On the security of the keyed sponge construction

Guido Bertoni[1] Joan Daemen[1]
Michaël Peeters[2] Gilles Van Assche[1]

[1]STMicroelectronics

[2]NXP Semiconductors

Symmetric Key Encryption Workshop (SKEW)
Lyngby, Denmark, February 16-17, 2011

Outline

The sponge construction

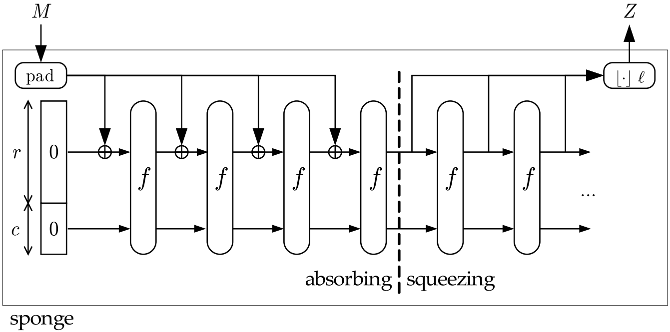

- f: a b-bit permutation with $b = r + c$

From hashing to encryption

- Hashing: $\textsc{sponge}(m) = h$
- Encryption as a stream cipher
  - Squeezing $\textsc{sponge}(K||\mathrm{IV})$, or
  - Random-access key stream block $k_i = \textsc{sponge}(K||\mathrm{IV}||i)$
- Authentication: $\textsc{sponge}(K||m) = \mathrm{MAC}$
  - Note: no need for HMAC construction
- Authenticated encryption using duplex
  - First call is $\textsc{duplex}.\mathrm{duplexing}(K)$
  - Further calls are equivalent to $\textsc{sponge}(K|| \ldots)$

Keyed sponge functions

Keyed sponge

$\textsc{KeyedSponge}[K](x) = \textsc{sponge}(K||x)$

- E.g., MAC $= \textsc{KeyedSponge}(m)$

Security against generic attacks

RO-differentiability advantage

- Provably secure against attacks with $< 2^{c/2}$ calls to f
  [Bertoni et al., Eurocrypt 2008]
- Proof assumes f is a random permutation
- So, SPONGE is secure if f has no exploitable properties

And for KEYEDSPONGE...

- Proof currently limited to $2^{c/2}$
  - **Can we go beyond?**

Indistinguishability setting

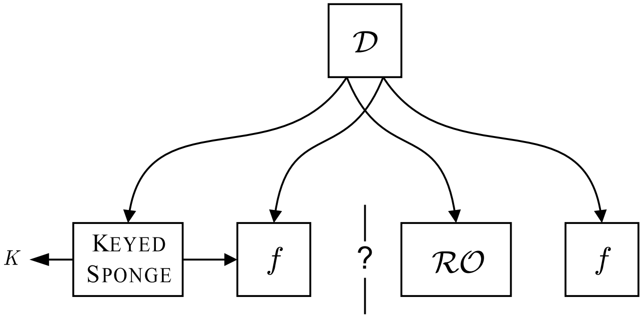

- M: online **data** complexity (blocks)
  - Calls to KEYEDSPONGE$[K]$ with unknown key K, or to $\mathcal{RO}$
- N: offline **time** complexity (calls to f)
  - Not involving the key

Indistinguishability theorem

Distinguishability upper bound

$$1 - \exp\left(-\frac{M^2/2 + 2MN}{2^c}\right) + P_{\text{key}}(N)$$

- $P_{\text{key}}(N)$: probability of guessing the key after N calls to f
  - i.e., of making a query to f with input in $\widehat{\text{absorb}}(K)$

If $M \ll 2^{c/2}$

Time complexity is about $\min(2^{c-1}/M, 2^{|K|})$

Limited data complexity

- If the (online) data complexity is limited to $M \leq 2^a$
  - … by the protocol, by the secure device …
- And the capacity is $c \geq |K| + a + 1$
- Then we get the security of the exhaustive key search

$$\min(2^{c-1}/M, 2^{|K|}) = 2^{|K|}$$

The new bound, illustrated

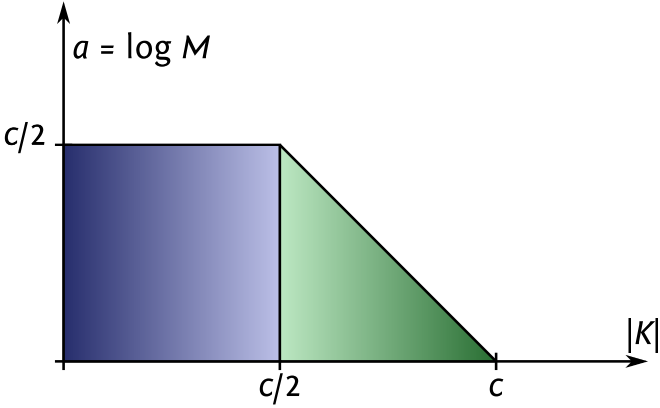

Building lightweight implementations

- Trade-off between security and efficiency
  - Security level determined by c
  - Efficiency: r input/output bits per call to f
- Example 1: QUARK [Aumasson et al., QUARK, ..., CHES 2010]

U-QUARK	$r = 8$	$c = 128$
D-QUARK	$r = 16$	$c = 160$
S-QUARK	$r = 32$	$c = 224$

- Example 2: KECCAK supports : $b \in \{25, 50, 100 \ldots 1600\}$
  - E.g., KECCAK$[r = 40, c = 160]$ is compact in hardware [Bertoni et al., KECCAK implementation overview]

Building implementations that are even lighter

Target example: 80-bit key with QUARK

- Old bound: D-QUARK ($r = 16$, $c = 160$)
  - $c = 2|K|$
- New bound: U-QUARK ($r = 8$, $c = 128$)
  - with data complexity restricted to 2^{47} blocks

If the distinguisher had no access to f...

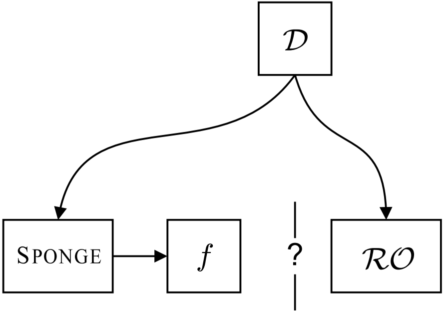

- Only distinguishing property: the **inner collisions** ($M^2 / 2^c$)
- No access to f: not very realistic...
  - [Bertoni et al., Sponge functions, 2007]

No inner clashes, please

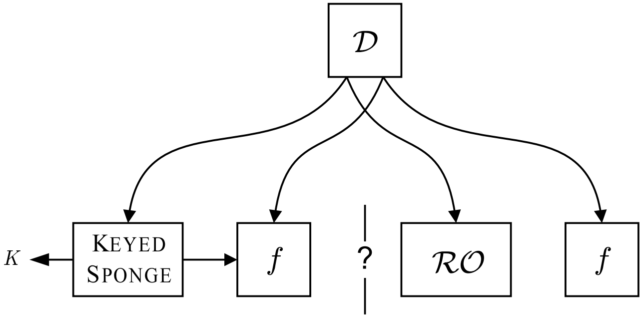

- Inner collisions in keyed sponge ($M^2 / 2^c$)
- Uniformity if no **inner clash** with queries to f ($MN / 2^c$)
  - Key guessing implies an inner clash

Conclusions

Thanks for your attention!

Q?