Radbound University

Tight Preimage Resistance of the Sponge Construction

Charlotte Lefevre, Bart Mennink

Radboud University (The Netherlands)

NIST LWC Workshop

May 11, 2022

ESCADA

A generalized sponge construction [Bertoni et al., 2007]

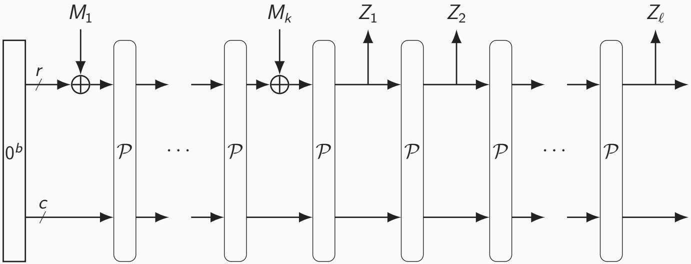

- $M_1 \| \cdots \| M_k$ is the message padded into r-bit blocks
- Variable-length digest, if n bits required, the digest is the first n bits of $Z_1 \| \cdots \| Z_\ell$

A generalized sponge construction [Bertoni et al., 2007]

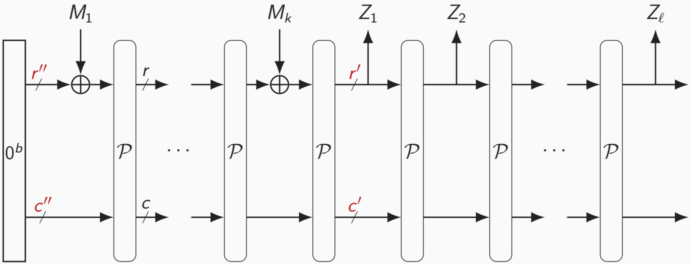

- $M_1 \| \cdots \| M_k$ is the message padded into r-bit blocks
- Variable-length digest, if n bits required, the digest is the first n bits of $Z_1 \| \cdots \| Z_\ell$
- The first message block can be larger, can squeeze at a larger rate [Guo et al., 2011, PHOTON]

Classical security requirements

Preimage

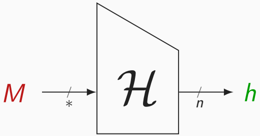

Given h, find M

Second Preimage

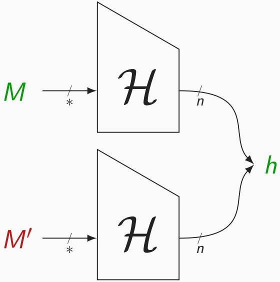

Given M, find $M' \neq M$

Collision

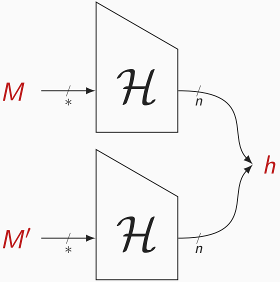

Find $M \neq M'$

Classical security requirements

Preimage

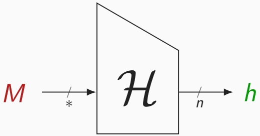

Given h, find M

Second Preimage

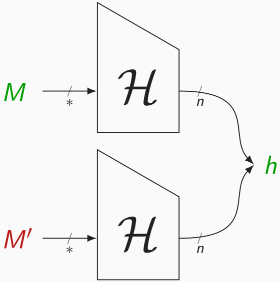

Given M, find $M' \neq M$

Collision

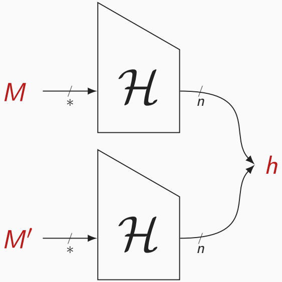

Find $M \neq M'$

- Insufficient for certain applications (e.g., $MAC(k, m) = \mathcal{H}(k \| m)$ with $\mathcal{H}$ = plain Merkle-Damgård)
- Hash function should behave like a random oracle

Indifferentiability [Maurer et al., 2004, Coron et al., 2005]

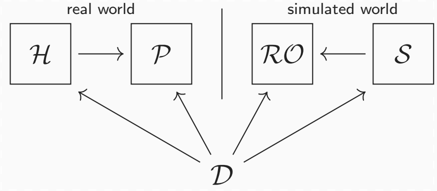

- $(\mathcal{H}^{\mathcal{P}}, \mathcal{P})$ for a random primitive $\mathcal{P}$ should behave like a random oracle $\mathcal{RO}$ paired with a simulator $\mathcal{S}$ that maintains construction-primitive consistency
- $\mathcal{H}$ is indifferentiable from $\mathcal{RO}$ for some simulator $\mathcal{S}$ whenever any $\mathcal{D}$ can distinguish the two worlds only with a negligible probability

Indifferentiability [Maurer et al., 2004, Coron et al., 2005]

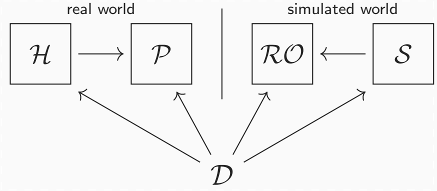

- $(\mathcal{H}^{\mathcal{P}}, \mathcal{P})$ for a random primitive $\mathcal{P}$ should behave like a random oracle $\mathcal{RO}$ paired with a simulator $\mathcal{S}$ that maintains construction-primitive consistency
- $\mathcal{H}$ is indifferentiable from $\mathcal{RO}$ for some simulator $\mathcal{S}$ whenever any $\mathcal{D}$ can distinguish the two worlds only with a negligible probability
- Indifferentiability $\Longrightarrow$ Pre/SecPre/Col security [Andreeva et al., 2010]

Indifferentiability of the sponge construction

- The (generalized) sponge construction was proven indifferentiable with a bound $\mathcal{O}\left(\frac{q}{2^{c/2}}\right)$, [Bertoni et al., 2008, Naito and Ohta, 2014] [1]

$\Longrightarrow$ The sponge is unlikely differentiable from a $\mathcal{RO}$ with less than $q \approx 2^{c/2}$ queries

- This implies the following security bounds:

Security property	Security bound
Indifferentiability	$\mathcal{O}\left(\frac{q}{2^{c/2}}\right)$
Col	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q^2}{2^n}\right)$
SecPre	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q}{2^n}\right)$
Pre	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q}{2^n}\right)$

indifferentiability
(q is number of primitive queries)

classical attacks against $\mathcal{RO}$
(q is number of oracle queries)

[1] As long as the first message block and squeezing rate are not too large

Indifferentiability of the sponge construction

- The (generalized) sponge construction was proven indifferentiable with a bound $\mathcal{O}\left(\frac{q}{2^{c/2}}\right)$, [Bertoni et al., 2008, Naito and Ohta, 2014] [1]

$\implies$ The sponge is unlikely differentiable from a $\mathcal{RO}$ with less than $q \approx 2^{c/2}$ queries

- This implies the following security bounds:

Security property	Security bound	Best attack cost	Tight?
Indifferentiability	$\mathcal{O}\left(\frac{q}{2^{c/2}}\right)$	$2^{c/2}$	Y
Col	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q^2}{2^n}\right)$	$\min\left\{2^{c/2}, 2^{n/2}\right\}$	Y
SecPre	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q}{2^n}\right)$	$\min\left\{2^{c/2}, 2^{n}\right\}$	Y
Pre	$\mathcal{O}\left(\frac{q}{2^{c/2}} + \frac{q}{2^n}\right)$	$\min\{2^{n-r'} + 2^{c/2}, 2^n\}$	N

- There is a gap in the first preimage security $\implies$ we fill it

[1] As long as the first message block and squeezing rate are not too large

Collision attack with $q \approx 2^{c/2}$ queries [Bertoni et al., 2011]

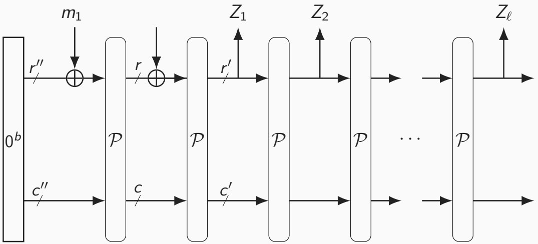

- Query $P(m_1 \| 0^c)$ for $2^{c/2}$ different m_1's, store them in a list L

Collision attack with $q \approx 2^{c/2}$ queries [Bertoni et al., 2011]

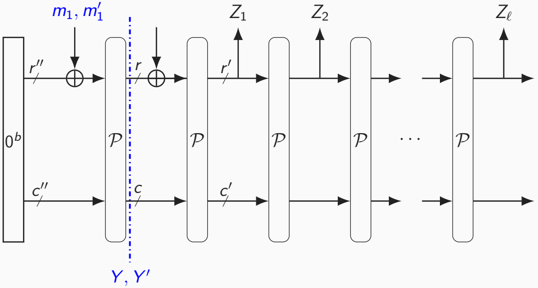

- Query $P(m_1 \| 0^c)$ for $2^{c/2}$ different m_1's, store them in a list L
- With high probability, there exists $Y \neq Y' \in L$ s.t., $\mathrm{inner}_c(Y) = \mathrm{inner}_c(Y')$

Collision attack with $q \approx 2^{c/2}$ queries [Bertoni et al., 2011]

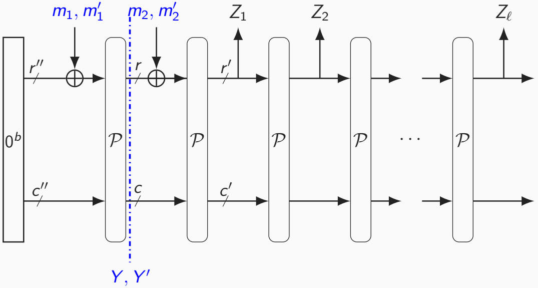

- Query $P(m_1 \| 0^c)$ for $2^{c/2}$ different m_1's, store them in a list L
- With high probability, there exists $Y \neq Y' \in L$ s.t., $\mathrm{inner}_c(Y) = \mathrm{inner}_c(Y')$

$\implies$ take $m_2 = \mathrm{outer}_r(Y), m_2' = \mathrm{outer}_r(Y')$

$\implies$ It gives $\mathcal{H}(\mathrm{unpad}(m_1 \| m_2)) = \mathcal{H}(\mathrm{unpad}(m_1' \| m_2'))$

Second preimage attack with $q \approx 2^{c/2}$ queries [Bertoni et al., 2011]

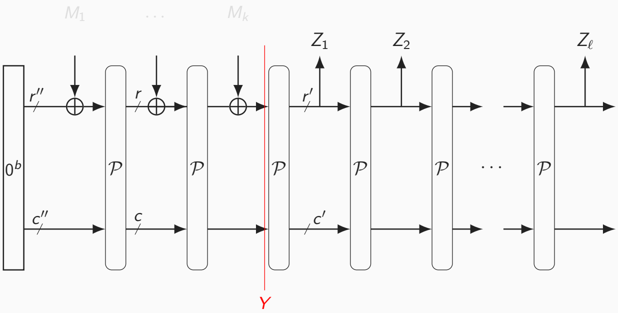

- Let M be the first preimage, $M_1 \| \dots \| M_k := \text{pad}(M)$
- Compute the state before the first squeeze, call it Y

Second preimage attack with $q \approx 2^{c/2}$ queries [Bertoni et al., 2011]

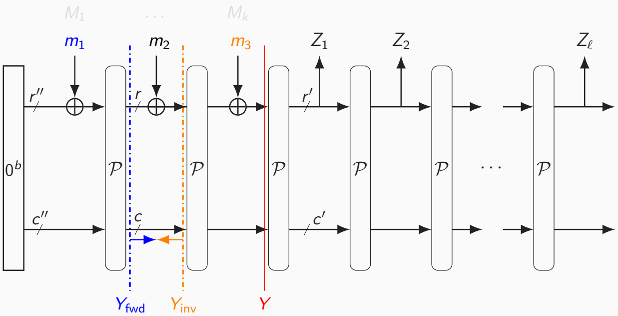

- Let M be the first preimage, $M_1 \| \dots \| M_k := \text{pad}(M)$
- Compute the state before the first squeeze, call it Y
- Reach Y with an inner forward/backward collision, compensate the outer part with $m_2 = \text{outer}_r(Y_{\text{inv}}) \oplus \text{outer}_r(Y_{\text{fwd}})$

$\Longrightarrow$ It gives $\mathcal{H}(\text{unpad}(m_1 \| m_2 \| m_3)) = \mathcal{H}(M)$

First preimage attack [Bertoni et al., 2011]

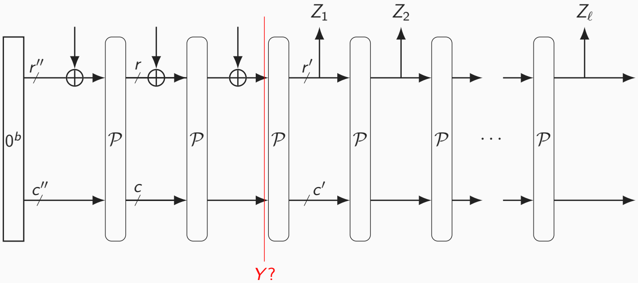

- Let $Z = Z_1 \| \cdots \| Z_\ell$ be the image
- Here, there is no intermediate state Y: we need to find it before applying the same attack
- More precisely, we need Y s.t. $\forall i = 1, \dots, \ell, \ \mathrm{outer}_{r'}(P^i(Y)) = Z_i$

First preimage attack [Bertoni et al., 2011]

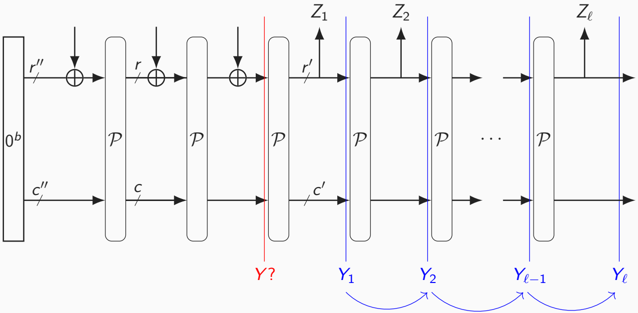

- Let $Z = Z_1 \| \cdots \| Z_\ell$ be the image
- Here, there is no intermediate state Y: we need to find it before applying the same attack
- More precisely, we need Y s.t. $\forall i = 1, \ldots, \ell$, $\mathrm{outer}_{r'}(P^i(Y)) = Z_i$
- One attempt succeeds with probability $\approx \frac{1}{2^{n-r'}} \implies$ this attack succeeds after $\approx 2^{n-r'} + 2^{c/2}$ queries

Improved first preimage resistance

- This first preimage attack succeeds after $\approx \min\{2^{n-r'} + 2^{c/2}, 2^n\}$ queries, while the bound from indifferentiability guarantees preimage security up to $\approx \min\left\{2^{c/2}, 2^n\right\}$ queries

$\Longrightarrow$ It does not match the attack when $c/2 \leq n - r'$

- Our contribution: we prove the preimage resistance with bound

$$\mathcal{O}\left(\frac{q}{2^n} + \min\left\{\frac{q}{2^{n-r'}}, \frac{q^2}{2^c}\right\}\right)$$

$\Longrightarrow$ It shows optimality of the attack

Improved first preimage resistance

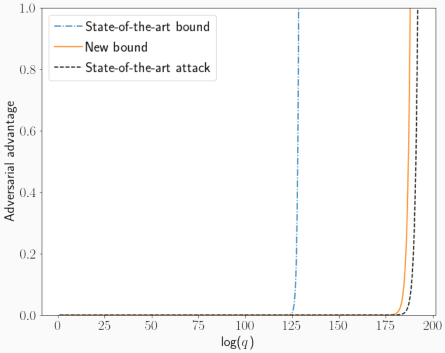

Ascon
$(b, c, r, r', n) = (320, 256, 64, 64, 256)$

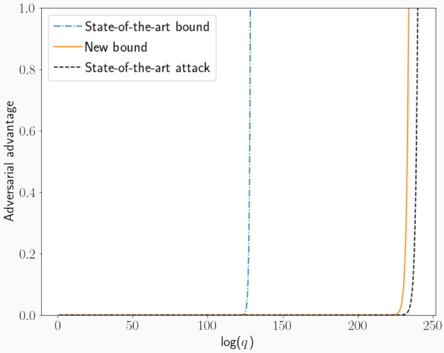

Spongent largest mode
$(b, c, r, r', n) = (272, 256, 16, 16, 256)$

Adversarial advantage upperbound according to number of queries

Improved first preimage resistance

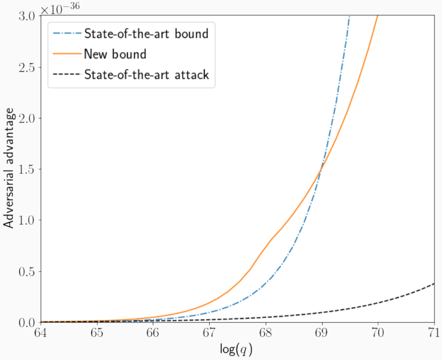

Ascon

$(b, c, r, r', n) = (320, 256, 64, 64, 256)$

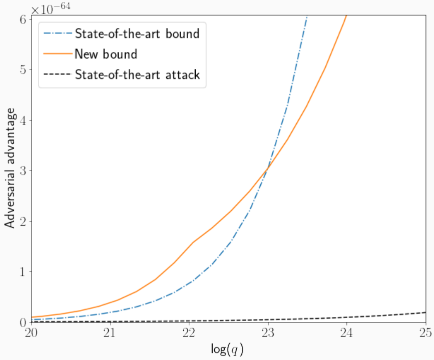

Spongent largest mode

$(b, c, r, r', n) = (272, 256, 16, 16, 256)$

Closeup

Proof idea

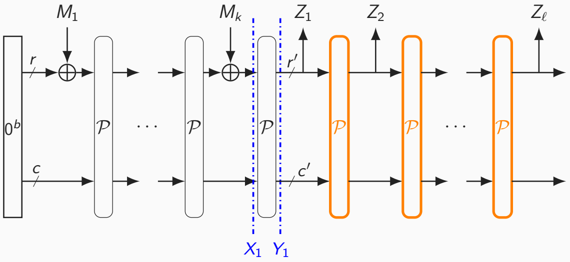

- To find a preimage, the adversary must find a **cascade** of $\ell - 1$ permutation evaluations giving $Z_1, \ldots Z_\ell$
- But this is not enough, this cascade must be reached from 0^b
- Depending on the direction of the query $X_1 \to Y_1$, there are two scenarios:

Proof idea

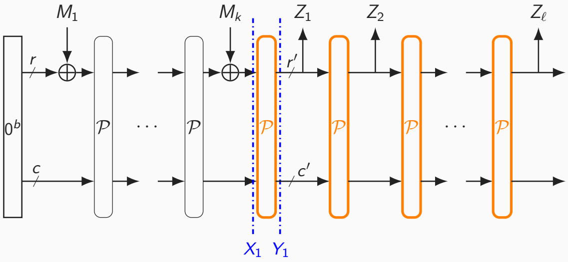

- To find a preimage, the adversary must find a **cascade** of $\ell - 1$ permutation evaluations giving $Z_1, \dots Z_\ell$
- But this is not enough, this cascade must be reached from 0^b
- Depending on the direction of the query $X_1 \to Y_1$, there are two scenarios:
  - Forward direction: the cascade is extended by one

Proof idea

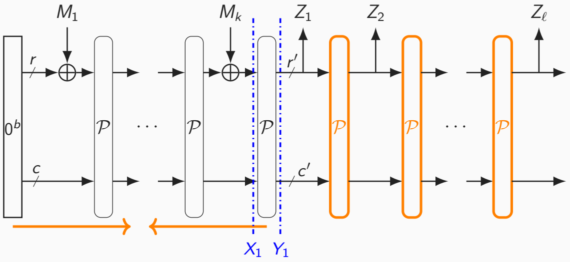

- To find a preimage, the adversary must find a **cascade** of $\ell - 1$ permutation evaluations giving $Z_1, \ldots Z_\ell$
- But this is not enough, this cascade must be reached from 0^b
- Depending on the direction of the query $X_1 \to Y_1$, there are two scenarios:
  - Forward direction: the cascade is extended by one
  - Inverse direction: an inner collision must have been found

Proof idea - probability computation

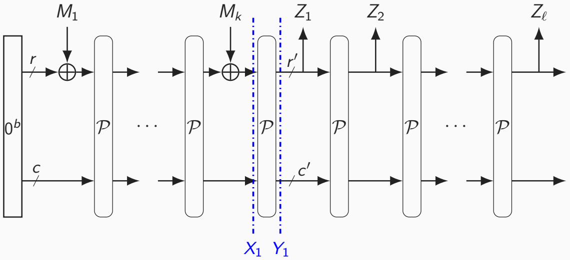

- Probability of inner collision upper bounded by $\frac{q(q+1)}{2^c}$
- Probability of finding a cascade:
  - Forward direction: adversary must guess a "good" $X_1 \implies \mathcal{O}\left(\frac{q}{2^n}\right)$
  - Inverse direction: more involved, since the queries can appear in any order, any direction within the cascade $\implies \mathcal{O}\left(\frac{q}{2^{n-r'}}\right)$

Impact on the generic security of a few schemes

Scheme	Parameters						Security bound		Note
	b	c	r	r'	n	ℓ	Old	New	
Spongent	272	256	16	16	256	16	2^{128}	2^{240}	ISO/IEC standard
PHOTON	288	256	32	32	256	8	2^{128}	2^{224}	ISO/IEC standard
ACE-Hash	320	256	64	64	256	4	2^{128}	2^{192}	NIST LWC round 2
KNOT Hash	256	224	32	128	256	2	2^{112}	2^{128}	NIST LWC round 2
	384	336	48	192	384	2	2^{168}	2^{192}	
	512	448	64	256	512	2	2^{224}	2^{256}	
SKINNY-tk2-Hash	256	224	32	128	256	2	2^{112}	2^{128}	NIST LWC round 2
Subterranean 2.0	257	248	9	32	256	8	2^{124}	2^{224}	NIST LWC round 2
Ascon-Hash	320	256	64	64	256	4	2^{128}	2^{192}	NIST LWC finalist
PHOTON-Beetle-Hash	256	224	32	128	256	2	2^{112}	2^{128}	NIST LWC finalist

Conclusion

- We derived a tight security bound for the first preimage of the sponge construction
- This bound has direct implications on the security of lightweight cryptographic sponges

Thank you for your attention!

Andreeva, E., Mennink, B., and Preneel, B. (2010).
Security Reductions of the Second Round SHA-3 Candidates.
In Burmester, M., Tsudik, G., Magliveras, S. S., and Ilic, I., editors, *Information Security - 13th International Conference, ISC 2010, Boca Raton, FL, USA, October 25-28, 2010, Revised Selected Papers*, volume 6531 of *Lecture Notes in Computer Science*, pages 39–53. Springer.

Bertoni, G., Daemen, J., Peeters, M., and Van Assche, G. (2007).
Sponge functions.
Ecrypt Hash Workshop 2007.

Bertoni, G., Daemen, J., Peeters, M., and Van Assche, G. (2008).
On the Indifferentiability of the Sponge Construction.
In Smart, N. P., editor, *Advances in Cryptology - EUROCRYPT 2008, 27th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Istanbul, Turkey, April 13-17, 2008. Proceedings*, volume 4965 of *Lecture Notes in Computer Science*, pages 181–197. Springer.

Bertoni, G., Daemen, J., Peeters, M., and Van Assche, G. (2011).
Cryptographic sponge functions.
`https://keccak.team/files/CSF-0.1.pdf`.

Coron, J., Dodis, Y., Malinaud, C., and Puniya, P. (2005).
Merkle-Damgård Revisited: How to Construct a Hash Function.
In Shoup, V., editor, *Advances in Cryptology - CRYPTO 2005: 25th Annual International Cryptology Conference, Santa Barbara, California, USA, August 14-18, 2005, Proceedings*, volume 3621 of *Lecture Notes in Computer Science*, pages 430–448. Springer.

Guo, J., Peyrin, T., and Poschmann, A. (2011).
The PHOTON Family of Lightweight Hash Functions.
In Rogaway, P., editor, *Advances in Cryptology - CRYPTO 2011 - 31st Annual Cryptology Conference, Santa Barbara, CA, USA, August 14-18, 2011. Proceedings*, volume 6841 of *Lecture Notes in Computer Science*, pages 222–239. Springer.

Maurer, U. M., Renner, R., and Holenstein, C. (2004).
Indifferentiability, Impossibility Results on Reductions, and Applications to the Random Oracle Methodology.
In Naor, M., editor, *Theory of Cryptography, First Theory of Cryptography Conference, TCC 2004, Cambridge, MA, USA, February 19-21, 2004, Proceedings*, volume 2951 of *Lecture Notes in Computer Science*, pages 21–39. Springer.

Naito, Y. and Ohta, K. (2014).
Improved Indifferentiable Security Analysis of PHOTON.
In Abdalla, M. and Prisco, R. D., editors, *Security and Cryptography for Networks - 9th International Conference, SCN 2014, Amalfi, Italy, September 3-5, 2014. Proceedings*, volume 8642 of *Lecture Notes in Computer Science*, pages 340–357. Springer.