On round elimination for special-sound multi-round identification and the generality of the hypercube for MPCitH

Andreas Hülsing [1,2] David Joseph [2] Christian Majenz [3]
Anand Kumar Narayanan [2]

[1] Eindhoven University of Technology, Eindhoven, The Nederlands.

[2] SandboxAQ, Palo Alto, CA, USA.

[3] Technical University of Denmark, Copenhagen, Denmark.

Fiat-Shamir in the Quantum-accessible Random Oracle Model (QROM)

A popular recipe for post-quantum signatures is

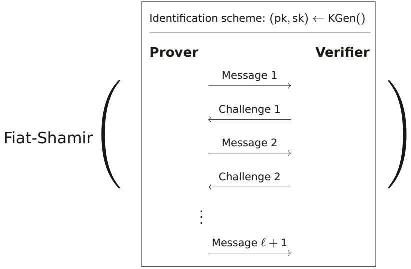

Fiat-Shamir: Replace each verifier challenge with an independent hash function evaluated at the messages so far.

Difficult to analyse, instead
ROM: Recast hash functions with query/oracle access to random functions (RO).

Fiat-Shamir in the Quantum-accessible Random Oracle Model (QROM)

A popular recipe for post-quantum signatures is

Fiat-Shamir

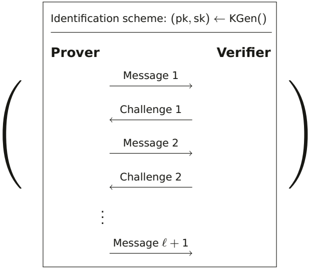

Fiat-Shamir: Replace each verifier challenge with an independent hash function evaluated at the messages so far.

Difficult to analyse, instead
ROM: Recast hash functions with query/oracle access to random functions (RO).

QROM (Boneh, Dagdelen, Fischlin, Lehmann, Schaffner, and Zhandry [Asiacrypt10]):

- the quantum adversary has quantum query access to the random functions.

Fiat-Shamir in the Quantum-accessible Random Oracle Model (QROM)

A popular recipe for post-quantum signatures is

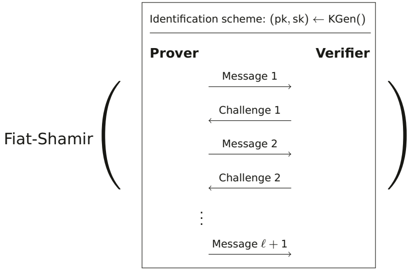

Fiat-Shamir: Replace each verifier challenge with an independent hash function evaluated at the messages so far.

Difficult to analyse, instead
ROM: Recast hash functions with query/oracle access to random functions (RO).

QROM (Boneh, Dagdelen, Fischlin, Lehmann, Schaffner, and Zhandry [Asiacrypt10]):

- the quantum adversary has quantum query access to the random functions.

eQROM: Adaptive access to an additional extraction interface of the RO.

Context

- Don, Fehr, Majenz, and Schaffner [Crypto22] proved optimal QROM security for commit-open 3-rounds.
- For more rounds (such as 5-round MPCitHs), only loose bounds were known.

Context

- Don, Fehr, Majenz, and Schaffner [Crypto22] proved optimal QROM security for commit-open 3-rounds.
- For more rounds (such as 5-round MPCitHs), only loose bounds were known.
- Aguilar-Melchor, Hülsing, Joseph, Majenz, Ronen, and Yue [Asiacrypt23] proved tight QROM security for SDitH (a 5-round MPCitH scheme) signatures, by reducing 5-rounds to 3-rounds and invoking DFMS-Crypto22.

Context

- Don, Fehr, Majenz, and Schaffner [Crypto22] proved optimal QROM security for commit-open 3-rounds.
- For more rounds (such as 5-round MPCitHs), only loose bounds were known.
- Aguilar-Melchor, Hülsing, Joseph, Majenz, Ronen, and Yue [Asiacrypt23] proved tight QROM security for SDitH (a 5-round MPCitH scheme) signatures, by reducing 5-rounds to 3-rounds and invoking DFMS-Crypto22.

Results

- We generalise round elimination to 5 (or more) rounds.

Context

- Don, Fehr, Majenz, and Schaffner [Crypto22] proved optimal QROM security for commit-open 3-rounds.
- For more rounds (such as 5-round MPCitHs), only loose bounds were known.
- Aguilar-Melchor, Hülsing, Joseph, Majenz, Ronen, and Yue [Asiacrypt23] proved tight QROM security for SDitH (a 5-round MPCitH scheme) signatures, by reducing 5-rounds to 3-rounds and invoking DFMS-Crypto22.

Results

- We generalise round elimination to 5 (or more) rounds.
- Hypercube optimization for most MPCitH based signatures.

Round elimination: Soundness preservation.

What do we bound?

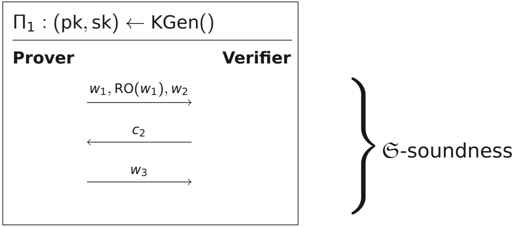

We bound a computational version of Don, Fehr, Majenz, and Schaffner [Eurocrypt22]'s $\mathfrak{S}$-soundness, a generalisation of query-bounded special soundness to arbitrary challenge patterns.

Round elimination: Soundness preservation.

When are the bounds meaningful/tight etc.?

- The bounds are tight for most 5-round commit-open schemes, including most MPCitH based on-ramp NIST signatures.

  A sufficient condition: Even if the first message is adversarially chosen and the first challenge is uniformly sampled afterwards, the remaining protocol has some form of special soundness with overwhelming probability.

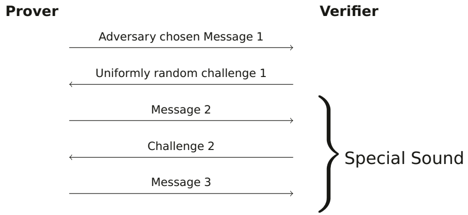

Round elimination: Soundness preservation.

When are the bounds meaningful/tight etc.?

- The bounds are tight for most 5-round commit-open schemes, including most MPCitH based on-ramp NIST signatures.

  A sufficient condition: Even if the first message is adversarially chosen and the first challenge is uniformly sampled afterwards, the remaining protocol has some form of special soundness with overwhelming probability.

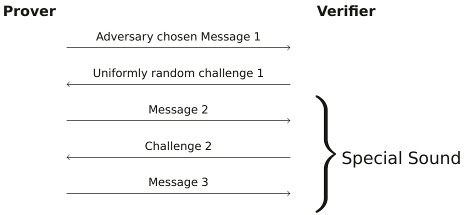

- The bounds can be trivial even for some 5-round cases, such as MQ-DSS, when the soundness does not “factor" through the two verifier challenge rounds.

Round elimination: Soundness preservation proof sketch for 5-rounds

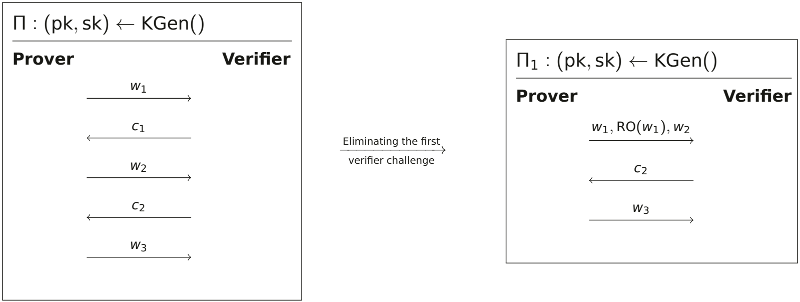

$$t \leftarrow \langle \mathrm{Adver}, \mathrm{Verifier} \rangle \longleftrightarrow t' \leftarrow \langle \mathrm{Adver}_q^{\mathrm{RO}}, \mathrm{Verifier} \rangle$$

Round elimination: Soundness preservation proof sketch for 5-rounds

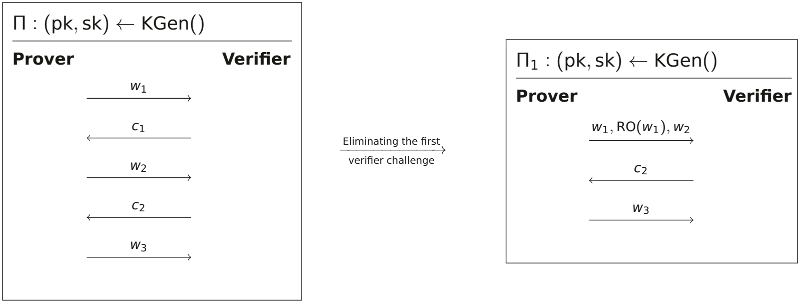

$$t \leftarrow \langle \mathsf{Adver}, \mathsf{Verifier} \rangle \longleftrightarrow t' \leftarrow \langle \mathsf{Adver}_q^{\mathsf{RO}}, \mathsf{Verifier} \rangle$$

Consider the maximum over the first message of the expectation

$$\mu := \max_{w_1} \mathbb{E}_{c_1} \left[\Pr\left[\text{cheating } \Pi \text{ conditioned on } (w_1, c_1) \right] \right]$$

Round elimination: Soundness preservation proof sketch for 5-rounds

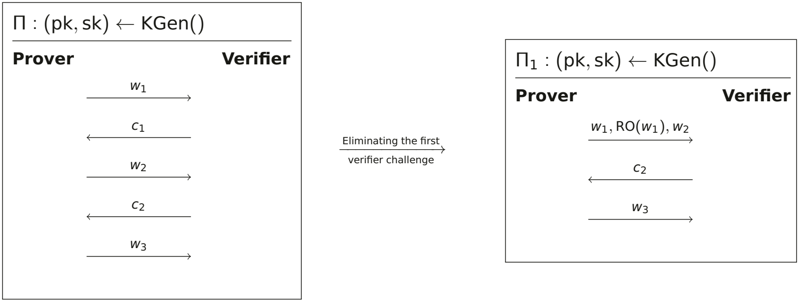

$$t \leftarrow \langle \text{Adver}, \text{Verifier} \rangle \longleftrightarrow t' \leftarrow \langle \text{Adver}_q^{\text{RO}}, \text{Verifier} \rangle$$

Consider the maximum over the first message of the expectation

$$\mu := \max_{w_1} \mathbb{E}_{c_1} \Big[\Pr\big[\text{cheating } \Pi \text{ conditioned on } (w_1, c_1) \big] \Big]$$

and the standard deviation

$$\sigma := \max_{w_1} \sqrt{\mathrm{Var}_{c_1} \Big[\Pr\big[\text{cheating } \Pi \text{ conditioned on } (w_1, c_1) \big] \Big]}$$

over the first challenge of cheating the remaining protocol Π.

Round elimination: Soundness preservation proof sketch

Can the adversary $\mathrm{Adver}_q^{\mathrm{RO}}$ search for a w_1 that enables cheating on the remaining round eliminated protocol Π_1 with high probability?

Round elimination: Soundness preservation proof sketch

Can the adversary $\mathsf{Adver}_q^{\mathsf{RO}}$ search for a w_1 that enables cheating on the remaining round eliminated protocol Π_1 with high probability?

Main Technical Theorem:

$$\mathbb{E}_{(w_1,\mathsf{RO}(w_1),w_2,c_2,w_3))\leftarrow\mathsf{Adver}_q^{\mathsf{RO}}}\left[\Pr\left[\text{cheating }\Pi_1\text{ conditioned on }(w_1,\mathsf{RO}(w_1))\right]\right]$$

$$\leq \mu + 3\sqrt{304}q\sigma + 608q^2\sigma^2\mu\log\left(\frac{1}{\sqrt{304}q\sigma}\right).$$

Round elimination: Soundness preservation proof sketch

Can the adversary $\mathsf{Adver}_q^{\mathsf{RO}}$ search for a w_1 that enables cheating on the remaining round eliminated protocol Π_1 with high probability?

Main Technical Theorem:

$$\mathbb{E}_{(w_1,\mathsf{RO}(w_1),w_2,c_2,w_3))\leftarrow\mathsf{Adver}_q^{\mathsf{RO}}}\left[\Pr\left[\text{cheating }\Pi_1\text{ conditioned on }(w_1,\mathsf{RO}(w_1))\right]\right] \leq \mu + 3\sqrt{304}q\sigma + 608q^2\sigma^2\mu\log\left(\frac{1}{\sqrt{304}q\sigma}\right).$$

Proof idea: Hardness of optimization/search in the QROM.

- The probability of cheating conditioned on w_1, is a function of $(w_1, \mathsf{RO}(w_1))$.

Round elimination: Soundness preservation proof sketch

Can the adversary $\mathsf{Adver}_q^{\mathsf{RO}}$ search for a w_1 that enables cheating on the remaining round eliminated protocol Π_1 with high probability?

Main Technical Theorem:

$$\mathbb{E}_{(w_1,\mathsf{RO}(w_1),w_2,c_2,w_3))\leftarrow\mathsf{Adver}_q^{\mathsf{RO}}}\left[\Pr\left[\text{cheating } \Pi_1 \text{ conditioned on } (w_1,\mathsf{RO}(w_1))\right]\right]$$
$$\leq \mu + 3\sqrt{304}q\sigma + 608q^2\sigma^2\mu\log\left(\frac{1}{\sqrt{304}q\sigma}\right).$$

Proof idea: Hardness of optimization/search in the QROM.

- The probability of cheating conditioned on w_1, is a function of $(w_1, \mathsf{RO}(w_1))$.
- Don, Fehr, Majenz, and Schaffner [Eurocrypt22] and Hövelmanns, Hülsing, and Majenz [Asiacrypt22] tell us how hard it is to search for an argument w_1 that finds large values of a function of $\mathsf{RO}(w_1)$.

Round elimination: Soundness preservation proof sketch

Can the adversary $\text{Adver}_q^{\text{RO}}$ search for a w_1 that enables cheating on the remaining round eliminated protocol Π_1 with high probability?

Main Technical Theorem:

$$\mathbb{E}_{(w_1, \text{RO}(w_1), w_2, c_2, w_3)) \leftarrow \text{Adver}_q^{\text{RO}}} \left[\Pr\left[\text{cheating } \Pi_1 \text{ conditioned on } (w_1, \text{RO}(w_1)) \right] \right] \\ \leq \mu + 3\sqrt{304} q\sigma + 608 q^2 \sigma^2 \mu \log\left(\frac{1}{\sqrt{304} q\sigma} \right).$$

Proof idea: Hardness of optimization/search in the QROM.

- The probability of cheating conditioned on w_1, is a function of $(w_1, \text{RO}(w_1))$.
- Don, Fehr, Majenz, and Schaffner [Eurocrypt22] and Hövelmanns, Hülsing, and Majenz [Asiacrypt22] tell us how hard it is to search for an argument w_1 that finds large values of a function of $\text{RO}(w_1)$.

Corollary: Soundness preservation for 5-round to 3-round

The “additional" advantage of a q-query polynomial time quantum adversary for the 3-round round elimination of a d-special sound 5-round parallel repeated scheme is at most

$$3\sqrt{304} q\sigma + 608 q^2 \sigma^2 \mu \log\left(\frac{1}{\sqrt{304} q\sigma} \right).$$

Round elimination: Zero-Knowledge Preservation.

We prove that honest-verifier zero-knowledge is preserved by round elimination.

Key tool:
The adaptive reprogramming lemma of Grilo, Hövelmanns, Hülsing, and Majenz [Asiacrypt21].

MPCitH

- MPCitH: Ishai, Kushilevitz, Ostrovsky, and Sahai [STOC07] introduced new zero-knowledge proofs of NP statements, using Multi-Party Computation.
- In our context, a public key pk defines a function f_{pk} to which the Prover claims knowledge of a (secret key/witness) satisfying assignment $x = sk$, such that

  $$f_{pk}(x) = 0.$$

  A special sound and HVZK 3-round scheme to verify this claim.

MPCitH

- MPCitH: Ishai, Kushilevitz, Ostrovsky, and Sahai [STOC07] introduced new zero-knowledge proofs of NP statements, using Multi-Party Computation.
- In our context, a public key *pk* defines a function f_{pk} to which the Prover claims knowledge of a (secret key/witness) satisfying assignment $x = sk$, such that

$$f_{pk}(x) = 0.$$

A special sound and HVZK 3-round scheme to verify this claim.

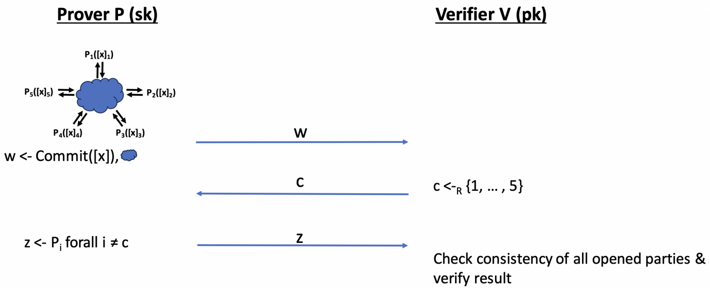

Hypercube aggregation for MPCitH

- The hypercube technique is an optimization introduced by Aguilar-Melchor, Gama, Howe, Hülsing, Joseph, and Yue [Eurocrypt23] to accelerate the signature and verification procedure of the SDitH signature scheme.

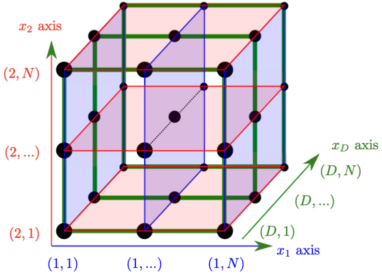

The $D \times N$ main party slices

Picture courtesy of AGHHJY-Eurocrypt23.

Hypercube aggregation for MPCitH

- The hypercube technique is an optimization introduced by Aguilar-Melchor, Gama, Howe, Hülsing, Joseph, and Yue [Eurocrypt23] to accelerate the signature and verification procedure of the SDitH signature scheme.

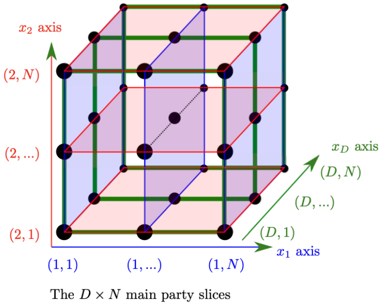

Picture courtesy of AGHHJY-Eurocrypt23.

- The hypercube has since been adopted by several of the MPCitH based NIST on-ramp signatures, carefully tailoring the optimization to their context.
- Hypercube improves the signing/verification times by an order of 4 to 12.

Hypercube for most 5-round MPCitH based signatures

We present an abstraction of 3-round MPCitHs with MPCs that are

- N-1 private in the semi-honest model,
- symmetric in the parties, meaning it looks the same if we permute the parties,
- and additive in all the inputs.

Hypercube for most 5-round MPCitH based signatures

We present an abstraction of 3-round MPCitHs with MPCs that are

- N-1 private in the semi-honest model,
- symmetric in the parties, meaning it looks the same if we permute the parties,
- and additive in all the inputs.

We transform any such 3-round MPCitH into one with the Hypercube optimisation such that soundness and HVZK are preserved.

- To achieve N^{-D} soundness error takes N^D parties. Using the hypercube technique, communicating the computation of ND parties suffice.

Hypercube for most 5-round MPCitH based signatures

We present an abstraction of 3-round MPCitHs with MPCs that are

- N-1 private in the semi-honest model,
- symmetric in the parties, meaning it looks the same if we permute the parties,
- and additive in all the inputs.

We transform any such 3-round MPCitH into one with the Hypercube optimisation such that soundness and HVZK are preserved.

- To achieve N^{-D} soundness error takes N^D parties. Using the hypercube technique, communicating the computation of ND parties suffice.

Efficient 5-round MPCitH (atleast 9/40 of the NIST on-ramp sigs)

- Most efficient MPCitH schemes tend to be 5-round, following the motif of Lindell & Nof [CCS17], and Baum & Nof [PKC20].
- Well suited for predicates that are mostly linear, with a non-linearity constraint for hardness. For example, syndrome decoding, rank-metric decoding, etc..
- First message commits to a linear MPC. Second message commits to a multiplicative MPC.

Hypercube for most 5-round MPCitH based signatures

We present an abstraction of 3-round MPCitHs with MPCs that are

- N-1 private in the semi-honest model,
- symmetric in the parties, meaning it looks the same if we permute the parties,
- and additive in all the inputs.

We transform any such 3-round MPCitH into one with the Hypercube optimisation such that soundness and HVZK are preserved.

- To achieve N^{-D} soundness error takes N^D parties. Using the hypercube technique, communicating the computation of ND parties suffice.

Efficient 5-round MPCitH (atleast 9/40 of the NIST on-ramp sigs)

- Most efficient MPCitH schemes tend to be 5-round, following the motif of Lindell & Nof [CCS17], and Baum & Nof [PKC20].
- Well suited for predicates that are mostly linear, with a non-linearity constraint for hardness. For example, syndrome decoding, rank-metric decoding, etc..
- First message commits to a linear MPC. Second message commits to a multiplicative MPC.

We first apply our round elimination to convert these 5-round schemes to 3-round schemes conforming to our abstraction, then apply the hypercube transform.