Proof Techniques for a Quantum World

Ritam Bhaumik, CRC, TII Abu Dhabi
(based on joint work with Jyotirmoy Basak, Amit Kumar Chauhan, Benôit Cogliati, Jordan Ethan, Ravindra Jejurikar, Ashwin Jha, Anandarup Roy, André Schrottenloher, and Suprita Talnikar)

GAPS 2025
September 2, 2025
Singapore

Outline

Outline

Outline

Outline

Outline

The game of Penultima

- A game of chess involving ‘spectators’
- Spectators create secret custom rules modifying how pieces move and capture
- Players find out which moves are legal through trial and error
- The goal is to figure out the rules (but also to win!)

Navigating a Quantum World

- Imagine that you are a symmetric cryptographer used to doing classical proofs
- The problem of writing proofs in the quantum world looks deceptively familiar
- But soon you learn about the new rules nobody told you about
- From then on it is a struggle to complete the proofs while respecting rules you do not fully know or understand

Symmetric, yet Post-Quantum?

- Natural question: what about the quantum experts well-versed in those new rules?
- Short answer: they don't really care about security proofs in symmetric cryptography
- It is a persistent myth that symmetric cryptography has nothing to fear from quantum adversaries
- Symmetric cryptographers are left to figure things out for themselves by floundering in the confusing quantum world

Outline

The Recording Conundrum

- Classical reduction proofs frequently rely on ‘transcripts’
- Transcripts save a record of all the queries and responses exchanged in the course of a game
- Such transcripts don’t work for a game involving quantum queries, as quantum states cannot be ‘cloned’
- This presents an immediate hurdle for translating classical proofs to post-quantum proofs

Standard Oracle

Standard trick of implementing a classical function f on a quantum channel so the operation is unitary:

$$\mathrm{stO}_f \left|x\right\rangle \left|y\right\rangle = \left|x\right\rangle \left|y \oplus f(x)\right\rangle$$

x: query register

y: response register

Equivalent formulation using truth tables:

$$\mathrm{stO} \left|x\right\rangle \left|y\right\rangle \left|T_f\right\rangle = \left|x\right\rangle \left|y \oplus T_f[x]\right\rangle \left|T_f\right\rangle$$

T_f: complete truth table of f (ignore efficiency)

Fourier Basis

Computational basis:

$$|0\rangle, |1\rangle, \ldots, |2^n - 1\rangle$$

(conventionally mapped to a canonical basis of $\mathbb{C}^{2^n}$)

Hadamard transform (ignore normalisation):

$$H_n |x\rangle = \sum_{z=0}^{2^n-1} (-1)^{x \cdot z} |z\rangle$$

Fourier basis:

$$H_n |0\rangle, H_n |1\rangle, \ldots, H_n |2^n - 1\rangle$$

Turning the Tables

$$
\begin{aligned}
U|y\rangle|z\rangle &:= |y \oplus z\rangle|z\rangle, \\
U|\hat{y}\rangle|\hat{z}\rangle &= \sum_{u,v=0}^{2^n-1}(-1)^{y\cdot u \oplus z\cdot v}U|u\rangle|v\rangle \\
&= \sum_{u,v=0}^{2^n-1}(-1)^{y\cdot u \oplus z\cdot v}|u \oplus v\rangle|v\rangle \\
&= \sum_{u,v=0}^{2^n-1}(-1)^{y\cdot(u\oplus v)\oplus(z\oplus y)\cdot v}|u \oplus v\rangle|v\rangle \\
&= |\hat{y}\rangle\left|\widehat{z \oplus y}\right\rangle
\end{aligned}
$$

Wherein lies the Magic (or so I think)

Now observe how the standard oracle acts on the Fourier basis:

$$\mathsf{stO}\,|x\rangle\,|\hat{y}\rangle\left|\widehat{T_f}\right\rangle = |x\rangle\,|\hat{y}\rangle\left|\widehat{T_{f\oplus\delta_{xy}}}\right\rangle$$

where

$$\begin{aligned}\delta_{xy}(z) &= y \text{ when } z = x,\\ &= 0 \text{ elsewhere}\end{aligned}$$

For all intents and purposes, it looks like the standard oracle modifies one cell in the truth table!

'Databases', at last!

- The truth table of a partial function defined at q points = a database with q entries
- A partial function defined at q points = a lazily sampled function queries q times
- Database = fancy rebranding of our old friend Transcript
- Modifying an empty cell of a truth table $\approx$ adding a new entry to the database
- With this shift in perspective, we can now leave the game untouched and still pretend that queries are being recorded!

Outline

Transition Capacity Formalism

- Properties are predicates satisfied by certain databases
- Examples include containing a collision pair or a zero-preimage
- **Transition capacity** is (loosely) the square root of the probability of acquiring a (new) property after the next query
- Example of a transition into a property could be a collision-free database gaining a collision on the next query

Limiting 'Bad' Transitions

- Consider a certain 'bad' property P (e.g., having a collision) and a database D not satisfying P
- Identify a set S of possible responses on the next query which can lead to D transitioning into P
- For the collision example, S would be the range of the partial function already sampled and stored in the database
- Then we can show that for the transition of D into P,

$$\text{transition capacity} \leq O\left(\sqrt{\frac{|S|}{2^n}}\right)$$

Two-Domain Distance Bounds

- Consider a distinguishing game between a real world and an ideal world (defined on different domains)
- Each world records all intermediate primitive queries into corresponding databases
- Suppose we identify bad properties for both worlds and show that as long as the databases in neither world transitions into bad, they continue to evolve identically
- Then the Q2 distinguishing advantage between the two worlds can be upper bounded by a sum transition capacities corresponding to bad transitions in either world in different stages of the game

Example: TNT

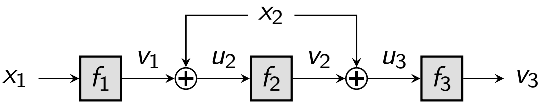

- Bad property: a collision at u_3, i.e., an entry (u_3, v_3) in the database of f_3 which 'corresponds' to two distinct queries (x_1, x_2) and (x_1', x_2')
- Transition to bad can occur when adding an entry (u_2, v_2) to f_2 for certain values of v_2

Example: TNT (cont'd)

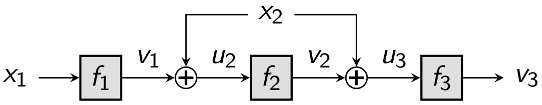

- Note that we don't have a way of recording which entries in other databases led to a particular entry (u_3, v_3)
- Thus for bounding bad transition capacities all possible cross-combinations need to be checked
- This leads to an unfortunate quadratic blowup which we currently don't know how to avoid

The Gap

- Consider 3-round Feistel, where we believe the right half of the output should behave like a qPRF output
- To apply the Two-Domain Distance Technique, we would need to classify collisions in x_3 (the input of f_3) as bad
- Now, $x_3 = x_1 \oplus f_2(x_2)$
- Because of the blowup, we need to consider the combination of all x_1 with all entries of the database for f_2
- But future values of x_1 come directly from the adversary :(

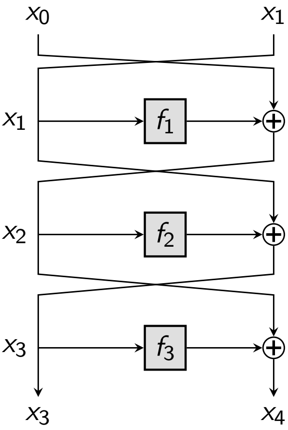

Verdict on Q2

- We have begun taking baby steps in understanding how symmetric provable security in the Q2 model should look like
- Numerous serious obstacles still lying ahead, e.g., we don't yet know how to lazily sample random permutations
- Bounds are also terrible, owing to the quadratic blowup from the previous slide and other factors
- Silver lining: proofs of a classical counting flavour finally beginning to take shape

Outline

Dialling it Down a Notch?

- Now let's return to a less ambitious but more practically useful security model
- In the Q1 model, the adversary has a quantum computer at home, so can make superposition queries to public primitives
- The communication channel is still classical, so superposition queries cannot be made to the keyed construction
- Question: how far can classical public-primitive proofs be lifted to the Q1 model?

Constructing Hybrids

- We divide the game into *epochs*—each (classical) construction query ends the current epoch and begins the next one
- The adversary is trying to distinguish between the real world and the ideal world, which differ only in the construction oracle
- What we would like to do: define hybrid games where the first i epochs take place in the ideal world and the remaining in the real world
- The problem: previous responses in the ideal world are not consistent with the primitive, and this may be detected in a later epoch while making quantum queries to the primitive

Reprogramming and Resampling

- *Reprogramming* an oracle is to modify it at certain points to output a pre-determined value
- Reprogramming F with a pair (x, y) sets $F(x) = y$ and leaves F unchanged at all other points
- *Resampling* F at a point x discards $F(x)$, freshly samples a value y, and sets $F(x) = y$
- Usually in resampling x is also chosen at random, so it is equivalent to reprogramming F with a random pair (x, y)
- There are results showing that reprogramming or resampling F at a small number of points is difficult to detect for an adversary even with superposition access

How Reprogramming Helps

- Going back to our hybrids, when switching from the ideal world to the real world, we can reprogram the primitive retroactively to be consistent with the ideal oracle responses
- This ensures that the construction oracle switch will not be detected in the future
- The results on reprogramming ensure that the primitive switch is itself is also likely to never be detected
- This result can be repeatedly invoked to bound the distance between the real and the ideal world
- (An additional step involving resampling is also needed to complete the reduction for each hybrid)

Illustration: Key-Alternating Feistel

- Suppose the random permutation (in the ideal world) outputs (x_3, x_4) on query (x_0, x_1)
- We can reprogram $f = (f_1, f_2, f_3)$ to be consistent with this output
- We first sample a random z and reprogram f_1 at $(x_1 \oplus k_1, z)$
- Then we reprogram f_2 at $(x_0 \oplus z \oplus k_2, x_1 \oplus x_3)$
- Finally we reprogram f_3 at $(x_3 \oplus k_3, x_0 \oplus z \oplus x_4)$

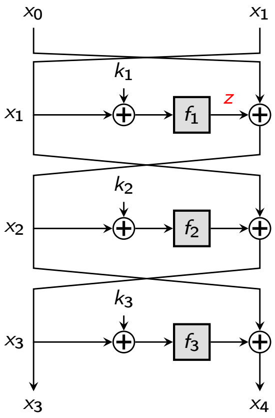

How Things Look at Present

- So far we have reproduced several classical security results for 3-round and 4-round Function-based Key-Alternating Feistel
- We are trying to extend this to Permutation-based KAF (reprogramming a permutation is trickier, as it involves swapping two points)
- Once some basic hurdles are cleared and some creases ironed out, our technique should be applicable to many results from classical provable security
- The Q1 situation looks more optimistic than the Q2 situation
- More advanced aspects like beyond-birthday-bound security proofs still to be explored

Thank You for Listening!

If you're still awake, I am happy to take some (easy) questions.