Security Proofs for PRØST

Martin M. Lauridsen

DTU Compute, Technical University of Denmark
http://martinlauridsen.info

1 Notation

Recall that PRØST-n is a permutation on $2n$ bits. In the following, we let τ denote the size of the tag. Furthermore, in the context of PRØST-APE, c denotes the capacity and r denotes the rate.

Let $\text{Perm}(2n)$ be the set of all permutations on $2n$ bits. We write $x \xleftarrow{\$} X$ to denote that x is sampled uniformly at random from X. Let $\tilde{P}_K$ denote the single-key Even-Mansour construction (SEM) using permutation P with key K.

2 Acknowledgments

We, the PRØST team, would like to thank Bart Mennink for his tireless assistance in these proofs. Without him, they would not have been possible.

3 Security Proofs

For all our security proofs, we assume PRØST to be an ideal permutation. Two of the suggested modes of operation for PRØST are block cipher-based (COPA and OTR), while the last (APE) is permutation-based. As such, the security proof for PRØST-APE is obtained via the security proof for the APE mode of operation, as this proof assumes any ideal permutation. For the block cipher-based modes, we remark that both the security proofs for the COPA and OTR modes of operation apply to *any* block cipher E, and indeed one needs only use the appropriate $\mathbf{Adv}^{\mathbf{sprp}}_E$ for the particular choice of E.

3.1 Security Definitions

In the following, we give the canonical definitions of security that allow us in Sections 3.3 through 3.5 to show the security of our proposals with respect to privacy and authenticity.

Definition 1 (Distinguisher). *We define a* distinguisher $\mathcal{D}$ *to be an algorithm which is given access to a list of oracles* $\mathcal{O}_1, \ldots, \mathcal{O}_k$*, denoted* $\mathcal{D}^{\mathcal{O}_1,\ldots,\mathcal{O}_k}$*. The oracles represent on one hand either the behavior of an instantiation of a concrete cryptographic primitive or an idealized versions of the primitive on the other hand. The goal of the distinguisher is to guess which is the case. We say (without loss of generality) that* $\mathcal{D}^{\mathcal{O}_1,\ldots,\mathcal{O}_k} = 1$ *if the distinguisher guesses that the oracles represent the former, and* $\mathcal{D}^{\mathcal{O}_1,\ldots,\mathcal{O}_k} = 0$ *otherwise.*

For the security of PRØST-COPA, we will need the notion of an *online cipher*. The notation is re-used from the security proof of COPA [ABL$^+$13].

Definition 2 (Online cipher). *A cipher* $\mathcal{E} : \mathcal{K} \times (\mathbb{F}_2^{2n})^+ \to (\mathbb{F}_2^{2n})^+$ *is said to be* online *if it has the property that*

1. *It is a permutation on every block of* $2n$ *bits and*
2. *The output blocks are identical for two different inputs with a common prefix, i.e.* $\mathcal{E}_K(X\|Y)$ *and* $\mathcal{E}_K(X\|Y')$ *are identical on the first* $|X|$ *bits for any* $X, Y, Y' \in (\mathbb{F}_2^{2n})^+$.

As such, an online cipher $\mathcal{E}_K$ *is a permutation on the blocks starting from block* i *and onwards, and is determined by the first* $i-1$ *blocks. We let* $\mathrm{OPerm}(2n)$ *denote the set of all such permutations* $\pi : (\mathbb{F}_2^{2n})^+ \to (\mathbb{F}_2^{2n})^+$. *For APE, we use a slightly different online permutation, see [ABB+14].*

Definition 3. *Let* E *be a block cipher. The* **sprp** *advantage of a distinguisher* $\mathcal{D}$ *is defined as*

$$\mathbf{Adv}_E^{\mathsf{sprp}}(\mathcal{D}) = \left| \Pr_K \left[\mathcal{D}^{E_K, E_K^{-1}} = 1 \right] - \Pr_\pi \left[\mathcal{D}^{\pi, \pi^{-1}} = 1 \right] \right|.$$

The probabilities are taken over $K \xleftarrow{\$} \mathcal{K}$, $\pi \xleftarrow{\$} \mathrm{Perm}(2n)$ *and any random choices made by* $\mathcal{D}$. *We denote by* $\mathbf{Adv}_E^{\mathsf{sprp}}(t, q)$ *the maximum advantage taken over all distinguishers* $\mathcal{D}$ *that run in time* t *and make* q *queries. Note: In the case where* E *is the Single-Key Even-Mansour construction using an ideal permutation* P, *the distinguisher has also access to the underlying permutation (in both directions) in both worlds (and thus access to four oracles). In the ideal world,* P *and* π *are independent.*

With respect to privacy, the attack model considered is chosen-plaintext attacks (IND-CPA).

Definition 4. *Let* $\mathcal{E}$ *be a block cipher-based AE scheme. The IND-CPA advantage of a distinguisher* $\mathcal{D}$ *is defined as*

$$\mathbf{Adv}_{\mathcal{E}}^{\mathsf{priv}}(\mathcal{D}) = \left| \Pr_K \left[\mathcal{D}^{\mathcal{E}_K} = 1 \right] - \Pr_{\$} \left[\mathcal{D}^{\$} = 1 \right] \right|.$$

The probabilities are taken over $K \xleftarrow{\$} \mathcal{K}$, $\$ \xleftarrow{\$} \mathrm{OPerm}(2n)$ *and any random choices made by* $\mathcal{D}$. *By* $\mathbf{Adv}_{\mathcal{E}}^{\mathsf{priv}}(t, q, \sigma, \ell)$ *we denote the maximum advantage taken over all distinguishers* $\mathcal{D}$ *that run in time* t *and make* q *queries, of length at most* ℓ *blocks, and of total length at most* σ *blocks.*

Definition 5. *Let* $\mathcal{E}$ *be a block cipher-based AE scheme. The* authenticity advantage *of a distinguisher* $\mathcal{D}$ *is defined as*

$$\mathbf{Adv}_{\mathcal{E}}^{\mathsf{auth}}(\mathcal{D}) = \left| \Pr_K \left[\mathcal{D}^{\mathcal{E}_K, \mathcal{E}_K^{-1}} = 1 \right] - \Pr_K \left[\mathcal{D}^{\mathcal{E}_K, \perp} = 1 \right] \right|,$$

where $\perp$ *in this context denotes a function that returns* $\perp$ *on every input. The probabilities are taken over* $K \xleftarrow{\$} \mathcal{K}$ *and any random choices made by* $\mathcal{D}$. *By* $\mathbf{Adv}_{\mathcal{E}}^{\mathsf{auth}}(t, q, \sigma, \ell)$ *we denote the maximum advantage taken over all distinguishers* $\mathcal{D}$ *that run in time* t *and make* q *queries, of length at most* ℓ *blocks, and of total length at most* σ *blocks. We assume that* $\mathcal{D}$ *does not make a decryption query* (A, C, T) *if it has already seen* $(C, T) = \mathcal{E}_K(A, M)$ *for some* M.

Definition 6. *Let* $\mathcal{E}$ *be a permutation-based AE scheme. The* IND-CPA advantage *of a distinguisher* $\mathcal{D}$ *is defined as*

$$\mathbf{Adv}_{\mathcal{E}}^{\mathsf{perm\text{-}priv}}(\mathcal{D}) = \left| \Pr_{K,\pi} \left[\mathcal{D}^{\mathcal{E}_K, \pi, \pi^{-1}} = 1 \right] - \Pr_\pi \left[\mathcal{D}^{\$, \pi, \pi^{-1}} = 1 \right] \right|.$$

The probabilities are taken over $K \xleftarrow{\$} \mathcal{K}$, $\pi \xleftarrow{\$} \mathrm{Perm}(2n), \$ \xleftarrow{\$} \mathrm{OPerm}(r)$, *and any random choices made by* $\mathcal{D}$. *By* $\mathbf{Adv}^{\mathsf{perm\text{-}priv}}_{\mathcal{E}}(q,m)$ *we denote the maximum advantage taken over all distinguishers* $\mathcal{D}$ *making* q *queries totaling* m *blocks.*

Definition 7. *Let* $\mathcal{E}$ *be a permutation-based AE scheme. The* authenticity advantage *of a distinguisher* $\mathcal{D}$ *is defined as*

$$\mathbf{Adv}^{\mathsf{perm\text{-}auth}}_{\mathcal{E}}(\mathcal{D}) = \left| \Pr_{K,\pi}\left[\mathcal{D}^{\mathcal{E}_K,\mathcal{E}_K^{-1},\pi,\pi^{-1}} = 1\right] - \Pr_{\pi}\left[\mathcal{D}^{\mathcal{E}_K,\perp,\pi,\pi^{-1}} = 1\right]\right|.$$

The probabilities are taken over $K \xleftarrow{\$} \mathcal{K}$, $\pi \xleftarrow{\$} \mathrm{Perm}(2n)$, *and any random choices made by* $\mathcal{D}$. *By* $\mathbf{Adv}^{\mathsf{perm\text{-}auth}}_{\mathcal{E}}(q,m)$ *we denote the maximum advantage taken over all distinguishers* $\mathcal{D}$ *making* q *queries totaling* m *blocks. We assume that* $\mathcal{D}$ *does not make a decryption query* (A,C,T) *if it has already seen* $(C,T) = \mathcal{E}_K(A,M)$ *for some* M.

3.2 Patarin's *H*-coefficient Technique

For our proof of Lemma 1, we will rely on a proof technique due to Patarin [Pat08]. We state here the result as we need it and refer to [CS14] for further details on the technique.

Let $\mathcal{D}$ be a distinguisher trying to distinguish between two systems X and Y. The interaction of $\mathcal{D}$ is captured by a transcript which is denoted τ. For $Z \in \{X,Y\}$, we let D_Z denote the probability distribution over transcripts when interacting with system Z. Let $\mathcal{T}$ be the set of all feasible transcripts which is partitioned into a set of *good* and *bad* transcripts s.t. $\mathcal{T} = \mathcal{T}_{\mathrm{good}} \cup \mathcal{T}_{\mathrm{bad}}$. Now, consider a fixed distinguisher $\mathcal{D}$ and let ε be s.t. for all $\tau \in \mathcal{T}_{\mathrm{good}}$ it holds that

$$\frac{\Pr[D_X = \tau]}{\Pr[D_Y = \tau]} \geq 1 - \varepsilon,$$

then the H-coefficient technique says that $\mathbf{Adv}(\mathcal{D}) \leq \varepsilon + \Pr[D_Y \in \mathcal{T}_{\mathrm{bad}}]$.

Lemma 1 (Security bound on SEM). *Let* P *be an ideal permutation s.t. an adversary can make at most* ρ *evaluations of* P *in time* t. *Then*

$$\mathbf{Adv}^{\mathsf{sprp}}_{\tilde{P}_K}(t,q) \leq \frac{2\rho q}{2^{2n}}.$$

Proof. Let $\mathcal{D}$ be any distinguisher which can evaluate P at most ρ times in time t. For the proof we use Patarin's H-coefficient technique. Let X denote the real world in which $\mathcal{D}$ interacts with the oracles $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3, \mathcal{O}_4 = \tilde{P}_K, \tilde{P}_K^{-1}, P, P^{-1}$ and let Y denote the ideal world where $\mathcal{D}$ interacts with $\mathcal{O}_1, \mathcal{O}_2, \mathcal{O}_3, \mathcal{O}_4 = \pi, \pi^{-1}, P, P^{-1}$ for $K \xleftarrow{\$} \mathcal{K}$ and $\pi \xleftarrow{\$} \mathrm{Perm}(2n)$.

The result of the interaction by $\mathcal{D}$ using q construction queries we denote $\tau_E = \{(s_i,t_t)\}_{i=1}^{q}$. Similarly, the result of the interaction using ρ queries to P we denote $\tau_P = \{(x_i,y_i)\}_{i=1}^{\rho}$. To ease the analysis, the key K is disclosed at the end of the experiment (in the ideal world, a dummy key K is disclosed). We define a transcript as a tuple (τ_E, τ_P, K) and a bad transcript is such a tuple where it holds that

$$K \in \{s \oplus x, y \oplus t \mid (s,t) \in \tau_E \wedge (x,y) \in \tau_P\}.$$

Bounding $\Pr[D_Y \in \mathcal{T}_{\text{bad}}]$. There are ρ pairs $(x, y) \in \tau_P$ and for each of them we consider each pair $(s, t) \in \tau_E$. This means there are at most $q \cdot \rho$ values for $s \oplus x$, any of which equals K with probability 2^{-2n}. A similar argument applies to the probability of there being a pair (x, y) and (s, t) s.t. $t \oplus y = K$. As such, we find $\Pr[D_Y \in \mathcal{T}_{\text{bad}}] \leq \frac{2q\rho}{2^{2n}}$.

Bounding $\Pr[D_X = \tau] / \Pr[D_Y = \tau]$ *for* $\tau \in \mathcal{T}_{\text{good}}$. Consider some $\tau \in \mathcal{T}_{\text{good}}$. Let Ω_X and Ω_y denote all possible oracles in the real world and ideal world, respectively. Correspondingly, let $\mathsf{comp}_X(\tau)$ (respectively $\mathsf{comp}_Y(\tau)$) denote transcripts in Ω_X (respectively Ω_Y) which are compatible with τ.

Since the key space has size 2^{2n}, and there are $2^{2n}!$ permutations on $2n$ bits, we have that $|\Omega_X| = 2^{2n} \cdot 2^{2n}!$ and $|\Omega_Y| = 2^{2n} \cdot (2^{2n}!)^2$. Now, $\tau \in \mathcal{T}_{\text{good}}$ implies that any tuple in τ defines a unique input/output pair to P. As $\tau_E \cup \tau_P$ consists of $q + \rho$ tuples, the number of compatible $2n$-bit permutations in the real world is $|\mathsf{comp}_X| = (2^{2n} - q - \rho)!$. Correspondingly, in the ideal world, the number of permutations compliant with P is $(2^{2n} - \rho)!$ while the number of permutations compliant with the construction queries is $(2^{2n} - q)!$. As such, $|\mathsf{comp}_Y| = (2^{2n} - q)!(2^{2n} - \rho)! \leq (2^{2n} - q - \rho)!2^{2n}!$.

By definition, we find that

$$\begin{aligned}
\Pr[D_X = \tau] &= \frac{(2^{2n} - q - \rho)!}{2^{2n} \cdot 2^{2n}!} \\
&= \frac{(2^{2n} - q - \rho)!2^{2n}!}{2^{2n} \cdot (2^{2n}!)^2} \\
&\geq \frac{|\mathsf{comp}_Y|}{|\Omega_Y|} = \Pr[D_Y = \tau].
\end{aligned}$$

As such, we see $\Pr[D_X = \tau] \geq \Pr[D_Y = \tau]$, so $\varepsilon \leq 0$ and we have $\mathbf{Adv}^{\mathsf{sprp}}_{\tilde{P}_K}(t, q) \leq \frac{2q\rho}{2^{2n}}$. □

3.3 PRØST-COPA

Theorem 1 (Privacy for PRØST-COPA). *Assume that* PRØST *is an ideal permutation and that an adversary can make at most ρ evaluations of the* PRØST *permutation in time t', where $t' \approx t$. Then*

$$\mathbf{Adv}^{\mathsf{priv}}_{\text{PRØST-COPA}}(t, q, \sigma, \ell) \leq \frac{39(\sigma + q)^2}{2^{2n}} + \frac{8\rho(\sigma + q)}{2^{2n}} + \frac{(\ell + 2)(q - 1)^2}{2^{2n}}.$$

Proof. The proof follows from combining the proof for privacy of COPA [ABL+13, Theorem 2] with Lemma 1. □

Theorem 2 (Authenticity for PRØST-COPA). *Assume that* PRØST *is an ideal permutation and that an adversary can make at most ρ evaluations of the* PRØST *permutation in time t', where $t' \approx t$. Then*

$$\mathbf{Adv}^{\mathsf{auth}}_{\text{PRØST-COPA}}(t, q, \sigma, \ell) \leq \frac{39(\sigma + q)^2}{2^{2n}} + \frac{8\rho(\sigma + q)}{2^{2n}} + \frac{(\ell + 2)(q - 1)^2}{2^{2n}} + \frac{2q}{2^{\tau}}.$$

Proof. The proof follows from combining the proof for authenticity of COPA [ABL+13, Theorem 3] with Lemma 1. □

3.4 PRØST-OTR

The security proof for OTR by Minematsu [Min14] is in the ideal model, i.e. it assumes the underlying block cipher to be an ideal primitive. In this section, we give a proof in the standard model. In particular, we model the encrypting of a block in OTR, with its masking, as an XE construction (see [Rog04]).

Theorem 3 (Privacy for PRØST-OTR). *Assume that* PRØST *is an ideal permutation and that an adversary can make at most ρ evaluations of the* PRØST *permutation in time t', where $t' \approx t$. Then*

$$\mathbf{Adv}^{\mathsf{priv}}_{\text{PRØST-OTR}}(t, q, \sigma, \ell) \leq \frac{6(\sigma+q)^2}{2^{2n}} + \frac{4\rho(\sigma+q)}{2^{2n}}.$$

Proof. The proof follows from combining three parts: The proof for privacy of OTR [Min14, Theorem 1] in the ideal model; the fact that the modeling of OTR using XE-blocks admits the term $\mathbf{Adv}^{\mathsf{sprp}}_{E}(t', 2\sigma)$ in the standard model (where $t' \approx t$); Lemma 1 which gives the term $\mathbf{Adv}^{\mathsf{sprp}}_{E}(t', 2\sigma)$. □

Theorem 4 (Authenticity for PRØST-OTR). *Assume that* PRØST *is an ideal permutation and that an adversary can make at most ρ evaluations of the* PRØST *permutation in time t', where $t' \approx t$. Then*

$$\mathbf{Adv}^{\mathsf{auth}}_{\text{PRØST-OTR}}(t, q, \sigma, \ell) \leq \frac{6(\sigma+q)^2}{2^{2n}} + \frac{4\rho(\sigma+q)}{2^{2n}} + \frac{q}{2^{\tau}}.$$

Proof. The proof follows from combining the proof for authenticity of OTR [Min14, Theorem 2] in the ideal model; the fact that the modeling of OTR using XE-blocks admits the term $\mathbf{Adv}^{\mathsf{sprp}}_{E}(t', 2\sigma)$ in the standard model (where $t' \approx t$); Lemma 1 which gives the term $\mathbf{Adv}^{\mathsf{sprp}}_{E}(t', 2\sigma)$. □

3.5 PRØST-APE

In this section we present security bounds for PRØST-APE. The proofs of security for the APE mode of operation assume an ideal permutation, thus under this assumption, the security bounds for the APE construction carry directly over to PRØST-APE. Note that the security bounds do not depend on the time t used by the distinguisher. Indeed, the bound holds for the strongest type of distinguishers, whose time complexity is unbounded; only the number of queries q and their total length m, made by the distinguisher matters.

Theorem 5 (Privacy for PRØST-APE). *Assume that* PRØST *is an ideal permutation. Then*

$$\mathbf{Adv}^{\mathsf{perm\text{-}priv}}_{\text{PRØST-APE}}(q, m) \leq \frac{m^2}{2^{2n}} + \frac{m(m+1)}{2^{c}}.$$

Proof. The proof is given in [ABB+14, Theorem 1].

Theorem 6 (Authenticity for PRØST-APE). *Assume that* PRØST *is an ideal permutation. Then*

$$\mathbf{Adv}^{\mathsf{perm\text{-}auth}}_{\text{PRØST-APE}}(q, m) \leq \frac{m^2}{2^{2n}} + \frac{2m(m+1)}{2^{c}}.$$

Proof. The proof is given in [ABB+14, Theorem 2].

References

[ABB+14] Elena Andreeva, Begül Bilgin, Andrey Bogdanov, Atul Luykx, Bart Mennink, Nicky Mouha, and Kan Yasuda. APE: Authenticated Permutation-Based Encryption for Lightweight Cryptography. In Carlos Cid and Christian Rechberger, editors, *Fast Software Encryption*, LNCS. Springer, 2014. `http://eprint.iacr.org/2013/791/`.

[ABL+13] Elena Andreeva, Andrey Bogdanov, Atul Luykx, Bart Mennink, Elmar Tischhauser, and Kan Yasuda. Parallelizable and Authenticated Online Ciphers. In Kazue Sako and Palash Sarkar, editors, *Advances in Cryptology – ASIACRYPT 2013*, volume 8269 of *LNCS*, pages 424–443. Springer, 2013. `eprint.iacr.org/2013/790/`.

[CS14] Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *Advances in Cryptology - EUROCRYPT 2014 - 33rd Annual International Conference on the Theory and Applications of Cryptographic Techniques, Copenhagen, Denmark, May 11-15, 2014. Proceedings*, volume 8441 of *Lecture Notes in Computer Science*, pages 327–350. Springer, 2014.

[Min14] Kazuhiko Minematsu. Parallelizable Rate-1 Authenticated Encryption from Pseudorandom Functions. In *EUROCRYPT*, pages 275–292, 2014.

[Pat08] Jacques Patarin. The "coefficients h" technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography, 15th International Workshop, SAC 2008, Sackville, New Brunswick, Canada, August 14-15, Revised Selected Papers*, volume 5381 of *Lecture Notes in Computer Science*, pages 328–345. Springer, 2008.

[Rog04] Phillip Rogaway. Efficient Instantiations of Tweakable Blockciphers and Refinements to Modes OCB and PMAC. In Pil Joong Lee, editor, *Advances in Cryptology - ASIACRYPT 2004, 10th International Conference on the Theory and Application of Cryptology and Information Security, Jeju Island, Korea, December 5-9, 2004, Proceedings*, volume 3329 of *Lecture Notes in Computer Science*, pages 16–31. Springer, 2004.