SDitH in the QROM

Carlos Aguilar-Melchor, Andreas Hülsing, David Joseph, Christian Majenz, Eyal Ronen, and Dongze Yue

Mathematics and Computer Science

TU/e EINDHOVEN UNIVERSITY OF TECHNOLOGY

Joint work with

Christian Majenz

Carlos Aguilar Melchor

David Joseph

Dongzhe Yue

Eyal Ronen

DTU

SANDBOXAQ

TEL AVIV UNIVERSITY אוניברסיטת תל אביב

TU/e

Syndrome Decoding in the Head (FJR22)

- Code-based signature scheme
- Using MPC in the Head (MPCitH)

Source:
Thibauld Feneuil, Antoine Joux, and Matthieu Rivain.
Syndrome Decoding in the Head: Shorter Signatures from Zero-Knowledge Proofs.
Crypto'22

Scheme Name	Year	|sgn|	|pk|	t_{sgn}	t_{verif}	Assumption
Wave	2019	2.07 K	3.2 M	300	-	SD over $\mathbb{F}_3$ (large weight) $(U, U+V)$-codes indisting.
Durandal - I	2018	3.97 K	14.9 K	4	5	Rank SD over $\mathbb{F}_{2^m}$
Durandal - II	2018	4.90 K	18.2 K	5	6	Rank SD over $\mathbb{F}_{2^m}$
LESS-FM - I	2020	15.2 K	9.77 K	-	-	Linear Code Equivalence
LESS-FM - II	2020	5.25 K	206 K	-	-	Perm. Code Equivalence
LESS-FM - III	2020	10.39 K	11.57 K	-	-	Perm. Code Equivalence
[GPS22]-256	2021	24.0 K	0.11 K	-	-	SD over $\mathbb{F}_{256}$
[GPS22]-1024	2021	19.8 K	0.12 K	-	-	SD over $\mathbb{F}_{1024}$
[FJR21] (fast)	2021	22.6 K	0.09 K	13	12	SD over $\mathbb{F}_2$
[FJR21] (short)	2021	16.0 K	0.09 K	62	57	SD over $\mathbb{F}_2$
[BGKM22] - Sig1	2022	23.7 K	0.1 K	-	-	SD over $\mathbb{F}_2$
[BGKM22] - Sig2	2022	20.6 K	0.2 K	-	-	(QC)SD over $\mathbb{F}_2$
Our scheme - Var1f	2022	15.6 K	0.09 K	-	-	SD over $\mathbb{F}_2$
Our scheme - Var1s	2022	10.9 K	0.09 K	-	-	SD over $\mathbb{F}_2$
Our scheme - Var2f	2022	17.0 K	0.09 K	13	13	SD over $\mathbb{F}_2$
Our scheme - Var2s	2022	11.8 K	0.09 K	64	61	SD over $\mathbb{F}_2$
Our scheme - Var3f	2022	11.5 K	0.14 K	6	6	SD over $\mathbb{F}_{256}$
Our scheme - Var3s	2022	8.26 K	0.14 K	30	27	SD over $\mathbb{F}_{256}$

Table 6. Comparison of our scheme with signatures from the literature (128-bit security). The sizes are in bytes and the timings are in milliseconds. Reported timings are from the original publications: Wave has been benchmarked on a 3.5 Ghz Intel Xeon E3-1240 v5, Durandal on a 2.8 Ghz Intel Core i5-7440HQ, while [FJR21] and our scheme on a 3.8 GHz Intel Core i7.

TU/e

Identification schemes (3-round, public coin)

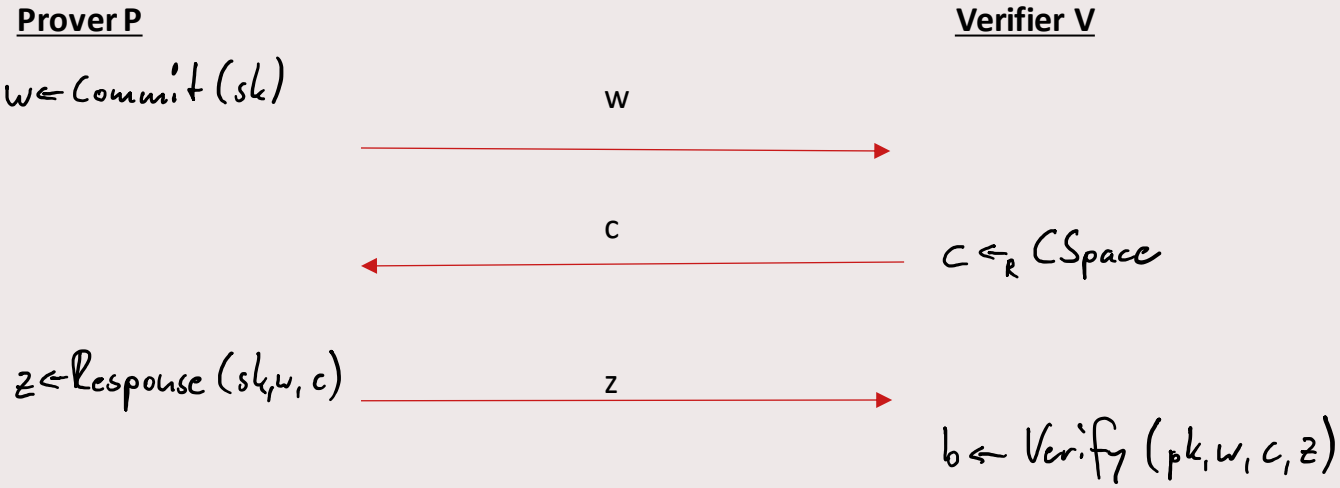

Identification schemes (5-round, public coin)

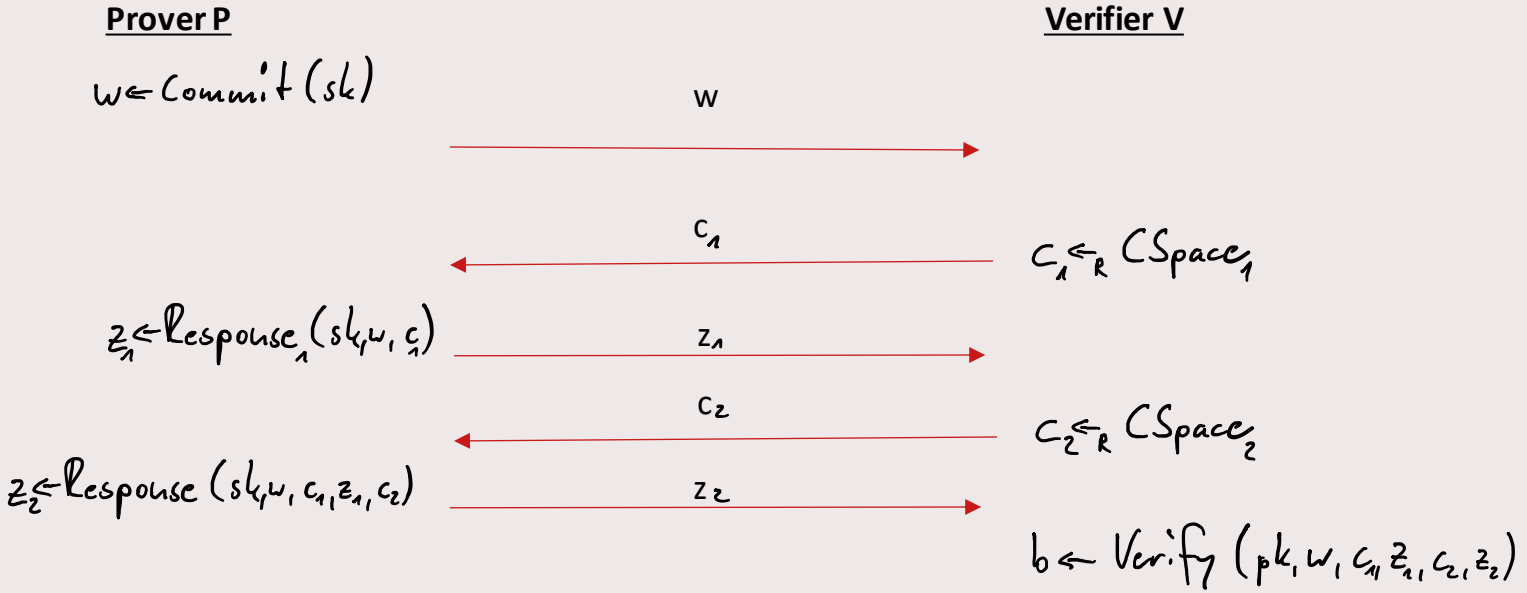

Security Properties

(special) soundness: There exists an efficient extractor E that given two transcripts with same w but different c, extracts sk.

Honest verifier zero-knowledge (HVZK): There exists an efficient simulator S that, given only the public key, outputs transcripts which are indistinguishable from transcripts of honest protocol runs

Identification schemes (3-round, public coin)

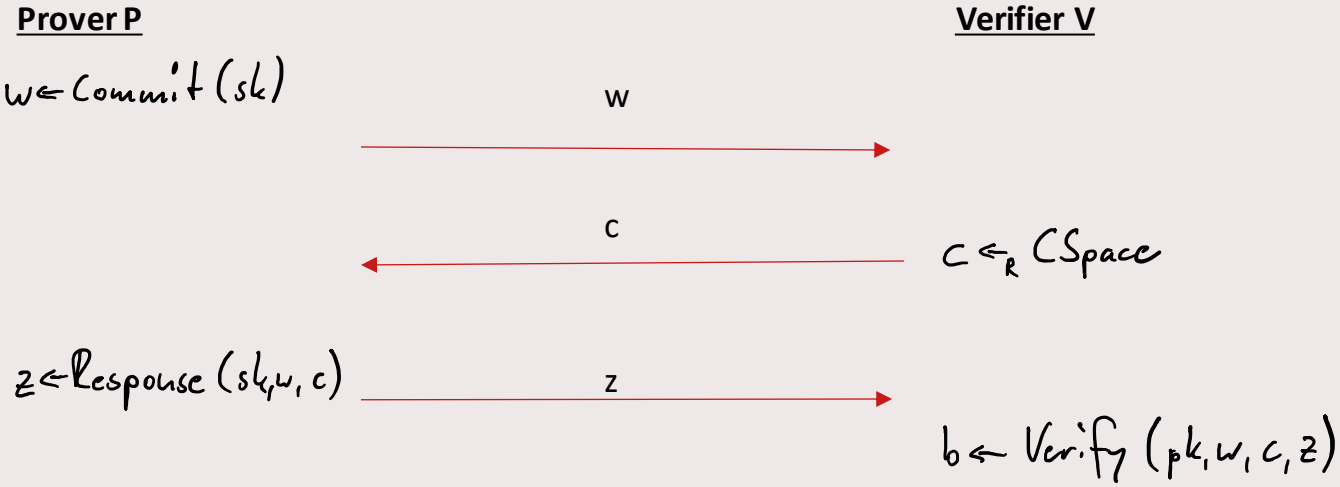

MPCitH for PQ-identification

(Y. Ishai, E. Kushilevitz, R. Ostrovsky, and A. Sahai. “Zero-knowledge from secure multiparty computation”. STOC'07)

Given OWF F: X -> Y

Create identification scheme IDS that proves knowledge of x such that

$$F(x) = y$$

for given y in zero-knowledge.

sk = x, pk = y

Used for (at least) 9 of 40 new NIST signature proposals.

MPCitH

KeyGen:

Sample x, set y = F(x)

MPCitH

Commit:

Secret share x: $x = \sum_1 x_i + x_N$ with $x_i \leftarrow_R \mathbb{F} \wedge x_N \leftarrow x - \sum x_i$

Sample random tapes: $r_i \leftarrow_R R$

Commit to shares & rand: $com_i = COM(x_i, r_i)$

Run MPC protocol π such that $\pi(x_i, r_i) = \alpha_i \wedge \sum_1 \alpha_i = 0$ iff $F(\sum x_i) = y$

Output $(com_i, \alpha_i)_{i=1}^{N}$

TU/e

MPCitH

Response:

Open all commitments except com_c and output openings.

$\Rightarrow$ Return $(x_i, r_i)_{i \neq c}$

MPCitH

Verify:

Check $com_i = COM(x_i, r_i) \wedge \alpha_i = \pi(x_i, r_i) \quad \forall\, i \neq c$

Verify $\sum_{i=1}^{N} \alpha_i = 0$

Return true if none of the above failed.

MPCitH Security

HVZK: Secrecy of inputs in MPC

Soundness: Cut & Choose - catch a cheating prover with probability 1- (1 / #parties)

Special soundness: Two valid openings for same commitments but different challenge reveal all secret shares (and as it opens all parties, none of them can have cheated without getting caught)

SDitH (FJR'22)

Apply MPCitH to Syndrome Decoding problem

Definition 4 (Coset Weights Syndrome Decoding problem). *Sample a uniformly random parity check matrix* $\mathbf{H} \in \mathbb{F}_{SD}^{(m-k)\times m}$, *and binary vector* $\mathbf{x} \in \mathbb{F}_{SD}^{m}$ *with* $wt(\mathbf{x}) = \omega$. *Let syndrome* $\mathbf{y} = \mathbf{Hx}$. *Then given only* $\mathbf{H}, \mathbf{y}$, *it is difficult to find* $\mathbf{x}' \in \mathbb{F}_{SD}^{m}$ *such that* $\mathbf{Hx}' = \mathbf{y}$ *with* $wt(\mathbf{x}') \leq \omega$.

SDitH (FJR'22)

Apply MPCitH to Syndrome Decoding problem

Definition 4 (Coset Weights Syndrome Decoding problem). *Sample a uniformly random parity check matrix* $\mathbf{H} \in \mathbb{F}_{SD}^{(m-k)\times m}$, *and binary vector* $\mathbf{x} \in \mathbb{F}_{SD}^{m}$ *with* $wt(\mathbf{x}) = \omega$. *Let syndrome* $\mathbf{y} = \mathbf{Hx}$. *Then given only* $\mathbf{H}, \mathbf{y}$, *it is difficult to find* $\mathbf{x}' \in \mathbb{F}_{SD}^{m}$ *such that* $\mathbf{Hx}' = \mathbf{y}$ *with* $wt(\mathbf{x}') \leq \omega$.

Advantage: Linear function.

TU/e

SDitH (FJR'22)

Apply MPCitH to Syndrome Decoding problem

Definition 4 (Coset Weights Syndrome Decoding problem). *Sample a uniformly random parity check matrix* $\mathbf{H} \in \mathbb{F}_{SD}^{(m-k)\times m}$, *and binary vector* $\mathbf{x} \in \mathbb{F}_{SD}^{m}$ *with* $wt(\mathbf{x}) = \omega$. *Let syndrome* $\mathbf{y} = \mathbf{Hx}$. *Then given only* $\mathbf{H}, \mathbf{y}$, *it is difficult to find* $\mathbf{x}' \in \mathbb{F}_{SD}^{m}$ *such that* $\mathbf{Hx}' = \mathbf{y}$ *with* $wt(\mathbf{x}') \leq \omega$.

Advantage: Linear function.
Disadvantage: Weight check.

SDitH – Weight check

- Uses "Polynomial zero-test"
- Uses polys Q, P, and public F as well as polynomial S derived from x such that

$$T = SQ - PF = 0 \text{ if } wt(x) \leq \omega$$

- Checking this is done by evaluating T at random points.
- Needs multiplication which needs one more round of interaction!

SDitH Identification scheme (5-round, public coin)

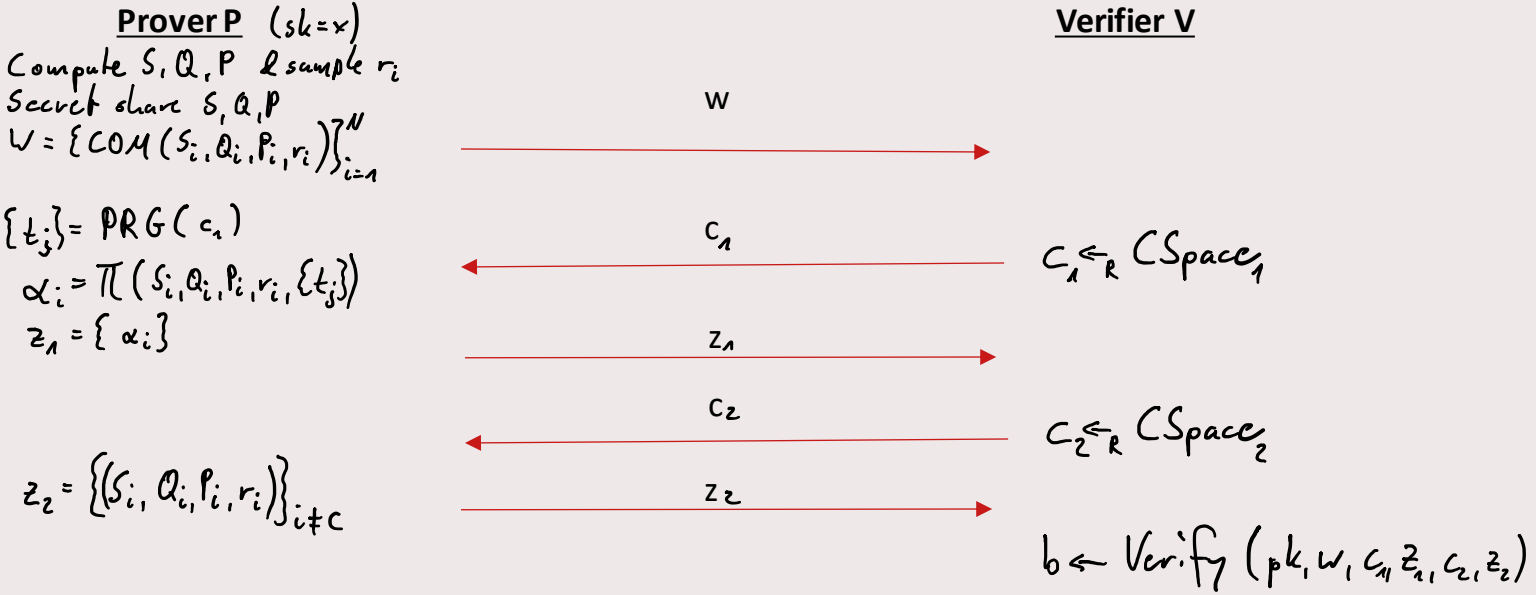

Tweaks

Use TreePRG for random x_i and r_i. (Log size opening)
Hypercube:

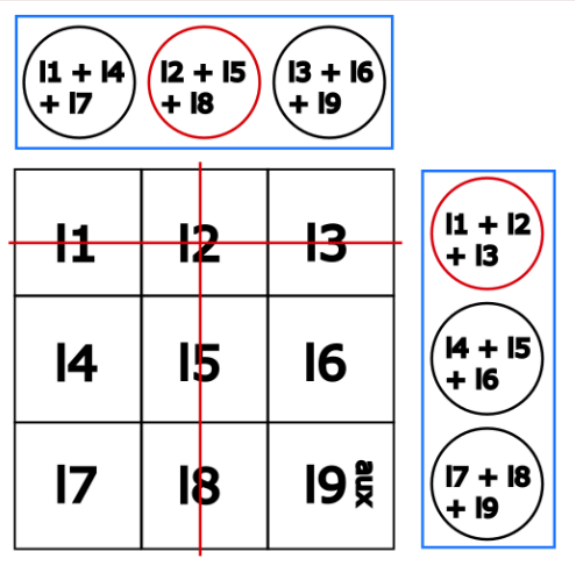

Carlos Aguilar-Melchor, Nicolas Gama, James Howe, Andreas Hülsing, David Joseph, and Dongze Yue
The Return of the SDitH.
EUROCRYPT, 2023

Signature Scheme

Fiat-Shamir transform

- S.KeyGen = IDS.KeyGen
- S.Sign(sk,m) = P.COMMIT + $P.RESPONSE_1$ + $P.RESPONSE_2$
  with $c_1 = H(w[, m])$, $c_2 = H(c_1,z_1,m)$
- S.Verify = V.verify with $c_1 = H(w[, m])$, $c_2 = H(c_1,z_1,m)$

TU/e

How to prove security?

- IDS: Done in [FJR'22]
- Signature against classical adversaries (ROM): Done in [FJR'22]
- Signature against quantum adversaries (QROM): ?

How to prove security?

- IDS: Done in [FJR'22]
- Signature against classical adversaries (ROM): Done in [FJR'22]
- Signature against quantum adversaries (QROM): ?

- Generic results on (5-round) FS have a horrible tightness loss
- Amazing (pretty) tight result for commit & open IDS

  J. Don, S. Fehr, C. Majenz, and C. Schaffner.
  Efficient NIZKs and Signatures from Commit-and-Open Protocols in the QROM.
  Crypto'22

  But: only for 3-round IDS

Wait, FJR'22 showed 2-special soundness.

We showed something about 2-special sound 5-round IDS in the MQDSS paper...

Observation

We can apply a "cheap FS transform" to the first challenge.

- Replace challenge by hash of commitment
- Security argument based on hard search problem

- Cheap? No extraction needed. Just information theoretic arguments (as everything is in the (Q)ROM).

Proof strategy

- Reduce to 3-rounds
- Prove HVZK in QROM -> standard
- Prove Soundness in QROM -> see below
- Apply known results:
  - A. B. Grilo, K. Hövelmanns, A. Hülsing, and Christian Majenz. *Tight adaptive reprogramming in the QROM.* Asiacrypt'21
    UF-NMA + HVZK ==QROM==> UF-CMA
  - J. Don, S. Fehr, C. Majenz, and C. Schaffner.
    Efficient NIZKs and Signatures from Commit-and-Open Protocols in the QROM.
    Crypto'22
    Sp. Sound. ==QROM==> UF-NMA

Computational version of special soundness

Definition 3 ((Query-bounded) distance-d special soundness for IDS with splittable challenge). *We define the advantage of a possibly quantum adversary* A *against the query bounded special soundness of a composed* IDS *with respect to extractor* Ext *in the (quantum-accessible) random oracle model as follows*

$$\begin{aligned}\mathrm{Adv}^{d-\mathsf{spS}}_{\mathsf{IDS},\mathsf{Ext}}(\mathsf{A}) := \Pr[&(\mathsf{sk},\mathsf{pk}) \leftarrow \mathsf{Keygen}(); ((\mathsf{w}_1,\mathsf{c}_1,\mathsf{z}_1),(\mathsf{w}_2,\mathsf{c}_2,\mathsf{z}_2)) \leftarrow \mathsf{A}^{\mathsf{RO}}(\mathsf{pk});\\ &\mathsf{sk}' \leftarrow \mathsf{Ext}^{\mathsf{RO}}((\mathsf{w}_1,\mathsf{c}_1,\mathsf{z}_1),(\mathsf{w}_2,\mathsf{c}_2,\mathsf{z}_2)) : \mathsf{Vrf}(\mathsf{pk},\mathsf{w}_i,\mathsf{c}_i,\mathsf{z}_i) = 1\\ &, i \in \{1,2\} \wedge (\mathsf{w}_1 = \mathsf{w}_2) \wedge d = \mathsf{Dist}(\mathsf{c}_1,\mathsf{c}_2) \wedge (\mathsf{sk}',\mathsf{pk}) \notin \mathsf{Keygen}()],\end{aligned}$$

Proven bound

Theorem 4. *Our identification scheme* Π *has query-bounded distance-d special soundness. More precisely, let* $\mathsf{A}^{\mathsf{Com},G}$ *be a distance-d special soundness adversary making at most* q_{Com} *and* q_G *queries to its oracles* Com *and* G*, respectively, and set* $q = q_{\mathsf{Com}} + q_G$ *and* $\tilde{q} = q + \tau \cdot N^D + 1$*. Then the bounds*

$$\mathrm{Adv}^{d-\mathsf{spS}}_{\mathsf{IDS},\mathsf{Ext}}(\mathsf{A}) \leq \begin{cases} (\tau N^D + 1)\frac{\tilde{q}^2}{2^c} + \tilde{q}\binom{\tau}{d}p^{t\cdot d} & \text{in the ROM} \\ (10\tau N^D + 47)\frac{\tilde{q}^3}{2^c} + 10\tilde{q}^2\binom{\tau}{d}p^{t\cdot d} & \text{in the QROM+} \end{cases}$$

hold, where c *is the output length of* Com.

TU/e

Proven bound

Theorem 4. *Our identification scheme* Π *has query-bounded distance-d special soundness. More precisely, let* $\mathsf{A}^{\mathsf{Com},G}$ *be a distance-d special soundness adversary making at most* q_{Com} *and* q_G *queries to its oracles* Com *and* G*, respectively, and set* $q = q_{\mathsf{Com}} + q_G$ *and* $\tilde{q} = q + \tau \cdot N^D + 1$*. Then the bounds*

$$\mathrm{Adv}^{d-\mathsf{spS}}_{\mathsf{IDS},\mathsf{Ext}}(\mathsf{A}) \leq \begin{cases} (\tau N^D + 1)\frac{\tilde{q}^2}{2^c} + \tilde{q}\binom{\tau}{d}p^{t\cdot d} & \text{in the ROM} \\ (10\tau N^D + 47)\frac{\tilde{q}^3}{2^c} + 10\tilde{q}^2\binom{\tau}{d}p^{t\cdot d} & \text{in the QROM+} \end{cases}$$

hold, where c *is the output length of* Com.

QROM+ - Phase 1

Compressed oracle [Zhandry'18]

QROM+ - Phase 2

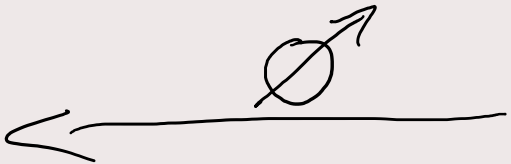

Compressed oracle [Zhandry'18]

Why do we need a QROM+?

- We build algorithm R for oracle search problem
- R runs A against soundness of IDS
- A solves search problems (reflected in queries)
- A's QROM queries cannot be seen by R

TU/e

Why is this unproblematic?

Search problems are not easier in QROM+!

- R as a whole (including A) has the knowledge
- It's as if R is oblivious
- Measurement does not give any new information

UF-NMA

Syndrome decoding

FS-transform

Corollary 2. *Let* A *be a* UF-NMA*-adversary against* $\mathsf{FS}[\Pi, \mathsf{RO}]$ *that makes* $q_{\mathsf{RO}} \geq \tau \cdot N^D + 1$, q_{Com} *and* q_G *quantum queries to* RO, Com *and* G *respectively. Then for all* $d = 0, 1, \ldots, \tau$ *we get*

$$\mathrm{Adv}^{\text{UF-NMA}}_{\mathsf{FS}[\mathsf{IDS},\mathsf{RO}]}(\mathsf{A}) \leq \epsilon_{\mathrm{SD}} + (32\tau N^D + 107)\frac{q^3}{2^c} + 10 \cdot q^2 \binom{\tau}{d} p^{t \cdot d} + 20q^2 \frac{1}{N^{D \cdot (\tau - d)}}.$$

Here, ϵ_{SD} *is the maximal success probability that an adversary with runtime* $\mathsf{TIME}(\mathsf{A}) + \mathsf{TIME}(\mathrm{CompOr}(q)) + \mathsf{TIME}(\mathsf{Ext}_d)$*, where* $\mathsf{TIME}(\mathrm{CompOr}(q))$ *is the runtime of a compressed oracle simulation for* q *queries, can solve syndrome decoding. Also* $q = q_{\mathsf{Com}} + q_{\mathsf{RO}} + q_G$ *is the total number of random oracle queries of* A*,* c *is the output length of* Com*, and the atomic polynomial zero test false-positive probability* p *is defined and bounded in Equation* (11) *and Equation* (12)*.*

special soundness + FS-transform

TU/e

UF-CMA

Corollary 3. *Let* A *be a* UF-CMA*-adversary against* FS[Π, RO] *that makes* $q_{\mathsf{RO}} \geq \tau \cdot N^D + 1$, q_{PRG}, q_{Com} *and* q_G *quantum queries to* RO, PRG, Com *and* G *respectively, and* q_{S} *(classical) signing queries. Then for all* $d = 0, 1, \ldots, \tau$,

$$\mathrm{Adv}^{\text{UF-CMA}}_{\mathsf{FS[IDS,RO]}}(\mathsf{A}) \leq \epsilon_{\mathrm{SD}} + (32\tau N^D + 107)q^3 2^{-c} + 10 \cdot q^2 \binom{\tau}{d} p^{t \cdot d} + 20q^2 \frac{1}{N^{D \cdot (\tau - d)}}$$

$$+ q_{\mathsf{S}}\tau \left(16 q_{\mathsf{Com}} 2^{-r/2} + \log(N^D - 1) \frac{(q_{\mathsf{PRG}} + q_{\mathsf{S}}\tau)^2}{2^n} \right) + \frac{3q_{\mathsf{S}}}{2} \sqrt{\frac{q_{\mathsf{RO}} + q_{\mathsf{S}} + 1}{2^n}}, \qquad (14)$$

Here ϵ_{SD} *is the maximal success probability that an adversary that runs in time* TIME(A) + TIME(CompOr(q)) + TIME(Ext_d)*, where* TIME(CompOr(q)) *is the runtime of a compressed oracle simulation for* q *queries, can solve syndrome decoding. Moreover,* $q = q_{\mathsf{Com}} + q_{\mathsf{RO}} + q_G$ *is the total number of random oracle queries of* A*,* c *is the output length of* Com*, and the atomic polynomial zero test false-positive probability* p *is defined in Equation (11) and bounded in Equation (12),* n *is the seed length of* TreePRG*,* r *is the length of commitment randomness.*

HVZK

UF-NMA

Reprogramming

TU/e

UF-CMA

Binding Com.

HVZK

Grover search for G & RO

Corollary 3. *Let* A *be a* UF-CMA*-adversary against* FS[Π, RO] *that makes* $q_{\mathsf{RO}} \geq \tau \cdot N^D + 1$, q_{PRG}, q_{Com} *and* q_G *quantum queries to* RO, PRG, Com *and* G *respectively, and* q_{S} *(classical) signing queries. Then for all* $d = 0, 1, \ldots, \tau$,

$$\mathrm{Adv}^{\text{UF-CMA}}_{\mathsf{FS[IDS,RO]}}(\mathsf{A}) \leq \epsilon_{\mathrm{SD}} + (32\tau N^D + 10\tau) q^3 2^{-c} + 10 \cdot q^2 \binom{\tau}{d} p^{t \cdot d} + 20 q^2 \frac{1}{N^{D \cdot (\tau - d)}}$$

} UF-NMA

$$+ q_{\mathsf{S}}\tau \left(16 q_{\mathsf{Com}} 2^{-r/2} + \log(N^D - 1) \frac{(q_{\mathsf{PRG}} + q_{\mathsf{S}}\tau)^2}{2^n} \right) + \frac{3 q_{\mathsf{S}}}{2} \sqrt{\frac{q_{\mathsf{RO}} + q_{\mathsf{S}} + 1}{2^n}}, \qquad (14)$$

Here ϵ_{SD} *is the maximal success probability that an adversary that runs in time* TIME(A) + TIME(CompOr(q)) + TIME(Ext_d)*, where* TIME(CompOr(q)) *is the runtime of a compressed oracle simulation for* q *queries, can solve syndrome decoding. Moreover,* $q = q_{\mathsf{Com}} + q_{\mathsf{RO}} + q_G$ *is the total number of random oracle queries of* A*,* c *is the output length of* Com*, and the atomic polynomial zero test false-positive probability* p *is defined in Equation (11) and bounded in Equation (12).* n *is the seed length of* TreePRG*,* r *is the length of commitment randomness.*

Multi-target attacks

Hiding Com

PRG

Reprogramming

TU/e

Results

Table 1: Implementation benchmarks of Hypercube-SDitH vs our tweaked scheme for NIST security level I. For the PoW, the parameter $k_{iter} = D$ is used.

Scheme	Aim	Signature Size (bytes)	Parameters				Sign Time (in ms)			Verify Time
			$\lvert\mathbb{F}_{\text{points}}\rvert$	t	D	τ	Offline	Online	Total	(in ms) Total
Hypercube-SDitH [2]	Short	8464	2^{24}	5	8	17	3.83	0.68	4.51	4.16
	Shorter	6760	2^{24}	5	12	12	44.44	0.60	45.04	42.02
Ours Vanilla	Short	8464	2^{24}	5	8	17	4.45	0.049	4.50	4.17
	Shorter	6760	2^{24}	5	12	12	44.98	0.080	45.06	42.02
Ours PoW	Short	7968	2^{24}	5	8	16	4.20	0.14	4.34	4.00
	Shorter	6204	2^{24}	5	12	11	41.06	1.49	42.55	39.75

Conclusion

- Security proof for SDitH and H-SDitH against quantum adversaries
- Bound is tight up to constants if multi-target mitigation is used
- Allows for online-offline signatures with very short online phase
- Techniques may apply to similar schemes
- (eprint) PoW can be used to optimize parameters

https://eprint.iacr.org/2023/756.pdf

Backup

PoW (increase cost of RO query)

Corollary 3. *Let* A *be a* UF-CMA*-adversary against* FS$[\Pi, \mathsf{RO}]$ *that makes* $q_{\mathsf{RO}} \geq \tau \cdot N^D + 1$, q_{PRG}, q_{Com} *and* q_G *quantum queries to* RO, PRG, Com *and* G *respectively, and* q_{S} *(classical) signing queries. Then for all* $d = 0, 1, \ldots, \tau$,

$$\mathrm{Adv}^{\text{UF-CMA}}_{\mathsf{FS}[\mathsf{IDS},\mathsf{RO}]}(\mathsf{A}) \leq \epsilon_{\mathrm{SD}} + (32\tau N^D + 107)q^3 2^{-c} + 10 \cdot q^2 \binom{\tau}{d} p^{t \cdot d} + 20q^2 \frac{1}{N^{D \cdot (\tau - d)}}$$
$$+ q_{\mathsf{S}}\tau \left(16 q_{\mathsf{Com}} 2^{-r/2} + \log(N^D - 1) \frac{(q_{\mathsf{PRG}} + q_{\mathsf{S}}\tau)^2}{2^n} \right) + \frac{3q_{\mathsf{S}}}{2} \sqrt{\frac{q_{\mathsf{RO}} + q_{\mathsf{S}} + 1}{2^n}}, \qquad (14)$$

Here ϵ_{SD} *is the maximal success probability that an adversary that runs in time* TIME(A) + TIME(CompOr(q)) + TIME(Ext_d)*, where* TIME(CompOr(q)) *is the runtime of a compressed oracle simulation for* q *queries, can solve syndrome decoding. Moreover,* $q = q_{\mathsf{Com}} + q_{\mathsf{RO}} + q_G$ *is the total number of random oracle queries of* A*,* c *is the output length of* Com*, and the atomic polynomial zero test false-positive probability* p *is defined in Equation* (11) *and bounded in Equation* (12)*,* n *is the seed length of* TreePRG*,* r *is the length of commitment randomness.*

TU/e