Multiforked Iterated Even-Mansour and a Note on the Tightness of IEM Proofs

Elena Andreeva[1], Amit Singh Bhati[2,3], Andreas Weninger[1]

[1] TU Wien, Austria
elena.andreeva@tuwien.ac.at, andreas.weninger@tuwien.ac.at
[2] COSIC, KU Leuven, Belgium
[3] 3milabs, Belgium
amitsingh.bhati@3milabs.tech

Keywords: Iterated Even-Mansour · forkcipher · provable security

Abstract. The Iterated Even-Mansour (IEM) construction was introduced by Bogdanov et al. at EUROCRYPT 2012 and can be seen as an abstraction or idealization of blockciphers like AES. IEM provides insights into the soundness of this blockcipher structure and the best possible security for any number of rounds. IEM with r permutations on n-bit blocks is secure up to $q \approx 2^{rn/(r+1)}$ queries to the cipher and each permutation.

Forkciphers, introduced at ASIACRYPT 2019 as expanding symmetric ciphers, have since found applications in encryption, authenticated encryption and key derivation. Kim et al. (ToSC 2020) proposed the first IEM-style forkcipher, FTEM, but their security proof is limited to a 2-round design with tweak processing based on XORing AXU hashes. This offers limited insight into practical forkciphers like ForkSkinny, which use 40 to 56 rounds and a different tweak schedule. No security results currently exist for forked IEM constructions with more than two rounds.

We propose a generalized forked IEM construction called GIEM which integrates any tweakey schedule (including tweak-dependent round keys or constant keys) and thus encompasses IEM, FTEM and similar IEM-related constructions.

We define three forkcipher-related instantiations, FEM (2 branches and no tweaks), FTEM-ITS (2 branches and idealized tweakey schedule) and MFTEM (unlimited branches and AXU-based tweakey schedule). We prove that each construction achieves security similar to the respective non-forked construction. This shows the soundness of the forking design strategy and can serve as a basis for new constructions with more than two branches.

In their work, Bogdanov et al. also propose an attack against IEM using $q \approx 2^{rn/(r+1)}$ queries, which is used in a number of follow-up works to argue the tightness of IEM-related security bounds. In this work, we demonstrate that the attack is ineffective with the specified query complexity. To salvage the purported tightness results, we turn to an attack by Gaži (CRYPTO 2013) against cascading block ciphers and provide the necessary parameters to apply it to IEM. This validates the tightness of the known IEM security bound.

1 Introduction

Blockciphers are a fundamental building block in cryptography. Their design has been an important research focus in the last decades. In 1997 Even and Mansour [EM97] introduced the concept of building a blockcipher based on a public permutation P and a secret key $k = (k_1, k_2)$ as

$$c = P(m \oplus k_1) \oplus k_2 \,.$$

The Even-Mansour cipher (EM) [EM97] is provably secure up to roughly $2^{n/2}$ adversarial queries against adversaries that do not exploit the specific instantiation of the permutation. To formally model this, the permutation is randomly drawn from the set of all permutations and the adversary can query it arbitrarily. The security proof of the Even-Mansour cipher gives a generic security guarantee that shifts the difficulty of designing a secure blockcipher to the arguably simpler concept of designing a secure permutation. In 2012, Bogdanov et al. [BKL$^+$12] extended the EM design to multiple rounds as

$$c = P_r(...P_2(P_1(m \oplus k_1) \oplus k_2)...) \oplus k_{r+1} \,,$$

where $P_1, \ldots, P_r$ are independently and randomly sampled permutations. This construction is referred to as a Key Alternating Cipher (KAC) or Iterated Even-Mansour cipher (IEM) in the literature. While the original EM design is the basis for creating new blockciphers, IEM has an additional relevance, as it follows the internal (permute-round key XOR) alternating structure of popular ciphers, such as AES. Bogdanov et al. first provably analyzed the IEM structure. Since practical ciphers with a similar structure do not use independent random permutations, the security of existing ciphers does not surpass that of IEM. Hence the results on IEM give the minimum number of rounds of any secure key-alternating cipher.

In terms of provable security, the added rounds allow IEM to achieve beyond birthday bound security as a pseudo-random permutation. Bogdanov et al. [BKL$^+$12] proved the security of IEM for information theoretic adversaries up to roughly $q \approx 2^{2n/3}$ queries when the number of rounds $r \geq 2$, and conjectured a security of $q \approx 2^{rn/(r+1)}$ in general.

The IEM work of Bogdanov et al. [BKL$^+$12] sparked a broad body of follow-up works. In 2014, Chen and Steinberger [CS14] proved the IEM conjectured security up to about $2^{rn/(r+1)}$ queries for any number of rounds r. Later, Hoang and Tessaro [HT16] confirm this asymptotic bound with a proof in their new framework called the *expectation method*. This method both enables a simpler proof and an improved final bound. While asymptotically the same, the bound by Hoang and Tessaro gives better concrete security. For AES-like parameters ($n = 128$ and $r = 10$) and $q = 2^{110}$ queries to the cipher and each primitive, it gives still a reasonable bound on the maximum advantage of 2^{-50} whereas Chen and Steinberger's bound is already vacuous. Further, the expectation method gives a bound on the multi-user security of IEM. Multi-user security assumes each user to have independent keys but the same permutations, which better captures how blockciphers are deployed in the real world.

In a different line of work, Cogliati et al. [CLS15] proposed a tweakable variant of IEM called Tweakable Even-Mansour (TEM). Tweakable block ciphers (TBCs) are a generalization of traditional block ciphers that, aside from the message and key, take an additional tweak input, which may be adversarially controlled. For each distinct tweak, the tweakable cipher is supposed to behave like an independent random permutation. TBCs are used for length-preserving encryption modes [HR03, HR04], message authentication codes [LRW02, LST12] and authenticated encryption modes [RBBK01, LRW02, Rog04]. Cogliati et al. [CLS15] generate the round keys of TEM with almost XOR-universal (AXU) functions on input the current tweak. They proved security up to roughly $q \approx 2^{rn/(r+1)}$ queries in a dedicated proof for the 1- and 2-round construction, and security up to $q \approx 2^{rn/(r+2)}$ queries when $r \geq 3$, which is equivalent to the expected security of (regular) IEM with half as many rounds.

A number of works on IEM-like constructions [BKL$^+$12, LPS12, Ste12, CS14, CLS15] argue about the tightness of their security bounds. Their arguments appear to be rooted in the attack by Bogdanov et al. [BKL$^+$12] which uses $2^{rn/(r+1)}$ queries to each permutation and the cipher. However, to the best of our knowledge, a detailed analysis of this attack is missing in the existing literature.

In 2020, Kim et al. [KLL20] proposed FTEM, a novel key-alternating structure similar to a forkcipher. Forkciphers were introduced by Andreeva et al. [ALP$^+$19]. They are expanding tweakable ciphers, that produce $2n$-bit ciphertext outputs from an n-bit input message. Intuitively, their output can be seen as the result of the application of two distinct tweakable blockciphers over the same input message. Forkciphers are used to build robust authenticated encryption schemes [ALP$^+$19, ADP$^+$20, ABV21, BAV24, BPA$^+$23, BBDL23], generalized CTR mode encryption [ABPV21], tweakable enciphering scheme [BVA24], pseudorandom number generators [AW23, BDA$^+$24], message authentication codes [BAMV24] and efficient PRFs [DGL22]. The work of Kim et al. [KLL20] follows the idea of Cogliati et al. [CLS15] and uses AXU functions to derive round keys from the tweak. They specify FTEM for 2 branches and any number of rounds. Although FTEM prescribes a way to build new forkciphers, it does not capture existing forkciphers with a different key schedule such as ForkSkinny [ALP$^+$19]. The mismatch stems from an optimization in FTEM to reduce the necessary number of AXU functions, which is not present in ForkSkinny. Widening the gap even further, the FTEM security proof is limited to two rounds whereas ForkSkinny variants have 40 to 56 rounds. This brings up the question of how one can achieve a maximally secure "forked" TEM, that also encompasses ForkSkinny. Kim et al. [KLL20] prove the security of FTEM up to roughly $q \approx 2^{2n/3}$ queries, thus matching IEM with the same number of rounds and gaining tweakability and forking "for free". They conclude their work with two open questions:

"What is the security of this construction with more than 2 rounds and can it be extended to a multiforkcipher with more than 2 branches?"

A multiforkcipher [ABPV21] is an expanding cipher with more than two output branches. Multiforkciphers have immediate applications, for example, in CTR-style encryption where already $b = 2$ branch (multi)forkciphers have been shown

to be approximately 20% more efficient than the regular blockcipher-based CTR encryption mode [ABPV21].

Name	r	b	**Tweakable**	**Security (Queries)**
IEM [CS14, HT16]	any	1	no	$2^{rn/(r+1)}$
TEM, $r \leq 2$ [CLS15]	2	1	yes (AXU)	$2^{2n/3}$
TEM [CLS15]	any	1	yes (AXU)	$2^{rn/(r+2)}$
FTEM [KLL20]	2	2	yes (AXU)	$2^{2n/3}$
FEM (this work)	any	2	no	$2^{rn/(r+1)}$
FTEM-ITS (this work)	any	2	yes (ideal)	$2^{rn/(r+1)}$
MFTEM (this work)	any	any	yes (AXU)	$\frac{1}{b^2}2^{rn/(r+2)}$

Fig. 1. Comparison of Iterated Even-Mansour variants. r denotes the number of rounds and b the number of (multiforkcipher) branches. For tweakable ciphers, we distinguish based on the tweakey schedule, using AXU (almost XOR universal) functions or an idealized tweakey schedule. Security is given in the number of adversarial queries that are roughly needed to attack successfully, so a higher number is better.

1.1 Contribution and Related Work

In support of additional analysis of forkciphers and similar iterative expanding structures, we introduce GIEM. GIEM is a generalization of IEM that allows an arbitrary choice for tweakey schedule functions, rounds and branches. As such GIEM can be instantiated to EM [EM97], IEM [BKL+12], TEM [CLS15] and FTEM [KLL20]. We provide 3 novel instantiations of GIEM, FEM (2 branches and no tweaks), FTEM-ITS (2 branches and idealized tweakey schedule) and MFTEM (unlimited branches and AXU-based tweakey schedule).

FEM is the forked version of IEM [BKL+12]. As mentioned, IEM starts from the message and then alternates between XORing the round key and applying a permutation, and after r rounds (i.e. permutations) gives a single ciphertext block as output. FEM also follows the IEM structure for $r/2$ rounds, after which the current state is used as the starting point for two branches, each of which performs the IEM structure for $r/2$ rounds again (with independent keys and permutations), resulting in two ciphertext blocks. We prove FEM to have security similar to IEM, i.e. up to $q \approx 2^{rn/(r+1)}$ queries. Due to the structural similarity to IEM, FEM gives direct insight into the security impact of the forking design strategy. Current forkciphers additionally take a tweak input, which is not the case for FEM.

This leads us to our second construction, MFTEM (Multi-Forked Tweakable iterated Even-Mansour), which is tweakable and answers the open questions of Kim et al. [KLL20]. MFTEM is a generalization of FTEM [KLL20], and is the first provably secure IEM-style construction with an arbitrary number of branches. In the two-branch case, MFTEM is similar to FEM, but MFTEM derives its round keys from the tweak input using AXU (almost XOR universal) functions. We

prove the security of MFTEM for any number and rounds, unlike the two-round proof by Kim et al. For b branches, we prove security up to $q \approx \frac{1}{b^2} 2^{rn/(r+2)}$ queries. Here any small constant b has a negligible impact, and in particular for $b = 1$, MFTEM matches the best known bounds for TEM [CLS15]. Similar to FTEM, the tweakey schedule MFTEM is incompatible with ForkSkinny [ALP+19]. Yet, MFTEM can serve as the basis for new multiforkciphers with more than 2 branches.

Finally, we propose FTEM-ITS (2 branches and idealized tweakey schedule) as our third instantiation of GIEM and as an idealization of ForkSkinny. FTEM-ITS requires a stronger assumption than MFTEM (idealized tweakey schedule) but provides better security, up to roughly $q \approx 2^{rn/(r+1)}$ queries.

Overall, our results show that the security impact of the forking design strategy on the constructions we analyzed results in only a negligible security loss. At the same time, the forked constructions can provide two output blocks at the cost of only 1.5 times as many permutation calls. When we add tweaking based on AXU functions (MFTEM with $b = 2$ branches), the security up to $q \approx 2^{rn/(r+2)}$ is worse than that of IEM ($q \approx 2^{rn/(r+1)}$), but this loss is the same as for (single-branch) TEM [CLS15] versus IEM. When tweaking is done with an idealized key schedule, i.e. FTEM-ITS, we achieve security close to IEM (asymptotically up to $q \approx 2^{rn/(r+1)}$ queries). Compared to prior work, our results bring tweaking and forking to IEM (with idealized key schedule) and forking to TEM (with AXU-based key schedule) without weakening the asymptotic security.

In terms of security proofs, our single proof for FEM and FTEM-ITS provides the largest novelty. Our proof strategy follows the expectation method by Hoang and Tessaro [HT16], as well as their proof approach of creating a graph. We define a forked version of this graph. At the heart of our proof we provide a novel argument to bound the difference between sampling in this forked graph versus the regular graph. Combining this with results from the Chen and Steinberger proof for IEM [CS14] gives the final asymptotically tight bound. On the other hand, our MFTEM security proof is based on the proof by Cogliati et al. [CLS15] and has the benefit of being very concise compared to our proof for FEM and FTEM-ITS. The proof uses the probability distribution of IEM with half as many rounds to find the probability of producing a specific state in the middle layer. As a result the asymptotic security of our tweaked construction only matches the security of IEM with half as many rounds, similar to Cogliati's proof.

To answer whether our security bounds are tight, we take a detailed look at the attack by Bogdanov et al. [BKL+12]. Surprisingly, our analysis reveals a problem that renders the attack ineffective. Our analysis relies on the number of queries that Bogdanov et al. defined. While the concrete problem we describe does not apply if the number of queries is multiplied by some constant, the attack by Bogdanov et al. remains without a formal proof. In order to rely on proven results, we instead show that the generic attack by Gaži [Gaž13] against cascades of block ciphers applies to IEM. We describe how to use the attack by providing the necessary parameters.

2 Preliminaries

Notation. For any bitstring X, $|X|$ denotes its length and $X_{[a:b]}$ denotes the bitstring which is comprised of the bits of X at indices a to b. $X_{[1:|X|]} = X$. n is the security parameter, $N = 2^n$. For any set S, let 2^S denote the power set of S and $|S|$ the size of S. For a set S, $\min S$ denotes the minimum element of S, for a tuple (or finite sequence) $I = (i_b)_{b \in B}$, $\min I$ denotes $\min\{i_b | b \in B\}$, similarly $\max S$ the maximum element. For two bitstrings A, B, let $A||B$ denote their concatenation. We write $x \leftarrow\!\$\ X$ to denote uniformly randomly sampling an element x from the set X.

Distributions. Given a finite event space Ω and probability distributions μ, ν defined on Ω, the total variation distance between μ and ν is $||\mu - \nu|| = \frac{1}{2}\sum_{x \in \Omega} |\mu(x) - \nu(x)|$.

Function Families. Let $\mathsf{Perm}(n)$ be the set of all permutations from $\{0,1\}^n$ to $\{0,1\}^n$, and $\mathsf{Func}(n)$ be the set of all functions from $\{0,1\}^n$ to $\{0,1\}^n$. Let $\mathcal{H}$ be a family of functions from some set $\mathtt{T}$ to $\{0,1\}^n$. $\mathcal{H}$ is *uniform* if for any $t \in \mathtt{T}$ and any $y \in \{0,1\}^n$, $\Pr[H \leftarrow\!\$\ \mathcal{H} : H(t) = y] = 2^{-n}$. $\mathcal{H}$ is *ϵ-almost universal* (ϵ-AU) if for any distinct $t, t' \in \mathtt{T}$, $\Pr[H \leftarrow\!\$\ \mathcal{H} : H(t) = H(t')] \leq \epsilon$. $\mathcal{H}$ is *ϵ-almost XOR-universal* (ϵ-AXU) if for any distinct $t, t' \in \mathtt{T}$ and any $y \in \{0,1\}^n$, $\Pr[H \leftarrow\!\$\ \mathcal{H} : H(t) \oplus H(t') = y] \leq \epsilon$.

Multiforkcipher. We slightly adapt the multiforkcipher notion [ABPV21] to allow an arbitrary key K and tweak $\mathtt{T}$ spaces and fix each function to always produce the maximum number of outputs. We define a multiforkcipher (MFC) F_b (with b branches) as a pair of deterministic algorithms, the forward $F_b : K \times \mathtt{T} \times \{0,1\}^n \to \{0,1\}^{bn}$ (i.e. from the key, tweak and message it produces b blocks of n bits each), and the backward (or inverse) $F_b^{-1} : K \times \mathtt{T} \times \{0,1\}^n \times [b] \to \{0,1\}^{bn}$. The extra F_b^{-1} input is called the input indicator. The MFC is said to be *correct*, if for every call $F_b^{-1}(k,t,y)$, where y is one of the output blocks of a $F_b(k,t,x)$ call, the F_b^{-1} output is x followed by the other output blocks of the F_b call. We write $F_b[\mathbf{P}]$ to denote an MFC F that is based on a tuple of permutations $\mathbf{P}$.

The advantage of an adversary $\mathcal{A}$ in distinguishing F_b from a random multi-forked permutation $\tilde{P}$ (c.f. Section 3.2 [ABPV21]) is defined as

$$\mathsf{Adv}^{\mathbf{prtmfp}}_{F_b}(\mathcal{A}) = |\Pr[\mathcal{A}^{\mathbf{prtmfp\text{-}real}_{F_b}} \Rightarrow 1] - \Pr[\mathcal{A}^{\mathbf{prtmfp\text{-}ideal}_{F_b}} \Rightarrow 1]|.$$

where $\mathbf{prtmfp\text{-}real}_{F_b}$ allows $\mathcal{A}$ to access $F_b(k,\cdot,\cdot), F_b^{-1}(k,\cdot,\cdot,\cdot)$, k is a random secret key. $\mathbf{prtmfp\text{-}ideal}_{F_b}$ gives access to $\tilde{P} : K \times \mathtt{T} \times \{0,1\}^n \to \{0,1\}^{bn}$, which is equivalent to a randomly drawn permutation for each tweak $t \in \mathtt{T}$ and each branch $1, \ldots, b$. When F is based on some internal permutations $\mathbf{P}$, then they are drawn randomly at the start of both games and can be queried by $\mathcal{A}$ directly as well. A more detailed MFC definition is given in Appendix A.

H-coefficient Technique [Pat09]. We consider the interactions of a distinguisher $\mathcal{A}$ with an abstract system S that answers $\mathcal{A}$'s queries. The resulting

interaction then generates a transcript $\tau = ((X_1, Y_1), ..., (X_q, Y_q))$ of query-response pairs. S is entirely described by the probabilities $\mathsf{ps}(\tau)$ that correspond to the system S responding with answers as indicated by τ when queries in τ are made. We will generally describe systems informally, or more formally in terms of a set of oracles they provide, and only use the fact that they define corresponding probabilities $\mathsf{ps}(\tau)$ without explicitly giving these probabilities. We say that a transcript is valid for system S if $\mathsf{ps}(\tau) > 0$.

For any systems $\mathbf{S}_1$ and $\mathbf{S}_0$, let $\Delta_{\mathcal{A}}(\mathbf{S}_1, \mathbf{S}_0)$ denote the distinguishing advantage of the adversary $\mathcal{A}$ against the "real" system $\mathbf{S}_1$ and the "ideal" system $\mathbf{S}_0$.

We now describe the H-coefficient technique of Patarin [Pat09]. Generically, it considers a deterministic distinguisher $\mathcal{A}$ that tries to distinguish a "real" system $\mathbf{S}_1$ from an "ideal" system $\mathbf{S}_0$. The adversary's interactions with those systems define transcripts X_1 and X_0, respectively, and a bound on the distinguishing advantage of $\mathcal{A}$ is given by the statistical distance $\mathsf{SD}(X_1, X_0)$.

Lemma 1 (see [Pat09]). *Suppose we can partition the set of valid transcripts for the ideal system into good and bad ones. Further, suppose that there exists $\epsilon \geq 0$ such that $1 - \frac{\mathsf{ps}_1(\tau)}{\mathsf{ps}_0(\tau)} \leq \epsilon$ for every good transcript τ. Then,*

$$\mathsf{SD}(X_1, X_0) \leq \epsilon + \Pr[X_0 \text{ is bad}].$$

Expectation Method. For the expectation method by Hoang and Tessaro [HT16], assume an adversary plays against a single-user security game $\mathsf{su\text{-}game}$ for a keyed construction Π that is built from some primitive P. The adversary is allowed q queries to the construction and p queries to each primitive, and shall output 0 if the construction queries were answered with the real Π or 1 if the queries where answered by the idealized system. Let $\mathbf{S}_0$ the system of the adversary interacting with the real system and $\mathbf{S}_1$ the random system. Further, let $\mathsf{mu\text{-}game}$ the multi-user security game corresponding to $\mathsf{su\text{-}game}$, where each user has its own key but the same primitive P. In $\mathsf{mu\text{-}game}$, the adversary has access to the same oracles as in $\mathsf{su\text{-}game}$ (e.g. for encryption and decryption), but each oracle has an additional input to select the index of the user whose key is used. A more detailed description of admissible security games is given by Hoang and Tessaro [HT16]. Let random variable S capture the key of Π (which we sample uniformly both in $\mathbf{S}_0$ and $\mathbf{S}_1$). $\mathsf{p}_{\mathbf{S}}(\tau, s)$ denotes the probability that $S = s$ and $\mathbf{S}$ behaves according to τ.

Lemma 2 (Lemma 2, Definition 1 and Lemma 3 in [HT16]). *If for any transcript τ for which $\mathsf{p}_{\mathbf{S}_1}(\tau) > 0$ there exists a partition $\Gamma_{\mathsf{good}}, \Gamma_{\mathsf{bad}}$ of the range $\mathcal{U}$ of S, as well as a function $g : \mathcal{U} \to [0, \infty)$ such that $\Pr[S \in \Gamma_{\mathsf{bad}}] \leq \delta$ and for all $s \in \Gamma_{\mathsf{good}}$,*

$$1 - \frac{\mathsf{p}_{\mathbf{S}_0}(\tau, s)}{\mathsf{p}_{\mathbf{S}_1}(\tau, s)} \leq g(s).$$

Then, for $\epsilon(p, q) := \delta + \mathbb{E}[g(S)]$,

$$\begin{aligned} \mathsf{Adv}_{\Pi}^{\text{su-game}}(p, q) \leq & \epsilon(p, q), \\ \mathsf{Adv}_{\Pi}^{\text{mu-game}}(p, q) \leq & 2\epsilon(p + qm, q). \end{aligned}$$

3 The GIEM Construction

We now describe our r-rounds b-branches GIEM (Generalized Iterated Even-Mansour) construction, where $r \geq 2$ is even, and $b \geq 1$.

Definition 1. *We define*

$$\mathbb{I} = \{(0,1), \ldots, (0,r/2), (1,r/2+1), \ldots, (1,r), (2,r/2+1), \ldots, (b,r)\},$$
$$\mathbb{J} = \{(0,1), \ldots, (0,r/2), (1,r/2+1), \ldots, (1,r+1), (2,r/2+1), \ldots, (b,r+1)\}.$$

Let $\mathtt{T}$ denote the tweakspace. For all $i \in \mathbb{I}$, let $P_i : \{0,1\}^n \to \{0,1\}^n$ be a permutation. For all $j \in \mathbb{J}$, let $h_j : \mathtt{T} \to \{0,1\}^n$ be some function, and $\mathbf{H} = (h_j)_{j \in \mathbb{J}}$. $\mathsf{GIEM}^{b,r}_{\mathbf{H}}$ is defined in Figure 2, and an example with $r = 4$ and $b = 2$ is given in Figure 3. When we want to make the internally used permutations explicit we write $\mathsf{GIEM}^{b,r}_{\mathbf{H}}[\mathbf{P}]$, where $\mathbf{P} = (P_i)_{i \in \mathbb{I}}$.

$\mathsf{GIEM}^{b,r}_{\mathbf{H}}(t, m)$

$S \leftarrow m$
for $j = 1$ **to** $r/2$ **do**
  $S \leftarrow P_{0,j}(S \oplus h_{0,j}(t))$
for $i = 1$ **to** b :
  $S_i \leftarrow S$
  for $j = r/2+1$ **to** r **do**
    $S_i \leftarrow P_{i,j}(S \oplus h_{i,j}(t))$
  $S_i \leftarrow S_i \oplus h_{i,r+1}(t)$
return $S_1 || \ldots || S_b$

$(\mathsf{GIEM}^{-1})^{b,r}_{\mathbf{H}}(t, c, \beta)$

$S \leftarrow c \oplus h_{\beta,r+1}(t)$
for $j = r$ **to** $r/2+1$ **do**
  $S \leftarrow P^{-1}_{\beta,j}(S) \oplus h_{\beta,j}(t)$
for $j = r/2$ **to** 1 **do**
  $S \leftarrow P^{-1}_{0,j}(S) \oplus h_{0,j}(t)$
$Y_1 || \ldots || Y_b \leftarrow \mathsf{GIEM}^{b,r}_{\mathbf{H}}.\mathsf{Enc}(t, S)$
return $S || Y_1 || \ldots || Y_{\beta-1} || Y_{\beta+1} || \ldots || Y_b$

Fig. 2. $\mathsf{GIEM}^{b,r}_{\mathbf{H}}$ algorithms.

Related Constructions. If $b = 2$ and $\mathbf{H} = (h_{0,1}, h_{0,2} \oplus h_{0,1}, \ldots, h_{b,r-1} \oplus h_{b,r}, h_{b,r})$ where all h_j are sampled randomly and independently from a uniform AXU-family, then $\mathsf{GIEM}^{2,r}_{\mathbf{H}}$ is equivalent to FTEM by Kim et al. [KLL20]. Contrary to the idea of Kim et al., we fix the branching point to after the $r/2$-th permutation, instead of allowing the common initial part to have a different length than the branches. The reason is that the security appears to be limited by the shortest path from the message to one ciphertext block or from one ciphertext block to another. Our proof strategy requires having at least $r/2$ permutations after the branching point.

If $b = 1$ instead, then $\mathsf{GIEM}^{1,r}_{\mathbf{H}}$ is equivalent to TEM by Cogliati et al. [CLS15]. If $b = 1$, $|\mathtt{T}| = 1$ and $\mathbf{H} = (k_{0,1}, \ldots k_{1,r+1})$ where each k_i is a random value from $\{0,1\}^n$ (or more precisely, and to match our notation, a constant function pointing to this value) then $\mathsf{GIEM}^{1,r}_{\mathbf{H}}$ is equivalent to the original IEM by Bogdanov et al. [BKL+12].

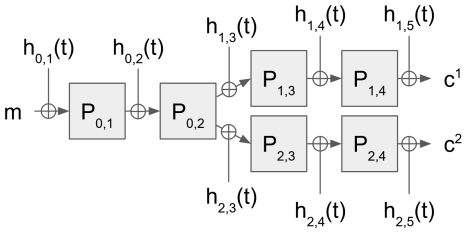

Fig. 3. $\mathsf{GIEM}^{b,r}_{\mathbf{H}}$ with $r = 4, b = 2$.

4 FEM and FTEM-ITS

We describe instantiations FEM and FTEM-ITS of GIEM. Let $\mathsf{FEM} = \mathsf{GIEM}^{b,r}_{\mathbf{H}}$, with $b = 2$ and each $h_{i,j} \in \mathbf{H}$ being a randomly sampled constant function, i.e. $\forall x : h_{i,j}(x) = k_{i,j}$ for some $k_{i,j} \in \{0,1\}^n$. Let the tweakspace $\mathtt{T} = \{0\}$. Since all queries must use the same tweak 0, FEM acts as a forkcipher with randomly sampled round keys and no tweaks, making it the 2-branch forked variant of IEM.

Let $\mathsf{FTEM\text{-}ITS} = \mathsf{GIEM}^{b,r}_{\mathbf{H}}$ with $b = 2$ and each $h \in \mathbf{H}$ being a randomly sampled function from an arbitrary tweakspace to $\{0,1\}^n$. Thus FTEM models a forkcipher with ideal tweakey schedule, giving independent round keys for each tweak.

Theorem 1. *For any adversary $\mathcal{A}$ making at most $q \leq N/2^{r+2}$ queries to each permutation and the cipher, we have*

$$\mathsf{Adv}^{\mathrm{prtmfp}}_{\mathsf{FEM}}(q) \leq 2 \cdot 10^r q^{r+1}/N^r,$$
$$\mathsf{Adv}^{\text{mu-prtmfp}}_{\mathsf{FEM}}(q) \leq (3r/2+1)^{r+1} \cdot 4 \cdot 10^r q^{r+1}/N^r.$$

The proof is provided in the next section.

Relation to tweaked forkciphers. Since the multi-user security assumes independent keys for each user, our above bound for multi-user security of untweaked forkciphers also applies to the (single-user) security of forkciphers where the tweakey schedule is assumed to be ideal. Concretely,

$$\mathsf{Adv}^{\mathrm{prtmfp}}_{\mathsf{FTEM\text{-}ITS}}(q) \leq (3r/2+1)^{r+1} \cdot 4 \cdot 10^r q^{r+1}/N^r\,.$$

Bound interpretation. First, note that the restriction to $q \leq N/2^{r+2}$ has little impact, for more queries q and AES-like parameters ($n = 128, r = 10$) the right-hand sides of the inequalities are larger than 1 anyways. The single-user advantage of FEM in Theorem 1 is similar to the bound for IEM of Hoang and Tessaro (Theorem 1 in [HT16]) with q primitive queries, except having $2 \cdot 10^r$ instead of 4^r as factor. Asymptotically, both bounds give security up to $q \approx N^{r/(r+1)}$ queries. The concrete security is also similar, an adversary attacking FEM with q queries

to the cipher and each permutation has less advantage than attacking IEM with $4q$ queries. Our multi-user security of FEM is obtained from the expectation method and holds for any number of users. Comparing multi-user to single-user security, our FEM bound has a similar loss as is the case for IEM (Theorem 2 in [HT16]) based on the number of permutations, which is larger in FEM, i.e. $3r/2$. Again, the asymptotic security of both is $q \approx N^{r/(r+1)}$ and attacking FEM in the multi-user setting with q queries has less advantage than the corresponding IEM bound [HT16] with $8q$ queries.[4] Thus the single-user and multi-user security of FEM is similar to that of regular IEM.

For FTEM-ITS there is no construction that is the direct equivalent without forking. The only tweakable IEM construction, TEM by Cogliati et al. [CLS15], has security up to $q \approx N^{\frac{r}{r+2}}$ queries. For FTEM-ITS we have security up to $q \approx N^{\frac{r}{r+1}}$ queries, which is better. However TEM is based on a specific key schedule with AXU functions, which is a weaker assumption than what we use for FTEM-ITS. Compared to regular IEM, the single-user security of FTEM-ITS for q queries is better than the single-user security of IEM [CS14, HT16] with $8rq$ queries, while FTEM-ITS provides additional tweaking and forking capabilities.

4.1 Proof of Theorem 1

We follow the IEM security proof by Hoang and Tessaro [HT16] in the expectation method (Lemma 2), adapted to our forked IEM structure. Let $\mathbf{S}_0, \mathbf{S}_1$ be the systems associated to the real and ideal game for prtmfp security. $\mathbb{I}, \mathbb{J}$ were defined in Definition 1. Transcripts τ contain two types of entries: For any $(i,j) \in \mathbb{I}$, $y = P_{i,j}(x)$ and $x = P_{i,j}^{-1}(y)$ queries result in entry $(\mathsf{Perm}, i, j, x, y)$, $\boldsymbol{c} = \mathsf{Enc}(m)$ and Dec queries result in entry $(\mathsf{Enc}, m, \boldsymbol{c})$, where $\boldsymbol{c} = (c^1, c^2)$ contains the ciphertext blocks of branch 1 and 2. (To maintain the same format, for forkcipher decryption queries of a ciphertext block at index $\beta \in \{1,2\}$, put the query input in the position of c^β and the output blocks in the position of the message and the other ciphertext block.) We note that the direction of the queries does not impact the probabilities $\mathsf{p}_{\mathbf{S}_0}(\tau), \mathsf{p}_{\mathbf{S}_1}(\tau)$ of a transcript occurring in $\mathbf{S}_0$ or $\mathbf{S}_1$. Let $f = r/2$, i.e. the final round after which the construction forks into two branches. In the above definitions, i corresponds to the branch (0 for the initial part before forking) and j to the round. For notational convenience we will also use indexing as described in the following definition.

Notation 1. *Throughout this proof, whenever we use an index (i,j) with $j \leq f$ and $i \in \{1,2\}$ it should be treated as $(0,j)$.*

Thus $P_{0,1} = P_{1,1} = P_{2,1}$, and we can always specify the permutation "to the left" of permutation $P_{i,j}$ as $P_{i,j-1}$.

[4] In our multi-user security bound, the factor $(3r/2+1)^{r+1}q^{r+1}$ matches the $(p+qt)^t$ term in Theorem 2 in [HT16] (for $t = r$ permutations and $p = q$ queries), except for having $3r/2$ rather than r permutations.

Overview, defining $G(s)$. Define $N = 2^n$. We have $q \leq N/4$ from the assumption in the theorem statement. Fix a transcript τ. We will use the expectation method (Lemma 2). Let S be the random variable for the key of FEM in $\mathbf{S}_0$. Let $\mathcal{K}$ be the keyspace, then S is uniformly distributed over $\mathcal{K}$. For each key $s = (L_j)_{j \in \mathbb{J}} \in \mathcal{K}$, define graph $G(s)$ as follows: Let $\mathbb{I}_0 = \mathbb{I} \cup \{(0,0)\}$. Its set of vertices is partitioned into sets $V_\gamma, \gamma \in \mathbb{I}_0$ each of 2^n elements which are labelled $(\gamma, 0) \ldots (\gamma, 2^n - 1)$. For each entry $(\mathsf{Perm}, i, j, x, y)$ in τ, add an edge between $((i, j-1), x \oplus L_{j-1})$ and $((i,j), y)$. For a path F in $G(s)$, let $|F|$ be the number of edges in F. An example of such a graph is given in Figure 4.

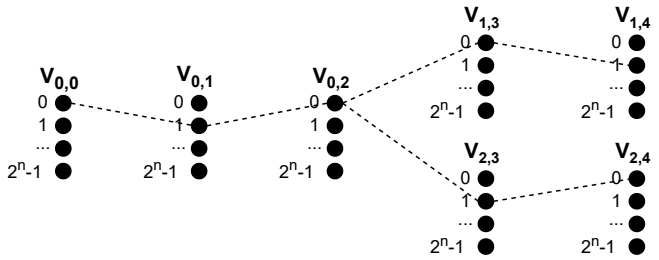

Fig. 4. Graph $G(s)$ for FEM with $r = 4$. Here the transcript τ contains exactly the Perm queries that would be needed to evaluate the encryption of $m = 0$, the edges thus form a "forked" path.
In the example, τ contains queries $(\mathsf{Perm}, 0, 1, 0, 1), (\mathsf{Perm}, 0, 2, 5, 0), (\mathsf{Perm}, 1, 3, 0, 0), (\mathsf{Perm}, 1, 4, 0, 1), (\mathsf{Perm}, 2, 3, 2, 1)$ *and* $(\mathsf{Perm}, 2, 4, 4, 0)$. $s_{0,1} = 0, s_{0,2} = 4, s_{1,3} = 0, s_{1,4} = 0, s_{2,3} = 2, s_{2,4} = 5$. *($s_{1,5}$ and $s_{2,5}$ are irrelevant for the graph.)*

Following the expectation method, we now classify the keys as good or bad. Intuitively, for bad keys the adversary made the necessary Perm queries to verify if encrypting message m can result in $\boldsymbol{c}$ for at least one $(\mathsf{Enc}, m, \boldsymbol{c})$ in τ.

Definition 2 (Good and bad keys). *We say a key s is* bad *if τ contains an entry $(\mathsf{Enc}, m, \boldsymbol{c})$ such that in $G(s)$, there are paths F_0, F_1 starting from two distinct elements in $\{(0,0,m), (1, r, c^1 \oplus L_{1,r+1}), (2, r, c^2 \oplus L_{2,r+1})\}$ and some $\gamma \in \mathbb{I}_0$ such that F_0 contains some vertex in V_γ and F_1 contains some vertex in V_γ. If a key is not bad then we say it is* good. *Let Γ_{bad} be the set of bad keys, and $\Gamma_{\mathsf{good}} = \mathcal{K} \setminus \Gamma_{\mathsf{bad}}$.*

We will show

$$\Pr[S \in \Gamma_{\mathsf{bad}}] \leq \frac{1}{2} \cdot 10^r \cdot \frac{q^{r+1}}{N^r}\,. \tag{1}$$

To see this, let $S = (S_j)_{j \in \mathbb{J}}$. First treat the case that F_0 starts from $(0,0,m)$ and F_1 in $(1, r, c_1 \oplus L_r)$. $S \in \Gamma_{\mathsf{bad}}$ means that τ contains entries

$$(\mathsf{Enc}, m, \boldsymbol{c}), (\mathsf{Perm}, 1, 1, u_1, v_1), \ldots, (\mathsf{Perm}, 1, r, u_r, v_r)$$

(indexing as in Notation 1) such that one of the following happens:

1. $u_1 = m \oplus S_{1,1}$ and $u_j = v_{j-1} \oplus S_{1,j}$ for every $j \in \{2, \ldots, r\}$, or
2. $v_r = c^i \oplus S_{1,r+1}$ and $u_j = v_{j-1} \oplus S_{1,j}$ for every $j \in \{2, \ldots, r\}$, or
3. $u_1 = m \oplus S_{1,1}$, $v_r = c^i \oplus S_{1,r+1}$ and there is an $\ell \in \{2, \ldots, r\}$ such that $u_j = v_{j-1} \oplus S_{1,j}$ for every $j \in \{2, \ldots, r\} \setminus \{\ell\}$.

There are q choices for the query $(\mathsf{Enc}, m, \boldsymbol{c})$ in τ and q choices for $(\mathsf{Perm}, 1, j, u_j, v_j)$ for each j, i.e. q^{r+1} choices in total for the list of queries. For each choice, there is only one $S_{1,j}$ for each $j \in \{1, \ldots, r\}$ to fulfill condition (1). Since all $S_{i,j}$ are uniform in $\{0,1\}^n$, the chance for condition (1) is at most q^{r+1}/N^r. For condition (2) and (3) for any ℓ the same is true, giving $(r+1)q^{r+1}/N^r$ in total. For the other 2 pairs of elements in $\{(0,0,m), (1,r,c_1 \oplus L_r), (2,r,c_2 \oplus L_r)\}$, the bound is analogous, resulting in $\Pr[S \in \Gamma_{\mathsf{bad}}] \leq 3(r+1)q^{r+1}/N^r \leq \frac{1}{2} \cdot 10^r \cdot \frac{q^{r+1}}{N^r}$.

Lemma 3. *There is a non-negative function $g : \mathcal{K} \to [0, \infty)$ such that for any $s \in \Gamma_{\mathsf{good}}$, it holds that*

$$1 - \frac{\mathsf{ps}_0(\tau, s)}{\mathsf{ps}_1(\tau, s)} \leq g(s) \quad \text{and} \quad \mathbb{E}[g(s)] \leq \frac{3}{2} \cdot 10^r \cdot \frac{q^{r+1}}{N^r}.$$

The proof is given in Section 4.2. Using Lemma 2 with Lemma 3 and Equation (1) proves Theorem 1.

4.2 Proof of Lemma 3

Defining $\alpha, \beta, Z^i_s(a,b)$. Fix some $s = (L_j)_{j \in \mathbb{J}}$. Let $Z^i_s(a,b)$ be the number of paths from vertices $V_{i,a}$ to vertices in $V_{i,b}$ of $G(s)$. Let the Enc entries of τ be $(\mathsf{Enc}, m_1, \boldsymbol{c}_1) \ldots (\mathsf{Enc}, m_q, \boldsymbol{c}_q)$. For $k \in \{1, \ldots, q\}$ let $\alpha_k[s]$ be the length of the longest path starting from $(0,0,m_k)$ and for $i \in \{1,2\}$, let $r - \beta^i_k[s]$ be the length of the longest path ending at $(i, r, c^i_k \oplus L_{i,r+1})$. Note that if s is good then $\alpha_k[s] \leq \beta^i_k[s]$ for every $k \in \{1, \ldots, q\}, i \in \{1,2\}$.

Defining $G_k, \mathsf{PathSample}$. Let G_0 be $G(s)$. For each $k \in \{1, \ldots, q\}$, let G_k be defined from G_{k-1} by running the following procedure, which we refer to as $\mathsf{PathSample}_k$. Intuitively, $\mathsf{PathSample}_k$ simulates the evaluation of the cipher on input m_k, represented as a forkpath through the graph G_k.

Definition 3. $\mathsf{PathSample}_k$*: Let $z_{0,0} \leftarrow m_k \oplus L_{0,1}$, $z_{0,1} \leftarrow P_{0,1}(z_{0,0})$ and for each $(i,j) \in \mathbb{I} \setminus \{(0,1)\}$, let $z_{i,j} \leftarrow P_{i,j}(z_{i,j-1} \oplus L_{i,j})$. For all $(i,j) \in \mathbb{I}$ connect vertices $((i,j-1), z_{i,j-1})$ and $((i,j), z_{i,j})$ if this edge is not yet in graph G_{k-1}.*

Defining $m_k \to^* \boldsymbol{c}_k, m_k \to^i \boldsymbol{c}_k$. Fix some $k \in \{1, \ldots, q\}$. Let G be a graph in the support of random variable G_{k-1}. We define a forkpath F as a triple of paths F_0, F_1, F_2 that have a common endpoint in $V_{0,r/2}$ and otherwise each path contains only vertices from branch $0, 1$ or 2, respectively. F is said to connect vertices x and y if $x, y \in F_0 \cup F_1 \cup F_2$. We say that G is *well-formed* if there is a forkpath in G connecting $((0,0), m_j)$ to $((r,1), c^1_j \oplus L_{1,r+1})$ and $((r,2), c^2_j \oplus L_{2,r+1})$ for every $j \in \{1, \ldots, k-1\}$. If G is well-formed then let $\Pr[m_k \to^i c^i_k]$ (for $i \in \{1,2\}$) be

the probability that in $\mathbf{S}_0$, if S agrees with s, and the permutation oracles behave according to G for every $j \in \{1, \ldots, k-1\}$, then querying m_k to Enc results in the block c_k^i in branch i. Let $\Pr[m_k \to^* \boldsymbol{c}_k] = \Pr[m_k \to^1 c_k^1 \wedge m_k \to^2 c_k^2]$, i.e. the probability that m_k results in both blocks of $\boldsymbol{c}_k$. Then from the definition of $\mathsf{PathSample}$, $\Pr[m_k \to^i c_k^i]$ is equal to the probability that the corresponding vertices are connected in G_k, i.e. $((0,0), m_k \oplus L_{0,1})$ connected to $((i,r), c_k^i \oplus L_{i,r+1})$.

Defining $\bar{\mathsf{U}}_G^i(a,b)$. Let G^* be the graph obtained from G by deleting the paths connecting $((0,0), x_j)$ to $((r,1), c_j^1)$ and $((0,0), x_j)$ to $((r,2), c_j^2)$ for every $j \in \{1, \ldots, k-1\}$. Let $U_G^i(a,b)$ be the number of paths F from vertices in $V_{i,a}$ to $V_{i,b}$ in G^*, such that there is no vertex in $V_{i,a-1}$ connected to the first vertex in F. Let $\bar{\mathsf{U}}_G^i(a,b)$ be the vertices of such paths F in $V_{i,a}$ (i.e. the endpoints of these paths). We have

$$U_G^i(a,b) \le Z_s^i(a,b), \tag{2}$$

since all paths contributing to $U_G^i(a,b)$ also contribute to $Z_s^i(a,b)$.

Preparing the expectation method. We turn to establishing a bound on $1 - \frac{\mathsf{p}_{\mathsf{s}_0}(\tau,s)}{\mathsf{p}_{\mathsf{s}_1}(\tau,s)}$.

Lemma 4. *We have*

$$1 - \frac{\mathsf{p}_{\mathsf{s}_0}(\tau,s)}{\mathsf{p}_{\mathsf{s}_1}(\tau,s)} \le \sum_{k=1}^{q} \left(1 - (N-k+1)^2 \Pr[m_k \to^* \boldsymbol{c}_k]\right) .$$

The proof is given in Appendix A.2, and follows similar steps to Hoang and Tessaro [HT16]. To continue, we define

$$\Theta_k = \Pr[m_k \to^* \boldsymbol{c}_k] - \Pr[m_k \to^1 c_k^1] \Pr[m_k \to^2 c_k^2] .$$

Intuitively, this is the difference between receiving the correct ciphertext from FEM (which has a single set of permutations before forking) versus two IEM instances that use independent permutations (from the set of permutations that are consistent with the permutation queries in the transcript). Then,

$$\Pr[m_k \to^* \boldsymbol{c}_k] = \Pr[m_k \to^1 c_k^1] \Pr[m_k \to^2 c_k^2] + \Theta_k . \tag{3}$$

We will bound $\Pr[m_k \to^i c_k^i]$ using existing results on (non-forked) IEM. The main challenge lies in establishing a lower bound for Θ_k. Note that event $m_k \to^* \boldsymbol{c}_k$ happens when the forkpath F that is created from following the permutations P connects m_k to $c_k^1 \oplus L_{1,r+1}$ and $c_k^2 \oplus L_{2,r+1}$. Since $m_k \to^* \boldsymbol{c}_k$ is conditioned on the permutations agreeing with the Perm queries, F must contain the paths associated to $\alpha_k[s], \beta_k^1[s], \beta_k^2[s]$. Consider some graph G with $a = \alpha_k[s], b_1 = \beta_k^1[s], b_2 = \beta_k^2[s]$. Since we fixed the transcript τ and the key s, the remaining randomness lies in the unqueried permutation inputs. In the case that $a > f$, the path F must contain the path associated to $\alpha_k[s]$. $m_k \to^1 c_k^1$ therefore only depends on sampling the permutations of branch 1, and $m_k \to^2 c_k^2$

only on permutations in branch 2, making them independent. Thus $\Theta_k = 0$. In the case that $f \geq b_1$, F again contains the path for $b_1 = \beta_k^1[s]$, $m_k \rightarrow^1 c_k^1$ now depends only on permutations in branch 0 and $m_k \rightarrow^2 c_k^2$ only on permutations in branch 2, again $\Theta_k = 0$. The case of $f \geq b_2$ is analogous. The rest of this proof is on proving a bound for Θ_k in general, which is worse than $\Theta_k = 0$. Thus from this point on we restrict our focus to

$$a \leq f < \min\{b_1, b_2\}\,.$$

Defining $u_{i,j}, w_{i,j}$. Fix any $k \in \{1, \ldots, q\}$ and $a = \alpha_k[s], b_1 = \beta_k^1[s], b_2 = \beta_k^2[s]$. We now define an alternative algorithm $\mathsf{PathSample}'_k$ to obtain G_k from G_{k-1}, and will show that it is equivalent to $\mathsf{PathSample}_k$. Let $\mathbb{M} = \{(0, a+1), \ldots, (0, f), (1, f+1), \ldots, (1, b_1), (2, f+1), \ldots, (2, b_2)\}$. For $(i, j) \in \mathbb{M}$ we define $u_{i,j}$ to be a randomly sampled vertex among the vertices in $V_{i,j}$, that do not have an edge to $V_{i,j-1}$. Then define $w_{0,a}$ as the endpoint of the path associated to $\alpha_k[s]$ in $V_{0,a}$ (recall our assumption $a \leq f$). For $(i, j) \in \mathbb{M}$, if there is an edge between $w_{i,j-1}$ and some vertex $v \in V_{i,j}$, then $w_{i,j} := v$. Otherwise, $w_{i,j} := u_{i,j}$. Note the vertices w_i define a forkpath F that has the same distribution as the original sampling algorithm, the first condition captures when $P_{i,j}$ was queried on the relevant input already, and the second condition captures lazy sampling of the unqueried permutation inputs. Similar to $\mathsf{PathSample}_k$, the edges of F are added to the graph G_{k-1} to obtain G_k.

Defining $\odot, \otimes, \Delta$. For $(i, x) \in \mathbb{M}, y \in \{x+1, \ldots, r\}$, let $\odot_{xy}^i = \Pr[u_{i,x} \in \bar{\mathsf{U}}_{x,y}^i]$ and $\otimes_{xy}^i = \Pr[w_{i,x} \in \bar{\mathsf{U}}_{x,y}^i]$. Intuitively, $\odot_{xy}^i$ is the probability that $\mathsf{PathSample}'$ "tries" to select a path from $V_{i,x}$ to $V_{i,y}$, whereas $\otimes_{xy}^i$ is the probability that such a path is in fact selected.[5] We sometimes write $\odot_{x,y}^i$ instead of $\odot_{xy}^i$ for clarity. For pairs $X = (x_1, x_2), Y = (y_1, y_2)$, we define the following symbols.

$$\begin{aligned}
\bar{\odot}_{XY} &= \Pr[u_{1,x_1} \in \bar{\mathsf{U}}_{x_1,y_1}^1 \wedge u_{2,x_2} \in \bar{\mathsf{U}}_{x_2,y_2}^2] \\
\bar{\otimes}_{XY} &= \Pr[w_{1,x_1} \in \bar{\mathsf{U}}_{x_1,y_1}^1 \wedge w_{2,x_2} \in \bar{\mathsf{U}}_{x_2,y_2}^2] \\
\dot{\odot}_{XY} &= \odot_{x_1,y_1}^1 \odot_{x_2,y_2}^2 \\
\dot{\otimes}_{XY} &= \otimes_{x_1,y_1}^1 \otimes_{x_2,y_2}^2 \\
\Delta_{XY} &= \bar{\otimes}_{XY} - \dot{\otimes}_{XY}
\end{aligned}$$

To give an intuition, one may interpret $\bar{\odot}_{XY}, \bar{\otimes}_{XY}$ as related to our forked graph, unlike the single-branch $\odot_{xy}^i, \otimes_{xy}^i$. Then $\dot{\odot}_{XY}, \dot{\otimes}_{XY}$ could be thought of as relating to two IEM instances with independent permutations. Indeed, if for $i \in \{1, 2\}$, $x \leq f < y_i$, then $u_{1,x} = u_{2,x} = u_{0,x}$ and thus for $X = (x, x), Y = (y_1, y_2)$, $\bar{\odot}_{XY} > 0$ requires the existence of a joint path through $V_{0,x}, \ldots, V_{0,f}$ which then splits into paths for each branch, whereas $\dot{\odot}_{XY} > 0$ is possible if there are two separate paths from a vertex in $V_{0,x}$ to V_{1,y_1} and V_{2,y_2}, respectively.

[5] This notation resembles the one by Chen and Steinberger [CS14] in their Lemma 1 proof. However they defined $\odot, \otimes$ as the corresponding events, i.e. our $\odot_{xy}^1$ is equal to $\Pr[\odot_{xy}]$ in their notation. They use unconventional arithmetic notation $(+, -, \ldots)$ for operations on these events, we chose our different definition to avoid this.

Characterizations and defining M, B, R. Recall we fixed $a = \alpha_k[s], b_1 = \beta_k^1[s], b_2 = \beta_k^2[s]$. For natural number x, define $[x] = \{a+1, \ldots, x-1\}$. For $X = (x_1, x_2), Y = (y_1, y_2)$ define $X \prec Y \Leftrightarrow x_1 < y_1 \wedge x_2 < y_2, X \preceq Y \Leftrightarrow x_1 \leq y_1, x_2 \leq y_2$. For notational convenience, we consider G with additional vertices $V_{i,r+1}$ for $i \in \{1,2\}$, with a single edge from c_k^i in $V_{i,r}$ to vertex $((i, r+1), 0)$. Let $M := N-k+1$. Then for all $\gamma \in \mathbb{I}_0$, we have $M = |V_\gamma|$ in G^*. Let $B = (b_1, b_2)$. Let $R = (r+1, r+1)$. Let $X = (x_1, x_2) \in [b_1+1] \times [b_2+1], Y = (y_1, y_2)$ with $X \prec Y \preceq R$. We now give a list of equations to characterize our $\odot, \otimes$ expressions, and prove all of them below. The following first equations allow us to translate our prior probabilities.

$$\Pr[m_k \to^i c_k^i] = \otimes_{b_i, r+1}^i \tag{4}$$

$$\Pr[m_k \to^* \boldsymbol{c}_k] = \bar{\otimes}_{BR} \tag{5}$$

The next equations will help us to compute the concrete values for $\otimes_{b_i,r+1}^i, \bar{\otimes}_{BR}$ by deconstructing the $\otimes$ expressions step-by-step.

$$\otimes_{xy}^i = \odot_{xy}^i (1 - \sum_{j \in [x]} \otimes_{jx}^i) \tag{6}$$

$$\dot{\otimes}_{XY} = \dot{\odot}_{XY} \left(1 - \sum_{i \in \{1,2\}} \sum_{j \in [x_i]} \otimes_{j,x_i}^i + \sum_{J \in [x_1] \times [x_2]} \dot{\otimes}_{JX} \right) \tag{7}$$

$$\bar{\otimes}_{XY} = \bar{\odot}_{XY} \left(1 - \sum_{i \in \{1,2\}} \sum_{j \in [x_i]} \otimes_{j,x_i}^i + \sum_{J \in [x_1] \times [x_2]} \bar{\otimes}_{JX} \right) \tag{8}$$

$$\bar{\otimes}_{XY} \leq \bar{\odot}_{XY} \tag{9}$$

$$\dot{\odot}_{XY} = \bar{\odot}_{XY} \quad \text{if } \max X > f \tag{10}$$

$$\dot{\odot}_{XY} \leq 4 \frac{Z_s^1(x_1, y_1) Z_s^2(x_2, y_2)}{N^2} \tag{11}$$

$$\dot{\odot}_{BR} \leq \frac{2}{M^2} \tag{12}$$

We justify the equations in Appendix A.2.

Continuing the proof. For the above characterizations, we fixed some arbitrary k. To continue the proof, we make this fixed k explicit by writing $\dot{\odot}_{BR}[k]$. Comparing the definitions, we find $\Theta_k = \Delta_{BR}[k]$ due to Equations (4) and (5). Using Lemma 4 and eq. (3) with Equation (4) results in

$$1 - \frac{\mathsf{ps}_0(\tau, s)}{\mathsf{ps}_1(\tau, s)} \leq g(s) \quad \text{with} \quad g(s) = \sum_{k=1}^{q} \left[1 - (N-k+1)^2 (\dot{\otimes}_{BR}[k] + \Delta_{BR}[k]) \right]. \tag{13}$$

To apply the expectation method, what is left is to compute an upper bound for $\mathbb{E}[g(S)]$. From the linearity of expectation,

$$\mathbb{E}[g(S)] = \sum_{k=1}^{q} \left[1 - (N-k+1)^2 (\mathbb{E}[\dot{\otimes}_{BR}[k]] + \mathbb{E}[\Delta_{BR}[k]]) \right]. \tag{14}$$

Bounding $\dot{\otimes}_{BR}[k]$. We turn to $\dot{\otimes}_{BR}[k] = \Pr[m_k \to^1 c_k^1] \cdot \Pr[m_k \to^2 c_k^2]$.

Lemma 5. *We have*

$$\mathbb{E}[\dot{\otimes}_{BR}[k]] \geq \frac{1}{(N-k+1)^2}\left(1 - 2\frac{4^r q^r}{N^r}\right).$$

The proof is given in Appendix A.2 and utilizes results by Chen and Steinberger [CS14] and Hoang and Tessaro [HT16].

Bounding $\Delta_{BR}[k]$. We now turn to $\Delta_{BR}[k]$. For notational convenience, we fix an arbitrary $k \in \{1, \ldots, q\}$ for now and omit the $[k]$. Let $X = (x_1, x_2) \in [b_1 + 1] \times [b_2 + 1], Y = (y_1, y_2)$ with $X \prec Y \preceq R$. By Equations (7) and (8),

$$\begin{aligned}
\Delta_{XY} &= \bar{\otimes}_{XY} - \dot{\otimes}_{XY} \\
&= \bar{\odot}_{XY}\left(1 - \sum_{i\in\{1,2\}}\sum_{w\in[x_i]} \otimes^i_{w,x_i} + \sum_{W\in[x_1]\times[x_2]} \bar{\otimes}_{WX}\right) \\
&\quad - \dot{\odot}_{XY}\left(1 - \sum_{i\in\{1,2\}}\sum_{w\in[x_i]} \otimes^i_{w,x_i} + \sum_{W\in[x_1]\times[x_2]} \dot{\otimes}_{WX}\right).
\end{aligned}$$

We make two observations. (1) In general, $\Delta_{XY} \geq -\dot{\otimes}_{XY} \geq -\dot{\odot}_{XY}$ (by Equation (9)). (2) Additionally, if $\max X > f$ then $\bar{\odot}_{XY} = \dot{\odot}_{XY}$ (Equation (10)) and

$$\Delta_{XY} = \dot{\odot}_{XY} \sum_{W\in[x_1]\times[x_2]} (\bar{\otimes}_{WX} - \dot{\otimes}_{WX}) = \dot{\odot}_{XY} \sum_{W\in[x_1]\times[x_2]} \Delta_{WX}.$$

We utilize these two rules to lower bound Δ. Define

$$\tilde{\Delta}_{XY} := \begin{cases} -\dot{\odot}_{XY} & \text{if } \max X \leq f, \\ \dot{\odot}_{XY} \sum_{W\in[x_1]\times[x_2]} \tilde{\Delta}_{WX} & \text{else.} \end{cases}$$

$\tilde{\Delta}$ is well defined, since applying the "else" case lowers the first index, i.e. $W \prec X$, meaning that eventually the first case is applied. We have $\Delta_{XY} \geq \tilde{\Delta}_{XY}$. The explicit form for $\tilde{\Delta}_{XY}$ is quite unwieldy, so instead we show several properties to characterize $\tilde{\Delta}_{XY}$. Firstly, there is a set $\mathfrak{S}$ which contains sets $\sigma \in \mathfrak{S}$, which in turn contain pairs $(C, D) \in \sigma$, such that $C = (c_1, c_2), D = (d_1, d_2)$ are pairs, and

$$\tilde{\Delta}_{XY} = -\sum_{\sigma\in\mathfrak{S}} \prod_{(C,D)\in\sigma} \dot{\odot}_{CD}.$$

To see that this set exists, note that applying the $\tilde{\Delta}$ definition to unwrap $\tilde{\Delta}_{XY}$ repeatedly multiplies $\dot{\odot}$ expressions, with one final negative $-\dot{\odot}$ expression, making any $\tilde{\Delta}_{XY}$ negative. Furthermore, after writing out all sum expressions explicitly and applying the distributive law, results in a sum of $\dot{\odot} \cdot \dot{\odot} \cdots \cdot \dot{\odot}$ products.

Lemma 6. *For any X, Y, there is a set $\mathfrak{S}$ such that*

$$\tilde{\Delta}_{XY} = -\sum_{\sigma\in\mathfrak{S}} \prod_{(C,D)\in\sigma} \dot{\odot}_{CD}$$

where all $\sigma \in \mathfrak{S}$ are of the form $\sigma = \{(C_1, C_2), (C_2, C_3), \ldots (C_{\ell-1}, C_\ell)\}$ (for some $\ell, C_1, \ldots, C_\ell$) and

$$\forall i : C_i \prec C_{i+1}, \tag{15}$$
$$\max C_1 \le f, \ \max C_2 > f \tag{16}$$
$$C_{\ell-1} = X, \ C_\ell = Y \tag{17}$$
$$|\sigma| \le \max X - f + 1 \tag{18}$$
$$|\mathfrak{S}| \le f 2^{x_1+x_2-f-1} \tag{19}$$

Proof. To see that σ follows the form $(C_1, C_2), (C_2, C_3), \ldots$, note that the second case in the $\tilde{\Delta}_{XY}$ definition always uses X as the second index for the inner $\tilde{\Delta}_{WX}$ expressions. Equation (15) is due to taking $W \in [x_1] \times [x_2]$, Equation (16) since unfolding $\tilde{\Delta}$ only stops with the first case of the $\tilde{\Delta}$ definition and Equation (17) is justified because $\tilde{\Delta}_{XY}$ always resolves to a $\dot{\odot}_{XY}$ expression when applying the definition of $\tilde{\Delta}$ the first time. For Equation (18), note that for any $\sigma \in \mathfrak{S}$, Equation (15) implies that each C_i must have both elements smaller than C_{i+1}. Now $\max C_{\ell-1} = \max X$, $\max C_{\ell-2} \le \max X - 1, \ldots$, in general $\max C_{\ell-1-z} \le \max X - z$. Thus $\max C_{\ell-1-(\max X - f)} \le f$ if it exists and from Equation (16) $C_{\ell-1-(\max X-f)} = C_1$. Comparing the C indices gives $1 \ge \ell - 1 - (\max X - f)$, hence $\ell - 1 \le 1 + \max X - f$. Since $|\sigma| = \ell - 1$ this proves the Equation (18). We will use a similar argument for Equation (19). Let $\Pi_1((a, b)) = a, \Pi_2((a, b)) = b$. Fix any $\sigma \in \mathfrak{S}$, let C_i the sets as in Lemma 6. From Equation (16), consider $\Pi_1(C_1) \le f < \Pi_1(C_2)$ as case (1). Let $E_1 = \{f + 1, \ldots, x_1 - 1\}, E_2 = \{1, \ldots, x_2 - 1\}$. Note that $\{\Pi_1(C_2), \ldots, \Pi_1(C_{\ell-2})\} \subseteq E_1$, due to $C_{\ell-1} = X$. Further $\Pi_1(C_1) \in \{1, \ldots, f\}$. Similarly $\{\Pi_2(C_1), \ldots, \Pi_2(C_{\ell-2})\} \subseteq E_2$. Therefore each $\sigma \in \mathfrak{S}$ where case (1) applies can be represented as an element from $E = \{1, \ldots, f\} \times 2^{E_1} \times 2^{E_2}$. With $|E_1| = x_1 - f - 1, |E_2| = x_2 - 1$, we have $|E| \le f 2^{x_1+x_2-f-2}$. Case (2), i.e. $\Pi_2(C_1) \le f < \Pi_2(C_2)$, is analogous. Each such $\sigma \in \mathfrak{S}$ can be mapped into set E' with $|E'| \le f 2^{x_1+x_2-f-2}$. The sum over both cases gives Equation (19). □

We use this bound to derive a lower bound of Δ_{BR}. Since $\mathfrak{S}$ depends on $\beta^* = (\beta^1_k[s], \beta^2_k[s])$, we write it as $\mathfrak{S}(\beta^*)$. When fixing β^* to a specific value $B = (b_1, b_2)$, we write $\mathfrak{S}(B)$. Let $R^i_{B,k}[s] = 1$ if $\beta^1_k[s] \le b_1$ and $\beta^2_k[s] \le b_2$, and let $R^i_{a,b,k}[s] = 0$ otherwise. Let $\bar{\mathsf{B}} = \{f + 1, \ldots, r\} \times \{f + 1, \ldots, r\}$. From $\Delta_{BR} \ge \tilde{\Delta}_{BR} \ge -\sum_{\sigma\in\mathfrak{S}(\beta^*)} \prod_{(C,D)\in\sigma} \dot{\odot}_{CD}$,

$$\Delta_{BR} \ge -\frac{2}{M^2} \sum_{\sigma\in\mathfrak{S}(\beta^*)} \prod_{(C,D)\in\sigma\setminus\{(B,R)\}} 4\frac{Z^1_s(c_1, d_1) Z^2_s(c_2, d_2)}{N^2} \tag{20}$$
$$\ge -\frac{2}{M^2} \sum_{B\in\bar{\mathsf{B}}} R_{B,k}[s] \cdot \sum_{\sigma\in\mathfrak{S}(B)} \prod_{(C,D)\in\sigma\setminus\{(B,R)\}} 4\frac{Z^1_s(c_1, d_1) Z^2_s(c_2, d_2)}{N^2} \tag{21}$$

where the first line follows from Equation (11), the fact that $(B, R) \in \sigma$ (Equation (17)) allowing us to factor out $\dot{\odot}_{BR}$ and apply Equation (12). For the second line it suffices that $R^i_{B,k}[s] = 1$ if $\beta^1_k[s] = b_1$ and $\beta^2_k[s] = b_2$. For any $c < c' < d$,

$$Z^i_s(c, d) \leq Z^i_s(c', d) \tag{22}$$

because any path that is counted by $Z_s(c, d)$ also contributes to $Z^i_s(c, k), Z^i_s(k, d)$. Take any $\sigma \in \mathfrak{S}(B)$ from Equation (21). Let $(C_1, \ldots, C_\ell)$ according to Lemma 6, i.e. $\sigma = \{(C_1, C_2), \ldots\}$. Let $(c^1_1, c^2_1) := C_1, (c^1_2, c^2_2) := C_2$. From Equation (16), there is an $i \in \{1, 2\}$ s.t. $c^i_1 \leq f < c^i_2$. Without loss of generality, assume $i = 2$. We will now *prune* branch 2, specifically we create $\phi(\sigma)$ as copy of $\sigma \setminus \{(B, R)\}$, except that $c^2_1 := f$. We then use ϕ inside Equation (21). Due to Equation (22),

$$\Delta_{BR} \geq -\frac{2}{M^2} \sum_{B \in \bar{\mathsf{B}}} R_{B,k}[s] \cdot \sum_{\sigma \in \mathfrak{S}(B)} \prod_{(C,D) \in \phi(\sigma)} 4 \frac{Z^1_s(c_1, d_1) Z^2_s(c_2, d_2)}{N^2} \tag{23}$$

As a result, for any fixed $\sigma \in \mathfrak{S}(B)$, the $\prod_{(C,D) \in \phi(\sigma)} 4 \frac{Z^1_s(c_1, d_1) Z^2_s(c_2, d_2)}{N^2}$ expression only contains independent random variables $Z^i_s(c_i, d_i)$. This is because the Perm queries are fixed in the transcript and thus $Z^i_s(c_i, d_i)$ depends only on the subkeys $s_{i,c_i+1}, \ldots, s_{i,d_1-1}$ ($d_i - c_i - 1$ subkeys). We have

$$\mathbb{E}[Z^i_s(c_i, d_i)] = q^{d_i - c_i} / N^{d_i - c_i - 1} \tag{24}$$

since there are q edges in $G(s)$ that can start the path from V_{i,c_i} and each will be connected to one of the q subsequent edges if the subkey is one of the q corresponding keys out of N keys. Due to pruning one of the branches for ϕ, the relevant subkeys for different $Z^i_s(c_i, d_i)$ expressions do not overlap. Further, for any fixed B, $R_{B,k}[s]$ is independent from the $Z^i_s(c_i, d_i)$, as $R_{B,k}$ depends on subkeys $(i, b_i + 1), \ldots, (i, r + 1)$ for $i \in \{1, 2\}$, which are not part of any $\phi(\sigma)$ due to removing (B, R). Since $R_{B,k}[s]$ depends on $2r - b_1 - b_2$ subkeys,

$$\mathbb{E}[R_{B,k}[s]] = (q/N)^{2r - b_1 - b_2}\,. \tag{25}$$

Starting from Equation (23),

$$\begin{aligned}
\mathbb{E}[\Delta_{BR}] &\geq -\frac{2}{M^2} \sum_{B \in \bar{\mathsf{B}}} \mathbb{E}\left[R_{B,k}[s]\right] \cdot \mathbb{E}\left[\sum_{\sigma \in \mathfrak{S}(B)} \prod_{(C,D) \in \phi(\sigma)} 4 \frac{Z^1_s(c_1, d_1) Z^2_s(c_2, d_2)}{N^2}\right] \\
&\geq -\frac{2}{M^2} \sum_{B \in \bar{\mathsf{B}}} \mathbb{E}\left[R_{B,k}[s]\right] \cdot \sum_{\sigma \in \mathfrak{S}(B)} \prod_{(C,D) \in \phi(\sigma)} 4 \frac{\mathbb{E}\left[Z^1_s(c_1, d_1)\right] \mathbb{E}\left[Z^2_s(c_2, d_2)\right]}{N^2}
\end{aligned}$$

due to linearity of expectation, and $R_{B,k}[s]$ and the Z_s depending on different subkeys. Now fix some $\sigma \in \mathfrak{S}(B)$, and let $(C_1, \ldots, C_\ell) := \phi(\sigma)$.

$$\prod_{(C,D) \in \phi(\sigma)} 4 \frac{\mathbb{E}\left[Z^1_s(c_1, d_1)\right] \mathbb{E}\left[Z^2_s(c_2, d_2)\right]}{N^2} = 4^{|\phi(\sigma)|} \prod_{(C,D) \in \phi(\sigma)} \frac{\mathbb{E}\left[Z^1_s(c_1, d_1)\right] \mathbb{E}\left[Z^2_s(c_2, d_2)\right]}{N^2}$$

Without loss of generality, assume we pruned branch 2 for ϕ. Either (a) $C_1 = (f, f)$ or (b) $C_1 = (h, f)$ with $h < f$ (Equation (16)). In case (a), we have $|\phi(\sigma)| \leq \min B - f$. Using Equation (24), since the $C_1 = (f, f)$ and $C_\ell = (b_1, b_2)$,

$$4^{|\phi(\sigma)|} \prod_{(C,D)\in\phi(\sigma)} \frac{\mathbb{E}\left[Z_s^1(c_1, d_1)\right] \mathbb{E}\left[Z_s^2(c_2, d_2)\right]}{N^2} \leq 4^{\min B - f}(q/N)^{b_1+b_2-2f}$$

In case (b), using Equations (18) and (24),

$$4^{|\phi(\sigma)|} \prod_{(C,D)\in\phi(\sigma)} \frac{\mathbb{E}\left[Z_s^1(c_1, d_1)\right] \mathbb{E}\left[Z_s^2(c_2, d_2)\right]}{N^2} \leq 4^{\max B - f + 1}(q/N)^{b_1+b_2-f-h}\,.$$

Using $h \leq f - 1$ and $\max B - f + 1 \leq (\min B - f) + f + 1$,

$$4^{\max B - f + 1}(q/N)^{b_1+b_2-f-h} \leq 4^{\min B - f}(q/N)^{b_1+b_2-2f}(4^{f+1}q/N)\,.$$

From assuming $4^{f+1}q/N \leq 1$, both case (a) and (b) have the same bound. Thus,

$$\mathbb{E}[\Delta_{BR}] \geq -\frac{2}{M^2} \sum_{B\in\bar{\mathsf{B}}} \mathbb{E}\left[R_{B,k}[s]\right] \cdot \sum_{\sigma\in\mathfrak{S}(B)} 4^{\min B - f}(q/N)^{b_1+b_2-2f}$$

Then,

$$\begin{aligned}
\mathbb{E}[\Delta_{BR}] &\geq -\frac{2}{M^2} \sum_{B\in\bar{\mathsf{B}}} (q/N)^{2r-b_1-b_2} \cdot \sum_{\sigma\in\mathfrak{S}(B)} 4^{\min B - f}(q/N)^{b_1+b_2-2f} \\
&\geq -\frac{2}{M^2} \sum_{B\in\bar{\mathsf{B}}} (q/N)^{2r-b_1-b_2} \cdot f 2^{b_1+b_2-f-1} 4^{\min B - f}(q/N)^{b_1+b_2-r} \\
&\geq -\frac{2}{M^2} \sum_{B\in\bar{\mathsf{B}}} (q/N)^{r} \cdot f 2^{5r/2-1} \geq -\frac{1}{M^2}(r/2)^3 2^{5r/2}(q/N)^r \\
\mathbb{E}[\Delta_{BR}] &\geq -\frac{1}{(N-k+1)^2} 10^r (q/N)^r
\end{aligned}$$

where the first line is obtained using Equation (25), the second using Equation (19), the third from $\min B, b_1, b_2 \leq r, f = r/2, |B| = (r/2)^2$ and the final line from the definition of M. Thus with Equation (14) and lemma 5, and $M_k = N - k + 1$,

$$\begin{aligned}
\mathbb{E}[g(S)] &\leq \sum_{k=1}^{q} \left[1 - (N-k+1)^2 (\mathbb{E}[\dot{\otimes}_{BR}[k]] + \mathbb{E}[\Delta_{BR}[k]])\right] \\
&\leq \sum_{k=1}^{q} \left[1 - M_k^2 \left(\frac{1}{M_k^2}(1 - 2\frac{4^r q^r}{N^r}) - \frac{1}{M_k^2} 10^r (q/N)^r\right)\right] \\
\mathbb{E}[g(S)] &\leq \frac{3}{2} \cdot 10^r \cdot \frac{q^{r+1}}{N^r}
\end{aligned}$$

This proves Lemma 3.

5 The MFTEM Construction

We define MFTEM (multi-forked tweakable iterated Even-Mansour) as a specific instantiation of GIEM. It is the b-branch generalization of FTEM by Kim et al. [KLL20] (which has 2 branches) and TEM by Cogliati et al. [CLS15] (which has 1 branch).

Let $\mathbf{H} = (h_{0,1}, \ldots, h_{b,r})$ where all h_j are functions from $\mathtt{T}$ to $\{0,1\}^n$. Let $\mathbf{H'} = (h_{0,1}, h_{0,1} \oplus h_{0,2}, \ldots, h_{b,r-1} \oplus h_{b,r}, h_{b,r})$. Then define $\mathsf{MFTEM}^{b,r}_{\mathbf{H}}[\mathbf{P}] := \mathsf{GIEM}^{b,r}_{\mathbf{H'}}[\mathbf{P}]$.

5.1 Security of MFTEM

Let q_e denote the number of queries by the adversary to the cipher and q the number of queries to each oracle.

Theorem 2. *Let* $\mathsf{MFTEM}^{b,r}_{\mathbf{H}}$ *as defined above and* $r \geq 2$ *be an even number. Let all* $h_i \in \mathbf{H}$ *be sampled from a uniform and* ϵ*-AXU function family.*

$$\mathsf{Adv}^{\mathrm{prtmfp}}_{\mathsf{MFTEM}^b_r}(q_e, q) \leq 2(b+1)\sqrt{q(2q\epsilon + \frac{2q_e}{N})^{r/2}}$$

Hence if we set $q = q_e$ and $\mathbf{H}$ to an ϵ-AXU function family with $\epsilon \approx \frac{1}{N}$ then MFTEM achieves security up to roughly $q \approx \frac{N^{\frac{r}{r+2}}}{b^2}$ queries.

Proof Sketch. We use a strategy similar to Cogliati et al. [CLS15] (specifically their Theorem 4), and are able to reuse some of their lemmas. Thus we only sketch the proof idea here. Our full proof is given in Appendix A.3.

The r round MFTEM can be seen as $r/2$ rounds of TEM, and each branch then as inverse call of the corresponding $r/2$ round TEM. We reuse Lemma 10 by Cogliati et al. [CLS15], which gives us the distribution of the internal state of MFTEM at the branching point when the inputs from the transcript are used. Similarly for the ciphertext blocks in the transcript, we obtain the distribution of previous internal MFTEM states. From this we derive a set of internal states, that have a high probability of (a) occurring from the input and (b) resulting in the ciphertext blocks from the transcript, and thereby probability of the transcript occurring in the real world.

6 Attacks against IEM

We show that the attack by Bogdanov et al. (Section 3.1 in [BKL+12]) is ineffective, i.e. succeeds with probability less than $\frac{1}{N/2}$ when using the specified number of queries. For completeness, Bogdanov et al. noted: *"To get a better reduction on key-candidates, a bit more [...] queries are sufficient"*. We argue that the way they presented their attack could at least make readers assume that the attack already works decently for q as specified. Indeed, the follow-up literature [LPS12, Ste12, CS14, CLS15] treated it as such, and we suspect it was

also what the authors assumed. Furthermore, only slightly increasing the number of queries, say by adding a constant, does not fix the problem. It remains unclear if multiplying the number of queries by a constant solves the problem (and which constant would be needed), due to the lack of a proof.

We instead point to a similar attack by Gaži [Gaž13] (Section 5). Both attacks make random queries at the start and then iterate over all possible keys to output the first candidate. The main difference is that the Bogdanov et al. attack accepts the first key candidate where no inconsistencies were found, which includes keys for which we cannot evaluate even a single message-ciphertext pair. On the other hand the Gaži attack only accepts a key if it has a minimum number of verifiable message-ciphertext pairs. This minimum number is given as a parameter. To achieve our desired asymptotic query complexity and success probability, we need to use slightly more queries than the Bogdanov et al. attack (a multiplicative factor depending on the number of rounds). Gaži provides a security proof, which aids to preserve the asymptotic tightness of IEM proofs.

6.1 Attack by Bogdanov et al.

Below we first describe the attack by Bogdanov et al. [BKL$^+$12]. We denote $N = 2^n$. For consistency with the notation by Bogdanov et al., we use t as the number of rounds in the construction, rather than r as in our constructions.

Definitions. IEM is defined as $E_k = P_t(\dots P_1(m \oplus k_0) \oplus k_1 \dots) \oplus k_t$ for random public permutations $P_1, \dots, P_t$. We define the following predicates.

- $\mathsf{Connect}$: For key k, queried message m and integer $r, 1 \leq r \leq t$, $\mathsf{Connect}(m, r)$ is true if and only if we have made queries (see attack below) of the form

  $$c := E(m), y_1 := P_1(m \oplus k_0), y_2 := P_2(y_1 \oplus k_1), ..., y_r := P_r(y_{r-1} \oplus k_{r-1}).$$

  In other words, there are some $c, y_1, y_2, \dots$ s.t. $(m, c) \in \mathcal{M}, (m \oplus k_0, y_1) \in \mathcal{P}_1, \dots$ (see set definitions in the attack below).
- Path: For key k and queried message m, $\mathsf{Path}(k, m) :\Leftrightarrow \mathsf{Connect}(k, m, t)$ (i.e. all relevant queries including to the final permutation have been made).
- $\mathsf{ConsistentPath}$: For key k and queried message m, $\mathsf{ConsistentPath}(k, m) :\Leftrightarrow \mathsf{Path}(k, m) \wedge y_t \oplus k_t = c$ (i.e. the cipher query result matches the permutation query results).
- $\mathsf{InconsistentPath}$: For key k and queried message m, $\mathsf{InconsistentPath}(k, m) :\Leftrightarrow \mathsf{Path}(k, m) \wedge y_t \oplus k_t \neq c$.

IEM Attack.

1. For each oracle $P_1, \dots P_t$ and E, make q random queries. (Bogdanov et al. [BKL$^+$12] use $q = N^{t/(t+1)}$.) The results are stored in $\mathcal{P}_i, \mathcal{M}$ as input output pairs, e.g. if $(x, y) \in \mathcal{P}_1$, then this means we queried $P_1(x)$ and received y.
2. For each key candidate $k = (k_0, k_1, \dots, k_t) \in \{0,1\}^{n(t+1)}$ do:
   (a) For each message m: If $\mathsf{InconsistentPath}(k, m)$: Continue to the next key candidate (since this key is definitely wrong).
   (b) If for no message $\mathsf{InconsistentPath}(k, m)$ holds, output k as the key guess.

Analysis on the IEM Attack. Note that in the IEM attack the first key candidate without an inconsistent path is output. Given the number of such candidates, the probability that the attack succeeds is thus 1 divided by the number of candidates.

Lemma 7. *The above attack, when executed with $q = N^{t/(t+1)}$, has probability of less than $\frac{1}{N/2}$ of success.*

Proof. For all k_0, let M_{k_0} denote the set of messages $m, \exists c : (m, c) \in \mathcal{M}$, for which we have $\exists (x, y) \in \mathcal{P}_1 : m \oplus k_0 = x$. Note that if any $M_{k_0} = \emptyset$, then any key k containing partial key k_0 will not have $\mathsf{Connect}(k, m, 1)$ for any message m, let alone have $\mathsf{Path}(k, m)$ or $\mathsf{InconsistentPath}(k, m)$ for any message m. Thus, the attack would give at least N^t candidates.

Observe that any m from $\mathcal{M}$ will be in M_{k_0} for exactly q different k_0 ($k_0 = m \oplus x$ for each $(x, y) \in \mathcal{P}_1$). Thus, the q elements of $\mathcal{M}$ cause a total of q^2 elements spread over some or all M_{k_0}. How they are distributed depends on the randomly selected queries of step 1 of the attack. If they are distributed evenly, then every M_{k_0} contains q^2/N elements, since $k_0 \in \{0,1\}^n$. In any case, there are at least $N/2$ sets M_{k_0} which have strictly less than $2q^2/N$ elements. We will show that each of these keys k_0 can be extended into (at least) one key candidate, which proves the lemma.

Let k_0 s.t. $|M_{k_0}| < 2q^2/N$ elements. By a similar argument as above, there is a subkey k_1, such that strictly fewer than $2\frac{q^2}{N}\frac{q}{N}$ messages m have $\mathsf{Connect}(k, m, 2)$ by a key k containing k_0, k_1. Repeating this for each of t permutations, we have a partial key $k = (k_0, \ldots, k_{t-1})$, where there are less than $2\frac{q^{t+1}}{N^t}$ messages m for which $\mathsf{Connect}(k, m, t)$ is true. Thus, for fewer than $2\frac{q^{t+1}}{N^t} = 2$ messages i.e. at most for one message m we have $\mathsf{Path}(k, m)$. If no such m exists, any final subkey $k_t \in \{0,1\}^n$ can be used to complete the key $k^* = (k_0, \ldots, k_t)$ without having $\mathsf{InconsistentPath}(k^*, m)$. On the other hand, if a single such m exists, where $(m, c) \in \mathcal{M}$, then we can set $k_t = y_t \oplus c$ which results in $\mathsf{ConsistentPath}(k^*, m)$. □

6.2 Instantiating Gaži's attack.

To apply Gaži's attack [Gaž13] (Section 5) to IEM, fix $k_1 := 1, \ldots, k_\ell := \ell$ (independent of k'), which effectively extracts several public random permutations from the ideal cipher. We define k' as being comprised of all IEM round keys $k' = (k'_0, \ldots k'_\ell)$ and define $Q_{i,k'}(x) := x \oplus k'_i$. To asymptotically match the desired query number, we use the trade-off that Gaži described at the end of Section 5.

7 Future Work

In this work we applied the IEM paradigm to study forkciphers and multiforkciphers. Future work may consider alternative expanding primitives, such as the 8-branch tweakable expanding pseudorandom function Butterknife [ACL+24].

Acknowledgements

Elena Andreeva and Andreas Weninger are supported in part by the Austrian Science Fund (FWF) SFB project SPyCoDe 10.55776/F85. Amit Singh Bhati was supported by CyberSecurity Research Flanders with reference number VR20192203, in part by the Research Council KU Leuven C1 on Security and Privacy for Cyber-Physical Systems and the Internet of Things with contract number C16/15/058 and by the Flemish Government through FWO Project G.0835.16 A security Architecture for IoT.

References

ABPV21. Elena Andreeva, Amit Singh Bhati, Bart Preneel, and Damian Vizár. 1, 2, 3, fork: Counter mode variants based on a generalized forkcipher. *IACR Trans. Symm. Cryptol.*, 2021(3):1–35, 2021.

ABV21. Elena Andreeva, Amit Singh Bhati, and Damian Vizár. Nonce-misuse security of the SAEF authenticated encryption mode. In *Selected Areas in Cryptography: 27th International Conference, Halifax, NS, Canada (Virtual Event), October 21-23, 2020, Revised Selected Papers 27*, pages 512–534. Springer, 2021.

ACL+24. Elena Andreeva, Benoit Cogliati, Virginie Lallemand, Marine Minier, Antoon Purnal, and Arnab Roy. Masked iterate-fork-iterate: a new design paradigm for tweakable expanding pseudorandom function. In *International Conference on Applied Cryptography and Network Security*, pages 433–459. Springer, 2024.

ADP+20. Elena Andreeva, Arne Deprez, Jowan Pittevils, Arnab Roy, Amit Singh Bhati, and Damian Vizár. New results and insighs on forkae. In *NIST LWC workshop*, 2020.

ALP+19. Elena Andreeva, Virginie Lallemand, Antoon Purnal, Reza Reyhanitabar, Arnab Roy, and Damian Vizár. Forkcipher: A new primitive for authenticated encryption of very short messages. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part II*, volume 11922 of *LNCS*, pages 153–182, Kobe, Japan, December 8–12, 2019. Springer, Heidelberg, Germany.

AW23. Elena Andreeva and Andreas Weninger. A forkcipher-based pseudo-random number generator. In Mehdi Tibouchi and XiaoFeng Wang, editors, *Applied Cryptography and Network Security*, pages 3–31, Cham, 2023. Springer Nature Switzerland.

BAMV24. Amit Singh Bhati, Elena Andreeva, Simon Müller, and Damian Vizar. Sonikku: Gotta Speed, Keed! A Family of Fast and Secure MACs. Cryptology ePrint Archive, Paper 2024/1980, 2024.

BAV24. Amit Singh Bhati, Elena Andreeva, and Damian Vizár. OAE-RUP: a strong online AEAD security notion and its application to SAEF. In *International Conference on Security and Cryptography for Networks*, pages 117–139. Springer, 2024.

BBDL23. Arghya Bhattacharjee, Ritam Bhaumik, Avijit Dutta, and Eik List. PAE: Towards more efficient and bbb-secure ae from a single public permutation. Cryptology ePrint Archive, Paper 2023/978, 2023. `https://eprint.iacr.org/2023/978`.

BDA+24. Amit Singh Bhati, Antonín Dufka, Elena Andreeva, Arnab Roy, and Bart Preneel. Skye: An Expanding PRF based Fast KDF and its Applications. In *Proceedings of the 19th ACM Asia Conference on Computer and Communications Security*, pages 1082–1098, 2024.

BKL+12. Andrey Bogdanov, Lars R. Knudsen, Gregor Leander, François-Xavier Standaert, John P. Steinberger, and Elmar Tischhauser. Key-alternating ciphers in a provable setting: Encryption using a small number of public permutations - (extended abstract). In David Pointcheval and Thomas Johansson, editors, *EUROCRYPT 2012*, volume 7237 of *LNCS*, pages 45–62, Cambridge, UK, April 15–19, 2012. Springer, Heidelberg, Germany.

BPA+23. Amit Singh Bhati, Erik Pohle, Aysajan Abidin, Elena Andreeva, and Bart Preneel. Let's Go Eevee! A Friendly and Suitable Family of AEAD Modes for IoT-to-Cloud Secure Computation. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 2546–2560, 2023.

BVA24. Amit Singh Bhati, Michiel Verbauwhede, and Elena Andreeva. Breaking, Repairing and Enhancing XCBv2 into the Tweakable Enciphering Mode GEM. *Cryptology ePrint Archive*, 2024. `https://eprint.iacr.org/2024/1554`.

CLS15. Benoit Cogliati, Rodolphe Lampe, and Yannick Seurin. Tweaking Even-Mansour ciphers. In Rosario Gennaro and Matthew J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 189–208, Santa Barbara, CA, USA, August 16–20, 2015. Springer, Heidelberg, Germany.

CS14. Shan Chen and John P. Steinberger. Tight security bounds for key-alternating ciphers. In Phong Q. Nguyen and Elisabeth Oswald, editors, *EUROCRYPT 2014*, volume 8441 of *LNCS*, pages 327–350, Copenhagen, Denmark, May 11–15, 2014. Springer, Heidelberg, Germany.

DGL22. Avijit Dutta, Jian Guo, and Eik List. Forking sums of permutations for optimally secure and highly efficient prfs. Cryptology ePrint Archive, Paper 2022/1609, 2022. `https://eprint.iacr.org/2022/1609`.

EM97. Shimon Even and Yishay Mansour. A construction of a cipher from a single pseudorandom permutation. *Journal of Cryptology*, 10(3):151–162, June 1997.

Gaž13. Peter Gaži. Plain versus randomized cascading-based key-length extension for block ciphers. In Ran Canetti and Juan A. Garay, editors, *CRYPTO 2013, Part I*, volume 8042 of *LNCS*, pages 551–570, Santa Barbara, CA, USA, August 18–22, 2013. Springer, Heidelberg, Germany.

HR03. Shai Halevi and Phillip Rogaway. A tweakable enciphering mode. In Dan Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 482–499, Santa Barbara, CA, USA, August 17–21, 2003. Springer, Heidelberg, Germany.

HR04. Shai Halevi and Phillip Rogaway. A parallelizable enciphering mode. In Tatsuaki Okamoto, editor, *CT-RSA 2004*, volume 2964 of *LNCS*, pages 292–304, San Francisco, CA, USA, February 23–27, 2004. Springer, Heidelberg, Germany.

HT16. Viet Tung Hoang and Stefano Tessaro. Key-alternating ciphers and key-length extension: Exact bounds and multi-user security. In Matthew Robshaw and Jonathan Katz, editors, *CRYPTO 2016, Part I*, volume 9814 of *LNCS*, pages 3–32, Santa Barbara, CA, USA, August 14–18, 2016. Springer, Heidelberg, Germany.

KLL20. Hwigyeom Kim, Yeongmin Lee, and Jooyoung Lee. Forking tweakable Even-Mansour ciphers. *IACR Trans. Symm. Cryptol.*, 2020(4):71–87, 2020.

LPS12. Rodolphe Lampe, Jacques Patarin, and Yannick Seurin. An asymptotically tight security analysis of the iterated Even-Mansour cipher. In Xiaoyun Wang and Kazue Sako, editors, *ASIACRYPT 2012*, volume 7658 of *LNCS*, pages 278–295, Beijing, China, December 2–6, 2012. Springer, Heidelberg, Germany.

LRW02. Moses Liskov, Ronald L. Rivest, and David Wagner. Tweakable block ciphers. In Moti Yung, editor, *CRYPTO 2002*, volume 2442 of *LNCS*, pages 31–46, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Heidelberg, Germany.

LST12. Will Landecker, Thomas Shrimpton, and R. Seth Terashima. Tweakable blockciphers with beyond birthday-bound security. In Reihaneh Safavi-Naini and Ran Canetti, editors, *CRYPTO 2012*, volume 7417 of *LNCS*, pages 14–30, Santa Barbara, CA, USA, August 19–23, 2012. Springer, Heidelberg, Germany.

Pat09. Jacques Patarin. The "coefficients h" technique. In Roberto Maria Avanzi, Liam Keliher, and Francesco Sica, editors, *Selected Areas in Cryptography*, pages 328–345, Berlin, Heidelberg, 2009. Springer Berlin Heidelberg.

RBBK01. Phillip Rogaway, Mihir Bellare, John Black, and Ted Krovetz. OCB: A block-cipher mode of operation for efficient authenticated encryption. In Michael K. Reiter and Pierangela Samarati, editors, *ACM CCS 2001*, pages 196–205, Philadelphia, PA, USA, November 5–8, 2001. ACM Press.

Rog04. Phillip Rogaway. Efficient instantiations of tweakable blockciphers and refinements to modes OCB and PMAC. In Pil Joong Lee, editor, *ASIACRYPT 2004*, volume 3329 of *LNCS*, pages 16–31, Jeju Island, Korea, December 5–9, 2004. Springer, Heidelberg, Germany.

Ste12. John Steinberger. Improved security bounds for key-alternating ciphers via hellinger distance. Cryptology ePrint Archive, Report 2012/481, 2012. `https://eprint.iacr.org/2012/481`.

A Appendix (Supplementary)

A.1 Supporting Definitions

Multiforkcipher and Security A multiforkcipher (MFC) F_b (with b branches) is a pair of deterministic algorithms, the forward $F_b : K \times \mathtt{T} \times \{0,1\}^n \to \{0,1\}^{bn}$, and the backward (or inverse) $F_b^{-1} : K \times \mathtt{T} \times \{0,1\}^n \times [b] \to \{0,1\}^{bn}$.

The forward evaluation algorithm takes a key k, a tweak T and an input block X. It outputs the blocks $Y_1, \ldots, Y_b$. The backward evaluation algorithm F_b^{-1} takes in a key k, a tweak T, a block Y and an input indicator β. It then outputs blocks $X, Y_1, \ldots, Y_{\beta-1}, Y_{\beta+1}, \ldots, Y_b$. We omit b from the notation if it is fixed and clear from the context, and we will also write $F_k(T, X)$ for $F_b(k, T, X)$ (F_k^{-1} analogously).

MFC is *correct* if for every $k \in K, T \in \mathtt{T}, X \in \{0,1\}^n$ and $\beta \in [b]$ it satisfies:

$$F_k^{-1}(T, F_k(T, X)_{[(\beta-1)n+1:\beta n]}, \beta) = X||Z$$

where Z is $F_k(T, X)_{[1:(\beta-1)n]}||F_k(T, X)_{[\beta n+1:bn]}$. We write $F[\mathbf{P}]$ to denote a multiforkcipher F based on a tuple of permutations $\mathbf{P} = (P_i)_{i \in I}$ for some index set I.

The advantage of an adversary $\mathcal{A}$ at distinguishing F_b from a random multi-forked permutation $\tilde{P}$ (c.f. Section [ABPV21]) is defined w.r.t. Figure 5 as

$$\mathsf{Adv}^{\mathbf{prtmfp}}_{F_b}(\mathcal{A}) = |\Pr[\mathcal{A}^{\mathbf{prtmfp\text{-}real}_{F_b}} \Rightarrow 1] - \Pr[\mathcal{A}^{\mathbf{prtmfp\text{-}ideal}_{F_b}} \Rightarrow 1]|.$$

Note that for $b = 1$ branch (thus β fixed to 1) and $|\mathtt{T}| = 1$ this game is equivalent to the pseudorandom permutation (PRP) game as it is used for IEM proofs [BKL+12, CS14].

Game $\mathbf{prtmfp\text{-}real}_{F_b}$	Game $\mathbf{prtmfp\text{-}ideal}_{F_b}$
$k \leftarrow K$	**for** $T \in \mathtt{T}$ **do** $\pi_{T,1}, \ldots, \pi_{T,b} \leftarrow\!\$\ \mathsf{Perm}(n)$
$\mathbf{P} \leftarrow\!\$\ \mathsf{Perm}(n)^{\lvert I\rvert}$	$\mathbf{P} \leftarrow\!\$\ \mathsf{Perm}(n)^{\lvert I\rvert}$
$b \leftarrow \mathcal{A}^{\mathcal{E},\mathcal{D},\mathbf{P}}$	$b \leftarrow \mathcal{A}^{\mathcal{E},\mathcal{D},\mathbf{P}}$
return b	**return** b
Oracle $\mathcal{E}(T, X)$	**Oracle** $\mathcal{E}(T, X)$
return $F_k[\mathbf{P}](T, X)$	**return** $\pi_{T,1}(X)\|\| \ldots \|\|\pi_{T,b}(X)$
Oracle $\mathcal{D}(T, Y, \beta)$	**Oracle** $\mathcal{D}(T, Y, \beta)$
return $F_k^{-1}[\mathbf{P}](T, Y, \beta)$	$X \leftarrow \pi^{-1}_{T,\beta}(Y)$
	return $X\|\|\pi_{T,1}(X)\|\| \ldots \|\|\pi_{T,\beta-1}(X)\|\|$ $\pi_{T,\beta+1}(X)\|\| \ldots \|\|\pi_{T,b}(X)$

Fig. 5. Games $\mathbf{prtmfp\text{-}real}_{F_b}$ and $\mathbf{prtmfp\text{-}ideal}_{F_b}$ defining security of the multifork-cipher F_b. When $\mathbf{P}$ is made available to the adversary, the inverse permutations are implicitly made available as well.

A.2 Details for the Proof of Theorem 1

Proof of Lemma 4 Let Prim_b be the event that S_b behaves according to Prim entries in (τ, s). Then $\Pr[\mathsf{Prim}_0] = \Pr[\mathsf{Prim}_1]$. On the one hand

$$\mathsf{ps}_1(\tau, s) = \frac{\Pr[\mathsf{Prim}_1]}{N^{|\mathbb{J}|} \cdot (N(N-1)\ldots(N-q+1))^2}$$

where $\frac{1}{N^{|\mathbb{J}|}}$ is from hitting the key s and $\frac{1}{(N(N-1)\ldots(N-q+1))^2}$ from both construction branches outputting the correct ciphertexts. On the other hand, for $\mathbf{S}_0$ to behave according to according to (τ, s), it means that (i) S must agree with s, and the system must behave according to Prim entries in (τ, s), and (ii) for every

$k = 1, \ldots, q-1$, if condition (i) holds and querying $m_1, \ldots, m_{k-1}$ to Enc oracles in $\mathbf{S}_0$ results in $\boldsymbol{c}_1, \ldots, \boldsymbol{c}_{k-1}$, respectively, then querying m_k results in $\boldsymbol{c}_k$. Then

$$\mathsf{ps}_{\mathbf{S}_0}(\tau, s) = \frac{\Pr[\mathsf{Prim}_0]}{N^{|\mathbb{J}|}} \prod_{k=1}^{q} \Pr[m_k \to^* \boldsymbol{c}_k].$$

Thus,

$$1 - \frac{\mathsf{ps}_{\mathbf{S}_0}(\tau, s)}{\mathsf{ps}_{\mathbf{S}_1}(\tau, s)} = 1 - \prod_{k=1}^{q} (N-k+1)^2 \Pr[m_k \to^* \boldsymbol{c}_k].$$

We use the fact that for all reals $c_k \in [0;1]$, $\prod_{k=1}^{q}(1-c_k) \geq 1 - \sum_{k=1}^{q} c_k$, and hence $1 - \prod_{k=1}^{q} c_k = 1 - \prod_{k=1}^{q}(1-(1-c_k)) \leq 1 - (1 - \sum_{k=1}^{q}(1-c_k)) = \sum_{k=1}^{q}(1-c_k)$ to obtain the claimed bound.

Proof of Lemma 5

Lemma 8 (Equation 21 in [HT16]). *We have*

$$\Pr[m_k \to^i c_k^i] \geq \frac{1}{N-k+1} \left(1 - \sum_{\sigma \in \mathfrak{B}(\alpha_k[s], \beta_k^i[s])} \prod_{(a,b)\in\sigma} \frac{Z_s^i(a,b)}{N-2q}\right).$$

For $0 \leq a < b \leq r$, let $R_{a,b,k}^i[s] = 1$ if $\alpha_k[s] \geq a$ and $\beta_k^i[s] \leq b$, and let $R_{a,b,k}^i[s] = 0$ otherwise. Similar to Hoang and Tessaro in their proof of Lemma 4 we use $q \leq N/4$, to find

$$\Pr[m_k \to^i c_k^i] \geq \frac{1}{N-k+1} \left(1 - \sum_{0\leq a<b\leq r} R_{a,b,k}^i[s] \cdot \sum_{\sigma\in\mathfrak{B}(a,b)} \prod_{(a,b)\in\sigma} \frac{2Z_s^i(a,b)}{N}\right).$$

To see that the inequality holds, it suffices that $R_{a,b,k}^i[s] = 1$ if $a = \alpha_k[s], b = \beta_k^i[s]$. Then from $(1-c)(1-d) \geq 1-c-d$ for all $c, d \in [0;1]$ and $\dot{\otimes}_{BR}[k] = \Pr[m_k \to^1 c_k^1] \cdot \Pr[m_k \to^2 c_k^2]$,

$$\dot{\otimes}_{BR}[k] \geq \frac{1}{(N-k+1)^2} \left(1 - \sum_{i\in\{1,2\}} \sum_{0\leq a<b\leq r} R_{a,b,k}^i[s] \cdot \sum_{\sigma\in\mathfrak{B}(a,b)} \prod_{(a,b)\in\sigma} \frac{2Z_s^i(a,b)}{N}\right).$$

By linearity of expectation,

$$\mathbb{E}[\dot{\otimes}_{BR}[k]] \geq \frac{1}{(N-k+1)^2} \left(1 - \sum_{i\in\{1,2\}} \mathbb{E}\left[\sum_{0\leq a<b\leq r} R_{a,b,k}^i[s] \cdot \sum_{\sigma\in\mathfrak{B}(a,b)} \prod_{(a,b)\in\sigma} \frac{2Z_s^i(a,b)}{N}\right]\right). \tag{26}$$

Lemma 9 (Lemma 5 in [HT16]). *We have*

$$\mathbb{E}\left[\sum_{0\leq a<b\leq r} R_{a,b,k}^i[s] \cdot \sum_{\sigma\in\mathfrak{B}(a,b)} \prod_{(a,b)\in\sigma} \frac{2Z_s^i(a,b)}{N}\right] \leq \frac{4^r q^r}{N^r}.$$

Using the above lemma and Equation (26) we obtain the claimed bound.

Proof of characterizations We now justify equations 4 to 12. As discussed before, $\mathsf{PathSample}'$ is equivalent to $\mathsf{PathSample}$. To justify Equation (4), note that from $b^i = \beta^i_k[s]$ we know there is a path from some vertex in V_{i,b^i} to $(i, r+1, c^i_k \oplus L_{i,r+1})$, and from there the additional edge to $V_{i,r+2}$ that we added. If the sampled path connects m to $c^i_k \oplus L_{i,r+1}$, then it must follow that path, i.e. w_{i,b^i} must select this path (which is the only path to $V_{i,r+2}$). The analogous is true for Equation (5).

For Equation (6), the event corresponding to $\otimes^i_{xy}$ can only happen if the event corresponding to $\odot^i_{xy}$ happened (i.e. $u_{i,x}$ must be the start of a path to a vertex in $V_{i,y}$) and $w_{i,x}$ was not "hijacked" by a previous path (thus probability $1 - \sum_{j\in[x]} \otimes^i_{jx}$). These two events are indeed independent, since the sampling of $u_{i,x}$ is independent of what happened on previous layers. This result was also used by Chen and Steinberger (Equation 30 in [CS14]). Equation (7) is the result of $\dot{\otimes}_{XY} = \otimes^1_{x_1y_1}\otimes^2_{x_2y_2}$ and multiplying out Equation (6). For Equation (8) a similar argument applies as for Equation (6), now the event corresponding to $\bar{\odot}_{XY}$ is a requirement. In this case, the expression in brackets is the result of the inclusion exclusion principle on the hijacking paths. Specifically, any of the events corresponding to $\otimes^i_{j,x_i}$ makes $\bar{\otimes}_{XY}$ impossible, but the sum over all $\otimes^i_{j,x_i}$ counts the events twice, where any $\otimes^1_{j_1,x_1}$ and $\otimes^2_{j_2,x_2}$ happen at the same time.

Equation (9) is due to any $w_{i,j}$ only being able to choose a vertex from $\bar{U}_{i,j}$ if $w_{i,j} = u_{i,j}$ (otherwise $w_{i,j}$ is the endpoint of an existing edge, which cannot be in $\bar{U}_{i,j}$ per definition of $\bar{U}_{i,j}$). For Equation (10), since $i_1, i_2 > f$, u_{1,i_1} and u_{2,i_2} are independent random variables, so the multiplication rule for independent events applies. For Equation (11), we have

$$\dot{\odot}_{XY} = \odot^1_{x_1y_1}\odot^2_{x_2y_2} = \prod_{i\in\{1,2\}} \frac{U^i_{G_{k-1}}(x_i, y_i)}{M-q} \leq 4\frac{U^1_{G_{k-1}}(x_1, y_1)}{N} \cdot \frac{U^2_{G_{k-1}}(x_2, y_2)}{N}$$

The first equation is the definition of $\dot{\odot}_{XY}$, the second is from the fact that u_{1,x_1} is sampled from the vertices without an edge to the left in G_{k-1}. From the N total vertices, $k-1$ are used from prior paths and up to q additional edges exist due to Perm queries, so in total $M-q$ possible choices. There are $U^i_{G_{k-1}}(x_i, y_i)$ such paths. The inequality follows from $M - q = N - k + 1 - q \geq N - 2q$ and $q \leq N/4$. Together with Equation (2) this proves Equation (11). Equation (12) uses a similar argument, and the fact that only the one edge we considered for notational convenience exists to a vertex from $V_{i,r+2}$ for each $i \in \{1, 2\}$.

A.3 Proof of Theorem 2

Preliminaries. We also identify h_i with k_i. Before starting our proof, we recall some lemmas from Cogliati et al. [CLS15].

Lemma 10 (Lemma 8 in [CLS15]). *Let Ω be some finite event space and μ^* be the uniform probability distribution on Ω. Let μ be a probability distribution on Ω such that $||\mu - \mu^*|| \leq \epsilon$. Then there is a set $S \subset \Omega$ such that:*

– $|S| \geq (1 - \sqrt{\epsilon})|\Omega|$,
– $\forall x \in S, \mu(x) \geq (1 - \sqrt{\epsilon})\mu^*(x)$.

Let $\mathsf{TEM}^r_{\mathbf{H}}[\mathbf{P}] := \mathsf{MFTEM}^{1,r}_{\mathbf{H}}[\mathbf{P}]$. For a given sequence of tweaks $t = (t_1, \ldots, t_{q_e})$, let

$$\Omega_t = \{(x_1, \ldots, x_{q_e}) \in (\{0,1\}^n)^{q_e} | \forall i \neq j : (t_i, x_i) \neq (t_j, x_j)\}.$$

Definition 4. *For a given set of tweaks t and cipher inputs x and any attainable queries transcript $Q_{\mathbf{P}}$, define the tuple of random variables*

$$\mathsf{TEM}^r_{\mathbf{H}}[\mathbf{P}](t, x) := (\mathsf{TEM}^r_{\mathbf{H}}[\mathbf{P}](t_1, x_1), \ldots, \mathsf{TEM}^r_{\mathbf{H}}[\mathbf{P}](t_{q_e}, x_{q_e}))$$

and let μ_{t,x,Q_P} denote the distribution of the tuple conditioned on the event $\mathbf{P} \vdash Q_{\mathbf{P}}$ (i.e. when the key $\mathbf{H}$ is uniformly random and the permutations $\mathbf{P}$ are uniformly random among permutations that are consistent with $Q_{\mathbf{P}}$).

Lemma 11 (Lemma 10 in [CLS15]). *Let μ^*_t denote the uniform distribution on Ω_t. Fix any attainable queries transcript $Q_{\mathbf{P}}$ and any $t \in \mathtt{T}^{q_e}, x \in \Omega_t$*

$$||\mu_{t,x,Q_{\mathbf{P}}} - \mu^*_t|| \leq q_e(2q_e\epsilon + \frac{2q}{N})^r$$

of Theorem 2. We will follow a similar strategy to [CLS15] in order to separate the full r-round forked construction into smaller, single branch parts with $r/2$ rounds each.

We use the H-coefficient technique, and there will be no bad transcripts. Let the query transcript $\tau = (Q_C, Q_{P_{0,1}}, \ldots, Q_{P_{b,r}})$ be arbitrary from all attainable transcripts but fixed. Let $\mathbf{P}_0 = (P_{0,1}, \ldots, P_{0,r/2})$ and for all $0 < i \leq b$ let $\mathbf{P}_i = (P_{i,r/2+1}, \ldots, P_{i,r})$. Define $\mathbf{P}_i^{-1} = (P^{-1}_{i,r}, \ldots, P^{-1}_{i,r/2+1})$. Similarly $\mathbf{H}_0 = (h_{0,1}, \ldots, h_{0,r/2})$ and for all $0 < i \leq b$ let $\mathbf{H}_i = (h_{i,r/2+1}, \ldots, h_{i,r})$ and $\mathbf{H}_i^{-1} = (h_{i,r}, \ldots, h_{i,r/2+1})$.

Let $\gamma_i = (\mathsf{TEM}^{-1})^{r/2}_{\mathbf{H}_i^{-1}}[P_i^{-1}]$. Note that $\mathsf{MFTEM}^{b,r}_{\mathbf{H}}[\mathbf{P}](x) = \gamma_1(z)|| \ldots ||\gamma_b(z)$ where $z = \mathsf{TEM}^{r/2}_{\mathbf{H}_0}(x)$.

From the transcript τ, let $t = (t_1, \ldots, t_{q_e})$ be the tweaks in Q_C, $m = (m_1, \ldots, m_{q_e})$ the messages and for all $1 \leq i \leq b$, $c^i = (c^i_1, \ldots, c^i_{q_e})$ the i-th ciphertext blocks. Let $\mu^0(z) = \mu^0_{t,m,Q_{\mathbf{P}_0}}(z)$ be the distribution of tuple $\mathsf{TEM}^{r/2}_{\mathbf{H}_0}[\mathbf{P}_0](t, m)$ like in Definition 4. Let $\mu^i(z) = \mu^i_{t,c^i,Q^{-1}_{\mathbf{P}_i}}(z)$ be the distribution of tuple $\mathsf{TEM}^{r/2}_{\mathbf{H}_i^{-1}}[\mathbf{P}_i^{-1}](t, c^i)$, where $Q^{-1}_{\mathbf{P}_i}$ is $Q_{\mathbf{P}_i}$ with the elements in the tuples in reverse order (to invert the permutations) and the permutations in reverse order. Intuitively, each of the $\mu^i(z)$, with $z = (z_1, \ldots, z_{q_e})$, gives the "local" probability that, for $1 \leq \ell \leq q_e$, the intermediate state z_ℓ occurred during the ℓ-th execution of MFTEM.

Let $r' = r/2$. We denote $\alpha = 2^{r'} \frac{q_e(N\epsilon q_e + q_p)^{r'}}{N^{r'}}$. We can apply the bound of Lemma 11 to each μ^i. Hence by Lemma 10 for each $0 \leq i \leq b$ there exists a set $S_i \subset \Omega_t$ of size at least $(1 - \sqrt{\alpha})|\Omega_t|$ such that for all $z \in S_i$: $\mu^i(z) \geq \frac{1-\sqrt{\alpha}}{|\Omega_t|}$.

Let $S = S_0 \cap \cdots \cap S_b$. Note that $|S| \geq (1 - (b+1)\sqrt{\alpha})|\Omega_t|$. Recall that $\gamma_i = (\mathsf{TEM}^{-1})^{r/2}_{\mathbf{H}_i^{-1}}[P_i^{-1}]$ and we will again use the tuples m, c^i of all messages and ciphertext blocks. We will also use c^0 to refer to m. Let $\gamma_0 = \mathsf{TEM}^{r/2}_{\mathbf{H}_0}[P_0]$. Since the permutations and keys are uniformly random and independent, we have

$$\begin{aligned}
\Pr[T_{re} = \tau'] &= \sum_{z \in (\{0,1\}^n)^{q_e}} \Pr[\mathbf{P} \vdash Q_{\mathbf{P}} \wedge \gamma_0(m) = z \wedge \bigwedge_{1 \leq i \leq b} \gamma_i(z) = c^i] \\
&\geq \sum_{z \in S} \prod_{i=0}^{b} \Pr[\mathbf{P} \vdash Q_{\mathbf{P}} \wedge \mathsf{TEM}^{r/2}_{\mathbf{H}_i}[\mathbf{P}_i](t, c^i) = z] \\
&\geq \sum_{z \in S} \prod_{i=0}^{b} \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \mu^i(z) \geq \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \sum_{z \in S} \prod_{i=0}^{b} \mu^i(z) \\
&\geq \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \sum_{z \in S} \frac{(1 - \sqrt{\alpha})^{b+1}}{|\Omega_t|^{b+1}} \geq \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \frac{|S|(1 - (b+1)\sqrt{\alpha})}{|\Omega_t|^{b+1}} \\
&\geq \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \frac{(1 - (b+1)\sqrt{\alpha})^2}{|\Omega_t|^b} \geq \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \frac{(1 - 2(b+1)\sqrt{\alpha})}{|\Omega_t|^b}
\end{aligned}$$

Let us turn to $\Pr[T_{id} = \tau']$. Let $\tilde{P}$ be a random multiforked permutation. All (multiforked) permutations are independent, thus we have $\Pr[T_{id} = \tau'] = \Pr[\mathbf{P} \vdash Q_{\mathbf{P}}] \Pr[\tilde{P} \vdash Q_C]$, $\Pr[\tilde{P} \vdash Q_C] = \frac{1}{|(\Omega_t)^b|}$. Finally, this results in $\frac{\Pr[T_{re} = \tau']}{\Pr[T_{id} = \tau']} \geq 1 - 2(b+1)\sqrt{\alpha}$.

□