On the (In)Security of the BUFF Transform

Jelle Don Serge Fehr **Yu-Hsuan Huang** Patrick Struck

CWI

Universität Konstanz

NIST Competition

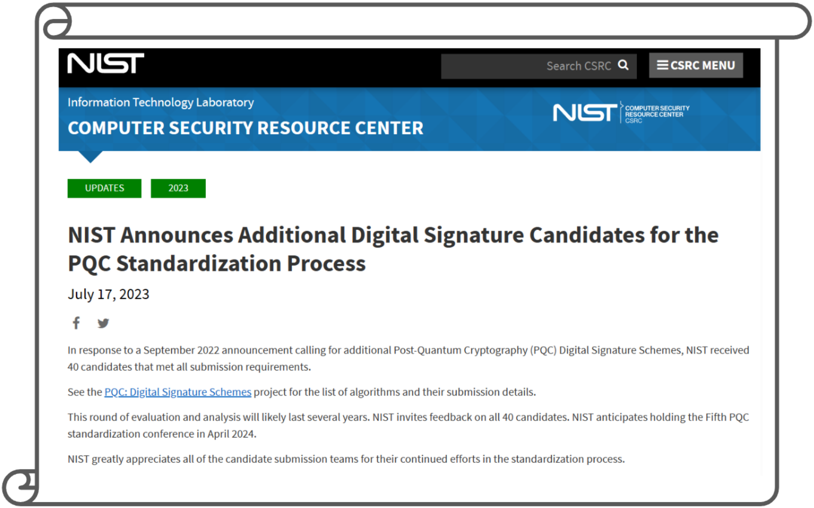

NIST

Search CSRC

CSRC MENU

Information Technology Laboratory

COMPUTER SECURITY RESOURCE CENTER

NIST COMPUTER SECURITY RESOURCE CENTER CSRC

UPDATES 2023

NIST Announces Additional Digital Signature Candidates for the PQC Standardization Process

July 17, 2023

In response to a September 2022 announcement calling for additional Post-Quantum Cryptography (PQC) Digital Signature Schemes, NIST received 40 candidates that met all submission requirements.

See the PQC: Digital Signature Schemes project for the list of algorithms and their submission details.

This round of evaluation and analysis will likely last several years. NIST invites feedback on all 40 candidates. NIST anticipates holding the Fifth PQC standardization conference in April 2024.

NIST greatly appreciates all of the candidate submission teams for their continued efforts in the standardization process.

Figure: NIST Additional PQ Signature Competition

Security Beyond Unforgeability

NIST asked for "additional desirable security properties":

- exclusive ownership (S-CEO, S-DEO, M-S-UEO)
- message-bound signatures (MBS)
- **non-resignability (NR)**

Security Beyond Unforgeability

NIST asked for "additional desirable security properties":

- exclusive ownership (S-CEO, S-DEO, M-S-UEO)
- message-bound signatures (MBS)
- **non-resignability (NR)**

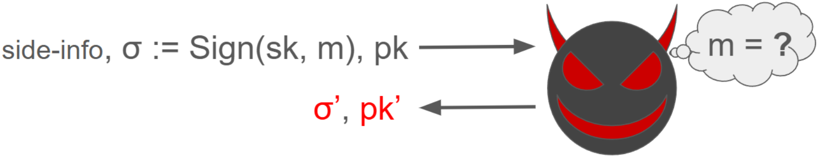

$$\Pr\left[\begin{matrix} pk \neq pk' \\ \mathrm{Ver}(pk', m, \sigma') = 1 \end{matrix}\right] < \text{small}$$

Security Beyond Unforgeability

NIST asked for "additional desirable security properties":

- exclusive ownership (S-CEO, S-DEO, M-S-UEO)
- message-bound signatures (MBS)
- **non-resignability (NR)**

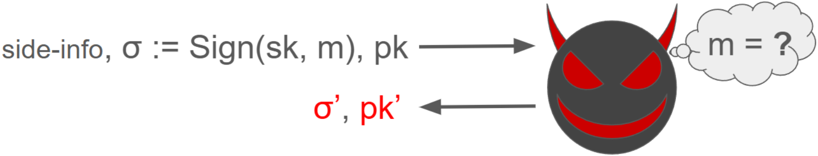

$$\Pr\left[\begin{array}{c} pk \neq pk' \\ \mathrm{Ver}(pk', m, \sigma') = 1 \end{array}\right] < \text{small}$$

uncertainty of m via **statistical/computational (HILL)** entropy

$$\mathrm{H}_\infty(m \mid \mathrm{pk}, \text{side-info}) \geq \text{high} \; .$$

Remark. $m \not\leftarrow (\mathrm{pk}, \sigma)$

Security Beyond Unforgeability

NIST asked for "additional desirable security properties":

- exclusive ownership (S-CEO, S-DEO, M-S-UEO)
- message-bound signatures (MBS)
- **non-resignability (NR)**

BUFF transformation [CDF+21],

$$\text{any signature } \mathcal{S} \mapsto \text{BUFF}[\mathcal{S}, H] \text{ with}$$

$$\sigma := \Big(Sign(sk, y), y\Big)\ , \text{ where } y := H(m, \text{pk})$$

- **claimed to give above securities**
- explicitly referred to by NIST
- relevant to Dilithium, Falcon, SPHINCS$^+$, HAWK and more.

Security Beyond Unforgeability

NIST asked for "additional desirable security properties":

- exclusive ownership (S-CEO, S-DEO, M-S-UEO)
- message-bound signatures (MBS)
- **non-resignability (NR)**

BUFF transformation [CDF$^+$21],

$$\text{any signature } \mathcal{S} \mapsto \text{BUFF}[\mathcal{S}, H] \text{ with}$$

$$\sigma := \Big(Sign(sk, y), y\Big) \text{ , where } y := H(m, \text{pk})$$

- **claimed to give above securities**
- explicitly referred to by NIST
- relevant to Dilithium, Falcon, SPHINCS$^+$, HAWK and more.

Plot-twist: NR as in [CDF$^+$21] is basically un-achievable!

Our Result, on the Negative side

In this work, we show:

1. Any "natural" signature scheme $\mathcal{S}$ is **not** NR.
2. $\forall \mathcal{S}$ and (sufficiently compressing) hash function H:

$$\text{BUFF}[\mathcal{S}, H] \text{ is } \textbf{not} \text{ NR.}$$

Our Result, on the Negative side

In this work, we show:

1. Any "natural" signature scheme $\mathcal{S}$ is **not** NR.
2. $\forall \mathcal{S}$ and (sufficiently compressing) hash function H:

$$\text{BUFF}[\mathcal{S}, H] \text{ is } \textbf{not} \text{ NR.}$$

Contradicting claimed BUFF security in [CDF+21]!

Our Result, on the Negative side

In this work, we show:

1. Any "natural" signature scheme $\mathcal{S}$ is **not** NR.
2. $\forall \mathcal{S}$ and (sufficiently compressing) hash function H:

   BUFF$[\mathcal{S}, H]$ is **not** NR.

   Contradicting claimed BUFF security in [CDF+21]!

A claim "random oracle is Φ-non-malleable" is false:

3. For any "sufficiently compressing" hash function H,

   $\exists$ attack that breaks Φ-non-malleability.

Our Result, on the Negative side

In this work, we show:

1. Any "natural" signature scheme $\mathcal{S}$ is **not** NR.
2. $\forall \mathcal{S}$ and (sufficiently compressing) hash function H:

   $$\text{BUFF}[\mathcal{S}, H] \text{ is } \textbf{not} \text{ NR.}$$

   Contradicting claimed BUFF security in [CDF+21]!

A claim "random oracle is Φ-non-malleable" is false:

3. For any "sufficiently compressing" hash function H,

   $$\exists \text{ attack that breaks } \Phi\text{-non-malleability.}$$

All of the above applies to both plain model and (Q)ROM.

Positive and More Negative Results

We then introduce a weakened notion:

4. $\mathsf{NR}^{\perp}$ in (Q)ROM where generic attacks no longer apply

    still meaningful for intended applications

Positive and More Negative Results

We then introduce a weakened notion:

4. $\mathsf{NR}^{\perp}$ in (Q)ROM where generic attacks no longer apply
   still meaningful for intended applications

To achieve $\mathsf{NR}^{\perp}$, we propose a **salted** variant \$-BUFF.

5. Under **statistical** entropy requirement:
   $\forall \mathcal{S}$: \$-BUFF[$\mathcal{S}, H$] is $\mathsf{NR}^{\perp}$ in (Q)ROM.
6. Under **HILL** entropy requirement: assuming CDH,
   $\exists \mathcal{S}$: \$-BUFF[$\mathcal{S}, H$] is **not** $\mathsf{NR}^{\perp}$ in (Q)ROM.

Positive and More Negative Results

We then introduce a weakened notion:

4. $\mathsf{NR}^{\perp}$ in (Q)ROM where generic attacks no longer apply

   still meaningful for intended applications

To achieve $\mathsf{NR}^{\perp}$, we propose a **salted** variant \$-BUFF.

5. Under **statistical** entropy requirement:

   $\forall \mathcal{S}$: \$-BUFF$[\mathcal{S}, H]$ is $\mathsf{NR}^{\perp}$ in (Q)ROM.

6. Under **HILL** entropy requirement: assuming CDH,

   $\exists \mathcal{S}$: \$-BUFF$[\mathcal{S}, H]$ is **not** $\mathsf{NR}^{\perp}$ in (Q)ROM.

   In fact **neither** is BUFF$[\mathcal{S}, H]$!

Addendum: responding to our work, [CDF+21] was updated to [CDF+23], but the security reasoning remains flawed.

Positive and More Negative Results

We then introduce a weakened notion:

4. $\mathsf{NR}^{\perp}$ in (Q)ROM where generic attacks no longer apply

   still meaningful for intended applications

To achieve $\mathsf{NR}^{\perp}$, we propose a **salted** variant \$-BUFF.

5. Under **statistical** entropy requirement:

   $\forall \mathcal{S}$: \$-BUFF$[\mathcal{S}, H]$ is $\mathsf{NR}^{\perp}$ in (Q)ROM.

6. Under **HILL** entropy requirement: assuming CDH,

   $\exists \mathcal{S}$: \$-BUFF$[\mathcal{S}, H]$ is **not** $\mathsf{NR}^{\perp}$ in (Q)ROM.

   In fact **neither** is BUFF$[\mathcal{S}, H]$!

Addendum: responding to our work, [CDF^{+}21] was updated to [CDF^{+}23], but the security reasoning remains flawed.

Take-away: non-resignability is brittle...

Overview

- Negative Results
- Positive (and More Negative) Results
- Conclusion

Non-resignability

Formally modelled via a two-staged game.

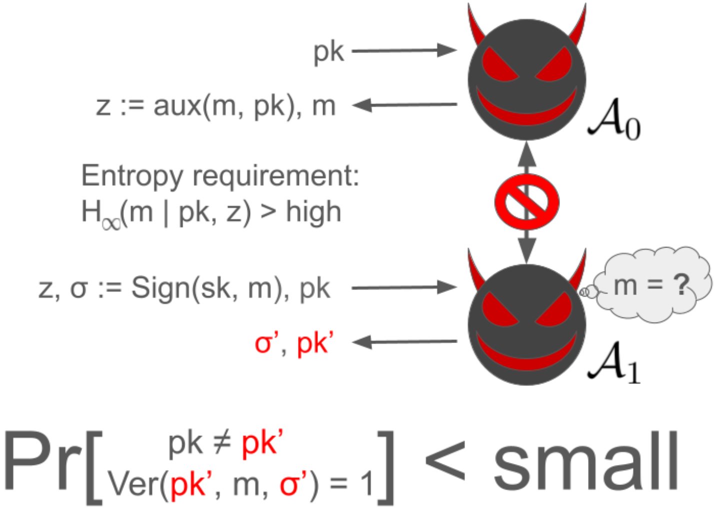

Non-resignability Attacked

Attackers can exploit side-info of *m*, while *m* remains hidden.

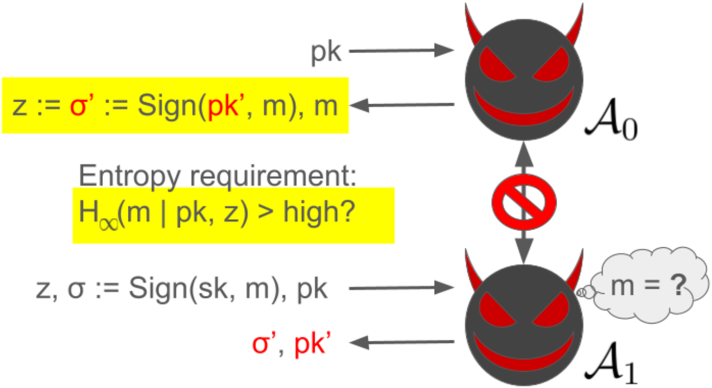

$$\Pr\left[\begin{array}{c} pk \neq pk' \\ \mathrm{Ver}(pk', m, \sigma') = 1 \end{array}\right] \approx 1$$

Non-resignability Attacked

Attackers can exploit side-info of m, while m remains hidden.

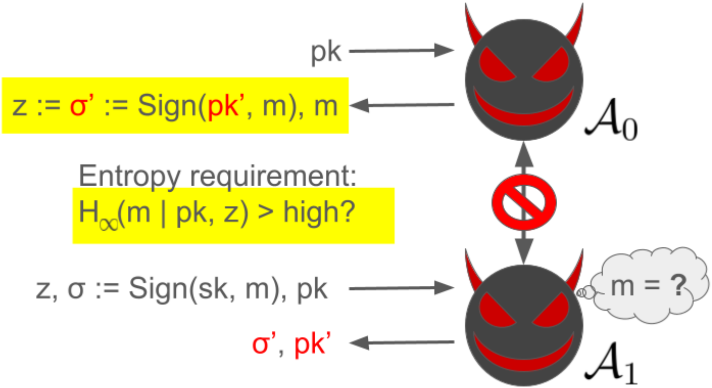

Case 1. $m \overset{\text{eff.}}{\leftarrow} (\mathsf{pk}, \sigma) \Rightarrow \mathcal{S}$ is trivially **not** NR

Non-resignability Attacked

Attackers can exploit side-info of m, while m remains hidden.

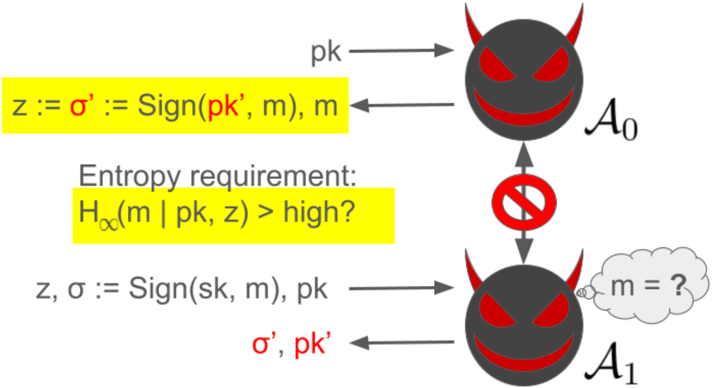

Case 1. $m \overset{\text{eff.}}{\leftarrow} (\mathsf{pk}, \sigma) \Rightarrow \mathcal{S}$ is trivially **not** NR
Case 2. $\mathsf{H}_\infty(m \mid \mathsf{pk}, \sigma) \geq \text{high}$
$\Rightarrow$ entropy cond. is satisfied $\Rightarrow$ the NR attack is valid

Wait a Minute...[1]

Claimed BUFF Security [CDF+21] $\rightarrow\leftarrow$ Generic NR attack

[1] Meme from https://emoji.gg/emoji/3803_Thonking with basic license.

What's Wrong?

[CDF$^+$21, Theorem 5.5] showed:

$$H \text{ is } \Phi\text{-non-malleable (for suitable } \Phi) \Rightarrow \mathsf{BUFF}[\mathcal{S}, H] \text{ is NR} .$$

What's Wrong?

[CDF$^+$21, Theorem 5.5] showed:

H is Φ-non-malleable (for suitable Φ) $\Rightarrow$ BUFF$[\mathcal{S}, H]$ is NR .

[BFS11, CDF$^+$21] claimed Φ-non-malleability of RO.

What's Wrong?

[CDF$^+$21, Theorem 5.5] showed:

H is Φ-non-malleable (for suitable Φ) $\Rightarrow$ BUFF[$\mathcal{S}, H$] is NR .

[BFS11, CDF$^+$21] claimed Φ-non-malleability of RO.

Any (sufficiently compressing) hash H is **not** Φ-non-malleable!

Overview

Negative Results
- Positive (and More Negative) Results
- Conclusion

Properly Re-define NR

Observation: side-info typically doesn't contain hashes.

Properly Re-define NR

Observation: side-info typically doesn't contain hashes.

a weakening $\text{NR}^{\perp}$ with **restricted side-info** in the (Q)ROM

The $\text{NR}^{\perp}$ game:

1: $m \leftarrow \mathcal{A}_0^H(\text{pk})$
2: $\sigma \leftarrow \text{Sign}^H(\text{sk}, m)$
3: $(\text{pk}', \sigma') \leftarrow \mathcal{A}_1^H(\text{pk}, \sigma, \text{aux}^{\cancel{H}}(m, \text{pk}))$
4: **return** $\text{Ver}^H(\text{pk}, m, \sigma') \ \wedge \ \text{pk} \neq \text{pk}'$

Properly Re-define NR

Observation: side-info typically doesn't contain hashes.

a weakening $\mathsf{NR}^{\perp}$ with **restricted side-info** in the (Q)ROM

The $\mathsf{NR}^{\perp}$ game:

1: $m \leftarrow \mathcal{A}_0^H(\mathsf{pk})$
2: $\sigma \leftarrow \mathsf{Sign}^H(\mathsf{sk}, m)$
3: $(\mathsf{pk}', \sigma') \leftarrow \mathcal{A}_1^H(\mathsf{pk}, \sigma, \mathsf{aux}^{\cancel{H}}(m, \mathsf{pk}))$
4: **return** $\mathsf{Ver}^H(\mathsf{pk}, m, \sigma') \ \wedge \ \mathsf{pk} \neq \mathsf{pk}'$

Definition 1. A signature is $\mathsf{NR}^{\perp}$, if $\forall(\mathcal{A}_0, \mathcal{A}_1, \mathsf{aux})$ under the (statistical/computational) entropy requirement $\Pr[1 \leftarrow \mathsf{NR}^{\perp}] \leq \mathsf{small}$.

Properly Re-define NR

Observation: side-info typically doesn't contain hashes.

a weakening $\mathsf{NR}^{\perp}$ with **restricted side-info** in the (Q)ROM

The $\mathsf{NR}^{\perp}$ game:

1: $m \leftarrow \mathcal{A}_0^H(\mathsf{pk})$
2: $\sigma \leftarrow \mathsf{Sign}^H(\mathsf{sk}, m)$
3: $(\mathsf{pk}', \sigma') \leftarrow \mathcal{A}_1^H(\mathsf{pk}, \sigma, \mathsf{aux}^{\cancel{H}}(m, \mathsf{pk}))$
4: **return** $\mathsf{Ver}^H(\mathsf{pk}, m, \sigma') \ \wedge \ \mathsf{pk} \neq \mathsf{pk}'$

The generic attack no longer applies to $\mathsf{NR}^{\perp}$:

$$\mathsf{aux}(m, \mathsf{pk}) := \xcancel{\mathsf{Sign}^H(\mathsf{sk}', m)}\ .$$

Redeeming $NR^{\perp}$

Does BUFF provide $NR^{\perp}$?

Redeeming $NR^{\perp}$

Does BUFF provide $NR^{\perp}$? it's not clear.

Redeeming $\mathsf{NR}^{\perp}$

Does BUFF provide $\mathsf{NR}^{\perp}$? it's not clear.

Instead, we consider a **salted** variant \$-BUFF:

$$\sigma := \left(\mathsf{Sign}(\mathsf{sk}, y_s), y_s, s\right) \text{ , where } s \leftarrow \{0,1\}^{\ell} \text{ and } y_s := H(m, \mathsf{pk}, s)$$

Redeeming $\mathsf{NR}^{\perp}$

Does BUFF provide $\mathsf{NR}^{\perp}$? it's not clear.

Instead, we consider a **salted** variant \$-BUFF:

$$\sigma := \left(\mathsf{Sign}(\mathsf{sk}, y_s), y_s, s\right) \text{ , where } s \leftarrow \{0,1\}^{\ell} \text{ and } y_s := H(m, \mathsf{pk}, s)$$

Under statistical entropy requirement: \$-BUFF$[\mathcal{S}, H]$ is $\mathsf{NR}^{\perp}$ $\forall \mathcal{S}$.

Redeeming $\mathsf{NR}^{\perp}$

Does BUFF provide $\mathsf{NR}^{\perp}$? it's not clear.

Instead, we consider a **salted** variant \$-BUFF:

$$\sigma := \left(\mathsf{Sign}(\mathsf{sk}, y_s), y_s, s\right), \text{ where } s \leftarrow \{0,1\}^{\ell} \text{ and } y_s := H(m, \mathsf{pk}, s)$$

Under statistical entropy requirement: \$-BUFF$[\mathcal{S}, H]$ is $\mathsf{NR}^{\perp}$ $\forall \mathcal{S}$.

Under only HILL entropy requirement:

- Assuming CDH, there is a strongly unforgeable signature $\mathcal{S}$, for which \$-BUFF$[\mathcal{S}, H]$ is not $\mathsf{NR}^{\perp}$.
- The same insecurity also applies to BUFF.

\$-BUFF[$\mathcal{S}, H$] is $\text{NR}^{\perp}$

Under Statistical Entropy Requirement

Following the proof strategy as in [CDF+21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM $\Rightarrow$ \$-BUFF[$\mathcal{S}, H$] is $\text{NR}^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

\$-BUFF[$\mathcal{S}$, H] is $\mathrm{NR}^{\perp}$

Under Statistical Entropy Requirement

Following the proof strategy as in [CDF+21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM $\Rightarrow$ \$-BUFF[$\mathcal{S}$, H] is $\mathrm{NR}^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

  ↑ the tricky part, previously undealt

\$-BUFF[$\mathcal{S}$, H] is NR$^{\perp}$

Under Statistical Entropy Requirement

Following the proof strategy as in [CDF^{+}21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM ⇒ \$-BUFF[$\mathcal{S}$, H] is NR$^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

  ↑ the tricky part, previously undealt

Sophisticated quantum argumentation:

- one-way-to-hiding lemma [AHU19]
- adaptive-reprogramming lemma [GHHM21]
- measure-and-reprogram technique [DFM20] but enhanced with a "stingy" simulator

\$-BUFF[$\mathcal{S}, H$] is NR$^{\perp}$

Under Statistical Entropy Requirement

Following the proof strategy as in [CDF$^+$21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM $\Rightarrow$ \$-BUFF[$\mathcal{S}, H$] is NR$^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

  ↑ the tricky part, previously undealt

Sophisticated quantum argumentation:

- one-way-to-hiding lemma [AHU19]
- adaptive-reprogramming lemma [GHHM21]
- measure-and-reprogram technique [DFM20] but enhanced with a "stingy" simulator

See our paper for more detail!

\$-BUFF[$\mathcal{S}$, H] is not $\mathrm{NR}^{\perp}$

Under HILL Entropy Requirement

Following the proof strategy as in [CDF+21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM $\not\Rightarrow$ \$-BUFF[$\mathcal{S}$, H] is $\mathrm{NR}^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

\$-BUFF[$\mathcal{S}$, H] is not $\mathrm{NR}^{\perp}$

Under HILL Entropy Requirement

Following the proof strategy as in [CDF+21]:

- Define \$-Φ-NM: a tailored variant of Φ-NM
- H is \$-Φ-NM $\not\Rightarrow$ \$-BUFF[$\mathcal{S}$, H] is $\mathrm{NR}^{\perp}$
- Prove that the random oracle H is \$-Φ-NM.

See full paper for simple CDH-based counterexample.

Overview

Negative Results

Positive (and More Negative) Results

- Conclusion

Conclusion

Defining/achieving non-resignability is much more subtle than what's believed.

Follow-up Questions

We've analyzed salted BUFF, what about the unsalted one?

- Is BUFF[$\mathcal{S}, H$] $\mathsf{NR}^{\perp}$ under statistical entropy requirement?
- Does BUFF[$\mathcal{S}, H$] satisfy any notion of NR computationally?

Follow-up Questions

We've analyzed salted BUFF, what about the unsalted one?

- Is BUFF[$\mathcal{S}, H$] $\text{NR}^{\perp}$ under statistical entropy requirement?
- Does BUFF[$\mathcal{S}, H$] satisfy any notion of NR computationally?

A follow-up work [DFH^{+}24]: Yes (to both)!

Follow-up Questions

We've analyzed salted BUFF, what about the unsalted one?

- Is BUFF$[\mathcal{S}, H]$ NR$^{\perp}$ under statistical entropy requirement?
- Does BUFF$[\mathcal{S}, H]$ satisfy any notion of NR computationally?

A follow-up work [DFH^{+}24]: Yes (to both)!

We've modelled the hash function as a RO:

- What about real-world hash functions, e.g. Sponge and/or Merkle-Damgard constructions?

That's It

Thank you for listening.

Eprint: `ia.cr/2023/1634`

References I

[AHU19] Andris Ambainis, Mike Hamburg, and Dominique Unruh. Quantum security proofs using semi-classical oracles. In Alexandra Boldyreva and Daniele Micciancio, editors, Advances in Cryptology – CRYPTO 2019, pages 269–295, Cham, 2019. Springer International Publishing.

[BFS11] Paul Baecher, Marc Fischlin, and Dominique Schröder. Expedient non-malleability notions for hash functions. In Aggelos Kiayias, editor, CT-RSA 2011, volume 6558 of LNCS, pages 268–283. Springer, Heidelberg, February 2011.

References II

[CDF+21] Cas Cremers, Samed Düzlü, Rune Fiedler, Marc Fischlin, and Christian Janson. BUFFing signature schemes beyond unforgeability and the case of post-quantum signatures. In 2021 IEEE Symposium on Security and Privacy, pages 1696–1714. IEEE Computer Society Press, May 2021. Cryptology ePrint Archive version available at `https://eprint.iacr.org/archive/2020/1525/20230116:141028` (Version 1.3).

[CDF+23] Cas Cremers, Samed Düzlü, Rune Fiedler, Marc Fischlin, and Christian Janson. BUFFing signature schemes beyond unforgeability and the case of post-quantum signatures, 2023. An updated version (Version 1.4) of [CDF+21], available at `https://eprint.iacr.org/archive/2020/1525/20231020:082812`.

References III

[DFH⁺24] Jelle Don, Serge Fehr, Yu-Hsuan Huang, Jyun-Jie Liao, and Patrick Struck. Hide-and-seek and the non-resignability of the BUFF transform. Cryptology ePrint Archive, Paper 2024/793, 2024. https://eprint.iacr.org/2024/793.

[DFM20] Jelle Don, Serge Fehr, and Christian Majenz. The measure-and-reprogram technique 2.0: Multi-round fiat-shamir and more. In Daniele Micciancio and Thomas Ristenpart, editors, Advances in Cryptology – CRYPTO 2020, pages 602–631, Cham, 2020. Springer International Publishing.

References IV

[GHHM21] Alex B. Grilo, Kathrin Hövelmanns, Andreas Hülsing, and Christian Majenz. Tight adaptive reprogramming in the qrom. In Mehdi Tibouchi and Huaxiong Wang, editors, Advances in Cryptology – ASIACRYPT 2021, pages 637–667, Cham, 2021. Springer International Publishing.