποδεικνύουν την εγκυρότητα μιας πρότασης, χωρίς να φανερώνουν κάποια επιπλέον πληροφορία.

Τα παραπάνω εργαλεία πρέπει να είναι και αποδοτικά ώστε να μπορούν να φανούν χρήσιμα και σε πραγματικά συστήματα. Έτσι οι ερευνητές προσπαθούν να ελαχιστοποιήσουν την πολυπλοκότητα επικοινωνίας τους. Η πολυπλοκότητα επικοινωνίας καθορίζεται τόσο από το μέγεθος όσο και από τον αριθμό των μηνυμάτων που ανταλλάσουν τα εμπλεκόμενα μέρη.

Εμβαθύνοντας στα προαναφερθέντα πρωτόκολλα, ένα σύστημα πλήρως ομομορφικής κρυπτογράφησης επιτρέπει στο ένα μέρος να στείλει το κρυπτογραφημένο του μήνυμα m κάτω από ένα δημόσιο κλειδί έτσι ώστε το άλλο να μπορεί ύστερα, έχοντας ένα κύκλωμα C, να υπολογίσει και να στείλει

$$\mathsf{Enc}(m) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(C(m)),$$

χωρίς να μάθει κάποια νέα πληροφορία για το μήνυμα m. Οι υπολογισμοί πάνω σε κρυπτογραφημένα δεδομένα έχουν πολλαπλές εφαρμογές, όπως στην περίπτωση όπου ένας υπολογιστικά αδύναμος client θέλει να ανεβάσει τα δεδομένα του σε έναν ισχυρότερο server που μπορεί να τρέξει περίπλοκα κυκλώματα, διατηρώντας την ιδιωτικότητα του. Είναι βέβαια σημαντικό να εξασφαλίσουμε πως η πολυπλοκότητα επικοινωνίας που εισάγει το FHE πρωτόκολλο δεν αναιρεί την βελτίωση της απόδοσης που επιφέρει η ανάθεση στον server. Πρόσφατα αποδείχθηκε πως υπάρχουν FHE πρωτόκολλα όπου η πολυπλοκότητα επικοινωνίας πλησιάζει εκείνη των ανασφαλή πρωτοκόλλων [BDGM19] (εκείνα όπου το πρώτο μέλος στέλνει το μήνυμα του μη κρυπτογραφημένο), θεωρώντας την δυσκολία του learning with errors (LWE) προβλήματος.

Σχετικά με τις αποδείξεις μηδενικής γνώσης, από την εισαγωγή τους το 1989 [GMR89] έχουν ιδιαίτερα σημαντική επίδραση στην κρυπτογραφία και τη θεωρητική πληροφορική γενικότερα. Έχουν μελετηθεί εκτενώς από τους κρυπτογράφους, με στόχο την κατανόηση των απαραίτητων υποθέσεων και την βελτιστοποίηση του αριθμού απαραίτητων γύρων. Αποδεικνύεται πως θεωρώντας standard computational assumptions, οποιαδήποτε NP πρόταση μπορεί να αποδειχθεί μέσα σε το πολύ τέσσερις γύρους αλληλεπίδρασης [GMW86, GK96].

Εντούτοις, σε αντίθεση με τα κλασσικά πρωτόκολλα, γνωρίζουμε πολύ λιγότερα για τα αντίστοιχα κβαντικά. Τα κβαντικά πλήρως ομομορφικά πρωτόκολλα επιτρέπουν αντίστοιχα κβαντικές πράξεις όπως

$$\mathsf{Enc}(|\psi\rangle) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(C(|\psi\rangle)),$$

όπου $|\psi\rangle$ είναι μια αυθαίρετη κβαντική κατάσταση και C κάποιος μοναδιακός (unitary) τελεστής. Παρόλο που αυτό το πρόβλημα έχει μελετηθεί λιγότερο από το αντίστοιχο κλασσικό, πιστεύουμε πως είναι ακόμη πιο σημαντικό δεδομένης της υπολογιστικής υπεροχής των κβαντικών server σε σχέση με κάποιον client. Ακόμη και μελλοντικά όπου οι χρήστες θα έχουν πρόσβαση σε κβαντικούς υπολογιστές, το πιθανότερο είναι πως οι περισσότερο δύσκολοι υπολογισμοί θα είναι πραγματοποιήσιμοι μόνο από ισχυρούς server. Υπάρχουν ορισμένες κατασκευές κβαντικών FHE [BJ15], ακόμη και με πλήρως κλασσικό client [Mah18a]. Παρόλα αυτά, το πρόβλημα κβαντικών FHE συστημάτων με πολυπλοκότητα επικοινωνίας ανεξάρτητη του μεγέθους του κυκλώματος παραμένει ανοικτό.

Ταυτόχρονα, όσον αφορά στις αποδείξεις μηδενικής γνώσης, πρωτόκολλα για την κλάση QMA έχουν μόλις προσφάτως κατασκευαστεί [BJSW16] και το βέλτιστο μέχρι στιγμής αποτέλεσμα (σχετικά με την πολυπλοκότητα επικοινωνίας) είναι από τους Bitansky και Shmueli [BS20], όπου κατασκεύασαν ένα επιχείρημα μηδενικής γνώσης (ZK argument, δηλ. με υπολογιστική ορθότητα) με σταθερό αριθμό γύρων (>4).

Έτσι, θέτουμε τα παρακάτω ερωτήματα:

Μπορούμε να κατασκευάσουμε κβαντικό FHE με ελάχιστη πολυπλοκότητα επικοινωνίας; Για να αποδείξουμε προτάσεις της κλάσης QMA οφείλουμε εκ φύσεως να προσθέσουμε επιπλέον γύρους αλληλεπίδρασης;

Σε αυτή την εργασία μελετάμε τα παραπάνω προβλήματα και κατασκευάζουμε κβαντικά πρωτόκολλα με πολυπλοκότητα επικοινωνίας αντίστοιχη με εκείνη των κλασσικών πρωτοκόλλων.

2 Κβαντικό Πλήρως Ομομορφικό Σχήμα Κρυπτογράφησης Ρυθμού-1

Αρχικά μελετάμε το Κβαντικό Πλήρως Ομομορφικό Σχήμα Κρυπτογράφησης Ρυθμού-1.

Γιατί αποτελεί μη-τετριμμένο πρόβλημα; Πριν εμβαθύνουμε στην κατασκευή μας, είναι σημαντικό να κατανοήσουμε γιατί τα υπάρχοντα πρωτόκολλα αποτυγχάνουν και εμφανίζουν μεγαλύτερο ρυθμό. Ως ρυθμός γενικά ορίζεται το κλάσμα του μεγέθους του αποτελέσματος των υπολογισμών σε μη κρυπτογραφημένη πληροφορία ως προς το μέγεθος του αποτελέσματος των υπολογισμών σε κρυπτογραφημένη πληροφορία. Στα πρωτόκολλα όπως τα [Mah18a, Bra18], ένα κρυπτοκείμενο που αποτελεί κρυπτογράφηση μιας ℓ-qubit

κβαντικής κατάστασης $|\psi\rangle$ είναι την μορφής

$$\mathsf{QOTP}((x_1, z_1, \ldots, x_\ell, z_\ell), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, (x_1, z_1, \ldots, x_\ell, z_\ell))$$

όπου το QOTP (Quantum One-Time Pad) εφαρμόζεται ξεχωριστά σε κάθε qubit και η συμβολοσειρά $\mathsf{otk} = (x_1, z_1, \ldots, x_\ell, z_\ell)$ είναι κρυπτογραφημένη bit ανά bit. Εύκολα μπορεί κανείς να παρατηρήσει πως το συγκεκριμενο κρυπτοσύστημα έχει ρυθμό αντίστροφα πολυωνυμικό, εξαιτίας του κλασσικού FHE σχήματος.

Μία προφανής λύση θα ήταν να υιοθετήσουμε μια υβριδική προσέγγιση και να δειγματοληπτούμε το QOTP κλειδί χρησιμοποιώντας έναν ψευδοτυχαίο γεννήτορα (PRG). Πιο συγκεκριμένα, θα μπορούσαμε να βελτιώσουμε τον ρυθμό υπολογίζοντας

$$\mathsf{QOTP}(\mathsf{PRG}(\mathsf{seed}), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{seed})$$

για κάποιο ομοιόμορφα δειγματοληπτημένο $\mathsf{seed} \leftarrow\!\!\$\, \{0,1\}^\lambda$. Τότε θα μπορούμε ακόμη να υπολογίσουμε μια συνάρτηση με κρυπτογραφημένη είσοδο, αφού υπάρχει η δυνατότητα να μετατρέψουμε τα κρυπτοκείμενα στην αρχική τους μορφή, υπολογίζοντας ομομορφικά τον ψευδοτυχαίο γεννήτορα.

Παρόλο που αυτή η προσέγγιση λειτουργεί για κρυπτοκείμενα που έχουν μόλις κρυπτογραφηθεί, δεν ισχύει το ίδιο ύστερα από ομομορφική εφαρμογή συναρτήσεων· ανάλογα με την πύλη που θα εφαρμοστεί στην κβαντική είσοδο, το QOTP κλειδί otk αλλάζει σε κάποια διαφορετική συμβολοσειρά otk'. Αν και υπάρχει τρόπος να ανανεώνουμε το κλασσικό μέρος του κρυπτοκειμένου με κάθε υπολογισμό [Mah18a], δεν συμβαδίζει με την υβριδική προσέγγιση που εξετάζουμε. Αυτό συμβαίνει αφού το τροποποιημένο otk' πιθανώς δεν ανήκει στο σύνολο αφίξεως του ψευδοτυχαίου γεννήτορα PRG, δηλαδή μπορεί να μην υπάρχει συμβολοσειρά seed' τέτοια ώστε $\mathsf{PRG}(\mathsf{seed}') = \mathsf{otk}'$. Συνεπώς παρατηρούμε ότι δεν μπορούμε να αποβάλλουμε την κλασσική κρυπτογράφηση $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk})$. Ακόμη και στην ιδανική περίπτωση όπου το κλασσικό FHE έχει βέλτιστο ρυθμό, εφόσον απαιτούνται τουλάχιστον δύο κλασσικά bit για να κρυπτογραφήσουν ένα qubit, ακόμη θα έχουμε πολυπλοκότητα μεγαλύτερη της επιθυμητής και φτάνουμε σε αδιέξοδο.

Spooky Υπολογισμοί. Για την λύση μας αξιοποιούμε την δομή μιας ιδιαίτερης περίπτωσης FHE συστήματος ώστε να μετατρέψουμε το κβαντικό FHE κρυπτοκείμενο σε υβριδική κατάσταση ρυθμού-1. Παρατηρούμε ότι σε οριμένα νέα FHE συστήματα [BDGM19] ομαδοποιούν k κλασσικά bit μορφής $c = (\mathbf{c}_0, c_1, \ldots, c_k) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^k$, για κάποιο modulus q και $n = \mathsf{poly}(\lambda)$. Η ενδιαφέρουσα για εμάς ιδιότητα είναι πως τα τελευταία k bits των κρυπτοκειμένων συσχετίζονται μη-τοπικά με το ιδιωτικό κλειδί sk. Πιο συγκεκριμένα, ο αλγόριθμος αποκρυπτογράφησης ανακτά το αρχικό μήνυμα υπολογίζοντας

$$\mathsf{Dec}(\mathsf{sk}, c) = F(\mathsf{sk}, \mathbf{c}_0) \oplus (c_1, \ldots, c_k)$$

για κάποια συνάρτηση F, της οποίας ο ακριβώς ορισμός ξεφεύγει από τα πλαίσια της περίληψης μας. Αυτή η ιδιότητα την οποία ονομάζουμε *spooky αποκρυπτογράφηση*,[1] είναι κρίσιμη για την λύση μας.

Η λύση μας. Με βάση τα παραπάνω, μπορούμε να μετατρέψουμε κβαντικά κρυπτοκείμενα (ακόμη και ύστερα από ομομορφικούς υπολογισμούς) της μορφής $(\mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}'))$ σε ρυθμού-1 μορφή ακολουθώντας τα παρακάτω βήματα:

[1] Το όνομα είναι εμπνευσμένο από ένα παρόμοιο φαινόμενο σε FHE συστήματα πολλαπλών κλειδιών [DHRW16].

- Μετατρέπουμε το $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}')$ σε FHE κρυπτοκείμενο με spooky αποκρυπτογράφηση χρησιμοποιώντας την τεχνική bootstrapping (δηλ. τρέχοντας τον αλγόριθμο αποκρυπτογράφησης ομομορφικά).

- Το κρυπτοκείμενο που προκύπτει το συμβολίζουμε ως εξής:

$$c = (\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^{2\ell}.$$

- Επιστρέφουμε $\mathbf{c}_0$ και $\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle)$.

Εφόσον $|\mathbf{c}_0| = \mathsf{poly}(\lambda)$, το μέγεθος του συμπιεσμένου κρυπτοκειμένου είναι ℓ qubits συν $\mathsf{poly}(\lambda)$ bit κλασσικής πληροφορίας. Ο ρυθμός που προκύπτει είναι βέλτιστος δεδομένου ότι οποιοδήποτε κρυπτοσύστημα δημοσίου κλειδιού έχει κρυπτοκείμενα μεγέθους τουλάχιστον λ bit, οπότε ένας πρόσθετος όρος της παραμέτρου ασφαλείας είανι αναπόφευκτος. Στην περίπτωση μας έχουμε έναν μεγαλύτερο πρόσθετο όρο, ο οποίος όμως είναι ασυμπτωτικά αμελητέος.

Για να δούμε ξεκάθαρα γιατί αυτή η διαδικασία μας δίνει ένα κρυπτοκείμενο που μπορούμε να αποκρυπτογραφήσουμε, ανακατατάσσοντας την παραπάνω εξίσωση έχουμε

$$F(\mathsf{sk}, \mathbf{c}_0) = (x'_1, z'_1, \ldots, x'_\ell, z'_\ell) \oplus (c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z})$$

όπου είναι το κατάλληλο one-time key της κβαντικής κατάστασης

$$\begin{aligned}
&\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle \\
&= \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \bigotimes_{i\in[l]} \left(X^{x'_i} Z^{z'_i}\right) \cdot |\psi'\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{c_{i,x}\oplus x'_i} Z^{c_{i,z}\oplus z'_i}\right) \cdot |\psi'\rangle\,.
\end{aligned}$$

3 Κβαντικό Πλήρως Ομοιομορφικό Σχήμα Κρυπτογράφησης Ρυθμού-1 μέσω του Συμπαγούς Dual-GSW

Αν και η παραπάνω λύση επιτυγχάνει βέλτιστο ρυθμό, παρατηρούμε ότι εισάγει ένα επιπλέον ιδιωτικό κλειδί στο σύστημα. Έτσι κατά την μετατροπή από leveled (ικανό να αποτιμήσει κυκλώματα οριοθετημένου βάθους) σε πλήρως ομομορφικό σύστημα εισάγεται ένα επιπλέον assumption κυκλικότητας. Αντίθετα με την συνηθισμένη περίπτωση όπου χρησιμοποιούμε την τεχνική bootstrapping και θεωρούμε ασφαλή την παρουσία ενός ιδιωτικού κλειδιού (circularity assumption), κατά την μετατροπή από το ένα κρυπτοσύστημα στο άλλο δημιουργείται ένας κύκλος με δύο κλειδιά την ασφάλεια του οποίου πρέπει να θεωρήσουμε. Αν και οι δύο περιπτώσεις δεν είναι συγκρίσιμες, γεννιέται το ερώτημα αν μπορούμε να πετύχουμε ρυθμό-1 με circularity assumption ενός κλειδιού. Σε αυτό το μέρος της εργασίας αποδεικνύουμε πως γίνεται, κατασκευάζοντας μία συμπαγή (packed) παραλλαγή του dual-GSW συστήματος [Mah18a] και αποδεικνύουμε πως είναι "κβαντικά ικανό" (υποστηρίζει δηλαδή τον υπολογισμό κβαντικών κυκλωμάτων). Ύστερα, χρησιμοποιώντας έναν αλγόριθμο συρρίκνωσης [BDGM19], προκύπτει ένα κβαντικά ικανό κρυπτοσύστημα ρυθμού-1, με την

ίδια *spooky αποκρυπτογράφηση* με παραπάνω. Συνεπώς, με παρόμοιες τεχνικές κατασκευάζουμε ξανά ένα κβαντικό πλήρες ομομορφικό κρυπτοσύστημα ρυθμού-1.

Συμπαγές Dual-GSW σύστημα. Το συμπαγές Dual-GSW σύστημα είναι ουσιαστικά το δυϊκό του κρυπτοσυστήματος των Hiromasa et al. [HAO15]. Στο (μη συμπαγές) dual-GSW, το κρυπτκείμενο ενός μηνύματος μ είναι της μορφής

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mu\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times(m+1)\log q}$$

where $\mathbf{A}' \in \mathbb{Z}_q^{(m+1)\times n}$, $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+1)\log q}$ και $\mathsf{sk}\cdot\mathbf{A}' = 0$, με sk να είναι το ιδιωτικό κλειδί του συστήματος. Στο συμπαγές κρυπτοσύστημα κρυπτογραφούμε ℓ-bit μηνύματα, οπότε θεωρούμε ως μηνύματα διαγώνιους πίνακες $\mathbf{M} \in \{0,1\}^{\ell\times\ell}$ που περιέχουν ℓ bits, και ορίζουμε τα κρυπτοκείμενα να είναι της μορφής

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y}\cdot\mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times(m+\ell)\log q},$$

όπου $\mathbf{Y} \in \{0,1\}^{(m+\ell)\times(m+\ell)}$ είναι μια κωδικοποίηση του μηνύματος, $\mathbf{A}' \in \mathbb{Z}_q^{(m+\ell)\times n}$, και $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+\ell)\log q}$.

Για να μπορέσουμε να υπολογίσουμε μια NAND πύλη χωρίς να αλλάξουμε την δομή του κρυπτοσυστήματος, και έτσι να διατηρήσουμε τις ομομορφικές ιδιότητες του, επιλέγουμε μια κωδικοποίηση όπου υποστηρίζει πρόσθεση και πολλαπλασιασμό στοιχείο-στοιχείο (point-wise) ενώ ταυτόχρονα ισχύει ότι $\mathbf{Y}\cdot\mathbf{A}' = 0$, ώστε να απαλειφθούν οι επιπλέον μικτοί όροι του πολλαπλασιασμού.

Για να το πετύχουμε αυτό, ορίζουμε το ιδιωτικό κλειδί ως $\left[\ \mathbf{E}_{sk}\ \middle|\ \mathbf{I}_l\ \right]$ με $\mathbf{E}_{sk} \in \{0,1\}^{\ell\times m}$ και ο πίνακας $\mathbf{Y}$ ορίζεται ως $\left[\dfrac{\mathbf{0}}{\mathbf{M}\cdot\mathsf{sk}}\right]$. Είναι σημαντικό να αναφέρουμε πως για να πετύχουμε την παραπάνω δομή του πίνακα $\mathbf{Y}$, ο αλγόριθμος παραγωγής κλειδιών του κρυπτοσυστήματος πρέπει να υπολογίζει και κρυπτογραφήσεις των $\mathbf{P}_i$ για $i \in \{0,\ldots,\ell\}$, όπου $\mathbf{P}_i$ είναι ένας διαγώνιος πίνακας με 1 στην θέση (i,i) και 0 οπουδήποτε αλλού. Έτσι, ο αλγόριθμος κρυπτογράφησης μπορεί να προσθέσει όλες τις κρυπτογραφήσεις που αντιστοιχούν στο μήνυμα που κρυπτογραφεί και να προσθέσει εκ νέου τυχαιότητα.

Για να γίνει κατανοητό γιατί το παρόν κρυπτοσύστημα είναι κβαντικά ικανό, παρατηρούμε πως αθροίζοντας τις στήλες $(m+i)\log q$ για $i \in \{1,\ldots,\ell\}$ στο κρυπτοκείμενό μας, έχουμε ως αποτέλεσμα το

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \left[\ \mathbf{0}\ \middle|\ \tfrac{q}{2}\mu_1 \cdots \tfrac{q}{2}\mu_\ell\ \right]^T \in \mathbb{Z}_q^{m+\ell}$$

όπου $(\mu_1,\ldots,\mu_\ell)$ είναι τα στοιχεία του πίνακα $\mathbf{M}$. Έτσι, απομονώνοντας τις πρώτες m σειρές του αποτελέσματος μαζί με την $(m+i)$-στη σειρά, παίρνουμε ένα dual-Regev κρυπτοκείμενο που κρυπτογραφεί το μ_i. Αυτό είναι ακριβώς το κρυπτοσύστημα στο οποίο μετατρέπει η Mahadev [Mah18a] το dual-GSW κρυπτοσύστημα (απομονώνοντας την τελευταία στηλη) και αποδεικνύει πως είναι κβαντικά ικανό. Ως αποτέλεσμα, μπορούμε να εφαρμόσουμε τον κρυπτογραφημένο CNOT υπολογισμό [Mah18a] υπολογίζοντας τον παράλληλα για κάθε ένα από τα ℓ κρυπτοκείμενα, και μετά να επιστρέψουμε στο συμπαγές σύστημα μέσω της μεθόδου bootstrapping ώστε να συνεχίσουμε τους ομομορφικούς υπολογισμούς.

4 Πρωτόκολλο Μηδενικής Γνώσης για την Κλάση QMA

Εδώ εξετάζουμε την κατασκευή επιχειρημάτων μηδενικής γνώσης (ZK arguments). Για να πετύχουμε μηδενική γνώση αξιοποιούμε την τεχνική των [BS20], που εισάγουν μια non-black-box τεχνική εξαγωγής που επιτρέπει στον προσομοιωτή να "μιμηθεί" τον έντιμο Prover χωρίς να γνωρίζει τον μάρτυρα (witness). Το πρωτόκολλό τους αποτελείται από σταθερό αριθμό γύρων (>4) και επιτυγχάνει υπολογιστική μηδενική γνώση. Εμείς επιχειρούμε να πετύχουμε στατιστική μηδενική γνώση ενώ ταυτόχρονα μειώνουμε τον αριθμό γύρων σε τέσσερις.

Απαραίτητα Κρυπτογραφικά Εργαλεία. Πριν παρουσιάσουμε την κατασκευή μας, μελετάμε ορισμένα εργαλεία που θα φανούν χρήσιμα. Το πρώτο είναι ένα κβαντικό πλήρως ομομορφικό σύστημα (QFHE), το οποίο όπως περιγράφηκε και νωρίτερα λειτουργεί όπως ένα FHE σύστημα, επιτρέποντας επιπλέον ομομορφικούς υπολογισμούς σε κβαντικά κυκλώματα και μηνύματα. Ακόμη, χρησιμοποιούμε compute-and-compare "συσκότιση" (obfuscation). Ένα compute-and-compare πρόγραμμα $\mathbf{CC}[f, s, z]$, όπου f είναι μία συνάρτηση και τα s, z συμβολοσειρές, έχει ως έξοδο την τιμή z για κάθε είσοδο x όπου $f(x) = s$ ενώ απορρίπτει κάθε άλλη είσοδο. Ένας compute-and-compare "συσκοτιστής" μετατρέπει το πρόγραμμα $\mathbf{CC}$ στο "συσκοτισμένο" πρόγραμμα $\widetilde{\mathbf{CC}}$ όπου είναι υπολογιστικά μη-διακριτό από ένα προσομοιωμένο "dummy" πρόγραμμα που απορρίπτει όλες τις εισόδους. Τέλος, χρησιμοποιούμε και ένα conditional disclosure of secrets (CDS) πρωτόκολλο. Ένα CDS πρωτόκολλο αποτελείται από δύο γύρους και έχει ως είσοδο μια πρόταση z και ένα μήνυμα m από τον αποστολέα. Ο παραλήπτης λαμβάνει το m μόνο αν η πρόταση είναι αληθής, ενώ σε αντίθετη περίπτωση το μήνυμα παραμένει κρυφό. Ταυτόχρονα, ο μάρτυρας w της πρότασης z του παραλήπτη παραμένει κρυφός από τον αποστολέα.

Witness-Indistinguishable Arguments Στη συνέχεια μελετάμε την κατασκευή μας για επιχειρήματα μη-διακριτότητας μάρτυρα (Witness-Indistinguishable Arguments - WI) δύο γύρων, η οποία καθώς αποτελεί την βάση των παρακάτω αποτελεσμάτων. Το πρωτόκολλο βασίζεται στην κατασκευή του Shmueli [Shm20], η οποία με την σειρά της βασίζεται στο Σ-πρωτόκολλο για την κλάση QMA των [BG20]. Αυτό το πρωτόκολλο αποτελείται από τρία μηνύματα: μια δέσμευση (commitment) α, μία πρόκληση β, και μία απάντηση γ. Η χρήσιμη σε εμάς ιδιότητα είναι πως ο υπολογισμός των β και γ γίνεται με αποκλειστικά κλασσικές μεθόδους. Στη δικιά μας κατασκευή χρησιμοποιούμε μια παραλλαγή με ένα επιπλέον μήνυμα (από τον verifier στον prover) ώστε να επιτύχουμε στατιστική μηδενική γνώση, την οποία αγνοούμε στα πλαίσια αυτής της περίληψης.

Η κύρια ιδέα του πρωτοκόλου είναι να χρησιμοποιηθεί ένα (levelled) πλήρως ομομορφικό κρυπτοσύστημα με ασφάλεια κυκλώματος σε κακόβουλο περιβάλλον (maliciously circuit private), ώστε να μειώσουμε τους γύρους στο Σ-πρωτόκολλο ως εξής: ο verifier στέλνει στον prover μία κρυπτογραφημένη πρόκληση β, και ο prover υπολογίζει πρώτα την δέσμευση α και ύστερα την απάντηση γ ομομορφικά (κρυπτογραφημένη). Ο verifier, γνωρίζοντας το ιδιωτικό κλειδί του ομομορφικού κρυπτοσυστήματος, μπορεί να αποκρυπτογραφήσει το κρυπτοκείμενο που λαμβάνει και να επιβεβαιώσει την ορθότητα των (α, β, γ). Αν και διαισθητικά η ορθότητα του παραπάνω είναι άμεσο επακόλουθο της ασφάλειας του ομομορφικού κρυπτοσυστήματος, για την απόδειξη απαιτούνται οι παρακάτω αλλαγές:

- Ο prover υπολογίζει μια δέσμευση με την τυχαιότητα που χρησιμοποιείται στον

ομομορφικό υπολογισμό. Έτσι ο verifier μπορεί (στην απόδειξη ορθότητας) να επιβεβαιώσει την εγκυρότητα του Σ-πρωτοκόλλου χωρίς να γνωρίζει το ιδιωτικό κλειδί του ομομορφικού κρυπτοσυστήματος. Για να επιτευχθεί αυτό ενώ ταυτόχρονα διατηρούμε στατιστικό WI, χρησιμοποιείται ένα ιδιαίτερο σχήμα δέσμευσης, η sometimes-binding statistically hiding (SBSH) δέσμευση. Σε ένα τέτοιο σχήμα δέσμευσης, υπάρχει μια (αμελητέα μικρή) πιθανότητα να έχει τέλεια δέσμευση. Σε τέτοια περίπτωση ο verifier μπορεί να εξάγει το δεσμευμένο μήνυμα. Η ορθοτητα αποδεικνύεται με την τεχνική leveraging.

- Όλη η παραπάνω διαδικασία επαναλαμβάνεται δύο φορές και ο prover αποδεικνύει πως τουλάχιστον σε μία απο τις δύο περιπτώσεις έγιναν ορθά οι υπολογισμοί, χρησιμοποιώντας ένα στατιστικό WI (για την κλάση NP). Αυτό αρκεί για να αποδείξουμε την μη-διάκριση του μάρτυρα του συνολικού πρωτοκόλλου μιας και στην απόδειξη μπορούμε να "ανταλλάξουμε" τον μάρτυρα σε κάβε βήμα ξεχωριστά.

Όλα τα παραπάνω μπορούν να κατασκευαστούν θεωρώντας την οιωνεί-πολυωνυμική (quasi-polynomial) δυσκολία του LWE προβλήματος. Εφόσον αυτό το πρωτόκολλο αποτελεί βάση για τα επόμενα, και εκείνα με την σειρά τους θα βασίζονται στην οιωνεί-πολυωνυμική δυσκολία του LWE προβλήματος

Η Τεχνική Προσομοίωσης [BS20]. Πριν συνεχίσουμε, περιγράφουμε την τεχνική προσομοίωσης των [BS20]. Για λόγους απλότητας, θεωρούμε για αρχή verifiers που δεν διακόπτουν την επικοινωνία και είναι εξηγήσιμοι (explainable), δηλαδή τα μηνύματα υποστηρίζονται από τους αλγορίθμους έντιμων verifier. Η ουσία του πρωτοκόλλου [BS20] είναι το εξαγώγιμο σχήμα δέσμευσης το οποίο λειτουργεί ως εξής:

- Ο αποστολέας δειγματοληπτεί δύο τυχαίες συμβολοσειρές s, td, καθώς και:
  - Ένα δημόσιο και ένα ιδιωτικό κλειδί $(\mathsf{pk}, \mathsf{sk})$ ενός QFHE κρυπτοσυστήματος και την κρυπτογράφηση της συμβολοσειράς td, $c_{td} = \mathsf{QFHE.Enc}(\mathsf{pk}, \mathsf{td})$.
  - Το "συσκοτισμένο" πρόγραμμα $\widetilde{\mathbf{CC}} \leftarrow Obf(\mathbf{CC}[f, s, (\mathsf{sk}, m)]$, με f να είναι η συνάρτηση αποκρυπτογράφησης του QFHE.

  Ο αποστολέας στέλνει τα $\mathsf{pk}, c_{\mathsf{td}}$ και $\widetilde{\mathbf{CC}}$ στο παραλήπτη.
- Ο παραλήπτης στέλνει μαντεύει μια τιμή y και την στέλνει κωδικοποιημένη μέσω του πρωτοκόλλου CDS.
- Ο αποστολέας απαντά με ένα μήνυμα κρυπτογραφημένο μέσω του CDS πρωτοκόλλου, ώστε αν $y = t$, τότε επιστρέφει την τιμή s. Εναλλακτικά επιστρέφει $\perp$.

Διαισθητικά, η παραπάνω διαδικασία προσφέρει δέσμευση αφού το μήνυμα στο "συσκοτισμένο" πρόγραμμα είναι μοναδικά ορισμένο, καθώς και μυστικότητα εφόσον ο receiver δεν μπορεί να μαντέψει σωστά την τιμή td, παρά με αμελητέα πιθανότητα. Ακόμη, ο προσομοιωτής μπορεί εξάγει τα (sk, m) και και να προσομοιώσει την οπτική του αποστολέα· αφού λάβει το πρώτο μήνυμα, υπολογίζει ομομορφικά το τελευταίο μήνυμα του αποστολέα χρησιμοποιώντας το κύκλωμα του, με είσοδο την κρυπτογραφημένη τιμή του td και την εσωτερική κατάσταση του αποστολέα. Το αποτέλεσμα του ομομορφικού υπολογισμού είναι το μήνυμα κρυπτογραφημένο μέσω του CDS, του οποίου η πρόταση είναι ορθή και άρα επιστρέφει την τιμή s, κρυπτογραφημένη

μέσω του QFHE. Αυτή η τιμή είναι ακριβώς η απαιτούμενη είσοδος για το $\widetilde{\mathbf{CC}}$ ώστε να επιστρέψει το μήνυμα m. Επιπλέον, το πρόγραμμα $\mathbf{CC}$ επιστρέφει μαζί και το ιδιωτικό κλειδί sk ώστε ο προσομοιωτής να μπορεί να αποκρυπτογραφήσει τα κρυπτογραφημένα μέσω QFHE μηνύματα και να παράξει ένα έγκυρο αντίγραφο επικοινωνίας T, χωρίς να πρέπει να χρησιμοποιήσει rewinding.

Μηδενική Γνώση σε 4 Γύρους. Χρησιμοποιώντας την παραπάνω τεχνική δέσμευσης, μπορούμε να αναβαθμίσουμε το WI πρωτόκολλο σε μηδενική γνώση ως εξής: Στον πρώτο γύρο ο verifier στέλνει μία δέσμευση σε μηδενική τιμή με τυχαιότητα r (ίδια με την τυχαιότητα που χρησιμοποιήθηκε στην παραγωγή των κλειδιών του QFHE κρυπτοσυστήματος). Στη συνέχεια υλοποιούμε την παραπάνω τεχνική εξαγωγής θέτοντας ως m την τυχαιότητα r. Μετά την λήξη της αλληλεπίδρασης τους, ο prover χρησιμοποιεί το WI πρωτόκολλο για να αποδείξει πως είτε γνωρίζει την τυχαιότητα r είτε $x \in \mathcal{L}$.

Για να απορρίψουμε κακόβουλες επιθέσεις όπου ο prover πλαστογραφεί έναν έγκυρο μάρτυρα για το πρωτόκολλο CDS από την κρυπτογράφηση του td, προσθέτουμε και μία SBSH δέσμευση της τιμής y, η οποία μπορεί να εξαχθεί με χαμηλή πιθανότητα, επιτρέποντας την αναγωγή της επίθεσης στην ασφάλεια του QFHE. Παράλληλα ελέγχουμε μέσω του CDS πρωτοκόλλου πως το σχήμα δέσμευσης είναι ορθά ορισμένο (δηλαδή ο prover συμπεριλαμβάνει και την τυχαιότητα της SBSH δέσμευσης ως μέρος του witness).

Ένα πρόβλημα που προκύπτει είναι πως τα περισσότερα CDS πρωτόκολλα 2 γύρων προσφέρουν υπολογιστική (έναντι στατιστικής) ασφάλειας, κάτι που αποτελεί εμπόδιο την προσπάθεια μας να πετύχουμε στατιστική μηδενική γνώση. Έτσι χρησιμοποιούμε το μετα-κβαντικό CDS πρωτόκολλο 3 γύρων με στατιστική ασφάλεια από την πλευρά του παραλήπτη, όπως φαίνεται στο [CM21].

Κακόβουλοι Verifiers. Το μόνο πρόβλημα που απομένει είναι πως θεωρήσαμε ότι οι verifiers δεν διακόπτουν την επικοινωνία και είναι εξηγήσιμοι. Για το πρώτο πρόβλημα θεωρούμε δύο προσομοιωτές (όπως στο [BS20]), έναν για την περίπτωση που διακόπτει και έναν για την περίπτωση που δεν διακόπτει την αλληλεπίδραση. Ύστερα κατασκευάζουμε έναν συνδυαστικό προσομοιωτή ο οποίος διαλέγει τυχαία ποιόν από τους δύο θα χρησιμοποιήσει. Το Watrous' rewinding λήμμα [Wat09] επιτρέπει στον προσομοιωτή να κάνει rewind μέχρι να μαντέψει σωστά, χωρίς να επηρεάζει τον verifier. Από την άλλη, για να επιβεβαιώσουμε πως ο verifier είναι εξηγήσιμος, προσθέτουμε στο πρωτόκολλο μας μια απόδειξη μηδενικής γνώσης (από τον verifier στον prover) που επικυρώνει ότι τα μηνύματα του ήταν έντιμα. Εφόσον στο πρωτόκολλο μας ο verifier είναι κλασσικός, αρκεί η απόδειξη μηδενικής γνώσης να είναι για την κλάση NP. Για να εξασφαλίσουμε στατιστική μηδενική γνώση όμως πρέπει να έχουμε στατιστική ορθότητα στο νέο ZK πρωτόκολλο, και άρα χρειαζόμαστε μια απόδειξη μηδενικής γνώσης *καθυστερημένης-εισόδου* (με στατιστική ορθότητα). Ταυτόχρονα πρέπει η απόδειξη να μην υπερβαίνει τους 3 γύρους για να μην προσθέση επιπλέον γύρω στην συνολική επικοινωνία. Εντούτοις, δεν γνωρίζουμε κάποιο πρωτόκολλο μηδενικής γνώσης 3 γύρων (πόσο μάλλον μετακβαντικό).

Παρόλα αυτά παρατηρούμε πως για την περίπτωση μας αρκεί ένα λιγότερο ισχυρό εργαλείο και μπορούμε να χρησιμοποιήσουμε *ενίοτε προσομοιώσιμη (sometimes simulatable)* μηδενική γνώση (SSim ZK) [CM21], όπου η προσομοίωση είναι πιθανή με αμελητέα μικρή πιθανότητα. Για την χρήση του παραπάνου εργαλείου πρέπει να ρυθμίσουμε τις παραμέτρους ασφαλείας των υπόλοιπων πρωτοκόλλων κατάλληλα για να αντισταθμίσουμε αυτή την εκθετική απώλεια, όμοια με το σχήμα δέσμευσης SBSH. Η SSim θυμίζει μηδενική

γνώση με υπερ-πολυωνυμική (superpolynomial) προσομοίωση (SPS) [Pas03a], με την κύρια διαφορά ότι στην SPS μηδενική γνώση ο προσομοιωτής τρέχει σε υπερ-πολυωνυμικό χρόνο, σε αντίθεση με την SSim μηδενική γνώση όπου ο προσομοιωτής τρέχει σε πολυωνυμικό χρόνο αλλά υπάρχει εκθετικά μικρή πιθανότητα επιτυχίας. Αυτή η διαφορά είναι μείζονος σημασίας για το πρωτόκολλο μας αφού κατά βάση δεν μπορούμε να κάνουμε rewind την κατάσταση του verifier και άρα απαιτούμε η προσομοίωση να είναι γραμμική.

5 Πρωτόκολλο Μηδενικής Γνώσης για την Κλάση QMA στο Χρονικό Μοντέλο

Τέλος, μελετάμε πώς να πετύχουμε μηδενική γνώση για την κλάση QMA σε δύο γύρους, μεταφέροντας το πρωτόκολλο στο χρόνικο μοντέλο. Πιο συγκεκριμένα, θεωρούμε πως τα μέλη της αλληλεπίδρασης μπορούν να μετρήσουν αξιόπιστα την πάροδο του χρόνου. Για την κατασκευή μας θεωρούμε την ύπαρξη μιας μη-παραλληλοποιήσιμης συνάρτησης F. Μία μη-παραλληλοποιήσιμη συνάρτηση είναι μια συνάρτηση η οποία μπορεί να υπολογιστεί σε χρόνο T, ενώ δεν είναι δυνατόν για έναν κακόβουλο αντίπαλο με βάθος μικρότερο του T να μαντέψει το αποτέλεσμα της με μια είσοδο x (δηλαδή δεν μπορεί να την τρέξει παράλληλα σε λιγότερο χρόνο)

Υπολογιστική Μηδενική Γνώση. Αρχικά κατασκευάζουμε ένα πρωτόκολλό με υπολογιστική μηδενική γνώση, όμοια με την [DS02] προσέγγιση. Λαμβάνουμε ως χρονική παράμετρο την τιμή T και θεωρούμε μια υπο-εκθετική μη-παραλληλοποιήσιμη συνάρτηση F, ασφαλή απέναντι σε αλγορίθμους βάθους μικρότερου του T. Ο prover υπολογίζει το κρυπτοκείμενο α κρυπτογραφώντας μία τυχαία συμβολοσειρά με ένα ομομορφικό κρυπτοσύστημα, καθώς και το αποτέλεσμα β του ομομορφικού του υπολογισμού με την συνάρτηση F. Στη συνέχεια, ο verifier στέλνει μια τυχαία τιμή x^* και ο prover στέλνει μια απόδειξη πως είτε $x \in \mathcal{L}$ είτε ξέρει μια κρυπτογράφηση α του x^*. Ο verifier αποδέχεται εάν ο prover απαντήσει εγκαίρως, η απόδειξη που στέλνει είναι έγκυρη και ο ομομορφικός υπολογισμός του α με την συνάρτηση F είναι ίση με β.

Διαισθητικά, το πρωτόκολλο είναι ασφαλές εφόσον ο prover δεν έχει χρόνο να υπολογίσει ομομορφικά εκ νέου το β. Ως αποτέλεσμα, η ορθότητα του πρωτοκόλλου αποδεικνύεται ανάγοντάς της στην αμφισβήτηση της μη-παραλληλοποιησιμότητας της συνάρτησης F. Η μηδενική γνώση αποδεικνύεται εύκολα δεδομένου ότι ο προσομοιωτής έχει την δυνατότητα να "σταματήσει τον χρόνο" (από την οπτική του verifier) και να προσομοιώσει την "σωστή" απάντηση. Είναι σημαντικό να αναφέρουμε πως η προσομοίωση είναι γραμμική και δεν αντιγράφει ούτε κάνει rewind την κατάσταση του verifier, κάνοντάς την κατάλληλη για κβαντικά πρωτόκολλα.

Στατιστική Μηδενική Γνώση. Κάνοντας κάποιες ισχυρότερες υποθέσεις, μελετάμε και μια διαφορετική προσέγγιση με στόχο την επίτευξη στατιστικής μηδενικής γνώσης. Ειδικότερα, θεωρούμε την ύπαρξη ενός μετακβαντικού time-lock puzzle. Ένα time-lock puzzle ουσιαστικά αποτελεί μια κρυπτογράφηση όπου μπορεί να αποκρυπτογραφηθεί μετά από χρόνο T, ενώ ομοίως με τις μη-παραλληλοποιήσιμες συναρτήσεις, δεν μπορεί κάποιος χρήστης να σπάσει την κρυπτογράφηση χρησιμοποιώντας παράλληλους υπολογισμούς. Στα πλαίσια αυτής της πείληψης θεωρούμε μόνο εξηγήσιμους verifier, και η μετατροπή και για κακόβουλους verifiers γινεται με γνωστές τεχνικές [BKP19, CDM20].

Στην παρούσα κατασκευή, ο verifier στέλνει μια δέσμευση σε μία μηδενική τιμή με τυχαιότητα r και ένα time-lock puzzle που κωδικοποιεί αυτή την τυχαιότητα. Ύστερα ο prover στέλνει μία WI απόδειξη, αποδεικνύοντας πως είτε γνωρίζει κάποια πρόταση $x \in \mathcal{L}$ είτε πως γνωρίζει την τυχαιότητα r. Ο verifier αποδέχεται αν ο prover απαντήσει έγκαιρα και η απόδειξη του είναι έγκυρη. Διαισθητικά, ένας κακόβουλος prover δεν μπορεί να λύσει το time-lock puzzle στον απαιτούμενο χρόνο, ενώ για να αποδείξουμε μηδενική γνώση επικαλούμαστε πως ο προσομοιωτής μπορεί ξανά να "παγώσει τον χρόνο" και να λύσει το time-lock puzzle, λαμβάνοντας την τυχαιότητα που μπορεί να την χρησιμοποιήσει σαν witness στην WI απόδειξη.

6 Σύνοψη και Μελλοντικές Επεκτάσεις

6.1 Σύνοψη

Με την άνοδο των κβαντικών υπολογιστών, η κβαντική κρυπτογραφία γίνεται όλο και περισσότερο δημοφιλής τα τελευταία χρόνια. Αν και ορισμένα κλασσικά κρυπτογραφικά πρωτόκολλα μετατρέπονται εύκολα σε κβαντικά, υπάρχει ακόμη μεγάλος αριθμός κβαντικών πρωτοκόλλων των οποίων η πολυπλοκότητα υστερεί συγκριτικά με τα αντίστοιχα κλασσικά. Σε αυτή την εργασία μελετήθηκε η πολυπλοκότητα επικοινωνίας των Κβαντικών Πλήρως Ομομορφικών Συστημάτων κρυπτογράφησης και των πρωτοκόλλων Μηδενικής Γνώσης, επιτυγχάνοντας αποδοτικότητα ανάλογη των κλασσικών πρωτοκόλλων.

Όσον αφορά τα πλήρως ομομορφικά συστήματα (FHE), παρουσιάζουμε δύο κατασκευές κβαντικών FHE συστημάτων κρυπτογράφησης ρυθμού-1, επιτυγχάνοντας βέλτιστη πολυπλοκότητα μεταξύ της πληροφορίας που μεταδίδεται από το ένα μέλος στο άλλο. Στην πρώτη κατασκευή, θεωρώντας ένα κβαντικό κρυπτοσύστημα με υβριδικά κρυπτοκείμενα (τα οποία περιέχουν κβαντική αλλά και κλασσική πληροφορία), χρησιμοποιούμε ένα κλασσικό FHE ρυθμού-1 ώστε να εναλλάσσουμε την κλασσική πληροφορία από το αρχικό κρυπτοσύστημα σε αυτό, εξασφαλίζοντας πολυπλοκότητα επικοινωνίας ίση με $(|\,|\psi\rangle\,|+|C(|\psi\rangle)|)\cdot(1+o(1))$. Για να συνεχίσουμε τους ομομορφικούς υπολογισμούς μετατρέπουμε ξανά την πληροφορία στο αρχικό κρυπτοσύστημα. Στο δεύτερο πρωτόκολλο, προσεγγίζουμε το πρόβλημα πιο ειδικά και κατασκευάζουμε ένα κλασσικό (μετα-κβαντικό) κρυπτοσύστημα για τα υβριδικά κρυπτοκείμενα το οποίο είναι από μόνο του ρυθμού-1. Κατά αυτόν τον τρόπο αποφεύγουμε τον κύκλο των δύο ιδιωτικών κλειδιών που δημιουργείται.

Σχετικά με τα πρωτόκολλα μηδενικής γνώσης, αρχικά κατασκευάζουμε ένα στατιστικό WI πρωτόκολλο για την κλάση QMA θεωρώντας την οιωνεί-πολυωνυμική δυσκολία του LWE προβλήματος. Στη συνέχεια μπορούμε και κατασκευάζουμε ένα πρωτόκολλο που επιτυγχάνει στατιστική μηδενική γνώση σε 4 γύρους. Επιπλέον, μεταφέροντας το πρωτόκολλο στο χρονικό μοντέλο, εξασφαλίζουμε μηδενική γνώση σε δύο γύρους (τόσο υπολογιστική όσο και στατιστική με κάποιες επιπλέον υποθέσεις)

6.2 Μελλοντικές Επεκτάσεις

Τα παραπάνω αποτελέσματα μπορούν να ως βάση για μελλοντικές δουλειές και επεκτάσεις.

Μία πιθανή επέκταση ειναι η κατασκευή ενός κβαντικού FHE κρυπτοσυστήματος ρυθμού-1 πολλαπλών κλειδιών (Rate-1 Multi-Key Quantum FHE scheme), συνδυάζοντας τα αποτελέσματα μας με μια κατασκευή κβαντικού FHE πολλαπλών κλειδιών [ABG+20] που

ομοίως με την δική μας δουλειά αξιοποιούν κρυπτοσύστηματα υβριδικής μορφής (όπως το [Mah18a]). Μία διαφορετική προσέγγιση θα ήταν η κατασκευή ενός verifiable FHE [ADSS17], όπου υπάρχει η δυνατότητα απόδειξης κάποιας ιδιότητας του μηνύματος ενώ αυτή παραμένει κρυπτογραφημένη και χωρίς να αποκαλύπτουμε καμία επιπλέον πληροφορία.

Όσον αφορά την μηδενική γνώση, ένα ανοιχτό πρόβλημα είναι να να επεκτείνουμε τα αποτελέσματα σε αποδείξεις μηδενικής γνώσης (αντί για επιχειρήματα), όπου το πρωτόκολλο θα έχει στατιστική ορθότητα. Ένα ακόμη άλυτο πρόβλημα είναι να επιτύχουμε την ίδια πολυπλοκότητα θεωρώντας την πολυωνυμική δυσκολία του LWE προβλήματος (έναντι της οιωνεί-πολυωνυμικής). Πρόσφατες εργασίες έχουν ακόμη εστιάσει σε black-box προσεγγίσεις μετακβαντικής ϵ-μηδενικής γνώσης σε σταθερό αριθμό γύρων [**?**, CCLY21a], όπου θα μπορούσε να μειωθεί σε 4 γύρους με όμοιες τεχνικές. Τέλος, απουσιάζει ακόμη από την βιβλιογραφία κάποιο πρωτόκολλο μηδενικής γνώσης 3 γύρων στο απλό μοντέλο, κάτι που θα ξεπερνούσε την πολυπλοκότητα επικοινωνίας των σημερινών πρωτοκόλλων μηδενικής γνώσης.

Ενδιαφέρον εμφανίζει και η εφαρμογή αυτών των κατασκευών σε κρυπτογραφικά πρωτόκολλα ψηφοφορίας, δεδομένου ότι σε παλαιότερες αλλά και πιο πρόσφατες κατασκευές χρησιμοποιούνται τεχνικές που αξιοποιούν ομομορφική κρυπτογράφηση και αποδείξεις μηδενικής γνώσης [SK94, KY02, GPZZ18, GPZZ21].

Chapter 1

Introduction

1 Backrounded Motivation

Cryptography has become a major part of our everyday lives, providing us with tools that help us realise secure communication between two or more parties. Apart from the most well known tools, like public key encryption schemes that allow secure communication via a public channel, there exist more nuanced and advanced protocols that can solve more complicated problems. Some of the most useful protocols with a profound impact in cryptography are:

- Fully-Homomorphic Encryption (FHE) schemes, which allow one to evaluate any function over encrypted data.
- Zero Knowledge (ZK) proofs which allow one to prove the veracity of a statement while revealing nothing beyond that.

The feasibility of the above results is not enough; in order for them to eventually be useful in real systems they need to be efficient. Thus cryptographers aim to minimize their communication complexity. Communication complexity in general specifies the size and number of messages sent between the involved parties.

Taking a closer look at the aforementioned protocols, an FHE scheme allows one party to send the encryption of its input m under a public key such that the other party can later, holding a circuit C, compute and send

$$\mathsf{Enc}(m) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(C(m)),$$

without learning any information about the message m (apart from its encrypted value). Among other applications, computation over encrypted data is useful in the case where a computationally constrained client uploads some data to a powerful server that can perform expensive computation, while preserving data privacy. In this setting, it is important to ensure that the communication overhead introduced by the FHE protocol does not nullify the efficiency gains of outsourcing the computation to a server. Recently, it was shown [BDGM19] that there exist FHE protocols where the communication complexity approaches that of the *insecure protocol* (where the first party sends its input m in plain), assuming the hardness of the learning with errors (LWE) problem.

Concerning ZK proofs, since their introduction [GMR89], they have had a profound impact on cryptography and theoretical computer science at large. Due to their foundational importance and large applicability, ZK proof systems have been the objective of a long series of work aiming at understanding the necessary assumptions and their round complexity: Under standard computational assumptions, any NP statement can be proven in as few as four rounds of interaction [GMW86, GK96].

In contrast to the classical case, much less in known for the above primitives in the quantum setting. In its most general form, quantum FHE allows the transformation

$$\mathsf{Enc}(|\psi\rangle) \xrightarrow{\mathsf{Eval}(C,\cdot)} \mathsf{Enc}(C(|\psi\rangle))$$

where $|\psi\rangle$ is some arbitrary quantum state and C is some unitary matrix. Despite the fact that this problem has received far less attention, we believe that this question is even more pressing than the classical case, due to the large gap between quantum capabilities of regular users and servers sitting on the cloud. Even in a future where regular users will be equipped with quantum-capable computers, it is likely that intensive quantum computations will be exclusive to large computer clusters. There have been some constructions of quantum fully homomorphic encryption (QFHE) schemes [BJ15], even assuming a completely classical client (Alice) [Mah18a]. However, to the best of our knowledge, the question of quantum FHE schemes with compact (i.e. independent of the size of the circuit) communication complexity has not been considered in the literature. Circling back to ZK proofs, their construction for QMA (the quantum analogue of NP) has been introduced only recently [BJSW16] and the best known result, in terms of round complexity is from the very recent work of Bitansky and Shmueli [BS20] where they presented a constant-round computational zero-knowledge argument system (i.e. with computational soundness).

Motivated by the unsatisfactory state of affairs, we ask the following questions:

Can we construct quantum FHE with minimal communication complexity?
Does proving QMA statements inherently introduce additional rounds of interaction?

In this work, we study these problems and we present protocols in a variety of settings that match the round complexity of their classical counterparts.

2 Thesis Contribution

In this work we initiate the study of the communication complexity of FHE for quantum circuits (quantum FHE) and ZK proofs for QMA.

Rate-1 QFHE. Our first main result is a protocol to compute any quantum circuit with communication complexity $(|\,|\psi\rangle\,|+|C(|\psi\rangle)|)\cdot(1+o(1))$ to compute some quantum circuit C over some state $|\psi\rangle$. This approaches the communication complexity of the insecure protocol, where the first party sends the state $|\psi\rangle$ in plain, and it is (asymptotically) optimal. As we discussed before, all known Quantum Fully Homomorphic Encryption (QFHE) schemes [Mah18a, Bra18] blow up the ciphertext by a polynomial factor $\mathsf{poly}(\lambda)$ for evaluated ciphertexts, i.e. they have low (inverse polynomial) rate. This means that

the communication complexity of the resulting protocol would be at least $|C(|\psi\rangle)| \cdot \mathsf{poly}(\lambda)$. Thus we reduce this gap by constructing a QFHE scheme with nearly optimal ciphertext expansion. Our protocol assumes the quantum hardness of the LWE problem (with polynomial modulo-to-noise ratio) in addition to a circular security assumption to apply the bootstrapping theorem [Gen09].

Theorem 2.1 (Informal)**.** *Assuming the quantum hardness of the LWE problem, there exists a (leveled) QFHE scheme with rate-1.*

Note that the above result can be combined with the other result presented in [CDM20], which transforms any QFHE scheme to a QFHE with scheme malicious circuit privacy. Combining both results we get a secure function evaluation scheme with nearly optimal communication security.

ZK for QMA. We begin by considering a weak version of zero-knowledge, namely, witness indistinguishability (WI), which only guarantees that a distinguisher cannot tell whether the prover used w_0 or w_1, where (w_0, w_1) are two valid witnesses for the given statement. While not immediately meaningful on its own, this notion and protocol will serve as the basis for our further results. We construct a 2-round protocol with statistical WI, assuming the quasi-polynomial hardness of the learning with errors (LWE) problem [Reg05]. This matches the round complexity of statistical WI protocols for NP [KKS18, BFJ+20, GJJM20].

Theorem 2.2 (Informal)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, there exists a 2-round statistical WI argument for QMA.*

Next, as our main result, we show how to compile the above into a fully-fledged 4-round *statistical* ZK argument for QMA. The protocol is a round compressed version of the [BS20] approach and, as such, also has a non-blackbox simulator.[1] In contrast to [BS20] our protocol achieves statistical ZK and relies on computational assumptions only to argue about soundness. On the flip side, we rely on the (quantum) *quasi-polynomial* security of the LWE problem and on the quantum fully-homomorphic encryption (QFHE). Our protocol matches the round complexity of the best known ZK proofs/arguments for NP against classical adversaries (albeit using non-blackbox simulation). We stress that, prior to our work, even post-quantum *statistical* ZK for NP was only known in polynomial rounds [Unr12, ACP20].

Theorem 2.3 (Informal)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem and a quasi-polynomially secure QFHE scheme, there exists a 4-round statistical ZK argument for QMA.*

Finally we consider the question of 2-round ZK in the timing model: Since 2-round ZK is known to be impossible [GO94] without additional assumptions, a common relaxation is to allow parties to reliably measure time during the execution of the protocol. In this context, we revisit the Dwork-Stockmeyer [DS02] approach and lift it to the quantum setting. In addition to quasi-polynomial LWE, we assume the existence of a post-quantum non-parallelizing function (e.g. repeated hashing).

[1] There is evidence [CCLY21b] that non-blackbox simulation is necessary for constant-round ZK against quantum adversaries.

Theorem 2.4 (Informal)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, an FHE scheme, and an average-case non-parallelizing function, there exists a 2-round computational ZK argument for QMA in the timing model.*

A shortcoming of the above approach is that it only achieves computational ZK. To overcome this issue, we propose a different route to construct statistical ZK in the timing model, which relies on slightly stronger assumptions (namely, post-quantum time-lock puzzles).

Theorem 2.5 (Informal)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem and a quasi-polynomially sequential post-quantum time-lock puzzle, there exists a 2-round statistical ZK argument for QMA in the timing model.*

3 Related Work

The (weakened) problem of quantum homomorphic encryption by allowing a quantum client has been studied extensively in the recent years [BJ15, OTF15, TKO$^+$14, LC18, NS17, DSS16] and has led to major advancements in delegated quantum computing.

The problem of secure (i.e. blind) computation of quantum circuits [BFK09, DNS10, DNS12] also has a strong tradition in the quantum cryptography literature. Blind computation is similar to quantum homomorphic encryption in the sense that they share the goal of carrying out a computation over encrypted data, but blind computation allows multiple rounds of interaction between the client and the server. To the best of our knowledge, the only two-round protocol was given in the recent work of Bartusek et al. [BCKM20]. In contrast to our work, the resulting communication complexity is proportional to the size of the circuit (i.e. it is not compact). On the flip side, they achieve the strong notion of simulation security and they assume any post-quantum two-round oblivious transfer, whereas we crucially rely on the LWE assumption.

A similar line of work has been focusing on *verifiability* of quantum computation (see [Mah18b] and references therein) where it is required that a malicious Bob must prove to Alice that he evaluated the "correct" circuit C (clearly, this notion only makes sense when the circuit C is public and the resources needed by Alice to check Bob's proof are less than those required to evaluate C).

We also mention a series of recent works [BG20, CVZ20, ACGH20, CCY20b, Shm20, BM21] that considers the problem of non-interactive ZK for QMA. All of these works require some notion of trusted setup, which is unavoidable for 1-round protocols. We also mention another line of work [Unr12, HSS11, LN11, ARU14, AL20] that studies the strong notion of arguments of knowledge in the quantum settings. Finally, in the multi-prover settings, it is known that NEXP [CFGS18] and MIP* [GSY19] admit perfect ZK interactive proofs (sound against entangled quantum provers).

4 Overview of Results and Techniques

Here we present an overview of the main technical ideas presented in the paper. For further details, we refer the reader to the technical sections.

4.1 Rate-1 Quantum Fully-Homomorphic Encryption

We first examine the description of the rate-1 QFHE scheme.

What Makes This a Non-Trivial Problem? Before describing our solution, it is instructive to understand why existing schemes fail to achieve good ciphertext expansions and have low (inverse polynomial) rate. In the schemes from [Mah18a, Bra18], a ciphertext encrypting an ℓ-qubit state $|\psi\rangle$ is of the form

$$\mathsf{QOTP}((x_1, z_1, \ldots, x_\ell, z_\ell), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, (x_1, z_1, \ldots, x_\ell, z_\ell))$$

where the QOTP is applied qubit-by-qubit and the classical string $\mathsf{otk} = (x_1, z_1, \ldots, x_\ell, z_\ell)$ is encrypted bit-by-bit. It is not hard to see that this scheme has inverse polynomial rate, due to the blow-up introduced by the (classical) FHE encryption.

One obvious solution to improve the rate would be to adopt the *hybrid encryption* approach and sample the QOTP key using a cryptographic PRG with polynomial stretch. That is, we could improve the rate of the ciphertexts by computing

$$\mathsf{QOTP}(\mathsf{PRG}(\mathsf{seed}), |\psi\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{seed})$$

for some uniformly sampled $\mathsf{seed} \leftarrow\!\!\$\, \{0,1\}^\lambda$. Note that we can still homomorphically compute a function in the resulting scheme, since one can always convert the ciphertexts back to their original form by evaluating the PRG homomorphically.

While this generic approach suffices for fresh ciphertexts, the troubles start once we begin to evaluate functions homomorphically: Depending on the gate that we apply to the quantum state, the one-time key otk changes accordingly to otk'. For the case of the encrypted CNOT operation, the modification is even non-deterministic. While [Mah18a] shows a way to update the classical component consistently, this method conflicts with our hybrid encryption strategy. This is because the modified otk' will most likely lie outside the support of the PRG and thus a string seed' such that $\mathsf{PRG}(\mathsf{seed}') = \mathsf{otk}'$ might simply not exist. Thus we are stuck with a classical encryption $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk})$, which brings us back to our original problem. Even assuming an ideal case where the classical FHE scheme has optimal rate, we still have a constant (> 2) ciphertext blow-up. Since two classical bits are necessary to encrypt a qubit [AMTDW00], we seem to have encountered a roadblock.

Spooky Interactions. On a high-level, our solution will leverage the structure of a special classical FHE scheme to refresh our QFHE ciphertext to the hybrid (i.e. rate-1) state. More in details, we observe that certain recent FHE schemes [BDGM19] pack k classical bits in ciphertexts of the form $c = (\mathbf{c}_0, c_1, \ldots, c_k) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^k$, for some modulus q and $n = \mathsf{poly}(\lambda)$. The interesting property for us is that the last k-bits of the ciphertexts are *non-locally* correlated with the secret key sk. Specifically, the decryption recovers the plaintext by computing

$$\mathsf{Dec}(\mathsf{sk}, c) = F(\mathsf{sk}, \mathbf{c}_0) \oplus (c_1, \ldots, c_k)$$

for some function F, whose exact description is irrelevant for us. This property, that we refer to as *spooky decryption*,[2] will be the key to our solution.

[2]The name is inspired by a similar phenomenon happening in multi-key FHE schemes [DHRW16].

The Solution. Equipped with the tool described above, we can convert evaluated QFHE ciphertexts of the form $(\mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle), \mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}'))$ back to a rate-1 form using the following procedure:

- Convert $\mathsf{QEnc}(\mathsf{pk}, \mathsf{otk}')$ into an FHE ciphertext with spooky decryption via bootstrapping (i.e. evaluating the decryption circuit of QEnc homomorphically).
- Parse the resulting ciphertext as

$$c = (\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^{2\ell}.$$

- Return $\mathbf{c}_0$ and $\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle)$.

Since $|\mathbf{c}_0| = \mathsf{poly}(\lambda)$, the size of the compressed ciphertext is ℓ qubits plus $\mathsf{poly}(\lambda)$ bits of classical information.This rate is optimal (up to polynomial additive terms), given that any public-key encryption scheme must have ciphertexts of size at least λ bits, so an additive term in the security parameter is unavoidable. This is the exact situation here, except that we have a larger additive term, which is however asymptotically insignificant.

To see why this procedure gives us a decryptable ciphertext, re-arrange the equation above to obtain

$$F(\mathsf{sk}, \mathbf{c}_0) = (x'_1, z'_1, \ldots, x'_\ell, z'_\ell) \oplus (c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z})$$

which is the correct one-time key of the quantum state

$$\begin{aligned}
&\bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP}(\mathsf{otk}', |\psi'\rangle \\
&= \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \bigotimes_{i\in[l]} \left(X^{x'_i} Z^{z'_i}\right) \cdot |\psi'\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{c_{i,x}\oplus x'_i} Z^{c_{i,z}\oplus z'_i}\right) \cdot |\psi'\rangle \,.
\end{aligned}$$

A Non-Generic Approach. The savvy reader might have noticed that the the above solution introduces an additional secret key in the scheme. In the transformation from leveled to fully homomorphic this results in a different circularity assumption: Instead of the plain circular security of the QFHE scheme, we now need to assume that semantic security is retained in the presence of a two-key cycle. While formally the two assumptions are incomparable, this motivates us to investigate on whether we can achieve full homomorphism and rate-1 under the plain one-key circularity. We show that this is fact the case, by constructing a packed version of the dual-GSW FHE scheme [Mah18a] and we prove that it is quantum capable (i.e. it supports the homomorphic evaluation of quantum circuits). Next, using the shrinking algorithm from [BDGM19], we end up with a rate-1 quantum capable scheme with the same *spooky decryption* introduced above. Thus, following a similar technique, we again obtain a rate-1 quantum fully homomorphic encryption scheme.

Packed Dual-GSW scheme. The construction of the packed dual-GSW scheme is essentially the dual of the scheme from Hiromasa et al. [HAO15]. Recall that, in the (non-packed) dual-GSW scheme, the ciphertext of a plaintext μ is of the form

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mu\mathbf{G} \in \mathbb{Z}_q^{(m+1)\times(m+1)\log q}$$

where $\mathbf{A}' \in \mathbb{Z}_q^{(m+1)\times n}$, $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+1)\log q}$ and $\mathsf{sk} \cdot \mathbf{A}' = 0$, with sk being the secret key of the scheme. The plaintext information is encoded in the last row of the ciphertext. In a packed scheme, we want to encrypt ℓ-bit messages, so we interpret the plaintext as a diagonal matrix $\mathbf{M} \in \{0,1\}^{\ell\times\ell}$ containing ℓ bits, and we define the ciphertext to be

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y} \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times(m+\ell)\log q}$$

where $\mathbf{Y} \in \{0,1\}^{(m+\ell)\times(m+\ell)}$ is an encoding of the message, $\mathbf{A}' \in \mathbb{Z}_q^{(m+\ell)\times n}$, and $\mathbf{S} \in \mathbb{Z}_q^{n\times(m+\ell)\log q}$.

In order to maintain the scheme's homomorphic properties and be able to compute a NAND gate without altering the structure of the ciphertext, we select a message encoding that preserves plaintext-point-wise addition and multiplication, as well as the relation $\mathbf{Y} \cdot \mathbf{A}' = 0$ to cancel out the mixed term of the multiplication. To achieve this, the secret key is defined as $\left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array}\right]$, for a matrix $\mathbf{E}_{sk} \in \{0,1\}^{\ell\times m}$ and $\mathbf{Y}$ is defined as $\left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{M} \cdot \mathsf{sk} \end{array}\right]$. Note that, in order to produce said form of $\mathbf{Y}$, the key-generation algorithm needs to provide encryptions of $\mathbf{P}_i$ for $i \in \{0,\ldots,\ell\}$, where $\mathbf{P}_i$ is a diagonal matrix with 1 in slot (i,i) and zero everywhere else. Then, the encryption algorithm sums all the encryptions corresponding to the input message and re-randomizes the result.

To see why the scheme is quantum capable, observe that by summing up columns $(m+i)\log q$ for $i \in \{1,\ldots,\ell\}$ in our ciphertext, we end up with

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \left[\begin{array}{c|ccc} \mathbf{0} & \frac{q}{2}\mu_1 & \cdots & \frac{q}{2}\mu_\ell \end{array}\right]^T \in \mathbb{Z}_q^{m+\ell}$$

where $(\mu_1,\ldots,\mu_\ell)$ are the entries in $\mathbf{M}$. Next, by isolating the first m rows of the result, alongside the $(m+i)$-th row, we obtain a dual-Regev ciphertext encrypting μ_i. This is the same scheme that Mahadev [Mah18a] converts dual-GSW to (by isolating the last column), and shows that it is quantum capable. Thus, we can apply the encrypted CNOT operation from [Mah18a] using each of the ℓ ciphertexts in parallel and then bootstrap back into the packed scheme to continue the homomorphic computations. We refer the reader to Section 5 for further details.

4.2 Zero Knowledge Arguments

To achieve ZK, we leverage the generic approach of [AL20, BS20], which introduces a non-black-box quantum extraction technique that allows the simulator to emulate the honest prover without knowing the witness. The extraction protocol consists of constant (> 4) number of rounds and the resulted ZK scheme for QMA in [BS20] achieves only computational ZK, so the challenge for us will be to lift this paradigm to the statistical ZK settings while at the same time squeezing the number of rounds down to 4.

Some Cryptographic Tools. Before presenting the construction we recall some necessary tools that we use. The first is a quantum fully homomorphic encryption (QFHE)

scheme, which as described earlier, works similarly to an FHE scheme, allowing us to additionally perform homomorphic evaluations of quantum circuits and inputs. We also use a compute-and-compare obfuscation. A compute-and-compare program $\mathbf{CC}[f, s, z]$ where f is a function, and s, z are strings, outputs z on every input x such that $f(x) = s$ and rejects the rest of the inputs. A compute-and-compare obfuscator compiles a $\mathbf{CC}$ program to the obfuscated program $\widetilde{\mathbf{CC}}$ and is computationally indistinguishable from a simulated dummy program, that rejects on all inputs. Finally, we use a conditional disclosure of secrets (CDS) protocol. This two-round protocol is parametrized by a statement z and a message m from the sender: The receiver is able to recover m if the statement is correct, whereas m stays hidden if this is not the case. Simultaneously, the witness w (held by the receiver) for z should be kept secret from the eyes of the sender.

Witness-Indistinguishable Arguments We continue by outlining our construction of a 2-round WI protocol for QMA, which will constitute the basis for the following results. Our protocol is based on the template from [Shm20], which in turn relies on the sigma protocol for QMA introduced in [BG20]. Such a protocol consists of the canonical three messages: A commitment α, a challenge β, and a response γ. The important property (also used in [Shm20]) is that the computation of β and γ is completely classical. In our protocol, we actually use a statistically zero-knowledge variant that has an additional first message (from the verifier to the prover), but for the sake of this overview we can ignore this aspect and simply consider a three message sigma protocol.

The basic idea of the protocol is to use a maliciously circuit private (levelled) homomorphic encryption to round-collapse the sigma protocol: The verifier sends to the prover an encrypted challenge β, then the prover computes in plain a commitment α and evaluates homomorphically the response function to return an encrypted version of γ. The verifier, who knows the secret key of the homomorphic encryption, can decrypt the incoming ciphertext and verify the validity of the transcript (α, β, γ). While intuitively the soundness follows from the semantic security of the homomorphic encryption scheme, turning this into a provably secure scheme requires some tweaks with some additional tools:

- The prover is let to compute a commitment to the random coins used in the homomorphic evaluation procedure. This allows the verifier (in the soundness proof) to check the validity of the transcript without knowing the secret key of the homomorphic encryption scheme. To achieve this while maintaining statistical WI, a special kind of sometimes-binding statistically hiding (SBSH) commitment is used. This is a standard statistically hiding commitment scheme, which has a certain (negligibly small) probability to be perfectly binding. When such event happens, the verifier can extract the committed message. Soundness is then argued by a standard complexity leveraging argument.

- A dual-track approach is used, where the above process is repeated twice and the prover shows that at least one of the two instances was computed correctly, via a statistical WI (for NP). This is sufficient to prove the overall WI of the protocol since the witness can be switched step-by-step for each branch.

All of the above building blocks can be instantiated assuming the quasi-polynomial hardness of the LWE problem. Since this protocol constitutes the basis of the upcoming ZK constructions, they will also be based on the quantum quasi-polynomial hardness of LWE.

The "Homomorphic Trapdoor" Technique. We now briefly recall the simulation technique from [BS20, AL20]. For simplicity, we consider a verifier that never aborts and that is explainable, i.e. it computes all its messages in the support of algorithms as dictated by the honest protocol. The crux of their protocol consists of the following extractable commitment scheme:

- The sender samples two random strings s, td in addition to:
  - A public and secret key $(\mathsf{pk}, \mathsf{sk})$ of a QFHE scheme and an encryption $c_{td} = \mathsf{QFHE.Enc}(\mathsf{pk}, \mathsf{td})$ of td.
  - The obfuscated program $\widetilde{\mathbf{CC}} \leftarrow Obf(\mathbf{CC}[f, s, (\mathsf{sk}, m)]$, where f is the decryption circuit of QFHE.

  The sender sends $\mathsf{pk}, c_{\mathsf{td}}, \widetilde{\mathbf{CC}}$ to the receiver.
- The receiver encodes a guess y via the CDS protocol.
- The sender responds with a message encrypted via the CDS protocol, such that, if the guess y is equal to td, then the message decrypts to s. Otherwise it returns $\perp$.

Intuitively, such a procedure is binding since the message in the obfuscated program is uniquely determined, and hiding since no receiver guesses td correctly, except with negligible probability. Furthermore, a simulator can extract the message (sk, m) and simulate the sender's view: After the simulator gets the first message, it homomorphically computes the sender's last message using the sender's circuit with inputs the encryption of td and the inner state of the sender. The result of the homomorphic computation is the message encrypted with the CDS, whose statement is satisfied and hence it returns s encrypted under QFHE. This is exactly the input needed for $\widetilde{\mathbf{CC}}$ in order to obtain m. Note that the simulator is able to also produce a valid transcript T without rewinding the adversary, since the $\mathbf{CC}$ program also returns sk, which can be used by the simulator to decrypt the QFHE-encrypted messages.

From WI to ZK in 4 Rounds. Given the above extractable commitment, one can boost a 2-round WI argument into a fully-fledged 4-round ZK protocol, as follows: The verifier in the first round sends a commitment to zero with randomness r (which is the same randomness used in the QFHE keys generation algorithm). Then, they perform the above quantum extraction technique with r as the message m. After the interaction, the prover utilizes the WI argument introduced before and sends a proof that either he knows the randomness r or that $x \in \mathcal{L}$.

To rule out mauling attacks where the prover could maul a QFHE encryption of td into a valid witness for the CDS protocol, we additionally include an SBSH commitment of y, which can be extracted with low probability, thus enabling a reduction against the semantic security of the QFHE scheme. Consistency is guaranteed by checking that

the SBSH commitment is well-formed within the CDS protocol (i.e. the prover includes also the randomness of the SBSH commitment as part of the witness).

One immediate problem is that most existing two-round CDS protocols only provide computational security for the receiver, which would result in us achieving only computational ZK. In order to achieve statistical security, we utilize the 3-round post quantum CDS protocol with statistical receiver privacy constructed in [CM21].

Malicious Verifiers. The only remaining problem is that the ZK protocols are simulatable under the assumption that the verifier is non-aborting and explainable. To deal with aborting verifiers, we (as done in [BS20]) define two simulators, an aborting and a non-aborting one, and we let the combined simulator guess which of the two he should use. Watrous' rewinding lemma [Wat09] allows the simulator to rewind until the guess was correct without disturbing the verifier's state. To ensure that the verifier is explainable, we augment the protocol with an additional ZK proof (from the verifier to the prover) that the messages where computed honestly. Note that in our protocol the verifier is completely classical, so ZK for NP suffices. In order to achieve statistical soundness and maintain the statistical ZK property though, we need a *delayed-input* ZK proof (with statistical soundness). This proof needs also to not exceed 3 rounds so as not to increase the rounds of the original protocol. Unfortunately, we do not have a 3-round ZK proof, let alone a post-quantum one.

We observe however, that for our case a weaker notion suffices and we can use *sometimes simulatable* zero-knowledge [CM21], where simulation is possible with some (negligibly) small probability. In order to be meaningfully used, one must set the security parameters of other primitives to account for this exponential loss, much like with SBSH commitments. Sometimes-simulatable (SSim) ZK is reminiscent of ZK with super-polynomial simulation (SPS) [Pas03a] but with a crucial difference: In SPS-ZK the simulator runs in super-polynomial time, whereas in SSim-ZK the simulator runs in polynomial time but only has an exponentially small success probability. This difference is important in our settings since (in general) we cannot rewind the state of the verifier and it is therefore important that the simulation is straight-line.

4.3 Zero Knowledge in the Timing Model

Finally, we investigate how to achieve ZK in QMA in two rounds, by moving the protocol to the timing model. In other words, we assume that the parties can reliably measure the lapse of time during the interaction. In order to achieve this, we assume the existence of a non-parallelizing function F. A non-parallelizing function is a function that can be computed in time T, while the result of the function with an input x cannot be predicted by an attacker with depth less than T (i.e it cannot be run quicker in parallel time).

Computational Zero-Knowledge. For our first construction we revisit the [DS02] approach. The protocol is parametrized by a time parameter T and and we assume a sub-exponentially non-parallelizing function F, secure against algorithms with depth less than T. The prover first computes an encryption α of a random string, and its homomorphic evaluation β with the function F. Then, after the verifier sends a random value x^*, the prover sends a proof that either $x \in \mathcal{L}$ or that he knows an encryption α

of x^*. Eventually, the verifier accepts if the prover responds in time, the proof is valid and the homomorphic evaluation of α with F is equal to β.

Intuitively, the protocol is secure because the prover doesn't have the time to homomorphically recompute β. Thus, soundness is proven by reducing to breaking the non-parallelizability of F. The zero-knowledge property is easily proven, having in mind that the simulator is allowed to "freeze time" (from the perspective of the verifier) while simulating the accepting transcript. Note that the simulation is straight-line and does not copy nor rewinds the state of the verifier, which makes it suitable for the quantum settings.

Statistical Zero-Knowledge. Assuming slightly stronger assumptions, we propose a different approach, achieving statistical ZK. In particular we assume the existence of a post-quantum time-lock puzzle. A time-lock puzzle essentially provides an encryption that is breakable after time T, but where one cannot gain a significant speedup with parallel computation (similar to the non-parallelizability). For the sake of this overview, we only consider explainable verifiers and the conversion to malicious verifiers can be done with standard techniques [BKP19, CDM20].

In our construction, the verifier sends a commitment to 0 with randomness r, along with a time-lock puzzle encrypting said randomness. Then the prover sends a WI proof proving that either it knows a statement $x \in \mathcal{L}$ or that it knows the randomness r. The verifier accepts if the prover responds in time and the proof is valid. Intuitively, a malicious prover cannot solve the time-lock puzzle in the necessary time, whereas in order to prove ZK, the simulator can again "freeze time" and solve the time-lock puzzle, acquiring the randomness and using it as a witness in the WI proof.

Chapter 2

Quantum Cryptography

Here we present some introductory notions to quantum computations and quantum cryptography. Part of the upcoming theoretical backround was taken from [Zha19, Zha18, Yue21]

In traditional (i.e. classical) cryptography, the most basic setting is where Alice wants to send a message to Bob, but suspects there is an eavesdropper Eve. If Alice and Bob share a secret key, Alice can encrypt a message m and send it to Bob, who in turn can run a decryption algorithm and get the message. Eve, without the secret key, cannot get any information about the message. Even in this very simple example, many assumptions are made. Specifically, we assume that both the related parties (Alice, Bob, Eve) and the communication channel obey Newtonian physics. We now know that this description can be incomplete; Newtonian physics fail when things are really big or really small. In the latter case, with the development of quantum computers, we need to also consider quantum mechanics.

1 Quantum Computing

Here we introduce some basic notions of quantum computing, necessary for understanding this work.

It might be useful to consider quantum information theory as a generalization of classical probability theory, where the probabilities are complex numbers instead of real positive numbers.

Quantum States. Let B be a finite set of classical basis states. A (pure) quantum state is a unit vector in $\mathbb{C}^{|B|}$. That means that the basis states form a complex Hilbert space and it only takes $|B|$ complex numbers to define a quantum state. These numbers are called amplitudes. When $B = \{0, 1\}$, the resulting quantum system is called a qubit, and can be thought as the quantum analog of a bit, that can be 0 or 1 with some probability.

Bra-Ket Notation. States are represented as column vectors. To denote a column vector ϕ we write $|\phi\rangle$, using the "ket" notation. For the notation of row vectors, we denote as $\langle\phi|$ the conjugate transpose of ϕ, using the "bra" notation. The inner product of a row vector $|\phi\rangle$ and a column vector $\langle\psi|$ is denoted as $\langle\phi|\psi\rangle$, and is called a "braket". Bearing this notation in mind, we think of the basis states as column vectors, where

$|x\rangle$ assigns weight 1 to x and 0 everywhere else. The entirety of these states form an orthonormal basis called the computational basis, and any quantum state can be written as a complex combination of the classical states $|x\rangle$. Such combinations are called superpositions. For example, in case where $B = \{0, 1\}$, the computation basis consists of $|0\rangle = \begin{pmatrix}1\\0\end{pmatrix}$ and $|1\rangle = \begin{pmatrix}0\\1\end{pmatrix}$. Any general qubit $|\psi\rangle$ can be written as a superposition of the classical basis states

$$|\psi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle\,,$$

where $\alpha, \beta \in \mathbb{C}$ are the amplitudes of the qubit.

Composite Quantum Systems. Given that a qubit is defined in the Hilbert space $\mathbb{C}^2$, the Hilbert space of two qubits is the tensor product space $\mathbb{C}^2 \otimes \mathbb{C}^2$, with computational basis $\{|0\rangle \otimes |0\rangle\,, |0\rangle \otimes |1\rangle\,, |1\rangle \otimes |0\rangle\,, |1\rangle \otimes |1\rangle\}$. Note that usually in literature (as well as in the upcoming computations) the tensor product is implied without being written, such that $|0\rangle \otimes |0\rangle \equiv |0\rangle\,|0\rangle \equiv |00\rangle$. Thus, to describe a system with two qubits $|\phi_0\rangle = \alpha_0\,|0\rangle + \beta_0\,|1\rangle$ and $|\phi_1\rangle = \alpha_1\,|0\rangle + \beta_1\,|1\rangle$ we get

$$|\phi_0\rangle\,|\phi_1\rangle = \alpha_0\alpha_1\,|00\rangle + \alpha_0\beta_1\,|01\rangle + \beta_0\alpha_1\,|10\rangle + \beta_0\beta_1\,|11\rangle\,.$$

The above can be generalised for more qubits or quantum states with a different set B. Any state $|\phi\rangle$ consisting of more than one qubit and cannot be written as a tensor product is called an entangled state.

Evolution of a Quantum System. As we mentioned before, quantum states are L_2-normalized vectors. Hence, computations on quantum states must preserve the L_2 norm. As a parallel with the classical probability theory, we notice that there the L_1 norm is preserved, while any transformation is computed with a stochastic matrix, preserving said norm. Similarly, in quantum computations we use unitary matrices. A unitary matrix U is a matrix whose inverse equals his conjugate transpose $U^\dagger$, or $UU^\dagger = I$. We denote the result of a unitary transformation of a quantum state $|\phi\rangle$ as $|\phi'\rangle = U\,|\phi\rangle$.

Measurements. A (pure) quantum state $|\phi\rangle$ can be measured. After the measurement, with probability $|\,\langle x|\phi\rangle\,|^2$ we observe x and the state "collapses" to the classical state $|x\rangle$. Note that any subsequent measurements will always output x. For example, consider a qubit $|\phi\rangle = \alpha\,|0\rangle + \beta\,|1\rangle$. Upon measurement, with probability $|\alpha|^2$ we observe 0 and the state collapse to $|0\rangle$. Similarly, with probability $|\beta|^2$ we observe 1 and the state collapse to $|1\rangle$.

In a composite quantum system, we can perform partial measurements. First, let's consider we have an un-entangled quantum system, i.e. a quantum state $|\phi\rangle = |\phi_0\rangle\,|\phi_1\rangle$ where $|\phi_0\rangle$ and $|\phi_1\rangle$ have amplitudes α_0, β_0 and α_1, β_1 respectively. In this case, the qubits can be measured independently. Thus, with probability $|\alpha_0|^2$ we observe 0 and the state collapses to $|0\rangle\,|\phi_1\rangle$, and similarly with probability $|\beta|^2$ we observe 1 and the state collapses to $|1\rangle\,|\phi_1\rangle$. More interest resides in the case where the quantum state is entangled. Assume the general state $|\phi\rangle = \sum_{i,j} a_{ij}\,|ij\rangle \in \mathbb{C}^2 \otimes \mathbb{C}^2$ and we want to measure the first qubit. Then we observe i with probability $p_i = \sum_j |\alpha_{i,j}|^2$ and the state collapses to $\sum_j \frac{a_{i,j}}{\sqrt{p_i}}\,|i, j\rangle$. That means that we have to cross out the inconsistent terms

after the measurement, i.e. the terms that their first qubit is not the same as the result of the partial measurement. The division with $\sqrt{p_i}$ is necessary in order to renormalize the state and still preserve the L_2-norm.

Mixed States. A quantum system can be in a pure state $|\phi\rangle$ with soame probability, for example $1/2$, and in a different pure state $|\psi\rangle$ with probability $1/2$. This can take place in events like a partial measurement on a product system. This probability distribution cannot be correctly described by a pure state alone. This we say that the system is in a mixed state. The statistical behavior of a mixed state can be captured by a density matrix. If the system is in a pure state $|\phi_i\rangle$ with probability p_i, then the density matrix for the system is defined as $\rho = \sum_i p_i |\phi_i\rangle \langle\phi_i|$. The density matrix for a pure state is given by the rank-1 matrix $|\phi\rangle \langle\phi|$.

No cloning. In the classical setting, the most obvious and trivial computation we can do is copy a string of bits. This however, is not the case with quantum information. Specifically, there is no quantum procedure that transforms $|\phi\rangle \rightarrow |\phi\rangle |,\rangle \forall\phi$. This is one of the most important theorems in quantum computation. It introduces many limitations since, combined with the irreversibility induced by measurements, we are restricted on the computations we can perform. Intuitively, since we cannot measure many different copies of a quantum state, we cannot learn with certainty what an arbitrary quantum state is. On the other hand, the no cloning theorem can serve as a cryptographic guarantee, since an eavesdropper cannot copy quantum messages and any changes will be detected.

2 Quantum Cryptography

In general we notice that quantum computations differ vastly from classical ones. It is proven that for any classical function f, we can deterministically define a Unitary that efficiantly performs said function. However, the different nature of quantum information allows the development of algorithms that perform better than their classical counterparts. For example, Shor's algorithm solves integer factorization, a problem that in the classical setting requires sub-exponential time, with complexity $O\left((\log N)^2(\log\log N)(\log\log\log N)\right)$. Hence, many assumptions made in cryptography can be broken with quantum methods.

There are two different approaches to combat this problem. One is post-quantum cryptography, which includes classical protocols that are secure also against quantum adversaries. In order to achieve that level of security, quantum-secure assumptions must be used, such as the ones found in lattice-based cryptography. Another approach is quantum cryptography, in which the protocol itself is also quantum and we can take advantage of the quantum properties to prove security. The protocols constructed in this work are part of quantum cryptography.

Chapter 3

Preliminaries

We denote by λ the security parameter. A function $f : \mathbb{N} \to [0,1]$ is negligible if for every constant $c \in \mathbb{N}$ there exists $N \in \mathbb{N}$ such that for all $n > N$, $f(n) < n^{-c}$. We recall some standard notation for classical Turing machines and Boolean circuits:

- We say that a Turing machine (or algorithm) is PPT if it is probabilistic and runs in polynomial time in λ.
- We sometimes think about PPT Turing machines as polynomial-size uniform families of circuits. A polynomial-size circuit family C is a sequence of circuits $C = \{C_\lambda\}_{\lambda\in\mathbb{N}}$, such that each circuit C_λ is of polynomial size $\lambda^{O(1)}$ and has $\lambda^{O(1)}$ input and output bits. We say that the family is uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ.
- For a PPT Turing machine (algorithm) M, we denote by $M(x;r)$ the output of M on input x and random coins r. For such an algorithm, and any input x, we write $m \in M(x)$ to denote that m is in the support of $M(x;\cdot)$. Finally we write $y \leftarrow\!\!\$\, M(x)$ to denote the computation of M on input x with some uniformly sampled random coins.

1 Quantum Adversaries

We recall some notation for quantum computation and we define the notions of computational and statistical indistinguishability for quantum adversaries. Various parts of what follows are taken almost in verbatim from [BS20].

- We say that a Turing machine (or algorithm) is QPT if it is quantum and runs in polynomial time.
- We sometimes think about QPT Turing machines as polynomial-size uniform families of quantum circuits (as they are equivalent models). We call a polynomial-size quantum circuit family $C = \{C_\lambda\}_{\lambda\in\mathbb{N}}$ uniform if there exists a polynomial-time deterministic Turing machine M that on input 1^λ outputs C_λ.
- Classical communication channels in the quantum setting are identical to classical communication channels in the classical setting, except that when a set of qubits

is sent through a classical communication channel, then the qubits decohere and are automatically measured in the standard basis.

- A quantum interactive algorithm (in the two-party setting) has input divided into two registers and output divided into two registers. For the input qubits, one register is for an input message from the other party, and a second register is for a potential inner state the machine holds. For the output, one register is for the message to be sent to the other party, and another register is for a potential inner state for the machine to keep for itself.

Throughout this work, we model efficient adversaries as quantum circuits with non-uniform quantum advices. This is denoted by $\mathcal{A}^* = \{\mathcal{A}^*_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$, where $\{\mathcal{A}^*_\lambda\}_{\lambda\in\mathbb{N}}$ is a polynomial-size non-uniform sequence of quantum circuits, and $\{\rho_\lambda\}_{\lambda\in\mathbb{N}}$ is some polynomial-size sequence of mixed quantum states. We now define the formal notion of computational indistinguishability in the quantum setting.

Definition 1.1 (Computational Indistinguishability)**.** *Two ensembles of quantum random variables $\mathcal{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ are said to be computationally indistinguishable (denoted by $\mathcal{X} \approx_c \mathcal{Y}$) if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$ and all non-uniform QPT distinguishers with quantum advice $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$, it holds that*

$$|\Pr[\mathcal{A}(X;\rho) = 1] - \Pr[\mathcal{A}(Y;\rho) = 1]| \leq \mu(\lambda)$$

where $X \leftarrow\!\!\$\, X_\lambda$ and $Y \leftarrow\!\!\$\, Y_\lambda$.

The trace distance between two quantum distributions (X_λ, Y_λ), denoted by $\mathsf{TD}(X_\lambda, Y_\lambda)$, is a generalization of statistical distance to the quantum setting and represents the maximal distinguishing advantage between two quantum distributions by an unbounded quantum algorithm. We define below the notion of statistical indistinguishability.

Definition 1.2 (Statistical Indistinguishability)**.** *Two ensembles of quantum random variables $\mathcal{X} = \{X_\lambda\}_{\lambda\in\mathbb{N}}$ and $\mathcal{Y} = \{Y_\lambda\}_{\lambda\in\mathbb{N}}$ are said to be statistically indistinguishable (denoted by $\mathcal{X} \approx_s \mathcal{Y}$) if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$, it holds that*

$$\mathsf{TD}(X_\lambda, Y_\lambda) \leq \mu(\lambda).$$

The Class QMA. A language $\mathcal{L} = (\mathcal{L}_{\text{yes}}, \mathcal{L}_{\text{no}})$ in QMA is defined by a tuple $(\mathcal{V}, p, \alpha, \beta)$, where p is a polynomial, $\mathcal{V} = \{V_\lambda\}_{\lambda\in\mathbb{N}}$ is a uniformly generated family of circuits such that for every λ, V_λ takes as input a string $x \in \{0,1\}^\lambda$ and a quantum state $|\psi\rangle$ on $p(\lambda)$ qubits and returns a single bit, and $\alpha, \beta : \mathbb{N} \to [0,1]$ are such that $\alpha(\lambda) - \beta(\lambda) \geq 1/p(\lambda)$. The language is then defined as follows.

- For all $x \in \mathcal{L}_{\text{yes}}$ of length λ, there exists a quantum state $|\psi\rangle$ of size at most $p(\lambda)$ such that the probability that V_λ accepts $(x, |\psi\rangle)$ is at least $\alpha(\lambda)$. We denote the (possibly infinite) set of quantum witnesses that make V_λ accept x by $\mathsf{R}_{\mathcal{L}}(x)$.

- For all $x \in \mathcal{L}_{\text{no}}$ of length λ, and all quantum states $|\psi\rangle$ of size at most $p(\lambda)$, it holds that V_λ accepts on input $(x, |\psi\rangle)$ with probability at most $\beta(\lambda)$.

2 Learning with Errors

We recall the definition of the learning with errors (LWE) problem [Reg05].

Definition 2.1 (Learning with Errors)**.** *The LWE problem is parametrized by a modulus $q = q(\lambda)$, polynomials $n = n(\lambda)$ and $m = m(\lambda)$, and an error distribution χ. The LWE problem is hard if it holds that*

$$(\mathbf{A}, \mathbf{A} \cdot \mathbf{s} + \mathbf{e}) \approx_c (\mathbf{A}, \mathbf{u})$$

where $\mathbf{A} \leftarrow\!\!\$\, \mathbb{Z}_q^{m \times n}$, $\mathbf{s} \leftarrow\!\!\$\, \mathbb{Z}_q^n$, $\mathbf{u} \leftarrow\!\!\$\, \mathbb{Z}_q^m$, and $\mathbf{e} \leftarrow\!\!\$\, \chi^m$.

As shown in [Reg05, PRS17], for any sufficiently large modulus q the LWE problem where χ is a discrete Gaussian distribution with parameter $\sigma = \xi q \geq 2\sqrt{n}$ (i.e. the distribution over $\mathbb{Z}$ where the probability of x is proportional to $e^{-\pi(|x|/\sigma)^2}$), is at least as hard as approximating the shortest independent vector problem (SIVP) to within a factor of $\gamma = \tilde{O}(n/\xi)$ in *worst case* dimension n lattices.

3 Pseudorandom Functions

We recall the standard notion of pseudorandom function (PRF) [GGM86].

Definition 3.1 (Pseudorandom Function)**.** *A pseudorandom function* $(\mathsf{PRF.Gen}, \mathsf{PRF.Eval})$ *consists of the following efficient algorithms.*

- $\mathsf{PRF.Gen}(1^\lambda)$*: On input the security parameter, the key generation algorithm returns a key k.*
- $\mathsf{PRF.Eval}(k, x)$*: On input a key k and a string $x \in \{0,1\}^\lambda$, the evaluation algorithm returns a string $y \in \{0,1\}^{e(\lambda)}$.*

The scheme must be pseudorandom in the following sense.

Definition 3.2 (Pseudorandomness)**.** *A pseudorandom function* $(\mathsf{PRF.Gen}, \mathsf{PRF.Eval})$ *is pseudorandom if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$ and all non-uniform QPT distinguishers with quantum advice $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$, it holds that*

$$\left|\Pr\left[\mathcal{A}(\rho)^{\mathsf{PRF.Eval}(k,\cdot)} = 1\right] - \Pr\left[\mathcal{A}(\rho)^{f(\cdot)} = 1\right]\right| \leq \mu(\lambda)$$

where $k \leftarrow\!\!\$\, \mathsf{PRF.Gen}(1^\lambda)$ and $f : \{0,1\}^\lambda \to \{0,1\}^{e(\lambda)}$ is a uniformly sampled truly random function.

4 Garbled Circuits

We recall the definition of a garbling scheme for circuits [Yao86, AIK04, BHR12].

Definition 4.1 (Garbled Circuit)**.** *A garbling scheme for circuits is a tuple of PPT algorithms* $(\mathsf{Garble}, \mathsf{GEval})$ *with the following syntax.*

- $\mathsf{Garble}\left(1^\lambda, C\right)$: Garble *takes as input a security parameter* 1^λ*, a circuit* C*, and outputs a* garbled circuit $\widetilde{C}$ *along with labels* $\{\ell_{i,b}\}_{i\in\{1,\ldots,n\},b\in\{0,1\}}$*, where* n *is the length of the input to* C*.*
- $\mathsf{GEval}\left(\widetilde{C}, \{\ell_{i,x_i}\}_{i\in\{1,\ldots,n\}}\right)$: *Given a garbled circuit* $\widetilde{C}$ *and a sequence of input labels* $\{\ell_{i,x_i}\}_{i\in\{1,\ldots,n\}}$*,* GEval *outputs a string* y*.*

We recall the notion of completeness.

Definition 4.2 (Completeness)**.** *A garbling scheme* $(\mathsf{Garble}, \mathsf{GEval})$ *is complete if for any circuit* C *and input* $x \in \{0,1\}^n$ *we have that:*

$$\Pr\left[C(x) = \mathsf{GEval}\left(\widetilde{C}, \{\ell_{i,x_i}\}_{i\in\{1,\ldots,n\}}\right)\right] = 1$$

where $(\widetilde{C}, \{\ell_{i,b}\}_{i\in\{1,\ldots,n\},b\in\{0,1\}}) \leftarrow \mathsf{Garble}\left(1^\lambda, C\right)$*.*

We define the notion of (statistical) simulation security, which is achievable for circuits in NC1.

Definition 4.3 (Security)**.** *A garbling scheme* $(\mathsf{Garble}, \mathsf{GEval})$ *is simulation secure if there exists a PPT simulator* GSim *such that for any circuit* C *and input* $x \in \{0,1\}^n$*, we have that*

$$\left(\widetilde{C}, \{\ell_{i,x_i}\}_{i\in\{1,\ldots,n\}}\right) \approx_s \mathsf{GSim}\left(1^\lambda, 1^{|C|}, 1^n, C(x)\right)$$

where $(\widetilde{C}, \{\ell_{i,b}\}_{i\in\{1,\ldots,n\},b\in\{0,1\}}) \leftarrow \mathsf{Garble}\left(1^\lambda, C\right)$*.*

5 Interactive Proofs and Sigma Protocols

We present the definitions of interactive proof systems and sigma protocols. Much of the following material is taken in verbatim from [Shm20]. We denote by (P, V) and interactive protocol between a prover P and a verifier V. The output of the verifier is denoted by $\mathsf{Out}(\mathsf{P}, \mathsf{V})$. For an honest verifier, the output is a classical bit that denotes acceptance or rejection. If the verifier is corrupted, the output can be an arbitrary quantum state. We define completeness in the following.

Definition 5.1 (Completeness)**.** *An interactive protocol* (P, V) *for a language* $\mathcal{L} \in \mathsf{QMA}$ *with relation* $\mathsf{R}_\mathcal{L}$ *is complete if there exists a polynomial* p *and a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$*, all* $x \in \mathcal{L}$*, and all* $|w\rangle \in \mathsf{R}_\mathcal{L}(x)$*, it holds that*

$$\Pr\left[\mathsf{Out}(\mathsf{P}(|w\rangle^{\otimes p(\lambda)}, x), \mathsf{V}(x)) = 1\right] \geq 1 - \mu(\lambda).$$

Next we define the notion of (non-adaptive) computational soundness.

Definition 5.2 (Computational Soundness)**.** *An interactive protocol* (P, V) *for a language* $\mathcal{L} \in \mathsf{QMA}$ *with relation* $\mathsf{R}_\mathcal{L}$ *is computationally sound if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$*, all* $x \notin \mathcal{L}$*, and all non-uniform QPT provers with quantum advice* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$*, it holds that*

$$\Pr\left[\mathsf{Out}(\mathcal{A}(x;\rho), \mathsf{V}(x)) = 1\right] \leq \mu(\lambda).$$

Sigma Protocols. We explicitly define sigma protocols (Σ), a special case of interactive protocols for QMA, and we define a special-zero knowledge guarantee that is satisfied by some protocols of interest.

Definition 5.3 (Sigma Protocol)**.** *A sigma protocol* $(\Sigma.\mathsf{Com}, \Sigma.\mathsf{Chal}, \Sigma.\mathsf{Resp})$ *consists of the following efficient algorithms.*

- $\Sigma.\mathsf{Com}(|w\rangle^{\otimes p(\lambda)}; r)$*: On input* $p(\lambda)$*-many copies of the witness and some (classical) random coins* $r \in \{0,1\}^{q(\lambda)}$*, the commitment algorithm returns a first commitment* $|\alpha\rangle$*.*
- $\Sigma.\mathsf{Chal}(x)$*: On input the instance* x*, the challenge algorithm returns a uniformly sampled (classical) string* $\beta \in \{0,1\}^{b(\lambda)}$*.*
- $\Sigma.\mathsf{Resp}(\beta, r)$*: On input the challenge* β *and the classical random coins* r*, the response algorithm returns a classical response* γ*.*

We highlight the fact that both the challenge and the response algorithm are completely classical: The only quantum computation needed is for the $\Sigma.\mathsf{Com}$ algorithm and for verifying that $x \in \mathcal{L}$, given the protocol transcript. We now define the notion of computational special zero-knowledge.

Definition 5.4 (Computational Special Zero-Knowledge)**.** *A sigma protocol* $(\Sigma.\mathsf{Com}, \Sigma.\mathsf{Chal}, \Sigma.\mathsf{Resp})$ *satisfies (computational) special zero-knowledge if there exists a QPT simulator* $\Sigma.\mathsf{Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all* $x \in \mathcal{L}$*, and all* $|w\rangle \in \mathsf{R}_{\mathcal{L}}(x)$*, it holds that*

$$(\Sigma.\mathsf{Com}(|w\rangle^{\otimes p(\lambda)}; r), \Sigma.\mathsf{Resp}(\beta, r)) \approx_c \Sigma.\mathsf{Sim}(x, \beta)$$

where $r \leftarrow\!\!\$\, \{0,1\}^{q(\lambda)}$ *and* $\beta \leftarrow\!\!\$\, \{0,1\}^{b(\lambda)}$*.*

The statistical notion is defined analogously, except that we require statistical indistinguishability between the two distributions. It was recently shown by Broadbent and Grilo [BG20] how to obtain a sigma protocol for QMA satisfying statistical soundness and special zero-knowledge, assuming a (classical) post-quantum non-interactive statistically binding bit commitment scheme [LS19, HW18]. Here we restate the main theorem of such a work.

Lemma 5.5 ([BG20])**.** *Assuming the post-quantum hardness of the LWE problem, there exists a sigma protocol* $(\Sigma.\mathsf{Com}, \Sigma.\mathsf{Chal}, \Sigma.\mathsf{Resp})$ *satisfying statistical soundness and computational special zero-knowledge.*

In this work we are also interested in the reverse guarantees, i.e. computational soundness and stastistical zero-knowledge. In [BG20], it was shown that instantiating the same protocol with a statistically hiding commitment results in a sigma protocol with the desired properties. However, (classical) statistically hiding commitments notoriously require two rounds of interaction and thus one needs to extend the syntax of the sigma protocol to have the verifier sampling the commitment key $\mathsf{ck} \leftarrow\!\!\$\, \Sigma.\mathsf{Gen}(1^\lambda)$, which is also given as an input to the $\Sigma.\mathsf{Com}$ algorithm. The definition of special zero-knowledge is extended accordingly. Since statistically hiding commitments can be constructed from any collision resistant hash function [HM96] (and in particular assuming LWE), we obtain the following implication.

Lemma 5.6 ([BG20])**.** *Assuming the post-quantum hardness of the LWE problem, there exists a sigma protocol* $(\Sigma.\mathsf{Gen}, \Sigma.\mathsf{Com}, \Sigma.\mathsf{Chal}, \Sigma.\mathsf{Resp})$ *satisfying computational soundness and statistical special zero-knowledge.*

6 Statistical ZAPs for NP

A ZAP protocol is a two-round witness-indistinguishable argument where the first message is instance-independent. We say that the protocol achieves *multi-theorem* security if the first round can be fixed once and for all and can be reused for an unbounded amount of second rounds. In the other hand, if the first round has to be re-initialized for each run of the protocol, we say that the ZAP achieves only *single-theorem* security. Additionally, we say that the protocol is *public coin* if the output of the protocol is publicly computable given the protocol transcript, and otherwise we say that the protocol is *private coin*. We being by defining the syntax of (public coin) statistical ZAPs for NP.

Definition 6.1 (ZAP Protocol for NP)**.** *A ZAP protocol* $(\mathsf{ZAP.Setup}, \mathsf{ZAP.Prove}, \mathsf{ZAP.Verify})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *consists of the following efficient algorithms.*

- $\mathsf{ZAP.Setup}(1^\lambda)$*: On input the security parameter* 1^λ*, the setup returns a common reference string* crs *and a trapdoor* td*.*
- $\mathsf{ZAP.Prove}(\mathsf{crs}, w, x)$*: On input a common reference string* crs*, a witness* w*, and a statement* x*, the proving algorithm returns a proof* π*.*
- $\mathsf{ZAP.Verify}(\mathsf{td}, \pi, x)$*: On input a trapdoor* td*, a proof* π*, and a statement* x*, the verification algorithm returns a bit* $\{0, 1\}$*.*

The definitions of completeness and computational soundness are identical to those given for general interactive proof systems (Section 5). Note that all definitions that we present here are for the single-theorem case. This is without loss of generality, since single-theorem soundness (witness indistinguishability, resp.) is equivalent to single-theorem soundness (witness indistinguishability, resp.) for public coin protocols. In the following we present the notion of (statistical) witness indistinguishability.

Definition 6.2 (Statistical Witness Indistinguishability)**.** *A ZAP protocol* $(\mathsf{ZAP.Setup}, \mathsf{ZAP.Prove}, \mathsf{ZAP.Verify})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is witness indistinguishable if for all* $\lambda \in \mathbb{N}$*, all* $x \in \mathcal{L}$*, all pairs of witnesses* $(w_0, w_1) \in \mathsf{R}_{\mathcal{L}}$*, and all common reference strings* crs *it holds that*

$$(\mathsf{crs}, \mathsf{ZAP.Prove}(\mathsf{crs}, w_0, x)) \approx_s (\mathsf{crs}, \mathsf{ZAP.Prove}(\mathsf{crs}, w_1, x)).$$

It was recently shown in [BFJ+20, GJJM20] that statistical ZAPs for NP exist assuming the quasi-polynomial (quantum) hardness of the LWE problem.

Lemma 6.3 ([BFJ+20, GJJM20])**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, there exists a public coin ZAP for* NP $(\mathsf{ZAP.Setup}, \mathsf{ZAP.Prove}, \mathsf{ZAP.Verify})$*.*

7 Sometimes-Binding Statistically Hiding Commitments

We introduce the notion of sometimes-binding statistically hiding (SBSH) commitments, as defined in [LVW20].

Definition 7.1 (SBSH Commitment)**.** *An SBSH commitment scheme* $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$ *consists of the following efficient algorithms.*

- $\mathsf{SBSH.Gen}(1^\lambda)$*: On input the security parameter* 1^λ*, the generation algorithm returns a partial commitment key* ck_0*.*
- $\mathsf{SBSH.Key}(\mathsf{ck}_0)$*: On input a partial key* ck_0*, the key agreement algorithm returns the complement of the key* ck_1*.*
- $\mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), m)$*: On input a commitment key* $(\mathsf{ck}_0, \mathsf{ck}_1)$ *and a message* m*, the commitment algorithm returns a partial commitment key* ck_1 *and a commitment* c*.*

The commitment must satisfy the notion of statistical hiding.

Definition 7.2 (Statistical Hiding)**.** *An SBSH commitment scheme* $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$ *is statistically hiding if for all* $\lambda \in \mathbb{N}$*, all partial keys* ck_0*, and all pairs of messages* (m_0, m_1)*, it holds that*

$$(\mathsf{ck}_0, \mathsf{ck}_1, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), m_0)) \approx_s (\mathsf{ck}_0, \mathsf{ck}_1, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), m_1))$$

where $\mathsf{ck}_1 \leftarrow\!\!\$\, \mathsf{SBSH.Key}(\mathsf{ck}_0)$*.*

Next we define the notion of sometimes-binding for an SBSH commitment scheme. We define the set $\mathsf{Binding}$ as the set of all commitment keys $(\mathsf{ck}_0, \mathsf{ck}_1)$ such that the any resulting commitment is perfectly binding. We present the definition of the property in the following.

Definition 7.3 (Sometimes Binding)**.** *An SBSH commitment scheme* $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$ *is* (ε, δ)*-sometimes binding if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$ *and all (stateful) QPT distinguishers* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$*, it holds that*

$$\Pr\left[\mathcal{A}(\mathsf{st}; \rho) = 1 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] = \varepsilon(\lambda) \cdot \Pr\left[\mathcal{A}(\mathsf{st}; \rho) = 1\right] + \delta(\lambda) \cdot \mu(\lambda)$$

where $\mathsf{ck}_0 \leftarrow\!\!\$\, \mathsf{SBSH.Gen}(1^\lambda)$ *and* $(\mathsf{st}, \mathsf{ck}_1) = \mathcal{A}(\mathsf{ck}_0; \rho)$*.*

We also require the existence of a polynomial-time extractor $\mathsf{SBSH.Ext}$ that, on input the random coins r used in the $\mathsf{SBSH.Gen}$ algorithm, extracts the committed message m from the protocol transcript if $(\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}$. The works of [KKS18, BFJ+20, GJJM20] construct SBSH commitment schemes (using a slightly different syntax) for quasi-polynomial (ε, δ) assuming the quasi-polynomial hardness of two-round statistically sender private oblivious transfer. Thus we can state the following lemma.

Lemma 7.4 ([BFJ+20, GJJM20])**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, there exists an* (ε, δ)*-sometimes binding SBSH commitment scheme* $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$*.*

8 Sometimes-Simulatable Zero-Knowledge

Here we introduce the notion of sometimes simulatability (SSim-ZK) as defined in [CM21]. This can be thought as the straight-line equivalent of super-polynomial simulation [Pas03b] and it is formally defined in the following.

Definition 8.1 (Sometimes Simulatability)**.** *An interactive protocol* (P, V) *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is* (ε, δ)*-sometimes simulatable if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$ *and all (stateful) QPT distinguishers* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$*, it holds that*

$$\Pr\left[\mathcal{A}(\mathsf{st};\rho) = 1 \wedge (\mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] = \varepsilon(\lambda) \cdot \Pr\left[\mathcal{A}(\mathsf{st};\rho) = 1\right] + \delta(\lambda) \cdot \mu(\lambda)$$

where $(\mathsf{zk}_0, \mathsf{zk}_1)$ *are the first two messages of the protocol and* $\mathsf{Simulation}$ *defines a set. Furthermore, we require the existence of a polynomial-time algorithm* Sim *such that, conditioned on the event* $(\mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}$*, it holds that*

$$\mathsf{Sim}(1^\lambda, r, x) \approx_c \mathsf{zk}_2$$

where r *are the random coins* r *used to compute* zk_0 *and* zk_2 *is the honestly computed third message.*

9 Compute-and-Compare Obfuscation

Here we define compute-and-compare circuits (CC) and obfuscators for said circuits (Obf). The definitions are taken in verbatim from [BS20].

Definition 9.1 (Compute-and-Compare Circuit)**.** *Let* $f : \{0,1\}^n \to \{0,1\}^\lambda$ *and let* $u \in \{0,1\}$*,* $z \in \{0,1\}^*$ *be strings. Then* $\mathbf{CC}[f, u, z](x)$ *is a circuit that returns* z *if* $f(x) = y$*, and* $\perp$ *otherwise.* $\mathbf{CC}$ *has a canonical description from which* f, u *and* z *can be read.*

For the following definition, Obf is a PPT algorithm that takes as input a CC circuit and outputs a new circuit $\widetilde{\mathbf{CC}}$.

Definition 9.2 (Correctness)**.** *A PPT algorithm* Obf *is a correct compute-and-compare obfuscator if for any circuit* $f : \{0,1\}^n \to \{0,1\}^\lambda$*,* $u \in \{0,1\}$*,* $z \in \{0,1\}^*$

$$\Pr\left[\forall x \in \{0,1\}^n : \widetilde{\mathbf{CC}}(x) = \mathbf{CC}[f,u,z](x) \;\middle|\; \widetilde{\mathbf{CC}} \leftarrow \mathsf{Obf}(\mathbf{CC}[f,u,z])\right] = 1$$

We define simulation security in the following.

Definition 9.3 (Simulation Security)**.** *A PPT algorithm* Obf *is a simulation secure compute-and-compare obfuscator if there exists a PPT Simulator* Sim *such for every two polynomials* $\ell_1(\cdot), \ell_2(\cdot)$*,*

$$\left\{\widetilde{\mathbf{CC}} \mid u \leftarrow \{0,1\}^\lambda, \widetilde{\mathbf{CC}} \leftarrow \mathsf{Obf}(\mathbf{CC})\right\}_{\lambda,f,z} \approx_c \left\{\mathsf{Sim}(1^{\ell_1(\lambda)}, 1^{\ell_2(\lambda)}, 1^\lambda)\right\}_{\lambda,f,z}$$

where $\lambda \in \mathbb{N}$*,* $f : \{0,1\}^\lambda$ *is a* $\ell_1(\lambda)$*-size circuit and* $z \in \{0,1\}^{\ell_2(\lambda)}$*.*

Constructions based on the quantum hardness of LWE can be found in [GKW17, WZ17, GKVW19].

10 Quantum One-Time Pad

We recall the quantum one-time pad (QOTP) construction [AMTDW00] for quantum states. We explicitly consider the scheme that allows one to encrypt an n-qubit quantum state with unconditional security.

Definition 10.1 (Quantum One-Time Pad)**.** *A quantum one-time pad* ($\mathsf{QOTP.Gen}$, $\mathsf{QOTP.Enc}, \mathsf{QOTP.Dec}$) *consists of the following efficient algorithms.*

- $\mathsf{QOTP.Gen}(1^n)$*: For all* $i = 1 \dots n$ *sample two classical bits* $(x_i, z_i) \leftarrow\!\!\$\, \{0,1\}^2$*. Return the one-time key* $\mathsf{otk} = (x_1, z_1, \dots, x_n, z_n)$*.*
- $\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$*: On input a one-time key* otk *and an* n*-qubit state* $|\psi\rangle$*, apply the Pauli transformation* $X^{x_i} Z^{z_i}$ *to the* i*-th qubit, for all* $i = 1 \dots n$*. Return the resulting state* $|\phi\rangle$*.*
- $\mathsf{QOTP.Dec}(\mathsf{otk}, |\phi\rangle)$*: On input a one-time key* otk *and an* n*-qubit state* $|\phi\rangle$*, apply the reverse Pauli transformation* $Z^{z_i} X^{x_i}$ *qubit-by-qubit to recover the original state.*

More explicitly, the (single qubit) Pauli transformation $X^{x_i} Z^{z_i}$ is the following unitary:

$$(\alpha_0 |0\rangle + \alpha_1 |1\rangle) \rightarrow (\alpha_0 |x_i\rangle + (-1)^{z_i} \alpha_1 |x_i \oplus 1\rangle).$$

As shown in [AMTDW00], the above scheme can be used to transform *any* n-qubit quantum state into a totally mixed state (no matter if some of its initial qubits are in an entangled state).

11 Pauli Operators

The Pauli Operators X, Y, Z are 2×2 matrices that are unitary and Hermitian. More specifically:

$$X = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} \; Y = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} \; Z = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}.$$

12 Homomorphic Encryption

In the following we define the main object of interest of our work, namely homomorphic encryption that allows one to evaluate classical and/or quantum circuits over encrypted data.

Classical Homomorphic Encryption

We recall the notion of classical homomorphic encryption [Gen09].

Definition 12.1 (Homomorphic Encryption)**.** *A homomorphic encryption scheme* ($\mathsf{FHE.Gen}$, $\mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec}$) *consists of the following efficient algorithms.*

- $\mathsf{FHE.Gen}(1^\lambda)$*: On input the security parameter, the key generation algorithm returns secret/public key pair* $(\mathsf{sk}, \mathsf{pk})$*.*

- $\mathsf{FHE.Enc}(\mathsf{pk}, m)$*: On input the public key* pk *and a message* m*, the encryption algorithm returns a ciphertext* c*.*
- $\mathsf{FHE.Eval}(\mathsf{pk}, C, c)$*: On input the public key* pk*, a (classical) circuit* C*, and a ciphertext* c*, the evaluation algorithm returns an evaluated ciphertext* $\tilde{c}$*.*
- $\mathsf{FHE.Dec}(\mathsf{sk}, c)$*: On input the secret key* sk *and a ciphertext* c*, the decryption algorithm returns a message* m*.*

We say that a scheme is fully homomorphic (FHE) if the evaluation algorithm supports all polynomial-size classical circuits (without posing an a-priori bound on the size of $|C|$). If the size of C needs to be fixed at the time of key generation, then we say that the scheme is levelled homomorphic. It is well-known that levelled FHE schemes can be based on the hardness of the (plain) LWE problem [BV11, BV14]. We recall the notion of single-hop evaluation correctness in the following and we refer the reader to [GHV10] for a more general definition of multi-hop evaluation correctness.

Definition 12.2 (Single-Hop Evaluation Correctness)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is correct if for all* $\lambda \in \mathbb{N}$*, all* $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{FHE.Gen}(1^\lambda)$*, all messages* m*, and all polynomial-size circuits* C*, it holds that*

$$\Pr\left[\mathsf{FHE.Dec}(\mathsf{sk}, \mathsf{FHE.Eval}(\mathsf{pk}, C, \mathsf{FHE.Enc}(\mathsf{pk}, m))) = C(m)\right] = 1$$

We recall the notion of semantic security for public-key encryption.

Definition 12.3 (Semantic Security)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is semantically secure if for all* $\lambda \in \mathbb{N}$ *and all pairs of messages* (m_0, m_1)*, it holds that*

$$\mathsf{FHE.Enc}(\mathsf{pk}, m_0) \approx_c \mathsf{FHE.Enc}(\mathsf{pk}, m_1)$$

where $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\!\$\, \mathsf{FHE.Gen}(1^\lambda)$*.*

Finally we define the notion of (malicious) statistical circuit privacy for FHE [OPP14].

Definition 12.4 (Statistical Circuit Privacy)**.** *A homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *is (malicious) statistically circuit private if there exists a pair of unbounded algorithms* $\mathsf{FHE.Ext}$ *and* $\mathsf{FHE.Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all public keys* pk^**, all ciphertexts* c^**, and all circuits* C*, it holds that*

$$\mathsf{FHE.Eval}(\mathsf{pk}^*, C, c^*) \approx_s \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}^*, c^*, C(x^*))$$

where $x^* = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}^*, c^*)$*.*

It is shown in [OPP14] that any FHE scheme can be converted into one with malicious circuit privacy generically, by additionally assuming a two-round statistically sender-private oblivious transfer. The latter can in turn be instantiated from LWE [BD18, DGI+19, BDGM19]. Taken together, these results give us the following implication.

Lemma 12.5 ([OPP14, BD18])**.** *Assuming the hardness of the circular LWE problem, there exists an FHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *with (malicious) statistical circuit privacy.*

Quantum Homomorphic Encryption

We extend the notion of classical FHE to the evaluation of quantum circuits [BJ15]. In this work we consider only quantum FHE (QFHE) schemes with completely classical key generation algorithms. We extend the syntax of classical FHE below.

Definition 12.6 (Quantum Homomorphic Encryption)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *consists of the following efficient algorithms.*

- $\mathsf{FHE.Gen}(1^\lambda)$*: Same as in Definition 12.1.*
- $\mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle)$*: On input the public key* pk *and a quantum state* $|\psi\rangle$*, the encryption algorithm returns a quantum ciphertext* $|\phi\rangle$*.*
- $\mathsf{FHE.QEval}(\mathsf{pk}, C, |\phi\rangle)$*: On input the public key* pk*, a quantum circuit* C*, and a quantum ciphertext* $|\phi\rangle$*, the evaluation algorithm returns an evaluated quantum ciphertext* $|\tilde{\phi}\rangle$*.*
- $\mathsf{FHE.QDec}(\mathsf{sk}, |\phi\rangle)$*: On input the secret key* sk *and a quantum ciphertext* $|\phi\rangle$*, the decryption algorithm returns a quantum state* $|\psi\rangle$*.*

Analogously to the classical case, we say that the scheme is fully homomorphic if the evaluation algorithm supports all polynomial-size quantum circuits. Next we define the notion of single-hop evaluation correctness for QFHE.

Definition 12.7 (Single-Hop Evaluation Correctness)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *is correct if for all* $\lambda \in \mathbb{N}$*, all* $(\mathsf{sk}, \mathsf{pk}) \in \mathsf{FHE.Gen}(1^\lambda)$*, all quantum states* $|\psi\rangle$*, and all polynomial-size quantum circuits* C*, it holds that*

$$\mathsf{FHE.QDec}(\mathsf{sk}, \mathsf{FHE.QEval}(\mathsf{pk}, C, \mathsf{FHE.QEnc}(\mathsf{pk}, |\psi\rangle))) \approx_s C(|\psi\rangle).$$

The notion of semantic security is defined analogously to the classical case, and we refer the reader to [BJ15] for a formal definition. We define the main notion of interest of this work, namely, malicious statistical circuit privacy for QFHE.

Definition 12.8 (Statistical Circuit Privacy)**.** *A quantum homomorphic encryption scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *is (malicious) statistically circuit private if there exists a pair of unbounded algorithms* $\mathsf{FHE.Ext}$ *and* $\mathsf{FHE.Sim}$ *such that for all* $\lambda \in \mathbb{N}$*, all public keys* pk^**, all quantum ciphertexts* $|\phi^*\rangle$*, and all quantum circuits* C*, it holds that*

$$\mathsf{FHE.QEval}(\mathsf{pk}^*, C, |\phi^*\rangle) \approx_s \mathsf{FHE.Sim}(1^\lambda, \mathsf{pk}^*, \alpha, C(|\psi^*\rangle))$$

where $(|\psi^*\rangle, \alpha) = \mathsf{FHE.Ext}(1^\lambda, \mathsf{pk}^*, |\phi^*\rangle)$*.*

Chapter 4

Rate-1 Quantum Fully Homomorphic Encryption

In the following we construct a QFHE scheme with rate approaching 1, as the security parameter (and consequently the message space) grows.

1 Definition

We begin by formally defining the notion of rate for a quantum homomorphic encryption scheme.

Definition 1.1 (Rate)**.** *We say that a quantum homomorphic encryption scheme* $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ *has rate* $\rho = \rho(\lambda)$*, if for all* pk *in the support of* $\mathsf{QFHE.Gen}(1^\lambda)$*, all supported quantum circuits* C *with sufficiently large output size, all polynomials* $\ell = \ell(\lambda)$*, all* ℓ*-qubit quantum states* $|\psi\rangle$*, and all states* $|\phi\rangle$ *where* $|\phi\rangle \in \mathsf{QFHE.QEnc}(\mathsf{pk}, |\psi\rangle)$*, it holds that*

$$\frac{|C(|\psi\rangle)|}{|\mathsf{QFHE.QEval}\,(pk, C, |\phi\rangle)|} \geq \rho$$

where $|\cdot|$ *is the size in qubits for quantum information and bits for classical information. We also say that a scheme has rate 1, if it holds that*

$$\lim_{\lambda\to\infty} \rho(\lambda) = 1$$

The notation $|\cdot|$ generally corresponds to the size of the input. In the classical setting, this translates to the number of bits that the information consists of. Similarly, in the quantum setting, we can extend the definition and measure the size in the basic unit of quantum information, a qubit. For constructing rate-1 QFHE schemes, it is convenient to define an additional ciphertext compression algorithm, together with a corresponding compressed decryption algorithm. The following are definitions from [BDGM19], extended to the quantum setting.

Definition 1.2 (Compression)**.** *Let* $\mathsf{QFHE} = (\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ *be a QFHE scheme and let* $\ell = \ell(\lambda)$ *be a polynomial. We say that* QFHE *supports* ℓ*-qubits ciphertext compression if there exist two algorithms* $\mathsf{QFHE.Compress}$ *and* $\mathsf{QFHE.CompressDec}$ *with the following syntax:*

- $\mathsf{QFHE.Compress}(\mathsf{pk}, |\phi\rangle)$*: Takes as input a public key* pk *and an encrypted* ℓ*-qubit state* $|\phi\rangle$ *and outputs a compressed ciphertext* $|\phi^*\rangle$*.*
- $\mathsf{QFHE.CompressDec}(\mathsf{sk}, |\phi^*\rangle)$*: Takes as input a secret key* sk *and a compressed ciphertext* $|\phi^*\rangle$ *and outputs an* ℓ*-qubit state* $|\psi\rangle$*.*

We require the following notion of correctness to hold for compressed ciphertexts.

Definition 1.3 (Compressed Correctness)**.** *A quantum homomorphic encryption scheme* $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec}, \mathsf{QFHE.Compress}, \mathsf{QFHE.CompressDec})$ *satisfies compressed correctness if for all* $\lambda \in \mathbb{N}$*, all* $\ell = \ell(\lambda)$*, all* $(\mathsf{sk}, \mathsf{pk})$ *in the support of* $\mathsf{FHE.Gen}(1^\lambda)$*, all* ℓ*-qubit quantum states* $|\psi\rangle$*, all* $|\phi\rangle$ *such that* $|\psi\rangle = \mathsf{QDec}(\mathsf{sk}, |\phi\rangle)$*, it holds that*

$$\mathsf{CompressDec}\left(\mathsf{sk}, \mathsf{Compress}(\mathsf{pk}, |\phi\rangle)\right) = |\psi\rangle\,.$$

The definition of rate is unchanged, except that we consider the size of compressed ciphertexts. For the case of classical FHE, it was recently shown by Brakerski et al. [BDGM19] that a leveled scheme with rate-1 exists under the standard LWE assumption (with polynomial modulo-to-noise ratio), which can be converted to fully homomorphic by an additional circularity assumption. The scheme satisfies an additional structural property that we call *spooky decryption* and we formally define below.

Lemma 1.4 ([BDGM19])**.** *Assuming the hardness of the circular LWE problem, there exists a rate-1 FHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ *and a function* F *such that for all ciphertexts* $c = (\mathbf{c}_0, c_1, \ldots, c_k) \in \mathbb{Z}_q^{n+1} \times \{0,1\}^k$ *it holds that*

$$\mathsf{FHE.Dec}(\mathsf{sk}, c) = F(\mathsf{sk}, \mathbf{c}_0) \oplus (c_1, \ldots, c_k).$$

Finally, the works of Mahadev [Mah18a] and Brakerski [Bra18] show that QFHE with classical keys can be constructed from the quantum hardness of the LWE problem. For the evaluation of unbounded circuits, an additional circularity assumption is required due to an application of the bootstrapping theorem [Gen09]. Both schemes follow the *hybrid encryption* approach where each ciphertext consists of (i) a QOTP of a given quantum state and (ii) a (classical) FHE encryption of the corresponding one-time key. This is captured by the following Lemma.

Lemma 1.5 ([Mah18a, Bra18])**.** *Assuming the quantum hardness of the circular LWE problem, there exists a QFHE scheme* $(\mathsf{FHE.Gen}, \mathsf{FHE.QEnc}, \mathsf{FHE.QEval}, \mathsf{FHE.QDec})$ *where (evaluated) ciphetexts are of the form*

$$\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{otk})$$

where $\mathsf{FHE.Enc}$ *is the encryption algorithm of a classical semi-honest circuit-private FHE scheme.*

2 Our Construction

Our scheme is again described as a generic transformation, assuming the existence of the following primitives:

- A rate-1 classical FHE scheme ($\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec}$) with spooky decryption (see Lemma 1.4).
- A quantum fully homomorphic encryption scheme ($\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec}$) with classical keys and hybrid ciphertexts of the form ($\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$, $\mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk})$) (see Lemma 1.5)

Our transformation is presented formally in Figure 4.1. As before, the scheme is fully homomorphic if both ingredients are also fully homomorphic and it is otherwise leveled homomorphic.

Analysis. We proceed by analyzing the security and the correctness of our scheme.

Lemma 2.1 (Security)**.** *Assuming that QFHE and FHE are semantically secure, the scheme in Figure 4.1 is semantically secure.*

Proof. Let $\mathcal{A}$ be a QPT adversary against the semantic security of the rate-1 QFHE scheme. Let $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ be a public key in support of the key generation algorithm where $\mathsf{ck} = \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk})$ and $(|\phi\rangle, c)$ be an honestly computed ciphertext, where

$$\mathsf{ck} = \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk}) \text{ and } (|\phi\rangle, c) = (\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.Enc}(\mathsf{qpk}, \mathsf{otk})).$$

We define a series of hybrid distributions and argue that they are indistinguishable from the original ciphertext. First, we substitute the computation of the compression key with an encryption of 0 (padded to the appropriate length), obtaining

$$\mathsf{FHE.Enc}(\mathsf{pk}, 0)$$

The resulting distribution is computationally indistinguishable due to the semantic security of FHE. Next, we substitute the classical part of the ciphertext with an encryption of 0 (padded to the appropriate length), obtaining the ciphertext

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QFHE.Enc}(\mathsf{qpk}, 0))$$

Computational indistinguishability follows from the semantic security of QFHE. Then, we replace the quantum one-time-padded state with a totally mixed ℓ-qubit state $|u\rangle$ and get

$$(|u\rangle, \mathsf{QFHE.Enc}(\mathsf{qpk}, 0)).$$

This distribution is indistinguishable from the above due to the information-theoretic security of the QOTP. $\mathcal{A}$'s advantage in this experiment is 0, given that the ciphertext consists of a maximally mixed state and an encryption of 0, whereas the public key no longer includes any information about the secret key. Since this last distribution is computationally indistinguishable from the original ciphertext, it follows that $\mathcal{A}$'s advantage in the original experiment is negligible. □

Next we show that the scheme satisfies single-hop evaluation correctness. We remark that, making an additional 2-key circularity assumption, we can extend the scheme to multi-hop (for any number of hops) homomorphic via the techniques outlined in Section 5.3 in [CDM20].

Rate-1 QFHE

- **Key Generation:** On input the security parameter 1^λ, the key generation algorithm samples two key pairs

$$(\mathsf{pk}, \mathsf{sk}) \leftarrow\!\$\, \mathsf{FHE.Gen}(1^\lambda) \text{ and } (\mathsf{qpk}, \mathsf{qsk}) \leftarrow\!\$\, \mathsf{QFHE.Gen}(1^\lambda).$$

  Then it samples a compression key $\mathsf{ck} \leftarrow\!\$\, \mathsf{FHE.Enc}(\mathsf{pk}, \mathsf{qsk})$. The secret key of the scheme is set to $(\mathsf{sk}, \mathsf{qsk})$ and the public key consists of $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$.

- **Encryption:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and a quantum state $|\psi\rangle$, the algorithm computes and outputs $(|\phi\rangle, c) \leftarrow\!\$\, \mathsf{QFHE.QEnc}(\mathsf{qpk}, |\psi\rangle)$.

- **Evaluation:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$, a quantum circuit C, and a ciphertext $\mathsf{ct} = (|\phi\rangle, c)$, the algorithm computes and outputs the evaluated ciphertext $(|\xi\rangle, \tilde{c}) = \mathsf{QFHE.QEval}(\mathsf{qpk}, C, \mathsf{ct})$.

- **Decryption:** On input the secret key $(\mathsf{sk}, \mathsf{qsk})$ and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \tilde{c})$, the algorithm returns $|\psi\rangle = \mathsf{QFHE.QDec}(\mathsf{qsk}, (|\xi\rangle, \tilde{c}))$.

- **Compression:** On input the public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \tilde{c})$, the compression algorithm key-switches from QFHE to FHE, by homomorphically decrypting the classical part of the ciphertext, computing

$$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) = \mathsf{FHE.Eval}\left(\mathsf{pk}, \mathsf{QFHE.Dec}(\cdot, \tilde{c}), \mathsf{ck}\right)$$

  Then, it computes an ℓ-qubit state

$$|\phi\rangle = \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot |\xi\rangle$$

  and outputs $(|\phi\rangle, \mathbf{c}_0)$.

- **Compressed Decryption:** On input the secret key $(\mathsf{sk}, \mathsf{qsk})$ and a compressed ciphertext $(|\phi\rangle, \mathbf{c}_0)$, where $|\phi\rangle$ is an ℓ-qubit state, the algorithm proceeds as follows. It computes $F(\mathsf{sk}, \mathbf{c}_0) = ((f_{1,x}, f_{1,z}), \ldots, (f_{\ell,x}, f_{\ell,z}))$ and outputs the ℓ-qubit state

$$|\psi\rangle = \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle .$$

Figure 4.1: Description of a rate-1 QFHE scheme.

Lemma 2.2 (Correctness)**.** *Assuming that the schemes* FHE *and* QFHE *are correct, the scheme in Figure 4.1 satisfies compressed correctness.*

Proof. Fix a public key $(\mathsf{pk}, \mathsf{qpk}, \mathsf{ck})$ and a secret key $(\mathsf{sk}, \mathsf{qsk})$ and an input ciphetext $(|\xi\rangle, \tilde{c})$ where

$$|\xi\rangle = \mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle)$$

for some quantum state $|\psi\rangle$, where $\mathsf{otk} = (x_1, z_1, \ldots, x_\ell, z_\ell)$ and $\tilde{c}$ is a classical encryption

of otk. Recall that the compression algorithm defines

$$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z}) = \mathsf{FHE.Eval}\,(\mathsf{pk}, \mathsf{QFHE.Dec}(\cdot, \tilde{c}), \mathsf{ck})$$

which is also a classical encryption of otk, and

$$\begin{aligned}
|\phi\rangle &= \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot |\xi\rangle \\
&= \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot \mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle) \\
&= \bigotimes_{i\in[l]} \left(X^{c_{i,x}\oplus x_i} Z^{c_{i,z}\oplus z_i}\right) \cdot |\psi\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) |\psi\rangle
\end{aligned}$$

by the spooky decryption property of the rate-1 FHE scheme. The compressed decryption algorithm then returns

$$\begin{aligned}
&\bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle \\
&= \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\psi\rangle \\
&= |\psi\rangle
\end{aligned}$$

which is the correct state. □

Parameters. We calculate the rate of the above scheme. Assuming that the plaintext $|\psi\rangle$ is an ℓ-qubit state, the compressed ciphertext consists of an ℓ-qubit state $|\phi\rangle$ and the classical information $\mathbf{c}_0 \in \mathbb{Z}_q^{n+1}$. Thus we obtain a rate of

$$\rho(\lambda) = \frac{\ell}{(n+1)\log(q) + \ell} = 1 - \frac{(n+1)\log(q)}{(n+1)\log(q) + \ell}.$$

Recall that q is some polynomial in λ and thus we can bound $\log(q) \leq \log(\lambda)^2$. Setting $\ell = \Omega(\lambda(n+1)\log(\lambda)^2)$ we obtain a rate of $\rho(\lambda) = 1 - O(1/\lambda)$.

Combining Lemma 2.1 and Lemma 2.2 we obtain the following result.

Theorem 2.3 (Rate-1 QFHE)**.** *Assuming the quantum hardness of the LWE problem, there exists a leveled QFHE scheme with rate-1. Additionally assuming that the scheme is circularly secure, there exists a QFHE scheme with rate-1.*

Chapter 5

Rate-1 QFHE via Packed Dual-GSW

In Section 4 we constructed a rate-1 QFHE scheme by key switching between the classical part of a quantum FHE scheme and a rate-1 FHE scheme. In this section, we present a different, and somewhat more direct, approach to construct rate-1 QFHE. Our generic approach required us to augment the encryption scheme with a two-key cycle in order to obtain full (as opposed to leveled) homomorphism. This non-generic approach has the advantage of requiring only a one-key cycle. While these two assumptions are formally incomparable, one-key circularity is arguably more studied and it is the same assumption that was used in [Mah18a].

1 Definitions

For our construction we will make use of the notion of a lattice trapdoor, along with the following theorem.

Theorem 1.1 ([MP12])**.** *There is an efficient algorithm* $\mathsf{GenTrap}(1^n, 1^m, q)$ *that, given* $n, m \geq 1$ *and* $q \geq 2$ *such that* $m = \Omega(n\ log(q))$*, returns a matrix* $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$ *and a trapdoor* τ_A *such that the distribution of* $\mathbf{A}$ *is negligibly (in* n*) close to the uniform distribution. Moreover, there is an efficient algorithm* Invert *that on input* $\mathbf{A}$*,* τ_A *and* $\mathbf{As} + \mathbf{e}$*, where* $\mathbf{s}$ *is arbitrary in* $\mathbb{Z}_q^n$ *and* $\|e\| \leq q/\left(O(n\ log(q))\right)$*, returns* $\mathbf{s}$ *and* $\mathbf{e}$ *with overwhelming probability.*

We recall the definition of the ciphertext shrinking algorithm.

Definition 1.2 ([BDGM19])**.** *Let* $(\mathsf{KeyGen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ *be an encryption scheme where the public key specifies a message space* $\mathbb{Z}_q^\ell$*, the secret key* sk *is a matrix in* $\mathbb{Z}_q^{\ell \times m}$ *and ciphertexts are of the form* $(\mathbf{c}_a, \mathbf{c}_b)$*. Assume publicly known functions* J*,* H*,* F *and let noisy decryption compute* $\mathsf{Dec}(\mathsf{sk}, (\mathbf{c}_a, \mathbf{c}_b)) = J(\mathbf{c}_b) - \mathsf{sk} \cdot H(\mathbf{c}_a)$*, where* $J(\mathbf{c}_b) \in \mathbb{Z}_q^\ell$*,* $H(\mathbf{c}_a) \in \mathbb{Z}_q^m$*. Then the following algorithms exist:*

- Shrink : *On input the public key and a ciphertext* $(\mathbf{c}_a, \mathbf{c}_b)$ *that encrypts* $\mathbf{m} = (m_1, \ldots, m_\ell)$*, it computes and outputs the shrunken ciphertext* $(\mathbf{c}_0, c_1, \ldots, c_\ell) \in \mathbb{Z}_q^{m+1} \times \{0,1\}^\ell$*.*

- ShrinkDec : *On input the secret key and a shrunken ciphertext* $(\mathbf{c}_0, c_1, \ldots, c_\ell)$*, it computes and outputs* $F(\mathsf{sk}, \mathbf{c}_a) \oplus (c_1, \ldots, c_\ell) = (m_1, \ldots, m_\ell)$.

The classical FHE used to encrypt classical information in a quantum scheme in Lemma 1.5 is referred to as a quantum capable scheme. Below we present the definition and requirements of a quantum capable scheme.

Definition 1.3 ([Mah18a])**.** *Let FHE be a classical leveled fully-homomorphic encryption scheme. FHE is quantum capable if there exists an encryption scheme AltHE such that:*

1. *There exists an algorithm FHE.Convert that takes as input an encryption c under FHE and outputs an encryption* $\hat{c}$ *under AltHE, where both ciphertexts encrypt the same value.*
2. *AltHE allows the operation* $\oplus_H$*, which is the XOR operation that can be performed homomorphically. This operation should also be easily invertible using only the public key of AltHE.*
3. *There exists a distribution D (which may depend on parameters of FHE) that satisfies the following conditions:*
   (a) *For all ciphertexts c that can arise during homomorphic evaluation,* $\{AltHE.Enc(pk, x; r) \mid (x, r) \leftarrow D\} \approx_s \{AltHE.Enc(pk, x; r) \oplus_H c \mid (x, r) \leftarrow D\}$*, where x is the plaintext and r is the chosen randomness.*
   (b) *There exists a bounded-error quantum polynomial time procedure to, given AltHE's public key, construct the superposition*

$$\sum_{x \in \{0,1\}, r} \sqrt{D(x, r)}\, |x, r\rangle$$

   (c) *Given a ciphertext* $c = AltHE.Enc(pk, x; r) \mid (x, r) \leftarrow D$*, the secret key and some trapdoor information, it must be possible to compute* (x, r)

2 Packed Dual GSW

The first step to construct a rate-1 QFHE scheme is to present a packed version of the dual-GSW FHE scheme.

Let λ be the security parameter of the scheme. For simplicity, we assume that q is a power of 2. Let m and n be polynomially bounded functions of λ, where $m = \Omega(n \log(q))$, ℓ be the number of bits encrypted in a packed ciphertext and $N = (m + \ell)\log(q)$. Also, let B be a positive integer that will constitute the bound of the distribution χ that the error will be sampled from (see [Mah18a] for more details). We require that $q \geq \omega(\mathsf{poly}(\lambda) \cdot B)$.

For our construction we will make use of two operations from [GSW13] as shown in [Mah18a]: The linear operator $\mathbf{G}$ and the inverse operator G^{-1}. $\mathbf{G}$ is the matrix $(1, 2, \ldots, 2^{\log(q)}) \otimes \mathbf{I}_{m+\ell}$, which converts a binary representation back to its original representation. The operator $\mathbf{G}$ is well defined even for non-binary vectors. The non linear

operator G^{-1} is the inverse of $\mathbf{G}$ and converts a vector (or each column of a matrix) to its binary representation. It is important to note that $\mathbf{G}G^{-1}$ results in the identity operation.

Our construction is inspired by the work of Hiromasa et al. [HAO15] and it is shown in Figure 5.1.

Ciphertext Conversion. To show the quantum capability of our scheme and to obtain a rate-1 scheme, we need some additional algorithms. Let dGSW be the packed dual GSW scheme we introduced in Figure 5.1. Consider the following scheme QCdGSW, which is identical to dGSW, except that it has an additional conversion algorithm that converts the ciphertext to an alternate scheme, as well as a different (noisy) decryption algorithm, as described below.

- **Conversion:** On input a ciphertext $\mathbf{C}$, the algorithm sums up columns $(m+i)\log(q)$ for $i \in \{1,\ldots,\ell\}$ and outputs the one column ciphertext

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \begin{bmatrix} \mathbf{0} \\ \hline \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{bmatrix} \in \mathbb{Z}_q^{m+\ell}$$

  It continues by extracting individual ciphertexts

$$\mathbf{c}_i^* = \left[\frac{\mathbf{A}}{-\mathbf{e}_{sk_i}\mathbf{A}}\right]\mathbf{s}^* + \mathbf{e}_i^* + \left[\frac{\mathbf{0}}{\frac{q}{2}\mu_i}\right] \in \mathbb{Z}_q^{m+1}$$

  for $i \in \{1,\ldots,\ell\}$, by keeping the first m rows and the $(m+i)$-th row of $\mathbf{C}^*$, where $\mathbf{e}_{sk_i}$ is the i-th row of $\mathbf{E}_{sk}$.

- **Noisy Decryption:** On input a secret key sk and a converted ciphertext $\mathbf{c}^*$, the algorithm creates two ciphertexts $\mathbf{c}_a^*$ and $\mathbf{c}_b^*$ by keeping the first m rows and the last ℓ rows of $\mathbf{c}^*$ respectively. Then it computes and outputs

$$\mathbf{c}_b^* + \mathbf{E}_{sk}\mathbf{c}_a^*.$$

Observe that the structure of the (noisy) decryption algorithm allows us to apply the shrinking algorithm from Definition 1.2.

3 Analysis

We proceed by proving that the scheme satisfies some properties of interest.

Homomorphic Evaluation. First we show that throughout homomorphic evaluations the ciphertext preserves the form

$$\mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y}\cdot\mathbf{G} \in \mathbb{Z}_q^{(m+\ell)\times N}$$

where $\mathbf{Y} = \left[\frac{\mathbf{0}}{\mathbf{M}\cdot\mathsf{sk}}\right] = \left[\frac{\mathbf{0}}{\mathbf{M}\cdot\mathbf{E}_{sk}} \middle| \frac{\mathbf{0}}{\mathbf{M}}\right] \in \{0,1\}^{(m+\ell)\times(m+\ell)}$ and $\mathbf{M} = \begin{bmatrix} \mu_1 & & \\ & \ddots & \\ & & \mu_\ell \end{bmatrix} \in \{0,1\}^{m\times m}$.

Packed Dual GSW

- **Key Generation:** On input the security parameter 1^λ, the key generation algorithm chooses $\mathbf{E}_{sk} \in \{0,1\}^{\ell \times m}$. Using the procedure $\mathsf{GenTrap}(1^n, 1^m, q)$ it samples a random trapdoor matrix $\mathbf{A} \in \mathbb{Z}_q^{m \times n}$. Let $\mathsf{sk} = \left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array} \right] \in \{0,1\}^{\ell \times (m+\ell)}$ and $\mathbf{A}' = \left[\begin{array}{c} \mathbf{A} \\ \hline -\mathbf{E}_{sk}\mathbf{A} \end{array} \right] \in \mathbb{Z}_q^{(m+\ell) \times n}$. Let $\mathbf{P}_i \in \{0,1\}^{\ell \times \ell}$ $(i \in \{1, \ldots, \ell\})$ be the matrix with 1 in the (i,i)-th position and zero everywhere else. For $i \in \{1, \ldots, \ell+1\}$ it samples $\mathbf{S}_i \leftarrow\$ \mathbb{Z}_q^{n \times N}$, $\mathbf{E}_i \leftarrow\$ \chi^{(m+\ell) \times N}$ and then calculates

$$\mathbf{X}_i = \mathbf{A}'\mathbf{S}_i + \mathbf{E}_i + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{P}_i \mathsf{sk} \end{array} \right] \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell) \times N}$$

for $i \in \{1, \ldots, \ell\}$, and

$$\mathbf{C}_I = \mathbf{A}'\mathbf{S}_{\ell+1} + \mathbf{E}_{\ell+1} + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{I}_\ell \mathsf{sk} \end{array} \right] \cdot \mathbf{G} \in \mathbb{Z}_q^{(m+\ell) \times N}.$$

It outputs the secret key of the scheme set to sk, the public key set to $(\mathbf{A}', \{\mathbf{X}_i\}_{i \in \{1....,\ell\}}, \mathbf{C}_I)$ and the trapdoor τ_A.

- **Encryption:** On input the public key $(\mathbf{A}', \{\mathbf{X}_i\}_{i \in \ell}, \mathbf{C}_I)$ and messages $(\mu_1, \ldots, \mu_\ell)$, the algorithm samples $\mathbf{S} \leftarrow\$ \mathbb{Z}_q^{n \times N}$, $\mathbf{E} \leftarrow\$ \chi^{(m+\ell) \times N}$ and it computes and outputs

$$\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{X}_i$$

- **Evaluation:** On input the public key $(\mathbf{A}', \{\mathbf{X}_i\}_{i \in \{0,\ldots,\ell\}}, \mathbf{C}_I)$, a circuit C, and a ciphertext $\mathbf{C}$, the algorithm computes and outputs the evaluated ciphertext $\mathbf{C}'$. In order to apply the NAND gate on input $\mathbf{C}_1, \mathbf{C}_2$, it computes $\mathbf{C}_I - \mathbf{C}_1\mathbf{G}^{-1}(\mathbf{C}_2)$.

- **Decryption:** On input a secret key sk and (without loss of generality) an evaluated ciphertext $\mathbf{C}$, the algorithm computes

$$\mathbf{M} = \begin{bmatrix} \mu'_1 & & \\ & \ddots & \\ & & \mu'_\ell \end{bmatrix} = \mathsf{sk} \cdot \mathbf{C}.$$

For each entry μ'_i it returns 0 if the result is closer to 0 than $q/2$ and 1 otherwise.

Figure 5.1: Description of the packed dual GSW scheme.

A freshly encrypted ciphertext corresponds to this structure since

$$\begin{aligned} \mathbf{C} &= \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{X}_i \\ &= \mathbf{A}' \left(\mathbf{S} + \sum_{i=1}^{\ell} \mu_i \mathbf{S}_i \right) + \left(\mathbf{E} + \sum_{i=1}^{\ell} \mu_i \mathbf{E}_i \right) + \left[\begin{array}{c} \mathbf{0} \\ \hline \sum_{i=1}^{\ell} \mu_i \mathbf{P}_i \cdot \mathsf{sk} \end{array} \right] \cdot \mathbf{G} \\ &= \mathbf{A}' \left(\mathbf{S} + \sum_{i=1}^{\ell} \mu_i \mathbf{S}_i \right) + \left(\mathbf{E} + \sum_{i=1}^{\ell} \mu_i \mathbf{E}_i \right) + \left[\begin{array}{c} \mathbf{0} \\ \hline \mathbf{M} \cdot \mathsf{sk} \end{array} \right] \cdot \mathbf{G} \end{aligned}$$

Now, we show that the same structure is maintained during a NAND operation.

$$\begin{aligned}\mathbf{C}_I - \mathbf{C}_1\mathbf{G}^{-1}(\mathbf{C}_2) &= \mathbf{C}_I - (\mathbf{A}'\mathbf{S}_1 + \mathbf{E}_1 + \mathbf{Y}_1 \cdot \mathbf{G})\,\mathbf{G}^{-1}(\mathbf{C}_2) \\ &= \mathbf{C}_I - \mathbf{A}'\mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{A}'\mathbf{S}_2 - \mathbf{Y}_1\mathbf{E}_2 - \mathbf{Y}_1\mathbf{Y}_2\mathbf{G} \\ &= \mathbf{C}_I - \mathbf{A}'\mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{E}_2 - \mathbf{Y}_1\mathbf{Y}_2\mathbf{G} \\ &= \mathbf{A}'\tilde{\mathbf{S}} + \tilde{\mathbf{E}} + \tilde{\mathbf{Y}}\mathbf{G}\end{aligned}$$

where $\tilde{\mathbf{S}} = (\mathbf{S}_I - \mathbf{S}_1\mathbf{G}^{-1}(\mathbf{C}_2))$, $\tilde{\mathbf{E}} = (\mathbf{E}_I - \mathbf{E}_1\mathbf{G}^{-1}(\mathbf{C}_2) - \mathbf{Y}_1\mathbf{E}_2)$ and

$$\begin{aligned}\tilde{\mathbf{Y}} &= \left(\left[\frac{\mathbf{0}}{\mathbf{I}_\ell \mathsf{sk}}\right] - \mathbf{Y}_1\mathbf{Y}_2\right)\mathbf{G} \\ &= \left(\left[\frac{\mathbf{0}}{\mathbf{I}_\ell \cdot \mathsf{sk}}\right] - \left[\frac{\mathbf{0}}{\mathbf{M}_1 \cdot \mathsf{sk}}\right]\left[\frac{\mathbf{0}}{\mathbf{M}_2 \cdot \mathsf{sk}}\right]\right)\mathbf{G} \\ &= \left[\frac{\mathbf{0}}{(\mathbf{I}_\ell - \mathbf{M}_1\mathbf{M}_2) \cdot \mathsf{sk}}\right]\mathbf{G}\end{aligned}$$

The second equality strands true because $\mathbf{G}G^{-1}(\mathbf{C}_2) = \mathbf{C}_2$ whereas the third equality because $\mathbf{Y}_1\mathbf{A}' = \left[\frac{\mathbf{0} \mid \mathbf{0}}{\mathbf{M}_1 \cdot \mathbf{E}_{sk} \mid \mathbf{M}_1}\right] \cdot \left[\frac{\mathbf{A}}{-\mathbf{E}_{sk}\mathbf{A}}\right] = \mathbf{0}$. The form of $\tilde{\mathbf{Y}}$ is correct since

$$\mathbf{Y}_1\mathbf{Y}_2 = \left[\frac{\mathbf{0} \mid \mathbf{0}}{\mathbf{M}_1 \cdot \mathbf{E}_{sk} \mid \mathbf{M}_1}\right] \cdot \left[\frac{\mathbf{0} \mid \mathbf{0}}{\mathbf{M}_2 \cdot \mathbf{E}_{sk} \mid \mathbf{M}_2}\right] = \left[\frac{\mathbf{0} \mid \mathbf{0}}{\mathbf{M}_1\mathbf{M}_2 \cdot \mathbf{E}_{sk} \mid \mathbf{M}_1\mathbf{M}_2}\right]$$

Note that in [Mah18a], the error parameter entries are sampled from a truncated discrete Gaussian distribution with bound B. Hence, in a freshly encrypted ciphertext the error entries are bounded by $(\ell+1)\cdot B$, in $\mathbf{E}_I$ by B, and $\mathbf{Y}_1$ has at most $(m+1)$ non-zero entries in each row. As a result the entries in $\tilde{\mathbf{E}}$ are bounded by

$$B((\ell+1)N + (\ell+1)(m+1)+1) = B\cdot((\ell+1)(N+m+1)+1) \leq B\cdot(\ell+1)(N+m+2).$$

If the scheme is used to compute an L-depth circuit, the error entries of $\tilde{\mathbf{E}}$ are bounded by $B' = B\cdot(\ell+1)(N+m+2)^L$ across all ciphertexts during computations.

Quantum Capability. Here we proceed to prove the quantum capability of our scheme.

Theorem 3.1. *The scheme* QCdGSW *described before is quantum capable.*

Proof. The individual ciphertexts $\mathbf{c}_i^*$ produced by the conversion algorithm are exactly the dual-Regev encryptions of $\mu_1, \ldots, \mu_\ell$, each with secret key $\left[\begin{array}{c|c} \mathbf{e}_{sk_i} & 1 \end{array}\right] \in \mathbb{Z}_q^{m+1}$. This constitutes the AltHE scheme in Definition 1.3, which is the same as the one used in [Mah18a]. As a result, the first two requirements in Definition 1.3 are satisfied immediately. Requirement 3c is also immediately satisfied, as the matrix A with trapdoor τ_A is intact in the converted ciphertexts.

For requirements 3a and 3b, we need to appropriately bound the error parameter. Specifically, in [Mah18a], the requirements hold if the parameter of the truncated discrete Gaussian distribution that the error was sampled from is super-polynomially larger than the error entries of $\mathbf{E}^*$. From the assumption parameters we know that

modulus q is super-polynomially larger than the original error entries. When proving the invariability of the ciphertext form, we proved that the error entries of an evaluated ciphertext are bounded by $B' = B \cdot (\ell+1)(N+m+2)^L$. This means that in a converted ciphertext, they are bounded by $\ell \cdot B'$. As a result, it is possible to define a large enough error bound so as to satisfy these requirements. □

Security and Correctness. We proceed by analyzing the security and the correctness of our scheme. The security is proven against the the standard LWE assumption in the presence of encryptions of the secret key (circular LWE).

Lemma 3.2 (Security)**.** *Assuming circular LWE, the* QCdGSW *scheme introduced above is semantically secure.*

Proof. Let $\mathcal{A}$ be a QPT adversary against the semantic security of QCdGSW. Let $(\mathbf{A}', \{\mathbf{X}_i\}_{i\in\{1,\ldots,\ell\}}, \mathbf{C}_I)$ be a public key in support of the key generation algorithm and $\mathbf{C} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{X}_i \in \mathbb{Z}_q^{(m+\ell)\times N}$ be an honestly computed ciphertext where

$$\mathbf{X}_i = \mathbf{A}'\mathbf{S}_i + \mathbf{E}_i + \left[\frac{\mathbf{0}}{\mathbf{P}_i\mathsf{sk}}\right] \cdot \mathbf{G},\ \ \mathbf{C}_I = \mathbf{A}'\mathbf{S}_I + \mathbf{E}_I + \left[\frac{\mathbf{0}}{\mathbf{I}_\ell\mathsf{sk}}\right] \cdot \mathbf{G} \text{ and } \mathbf{A}' = \left[\frac{\mathbf{A}}{-\mathbf{E}_{sk}\mathbf{A}}\right].$$

We define a series of hybrid distributions and argue that they are indistinguishable from the original ciphertext. First we substitute the matrix $\mathbf{A}$ with a uniformly random matrix $\mathbf{U}$, getting

$$\mathbf{A}' = \left[\frac{\mathbf{U}}{-\mathbf{E}_{sk}\mathbf{U}}\right]$$

This distribution is statistically indistinguishable from the original due to Theorem 1.1. Then, we substitute $\mathbf{A}'$ with a uniformly random matrix $\mathbf{U}'$. The resulting distribution is statistical indistinguishable from the previous one due to the the leftover hash lemma [HILL99].

We proceed with the next ℓ hybrid distributions. For $i \in \{1, \ldots, \ell\}$, Hybrid_{i+2} is the distribution where we substitute each

$$\mathbf{X}_i = \mathbf{U}'\mathbf{S}_i + \mathbf{E}_i + \left[\frac{\mathbf{0}}{\mathbf{P}_i\mathsf{sk}}\right] \cdot \mathbf{G}$$

with a uniform matrix $\mathbf{V}_i$. Each distribution is computationally indistinguishable from the previous one on account of the hardness of the circular LWE. Similarly, for the next hybrid distribution we replace $\mathbf{C}_I$ with a uniformly random matrix $\mathbf{V}_I$. At last, we substitute the ciphertext

$$\mathbf{C} = \mathbf{U}'\mathbf{S} + \mathbf{E} + \sum_{i=1}^{\ell} \mu_i \cdot \mathbf{V}_i$$

with a uniformly random matrix $\mathbf{W}$. The resulting distribution is again computationally indistinguishable from the previous by an invocation of the hardness of the circular LWE.

$\mathcal{A}$'s advantage in the last experiment is 0, given that the ciphertext consists of a uniformly random matrix and the public key no longer includes any information of the secret key. Since this last distribution is computationally indistinguishable from the original ciphertext, it follows that $\mathcal{A}$'s advantage in the original experiment is negligible. □

Next we proceed to show the single-hop evaluation correctness of the scheme. We can extend the scheme to multi-hop (for any number of hops) homomorphic via bootstrapping.

Lemma 3.3 (Correctness)**.** *Assuming that the algorithms* Shrink *and* ShrinkDec *in Definition 1.2 are correct, then the* QCdGSW *scheme introduced above satisfies decryption correctness.*

Proof. Fix a public key $\mathsf{pk} = \mathbf{A}'$ and a secret key $\mathsf{sk} = \left[\begin{array}{c|c} \mathbf{E}_{sk} & \mathbf{I}_l \end{array}\right]$. It was proven that a QCdGSW ciphertext has form

$$\mathbf{A}'\mathbf{S} + \mathbf{E} + \mathbf{Y} \cdot \mathbf{G} = \mathbf{A}'\mathbf{S} + \mathbf{E} + \left[\begin{array}{c|c} \mathbf{0} & \mathbf{0} \\ \hline \mathbf{M} \cdot \mathbf{E}_{sk} & \mathbf{M} \end{array}\right] \cdot \mathbf{G}$$

It is easily shown that column $(m+i)\log(q)$ of $\mathbf{Y} \cdot \mathbf{G}$ is a column with zeros everywhere else and $\frac{q}{2}\mu_i$ in the $(m+i)$-th row, for $i \in \{1, \ldots, \ell\}$. Hence, adding those columns results into the converted ciphertext

$$\mathbf{c}^* = \mathbf{A}'\mathbf{s}^* + \mathbf{e}^* + \left[\begin{array}{c} \mathbf{0} \\ \hline \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{array}\right] \in \mathbb{Z}_q^{m+\ell}$$

It is important to notice that the matrix $\mathbf{A}'$ remains intact after this transformation. Recall that the decryption algorithm creates $\mathbf{c}_a^*$ and $\mathbf{c}_b^*$ by keeping the first m rows and the last ℓ rows of $\mathbf{c}^*$ respectively. Then it outputs

$$\begin{aligned} \mathbf{c}_b^* + \mathbf{E}_{sk}\mathbf{c}_a^* &= -\mathbf{E}_{sk}\mathbf{A} \cdot \mathbf{s}^* + \mathbf{e}_b^* + \left[\begin{array}{c} \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{array}\right] + \mathbf{E}_{sk} \cdot \mathbf{A}\mathbf{s}^* + \mathbf{E}_{sk} \cdot \mathbf{e}_a^* \\ &= (\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*) + \left[\begin{array}{c} \frac{q}{2}\mu_1 \\ \vdots \\ \frac{q}{2}\mu_\ell \end{array}\right] \end{aligned}$$

which consist of the correct (noisy) plaintexts.

The rounding (in order to get the plaintexts) operates correctly as long as the entries in $(\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*)$ are bounded by $\frac{q}{4}$. The entries in $\mathbf{e}^*$ are bounded by $\ell B'$ and $\mathbf{E}_{sk}$ contributes at most by a factor m. As a result, the entries in $(\mathbf{E}_{sk}\mathbf{e}_a^* + \mathbf{e}_b^*)$ are bounded by $\ell(m+1)B' = \ell(m+1)(\ell+1)B \cdot (N+m+2)^L$. Given that q is super-polynomially larger than B, we can meet the necessary requirement. □

3.1 Rate-1 Quantum FHE

The rate-1 scheme we construct is identical to a QFHE with classical keys and hybrid ciphertexts of the form

$$(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QCdGSW.Enc}(\mathsf{pk}, \mathsf{otk}))$$

(see Lemma 1.5), with the addition of a Compression and a Compressed Decryption algorithm, as described below, in order to support ℓ-qubits ciphertext compression. We use the same techniques as the rate-1 QFHE scheme described Section 4.

- **Compression:** On input the public key and (without loss of generality) an evaluated ciphertext $(|\xi\rangle, \mathbf{c}^*)$, where $|\xi\rangle$ is a quantum one time padded ℓ-qubit state and $\mathbf{c}^* \in \mathbb{Z}_q^{m+2\ell}$ is a converted packed dual GSW ciphertext, the compression algorithm uses the Shrink algorithm described in Definition 1.2 and computes

$$(\mathbf{c}_0, c_{1,x}, c_{1,z}, \ldots, c_{\ell,x}, c_{\ell,z})) \in \mathbb{Z}_q^{m+1} \times \{0,1\}^{2\ell}$$

  Then, it computes an ℓ-qubit state

$$|\phi\rangle = \bigotimes_{i\in[l]} (X^{c_{i,x}} Z^{c_{i,z}}) \cdot |\xi\rangle$$

  and outputs $(|\phi\rangle, \mathbf{c}_0)$.

- **Compressed Decryption:** On input the secret key sk and a compressed ciphertext $(|\phi\rangle, \mathbf{c}_0)$, where $|\phi\rangle$ is an ℓ-qubit state, the algorithm computes $F(\mathsf{sk}, \mathbf{c}_0) = ((f_{1,x}, f_{1,z}), \ldots, (f_{\ell,x}, f_{\ell,z}))$ and outputs the ℓ-qubit state

$$|\psi\rangle = \bigotimes_{i\in[l]} \left(X^{f_{i,x}} Z^{f_{i,z}}\right) \cdot |\phi\rangle .$$

Parameters. Assume the quantum fully homomorphic encryption scheme $(\mathsf{QFHE.Gen}, \mathsf{QFHE.QEnc}, \mathsf{QFHE.QEval}, \mathsf{QFHE.QDec})$ with classical keys and hybrid ciphertexts of the form $(\mathsf{QOTP.Enc}(\mathsf{otk}, |\psi\rangle), \mathsf{QCdGSW.Enc}(\mathsf{pk}, \mathsf{otk}))$. We can see that the Compression algorithm is identical with one described in Section 4, and similarly we obtain a rate of $\rho(\lambda) = 1 - O(1/\lambda)$.

Chapter 6

Zero-Knowledge for QMA

In the following we present a 4-Round statistical zero-knowlege protocol for QMA. Before delving into the description of our protocol, we introduce a few cryptographic tools that are going to be useful for our main protocol.

1 Witness-Indistinguishable Arguments for QMA

This section is devoted to the definition and description of our 2-round witness indistinguishable (WI) argument for QMA.

1.1 Definition

We recall the definition of 2-round WI for QMA. We consider a variant where the first message is instance-independent and we define directly this notion.

Definition 1.1 (2-Round WI for QMA)**.** *A WI protocol* $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ *for a language* $\mathcal{L} \in \mathsf{QMA}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *consists of the following efficient algorithms.*

- $\mathsf{WI.Setup}(1^\lambda)$*: On input the security parameter* 1^λ*, the setup returns a classical common reference string* crs *and a classical trapdoor* td*.*
- $\mathsf{WI.Prove}(\mathsf{crs}, |w\rangle^{\otimes p(\lambda)}, x)$*: On input a common reference string* crs*,* $p(\lambda)$*-many copies of the witness* $|w\rangle$*, and a statement* x*, the proving algorithm returns a quantum state* $|\pi\rangle$*.*
- $\mathsf{WI.Verify}(\mathsf{td}, |\pi\rangle, x)$*: On input a trapdoor* td*, a quantum state* $|\pi\rangle$*, and a statement* x*, the verification algorithm returns a classical bit* $\{0, 1\}$*.*

For the definition of completeness we refer the reader to Section 5. In the following we define the notion of (non-adaptive) multi-theorem computational soundness for private-coin ZAPs.

Definition 1.2 (Computational Soundness)**.** *A WI protocol* $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ *for a language* $\mathcal{L} \in \mathsf{QMA}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is computationally sound if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$*, all* $x \notin \mathcal{L}$*, and all non-uniform QPT provers with quantum advice* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$*, it holds that*

$$\Pr\left[|\pi\rangle = \mathcal{A}(\mathsf{crs}, x; \rho) \wedge \mathsf{WI.Verify}(\mathsf{td}, |\pi\rangle, x) = 1\right] \leq \mu(\lambda)$$

where $(\mathsf{crs}, \mathsf{td}) \leftarrow\!\!\$\, \mathsf{WI.Setup}(1^\lambda)$.

We now define the notion of (statitical) witness indistinguishability.

Definition 1.3 (Statistical Witness Indistinguishability)**.** *A WI protocol* $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ *for a language* $\mathcal{L} \in \mathsf{QMA}$ *with relation* $\mathsf{R}_\mathcal{L}$ *is statistically witness indistinguishable if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$ *and all (stateful) admissible distinguishers* $\mathcal{A}$*, it holds that*

$$\left|\Pr\left[\mathcal{A}(\mathsf{crs}, \mathsf{st})^{\mathsf{WI.Prove}^0(\mathsf{crs},\cdot,\cdot,\cdot)} = 1\right] - \Pr\left[\mathcal{A}(\mathsf{crs}, \mathsf{st})^{\mathsf{WI.Prove}^1(\mathsf{crs},\cdot,\cdot,\cdot)} = 1\right]\right| \leq \mu(\lambda).$$

where $(\mathsf{st}, \mathsf{crs}) = \mathcal{A}(1^\lambda)$ *and the oracle* $\mathsf{WI.Prove}^b$ *takes as input a statement* x *and* $p(\lambda)$*-many copies of two witnesses* $|w_0\rangle$ *and* $|w_1\rangle$ *and returns* $\mathsf{WI.Prove}(\mathsf{crs}, |w_b\rangle^{\otimes p(\lambda)}, x)$*. We say that the distinguisher* $\mathcal{A}$ *is admissible if it holds that* $(|w_0\rangle^{\otimes p(\lambda)}, |w_1\rangle^{\otimes p(\lambda)}) \in \mathsf{R}_\mathcal{L}(x)$*.*

1.2 Construction

In the following we describe our protocol for statistical WI for QMA. Let $\varepsilon(\lambda)$ be a (fixed) negligible function. We assume the existence of the following building blocks (all secure against quantum adversaries):

- A sigma protocol $(\Sigma.\mathsf{Gen}, \Sigma.\mathsf{Com}, \Sigma.\mathsf{Chal}, \Sigma.\mathsf{Resp})$ for QMA satisfying statistical special zero-knowledge and with $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ soundness error.
- A public coin ZAP $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ for NP with statistical witness indistinguishability and $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ soundness error.
- A pseudorandom function $(\mathsf{PRF.Gen}, \mathsf{PRF.Eval})$ with distinguishing advantage $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$.
- A maliciously circuit private classical (levelled) FHE scheme $(\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec})$ with distinguishing advantage $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$.
- An SBSH commitment scheme $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$ that satisfies $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding.

Where $\mu(\lambda)$ is some negligible function and κ is the security parameter of the primitives with super-polynomially bounded disitinguishing advantage. Our protocol is formally described in Figure 6.1. Completeness of the protocol follows by a standard argument.

Soundness. We show that our protocol satisfies (non-adaptive) soundness. We also note that the proof can be lifted to the adaptive setting (i.e. where the prover can choose the challenge statement adaptively) using complexity leveraging, albeit at the cost of a stronger assumption for the security of the underlying primitives.

Theorem 1.4 (Soundness)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, the WI argument described in Figure 6.1 satisfies computational soundness.*

Statistical WI Arguments for QMA

- **Setup:** The setup algorithm samples a PRF key $k \leftarrow\$ \mathsf{PRF.Gen}(1^\kappa)$ and an FHE key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow\$ \mathsf{FHE.Gen}(1^\kappa)$. Additionally it samples a commitment key $\mathsf{ck} \leftarrow\$ \Sigma.\mathsf{Gen}(1^\kappa)$, an SBSH commitment key $\mathsf{ck}_0 \leftarrow\$ \mathsf{SBSH.Gen}(1^\lambda)$, and a common reference string $\mathsf{crs}_{\mathsf{ZAP}} \leftarrow\$ \mathsf{ZAP.Setup}(1^\kappa)$. The algorithm computes $c_k \leftarrow\$ \mathsf{FHE.Enc}(\mathsf{pk}, k)$ and sets the common reference string and the trapdoor as

$$\mathsf{crs} = (\mathsf{pk}, c_k, \mathsf{ck}, \mathsf{ck}_0, \mathsf{crs}_{\mathsf{ZAP}}) \text{ and } \mathsf{td} = (\mathsf{sk}, k).$$

- **Prove:** On input $2p(\lambda)$-many copies of the witness $|w\rangle^{2p(\lambda)}$ and a statement x, the proving algorithm does the following. First, it samples a commitment key $\mathsf{ck}_1 \leftarrow\$ \mathsf{SBSH.Key}(\mathsf{ck}_0)$, then for $b \in \{0,1\}$, it samples a classical string $r_{\Sigma,b} \leftarrow\$ \{0,1\}^\kappa$ and computes the first $|\alpha_b\rangle = \Sigma.\mathsf{Com}(|w\rangle^{\otimes p(\lambda)}, \mathsf{ck}; r_{\Sigma,b})$. Then it evaluates homomorphically the response function of the sigma protocol sampling the challenge from the PRF, i.e. it computes

$$c_{\gamma,b} = \mathsf{FHE.Eval}(\mathsf{pk}, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(\cdot, x\|b), r_{\Sigma,b}), c_k; r_{\mathsf{FHE},b}).$$

  where $r_{\mathsf{FHE},b}$ are some classical random coins. In addition, it computes an SBSH commitment to $r_{\Sigma,b}$ as $c_{r,b} = \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), r_{\Sigma,b}; r_{\mathsf{SBSH},b})$, where $r_{\mathsf{SBSH},b}$ are also uniformly sampled coins. Finally it computes a statistical ZAP π for the classical statement

$$\left\{ \exists\ (b, w_\Sigma, w_{\mathsf{FHE}}, w_{\mathsf{SBSH}}) \text{ s.t. } \begin{array}{l} c_{\gamma,b} = \mathsf{FHE.Eval}(\mathsf{pk}, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(\cdot, x\|b), w_\Sigma), c_k; w_{\mathsf{FHE}}) \\ \wedge\ c_{r,b} = \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), w_\Sigma; w_{\mathsf{SBSH}}) \end{array} \right\}$$

  using $(0, r_{\Sigma,0}, r_{\mathsf{FHE},0}, r_{\mathsf{SBSH},0})$ as a witness. The output of the algorithm is defined as

$$|\pi\rangle = (\mathsf{ck}_1, |\alpha_0\rangle, |\alpha_1\rangle, c_{\gamma,0}, c_{\gamma,1}, c_{r,0}, c_{r,1}, \pi).$$

- **Verify:** The verification algorithm checks that the ZAP π against the common reference string $\mathsf{crs}_{\mathsf{ZAP}}$, then for $b \in \{0,1\}$ does the following. It recomputes the challenge for the sigma protocol $\beta_b = \mathsf{PRF.Eval}(k, x\|b)$ and it recovers the response $\gamma_b = \mathsf{FHE.Dec}(\mathsf{sk}, c_{\gamma,b})$ by decrypting the corresponding FHE ciphertext. Then it checks whether $(|\alpha_b\rangle, \beta_b, \gamma_b)$ is a valid transcript for the sigma protocol. If all of the above conditions are satisfied, the algorithm returns 1, otherwise it returns 0.

Figure 6.1: Description of a statistical WI argument for QMA.

Proof. We are going to show that the prover success probability is bounded by a negligible function $\varepsilon(\lambda)$. Let $x \notin \mathcal{L}$ be the challenge statement and let Cheat be the event where the prover causes the verifier to accept x. Assume towards contradiction that

$$\Pr[\mathsf{Cheat}] \geq \varepsilon(\lambda).$$

Then, by the $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding property of the SBSH commitment scheme, we have that

$$\Pr[\mathsf{Cheat} \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

for some negligible function $\mu(\lambda)$. Let $r_0^* = \mathsf{SBSH.Ext}(r, \mathsf{ck}_0, \mathsf{ck}_1, c_{\tau,0})$ and $r_1^* = \mathsf{SBSH.Ext}(r, \mathsf{ck}_0, \mathsf{ck}_1, c_{\tau,1})$ denote the outputs of the extractor on such a transcript, where r denote the random coins used in the $\mathsf{SBSH.Gen}$ algorithm. We now gradually change the verification procedure and we argue that the probability that the above event happens does not decrease significantly.

- The verifier no longer decrypts the FHE ciphertext, instead, for $b \in \{0,1\}$, it computes $\gamma_b = \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(k, x\|b), r_b^*)$ and checks whether the transcript $(|\alpha_b\rangle, \mathsf{PRF.Eval}(k, x\|b), \gamma_b)$ is accepting. If at least one of the two transcripts is accepting and the ZAP π correctly verifies, then the verifier returns 1, otherwise it returns 0. Let Cheat_1 be the event that the prover causes the modified verifier to accept on some $x \notin \mathcal{L}$. We want to argue that

  $$\Pr\left[\mathsf{Cheat}_1 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

  for some negligible function $\mu(\lambda)$. To show this, it suffices to consider the case where the prover passes the original verification procedure but fails the modified one. This implies that the prover has computed two inconsistent commitments $(c_{\tau,0}, c_{\tau,1})$ but the ZAP π correctly verifies. Thus, if the inequality above does not hold, then we obtain a contradiction against the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-soundness of the ZAP argument.

- The verifier computes c_k as an encryption of 0 (padded to the appropriate length), i.e. it computes $c_k \leftarrow\!\$\, \mathsf{FHE.Enc}(\mathsf{pk}, 0)$. Let Cheat_2 be the event that the prover causes the modified verifier to accept on some $x \notin \mathcal{L}$. Recall that the modified verifier no longer uses the FHE secret key in its routine. Thus, by the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-semantic security of the FHE scheme we have that

  $$\Pr\left[\mathsf{Cheat}_2 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda)).$$

- Instead of computing $\beta_b = \mathsf{PRF.Eval}(k, x\|b)$, the verifier samples (β_0, β_1) uniformly. By the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-pseudorandomness of the pseudorandom function, we have that

  $$\Pr\left[\mathsf{Cheat}_3 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

  where Cheat_3 denotes the event that the prover causes the modified verifier to accept on some $x \notin \mathcal{L}$.

The last inequality implies that either of the sigma protocols $(|\alpha_0\rangle, \beta_0, \gamma_0)$, $(|\alpha_1\rangle, \beta_1, \gamma_1)$ is accepting for some $x \notin \mathcal{L}$, where β_0 and β_1 are sampled uniformly and independently of $|\alpha_0\rangle$ and $|\alpha_1\rangle$, with probability at least $\varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$. This contradicts the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-soundness of the sigma protocol and concludes our proof. □

Witness Indistinguishability. We show that our protocol satisfies statistical witness indistinguishability.

Theorem 1.5 (Statistical Witness Indistinguishability)**.** *The WI argument described in Figure 6.1 satisfies statistical witness indistinguishability.*

Proof. We begin by fixing the challenge bit $b = 0$ and we gradually modify the experiment through a series of hybrids that we show to be statistically close.

- Hybrid $\mathcal{H}_0$: This is the original experiment with the challenge bit fixed to $b = 0$, i.e. the oracle always uses the witness $|w_0\rangle$.
- Hybrid $\mathcal{H}_1$: In this hybrid we modify the answers to all queries of the adversary to compute $c_{\tau,1}$ as a commitment to 0, i.e. $c_{\tau,1} \leftarrow\!\!\$\, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), 0)$. Note that the randomness of the commitment is never used in the proof and thus, by the statistically hiding property of the SBSH commitment we have that

$$\mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), r_{\Sigma,1}) \approx_s \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), 0).$$

  It follows that the two hybrids are statistically indistinguishable.
- Hybrid $\mathcal{H}_2$: In this hybrid we first run the (unbounded) extractor given by the malicious circuit privacy of the FHE scheme $k^* = \mathsf{FHE.Ext}(1^\kappa, \mathsf{pk}, c_k)$, then we compute the evaluated cirphertext as $c_{\gamma,1} \leftarrow\!\!\$\, \mathsf{FHE.Sim}(1^\kappa, \mathsf{pk}, c_k, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(k^*, x\|1), r_{\Sigma,1}))$. By the statistical circuit privacy of the FHE scheme we have that

$$\begin{aligned}&\mathsf{FHE.Eval}(\mathsf{pk}, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(\cdot, x\|1), r_{\Sigma,1}), c_k)\\ &\approx_s \mathsf{FHE.Sim}(1^\kappa, \mathsf{pk}, c_k, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(k^*, x\|1), r_{\Sigma,1}))\end{aligned}$$

  and thus the two hybrids are statistically close.
- Hybrid $\mathcal{H}_3$: In this hybrid we compute $\beta_1 = \mathsf{PRF.Eval}(k^*, x\|1)$ and we use the challenge to simulate the response for the sigma protocol. I.e. we compute $(|\alpha_1\rangle, \gamma_1) \leftarrow\!\!\$\, \Sigma.\mathsf{Sim}(x, \beta_1)$ and we set $c_{\gamma,1} \leftarrow\!\!\$\, \mathsf{FHE.Sim}(1^\kappa, \mathsf{pk}, c_k, \gamma_1)$. Note that the only difference with respect to the previous hybrid is that we do compute a simulated transcript of the sigma protocol instead of an honest one. By the statistical special zero-knowledge property of the sigma protocol we have that

$$(\Sigma.\mathsf{Com}(|w_0\rangle^{\otimes p(\lambda)}; r_{\Sigma,1}), \Sigma.\mathsf{Resp}(\beta_1, r_{\Sigma,1})) \approx_s \Sigma.\mathsf{Sim}(x, \beta_1)$$

  and therefore the two hybrids are statistically close.
- Hybrid $\mathcal{H}_4$: In this hybrid we switch the computation of $|\alpha_1\rangle$ and $c_{\gamma,1}$ to use again an honest witness, except that we use $|w_1\rangle$ instead of $|w_0\rangle$. Specifically we compute the commitment of the sigma protocol as $|\alpha_1\rangle \leftarrow\!\!\$\, \Sigma.\mathsf{Com}(|w_1\rangle^{\otimes p(\lambda)}; r_{\Sigma,1})$ and the simulated ciphertext as $c_{\gamma,1} \leftarrow\!\!\$\, \mathsf{FHE.Sim}(1^\kappa, \mathsf{pk}, c_k, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(k^*, x\|1), r_{\Sigma,1}))$. The two hybrids are statistically indistinguishable by the statistical special zero-knowledge property of the sigma protocol (same argument as $\mathcal{H}_2 \approx_s \mathcal{H}_3$).
- Hybrid $\mathcal{H}_5$: In this hybrid we switch back to a correctly evaluated FHE ciphertext, i.e. we compute $c_{\gamma_1} \leftarrow\!\!\$\, \mathsf{FHE.Eval}(\mathsf{pk}, \Sigma.\mathsf{Resp}(\mathsf{PRF.Eval}(\cdot, x\|1), r_{\Sigma,1}), c_k)$. By the (malicious) statistical circuit privacy of the FHE scheme, the two hybrids are statistically close (same argument as $\mathcal{H}_1 \approx_s \mathcal{H}_2$).
- Hybrid $\mathcal{H}_6$: In this hybrid we revert the changes to the SBSH commitment, i.e. we compute $c_{\tau,1} \leftarrow\!\!\$\, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), r_{\Sigma,1})$. By the statistical hiding of the SBSH commitment we have that the two hybrids are statistically indistinguishable (same argument as $\mathcal{H}_0 \approx_s \mathcal{H}_1$).

- Hybrid $\mathcal{H}_7$: This hybrid is identical to the previous one, except that we compute the statistical ZAP argument using $(1, r_{\Sigma,1}, r_{\mathsf{FHE},1}, r_{\mathsf{SBSH},1})$. Note that the the messages are indeed well-formed and thus statistical indistinguishability follows by the statistical witness indistinguishability of the ZAP argument system.
- Hybrids $\mathcal{H}_8 \dots \mathcal{H}_{13}$: In this series of hybrids we change how we compute $(|\alpha_0\rangle, c_{\gamma,0}, c_{\tau,0})$ analogously as we did in hybrids $\mathcal{H}_1 \dots \mathcal{H}_6$, i.e. using $|w_1\rangle$ instead of $|w_0\rangle$. Note that the underlying random coins are no longer used in the computation of the ZAP argument and thus to indistinguishability follows along the same lines as what we discussed above.

Observe that hybrid $\mathcal{H}_{13}$ is identical to $\mathcal{H}_0$ except that the challenge bit is fixed to $b = 1$ and in particular the oracle uses the witness $|w_1\rangle$ to compute the ZAP argument. It follows that our protocol satisfies statistical witness indistinguishability. □

2 Post-Quantum Conditional Disclosure of Secrets

Conditional disclosure of secrets (CDS) [AIR01] for a language $\mathcal{L}$ in NP with relation $\mathsf{R}_{\mathcal{L}}$ is the interactive analogue of witness encryption [GGSW13]: Given a statement x and a message m from the sender, the receiver is able to recover m if $x \in \mathcal{L}$, whereas m stays hidden if this is not the case. Furthermore, the witness w for x should be kept secret from the eyes of the sender.

Definition. We recall the definition of a CDS protocol. In this work we consider two variants: A 3-round statistically-receiver private (SRP) CDS and a 2-round statistically sender private (SSP) CDS. The syntax below is defined for the 3-round variant and the 2-round protocol can be defined analogously by omitting the first algorithm.

Definition 2.1 (CDS Protocol for NP)**.** *A CDS protocol* $(\mathsf{Setup}, \mathsf{R}, \mathsf{S}, \mathsf{D})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *consists of the following efficient algorithms.*

- $\mathsf{Setup}(1^\lambda)$*: On input the security parameter* 1^λ*, the setup returns a first message* ct_0*.*
- $\mathsf{R}(\mathsf{ct}_0, w)$*: On input a first message* ct_0 *and a witness* w*, the receiver algorithm returns a second message* ct_1 *and a key* k*.*
- $\mathsf{S}(\mathsf{ct}_1, x, m)$*: On input a second message* ct_1*, a statement* x*, and a message* m*, the sender algorithm returns a third message* ct_2*.*
- $\mathsf{D}(\mathsf{ct}_2, k)$*: On input a third message* ct_2 *and a key* k*, the decryption algorithm returns a message* m

We define completeness for a CDS protocol.

Definition 2.2 (Completeness)**.** *A CDS protocol* $(\mathsf{Setup}, \mathsf{R}, \mathsf{S}, \mathsf{D})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is complete if for all* $\lambda \in \mathbb{N}$*, all* $x \in \mathcal{L}$*, all* $x \in \mathsf{R}_{\mathcal{L}}(x)$*, and all messages* m *it holds that*

$$\Pr\left[\mathsf{D}(\mathsf{S}(\mathsf{ct}_1, x, m), k) = m\right] = 1.$$

where $\mathsf{ct}_0 \leftarrow\!\!\$\, \mathsf{Setup}(1^\lambda)$ *and* $(\mathsf{ct}_1, k) \leftarrow\!\!\$\, \mathsf{R}(\mathsf{ct}_0, w)$*.*

Next we define the notion of (computational and statistical) receiver privacy.

Definition 2.3 (Receiver Privacy)**.** *A CDS protocol* $(\mathsf{Setup}, \mathsf{R}, \mathsf{S}, \mathsf{D})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is computationally (statistically, resp.) receiver private if for all* $\lambda \in \mathbb{N}$*, all strings* w*, and all first messages* ct_0 *the following distributions are computationally (statistically, resp.) indistinguishable*

$$(\mathsf{ct}_0, c_0) \approx (\mathsf{ct}_0, c_1)$$

where $(c_0, k_0) \leftarrow\!\!\$\, \mathsf{R}(\mathsf{ct}_0, 0)$ *and* $(c_1, k_1) \leftarrow\!\!\$\, \mathsf{R}(\mathsf{ct}_0, 1)$*.*

In recent works [CM21], there exists a simple construction of post-quantum SRP-CDS for NP assuming an $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes extractable SRP oblivious transfer and a simulation secure garbling scheme $(\mathsf{Garble}, \mathsf{GEval})$ for NC1 circuits.[1] Hence, we get the following.

Lemma 2.4 ([CM21])**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem, there exists an SRP-CDS scheme* $(\mathsf{R}, \mathsf{S}, \mathsf{D})$ *with statistical receiver privacy and computational sender privacy.*

Finally we define the notion of (computational and statistical) sender privacy.

Definition 2.5 (Sender Privacy)**.** *A CDS protocol* $(\mathsf{Setup}, \mathsf{R}, \mathsf{S}, \mathsf{D})$ *for a language* $\mathcal{L} \in \mathsf{NP}$ *with relation* $\mathsf{R}_{\mathcal{L}}$ *is computationally (statistically, resp.) sender private if there exists a negligible function* μ *such that for all* $\lambda \in \mathbb{N}$*, all* $x \notin \mathcal{L}$*, and all non-uniform QPT (unbounded, resp.) receivers with quantum advice* $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda \in \mathbb{N}}$*, it holds that*

$$|\Pr[\mathcal{A}(\mathsf{S}(\mathsf{ct}_1, x, m), \mathsf{st}; \rho) = 1] - \Pr[\mathcal{A}(\mathsf{S}(\mathsf{ct}_1, x, 0), \mathsf{st}; \rho) = 1]| \leq \mu(\lambda).$$

where $\mathsf{ct}_0 \leftarrow\!\!\$\, \mathsf{Setup}(1^\lambda)$ *and* $(\mathsf{st}, \mathsf{ct}_1) = \mathcal{A}(\mathsf{ct}_0; \rho)$*.*

It is well-known that a 2-round SSP-CDS can be built from any 2-round oblivious transfer and information-theoretically secure randomized encodings [IK00]. Thus we have the following fact.

Lemma 2.6 ([BD18])**.** *Assuming the post-quantum hardness of the LWE problem, there exists an SSP-CDS scheme* $(\mathsf{R}, \mathsf{S}, \mathsf{D})$ *with computational receiver privacy and statistical sender privacy.*

3 4-Round Zero-Knowledge for QMA

We assume the existence of the following building blocks (all secure against quantum adversaries):

- A circuit-private classical QFHE scheme $(\mathsf{QFHE.Gen}, \mathsf{QFHE.Enc}, \mathsf{QFHE.Eval}, \mathsf{QFHE.Dec})$ with distinguishing advantage $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$.
- A non-interactive perfectly binding commitment Com with hiding advantage $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$.

[1] Note that NC1 circuits suffice here, since it is well known that the validity of any NP statement can be verified by an NC1 circuit.

- An SBSH commitment scheme ($\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com}$) that satisfies $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding.
- A 2-round WI argument ($\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify}$) for QMA, with statistical witness indistinguishability and $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ soundness error.
- a 3-round statistically receiver private conditional disclosure of secrets scheme ($\mathsf{SRP\text{-}CDS.Setup}$, $\mathsf{SRP\text{-}CDS.R}, \mathsf{SRP\text{-}CDS.S}, \mathsf{SRP\text{-}CDS.D}$) for NP, with $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ computational sender privacy.
- A CC obfuscator Obf with $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ simulatability.
- A 3-round sometimes simulatable statistical ZK protocol ($\mathsf{SSim\text{-}ZK.Setup}, \mathsf{SSim\text{-}ZK.R}$, $\mathsf{SSim\text{-}ZK.S}$) that satisfies $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes simulatability.

Our protocol is formally described in Figure 6.2.

[Soundness] At first we show that the protocol satisfies computational soundness.

Theorem 3.1. *[Soundness] Assuming the quantum quasi-polynomial hardness of the LWE problem and the existence of a quantum quasi-polynomial semantically secure FHE, the protocol described in Figure 6.2 satisfies computational soundness.*

Proof. Let P^* be a malicious prover that produces an accepting state $(\mathsf{ck}_3, \mathsf{cmt}_{\mathsf{otk}}, |\psi\rangle, ct'_S, |\pi_2\rangle)$ for some statement $x \notin \mathcal{L}$. We define the aforementioned event as Cheat and assume towards contradiction that the probability of P^* succeeding in cheating is

$$\Pr[\mathsf{Cheat}] \geq \varepsilon(\lambda).$$

Then, by the $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding property of the SBSH commitment scheme and the $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes simulatability of the $\mathsf{SSim\text{-}ZK}$, we have that

$$\Pr[\mathsf{Cheat} \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge (\mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2$$

for some negligible function $\mu(\lambda)$. Let $\tilde{y} = \mathsf{SBSH.Ext}(r_{\mathsf{Gen}}, \mathsf{ck}_0, \mathsf{ck}_1, \mathsf{cmt}_y)$ and $z\tilde{k}_2 \leftarrow \mathsf{SSim\text{-}ZK.Sim}(1^\lambda,$
$r_{\mathsf{Setup}}, z_2)$, where r_{Gen} and r_{Setup} are the randomnesses used in the respective first messages. We can now gradually change the procedure and we argue that the probability that the above defined event happens does not decrease significantly.

- The verifier computes and sends a simulated $\tilde{\mathsf{zk}}_2$ instead of zk_2. If we define Cheat_1 as the event that this modified version accepts, we want to argue that

  $$\Pr[\mathsf{Cheat}_1 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2.$$

  Observe that the events Cheat and Cheat_1 only differ in case $\tilde{\mathsf{zk}}_2 \neq \mathsf{zk}_2$. If we assume that the inequality doesn't hold we get a contradiction against the sometimes simulatability of $\mathsf{SSim\text{-}ZK}$.

4-Round (Statistical) ZK Argument for QMA

- **First Message ($\mathbf{V} \to \mathbf{P}$):** The verifier samples $(\mathsf{td}, r, s) \leftarrow\!\$\, \{0,1\}^\kappa$ and computes:
  - $(\mathsf{pk}, \mathsf{sk}) \leftarrow\!\$\, \mathsf{QFHE.Gen}(1^\kappa; r)$ and $c_{td} = \mathsf{QFHE.Enc}(\mathsf{pk}, \mathsf{td})$.
  - the obfuscated program $\widetilde{\mathbf{CC}} \leftarrow Obf(\mathbf{CC}[\mathsf{QFHE.Dec}(\mathsf{sk}, \cdot), s, (r, \mathsf{sk})])$.
  - a commitment $c = \mathsf{Com}(0; r)$ and an SBSH commitment key $\mathsf{ck}_0 \leftarrow\!\$\, \mathsf{SBSH.Gen}(1^\lambda)$.
  - the first message of the 3-round CDS, $\mathsf{ct}_0 \leftarrow\!\$\, \mathsf{SRP\text{-}CDS.Setup}(1^\kappa)$.
  - the first message of SSim-ZK, $\mathsf{zk}_0 \leftarrow\!\$\, \mathsf{SSim\text{-}ZK.Gen}(1^\lambda)$.

  It sends $\left(\mathsf{pk}, c, c_{\mathsf{td}}, \widetilde{\mathbf{CC}}, \mathsf{ct}_0, \mathsf{ck}_0, \mathsf{zk}_0\right)$ to the prover.
- **Second Message ($\mathbf{P} \to \mathbf{V}$):** The prover samples $y \leftarrow\!\$\, 0^\kappa$ and computes:
  - a commitment key $\mathsf{ck}_1 \leftarrow\!\$\, \mathsf{SBSH.Key}(\mathsf{ck}_0)$ and $\mathsf{cmt}_y \leftarrow \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), y; r_{\mathsf{cmt}_y})$
  - $(\mathsf{zk}_1, \mathsf{td}_{\mathsf{SSim\text{-}ZK}}) \leftarrow \mathsf{SSim\text{-}ZK.R}(\mathsf{zk}_0)$.
  - the second message of the SRP-CDS, $(\mathsf{ct}_R, k) \leftarrow \mathsf{SRP\text{-}CDS.R}(\mathsf{ct}_0, (y, r_{\mathsf{cmt}_y}))$.

  It sends to the verifier $(\mathsf{ck}_1, \mathsf{zk}_1, \mathsf{ct}_R)$.
- **Third Message ($\mathbf{V} \to \mathbf{P}$):** The verifier computes
  - $\mathsf{ct}_S \leftarrow \mathsf{SRP\text{-}CDS.S}(\mathsf{ct}_R, z_1, s)$, where the statement z_1 attests to $y = \mathsf{td}$ and $\mathsf{cmt}_y = \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), y; r_{\mathsf{cmt}_y})$.
  - the first message of a WI argument, $(\mathsf{crs}_{wi}, \mathsf{td}_{wi}) \leftarrow\!\$\, \mathsf{WI.Setup}(1^\lambda)$.
  - $\mathsf{zk}_2 \leftarrow \mathsf{SSim\text{-}ZK.S}(\mathsf{zk}_1, z_2, r_{\mathsf{SSim\text{-}ZK}})$, where z_2 is the statement that all of the verifier's messages so far are explainable, with the random coins $r_{\mathsf{SSim\text{-}ZK}}$ as witness.

  It sends to the prover $(\mathsf{ct}_S, \mathsf{crs}_{wi}, \gamma, \mathsf{ck}_2)$.
- **Fourth Message ($\mathbf{P} \to \mathbf{V}$):** The prover first verifies $\mathsf{SSim\text{-}ZK.Verify}(\mathsf{td}_{\mathsf{SSim\text{-}ZK}}, \mathsf{zk}_2)$. If the verification is not successful it aborts. Else, on input $p(\lambda)$-many copies of the witness $|w\rangle^{\otimes p(\lambda)}$ and a statement x, it sends a WI proof $|\pi\rangle$:

$$\{x \in \mathcal{L} \vee \exists\ r : c = \mathsf{Com}(0; r)\}\,.$$

- **Verify:** The verifier accepts if $\mathsf{WI.Verify}(\mathsf{td}_{wi}, |\pi\rangle\,, x) = 1$.

Figure 6.2: Description of a 4-round statistical ZK argument for QMA (plain model)

- The verifier's third message of the SRP-CDS, ct_S, returns always zero. If we define Cheat_2 as the event that this modified version accepts, we want to prove that

$$\Pr\left[\mathsf{Cheat}_2 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2.$$

  The proof is presented in Lemma 3.2.

- The verifier's obfuscated program in the first message always returns 0. If we define Cheat_3 as the event that this modified version accepts we want to prove that

$$\Pr\left[\mathsf{Cheat}_3 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2.$$

  This is true due to the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ compute and compare obfuscation security

- The verifier's commitment c in the first message is changed to $c = \mathsf{Com}(1, r)$ instead of a commitment to zero. If we define Cheat_4 as the event that this modified version accepts we want to prove that

$$\Pr\left[\mathsf{Cheat}_4 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2.$$

  The inequality holds due to the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ hiding of the commitment.

This last inequality implies that the WI proof is accepting with probability at least $\varepsilon(\lambda)^4 \cdot (1+\mu(\lambda))^2$, when neither of the clauses is satisfied. This contradicts the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ soundness of the WI proof and concludes our proof. □

Lemma 3.2. *Given the definition of the events* Cheat_1 *and* Cheat_2 *in Theorem 3.1 and assuming that*

$$\Pr\left[\mathsf{Cheat}_1 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2,$$

then

$$\Pr\left[\mathsf{Cheat}_2 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right] \geq \varepsilon(\lambda)^4 \cdot (1 + \mu(\lambda))^2.$$

Proof. Consider the interaction where the verifier extracts $\tilde{y}$ using the $\mathsf{SBSH.Ext}$ algorithm, and if $\tilde{y} = \mathsf{td}$ the verifier aborts (denote this event by Abort); otherwise it continues with the interaction. If the event does not happen, then the desired inequality follows by a reduction against the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ computational sender privacy of the SRP-CDS.

We are now going to show that the probability that Abort happens is negligibly smaller than $\Pr\left[\mathsf{Cheat}_2 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding} \wedge \mathsf{zk}_0, \mathsf{zk}_1) \in \mathsf{Simulation}\right]$. In order to do that we consider the following hybrid distributions:

- Hybrid $\mathcal{H}_a$: This is the protocol we presented above.
- Hybrid $\mathcal{H}_b$: This hybrid process is identical to the above except that the the CC obfuscated program $\widetilde{\mathbf{CC}}$ returns always 0. These processes are computationally indistinguishable given the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ security of $\widetilde{\mathbf{CC}}$.
- Hybrid $\mathcal{H}_c$: This hybrid process is identical to the above except that the verifier, instead of sending an encryption of td in the first message, it sends $\mathsf{QFHE.Enc}(\mathsf{pk}, 0)$. Computational indistinguishability follows from the $\varepsilon(\lambda)^4 \cdot \mu(\lambda)$ semantic security of QFHE.

In the last hybrid process, we have no information about td, and hence the probability to guess $y = \mathsf{td}$ is negligible. This concludes our proof. □

Quantum Rewinding Lemma Before we move on to zero-knowledge, recall the definition of the Quantum Rewinding Lemma (Lemma 9 from [Wat09]), which constructs a quantum algorithm for amplifying the success probability of quantum sampler circuits under certain conditions. The below definition is taken directly from the modified version in [BS20].

Lemma 3.3. *[Quantum Rewinding Lemma] There is a quantum algorithm* R *that gets as input:*

- *A general quantum circuit* Q *with* n *input qubits that outputs a classical bit* b *and an additional* m *output qubits.*
- *An* n*-qubit state* $|\psi\rangle$*.*
- *A number* $t \in \mathbb{N}$*.*

R *executes in time* $t \cdot \mathsf{poly}(|\mathsf{Q}|)$ *and outputs a distribution over* m*-qubit states* $D_\psi \coloneqq \mathsf{R}(\mathsf{Q}, |\psi\rangle, t)$ *with the following guarantees.*

For an n*-qubit state* $|\psi\rangle$*, denote by* Q_ψ *the conditional distribution of the output distribution* $\mathsf{Q}(|\psi\rangle)$*, conditioned on* $b = 0$*, and denote by* $p(\psi)$ *the probability that* $b = 0$*. If there exists* $p_0, q \in (0,1)$*,* $\epsilon \in (0, \frac{1}{2})$ *such that:*

- *Amplification executes for enough time:* $t \geq \frac{\log(1/\epsilon)}{4 \cdot p_0(1-p_0)}$*,*
- *There is some minimal probability that* $b = 0$*: For every* n*-qubit state* $|\psi\rangle$*,* $p_0 \leq p(\psi)$*,*
- $p(\psi)$ *is input independant, up to* ϵ *distance: For every* n*-qubit state,* $|\psi\rangle, |p(\psi) - q| < \epsilon$*, and*
- q *is closer to* $\frac{1}{2}$*:* $p_0(1-p_0) \leq q(1-q)$*,*

then for every n*-qubit state* $|\psi\rangle$*,*

$$\mathsf{TD}\left(\mathsf{Q}_\psi, D_\psi\right) \leq 4\sqrt{\epsilon}\frac{\log(1/\epsilon)}{p_0(1-p_0)}$$

where TD *denotes the trace distance.*

Zero-Knowledge. Here we show that the scheme satisfies statistical zero-knowledge. In order to prove ZK, we follow the technique presented in [BS20], so as to simulate aborting verifiers as well. More specifically, we describe two simulators Sim_a and Sim_{na}, that on input (x, V^*, ρ) simulate different types of interactions. Sim_a tries to simulate an aborting interaction and Sim_{na} a non-aborting interaction. Formally, an aborting interaction is an interaction where the verifier aborts or fails to prove the $\mathsf{SSim\text{-}ZK}$, whereas a non-aborting interaction is one where the verifier doesn't abort before the fourth message and also the $\mathsf{SSim\text{-}ZK}$ is succesful. Then, we describe a combined Simulator Sim_{comb}, which randomly chooses $b \leftarrow\!\!\$\, \{a, na\}$ and uses Sim_b to simulate the interaction. We prove that the output of Sim_{comb} is indistinguishable from the output of the real interaction, as long as it doesn't fail (i.e. picks the correct b), which happens with probability negligibly close to $\frac{1}{2}$. Lastly, it is proven that Sim_{comb} satisfies

the required conditions for applying Watrous' quantum rewinding lemma, so that the success probability can be amplified negligibly close to 1.

The simulator $\mathsf{Sim}_a(x, V^*, \rho)$ proceeds as follows:

- In the second message, it computes $\mathsf{cmt}_y = \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), 0^\lambda; r_0)$.
- If at some point before the fourth message the verifier aborts or fails to prove the SSimZK proof, Sim_a outputs the verifier's output. Otherwise it outputs Fail.

The simulator $\mathsf{Sim}_{na}(x, V^*, \rho)$ proceeds as follows:

- It encrypts the the inner state of the verifier at that point $\rho^{(1)}$ under QFHE with public key pk ($\mathsf{ct}_{\rho^{(1)}} = \mathsf{QFHE.Enc}(\mathsf{pk}, \rho^{(1)})$) and proceeds to compute

$$\mathsf{ct}_{\mathsf{cds}_R} \leftarrow \mathsf{QFHE.Eval}(\mathsf{pk}, \mathsf{SRP\text{-}CDS.R}(\mathsf{ct}_0, (\cdot, r_{\mathsf{cmt}_y})), \mathsf{td}) \text{ and}$$

$$\mathsf{ct}_{\mathsf{cmt}_y} \leftarrow \mathsf{QFHE.Eval}(\mathsf{pk}, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), \cdot\,; r_{\mathsf{cmt}_y}), \mathsf{td}).$$

- Then, Sim_{na} continues to homomorphically evaluate the verifier's response

$$(\mathsf{ct}_{arg3}, \mathsf{ct}_{\rho^{(2)}}) \leftarrow \mathsf{QFHE.Eval}(\mathsf{pk}, V^*, ((\mathsf{ct}_{\mathsf{cds}_R}, \mathsf{ct}_{\mathsf{cmt}_y}), \mathsf{ct}_{\rho^{(1)}})),$$

  where ct_{arg3} is the third message of the verifier V^* and $\mathsf{ct}_{\rho^{(2)}}$ is the new inner state of V^*, both encrypted.

- From the encrypted result of the $\mathsf{SRP\text{-}CDS.V}$, given that y was equal to td (under encryption), it gets $\mathsf{QFHE.Enc}(s)$ by running $\mathsf{SRP\text{-}CDS.D}$ homomorphically. Thus, it can compute $(r, \mathsf{sk}') \leftarrow \widetilde{\mathbf{CC}}(\mathsf{QFHE.Enc}(s))$.
- Subsequently, Sim_{na} checks the validity of $(\mathsf{pk}', \mathsf{sk}) = \mathsf{QFHE.Gen}(1^\lambda; r)$. If $\mathsf{pk}' \neq \mathsf{pk}$ or $\mathsf{sk} \neq \mathsf{sk}'$ then it halts the simulation. Otherwise it obtains the inner state of the verifier by decrypting, using the acquired secret key. The simulator also simulates the missing transcript in the second message with the same values and randomnesses used in the homomorphic computations.
- Lastly, Sim_{na} continues with the protocol by computing and sending the WI proof. It uses as witness the randomness r.

The simulator $\mathsf{Sim}_{comb}(x, V^*, \rho)$ proceed as follows:

- First, it samples $b \leftarrow\!\!\$ \{a, na\}$.
- Then it runs $\mathsf{Sim}_b(x, V^*, \rho)$.

At last, $\mathsf{Sim}(x, V^*, \rho)$ proceeds as follows:

- Generates the circuit $\mathsf{Sim}_{comb,x,V^*}$, which is the implementation of Sim_{comb} with the inputs x, V^* hardwired, so that ρ is the only input.
- The output of the simulation is $\mathcal{R}(\mathsf{Sim}_{comb,x,V^*}, \rho, \lambda)$, where $\mathcal{R}$ is the algorithm from Lemma 3.3.

Proposition 3.4. *[Similarity of Aborting Plan] Let $V^* = V^*_\rho$ be an unbounded quantum verifier and let $\mathsf{OUT}_{V^*_a}$ be the verifier's output at the end of the protocol such that if V^* does not abort the output is Fail. We show that*

$$\left\{\mathsf{OUT}_{V^*_a}\left(P(|w\rangle^{\otimes p(\lambda)}, x), V^*(\rho, x)\right)\right\}_{\lambda,x,w} \approx_s \{\mathsf{Sim}_a(x, V^*, \rho))\}_{x,w},$$

where $\lambda \in \mathbb{N}, x \in \mathcal{L} \cap \{0,1\}^\lambda$ and $|w\rangle \in \mathcal{R}_\mathcal{L}(x)$.

Proof. The two distributions are identical, since both Sim_a and the prover act exactly the same up to the fourth message. In an aborting interaction the verifier would have aborted before this message. In the case of a non-aborting interaction, both outputs would be Fail. □

Proposition 3.5. *[Similarity of Non-Aborting Plan] Let V^* be an unbounded quantum verifier and let $\mathsf{OUT}_{V^*_{na}}$ be the verifier's output at the end of the protocol such that if V^* aborts the output is Fail. We show that*

$$\left\{\mathsf{OUT}_{V^*_{na}}\left(P(|w\rangle^{\otimes p(\lambda)}, x), V^*(\rho, x)\right)\right\}_{\lambda,x,w} \approx_s \{\mathsf{Sim}_{na}(x, V^*, \rho))\}_{x,w},$$

where $\lambda \in \mathbb{N}, x \in \mathcal{L} \cap \{0,1\}^\lambda$ and $|w\rangle \in \mathcal{R}_\mathcal{L}(x)$.

Proof. We consider the following hybrid distributions, which we prove that are statistically indistinguishable:

- Hybrid $\mathcal{H}_0$: This is the output distribution of Sim_{na}
- Hybrid $\mathcal{H}_1$: This process is identical to the above except that in the WI proof the simulator proves the first statement ($x \in \mathcal{L}$). Indistinguishability follows from the witness-indistinguishability property of the WI proofs.
- Hybrid $\mathcal{H}_2$: This hybrid process is identical to the above except that if the verifier's messages are not explainable and its $\mathsf{SSim\text{-}ZK}$ proof fails, then the process chooses to fail and outputs Fail. Otherwise, after performing the homomorphic computations, instead of getting the sk from the $\widetilde{\mathbf{CC}}$, it computes it inefficiently. It also computes s inefficiently and if $\mathsf{QFHE.Dec}(\mathsf{QFHE.Enc}(s)) \neq s$ (where $\mathsf{QFHE.Enc}(s)$ is part of ct_{arg_3}) it outputs Fail. Else, it continues with the simulation.

  Statistical indistinguishability will follow from the perfect correctness of the CC obfuscation, the perfect correctness of the QFHE and from the soundness of the $\mathsf{SSim\text{-}ZK}$ that the verifier sends. Assume that the two distributions are distinguishable and fix a partial transcript T' and a verifier's inner state $\rho^{(1)}$ that maximize the distinguishability.

  - In case T' is not expainable, the $\mathsf{SSim\text{-}ZK}$ will fail, and so will the hybrid process, resulting in a contradiction (since both outputs would be Fail).
  - In case T' is expainable, in both hybrids we can check if s is correct after obtaining the sk, either inneficiently or through $\widetilde{\mathbf{CC}}$. Hence, the statistical distance between them is bounded by the probability that the check in one hybrid process fails and succeeds in the other, which in turn is bounded by the

result of the SRP-CDS not being equal to s. Given the statistical correctness of the QFHE scheme (under which the homomorphic evaluations are performed), this leads to a contradiction.

- Hybrid $\mathcal{H}_3$: In this hybrid distribution we get rid of the homomorphic evaluation altogether. If the verifier's messages are explainable (and thus specifically $\mathsf{QFHE.Enc}(\mathsf{td})$) then the simulator sends cmt_y and ct_R in the clear (similar as in the original protocol, using td in place of $y = 0^\kappa$). If the verifier's SRP-CDS decrypted is equal to the precomputed s then the process continues. Otherwise, the process outputs Fail.

  Statistical indistinguishability follows directly from the perfect correctness of the QFHE and the soundness of the SSim-ZK. Assume that the two distributions are distinguishable and fix a partial transcript T' and a verifier's inner state $\rho^{(1)}$ that maximize the distinguishability.

  - In case T' is not explainable, the SSim-ZK as well as the hybrid process would output Fail, resulting in a contradiction.
  - In case T' is explainable, the difference in the distributions is that in one the verifier's response is computed homomorphically and in the other in the clear. By the QFHE correctness, this leads to a contradiction.

- Hybrid $\mathcal{H}_4$: This process is identical to the previous except that the simulator does not check the verifier's SRP-CDS response and always continues with the process. Assume that the two distributions are distinguishable and fix a partial transcript T'.

  - If T' is not explainable then both hybrids would output Fail and are thus identical.
  - If T' is explainable and the result of the SRP-CDS is equal to s then the hybrids are identical. Alternatively, if the result of the SRP-CDS is not equal to s, then $\mathcal{H}_3$ would output Fail, but, in the current hybrid, due to the correctness of the SRP-CDS, T' should not be explainable. Given the soundness of the SSim-ZK we reach a contradiction.

- Hybrid $\mathcal{H}_5$: This hybrid process is identical to the previous except that the prover always sends $\mathsf{cmt}_y \leftarrow \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), 0^\lambda; r_{\mathsf{cmt}_y})$ in the second message. Statistical indistinguishability follows from the statistical hiding of the commitment and the SRP-CDS statistical privacy. Assume towards contradiction that the distributions are distinguishable and fix a partial transcript T'.

  - If T' is not explainable then both hybrids would output Fail thanks to the soundness of the SSim-ZK.
  - In case T' is explainable, we reach a contradiction due to the statistical security of the SBSH commitment and the statistical privacy of SRP-CDS.

- Hybrid $\mathcal{H}_6$: This hybrid process is identical to the previous except that instead of getting s and sk inefficiently and verifying V^*'s messages (with the Gen algorithm),

it always sends $\mathsf{ct}_R \leftarrow \mathsf{SRP\text{-}CDS.R}(\mathsf{ct}_0, (0^\lambda, r_{\mathsf{cmt}_y}))$ in the second message. Assume towards contradiction that the distributions are distinguishable and fix partial transcript T'.

- If T' is not explainable then both hybrids would output Fail thanks to the soundness of the $\mathsf{SSim\text{-}ZK}$.
- In case T' is explainable we reach a contradiction due to the statistical privacy of $\mathsf{SRP\text{-}CDS}$.

Note that this last process is exactly the output of the interaction with a prover.

□

Next we prove that the output of of a successful Sim_{comb} is indistinguishable from a real interaction. The Proposition is identical to Proposition 3.4 in [BS20], with the necessary changes in order to argue statistical zero knowledge.

Proposition 3.6. *[The output of of a successful Sim_{comb} is Indistinguishable from Real Interaction] Let V^* be a verifier. For $x \in \mathcal{L}$, let $\widetilde{\mathsf{Sim}_{comb}}(x, V^*_\lambda, \rho_\lambda)$ denote the conditional distribution of $\mathsf{Sim}_{comb}(x, V^*_\lambda, \rho_\lambda)$, conditioned on the simulation being successful. Then,*

$$\left\{\mathsf{OUT}\left(P(|w\rangle^{\otimes p(\lambda)}, x), V^*(\rho, x)\right)\right\}_{\lambda,x,w} \approx_s \left\{\widetilde{\mathsf{Sim}_{comb}}(x, V^*, \rho))\right\}_{x,w},$$

where $\lambda \in \mathbb{N}, x \in \mathcal{L} \cap \{0,1\}^\lambda$ and $|w\rangle \in \mathcal{R}_{\mathcal{L}}(x)$.

Proof. Denote the following conditional distributions.

- $\mathsf{A}_{\mathsf{Sim}}$: A conditional distribution of $\mathsf{Sim}_a(x, V^*, \rho)$, conditioned on that the output is not Fail (might be an empty distribution, if $a(x, \rho) = 0$).
- $\mathsf{S}_{\mathsf{Sim}}$: A conditional distribution of $\mathsf{Sim}_{na}(x, V^*, \rho)$, conditioned on that the output is not Fail (might be an empty distribution, if $b(x, \rho) = 1$).
- $\mathsf{A}_{\langle P, V^*\rangle} = \left\{\mathsf{A}_{\langle P, V^*\rangle\lambda}\right\}_{\lambda\in\mathbb{N}}$: A conditional distribution of $\mathsf{OUT}_{V^*_a}\langle P, V^*\rangle$, conditioned on that the output is not Fail (might be an empty distribution, if $c(x, \rho, |w\rangle^{\otimes p(\lambda)}) = 0$).
- $\mathsf{S}_{\langle P, V^*\rangle} = \left\{\mathsf{S}_{\langle P, V^*\rangle\lambda}\right\}_{\lambda\in\mathbb{N}}$:A conditional distribution of $\mathsf{OUT}_{V^*_{na}}\langle P, V^*\rangle$, conditioned on that the output is not Fail (might be an empty distribution, if $c(x, \rho, |w\rangle^{\otimes p(\lambda)}) = 1$).

where the probabilities a, b, c are defined as follows:

- $a(x, \rho)$: The probability that the simulation of $\mathsf{Sim}_a(x, V^*, \rho)$ was aborting.
- $b(x, \rho)$: The probability that the simulation of $\mathsf{Sim}_{na}(x, V^\cdot \rho)$ was aborting.
- $c(x, \rho, |w\rangle^{\otimes p(\lambda)})$: The probability that the interaction $\left(P(|w\rangle^{\otimes p(\lambda)}), V^*(\rho)\right)(x)$ was aborting.

Proof. □

The distribution $\widetilde{\mathsf{Sim}_{comb}}(x, V^*, \rho)$ is the distribution generated by outputting a sample from $\mathsf{A}_{\mathsf{Sim}}$ with probability $\frac{a(x,\rho}{1+a(x,\rho)-b(x,\rho)}$, and a sample from $\mathsf{S}_{\mathsf{Sim}}$ with probability $\frac{1-b(x,\rho}{1+a(x,\rho)-b(x,\rho)}$. In addition, the distribution $\mathsf{OUT}_{V^*}(P, V^*)(x)$ is the distribution generated by outputting a sample from $\mathsf{A}_{\langle P,V^*\rangle}$ with probability $c(x, \rho, |w\rangle^{\otimes p(\lambda)})$ and from $\mathsf{S}_{\langle P,V^*\rangle}$ with probability $1 - c(x, \rho, |w\rangle^{\otimes p(\lambda)})$. We will show that the two distributions are *statistically* indistinguishable by a hybrid argument. Consider the following distributions:

- Hyb_0: This is the distribution $\widetilde{\mathsf{Sim}_{comb}}(x, V^*, \rho)$.
- Hyb_1: This process is identical to the above with the exception that instead of sampling from $\mathsf{A}_{\mathsf{Sim}}$ with probability $\frac{a(x,\rho}{1+a(x,\rho)-b(x,\rho)}$ and from $\mathsf{S}_{\mathsf{Sim}}$ with probability $\frac{1-b(x,\rho}{1+a(x,\rho)-b(x,\rho)}$, it samples from $\mathsf{A}_{\mathsf{Sim}}$ with probability $a(x, \rho)$ and from $\mathsf{S}_{\mathsf{Sim}}$ with probability $1 - a(x, \rho)$.
- Hyb_2: This process is identical to the above, but the probability $a(x, \rho)$ is changed to $c(x, \rho, |w\rangle^{\otimes p(\lambda)})$.
- Hyb_3: This process is identical to the above except that with probability $c(x, \rho, |w\rangle^{\otimes p(\lambda)})$ the process outputs a sample from $\mathsf{A}_{\langle P,V^*\rangle}$ instead of $\mathsf{A}_{\mathsf{Sim}}$.
- Hyb_4: This process is identical to the above except that with probability $1 - c(x, \rho, |w\rangle^{\otimes p(\lambda)})$ the process outputs a sample from $\mathsf{S}_{\langle P,V^*\rangle}$ instead of $\mathsf{S}_{\mathsf{Sim}}$.

Following the proof from [BS20] while using propositions 3.4 and 3.5 we prove statistical indistinguishability between the above hybrid distributions. □

Theorem 3.7 (Zero Knowledge)**.** *Let $V^* = V^*_\rho$ be an unbounded quantum verifier. The protocol described in Figure 6.2 satisfies statistical zero-knowledge:*

$$\left\{\mathsf{OUT}\left(P(|w\rangle^{\otimes p(\lambda)}, x), V^*(\rho, x)\right)\right\}_{\lambda,x,w} \approx_s \{\mathsf{Sim}(x, V^*, \rho))\}_{x,w},$$

where $\lambda \in \mathbb{N}, x \in \mathcal{L} \cap \{0,1\}^\lambda$ and $|w\rangle \in \mathcal{R}_{\mathcal{L}}(x)$.

Proof. The proof is identical with the proof of Proposision 3.5 in [BS20], where the authors use Watrous' Rewinding Lemma for $\mathsf{Sim}_{comb,x,V^*}$, which has probability success negligibly close to $1/2$. If we denote the success probability for input ρ by $p(\rho)$ and denote $\epsilon := \mathsf{negl}(\lambda)+2^{-\lambda^{\frac{3}{4}}}, p_0 := \frac{1}{4}$ and $q := \frac{1}{2}$, the conditions for the Quantum Rewinding Lemma [cite it] are satisfied.

Thus the trace distance between $\mathsf{Sim}_{comb}(x, V^*, \rho)$ and $\mathsf{R}(\mathsf{Sim}_{comb,x,V^*,\rho}(x, V^*, \rho)) = \mathsf{Sim}(x, V^*, \rho)$ is bounded by a negligible function. Finally, observe that as proven in Proposition 3.6, $\mathsf{Sim}_{comb}(x, V^*, \rho))$ is statistically indistinguishable from $\mathsf{OUT}_{V^*}\langle P(|w\rangle^{\otimes p(\lambda)}), V^*(\rho)\rangle(x)$, which concludes the proof.

□

Chapter 7

Zero-Knowledge for QMA in the Timing Model

In this section we present two zero-knowledge arguments for QMA languages in the timing model: The first satisfies computational zero-knowledge, whereas the latter satisfies statistical zero-knowledge but requires slightly stronger assumptions.

1 Computational Zero-Knowledge

We recall the definition of average-case non-parallelizing functions. Non-parallelizing functions can be instantiated via repeated hashing or via the universal construction of [JMR20], additionally assuming an FHE scheme.

Definition 1.1 (Average-Case Non-Parallelizing Functions [BGJ+16])**.** *A function family $\{F_{\lambda,T} : \mathcal{X}_{\lambda,T} \to \mathcal{Y}_{\lambda,T}\}_{\lambda,T\in\mathbb{N}}$ is T-non-parallelizing with gap $\zeta < 1$, if for all $x \in \mathcal{X}$, $F_{\lambda,T}(x)$ can be computed in time T and there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$ and all non-uniform QPT algorithm with quantum advice $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$ of depth at most T^ζ, it holds that*

$$\Pr\left[\mathcal{A}(x;\rho) = F_{\lambda,T}(x) \mid x \leftarrow\!\!\$\, \mathcal{X}_\lambda\right] = \mu(\lambda).$$

In this work we are interested in an even stronger variant, where we assume that the above holds also against sub-exponential size (but still depth bounded) adversaries, and we refer to this variant as *sub-exponential* average-case non-parallelizing functions.

Our Protocol. We are now ready to describe our 2-round ZK argument. We assume the existence of the following building blocks (all secure against quantum adversaries):

- A circuit-private classical FHE scheme ($\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec}$).
- A sub-exponentially secure average-case T-non-parallelizing function $F : \mathcal{X} \to \mathcal{Y}$ secure against algorithms of size $O(2^\lambda)$ and depth less than T^ζ.
- A 2-round WI argument ($\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify}$) for QMA.

Our protocol is parametrized by a time-parameter T and it is formally described in Figure 7.1. Completeness follows immediately from the completeness of the 2-round WI argument.

2-Round (Computational) ZK Arguments for QMA

- **Prover Precomputation:** Sample an FHE key pair $(\mathsf{sk}, \mathsf{pk}) \leftarrow\!\$\, \mathsf{FHE.Gen}(1^\lambda)$ and an input $x' \leftarrow\!\$\, \mathcal{X}$. Compute $\alpha \leftarrow\!\$\, \mathsf{FHE.Enc}(\mathsf{pk}, x')$ and $\beta = \mathsf{FHE.Eval}(\mathsf{pk}, F, \alpha)$.
- **First Message (V $\rightarrow$ P):** Sample a uniform input $x^* \leftarrow\!\$\, \mathcal{X}$ and a first message $(\mathsf{crs}, \mathsf{td}) \leftarrow\!\$\, \mathsf{WI.Setup}(1^\lambda)$ and return (x^*, crs).
- **Second Message (P $\rightarrow$ V):** On input $p(\lambda)$-many copies of the witness $|w\rangle^{\otimes p(\lambda)}$ and a statement z, compute a WI proof $|\pi\rangle$ for the statement

$$\mathsf{stmt} = \left\{ x \in \mathcal{L} \vee \mathsf{pk} \in \mathsf{FHE.Gen}(1^\lambda) \wedge \alpha \in \mathsf{FHE.Enc}(\mathsf{pk}, x^*) \right\}$$

  using $|w\rangle^{\otimes p(\lambda)}$ as the witness. The prover sends $(|\pi\rangle, \mathsf{pk}, \alpha, \beta)$ to the verifier.
- **Verify:** The verifier accepts if the following conditions are satisfied.
  1. The prover responds before time T^ζ.
  2. $\mathsf{WI.Verify}(\mathsf{td}, |\pi\rangle, \mathsf{stmt}) = 1$.
  3. $\mathsf{FHE.Eval}(\mathsf{pk}, F, \alpha) = \beta$.

Figure 7.1: Description of a 2-round (computational) ZK argument for QMA (timing model)

Soundness. In the following we show that the protocol satisfies computational soundness.

Theorem 1.2 (Soundness)**.** *Assuming that F is sub-exponentially average-case non-parallelizable and that* $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ *is computationally sound, the protocol described in Figure 7.1 satisfies single-theorem computational soundness.*

Proof. Consider a prover (running in time less than T^ζ) that produces an accepting state $(|\pi\rangle, \mathsf{pk}, \alpha, \beta)$ for some statement $x \notin \mathcal{L}$. By the computational soundness of the WI proof, it must be the case that

$$\mathsf{pk} \in \mathsf{FHE.Gen}(1^\lambda) \wedge \alpha \in \mathsf{FHE.Enc}(\mathsf{pk}, x^*) \tag{7.1}$$

as otherwise it would produce a valid WI second message for a false statement. We can use the prover to define an algorithm that breaks the (sub-exponential) non-parellelizability of F as follows: The reduction sets x^* to be the challenge input and proceeds with the protocol in the same way as the verifier would. Once the prover returns $(\mathsf{pk}, \alpha, \beta)$, the reduction recovers the sk in time $O(2^\lambda)$, by e.g. testing all random strings of the $\mathsf{FHE.Gen}$ algorithm in parallel. Then it and uses sk to decrypt β and returns whatever the decrypted message is.

Observe that, by Equation (1), α is indeed an encryption of x^*. By the evaluation correctness of the FHE scheme, we have that

$$\mathsf{FHE.Dec}(\mathsf{sk}, \beta) = \mathsf{FHE.Dec}(\mathsf{sk}, \mathsf{FHE.Eval}(\mathsf{pk}, F, \alpha)) = F(x^*).$$

Thus, the reduction returns the correct output. What is left to be shown is that the depth of the reduction is asymptotically smaller than T^{ζ}. Observe that the process of recovering sk can be computed by a circuit of depth $O(\lambda)$, by testing all random coins of the $\mathsf{FHE.Gen}$ in parallel and then selecting the matching secret key with a binary tree. The depth of the decryption procedure is bounded by a fixed polynomial in λ and is in particular independent of T. Thus, the depth of the reduction is only an additive term $\mathsf{poly}(\lambda)$ higher than the depth of the prover. For a large enough T, this contradicts the non-parallelizability of F. □

Zero-Knowledge. Finally, we show that the scheme satisfies zero-knowledge in the timing model. Recall that in the timing model [DS02] the simulator is allowed to "freeze time" while simulating the accepting transcript.

Theorem 1.3 (Zero-Knowledge)**.** *Assuming that* ($\mathsf{FHE.Gen}, \mathsf{FHE.Enc}, \mathsf{FHE.Eval}, \mathsf{FHE.Dec}$) *is semantically secure, the protocol described in Figure 7.1 satisfies computational zero-knowledge in the timing model.*

Proof. The simulator computes α as an encryption of x^*, then it computes β as $\mathsf{FHE.Eval}(\mathsf{pk}, F, \alpha)$ and uses the corresponding random coins as a witness to compute the WI argument. Recall that the simulator is allowed to perform computations without letting time elapsing, from the perspective of the verifier. To show that the simulation is computationally indistinguishable from the real proof, we consider the following hybrid distributions.

- Hybrid $\mathcal{H}_0$: This is the honestly computed proof $(|\pi\rangle, \mathsf{pk}, \alpha, \beta)$.
- Hybrid $\mathcal{H}_1$: Here we change α to be the encryption of x^*, instead of x'. Computational indistinguishability follows immediately from the semantic security of the FHE scheme.
- Hybrid $\mathcal{H}_2$: Here we use the random coins used to sample pk and to encrypt α to compute the WI proof, as opposed to the witness $|w\rangle^{\otimes p(\lambda)}$. By the statistical indistinguishability of the WI argument, the distributions are statistically close.

The proof is concluded by observing that the distribution induced by $\mathcal{H}_2$ is the same as the one induced by the simulator. □

2 Statistical Zero-Knowledge

We show a different protocol that achieves statistical zero-knowledge at the cost of requiring slightly stronger assumptions, namely the existence of a post-quantum time-lock puzzle. At present, we only know how to construct (presumably) post-quantum time-lock puzzles from succinct randomized encodings [BGJ+16].

Time-Lock Puzzles. We recall the definition of time-lock puzzles [RSW96] in the following.

Definition 2.1 (Time-Lock Puzzles)**.** *A time-lock puzzle* ($\mathsf{TLP.Gen}, \mathsf{TLP.Solve}$) *consists of the following efficient algorithms.*

- $\mathsf{TLP.Gen}(1^\lambda, T, m)$: *On input the security parameter, a time parameter T, and a message m, the puzzle generation algorithm returns a puzzle Z.*
- $\mathsf{TLP.Solve}(Z)$: *On input a puzzle Z, the solving algorithm returns a message m.*

In terms of efficiency, we only require that the algorithm $\mathsf{TLP.Gen}$ runs in time polynomial in λ and at most logarithmic in T. Whereas for correctness, we require that for all $\lambda \in \mathbb{N}$, all polynomials T, all messages m it holds that

$$\mathsf{TLP.Solve}(\mathsf{TLP.Gen}(1^\lambda, T, m)) = m$$

and the algorithm $\mathsf{TLP.Solve}$ runs in time linear in T. We recall the definition of security below.

Definition 2.2 (Sequentiality [BGJ+16])**.** *A time-lock puzzle* $(\mathsf{TLP.Gen}, \mathsf{TLP.Solve})$ *is T-sequential with gap $\zeta < 1$ if there exists a negligible function μ such that for all $\lambda \in \mathbb{N}$ and all non-uniform QPT algorithm with quantum advice $\mathcal{A} = \{\mathcal{A}_\lambda, \rho_\lambda\}_{\lambda\in\mathbb{N}}$ of depth at most T^ζ, it holds that*

$$\Pr\left[\mathcal{A}(Z;\rho) = b \;\middle|\; b \leftarrow\!\$\, \{0,1\}; Z \leftarrow\!\$\, \mathsf{TLP.Gen}(1^\lambda, T, b)\right] = 1/2 + \mu(\lambda).$$

Our Protocol. Let $\varepsilon(\lambda)$ be a (fixed) negligible function. We assume the existence of the following building blocks (all secure against quantum adversaries):

- A perfectly binding commitment Com which is hiding with $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ advantage.
- A 2-round WI argument $(\mathsf{WI.Setup}, \mathsf{WI.Prove}, \mathsf{WI.Verify})$ for QMA with statistical witness indistinguishability and $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ soundness error.
- A two-round statistically sender private conditional disclosure of secrets scheme $(\mathsf{SSP\text{-}CDS.R}, \mathsf{SSP\text{-}CDS.S}, \mathsf{SSP\text{-}CDS.D})$ for NP with $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ receiver security.
- A time-lock puzzle $(\mathsf{TLP.Gen}, \mathsf{TLP.Solve})$ T-sequential with advantage bounded by $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$.
- An SBSH commitment scheme $(\mathsf{SBSH.Gen}, \mathsf{SBSH.Key}, \mathsf{SBSH.Com})$ that satisfies $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding.

Where $\mu(\lambda)$ is some negligible function. Note that, with the exception of the time-lock puzzles, all other building blocks can be instantiated assuming the quantum hardness of quasi-polynomial LWE. We define T to be the time parameter of the scheme and we describe our protocol in Figure 7.2.

Soundness. We show that our protocol satisfies (non-adaptive) soundness.

Theorem 2.3 (Soundness)**.** *Assuming the quantum quasi-polynomial hardness of the LWE problem and quasi-polynomially T-sequential time-lock puzzles, the ZK argument described in Figure 7.2 satisfies computational soundness.*

Proof. We show that the success probability of the prover is bounded by a negligible function $\varepsilon(\lambda)$. Let $x \notin \mathcal{L}$ be the false statement and let Cheat be the event where the prover causes the verifier to accept x. Assume towards contradiction that

$$\Pr[\mathsf{Cheat}] \geq \varepsilon(\lambda).$$

2-Round (Statistical) ZK Arguments for QMA

- **First Message ($\mathbf{V} \to \mathbf{P}$):** Sample an SBSH commitment key $\mathsf{ck}_0 \leftarrow\!\$\, \mathsf{SBSH.Gen}(1^\lambda)$ and the first messages of two WI arguments $(\mathsf{crs}_1, \mathsf{td}_1), (\mathsf{crs}_2, \mathsf{td}_2) \leftarrow\!\$\, \mathsf{WI.Setup}(1^\lambda)$. Sample a uniform $c \leftarrow\!\$\, \mathsf{Com}(1^\lambda, 0; r)$ and compute $Z \leftarrow\!\$\, \mathsf{TLP.Gen}(1^\lambda, T, r; \tilde{r})$ and compute the first message of a CDS $\mathsf{ct} \leftarrow\!\$\, \mathsf{SSP\text{-}CDS.R}(1^\lambda, (r, \tilde{r}))$. Return $(\mathsf{ck}_0, \mathsf{crs}_1, \mathsf{crs}_2, c, Z, \mathsf{ct})$.
- **Second Message ($\mathbf{P} \to \mathbf{V}$):** On input $p(\lambda)$-many copies of the witness $|w\rangle^{\otimes p(\lambda)}$ and a statement x, compute a WI proof $|\pi_1\rangle$ (with respect to crs_1) for the statement

$$\mathsf{stmt}_1 = \left\{ x \in \mathcal{L} \vee \exists\ r : c = \mathsf{Com}(1^\lambda, 0 : r) \right\}.$$

  Sample $\mathsf{otk} \leftarrow\!\$\, \mathsf{QOTP.Gen}(1^\lambda)$ and calculate $|\psi\rangle = \mathsf{QOTP.Enc}(\mathsf{otk}, |\pi_1\rangle)$. Then sample $\mathsf{ck}_1 \leftarrow\!\$\, \mathsf{SBSH.Key}(\mathsf{ck}_0)$ and compute $\mathsf{cmt} \leftarrow\!\$\, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), \mathsf{otk})$. Compute the CDS message ct' for otk, conditioned the statement

$$\mathsf{stmt}_0 = \left\{ \exists\ r : Z \in \mathsf{TLP.Gen}(1^\lambda, T, r) \wedge c = \mathsf{Com}(1^\lambda, 0; r) \right\}.$$

  Finally compute WI proof $|\pi_2\rangle$ (with respect to crs_2) for the statement

$$\mathsf{stmt}_2 = \left\{ x \in \mathcal{L} \vee \mathsf{cmt} \in \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), \mathsf{otk}) \wedge \mathsf{ct}' \in \mathsf{SSP\text{-}CDS.S}(\mathsf{ct}, \mathsf{stmt}_0, \mathsf{otk}) \right\}$$

  using the witness for the second branch. Return $(\mathsf{ck}_1, \mathsf{cmt}, \mathsf{ct}', |\psi\rangle, |\pi_2\rangle)$.
- **Verify:** The verifier computes $\mathsf{otk}' = \mathsf{SSP\text{-}CDS.D}(\mathsf{ct}', (r, \tilde{r}))$ and $|\pi_1\rangle = \mathsf{QOTP.Dec}(\mathsf{otk}', |\psi\rangle)$. The verifier accepts if the following conditions are satisfied.
  1. The prover responds before time T^ζ.
  2. $\mathsf{WI.Verify}(\mathsf{td}_1, |\pi_1\rangle, \mathsf{stmt}_1) = 1$.
  3. $\mathsf{WI.Verify}(\mathsf{td}_2, |\pi_2\rangle, \mathsf{stmt}_2) = 1$.

Figure 7.2: Description of a 2-round (statistical) ZK argument for QMA (timing model)

Then, by the $(\varepsilon(\lambda), \varepsilon(\lambda)^2)$-sometimes binding property of the SBSH commitment scheme, we have that

$$\Pr\left[\mathsf{Cheat} \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

for some negligible function $\mu(\lambda)$. Let $\tilde{\mathsf{otk}} = \mathsf{SBSH.Ext}(r, \mathsf{ck}_0, \mathsf{ck}_1, \mathsf{cmt})$ be the output of the extractor, where r denote the random coins used in the $\mathsf{SBSH.Gen}$ algorithm. We now gradually change the verification procedure and we argue that the probability that the above defined event happens does not decrease significantly.

- The verifier computes $|\pi_1\rangle = \mathsf{QOTP.Dec}(\tilde{\mathsf{otk}}, |\psi\rangle)$, instead of recovering otk' from the CDS protocol. Let us now define Cheat_1 as the event where the modified verifier accepts. We want to argue that

$$\Pr\left[\mathsf{Cheat}_1 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

for some negligible function $\mu(\lambda)$. Note that the events Cheat and Cheat_1 only differ in the case where $\tilde{\mathsf{otk}} \neq \mathsf{otk}'$. Thus if the inequality above does not hold, we obtain a contradiction against the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-soundness of the WI argument.

- The verifier computes $\mathsf{ct} \leftarrow\!\!\$\, \mathsf{SSP\text{-}CDS.R}(1^\lambda, 0)$ and we define Cheat_2 as the event where the modified verifier accepts. We can show that

$$\Pr\left[\mathsf{Cheat}_2 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

  by a reduction against the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$-receiver hiding of the CDS scheme.

- The verifier computes $Z \leftarrow\!\!\$\, \mathsf{TLP.Gen}(1^\lambda, T, 0)$ and we define Cheat_3 as the event where the modified verifier accepts. We can show that

$$\Pr\left[\mathsf{Cheat}_3 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

  by a reduction against the T-sequentiality of the time-lock puzzle with advantage $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$.

- The verifier computes $c \leftarrow\!\!\$\, \mathsf{Com}(1^\lambda, 1)$ and we define Cheat_4 as the event where the modified verifier accepts. We have that

$$\Pr\left[\mathsf{Cheat}_4 \wedge (\mathsf{ck}_0, \mathsf{ck}_1) \in \mathsf{Binding}\right] \geq \varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$$

  by the $\varepsilon(\lambda)^2 \cdot \mu(\lambda)$ hiding of the commitment scheme Com.

The last inequality implies that the prover produces a valid $|\pi_1\rangle$ for a false statement with probability greater than $\varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$, which is a contradiction to the $\varepsilon(\lambda)^2 \cdot (1 + \mu(\lambda))$-soundness of the WI argument and concludes the proof. □

Zero-Knowledge. We now argue that the protocol is zero-knowledge in the timing model.

Theorem 2.4 (Zero-Knowledge)**.** *The protocol described in Figure 7.2 satisfies statistical zero-knowledge in the timing model.*

Proof. The simulator recovers a randomness r from Z (by computing $\mathsf{TLP.Solve}$) and checks whether $c = \mathsf{Com}(1^\lambda, 0; r)$. If this is the case it uses it as the witness to compute $|\pi_1\rangle$, otherwise it sets $|\pi_1\rangle$ to be the all 0 state (padded to the appropriate length). The simulator then proceeds as in the real protocol.

To show that the transcript produced by the simulator is statistically close to the one produced by the real prover, we consider the following hybrid distributions.

- Hybrid $\mathcal{H}_0$: This is the simulated transcript.

- Hybrid $\mathcal{H}_1$: We change the simulation to compute $|\pi_1\rangle$ using the real witness of the statement z, but only in the case where $c = \mathsf{Com}(1^\lambda, 0; r)$. By the statistical witness indistinguishability of the WI argument, this change is statistically indistinguishable.

- Hybrid $\mathcal{H}_2$: Here we compute $|\pi_2\rangle$ using the real witness of the statement z. This change is statistically indistinguishable to the eyes of the verifier by the statistical witness indistinguishability of the WI argument.
- Hybrid $\mathcal{H}_3$: If $c \neq \mathsf{Com}(1^\lambda, 0; r)$ we compute the CDS second message ct' with the message fixed to 0 (padded to the appropriate length), instead of otk. Note that the condition $c \neq \mathsf{Com}(1^\lambda, 0; r)$ implies that stmt is false, and therefore the distribution induced by this hybrid is statistically close to that of the previous one.
- Hybrid $\mathcal{H}_4$: If $c \neq \mathsf{Com}(1^\lambda, 0; r)$ we compute $\mathsf{cmt} \leftarrow\!\!\$\, \mathsf{SBSH.Com}((\mathsf{ck}_0, \mathsf{ck}_1), 0)$. This change is statistically indistinguishable by the statistical hiding property of the SBSH commitment.
- Hybrid $\mathcal{H}_5$: If $c \neq \mathsf{Com}(1^\lambda, 0; r)$ we compute $|\psi\rangle = \mathsf{QOTP}(\mathsf{otk}, |\pi_1\rangle)$, where $|\pi_1\rangle$ is computed using the real witness for z. To the eyes of the distinguisher $|\psi\rangle$ is now maximally mixed and therefore this distribution is identical to that of the previous hybrid.
- Hybrid $\mathcal{H}_6$: We revert the change done in $\mathcal{H}_4$.
- Hybrid $\mathcal{H}_7$: We revert the change done in $\mathcal{H}_3$.
- Hybrid $\mathcal{H}_8$: We revert the change done in $\mathcal{H}_2$.

The proof is concluded by observing that $\mathcal{H}_8$ is identical to the output of the honest prover. □

Chapter 8

Conclusions and Future Work

1 Conclusions

With the rise of quantum computers, quantum cryptography has been gaining increasingly more attention in the past years. While some cryptographic tools from the classical setting are easy to generalize, many quantum protocols are far from reaching their classical counterparts' efficiency. In this work we study communication complexity of Quantum FHE protocols and ZK arguments for QMA we are able to match it with the one of their corresponding classical protocols.

In the case of the Fully Homomorphic Encryption Scheme, we present two constructions of Rate-1 quantum FHE, where the amount of information exchanged between the parties is optimal. In the first one, assuming a quantum encryption scheme with hybrid ciphertexts (with both quantum and classical information) we utilize a classical rate-1 FHE and key-switch the classical information from the original scheme to the rate-1 one, managing to achieve communication complexity $(|\,|\psi\rangle\,| + |C(|\psi\rangle)|) \cdot (1 + o(1))$. In order to continue doing homomorphic ea=valuations, we can switch back to the original scheme. In the second construction, we approach the problem non-generically, and construct a scheme to hold the classical information in the hybrid ciphertext that is itself rate-1, avoiding the two-key cycle needed to obtain full (as opposed to leveled) homomorphism.

Regarding the zero-knowledge arguments, we first create a 2-round statistical WI argument for QMA based on the quasi-polynomial hardness of LWE. Based on that, we are able to create a protocol that achieves statistical ZK in only four rounds with the same assumption. In addition, by transferring the protocol to the timing model, we were able to achieve ZK in two rounds (both computationally and statistically with stronger assumptions).

2 Future Work

The above results can be used as motivation for further work, transferring them to different settings or assuming different parameters.

One way to do that would be to construct a Rate-1 Multi-Key Quantum FHE scheme, combining our results with the quantum multi-key FHE in [ABG+20] that also utilizes hybrid encryption schemes (as in [Mah18a]). Another approach would be

to extend out results to verifiable FHE [ADSS17], where we are able to prove some property about the message while it remains encrypted and not reveal any additional information.

Concerning ZK, extending our results to ZK proofs (instead of arguments), where the protocol provides statistical soundness is an open problem. It is also interesting to produce the same communication complexity assuming the polynomial security of LWE (instead of quasi-polynomial), which is one of the optimal assumptions in quantum cryptography so far. Recent works have also focused on black-box approaches to post-quantum ϵ-zero-knowledge in constant rounds [CCY20a, CCLY21a], which could be reduced to 4-rounds following our techniques. Lastly, there is yet to be constructed a 3-round ZK argument in the plain model, surpassing the so far optimal communication complexity of ZK protocols.

We would also be interested to investigate how such protocols might be applied to cryptographic voting, as homomorphic techniques and zero-knowledge proofs are an important ingredient of several earlier and recent protocols [SK94, KY02, GPZZ18, GPZZ21].

Bibliography

[ABG+20] Amit Agarwal, James Bartusek, Vipul Goyal, Dakshita Khurana, and Giulio Malavolta. Post-quantum multi-party computation. Cryptology ePrint Archive, Report 2020/1395, 2020. https://eprint.iacr.org/2020/1395.

[ACGH20] Gorjan Alagic, Andrew M. Childs, Alex B. Grilo, and Shih-Han Hung. Non-interactive classical verification of quantum computation. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 153–180. Springer, Heidelberg, November 2020.

[ACP20] Prabhanjan Ananth, Kai-Min Chung, and Rolando L. La Placa. On the concurrent composition of quantum zero-knowledge. Cryptology ePrint Archive, Report 2020/1528, 2020. https://eprint.iacr.org/2020/1528.

[ADSS17] Gorjan Alagic, Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum fully homomorphic encryption with verification. pages 438–467, 11 2017.

[AIK04] Benny Applebaum, Yuval Ishai, and Eyal Kushilevitz. Cryptography in NC^0. In *45th FOCS*, pages 166–175. IEEE Computer Society Press, October 2004.

[AIR01] William Aiello, Yuval Ishai, and Omer Reingold. Priced oblivious transfer: How to sell digital goods. In Birgit Pfitzmann, editor, *EUROCRYPT 2001*, volume 2045 of *LNCS*, pages 119–135. Springer, Heidelberg, May 2001.

[AL20] Prabhanjan Ananth and Rolando L. La Placa. Secure quantum extraction protocols. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 123–152. Springer, Heidelberg, November 2020.

[AMTDW00] Andris Ambainis, Michele Mosca, Alain Tapp, and Ronald De Wolf. Private quantum channels. In *Proceedings 41st Annual Symposium on Foundations of Computer Science*, pages 547–553. IEEE, 2000.

[ARU14] Andris Ambainis, Ansis Rosmanis, and Dominique Unruh. Quantum attacks on classical proof systems: The hardness of quantum rewinding.

In *55th FOCS*, pages 474–483. IEEE Computer Society Press, October 2014.

[BCKM20] James Bartusek, Andrea Coladangelo, Dakshita Khurana, and Fermi Ma. On the round complexity of two-party quantum computation. Cryptology ePrint Archive, Report 2020/1471, 2020. `https://eprint.iacr.org/2020/1471`.

[BD18] Zvika Brakerski and Nico Döttling. Two-message statistically sender-private OT from LWE. In Amos Beimel and Stefan Dziembowski, editors, *TCC 2018, Part II*, volume 11240 of *LNCS*, pages 370–390. Springer, Heidelberg, November 2018.

[BDGM19] Zvika Brakerski, Nico Döttling, Sanjam Garg, and Giulio Malavolta. Leveraging linear decryption: Rate-1 fully-homomorphic encryption and time-lock puzzles. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019, Part II*, volume 11892 of *LNCS*, pages 407–437. Springer, Heidelberg, December 2019.

[BFJ+20] Saikrishna Badrinarayanan, Rex Fernando, Aayush Jain, Dakshita Khurana, and Amit Sahai. Statistical ZAP arguments. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 642–667. Springer, Heidelberg, May 2020.

[BFK09] Anne Broadbent, Joseph Fitzsimons, and Elham Kashefi. Universal blind quantum computation. In *50th FOCS*, pages 517–526. IEEE Computer Society Press, October 2009.

[BG20] Anne Broadbent and Alex B. Grilo. QMA-hardness of consistency of local density matrices with applications to quantum zero-knowledge. In *61st FOCS*, pages 196–205. IEEE Computer Society Press, November 2020.

[BGJ+16] Nir Bitansky, Shafi Goldwasser, Abhishek Jain, Omer Paneth, Vinod Vaikuntanathan, and Brent Waters. Time-lock puzzles from randomized encodings. In Madhu Sudan, editor, *ITCS 2016*, pages 345–356. ACM, January 2016.

[BHR12] Mihir Bellare, Viet Tung Hoang, and Phillip Rogaway. Foundations of garbled circuits. In Ting Yu, George Danezis, and Virgil D. Gligor, editors, *ACM CCS 2012*, pages 784–796. ACM Press, October 2012.

[BJ15] Anne Broadbent and Stacey Jeffery. Quantum homomorphic encryption for circuits of low T-gate complexity. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part II*, volume 9216 of *LNCS*, pages 609–629. Springer, Heidelberg, August 2015.

[BJSW16] Anne Broadbent, Zhengfeng Ji, Fang Song, and John Watrous. Zero-knowledge proof systems for QMA. In Irit Dinur, editor, *57th FOCS*, pages 31–40. IEEE Computer Society Press, October 2016.

[BKP19] Nir Bitansky, Dakshita Khurana, and Omer Paneth. Weak zero-knowledge beyond the black-box barrier. In Moses Charikar and Edith Cohen, editors, *51st ACM STOC*, pages 1091–1102. ACM Press, June 2019.

[BM21] James Bartusek and Giulio Malavolta. Candidate obfuscation of null quantum circuits and witness encryption for qma. Cryptology ePrint Archive, Report 2021/421, 2021. `https://eprint.iacr.org/2021/421`.

[Bra18] Zvika Brakerski. Quantum FHE (almost) as secure as classical. In Hovav Shacham and Alexandra Boldyreva, editors, *CRYPTO 2018, Part III*, volume 10993 of *LNCS*, pages 67–95. Springer, Heidelberg, August 2018.

[BS20] Nir Bitansky and Omri Shmueli. Post-quantum zero knowledge in constant rounds. In Konstantin Makarychev, Yury Makarychev, Madhur Tulsiani, Gautam Kamath, and Julia Chuzhoy, editors, *52nd ACM STOC*, pages 269–279. ACM Press, June 2020.

[BV11] Zvika Brakerski and Vinod Vaikuntanathan. Efficient fully homomorphic encryption from (standard) LWE. In Rafail Ostrovsky, editor, *52nd FOCS*, pages 97–106. IEEE Computer Society Press, October 2011.

[BV14] Zvika Brakerski and Vinod Vaikuntanathan. Lattice-based FHE as secure as PKE. In Moni Naor, editor, *ITCS 2014*, pages 1–12. ACM, January 2014.

[CCLY21a] Nai-Hui Chia, Kai-Min Chung, Qipeng Liu, and Takashi Yamakawa. On the impossibility of post-quantum black-box zero-knowledge in constant rounds. *ArXiv*, abs/2103.11244, 2021.

[CCLY21b] Nai-Hui Chia, Kai-Min Chung, Qipeng Liu, and Takashi Yamakawa. On the impossibility of post-quantum black-box zero-knowledge in constant rounds. Cryptology ePrint Archive, Report 2021/376, 2021. `https://eprint.iacr.org/2021/376`.

[CCY20a] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. Black-box approach to post-quantum zero-knowledge in constant round, 11 2020.

[CCY20b] Nai-Hui Chia, Kai-Min Chung, and Takashi Yamakawa. Classical verification of quantum computations with efficient verifier. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part III*, volume 12552 of *LNCS*, pages 181–206. Springer, Heidelberg, November 2020.

[CDM20] Orestis Chardouvelis, Nico Döttling, and Giulio Malavolta. Rate-1 secure function evaluation for bqp. Cryptology ePrint Archive, Report 2020/1454, 2020. `https://eprint.iacr.org/2020/1454`.

[CFGS18] Alessandro Chiesa, Michael A. Forbes, Tom Gur, and Nicholas Spooner. Spatial isolation implies zero knowledge even in a quantum world. In Mikkel Thorup, editor, *59th FOCS*, pages 755–765. IEEE Computer Society Press, October 2018.

[CM21] Orestis Chardouvelis and Giulio Malavolta. The round complexity of quantum zero-knowledge. Cryptology ePrint Archive, Report 2021/918, 2021. https://eprint.iacr.org/2021/918.

[CVZ20] Andrea Coladangelo, Thomas Vidick, and Tina Zhang. Non-interactive zero-knowledge arguments for QMA, with preprocessing. In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part III*, volume 12172 of *LNCS*, pages 799–828. Springer, Heidelberg, August 2020.

[DGI+19] Nico Döttling, Sanjam Garg, Yuval Ishai, Giulio Malavolta, Tamer Mour, and Rafail Ostrovsky. Trapdoor hash functions and their applications. In Alexandra Boldyreva and Daniele Micciancio, editors, *CRYPTO 2019, Part III*, volume 11694 of *LNCS*, pages 3–32. Springer, Heidelberg, August 2019.

[DHRW16] Yevgeniy Dodis, Shai Halevi, Ron D. Rothblum, and Daniel Wichs. Spooky encryption and its applications. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part III*, volume 9816 of *LNCS*, pages 93–122. Springer, Heidelberg, August 2016.

[DNS10] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Secure two-party quantum evaluation of unitaries against specious adversaries. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 685–706. Springer, Heidelberg, August 2010.

[DNS12] Frédéric Dupuis, Jesper Buus Nielsen, and Louis Salvail. Actively secure two-party evaluation of any quantum operation. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 794–811. Springer, Heidelberg, August 2012.

[DS02] Cynthia Dwork and Larry J. Stockmeyer. 2-round zero knowledge and proof auditors. In *34th ACM STOC*, pages 322–331. ACM Press, May 2002.

[DSS16] Yfke Dulek, Christian Schaffner, and Florian Speelman. Quantum homomorphic encryption for polynomial-sized circuits. Cryptology ePrint Archive, Report 2016/559, 2016. https://eprint.iacr.org/2016/559.

[Gen09] Craig Gentry. Fully homomorphic encryption using ideal lattices. In Michael Mitzenmacher, editor, *41st ACM STOC*, pages 169–178. ACM Press, May / June 2009.

[GGM86] Oded Goldreich, Shafi Goldwasser, and Silvio Micali. How to construct random functions. *Journal of the ACM*, 33(4):792–807, October 1986.

[GGSW13] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th ACM STOC*, pages 467–476. ACM Press, June 2013.

[GHV10] Craig Gentry, Shai Halevi, and Vinod Vaikuntanathan. i-Hop homomorphic encryption and rerandomizable Yao circuits. In Tal Rabin, editor, *CRYPTO 2010*, volume 6223 of *LNCS*, pages 155–172. Springer, Heidelberg, August 2010.

[GJJM20] Vipul Goyal, Abhishek Jain, Zhengzhong Jin, and Giulio Malavolta. Statistical zaps and new oblivious transfer protocols. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 668–699. Springer, Heidelberg, May 2020.

[GK96] Oded Goldreich and Ariel Kahan. How to construct constant-round zero-knowledge proof systems for NP. *Journal of Cryptology*, 9(3):167–190, June 1996.

[GKVW19] Rishab Goyal, Venkata Koppula, Satyanarayana Vusirikala, and Brent Waters. On perfect correctness in (lockable) obfuscation. Cryptology ePrint Archive, Report 2019/1010, 2019. `https://eprint.iacr.org/2019/1010`.

[GKW17] Rishab Goyal, Venkata Koppula, and Brent Waters. Lockable obfuscation. In Chris Umans, editor, *58th FOCS*, pages 612–621. IEEE Computer Society Press, October 2017.

[GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM Journal on Computing*, 18(1):186–208, 1989.

[GMW86] Oded Goldreich, Silvio Micali, and Avi Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design (extended abstract). In *27th FOCS*, pages 174–187. IEEE Computer Society Press, October 1986.

[GO94] Oded Goldreich and Yair Oren. Definitions and properties of zero-knowledge proof systems. *Journal of Cryptology*, 7(1):1–32, December 1994.

[GPZZ18] Panagiotis Grontas, Aris Pagourtzis, Alexandros Zacharakis, and Bingsheng Zhang. Towards everlasting privacy and efficient coercion resistance in remote electronic voting. In *Financial Cryptography and Data Security - FC 2018 International Workshops, BITCOIN, VOTING, and WTSC*, volume 10958 of *Lecture Notes in Computer Science*, pages 210–231. Springer, 2018.

[GPZZ21] Panagiotis Grontas, Aris Pagourtzis, Alexandros Zacharakis, and Bingsheng Zhang. Publicly auditable conditional blind signatures. *Journal of Computer Security*, 29(2):229–271, 2021.

[GSW13] Craig Gentry, Amit Sahai, and Brent Waters. Homomorphic encryption from learning with errors: Conceptually-simpler, asymptotically-faster, attribute-based. In Ran Canetti and Juan A. Garay, editors,

CRYPTO 2013, Part I, volume 8042 of *LNCS*, pages 75–92. Springer, Heidelberg, August 2013.

[GSY19] Alex Bredariol Grilo, William Slofstra, and Henry Yuen. Perfect zero knowledge for quantum multiprover interactive proofs. In David Zuckerman, editor, *60th FOCS*, pages 611–635. IEEE Computer Society Press, November 2019.

[HAO15] Ryo Hiromasa, Masayuki Abe, and Tatsuaki Okamoto. Packing messages and optimizing bootstrapping in gsw-fhe. In Jonathan Katz, editor, *Public-Key Cryptography – PKC 2015*, pages 699–715, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.

[HILL99] Johan HÅstad, Russell Impagliazzo, Leonid A. Levin, and Michael Luby. A pseudorandom generator from any one-way function. *SIAM J. Comput.*, 28(4):1364–1396, March 1999.

[HM96] Shai Halevi and Silvio Micali. Practical and provably-secure commitment schemes from collision-free hashing. In Neal Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 201–215. Springer, Heidelberg, August 1996.

[HSS11] Sean Hallgren, Adam Smith, and Fang Song. Classical cryptographic protocols in a quantum world. In Phillip Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 411–428. Springer, Heidelberg, August 2011.

[HW18] Susan Hohenberger and Brent Waters. Synchronized aggregate signatures from the RSA assumption. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part II*, volume 10821 of *LNCS*, pages 197–229. Springer, Heidelberg, April / May 2018.

[IK00] Yuval Ishai and Eyal Kushilevitz. Randomizing polynomials: A new representation with applications to round-efficient secure computation. In *41st FOCS*, pages 294–304. IEEE Computer Society Press, November 2000.

[JMR20] Samuel Jaques, Hart Montgomery, and Arnab Roy. Time-release cryptography from minimal circuit assumptions. Cryptology ePrint Archive, Report 2020/755, 2020. `https://eprint.iacr.org/2020/755`.

[KKS18] Yael Tauman Kalai, Dakshita Khurana, and Amit Sahai. Statistical witness indistinguishability (and more) in two messages. In Jesper Buus Nielsen and Vincent Rijmen, editors, *EUROCRYPT 2018, Part III*, volume 10822 of *LNCS*, pages 34–65. Springer, Heidelberg, April / May 2018.

[KY02] Aggelos Kiayias and Moti Yung. Self-tallying elections and perfect ballot secrecy. In *PKC 2002*, volume 2274 of *LNCS*, pages 141–158. Springer, 2002.

[LC18] Ching-Yi Lai and Kai-Min Chung. On statistically-secure quantum homomorphic encryption. *Quantum Information and Computation*, 18:785–794, 08 2018.

[LN11] Carolin Lunemann and Jesper Buus Nielsen. Fully simulatable quantum-secure coin-flipping and applications. In Abderrahmane Nitaj and David Pointcheval, editors, *AFRICACRYPT 11*, volume 6737 of *LNCS*, pages 21–40. Springer, Heidelberg, July 2011.

[LS19] Alex Lombardi and Luke Schaeffer. A note on key agreement and non-interactive commitments. Cryptology ePrint Archive, Report 2019/279, 2019. `https://eprint.iacr.org/2019/279`.

[LVW20] Alex Lombardi, Vinod Vaikuntanathan, and Daniel Wichs. Statistical ZAPR arguments from bilinear maps. In Anne Canteaut and Yuval Ishai, editors, *EUROCRYPT 2020, Part III*, volume 12107 of *LNCS*, pages 620–641. Springer, Heidelberg, May 2020.

[Mah18a] Urmila Mahadev. Classical homomorphic encryption for quantum circuits. In Mikkel Thorup, editor, *59th FOCS*, pages 332–338. IEEE Computer Society Press, October 2018.

[Mah18b] Urmila Mahadev. Classical verification of quantum computations. In Mikkel Thorup, editor, *59th FOCS*, pages 259–267. IEEE Computer Society Press, October 2018.

[MP12] Daniele Micciancio and Chris Peikert. Trapdoors for lattices: Simpler, tighter, faster, smaller. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 700–718. Springer, Heidelberg, April 2012.

[NS17] Michael Newman and Yaoyun Shi. Limitations on transversal computation through quantum homomorphic encryption. *Quantum Information and Computation*, 18, 04 2017.

[OPP14] Rafail Ostrovsky, Anat Paskin-Cherniavsky, and Beni Paskin-Cherniavsky. Maliciously circuit-private FHE. In Juan A. Garay and Rosario Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 536–553. Springer, Heidelberg, August 2014.

[OTF15] Yingkai Ouyang, Si-Hui Tan, and Joseph Fitzsimons. Quantum homomorphic encryption from quantum codes. *Physical Review A*, 98, 08 2015.

[Pas03a] Rafael Pass. On deniability in the common reference string and random oracle model. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 316–337. Springer, Heidelberg, August 2003.

[Pas03b] Rafael Pass. Simulation in quasi-polynomial time, and its application to protocol composition. In Eli Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 160–176. Springer, Heidelberg, May 2003.

[PRS17] Chris Peikert, Oded Regev, and Noah Stephens-Davidowitz. Pseudorandomness of ring-LWE for any ring and modulus. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *49th ACM STOC*, pages 461–473. ACM Press, June 2017.

[Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th ACM STOC*, pages 84–93. ACM Press, May 2005.

[RSW96] Ronald L. Rivest, Adi Shamir, and David A. Wagner. Time-lock puzzles and timed-release crypto. Technical report, 1996.

[Shm20] Omri Shmueli. Multi-theorem (malicious) designated-verifier NIZK for QMA. Cryptology ePrint Archive, Report 2020/928, 2020. `https://eprint.iacr.org/2020/928`.

[SK94] Kazue Sako and Joe Kilian. Secure voting using partially compatible homomorphisms. In *Advances in Cryptology - CRYPTO '94, 14th Annual International Cryptology Conference*, volume 839 of *Lecture Notes in Computer Science*, pages 411–424. Springer, 1994.

[TKO+14] Si-Hui Tan, Joshua Kettlewell, Yingkai Ouyang, Lin Chen, and Joseph Fitzsimons. A quantum approach to homomorphic encryption. *Scientific Reports*, 6, 11 2014.

[Unr12] Dominique Unruh. Quantum proofs of knowledge. In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 135–152. Springer, Heidelberg, April 2012.

[Wat09] John Watrous. Zero-knowledge against quantum attacks. *SIAM J. Comput.*, 39(1):25–58, May 2009.

[WZ17] Daniel Wichs and Giorgos Zirdelis. Obfuscating compute-and-compare programs under LWE. In Chris Umans, editor, *58th FOCS*, pages 600–611. IEEE Computer Society Press, October 2017.

[Yao86] Andrew Chi-Chih Yao. How to generate and exchange secrets (extended abstract). In *27th FOCS*, pages 162–167. IEEE Computer Society Press, October 1986.

[Yue21] Henry Yuen. Crash course. The 11th BIU Winter School on Cryptography Lecture, 2021. `http://cyber.biu.ac.il/event/the-11th-biu-winter-school-on-cryptography/`.

[Zha18] Mark Zhandry. Quantum cryptography. University Lecture, 2018. `https://www.cs.princeton.edu/~mzhandry/2018-Fall-COS597A/`.

[Zha19] Mark Zhandry. Quantum lightning never strikes the same state twice. In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part III*, volume 11478 of *LNCS*, pages 408–438. Springer, Heidelberg, May 2019.