Invertible Universal Hashing and the TET Encryption Mode

Shai Halevi

IBM T.J. Watson Research Center,
Hawthorne, NY 10532, USA
shaih@alum.mit.edu

Abstract. This work describes a mode of operation, TET, that turns a regular block cipher into a length-preserving enciphering scheme for messages of (almost) arbitrary length. When using an n-bit block cipher, the resulting scheme can handle input of any bit-length between n and 2^n and associated data of arbitrary length.
The mode TET is a concrete instantiation of the generic mode of operation that was proposed by Naor and Reingold, extended to handle tweaks and inputs of arbitrary bit length. The main technical tool is a construction of invertible "universal hashing" on wide blocks, which is as efficient to compute and invert as polynomial-evaluation hash.

1 Introductions

Adding secrecy protection to existing (legacy) protocols and applications raises some unique problems. One of these problems is that existing protocols sometimes require that the encryption be "transparent", and in particular preclude length-expansion. One example is encryption of storage data "at the sector level", where both the higher-level operating system and the lower-level disk expect the data to be stored in blocks of 512 bytes, and so any encryption method would have to accept 512-byte plaintext and produce 512-byte ciphertext.

Clearly, insisting on a length-preserving (and hence deterministic) transformation has many drawbacks. Indeed, even the weakest common notion of security for "general purpose encryption" (i.e., semantic security [GM84]) cannot be achieved by deterministic encryption. Still, there may be cases where length-preservation is a hard requirement (due to technical, economical or even political constrains), and in such cases one may want to use some encryption scheme that gives better protection than no encryption at all. The strongest notion of security for a length-preserving transformation is "strong pseudo-random permutation" (SPRP) as defined by Luby and Rackoff [LR88], and its extension to "tweakable SPRP" by Liskov et al. [LRW02]. A "tweak" is an additional input to the enciphering and deciphering procedures that need not be kept secret. This report uses the terms "tweak" and "associated data" pretty much interchangeably, except that "associated data" hints that it can be of arbitrary length.

Motivated by the application to "sector level encryption", many modes of operation that implement tweakable SPRP on wide blocks were described in the literature in the last few years. Currently there are at least eight such proposals, following three different approaches: The "encrypt-mix-encrypt" approach is used for CMC, EME and EME* [HR03,HR04,Hal04], the "hash-ECB-hash" (due to Naor and Reingold [NR97]) is used in PEP [CS06b], and the "hash-CTR-hash" approach is used by XCB [FM04], HCTR [WFW05] and HCH [CS06a] (and some variation of the last approach is used in ABL4 [MV04]). Among these proposals, the "encrypt-mix-encrypt" modes are the most efficient (at least in software), the "hash-CTR-hash" modes are close behind, and PEP and ABL4 are considerably less efficient (more on efficiency in Section 3.5).

This work presents a ninth mode called TET (for linear-Transformation; ECB; linear-Transformation). TET belongs to the "hash-ECB-hash" family, but in terms of efficiency it is similar to the modes of the "hash-CTR-hash" family, thus complementing the current lineup of modes. We also mention that TET may have some practical advantages with respect to intellectual-property concerns, see further discussion in the appendix.

The main technical contribution of this work is a construction of an efficient invertible universal hashing for wide blocks, which is needed in the "hash-ECB-hash" approach. Given the wide range of applications of universal hashing in general, this invertible universal hashing may find applications beyond the TET mode itself. Another small contribution is a slight modification of the OMAC construction for pseudorandom function due to Iwata and Korasawa [IK03]. (In TET we use that pseudorandom function to handle the message-length and the tweak). This construction too can find other applications.

The Naor-Reingold construction and TET*.* Recall that the Naor-Reingold construction from [NR97] involves a layer of ECB encryption, sandwiched between two layers of universal hashing, as described in Figure 1. The universal hashing layers must be invertible (since they need to be inverted upon decryption), and their job is to ensure that different queries of the attacker will almost never result in "collisions" at the ECB layer. Namely, for any two plaintext (or ciphertext) vectors $\boldsymbol{p} = \langle p_1, \ldots, p_m \rangle$, $\boldsymbol{q} = \langle q_1, \ldots, q_m \rangle$ and two indexes i, j (such that $(\boldsymbol{p}, i) \neq (\boldsymbol{q}, j)$) it should hold with high probability (over the hashing key) that the i'th block of hashing $\boldsymbol{p}$ is different from the j'th block of hashing $\boldsymbol{q}$.

As mentioned above, the main contribution of this note is a construction of an invertible universal hashing on wide blocks, which is as efficient to compute and invert as polynomial-evaluation hash. In a nutshell, the hashing family works on vectors in $\mathrm{GF}(2^n)^m$, and it is keyed by a single random element $\tau \in_R \mathrm{GF}(2^n)$, which defines the following $m \times m$ matrix:

$$A_\tau \stackrel{\text{def}}{=} \begin{pmatrix} \tau & \tau^2 & & \tau^m \\ \tau & \tau^2 & & \tau^m \\ & & \ddots & \\ \tau & \tau^2 & & \tau^m \end{pmatrix}$$

Set $\sigma \stackrel{\text{def}}{=} 1 + \tau + \tau^2 + \ldots + \tau^m$, we observe that if $\sigma \neq 0$ then the matrix $M_\tau = A_\tau + I$ is invertible and its inverse is $M_\tau^{-1} = I - (A_\tau / \sigma)$. Thus multiplying

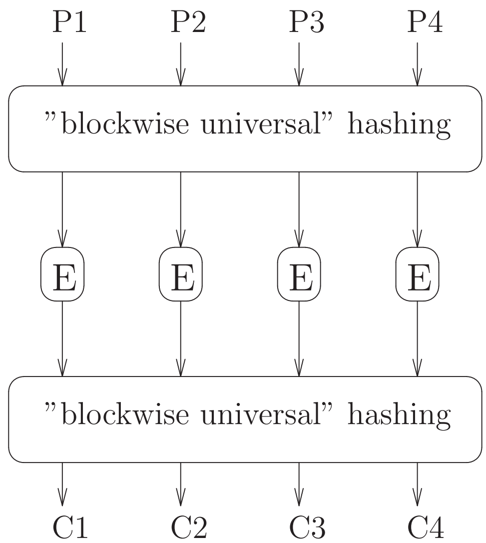

Fig. 1. The Naor-Reingold generic mode: the universal hashing must be invertible, and its job is to prevent collisions in the ECB layer.

by M_τ for a random τ (subject to $\sigma \neq 0$) is an invertible universal hashing, and computing and inverting this hash function is about as efficient as computing polynomial evaluation.

The starting point of this work is an implementation of the generic Naor-Reingold mode of operation using the above for the universal hashing layers. We then extend that mode to handle associated data and input of arbitrary length, thus getting the TET mode. Specifically, TET takes a standard cipher with n-bit blocks and turns it into a tweakable enciphering scheme with message space $\mathcal{M} = \{0,1\}^{n..2^n-1}$ (i.e., any string of at least n and at most $2^n - 1$ bits) and tweak space $\mathcal{T} = \{0,1\}^*$. The key for TET consists of two keys of the underlying cipher (roughly one to process the tweak and another to process the data). As we mentioned, TET offers similar performance characteristics to XCB, HCTR and HCH (making it significantly more efficient than PEP and ABL4, and almost as efficient as CMC, and EME/EME*).

A word on notations. Below we use $\oplus$ to denote bit-wise exclusive or (which is the same as addition in $GF(2^n)$), and we use $+/-$ to denote addition/subtraction in other fields or domains (e.g., integer addition). The sum operator $\sum$ is always used to denote finite-field addition.

Organization. Some standard definitions are recalled in Appendix A (which is taken almost verbatim from [HR04,Hal04]). Section 2 describes the hashing scheme that underlies TET, Section 3 describes the TET mode itself, and

Section 4 examines the security for this mode. In Appendix B we briefly discuss intellectual-property issues.

Acknowledgments. I would like to thank the participants of the IEEE SISWG working group for motivating me to write this note. I also thank Doug Whiting and Brian Gladman for some discussions about this mode, and the anonymous CRYPTO reviewers for their comments.

2 The underlying hashing scheme

The universality property that is needed for the Naor-Reingold mode of operation is defined next.

Definition 1. *Let $H : \mathcal{K} \times \mathcal{D} \rightarrow \mathcal{R}^m$ be a hashing family from some domain $\mathcal{D}$ to m-vectors over the range $\mathcal{R}$, with keys chosen uniformly from $\mathcal{K}$. We denote by $H_k(x)$ the output of H (which is an m-vector over $\mathcal{R}$) on key $k \in \mathcal{K}$ and input $x \in \mathcal{D}$. We also denote by $H_k(x)_i$ the i'th element of that output vector.*

For a real number $\epsilon \in (0, 1)$, we say that $\mathcal{H}$ is "ϵ-blockwise-universal" if for every $x, x' \in \mathcal{D}$ and integers $i, i' \leq m$ such that $(x, i) \neq (x', i')$, it holds that $\Pr_k[H_k(x)_i = H_k(x')_{i'}] \leq \epsilon$, where the probability is taken over the uniform choice of $k \in \mathcal{K}$.

We say that $\mathcal{H}$ is "ϵ-xor-blockwise-universal" if in addition for all fixed $\Delta \in \mathrm{GF}(2^n)$ it holds that $\Pr_k[H_k(x)_i \oplus H_k(x')_{i'} = \Delta] \leq \epsilon$.

It was proven in [NR99] that the construction from Figure 1 is a strong PRP on wide blocks provided that the hashing layers are blockwise universal and invertible, and the underlying cipher E is a strong PRP on narrow blocks.

2.1 BPE: A blockwise universal hashing scheme

To get an invertible blockwise universal hash function, Naor and Reingold proposed in [NR97] to use an unbalanced Feistel network with standard universal hashing. For example, use polynomial-evaluation hash function applied to the first $m-1$ blocks, xor the result to the last block, and then derive $m-1$ "pairwise independent" values from the last block and xor them back to the first $m-1$ blocks. This solution, however, is somewhat unsatisfying in that it entails inherent asymmetry (which is likely to raise problems with implementations). Below we propose a somewhat more elegant blockwise universal hashing based on a simple algebraic trick.

Although for TET we only need a hashing scheme over the field $GF(2^n)$, we describe here the scheme over an arbitrary field. Let $\mathcal{F}$ be a field (with more than $m+2$ elements) and consider an $m \times m$ matrix over $\mathcal{F}$, $M_\tau \stackrel{\text{def}}{=} A_\tau + I$ for an element $\tau \in \mathcal{F}$, where

$$A_\tau \stackrel{\text{def}}{=} \begin{pmatrix} \tau & \tau^2 & & \tau^m \\ \tau & \tau^2 & & \tau^m \\ & & \ddots & \\ \tau & \tau^2 & & \tau^m \end{pmatrix} \tag{1}$$

It is easy to check that the determinant of M_τ is $\sigma \stackrel{\text{def}}{=} \sum_{i=0}^{m} \tau^i$, and so M_τ is invertible if and only if $\sigma \neq 0$. We observe that when it is invertible, the structure of M_τ^{-1} is very similar to the structure of M_τ itself.

Observation 1 *Let $\mathcal{F}$ be a field and let $\tau \in \mathcal{F}$ be such that $\sigma \stackrel{\text{def}}{=} \sum_{i=0}^{m} \tau^i \neq 0$, let A_τ be an $m \times m$ matrix with $A_{i,j} = \tau^j$, and let $M_\tau \stackrel{\text{def}}{=} A_\tau + I$. Then $M_\tau^{-1} = I - (A_\tau/\sigma)$.*

Proof. We first note that $A_\tau^2 = A_\tau(\sigma - 1)$, since for all i, j we have

$$(A_\tau^2)_{i,j} \;=\; \sum_{k=1}^{m} \tau^{k+j} \;=\; \tau^j \left(\sum_{k=1}^{m} \tau^k \right) \;=\; \tau^j \left(\sum_{k=0}^{m} \tau^k - 1 \right) \;=\; (A_\tau)_{i,j} \cdot (\sigma - 1)$$

Therefore, assuming $\sigma \neq 0$ we get

$$(A_\tau + I) \cdot (I - \frac{A_\tau}{\sigma}) \;=\; A_\tau + I - \frac{A_\tau^2}{\sigma} - \frac{A_\tau}{\sigma} \;=\; I + \frac{A_\tau \sigma - A_\tau(\sigma - 1) - A_\tau}{\sigma} \;=\; I$$

It follows that computing $\mathbf{y} = M_\tau \mathbf{x}$ and $\mathbf{x} = M_\tau^{-1}\mathbf{y}$ can be done as efficiently as computing polynomial-evaluation hash. Namely, to compute $\mathbf{y} = M_\tau \mathbf{x}$ we first compute $s = \sum_{i=1}^{m} x_i \tau^i$ and set $y_i = x_i + s$, and to invert $\mathbf{x} = M_\tau^{-1}\mathbf{y}$ we re-compute s as $s = \sum_{i=1}^{m} y_i(\tau^i/\sigma)$ and set $x_i = y_i - s$. Moreover, since τ and σ depend only the hashing key, one can speed up the multiplication by τ and τ/σ by pre-computing some tables (cf. [Sho96]).

The blockwise-universal family BPE. Given the observation from above, we define the hashing family BPE (for Blockwise Polynomial-Evaluation) and its inverse BPE^{-1} as follows: Let $\mathcal{F}$ be a finite field with $m + 3$ or more elements.

Input: An m-vector of elements from $\mathcal{F}$, $\mathbf{x} = \langle x_1, \ldots, x_m \rangle \in \mathcal{F}^m$.
Keys: Two elements $\tau, \beta \in \mathcal{F}$, such that $\sum_{i=0}^{m} \tau^m \neq 0$.

Output: Let α be some fixed primitive element of $\mathcal{F}$, and denote by $\mathbf{b} \stackrel{\text{def}}{=} \langle \beta, \alpha\beta, \ldots, \alpha^{m-1}\beta \rangle$ the m-vector over $\mathcal{F}$ whose i'th entry is $\alpha^{i-1}\beta$. The two hash functions $\mathrm{BPE}_{\tau,\beta}(\mathbf{x})$ and $\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x})$ are defined as

$$\mathrm{BPE}_{\tau,\beta}(\mathbf{x}) \stackrel{\text{def}}{=} M_\tau \mathbf{x} + \mathbf{b} \text{ and } \mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}) \stackrel{\text{def}}{=} M_\tau^{-1}(\mathbf{x} - \mathbf{b}) \tag{2}$$

By construction if follows that $\mathrm{BPE}^{-1}_{\tau,\beta}(\mathrm{BPE}_{\tau,\beta}(\mathbf{x})) = \mathbf{x}$ for all $\mathbf{x}$ and all τ, β (provided that $\sum_{i=0}^{m} \tau^m \neq 0$). We now prove that these two families (BPE and its inverse) are indeed "blockwise universal".

Claim. Fix a finite field $\mathcal{F}$ and an integer $m \leq |\mathcal{F}| - 3$, and also fix $\mathbf{x}, \mathbf{x}' \in \mathcal{F}^m$ and indexes $i, i' \leq m$ such that $(x, i) \neq (x', i')$, and any $\delta \in \mathcal{F}$.

(i) If $i \neq i'$ then $\Pr_{\tau,\beta}\left[[\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_{i'} - [\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_{i'} = \delta\right] = 1/|\mathcal{F}|$ and similarly $\Pr_{\tau,\beta}\left[[\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_{i'} - [\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_{i'} = \delta\right] = 1/|\mathcal{F}|$.

(ii) If $i = i'$ and $\mathbf{x} \neq \mathbf{x}'$ then both $\Pr_{\tau,\beta}\left[[\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_{i'} - [\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_{i'} = \delta\right]$ and $\Pr_{\tau,\beta}\left[[\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_{i'} - [\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_{i'} = \delta\right]$ are bounded by $\frac{m}{|\mathcal{F}|-g}$, where $g = GCD(m+1, |\mathcal{F}|-1)$ if the characteristic of the field $\mathcal{F}$ divides $m+1$, and $g = GCD(m+1, |\mathcal{F}|-1) - 1$ otherwise.

Proof. **Case (i), $i \neq i'$.** In this case we have $[\mathrm{BPE}_{\tau,\beta}(\mathbf{x})]_i - [\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_{i'} = (\alpha^{i-1} - \alpha^{i'-1})\beta + ((M_\tau \mathbf{x})_i - (M_\tau \mathbf{x}')_{i'})$ which is equal to any fixed δ with probability exactly $1/|\mathcal{F}|$ over the choice of $\beta \in_R \mathcal{F}$ (since α is primitive and so $\alpha^{i-1} \neq \alpha^{i'-1}$). Similarly

$$
\begin{aligned}
[\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x})]_i - [\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_{i'} \;&=\; \left((I - \frac{A_\tau}{\sigma})(\mathbf{x}-\mathbf{b})\right)_i - \left((I - \frac{A_\tau}{\sigma})(\mathbf{x}'-\mathbf{b})\right)_{i'} \\
=\; &((\frac{A_\tau}{\sigma}\mathbf{b})_i - \mathbf{b}_i) - ((\frac{A_\tau}{\sigma}\mathbf{b})_{i'} - \mathbf{b}_{i'}) + \left((I - \frac{A_\tau}{\sigma})\mathbf{x}\right)_i - \left((I - \frac{A_\tau}{\sigma})\mathbf{x}'\right)_{i'} \\
=\; &(\alpha^{i'-1} - \alpha^{i-1})\beta + \left((I - \frac{A_\tau}{\sigma})\mathbf{x}\right)_i - \left((I - \frac{A_\tau}{\sigma})\mathbf{x}'\right)_{i'}
\end{aligned}
$$

where the last equality follows since $(A_\tau \mathbf{b})_i = (A_\tau \mathbf{b})_{i'}$ (because all the rows of A_τ are the same). Again, this sum equals δ with probability exactly 2^{-n}.

Case (ii), $i = i'$ and $\mathbf{x} \neq \mathbf{x}'$. In this case we have $[\mathrm{BPE}_{\tau,\beta}(\mathbf{x})]_i - [\mathrm{BPE}_{\tau,\beta}(\mathbf{x}')]_i - \delta = (x_i - x'_i - \delta) + \sum_{j=1}^{m}(x_j - x'_j)\tau^j$, which is zero only when τ is a root of this specific non-zero degree-m polynomial. Similarly for $\mathrm{BPE}^{-1}_{\tau,\beta}$ we have

$$
\begin{aligned}
[\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x})]_i - [\mathrm{BPE}^{-1}_{\tau,\beta}(\mathbf{x}')]_i - \delta \;&=\; \left((I - \frac{A_\tau}{\sigma})(\mathbf{x}-\mathbf{b})\right)_i - \left((I - \frac{A_\tau}{\sigma})(\mathbf{x}'-\mathbf{b})\right)_i - \delta \\
=\; &\left((I - \frac{A_\tau}{\sigma})\mathbf{x}\right)_i - \left((I - \frac{A_\tau}{\sigma})\mathbf{x}'\right)_i - \delta \;=\; (x_i - x'_i - \delta) + \sum_{j=1}^{m}\frac{\tau^j}{\sigma}(x_j - x'_j) \\
\overset{*}{=}\; &\frac{1}{\sigma}\left((x_i - x'_i - \delta)(\sum_{j=0}^{m}\tau^j) + \sum_{j=1}^{m}\tau^j(x_j - x'_j)\right) \\
=\; &\frac{1}{\sigma}\left((x_i - x'_i - \delta) + \sum_{j=1}^{m}\tau^j((x_j - x'_j) + (x_i - x'_i - \delta))\right)
\end{aligned}
$$

where the equality $\overset{*}{=}$ holds since $\sigma = \sum_{i=0}^{m}\tau^j$. The last expression is zero when τ is a root of the parenthesized polynomial. That polynomial is non-zero since (a) if $x_i - x'_i \neq \delta$ then it has non-zero constant term, and (b) if $x_i - x'_i = \delta$ then there is some index j such that $x_j \neq x'_j$, and thus the coefficient $((x_j - x'_j) + (x_i - x'_i - \delta))$ of τ^j is non-zero.

We conclude that for both $\mathrm{BPE}_{\tau,\beta}$ and $\mathrm{BPE}^{-1}_{\tau,\beta}$, a collision in this case implies that τ must be a root of some fixed non-zero degree-m polynomial. Such polynomials have at most m roots, and τ is chosen at random in $\mathrm{GF}(2^n)$ subject to the constraint that $\sigma \neq 0$. Since σ itself is a non-zero degree-m polynomial, then there are at least $2^n - m$ elements $\tau \in \mathrm{GF}(2^n)$ for which $\sigma \neq 0$, and so the collision probability is at most $m/(2^n - m)$.

Moreover, for most values of m we can actually show that there are fewer than m values of τ for which $\sigma = 0$. Specifically, we note that $\sigma = (\tau^{m+1} - 1)/(\tau - 1)$, so $\sigma = 0$ implies that also $\tau^{m+1} - 1 = 0$, which means that τ is an $m+1$'st root of unity in $\mathcal{F}$. We know that the number of $m+1$'st roots of unity in $\mathcal{F}$ is exactly $GCD(m+1, |\mathcal{F}| - 1)$, and one of them is the trivial root $\tau = 1$. The trivial root $\tau = 1$ is also a root of σ if and only if the characteristic of $\mathcal{F}$ divides $m+1$ (since there are $m+1$ terms in the sum that defines σ), and all the other $m+1$'st roots of unity are also root of σ. Hence τ is chosen at random from a set of size $|\mathcal{F}| - g$, where $g = GCD(m+1, |\mathcal{F}| - 1)$ if the characteristic of $\mathcal{F}$ divides $m+1$ and $g = GCD(m+1, 2^n - 1) - 1$ otherwise.

A variant of BPE. It is easy to see that the same claim can be proven also for the variant of BPE that subtracts the vector $\mathbf{b}$ before multiplying by M_τ, namely if we define

$$\widetilde{\mathrm{BPE}}_{\tau,\beta}(\mathbf{x}) \stackrel{\mathrm{def}}{=} M_\tau(\mathbf{x} - \mathbf{b}) \text{ and } \widetilde{\mathrm{BPE}}^{-1}_{\tau,\beta}(\mathbf{x}) \stackrel{\mathrm{def}}{=} M_\tau^{-1}\mathbf{x} + \mathbf{b} \tag{3}$$

then also the hash families $\widetilde{\mathrm{BPE}}$ and $\widetilde{\mathrm{BPE}}^{-1}$ are ϵ-blockwise universal for the same ϵ.

Variable input length. Claim 2.1 refers only to the fixed-input length scenario, where $\mathrm{BPE}_{\tau,\beta}$ is applied always to inputs of the same length. Similar arguments can be used to show universality of BPE, BPE^{-1}, $\widetilde{\mathrm{BPE}}$, and $\widetilde{\mathrm{BPE}}$ also in the variable-input-length scenario, where the same τ and β are used for all the different input lengths.

Claim 2.1 refers only to the fixed-input length scenario, where $\mathrm{BPE}_{\tau,\beta}$ is applied always to inputs of the same length. Similar arguments can show that the four variations BPE, BPE^{-1}, $\widetilde{\mathrm{BPE}}$, and $\widetilde{\mathrm{BPE}}$ are also ϵ-blockwise universal in the variable-input-length scenario, where the same τ and β are used for all the different input lengths.

One complication is that in the variable-input-length scenario, the element $\tau \in \mathcal{F}$ must be chosen such that *for all* m it holds that $1 + \tau + \ldots + \tau^m \neq 0$. This can be achieved by choosing τ as a primitive element in $\mathcal{F}$, which means that it is not an $m+1$'ts root of unity for any $m < |\mathcal{F}| - 2$, and therefore also not a root of $1 + \tau + \ldots + \tau^m$. As the number of primitive elements in $\mathcal{F}$ is $\phi(|\mathcal{F}| - 1)$ (where ϕ is Euler's totient function), it follows that in this case we choose τ from a set of size exactly $\phi(|\mathcal{F}| - 1)$. Hence the collision probability for any $\mathbf{x}, \mathbf{x}'$ is bounded by $\epsilon = m/\phi(|\mathcal{F}| - 1)$ where m is the length of the longer of $\mathbf{x}, \mathbf{x}'$.

3 The TET mode of operation

The BPE hashing scheme immediately implies a mode of operation for implementing a fixed-input-length, non-tweakable enciphering scheme for block-sizes that are a multiple of n bits: namely the Naor-Reingold construction from [NR97] with BPE for the hashing layers (over the field $\mathrm{GF}(2^n)$, where n is the block size of the underlying cipher). In this section I describe how to extend this construction to get a tweakable scheme that supports arbitrary input lengths (and remains secure also when using the same key for different input lengths).

3.1 Tweaks and variable input length

Incorporating a tweak into the basic mode turns out to be almost trivial: Instead of having the element β be part of the key, we derive it from the tweak using the underlying cipher. For example, if we are content with n-bit tweaks then we can just set $\beta \leftarrow E_K(T)$ where k is the cipher key and T is the tweak. Intuitively, this is enough since the multiples of β will be used to mask the input values before they arrive at the ECB layer, so using different pseudo-random values of β for different tweak values means that the ECB layer will be applied on different blocks.

To handle longer tweaks we can replace the simple application of the underlying cipher E with a variable-input-length cipher-based pseudo-random function (e.g., CBC-MAC, PMAC, etc.), using a key which is independent of the cipher key that is used for the ECB layer. In Section 3.3 I describe a particular CBC-MAC-like implementation that suits our needs.

The same fix can be applied also to handle variable input length: namely we derive β from both the tweak and the input length. If we are content with input length of no more than 2^ℓ and tweaks of size $n - \ell$ bits, then we can use $\beta \leftarrow E_K(L, T)$ where T is the tweak value and L is the input length, or else we can use $\beta \leftarrow \mathrm{PRF}_K(L, T)$ for some variable-input-length pseudo-random function. As noted above, using the same hashing key for different input lengths implies that the element τ must satisfy $\sigma_m = 1 \oplus \tau \oplus \ldots \oplus \tau^m \neq 0$ for every possible input length m, and this can be ensured by choosing τ as a random primitive element in $\mathrm{GF}(2^n)$.

3.2 Partial blocks

It appears harder to extend the mode to handle inputs whose length is not a multiple of n bits. Ideally, we would have liked an elegant way of extending BPE to handle such lengths, and then handle partial blocks in the ECB layer using ciphertext-stealing (cf. [MM82, Fig.2-23]). Unfortunately, I do not know how to extend BPE to handle input length that is not a multiple of n bits while maintaining invertability (except going back to the unbalanced Feistel idea).

Instead, I borrowed a technique that was used also in EME*: When processing an input whose length is not a multiple of n bits, one of the block cipher applications in the ECB layer is replaced with two consecutive applications of the cipher, and the middle value (between the two calls to the underlying cipher) is xor-ed to the partial block. (In addition, the partial block is added to the polynomial-evaluation, so that its value effects all the other blocks.)

In more details, let $\mathbf{x} = \langle x_1, \ldots, x_m \rangle$ be all the full input blocks and let x_{m+1} be a partial block, $\ell = |x_{m+1}|$, $0 < \ell < n$. Instead of just computing $\mathbf{y} = \mathrm{BPE}(\mathbf{x})$, we set the i'th full block to $y_i \leftarrow \mathrm{BPE}(\mathbf{x})_i \oplus (x_{m+1}10..0)$, while leaving x_{m+1} itself unchanged. Then we apply the ECB layer, computing $z_i \leftarrow E_K(y_i)$ for the first $m-1$ full blocks, and computing $u \leftarrow E_K(y_m)$ and $z_m \leftarrow E_K(u)$ for the last full block. The first bits of u are then xor-ed into the partial block, setting $w_{m+1} = x_{m+1} \oplus u|_{1..\ell}$. Then we do the final BPE layer (adding $(w_{m+1}10..0)$ to each full block), thus getting $w_i \leftarrow \mathrm{BPE}(\mathbf{z})_i \oplus (w_{m+1}10..0)$ and the TET output is the vector $w_1, \ldots, w_m, w_{m+1}$.

3.3 The PRF function

It is clear that any secure pseudo-random function can be used to derive the element β. We describe now a specific PRF, which is a slight adaptation of the OMAC construction of Iwata and Korasawa [IK03], that seems well suited for our application. The slight modification to OMAC can be thought of as constructing a "tweakable PRF", with an on-line/off-line optimization for the tweak.[1] (In our case, the input-length of TET is the "tweak" for the PRF and the tweak of TET is the input to the PRF.)

We assume that the input length of TET is less than 2^n bits, and we denote by L the input length in bits encoded as an n-bit integer. Also denote the tweak for TET (which is the input to the PRF) by $T = \langle T_1, \ldots, T_{m'} \rangle$ where $|T_1| = \cdots = |T_{m'-1}| = n$ and $1 \le |T_{m'}| \le n$.

To compute $\beta \leftarrow \mathrm{PRF}_K(L, T)$ we first compute $X \leftarrow E_K(L)$, then compute β as a CBC-MAC of T, but before the last block-cipher application we xor either the value αX or the value $\alpha^2 X$ (depending on whether the last block is a full block or a partial block). In more details, we set $V_0 = 0$ and then $V_i \leftarrow E_K(V_{i-1} \oplus T_i)$ for $i = 1, \ldots, m' - 1$. Then, if the last block is a full block ($|T_{m'}| = n$) then we set $\beta \leftarrow E_K(\alpha X \oplus V_{m'-1} \oplus T_{m'})$, and if the last block is a partial block ($|T_{m'}| < n$) then we set $\beta \leftarrow E_K(\alpha^2 X \oplus V_{m'-1} \oplus (T_{m'}10..0))$.

Notice that the only difference between this function and the OMAC construction is that OMAC does not have the additional input L and it sets $X \leftarrow E_K(0)$. Proving that this is a secure pseudo-random function is similar to the proof of OMAC [IK03], and is omitted here.

We point out that on one hand, the length L is needed only before processing the last tweak block, so this pseudo-random function is suited for streaming applications where the length of messages is not known in advance.[2] On the other hand, if used with a fixed input length (where L is known ahead of time) then the computation of X can be done off line, in which case we save one block-cipher application during the on-line phase.

---

[1] Formally there is not much difference between a "tweakable" and "non-tweakable" PRF, one can always process the tweak by concatenating it to the input. But here it is convenient to make the distinction since we can offer some tweak-specific performance optimization.

[2] As explained in Section 3.5, TET is not a very good fit for such cases, but this PRF functions can perhaps be used in applications other than TET.

3.4 Some other details

To get a fully-specified mode of operation one needs to set many other small details. Below I explain my choices for the details that I set, and describe those that are still left unspecified.

If the input length is	then these elements are bad values for τ	Bad key probability
512 bytes	$\alpha^{(2^{128}-1)/3},\ \alpha^{2\cdot(2^{128}-1)/3}$	2^{-127}
1024 bytes	$\alpha^{i\cdot(2^{128}-1)/5}$ for $i=1,2,3,4$	2^{-126}
4096 bytes	$\alpha^{i\cdot(2^{128}-1)/257}$ for $i=1,2,\ldots,256$	2^{-120}
65536 bytes	$\alpha^{i\cdot(2^{128}-1)/17}$ for $i=1,2,\ldots,16$	2^{-124}

Table 1. Bad τ values for various input lengths, assuming $n = 128$

The element $\alpha \in \mathrm{GF}(2^n)$. Recall that BPE uses a fixed primitive element $\alpha \in \mathrm{GF}(2^n)$. If the field $\mathrm{GF}(2^n)$ is represented with a primitive polynomial, then this fixed element should be set as the polynomial x (or $1/x$), in which case a multiplication by α can be implemented with an n-bit shift and a conditional xor.[3]

The two hashing layers. I chose to use the same hashing keys τ, β for both hashing layers. The security of the mode does not seem to be effected by this. On the other hand, having different keys for the two hashing layers adds a considerable burden to an implementation, especially it if optimizes the GF multiplications by preparing some tables off line.

The hashing key τ. I also chose to derive the hashing key τ from the same cipher key as the hashing key β, rather than being a separate key. (This decision is rather arbitrary, I made it because I could not see any reason to keep τ as a separate key.) Specifically, it can be set as $\tau \leftarrow \mathrm{PRF}_K(0, 0^n) = E_K(\alpha \cdot E_K(0))$. Note that this is not a duplicate of any $\mathrm{PRF}_K(L, T)$, since the input length L is always at least n bits.[4]

Of course, τ must be chosen so that for any message length m it holds that $\sigma_m \neq 0$ (where $\sigma_m = \sum_{i=0}^{m} \tau^m$). Hence if setting $\tau \leftarrow \mathrm{PRF}_K(0, 0)$ results in a bad value for τ then we can keep trying $\mathrm{PRF}_K(0, 1)$, $\mathrm{PRF}_K(0, 2)$, etc. When using TET with fixed input length (containing m complete blocks), we can just include a list of all the "bad τ values" for which $\sigma_m = 0$ with the implementation. This list is fairly easy to construct: Denoting $g = GCD(m+1, 2^n-1)$, when m is even the lists consists of $\alpha^{i\cdot(2^n-1)/g}$ for $i = 1, 2, \ldots, g-1$

[3] The choice between setting $\alpha = x$ or $\alpha = 1/x$ depends on the endianess of the field representation, and it should be made so that multiplication by α requires left shift and not right shift.

[4] Setting $\tau \leftarrow E_K(0)$ would work just as well in this context, but the effort in proving it is too big for the minuscule saving in running time.

(where α is a primitive element). When m is odd it consists of the same elements and also of the element $\alpha^0 = 1$. In Table 1 we list the "bad τ values" for various input lengths assuming $n = 128$.

The approach of having a fixed list of "bad τ values" may not work as well when using TET with variable-input length. One way to handle this case is to insist on τ being a primitive element in $\mathrm{GF}(2^n)$, in which case we know that $\sigma_m \neq 0$ for all length m. (We can efficiently test is τ is a primitive element given the prime factorization of $2^n - 1$). But a better way of handling variable length is to allow different τ's for different input lengths. Specifically, when handling a message of with m full blocks, we try $\mathrm{PRF}_K(0,0)$, $\mathrm{PRF}_K(0,1)$, ... and set τ to the first value for which $\sigma_m \neq 0$. It is not hard to see that this is just as secure as insisting on the same τ for all lengths (since we only use τ to argue about collisions between messages of the same length).

Ordering the blocks for polynomial-evaluation. I chose to order the blocks at the input of BPE in "reverse order", evaluating the polynomial as $\sum_{i=1}^{m} x_i \tau^{m-i+1}$. The reason is to allow processing to start as soon as possible in the case where the input arrives one block at a time. We would like to use Horner's rule when computing $\mathrm{BPE}(\mathbf{x})$, processing the blocks in sequence as

$$s = (\ldots((x_1\tau \oplus x_2)\tau \oplus x_3)\tau \ldots \oplus x_m)\tau$$

which means that x_1 is multiplied by τ^m, x_2 is multiplied by τ^{m-1}, etc. Similarly when computing $\mathrm{BPE}^{-1}(\mathbf{y})$ we would implement the polynomial-evaluation as

$$s = (\ldots((y_1\tau \oplus y_2)\tau \oplus y_3)\tau \ldots \oplus y_m)(\tau/\sigma)$$

which means that y_1 is multiplied by τ^m/σ, y_2 is multiplied by τ^{m-1}/σ, etc.

The hashing direction. For each of the two hashing layers, one can use either of BPE, BPE^{-1}, $\widetilde{\mathrm{BPE}}$, or $\widetilde{\mathrm{BPE}}^{-1}$. For the encryption direction, I chose to use $\widetilde{\mathrm{BPE}}^{-1}$ for the first hashing layer and BPE^{-1} for the second layer. This means that on decryption we use BPE as the first hashing layer and $\widetilde{\mathrm{BPE}}$ for the second layer.

I chose the inverse hash function on encryption and the functions themselves on decryption because inverting the functions may be less efficient than computing them in the forward direction (since one needs to multiply also by τ/σ). In a typical implementation for storage, one would use encryption when writing to storage and decryption when reading back from storage. As most storage is optimized for read (at the expense of the less-frequent write operations), it makes sense to allocate the faster operations for read in this case too.

As for the choice between BPE and $\widetilde{\mathrm{BPE}}$, I chose to add the vector $\mathbf{b}$ in the middle, right before and after the ECB layer. The rationale here is that it is possible to do the computation $\beta \leftarrow \mathrm{PRF}_K(L, T)$ concurrently with the multiplication by M_τ (or its inverse).

function $\mathrm{PRF}_K(L, T_1 \cdots T_{m'})$ // $|L| = |T_1| = \cdots = |T_{m'-1}| = n,\ 1 \le |T_{m'}| \le n$

001 $V_0 \leftarrow 0,\ X \leftarrow E_K(L)$
002 **for** $i \leftarrow 1$ **to** $m' - 1$ **do** $V_i \leftarrow E_K(V_{i-1} \oplus T_i)$
003 **if** $|T_{m'}| = n$ **then return** $E_K(V_{m'-1} \oplus T_{m'} \oplus \alpha X)$
004 **else return** $E_K(V_{m'-1} \oplus T_{m'} \oplus \alpha^2 X)$

Algorithm $\mathrm{TET}_{K_1,K_2}(T; P_1 \cdots P_m P_{m+1})$
// $|P_1| = \cdots = |P_m| = n,\ 0 \le |P_{m+1}| < n$

101 $L \leftarrow mn + |P_{m+1}|$ // input size (bits)
102 $i = 0$
103 $\tau \leftarrow \mathrm{PRF}_{K_1}(0, i),\ \sigma \leftarrow 1 \oplus \tau \oplus \ldots \oplus \tau^m$
104 **if** $\sigma = 0$ **then** $i \leftarrow i + 1$, **goto** 103
105 $\beta \leftarrow \mathrm{PRF}_{K_1}(L, T),\ SP \leftarrow 0,\ SC \leftarrow 0$

110 **for** $i \leftarrow 1$ **to** m **do** $SP \leftarrow (SP \oplus P_i) \cdot \tau$
111 $SP \leftarrow SP/\sigma$
112 **if** $|P_{m+1}| > 0$ **then**
113 $\quad SP \leftarrow SP \oplus P_{m+1}$ padded with 10..0

120 **for** $i \leftarrow 1$ **to** m **do**
121 $\quad PP_i \leftarrow P_i \oplus SP$
122 $\quad PPP_i \leftarrow PP_i \oplus \alpha^{i-1}\beta$
123 **for** $i \leftarrow 1$ **to** $m - 1$ **do**
124 $\quad CCC_i \leftarrow E_{K_2}(PPP_i)$
125 **if** $|P_{m+1}| > 0$ **then**
126 $\quad MM \leftarrow E_{K_2}(PPP_m)$
127 $\quad CCC_m \leftarrow E_{K_2}(MM)$
128 $\quad C_{m+1} \leftarrow P_{m+1} \oplus (MM$ truncated$)$
129 **else** $CCC_m \leftarrow E_{K_2}(PPP_m)$

130 **for** $i \leftarrow 1$ **to** m **do**
131 $\quad CC_i \leftarrow CCC_i \oplus \alpha^{i-1}\beta$
132 $\quad SC \leftarrow (SC \oplus CC_i) \cdot \tau$
133 $SC \leftarrow SC/\sigma$
134 **if** $|P_{m+1}| > 0$ **then**
135 $\quad SC \leftarrow SC \oplus C_{m+1}$ padded with 10..0

140 **for** $i \leftarrow 1$ **to** m **do**
141 $\quad C_i \leftarrow CC_i \oplus SC$

150 **return** $C_1 \ldots C_m C_{m+1}$

Algorithm $\mathrm{TET}^{-1}_{K_1,K_2}(T; C_1 \cdots C_m C_{m+1})$
// $|C_1| = \cdots = |C_m| = n,\ 0 \le |C_{m+1}| < n$

201 $L \leftarrow mn + |C_{m+1}|$ // input size (bits)
202 $i = 0$
203 $\tau \leftarrow \mathrm{PRF}_{K_1}(0, i),\ \sigma \leftarrow 1 \oplus \tau \oplus \ldots \oplus \tau^m$
204 **if** $\sigma = 0$ **then** $i \leftarrow i + 1$, **goto** 203
205 $\beta \leftarrow \mathrm{PRF}_{K_1}(L, T),\ SP \leftarrow 0,\ SC \leftarrow 0$

210 **for** $i \leftarrow 1$ **to** m **do** $SC \leftarrow (SC \oplus C_i) \cdot \tau$
212 **if** $|C_{m+1}| > 0$ **then**
213 $\quad SC \leftarrow SC \oplus C_{m+1}$ padded with 10..0

220 **for** $i \leftarrow 1$ **to** m **do**
221 $\quad CC_i \leftarrow C_i \oplus SC$
222 $\quad CCC_i \leftarrow CC_i \oplus \alpha^{i-1}\beta$
223 **for** $i \leftarrow 1$ **to** $m - 1$ **do**
224 $\quad PPP_i \leftarrow E^{-1}_{K_2}(CCC_i)$
225 **if** $|C_{m+1}| > 0$ **then**
226 $\quad MM \leftarrow E^{-1}_{K_2}(CCC_m)$
227 $\quad PPP_m \leftarrow E^{-1}_{K_2}(MM)$
228 $\quad P_{m+1} \leftarrow C_{m+1} \oplus (MM$ truncated$)$
229 **else** $PPP_m \leftarrow E^{-1}_{K_2}(CCC_m)$

230 **for** $i \leftarrow 1$ **to** m **do**
231 $\quad PP_i \leftarrow PPP_i \oplus \alpha^{i-1}\beta$
232 $\quad SP \leftarrow (SP \oplus PP_i) \cdot \tau$
234 **if** $|C_{m+1}| > 0$ **then**
235 $\quad SP \leftarrow SP \oplus P_{m+1}$ padded with 10..0

240 **for** $i \leftarrow 1$ **to** m **do**
241 $\quad P_i \leftarrow PP_i \oplus SP$

250 **return** $P_1 \ldots P_m P_{m+1}$

Fig. 2. Enciphering and deciphering under TET, with plaintext $P = P_1 \ldots P_m P_{m+1}$, ciphertext $C = C_1 \cdots C_m C_{m+1}$, and tweak T. The element $\alpha \in \mathrm{GF}(2^n)$ is a fixed primitive element.

Given the choices above, the specification of the TET mode is given in Figure 2. Other details that are not specified here are the choice of the underlying cipher and the block-size n, and the representation of the field $\mathrm{GF}(2^n)$ (including endianess issues).

3.5 Performance of TET

As specified above, the TET mode can be used with variable input length, and in the long version of this note [Hal07] we prove that it is secure when used in this manner. However, its efficiency (at least in software) depends crucially on pre-processing that is only possible when used with fixed input length (or at least with a small number of possible lengths). The reason is that on encryption one needs to multiply by τ/σ, which depends on the message length (since $\sigma = \sum_{i=0}^{m} \tau^i$). When used with fixed input length, the value τ/σ can be computed off line, and some tables can be derived to speed up the multiplication by τ/σ. When used with variable input length, however, the value τ/σ must be computed on-line, which at least for software implies a considerable cost. Hence, TET is not very appealing as a variable-input-length mode.

We stress, however, that the motivating application for TET, namely "sector-level encryption", is indeed a fixed-input-length application. Also, there are some limited settings where one can use variable input length without suffering much from the drawback above. For example, a "write once / read many times" application, where the data is encrypted once and then decrypted many times, would only need to worry about computing σ in the initial encryption phase (since σ is not used during decryption). Also, the same value of σ is used for every bit-length from mn to $(m+1)n-1$, so length variability within this limited range in not effected.[5]

Below we analyze the performance characteristics of TET only for fixed input length. With this assumption, the computation of the PRF function from above takes exactly m' applications of the cipher, where m' is the number of blocks of associated data (full or partial). (This is because the computation of the mask value $X \leftarrow E_K(L)$ can be done off line.) Then we need either m or $m-1$ GF-multiplies for the polynomial evaluation (depending if we have m or $m-1$ full blocks), followed by m block-cipher applications for the ECB layer, and again m or $m-1$ GF multiplies. Altogether, we need $m+m'$ block-cipher applications and either $2m$ or $2m-2$ GF multiplies. (The shift and xor operations that are also needed are ignored in this description, since they are insignificant in comparison.)

Table 2 compares the number of block-cipher calls and GF multiplies in CMC, EME*, XCB, HCH, and TET.[6] It is expected that software efficiency

[5] For example, an implementation can handle both 512-byte blocks and 520-byte blocks with a single value of σ (assuming block length of $n = 128$ bits).

[6] The other modes are not included since EME is essentially a special case of EME*, PEP an ABL4 are significantly less efficient than the others, and HCTR is almost identical to HCH.

Mode	CMC	EME*	XCB	HCH	TET
Block-cipher calls	$2m+1$	$2m+1+\lceil m/n \rceil$	$m+1$	$m+3$	$m+1$
GF multiplies	–	–	$2(m+3)$	$2(m-1)$	$2m$ or $2(m-1)$

Table 2. Workload for enciphering an m-block input with a 1-block tweak.

will be proportional to these numbers. (As far as I know, the current "common wisdom" is that computing a GF(2^{128}) multiplication in software using the approach from [Sho96] with reasonable-size tables, is about as fast as a single application of AES-128.)

As for hardware implementations, all the modes except CMC are parallelizable and pipelinable, so they can be made to run as fast as needed using sufficiently large hardware. Table 3 describes a somewhat speculative efficiency comparison of hypothetical "fully pipelined" implementations of the modes from above (except CMC). In that table I assume (following [YMK05]) that a one-cycle GF(2^{128}) multiplication is about three times the size of a module for computing the AES round function, and that AES-128 is implemented as 10 such modules. A few other relevant characteristics of these modes are discussed next.

Mode	EME*	XCB	HCH	TET
Latency	$m+30$	$m+13$	$m+31$	$2m+11$
Time	$2m+10(\lceil m/n \rceil+2)$	$2m+27$	$2m+31$	$2m+11$
Size	10	13	13	13

Table 3. Hardware efficiency: A speculative comparison of pipelined implementations for m-block input and 1-block tweak. Latency is number of cycles until first output block, time is number of cycles until last output block, and size is measured in the equivalent of number of AES-round modules.

Any input length. All of these modes except CMC support any input length from n bits and up. CMC supports only input length which is a multiple of n bits (but it should be relatively straightforward to extended it using ciphertext-stealing).

Associated data. The modes EME*, XCB and TET support tweaks of arbitrary length. CMC and HCH support only n-bit tweaks (but it is straightforward to extended them to support arbitrary-size tweaks).

Security proofs. The security of XCB was not proved formally, only a sketch was given, and CMC was only proven secure with respect to fixed input-length. The other modes were proven secure with respect to variable input-length. Providing the missing proof seems fairly straightforward to me (but one never knows for sure until the proof is actually written).

Number of keys. Although in principle it is always possible to derive all the needed key material from just one key, different modes are specified with different requirements for key material. In fact using two keys as in TET (one for the ECB layer and one for everything else) offers some unexpected practical advantages over using just one key.

Specifically, implementations sometimes need to be "certified" by standard bodies (such as NIST), and one criterion for certification is that the implementation uses an "approved mode of operation" for encryption. Since standard bodies are slow to approve new modes, it may be beneficial for a TET implementation to claim that it uses the "approved" ECB mode with pre-and post-processing, and moreover the pre- and post-processing is independent of the ECB key. Also, detaching the ECB key from the key that is used elsewhere may make it easier to use a hardware accelerator that supports ECB mode.

4 Security of TET

We relate the security of TET to the security of the underlying primitives from which it is built as follows:

Theorem 1. [TET security] *Fix $n, s \in \mathbb{N}$. Consider an adversary attacking the* TET *mode with a truly random permutation over $\{0,1\}^n$ in place of the block cipher and a truly random function instead of* PRF*, such that the total length of all the queries that the attacker makes is at most s blocks altogether.*

The advantage of this attacker in distinguishing TET *from a truly random tweakable length-preserving permutation is at most $1.5s^2/\phi(2^n-1)$ (where ϕ is Euler's totient function). Using the notations from Appendix A, we have*

$$\mathbf{Adv}^{\pm\widetilde{\mathrm{prp}}}_{\mathrm{TET}}(s) \le \frac{3s^2}{2\phi(2^n-1)}$$

The proof appears in the long version of this note [Hal07]. The intuition is that as long as there are no block collisions in the hash function, then the random permutation in the ECB layer will be applied to new blocks, so it will will output random blocks and the answer that the attacker will see is therefore random. □

Corollary 1. *With the same setting as in Theorem 1, consider an attacker against* TET *with a specific cipher E and a specific* PRF *F, where the attack uses at most total of s' blocks of associated data. Then*

$$\mathbf{Adv}^{\pm\widetilde{\mathrm{prp}}}_{\mathrm{TET}[E]}(t,s,s') \le \frac{3s^2}{2\phi(2^n-1)} + \ 2(\ \mathbf{Adv}^{\pm\mathrm{prp}}_{E}(t',s) + \mathbf{Adv}^{\mathrm{prf}}_{\mathrm{PRF}}(t',s'))$$

where $t' = t + O(n(s+s'))$. □

5 Conclusions

We presented a new method for invertible "blockwise universal" hashing which is about as efficient as polynomial-evaluation hash, and used it in a construction of a tweakable enciphering scheme called TET. This complements the current lineup of tweakable enciphering schemes by providing a scheme in the family of "hash-ECB-hash" which is as efficient as the schemes in the "hash-CTR-hash" family. We also expect that the hashing scheme itself will find other uses beyond TET.

References

[CS06a] Debrup Chakraborty and Palash Sarkar. HCH: A new tweakable enciphering scheme using the hash-encrypt-hash approach. In *INDOCRYPT'06*, volume 4329 of *Lecture Notes in Computer Science*, pages 287–302. Springer, 2006.

[CS06b] Debrup Chakraborty and Palash Sarkar. A new mode of encryption providing a tweakable strong pseudo-random permutation. In *The 13th International Workshop on Fast Software Encryption – FSE'06*, volume 4047 of *Lecture Notes in Computer Science*, pages 293–309. Springer, 2006.

[FM04] Scott R. Fluhrer and David A. McGrew. The extended codebook (XCB) mode of operation. Technical Report 2004/278, IACR ePrint archive, 2004. `http://eprint.iacr.org/2004/278/`.

[GM84] Shafi Goldwasser and Silvio Micali. Probabilistic encryption. *Journal of Computer and System Sciences*, 28(2):270–299, April 1984.

[Hal04] Shai Halevi. EME*: extending EME to handle arbitrary-length messages with associated data. In *INDOCRYPT'04*, volume 3348 of *LNCS*, pages 315–327. Springer, 2004.

[Hal07] Shai Halevi. Invertible Universal Hashing and the TET Encryption Mode. In *Advances in Cryptology – CRYPTO '07*, 2007. Long version available on-line at `http://eprint.iacr.org/2007/014/`.

[HR03] S. Halevi and P. Rogaway. A tweakable enciphering mode. In *Advances in Cryptology – CRYPTO '03*, volume 2729 of *LNCS*, pages 482–499. Springer, 2003.

[HR04] Shai Halevi and Phil Rogaway. A parallelizable enciphering mode. In *The RSA conference – Cryptographer's track, RSA-CT'04*, volume 2964 of *Lecture Notes in Computer Science*, pages 292–304. Springer-Velrag, 2004.

[IK03] Tetsu Iwata and Kaoru Kurosawa. "OMAC: One-Key CBC MAC". In *Fast Software Encryption - FSE'03*, volume 2887 of *Lecture Notes in Computer Science*, pages 129–153. Springer, 2003.

[LR88] Michael Luby and Charles Rackoff. How to construct pseudorandom permutations from pseudorandom functions. *SIAM Journal of Computing*, 17(2), April 1988.

[LRW02] Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In *Advances in Cryptology – CRYPTO '02*, volume 2442 of *Lecture Notes in Computer Science*, pages 31–46. Springer, 2002.

[MM82] Carl H. Meyr and Stephen M. Matyas. *Cryptography: A New Dimension in Computer Data Security*. John Wiley & Sons, 1982.

[MV04] David A. McGrew and John Viega. Arbitrary block length mode. Manuscript. Available on-line from `http://grouper.ieee.org/groups/1619/email/pdf00005.pdf`, 2004.

[NR97] Moni Naor and Omer Reingold. A pseudo-random encryption mode. Manuscript, available from `http://www.wisdom.weizmann.ac.il/~naor/`, 1997.

[NR99] Moni Naor and Omer Reingold. On the construction of pseudo-random permutations: Luby-Rackoff revisited. *Journal of Cryptology*, 12(1):29–66, 1999.

[Sho96] V. Shoup. On fast and provably secure message authentication based on universal hashing. In *Advances in Cryptology – CRYPTO '96*, volume 1109 of *Lecture Notes in Computer Science*, pages 74–85. Springer-Verlag, 1996.

[WFW05] Peng Wang, Dengguo Feng, and Wenling Wu. HCTR: A variable-input-length enciphering mode. In *Information Security and Cryptology – CISC'05*, volume 3822 of *Lecture Notes in Computer Science*, pages 175–188. Springer, 2005.

[YMK05] Bo Yang, Sambit Mishra, and Ramesh Karri. "A High Speed Architecture for Galois/Counter Mode of Operation (GCM)". Technical Report 2005/146, IACR ePrint archive, 2005. `http://eprint.iacr.org/2005/146/`.

A Preliminaries

A *tweakable enciphering scheme* is a function $\mathbf{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ where $\mathcal{M} = \bigcup_{i\in I}\{0,1\}^i$ is the *message space* (for some nonempty index set $I \subseteq \mathbb{N}$) and $\mathcal{K} \neq \emptyset$ is the *key space* and $\mathcal{T} \neq \emptyset$ is the *tweak space*. We require that for every $K \in \mathcal{K}$ and $T \in \mathcal{T}$ we have that $\mathbf{E}(K,T,\cdot) = \mathbf{E}_K^T(\cdot)$ is a length-preserving permutation on $\mathcal{M}$. The inverse of an enciphering scheme $\mathbf{E}$ is the enciphering scheme $\mathbf{D} = \mathbf{E}^{-1}$ where $X = \mathbf{D}_K^T(Y)$ if and only if $\mathbf{E}_K^T(X) = Y$. A *block cipher* is the special case of a tweakable enciphering scheme where the message space is $\mathcal{M} = \{0,1\}^n$ (for some $n \geq 1$) and the tweak space is the singleton set containing the empty string. The number n is called the *blocksize*. By $\mathrm{Perm}(n)$ we mean the set of all permutations on $\{0,1\}^n$. By $\mathrm{Perm}^{\mathcal{T}}(\mathcal{M})$ we mean the set of all functions $\pi : \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ where $\pi(T,\cdot)$ is a length-preserving permutation.

An *adversary* A is a (possibly probabilistic) algorithm with access to some oracles. Oracles are written as superscripts. By convention, the running time of an algorithm includes its description size. The notation $A \Rightarrow 1$ describes the event that the adversary A outputs the bit one.

Security measure. For a tweakable enciphering scheme $\mathbf{E} : \mathcal{K} \times \mathcal{T} \times \mathcal{M} \to \mathcal{M}$ we consider the advantage that the adversary A has in distinguishing $\mathbf{E}$ and its inverse from a random tweakable permutation and its inverse: $\mathbf{Adv}_{\mathbf{E}}^{\widetilde{\pm\mathrm{prp}}}(A) =$

$$\Pr\left[K \xleftarrow{\$} \mathcal{K} : \ A^{\mathbf{E}_K(\cdot,\cdot)\,\mathbf{E}_K^{-1}(\cdot,\cdot)} \Rightarrow 1\right] - \Pr\left[\pi \xleftarrow{\$} \mathrm{Perm}^{\mathcal{T}}(\mathcal{M}) : \ A^{\pi(\cdot,\cdot)\,\pi^{-1}(\cdot,\cdot)} \Rightarrow 1\right]$$

The notation shows, in the brackets, an experiment to the left of the colon and an event to the right of the colon. We are looking at the probability of the indicated event after performing the specified experiment. By $X \xleftarrow{\$} \mathcal{X}$ we mean to choose X at random from the finite set $\mathcal{X}$. In writing $\pm\widetilde{\text{prp}}$ the tilde serves as a reminder that the PRP is tweakable and the $\pm$ symbol is a reminder that this is the "strong" (chosen plaintext/ciphertext attack) notion of security. For a block cipher, we omit the tilde.

Without loss of generality we assume that an adversary never repeats an encipher query, never repeats a decipher query, never queries its deciphering oracle with (T, C) if it got C in response to some (T, M) encipher query, and never queries its enciphering oracle with (T, M) if it earlier got M in response to some (T, C) decipher query. We call such queries *pointless* because the adversary "knows" the answer that it should receive.

When $\mathcal{R}$ is a list of resources and $\mathbf{Adv}_{\Pi}^{\text{xxx}}(A)$ has been defined, we write $\mathbf{Adv}_{\Pi}^{\text{xxx}}(\mathcal{R})$ for the maximal value of $\mathbf{Adv}_{\Pi}^{\text{xxx}}(A)$ over all adversaries A that use resources at most $\mathcal{R}$. Resources of interest are the running time t, the number of oracle queries q, and the total number of n-bit blocks in all the queries s. The name of an argument (e.g., t, q, s) will be enough to make clear what resource it refers to.

B Intellectual-Property Issues

The original motivation for devising the TET mode was to come up with a reasonably efficient mode that is "clearly patent-free". The IEEE security-in-storage working group (SISWG) was working on a standard for length-preserving encryption for storage, and some of the participants expressed the wish to have such a mode. (**Disclaimer:** Not being a patent lawyer, I can only offer my educated guesses for the IP status of the various modes. The assessment below reflects only my opinion about where things stand.)

The modes CMC/EME/EME* from the "encrypt-mix-encrypt" family are all likely to be patent-encumbered, due to US Patent Application 20040131182A1 from the University of California (which as of this writing was not yet issued). Similarly, the XCB mode – which is the first proposed mode in the "hash-CTR-hash" family – is likely to be patent-encumbered due to a US patent application US20070081668A1 from Cisco Systems (also still not issued as of this writing). The status of the other members of the "hash-CTR-hash" family is unclear: they may or may not be covered by the claims of the Cisco patent when it is issued.

This state of affairs left the "hash-ECB-hash" approach as the best candidate for finding patent-free modes: this approach is based on the paper of Naor and Reingold [NR97] that pre-dates all these modes by at least five years, and for which no patent was filed. Specifically for TET, I presented the basic construction from Section 2 in an open meeting of the IEEE SISWG in October of 2002 (as well as in an email message that I sent the SISWG mailing list on October 9, 2002), and I never filed for any patent related to this construction. Thus it seems unlikely to me that there are any patents that cover either the general "hash-ECB-hash" approach or TET in particular.