Short Non-Malleable Codes from Related-Key Secure Block Ciphers

Pierre Karpman
(joint work with Serge Fehr and Bart Mennink)

Université Grenoble Alpes

Séminaire MAD/CASYS — Grenoble
2017–10–12

Non-Malleable Codes (simple def.)

Non-Malleable Code (informal)

An NMC is a pair $(\mathsf{Enc}, \mathsf{Dec})$ where Enc is an *unkeyed* randomized mapping and we have:

1. $\forall m,\ \mathsf{Dec}(\mathsf{Enc}(m)) = m$
2. $\forall\, \mathsf{T} \in \mathcal{T},\ \mathsf{Dec}(\mathsf{T}(\mathsf{Enc}(m))) \approx \mathsf{Dec}(\mathsf{T}(\mathsf{Enc}(m')))$

for some function space $\mathcal{T}$.

- Introduced by Dziembowski, Pietrzak and Wichs (2010)

Non-Malleable Codes (why?)

Application example: *tamper-resilient* cryptography:

- Store secrets S in coded form
- Only protect the decoding
- $\Rightarrow$ The circuit never runs on $S' \neq S$ correlated with S

$\Rightarrow$ Decrease the "attack surface" & protection complexity/overhead

Non-Malleable Codes (feasibility)

- Restrictions on $\mathcal{T}$ necessary. Cannot include, say $(x \mapsto \mathsf{Enc}(\mathsf{Dec}(x) + 1)$
- Special "trick" to include Id and variants: use $\dot{\mathsf{Dec}}^{\mathsf{Enc}(x)}(\alpha)$, answers **same** if $\alpha = \mathsf{Enc}(x)$

An approach for $\mathcal{T}$: *split-state tampering* only:

Split-state tampering model

$\mathsf{Enc} : \{0,1\}^\kappa \times \mathcal{M} \to \{0,1\}^{\ell_\mathrm{L}} \times \{0,1\}^{\ell_\mathrm{R}}$
$\mathcal{T} = \{\mathsf{T} = \mathsf{T_L} \,\|\, \mathsf{T_R} : \{0,1\}^{\ell_\mathrm{L}} \times \{0,1\}^{\ell_\mathrm{R}} \to \{0,1\}^{\ell_\mathrm{L}} \times \{0,1\}^{\ell_\mathrm{R}}\}$

- Constructions exist in this model (computational or information-theoretic)

Non-Malleable v. Error-Correcting

- Possible to have NMCs with $\mathcal{T} \ni (x \mapsto 0)$ ("ultimate" error pattern)
- If correction is not possible, decoding must fail "catastrophically" ("all-or-nothing")

Formalizing security (in short)

Tamp

$\mathsf{Tamp}^{\mathsf{T}}(m) := \dot{\mathsf{Dec}}^{\mathsf{Enc}_K(m)} \circ \mathsf{T} \circ \mathsf{Enc}_K(m)$

For $K \xleftarrow{\$} \{0,1\}^{\kappa}$

$\mathbf{Adv}_{\mathsf{NMC}}$

$\mathbf{Adv}_{\mathsf{NMC}}(t) :=$

$$\max_{m,m'} \max_{A,\mathsf{T}} \left| \Pr[A(\mathsf{Tamp}^{\mathsf{T}}(m)) = 1] - \Pr[A(\mathsf{Tamp}^{\mathsf{T}}(m')) = 1] \right|$$

for A running in time t

A simple construction

Let $\mathcal{E} : \{0,1\}^\kappa \times \mathcal{M} \to \mathcal{M}$ be a block cipher. Define RKNMC$[\mathcal{E}]$ as:

- $\mathsf{Enc}_k := (m \mapsto k || \mathcal{E}_k(m))$
- $\mathsf{Dec} := (c_L || c_R \mapsto \mathcal{E}^{-1}_{c_L}(c_R))$

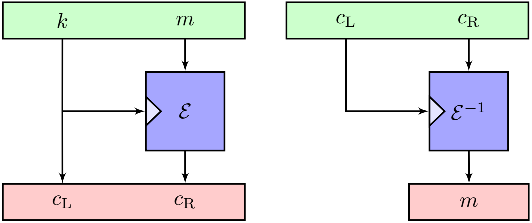

- Provides $\kappa/2$ bits of security, for "'good $\mathcal{E}$"

Related-work

- $m \mapsto (k, r) \| (\mathcal{E}_k(m), \mathcal{H}_z(r, k))$ (Kiayias & al., 2016)
  - Codewords of length $|m| + 9\kappa + 2\log^2(\kappa)$ or $|m| + 18\kappa$
  - Proof under KEA, with CRS
- $m \mapsto \mathsf{sk} \| (\mathsf{pk}, \mathcal{E}_{\mathsf{pk}}(m), \pi)$ (Liu and Lysyanskaya, 2012)
  - Codewords of length $|m| + \mathcal{O}(\kappa^2)$
  - Proof uses CRS

Related-work

Figure: KEA & CRS?

Broken instantiations

Take $\mathsf{EM}_k(m) := \mathcal{P}(m \oplus k) \oplus k$

- Secure in the ideal permutation model (Even & Mansour, 1991)
- But not *related-key* secure: $\mathsf{EM}_{k\oplus\Delta}(m \oplus \Delta) = \mathsf{EM}_k(m) \oplus \Delta$

So:

- Let $\mathsf{T_L} = \mathsf{T_R} = (x \mapsto x \oplus \Delta)$
- Then $\mathsf{Tamp}^{\mathsf{T}}(m) = \mathsf{EM}^{-1}_{k\oplus\Delta}(\mathsf{EM}_k(m) \oplus \Delta) = m \oplus \Delta$
- $\Rightarrow$ RKNMC[EM] is trivially insecure

Simulating Tamp from related-key queries

- Related-key attacks: the adversary can query $\mathcal{O}_k$, $\mathcal{O}_k^{-1}$, $\mathcal{O}_{\varphi(k)}$, $\mathcal{O}_{\varphi(k)}^{-1}$ for unknown k, chosen $\varphi \in \Phi$ w/ $\mathcal{O} = \mathcal{E}$ or $\mathcal{O} = \mathfrak{P}$
  - Objective: distinguish the two worlds
- Take $\mathrm{T} = \varphi \| \mathrm{T_R}$. For any m, the RK adversary can query $x := \mathcal{O}_k(m)$, $y := \mathcal{O}_{\varphi(k)}^{-1}(\mathrm{T_R}(x))$, run an NMC adversary $A(\mathrm{T}, m, \$)$ on y
- $\Rightarrow$ $\mathbf{Adv}_{\mathrm{RK}}$ w.r.t. φ is at least *not (much) less* than $\mathbf{Adv}_{\mathrm{NMC}}$ w.r.t. $\mathrm{Tamp}^{\mathrm{T}}$, $\mathrm{T} = \varphi \| \cdot$.

Related-key issues

- Problem: *generic* absence of RK security for unrestricted φ
- For instance, take $\varphi : x \mapsto 0$
- But $\mathsf{T_L} : x \mapsto 0$ *is* allowed
- $\Rightarrow$ upper-bounding $\mathbf{Adv}_{\mathsf{NMC}}$ by the $\mathbf{Adv}_{\mathsf{RK}}$ seems meaningless :(
- A condition for meaningful $\mathbf{Adv}_{\mathsf{RK}}$: $\varphi(K)$ "hard to guess" for uniform K

Related-key: the BK03 bound

- $\mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') :=$
$$\max_{P \subseteq \Phi, X \subseteq \mathcal{K}, \#P \leq r, \#X \leq r'} \Pr[\{\varphi(K) : \varphi \in P\} \cap X \neq \varnothing : K \xleftarrow{\$} \mathcal{K}]$$
- $\mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r) :=$
$$\max_{P \subseteq \Phi, \#P \leq r} \Pr[\#\{\varphi(K) : \varphi \in P\} < \#P : K \xleftarrow{\$} \mathcal{K}]$$

RK security of an ideal cipher

$$\mathbf{Adv}^{\text{prp-rka}}_{\Phi}(r, r') \leq \mathbf{InSec}^{\mathrm{up}}_{\Phi}(r, r') + \mathbf{InSec}^{\mathrm{cr}}_{\Phi}(r)$$

(Fun facts $\rightsquigarrow$ blackboard)

Switching to single-key security

- Take again $\mathsf{T_L} : x \mapsto 0$
- Then anyone with access to $\mathcal{O}_k$ may query $x := \mathcal{O}_k(m)$, $y := \mathcal{E}_0^{-1}(x)$
- $\Rightarrow$ $\mathbf{Adv}_{\mathsf{NMC}}$ w.r.t. such $\mathsf{T_L}$ reduces to *single key* security $\mathbf{Adv}_{\mathsf{PRP}}$ of $\mathcal{E}$!

More with single keys

- Take $\mathsf{T_L} : \{0,1\}^\kappa \to \{k_0, k_1, \ldots, k_w\} \subset \{0,1\}^\kappa$
- ... with $\mathcal{K}_i := \{{\mathsf{T_L}}^{-1}(k_i)\}$ all large (say size $\geq 2^{\kappa/2}$)
- If $\forall i$, $\mathcal{E}^{\mathcal{K}_i} : \mathcal{K}_i \times \mathcal{M} \to \mathcal{M}$ "is secure", $\mathbf{Adv}_{\mathsf{NMC}}$ is small w.r.t. $\mathsf{Tamp}^{\mathsf{T_L}\|\cdot}$
- (Anyone with access to $\mathcal{O}^{\mathcal{K}_i}$ can query $x := \mathcal{O}^{\mathcal{K}_i}(m)$, $y := \mathcal{E}_{k_i}^{-1}(x)$)
- Formalized through "PRP-with-leakage" notion

Main proof intuition

- Get a collection of reductions to RK, PRP-with-leakage
- Show that $\forall\, \mathsf{T_L}$, one reduction gives a "strong" bound

Technical definitions (1)

PRP-with-leakage

$$\mathbf{Adv}_{\mathcal{E}}^{\text{prp-leak}}(q,t) = \max_{A_{q,t}} \max_{\varphi} \left| \Pr\left[A_{q,t}^{\mathcal{O}_\varphi(0)}() = 1\right] - \Pr\left[A_{q,t}^{\mathcal{O}_\varphi(1)}() = 1\right] \right|$$

- $\mathcal{O}_\varphi(b)$ picks k, *aborts* if $\varphi(k)$ cannot be guessed w.p. $> 2^{-\kappa/2}$
- Otherwise gives $\varphi(k)$ to A, answers further queries as:
- $\mathcal{E}_k(\cdot)$ $(b = 0)$
- $\mathfrak{L}(\cdot)$ $(b = 1)$

For a "good" $\mathcal{E}$, expected $\mathbf{Adv}_{\mathcal{E}}^{\text{prp-leak}}(q,t) \approx t \cdot 2^{-\kappa/2}$

Technical definitions (2)

Fixed-RK

$$\mathbf{Adv}^{\mathsf{frk}}_{\mathcal{E}}(q,t) = \max_{A_{q,t}} \max_{\varphi} \left| \Pr\left[A_{q,t}^{\mathcal{O}_{\varphi}(0)}() = 1\right] - \Pr\left[A_{q,t}^{\mathcal{O}_{\varphi}(1)}() = 1\right] \right|$$

- $\mathcal{O}_{\varphi}(b)$ picks k, *aborts* if $\varphi(k)$ <u>can</u> be guessed w.p. $> 2^{-\kappa/2}$
- If still alive, answers further queries as:
- $\mathcal{E}^{\pm}_{k}(\cdot)$, $\mathcal{E}^{\pm}_{\varphi(k)}(\cdot)$ $(b = 0)$
- $\mathfrak{E}^{\pm}_{k}(\cdot)$, $\mathfrak{E}^{\pm}_{\varphi(k)}(\cdot)$ $(b = 1)$

For a "good" $\mathcal{E}$, expected $\mathbf{Adv}^{\mathsf{frk}}_{\mathcal{E}}(q,t) \approx t \cdot 2^{-\kappa/2}$

Part of the proof

Define Λ as "$\mathcal{O}$ does not abort in FRK (and does in PRP-leak)"; then:

$$
\begin{aligned}
&\Pr[A(\mathrm{Tamp}^{\mathsf{T}}_{\mathrm{RKNMC}}(m)) = 1 \wedge \neg\Lambda] \\
&\quad = \Pr\Big[A \circ \dot{\mathrm{D}}\mathrm{ec}^{\mathrm{Enc}_K(m)} \circ \mathsf{T} \circ \mathrm{Enc}_K(m) = 1 \wedge \neg\Lambda\Big] \\
&\quad = \Pr\Big[A \circ \dot{\mathcal{D}}^{K,\mathcal{E}_K(m)}_{\mathsf{T}_{\mathsf{L}}(K)} \circ \mathsf{T}_{\mathsf{R}} \circ \mathcal{E}_K(m) = 1 \wedge \neg\Lambda\Big] \\
&\quad = \Pr\Big[A \circ \mathcal{D}_{\mathsf{T}_{\mathsf{L}}(K)} \circ \mathsf{T}_{\mathsf{R}} \circ \mathcal{E}_K(m) = 1 \wedge \neg\Lambda\Big] \pm 2^{-\kappa/2} \\
&\quad = \Pr\Big[A \circ \mathcal{D}_{\mathsf{T}_{\mathsf{L}}(K)} \circ \mathsf{T}_{\mathsf{R}} \circ \mathfrak{E}(m) = 1 \wedge \neg\Lambda\Big] \pm \mathbf{Adv}^{\text{prp-leak}}_{\mathcal{E}}(1, 2t+1) \\
&\qquad \pm 2^{-\kappa/2}
\end{aligned}
$$

Final result

- Similar argument holds w.r.t. F-RK when Λ is true
- Λ is independent of m

$\Rightarrow$

Theorem

$\mathbf{Adv}_{\mathrm{RKNMC}}(t) \leq$
$2\max\left\{\mathbf{Adv}_{\mathcal{E}}^{\text{prp-leak}}(1, 2t+1) + 2^{-\kappa/2}, \mathbf{Adv}_{\mathcal{E}}^{\text{f-rk}}(4, 2t) + \varepsilon + 2^{-n}\right\}$

N.B.: there is a generic attack w. $\mathbf{Adv}(t) \approx t^2/2^\kappa$

Instantiation matters

Need block ciphers secure w.r.t. PRP-with-leakage and Fixed-RK
$\rightsquigarrow$ No known RK attack with ONE RK-query
$\rightsquigarrow$ No known large weak key classes

- Fixed message-length: e.g. AES-128 ($|m| = 128$, $\kappa = 64$); SHACAL-2 ($|m| = 256$, $\kappa = 256$)
- Variable message-length: VILBC, e.g. MisterMonsterBurrito + IEM

Fin