Forward-Security Under Continual Leakage with Deterministic Key Updates

Suvradip Chakraborty*, Harish Karthikeyan**, Adam O'Neill*** , and C. Pandu Rangan†

Abstract. In the setting of continual-leakage (CL) — Brakerski *et al.*, Dodis *et al.*, FOCS 2010 — the secret key of a cryptographic scheme evolves according to time periods; the adversary gets some bounded leakage function of its choice applied to the current secret key in each time period. This model necessitates a *randomized* key update procedure, which is a major source of difficulty for example in handling leakage on updates. However, the reason why a randomized key update procedure is required is arguably unsatisfying.

Our goal is to provide a general security model for continual leakage with deterministic key updates, and constructions that improve in various respects on prior work. For our basic security model we take an *entropy-based* approach, leading to a model we call *entropic continual leakage* (ECL). In the ECL model, the adversary is allowed to make a limited total number of leakage queries that, as in CL, can depend arbitrarily on other keys, but an *unlimited* total number of what we call "local" leakage queries. The latter does not decrease computational entropy of other keys.

Another benefit of allowing deterministic key updates is that we can readily incorporate forward security (FS) in our constructions, recently pointed out by Bellare *et al.* (CANS 2017) to be an important security hedge in this context. Accordingly, we also introduce the FS+ECL model. We provide FS+ECL constructions of public-key encryption (based on non-interactive key exchange, NIKE) and digital signatures (based on identification schemes). Our constructions improve over the assumptions or leakage rates of prior work. In particular, our result on NIKE improves on the very recent result of Li *et al.* (CRYPTO 2020), going from bounded-leakage to FS+ECL. Of independent interest, even disregarding leakage this gives the first FS-NIKE from indistinguishability obfuscation. Finally, as another result of independent interest, we present a PKE scheme in the FS+CL model (with randomized update) that improves on both the assumptions and leakage rates compared to the scheme of Bellare *et al*.

1 Introduction

Continual Leakage Resilience. When a cryptographic algorithm (or any algorithm) is implemented and run, it must be done on some *physical* system. This introduces *side channel* attacks where the adversary obtains some leakage about secrets, such as execution time, power consumption, and even sound waves [25, 32, 33]. The cryptographic community has responded by extending the attack model, which had previously only considered black-box access to the algorithms, so that in the extended model the adversary gets some "bounded" leakage about the secrets [3, 7, 31, 39]. Because there is no

---

  * Institute of Science and Technology Austria. `suvradip.chakraborty@ist.ac.at`
  ** Dept. of Computer Science, New York University. `karthik@cs.nyu.edu`
 *** Dept. of Computer Science, University of Massachusetts Amherst. `adamo@cs.umass.edu`
  † Dept. of Computer Science, Indian Institute of Technology Madras. `rangan@cse.iitm.ac.in`

inherent reason for the leakage to stay bounded throughout the entire lifetime of the key (if the adversary can get a large fraction of the key, what prevents it from completely recovering it?), works further extended the model to consider "continual" leakage (CL) attacks [9,16]. In this model, the life of a secret key is divided into time periods/epochs, and in time period $t+1$ one runs an update algorithm on the secret key of time period t to derive the new secret key for time period $t+1$. (The old secret key is erased.) In each time period the adversary queries for a function with bounded output length applied to the current secret key.

Motivation For Deterministic Updates. In the continual leakage model, the update algorithm needs to be *randomized*; otherwise, simply consider an adversary that request bits of some future secret key, one by one, in earlier rounds. Randomization seems like a simple enough requirement, but it leads to a number of issues. For one, a challenging problem has been then handling *leakage on updates* (*i.e.*, leakage on the coins of the update algorithm), which has required complex schemes or heavy-weight tools to address this [14,18,34]. Further, it is difficult to integrate with machinery for *forward security* (FS) [5], which is an important hedge against full key exposure in this context [6], since FS schemes typically have deterministic update. In particular, looking ahead it will allow us to use non-tree-based techniques to achieve FS (all constructions in [6] are tree-based), which has various advantages.

However, we observe that the reason above that the update algorithm of a CL resilient scheme requires randomized update is contrived in that the leakage function not only repeatedly computes the update algorithm itself, it leaks an entire future key. We assert that in the real world this would not happen, as exemplified by typical sources of leakage "controlled by nature." Could it not be possible to have deterministic update in the continual leakage setting and achieve reasonable security? What security notion could one aim for?

Prior Work and Our Approach. One approach is to assume that the leakage function lies in a strictly less powerful complexity class than the update function of the scheme. This was previously done for key-evolution schemes by Dziembowski *et al.* [20], who assume space-bounded leakage functions. However, such a requirement is hard to guarantee in practice. For example, timing or cache attacks may be a function of the entire memory. Furthermore, it is complicated in the following sense: If leakage is an arbitrary efficiently computable function of the secret key then leakage from cryptographic operations like signing and encrypting can be captured directly. However, when the leakage function is too weak leakage from such computation has to be captured *separately*.

1.1 The ECL and the FS+ECL Model.

This motivates us to find an approach that more directly captures the above intuition about contrived leakage functions. To this end, we classify leakage functions according to whether, intuitively speaking, they *compute the update function or not*, limiting the amount of the former. Namely, we categorize the leakage functions queried by the adversary into (i) *local* and (ii) *non-local* queries. The above categorization is based on the

property of entropy loss of the keys due to the leakage functions. In particular, we ask whether the outcome of a given leakage query (on one key) reduces the entropy of other keys (which are derived from that key using a (deterministic) key update procedure) or not. Very roughly, we call a leakage query that does not reduce the entropy of other keys "local". Of course, to allow deterministic key updates, such a notion of entropy must be computational, as each leakage query *necessarily* reduces information-theoretic entropy of the other (deterministically defined) keys. The formal definition is given in Definition 8. On the other hand, as the name suggests, non-local queries may potentially reduce the entropy of the other keys. In our corresponding model, which we call the "Entropic" Continual Leakage (ECL) model, there is no restriction on the type or the number of such queries the adversary can make over time. On the other hand, to guarantee security we place the following restrictions on these queries as: (i) The output length of each local leakage query is upper bounded by a parameter $\lambda(\kappa)$, where κ is the security parameter, (ii) the set of all non-local leakage queries made by the adversary should not reduce the entropy of any key beyond $\alpha(\kappa)$, which is also a parameter of the model. This nevertheless allows the leakage to still have some dependence on future keys as a hedge (which is not possible in general with deterministic update functions).

Incorporating FS. As enabled to allowing deterministic updates, all of our constructions actually achieve a much stronger definition that incorporates another hedge, namely forward security (FS). FS was recently studied as a hedge for continual leakage security by Bellare *et al.* [6]. FS refers to the fact that if the full secret key is recovered in some time period then security in prior time periods is maintained. As argued by [6], it is important to consider forward security in this context, since the update may not happen soon enough in some time period to avoid full leakage of the current secret key. To address this, Bellare *et al.* proposed the Forward security under Continual Leakage (FS+CL) model, where the adversary gets bounded leakage on the current secret key in each time period, and FS in the sense that if full key exposure occurs in one time period then executions in previous time periods remain secure. Note that the FS+CL model is actually stronger than either CL or FS individually. We leave it as an interesting open problem to study constructions in the ECL model itself.

1.2 Our Contributions.

In this work, we propose the "Entropic" Continual Leakage (ECL) and Forward security under "Entropic" Continual Leakage (FS+ECL) model and present new constructions of public key encryption (PKE) and digital signature schemes in the FS+ECL model. Besides, we also improve upon the construction of PKE scheme in the FS+CL model (due to [6]) in terms of the required hardness assumptions. We expand on each of these contributions in greater detail below.

The FS+ECL Model. We propose the FS+*E*CL model. The FS+ECL model is very similar to the ECL model (informally discussed above) with the following exception: we require that the security of the underlying primitive be maintained, given the set of

3

local and non-local leakage queries asked by the adversary till the period of break-in/exposure (i.e., when the adversary breaks in and recovers the full secret key of that time period). This is because, after recovering a key in full the adversary can compute the future keys by itself. This model raises some interesting questions: Can one construct schemes with better efficiency or under more standard assumptions in the FS+ECL model compared to the FS+CL constructs of Bellare *et al.* [6]? Intuitively, this should be the case because deterministic updates allow exploiting a host of FS machinery (the FS schemes often have deterministic update) already developed. We show that improvements are possible in the FS+ECL setting for both signatures and encryption.

PKE scheme in the FS+ECL Model. We define and construct key evolving encryption (KEE) scheme in the FS+ECL model. We obtain these results by considering another primitive, namely *non-interactive key exchange* (NIKE) in the FS+ECL model, and then show a generic transformation from a FS+ECL-secure NIKE scheme to a FS+ECL-secure PKE scheme. The resulting PKE scheme inherits the *same* leakage rate as of the underlying NIKE scheme.

We show a construction of FS+ECL-secure NIKE in the common reference string (CRS) model from indistinguishability obfuscation ($i\mathcal{O}$) and either DDH or LWE, having an optimal leakage rate, i.e., $1 - o(1)$. As shown in the recent work of [36], it is *impossible* to construct even bounded leakage-resilient NIKE in the plain model via a black-box reduction under any game-based assumption when the leakage is super-logarithmic in the security parameter. Hence, similar to [36], we consider the CRS model to bypass the above black-box impossibility result. We remark that, before this work, construction of even FS-NIKE was not known from $i\mathcal{O}$. The previously known construction of FS-NIKE by [41] relied on graded multilinear maps. However, the candidate constructions of such maps are prone to various attacks (see [11,12,37]). Moreover, it is not clear how to effectively leverage this construction to achieve leakage-resilience. Similar to [41], our construction of FS+ECL NIKE supports an a-priori *bounded* (but an arbitrary polynomial) number of time periods.

Theorem 1. *(Informal) Assuming indistinguishability obfuscation for circuits and either DDH or LWE, there exists forward-secure entropic leakage-resilient (FS+ECL) NIKE with optimal leakage rate $1 - o(1)$.*

Theorem 2. *(Informal) Assuming a forward-secure entropic leakage-resilient (FS+ECL) NIKE with leakage rate $1 - o(1)$ and strong existentially unforgeable one-time signature scheme, there exists PKE scheme in the FS+ECL model with leakage rate $1 - o(1)$.*

Combining the above with the recent construction of $i\mathcal{O}$ from the circular-secure LWE assumption [8], we obtain a construction of FS+ECL NIKE that is conjectured to be post-quantum secure. Even disregarding leakage, this gives the first post-quantum FS-NIKE.

Theorem 3. *(Informal) Assuming sub-exponential hardness of circular-secure LWE, we obtain the first construction of FS-NIKE that is conjectured post-quantum secure.*

While our FS+ECL NIKE scheme relies on $i\mathcal{O}$ and one-way functions, the reliance on $i\mathcal{O}$ does not seem to be inherent; yet there does not seem to be straightforward way to get a construction without using it.

Signature scheme in the FS+ECL Model. We define and construct key evolving signature (KES) schemes in the FS+ECL model. We observe that it is *impossible* for a signature scheme to meet the standard notion of existential unforgeability against chosen message attack (EUF-CMA) in the FS+ECL setting. Hence, we consider the weaker notion of *entropic unforgeability*, introduced in the context of bounded retrieval model by [4] to define the security of signature schemes in the FS+ECL model.

To this end, we introduce another primitive, namely *identification* (ID) scheme in the FS+ECL model, and then show how to transform such an ID scheme to a FS+ECL-secure signature scheme.

We show a construction of FS+ECL-secure ID scheme from the hardness of the (standard) RSA problem, enjoying optimal leakage rate, i.e., $1-o(1)$. Our ID scheme supports an a-priori bounded (but an arbitrary polynomial) number of time periods.

Theorem 4. *(Informal) Assuming the hardness of the (standard) RSA assumption, there exists Identification scheme secure in the FS+ECL model having optimal leakage rate of $1 - o(1)$.*

Finally, we show how to use the Fiat-Shamir transform to convert an FS+ECL-secure ID scheme to a FS+ECL-secure signature. The resulting signature scheme is secure in the Random Oracle (RO) model and has leakage rate half as that of the underlying FS+ECL-secure ID scheme, i.e., $1/2 - o(1)$.

PKE scheme in the FS+CL Model [6]. Finally, as a result of independent interest, we consider the problem of constructing an FS+CL encryption scheme. Bellare *et al.* [6] showed a construction of FS+CL secure encryption scheme from extractable witness encryption (EWE), which is believed to be a suspect assumption [22]. It was left as an open problem in [6] to construct a FS+CL-secure PKE scheme from standard assumptions. We significantly improve upon the construction of [6] and show how to construct a FS+CL-secure PKE scheme from standard assumptions (namely, *static* assumptions over composite-order bilinear maps). Besides, the leakage rate of our construction is also *optimal*, i.e., $1 - o(1)$, whereas the PKE scheme of [6] could only tolerate a leakage rate of $1/\mathsf{poly}(\kappa)$, where κ is the security parameter.

1.3 Technical Overview.

Modeling FS+ECL-secure NIKE. We observe that constructing a NIKE scheme even in the setting of bounded leakage (letting alone CL or ECL model) is *impossible* in general. This is because, the adversary while leaking from the secret key of one party involved in the NIKE protocol can encode the (description of the) shared-key derivation function along with the public key of the other party hard-coded in the function to leak directly from the shared key. This intuition has been formalized in the recent work of Li et al. [36], where they show that in the public-key setting it is impossible to construct

a bounded leakage-resilient NIKE in the plain model, and prove security via a black-box reduction from any *single-stage game assumption*[1,2]. To bypass the above result, they provide constructions of *bounded* leakage-resilient NIKE in the *common reference string* (CRS) model and in the *preprocessing* model (assuming the existence of leak-free randomness to run the setup and preprocessing phases respectively). Similar to [36], we also turn to constructing FS+ECL secure NIKE in the CRS model.

Constructing FS+ECL-secure NIKE in the CRS Model. The starting point of our construction of FS+ECL-secure NIKE scheme is the recent work of [36] (henceforth called the *LMQW protocol*), who construct bounded leakage-resilient NIKE in the CRS model based on indistinguishability obfuscation ($i\mathcal{O}$).

The main idea of the LMQW construction is as follows: Each user samples a random string s as its secret key and sets its public key as $x = G(s)$, where G is a function whose description is a part of the CRS and can be indistinguishably created in either *lossy* or in *injective* mode. In the real construction, the function G is set to be injective. To generate a shared key with an user j, user i inputs its own key pair (x_i, s_i) and the public key x_j of user j to an obfuscated program $\widehat{C}$ (which is also included as part of the CRS) which works as follows: The circuit C (which is obfuscated) simply checks if s_i is a valid pre-image of either x_i or x_j under G, i.e, it checks if either $x_i = G(s_i)$ or $x_j = G(s_i)$. If so, it runs $\mathsf{PRF}_K(x_i, x_j)$ (where the PRF key K is embedded inside the obfuscated program) and else it outputs $\perp$.

Lifting the LMQW protocol to the FS setting: It is easy to see that the LMQW protocol is *not* forward-secure. This is because, each public key is an injective function of its corresponding secret key, and hence if a secret key s is updated to s', the public key no longer stays the same. We now describe how to modify the above construction to achieve security in the FS+ECL setting. The main idea of our construction is as follows: Similar to the LMQW protocol, the (initial) secret key of each user i in our construction is also a random string $s_i^{(1)}$. The CRS also contains the description of the obfuscated program $\widehat{C}$ and the function G (as described above). However, the public key of each party i is now an obfuscation of a circuit C_i (denoted by $\widehat{C_i}$), which has the initial/root secret key $s_i^{(1)}$ (corresponding to base time period 1) of party i embedded in it. The circuit C_i works as follows: (1) It takes as input the secret key $s_j^{(t)}$ of user j (corresponding to some time period t) and updates the secret key $s_i^{(1)}$ of user i (hard-coded in it) to $s_i^{(t)}$ (corresponding to time period t) by running the (deterministic) NIKE update function (to be defined shortly) $t - 1$ times, (2) computes $x_i^{(t)} = G(s_i^{(t)})$ and $x_j^{(t)} = G(s_j^{(t)})$, and finally (3) internally invokes the obfuscated circuit $\widehat{C}$ (included as part of CRS) on input the tuple $(s_j^{(t)}, x_i^{(t)}, x_j^{(t)})$. To generate the shared key with an user i corresponding to

---

[1] This is a more general class than falsifiable assumptions [38], where the challenger is no longer required to be efficient.

[2] Their impossibility result even rules out the notion of *unpredictable* NIKE in the bounded leakage setting, where the adversary has to predict the entire exchanged key, rather than just distinguish it from uniform.

some time period (say t), user j runs $\widehat{C_i}$ with input its secret key $s_j^{(t)}$ corresponding to time period t to obtain the shared key $\mathsf{PRF}_K(x_i^{(t)}, x_j^{(t)})$. It is easy to see that user i also derives the same shared key for time period t by running the program $\widehat{C_j}$ (public key of user j) on input its own $s_i^{(t)}$ corresponding to time period t.

The security proof of our construction follows the proof technique of the LMQW protocol with some differences. The main idea of the proof of the LMQW protocol follows the punctured programming paradigm [42], where they puncture the PRF key K at the point (x_i, x_j) and program a random output y. However, instead of hard-coding y directly they hard-core $y \oplus s_i$ and $y \oplus s_j$, i.e, the one-time pad encryptions of y under s_i and s_j respectively. This allows the obfuscated program to decrypt y given either s_i or s_j as input. At this point, they switch the function G to be in *lossy* mode and argue that the shared key y retain high min-entropy, even given the obfuscated program with hard-coded ciphertexts, the public keys and leakages on the secret keys s_i and s_j. The entropic key k is then converted into a uniformly random string by using an appropriate extractor.

However, for our construction, we *cannot* argue the last step of the above proof, i.e., the shared key y (for time period t) retains enough entropy given all the public information (CRS and public keys) and entropic leakage on the keys. This is because the public keys $\hat{C}_i$ and $\hat{C}_j$ completely determine the keys $s_i^{(t)}$ and $s_j^{(t)}$ respectively, even after switching the function G to be in lossy mode. This is because the obfuscated programs $\hat{C}_i$ and $\hat{C}_j$ contains the base secret keys $s_i^{(1)}$ and $s_j^{(1)}$ hard-coded in them, and hence, given the public keys the secret keys have no entropy left. To this end, we switch the public key $\hat{C}_i$ to an obfuscation of a functionally equivalent program that, instead of embedding the base secret key $s_i^{(1)}$ embeds all possible public keys $(x_i^{(1)}, \cdots, x_i^{(T)})$ in it, where T is the total number of time period supported by our scheme and $x_i^{(j)} = G(s_i^{(j)})$ for $j \in [T]$. Since, the function G is lossy the shared key y still retains enough entropy, even given the public key. A similar argument can be made for party j. By setting the parameters appropriately, we can prove FS+ECL security of our NIKE construction.

From FS+ECL-secure NIKE to FS+ECL-secure PKE scheme. We provide a generic transformation from a FS+ECL-secure NIKE scheme to a FS+ECL-secure PKE scheme. The transformation follows the blueprint of [21], who showed how to transform a secure NIKE scheme (in an appropriate security model) to a CCA-secure KEM scheme, additionally using one-time signatures. We show how to appropriately modify their proof technique to deal with forward security and entropic continual leakage simultaneously. As a corollary, even disregarding leakage, we get the first non-tree based forward-secure PKE scheme.

Modeling FS+ECL-secure Signatures schemes. Next, we turn to constructing FS+ECL-secure signature schemes. Unfortunately, we observe that one *cannot* achieve the standard notion of *existential unforgeability* for signatures in the FS+ECL setting. This is because, the adversary can use the leakage function on some secret signing key to leak a signature under some future key (for an arbitrary hard-coded choice of random coins

used by the signing algorithm) and then present the leaked signature as a forgery. This is possible since the update algorithm is deterministic. Also note that, such a leakage is valid in the FS+ECL model, since the secret keys still retain enough computational entropy, even given the signatures.

To bypass the above impossibility, we consider the weaker notion of *entropic unforgeability* as defined by [4] in the context of bounded retrieval model (BRM). An attack against entropic unforgeability in the FS+ECL model is valid, if the adversary manages to forge a message (for some time period t) which is chosen from some distribution of significant entropy, *after* the adversary finishes interacting with the leakage oracle.

Constructing FS+ECL Signature Schemes. We construct FS+ECL signature schemes by defining and constructing another notion of *FS+ECL-secure Identification* (ID) schemes.

FS+ECL-secure Identification schemes. We consider three move public-coin ID schemes, in which the prover sends the first message a, the verifier sends a random challenge c, and the prover finally responds with some message z. The transcript of the ID protocol is denoted by (a, c, z). Our definition of FS+ECL-secure ID schemes generalizes the prior definitions of ID schemes, which were either forward-secure [1,5] or only leakage-resilient [4]. In particular, in the FS+ECL security model for ID schemes, the adversary gets oracle access to polynomially many copies of the prover. In addition, it may also learn entropic continual leakage from the secret key of the prover. Finally, after this phase, the adversary enters into an impersonation phase, where it looses access to both the prover and the leakage oracles. On a first thought, it might appear that, similar to signature schemes it may not be possible to construct a secure ID scheme in the FS+ECL model. In particular, the adversary may use the leakage oracle to leak a "forged" transcript of the ID scheme corresponding to some time period (say t) to win the security game against the identification scheme. However, we note that, the above attack does not work for ID schemes. This is because in each run of the protocol the verifier generates a random challenge instance, and hence the probability that the challenge string contained in the leaked transcripts (before the impersonation phase) equals the challenge instance sampled freshly by the verifier in the impersonation phase is negligible.

Our construction of FS+ECL-secure ID scheme can be conceptually viewed as a modular combination of two basic steps: (a) prove the (entropic) leakage-resilience of the generalized Guillou-Quisquater (GQ) identification scheme [15, 24] based on the hardness of the RSA problem, and (b) use a variant of the Itkis-Reyzin (IR) transform to convert the generalized GQ scheme to a forward-secure version of itself. We actually show that the generalized GQ scheme is a secure 3-round public-coin Σ-protocol satisfying two additional properties – (a) each statement (or theorem) has exponentially many witnesses, and (b) the uncertainty of the witness conditioned on the statement is high. These properties help us to prove entropic leakage-resilience.

From FS+ECL ID schemes to FS+ECL Signatures. Finally, we show how to transform a FS+ECL-secure ID scheme to a FS+ECL-secure signature scheme using the generalized Fiat-Shamir (FS) transform. This shows the applicability of the FS transform even in the FS+ECL setting.

FS+CL PKE Scheme from Simpler Assumptions As a result of independent interest, we consider the problem of constructing a FS+CL-secure PKE scheme (as proposed by Bellare *et al.* [6]). As an intermediate step to achieving this, we abstract out a new notion of *continuous leakage-rate binary tree encryption* (CLR-BTE), which is a continuous leakage-resilient analogue of the notion of binary tree encryption (BTE), introduced by Canetti, Halevi and Katz (CHK) [10]. We also observe that the construction of CLR-BTE follows in a straightforward manner from the CLR-HIBE scheme of Lewko et al. [35] (which is based on static assumptions over composite-order bilinear groups). For appropriate choice of parameters, the CLR-HIBE scheme achieves the optimal leakage rate of $1-o(1)$. Finally, we show how to transform a CLR-BTE scheme to a FS+CL-secure PKE scheme using the Canetti-Halevi-Katz (CHK) transform.

This approach of constructing FS+CL secure encryption scheme was already suggested in [6]. However, it was intuitively claimed in [6] that the above approach does not work, due to the following reason: "The problem is that FS+CL security of the resulting scheme requires that multiple nodes of the BTE construction can be leaked on jointly, whereas the CL security of HIBE only buys us leakage on each such node individually." Surprisingly, we prove the contrary and show that, indeed, it is possible to *simulate* the joint leakage by just leaking on a *single* node. This requires a careful analysis of the CHK transform in the FS+CL setting.

1.4 Perspective.

Our work is theoretically motivated, showing feasibility of qualitatively better results in our more realistic FS+ECL model. However, FS+ECL has practical motivation as well, *e.g.* lightweight devices without a good source of randomness for key update. We leave it for future work to explore more practical aspects of the FS+ECL model, and hopefully efficient constructions.

2 Preliminaries

2.1 Notations

Let $x \in \mathcal{X}$ denote an element x in the support of $\mathcal{X}$. For a probability distribution $\mathcal{X}$, let $|\mathcal{X}|$ denote the size of the support of $\mathcal{X}$, i.e., $|\mathcal{X}| = |\{x \mid \Pr[\mathcal{X} = x] > 0\}|$. If x is a string , we denote $|x|$ as the length of x. Let $x \leftarrow \mathcal{X}$ be the process of sampling x from the distribution $\mathcal{X}$. For $n \in \mathbb{N}$, we write $[n] = \{1, 2, \cdots, n\}$. When A is an algorithm, we write $y \leftarrow A(x)$ to denote a run of A on input x and output y; if A is randomized, then y is a random variable and $A(x; r)$ denotes a run of A on input x and randomness r. An algorithm A is probabilistic polynomial-time (PPT) if A is randomized and for any input $x, r \in \{0, 1\}^*$, the computation of $A(x; r)$ terminates in at most $poly(|x|)$ steps. For a set S, we let U_S denote the uniform distribution over S. For an integer $\alpha \in \mathbb{N}$, let U_α denote the uniform distribution over $\{0, 1\}^\alpha$, the bit strings of length α. Throughout this paper, we denote the security parameter by κ, which is implicitly taken as input by all the algorithms. For two random variables X and Y drawn from a finite set $\mathcal{X}$,

let $\delta(X, Y) = \frac{1}{2}|\sum_{x \in \mathcal{X}} \Pr(X = x) - \Pr(Y = x)|$ denote the statistical distance between them.

2.2 Different Notions of Entropy

In this section, we recall some the definitions of information-theoretic and computational notions of entropy that are relevant to this work and also state the results related to them.

Unconditional (Information-theoretic) Entropy

Definition 1 (Min-entropy). *The min-entropy of a random variable X, denoted as $H_\infty(X)$ is defined as $H_\infty(X) \stackrel{\text{def}}{=} - \log\left(\max_x \Pr[X = x]\right)$.*

This is a standard notion of entropy used in cryptography, since it measures the worst-case predictability of X.

Definition 2 (Conditional Min-entropy [19]). *The average-conditional min-entropy of a random variable X conditioned on a (possibly) correlated variable Z, denoted as $\widetilde{H}_\infty(X|Z)$ is defined as*

$$\widetilde{H}_\infty(X|Z) = - \log\left(\mathbb{E}_{z \leftarrow Z}\left[\max_x \Pr[X = x | Z = z]\right]\right) = - \log\left(\mathbb{E}_{z \leftarrow Z}\left[2^{-H_\infty(X|Z=z)}\right]\right).$$

This measures the worst-case predictability of X by an adversary that may observe a correlated variable Z.

Lemma 1 (Chain Rule for min-entropy [19]). *For any random variable X, Y and Z, if Y takes on values in $\{0, 1\}^\ell$, then*

$$\widetilde{H}_\infty(X|Y, Z) \geq \widetilde{H}_\infty(X|Z) - \ell \quad \text{and} \quad \widetilde{H}_\infty(X|Y) \geq \widetilde{H}_\infty(X) - \ell.$$

One may also define a more general notion of conditional min-entropy $\widetilde{H}_\infty(X|\mathcal{E})$, where the conditioning happens over an arbitrary experiment $\mathcal{E}$, and not just a "one-time" random variable Y [4].

Computational Entropy a.k.a pseudo-entropy Computational entropy or pseudo-entropy is quantified with two parameters- *quality* (i.e., how much distinguishable a random variable is from a source with true min-entropy to a size-bounded (poly-time) distinguisher)) and *quantity* (i.e., number of bits of entropy).

Definition 3 (Hill Entropy [27, 29]). *A distribution $\mathcal{X}$ has **HILL entropy** at least k, denoted by $H_{\epsilon,s}^{\mathsf{HILL}}(\mathcal{X}) \geq k$, if there exists a distribution $\mathcal{Y}$, where $H_\infty(\mathcal{Y}) \geq k$, such that $\forall \mathcal{D} \in \mathcal{D}_s^{\mathsf{rand},\{0,1\}}, \delta^{\mathcal{D}}(\mathcal{X}, \mathcal{Y}) \leq \epsilon$.*
*For a distribution $(\mathcal{X}, \mathcal{Z})$ we say that $\mathcal{X}$ has **conditional HILL entropy** at least k conditioned on $\mathcal{Y}$, denoted $H_{\epsilon,s}^{\mathsf{HILL}}(\mathcal{X}|\mathcal{Y}) \geq k$, if there exists a joint distribution $(\mathcal{Z}, \mathcal{Y})$ such that $\widetilde{H}_\infty(\mathcal{Z}|\mathcal{Y}) \geq k$, and $\forall \mathcal{D} \in \mathcal{D}_s^{\mathsf{rand},\{0,1\}}, \delta^{\mathcal{D}}\left((\mathcal{X}, \mathcal{Y}), (\mathcal{Z}, \mathcal{Y})\right) \leq \epsilon$.*

2.3 Puncturable Pseudo-Random Functions

In this section, we define the syntax and security properties of a puncturable pseudo-random function family. We follow the definition given in [28].

A *puncturable* family of PRF $\mathsf{pPRF} : \mathcal{K} \times \mathcal{X} \to \mathcal{Y}$ is given by a triple of polynomial time algorithms $(\mathsf{pPRF.setup}, \mathsf{pPRF.puncture}, \mathsf{pPRF.eval})$ and equipped with an additional (punctured) key space $\mathcal{K}_p$ defined as follows:

- $\mathsf{pPRF.setup}(1^\kappa)$: This is a randomized algorithm that takes the security parameter κ as input and outputs a description of the key space $\mathcal{K}$, the punctured key space $\mathcal{K}_p$ and the PRF pPRF.

- $\mathsf{pPRF.puncture}(K, x)$: This is also a randomized algorithm that takes as input a (master) PRF key $K \in \mathcal{K}$, and an input $x \in \mathcal{X}$, and outputs a key $K_x \in \mathcal{K}_p$, often denoted as the punctured key (punctured at the point x). Without loss of generality, we can think of F.puncture as a deterministic algorithm also. This is because we can de-randomize the algorithm by generating its random bits using a PRF keyed by a part of the master key K and given the point x as input.

- $\mathsf{pPRF.eval}(K_x, x')$: This is deterministic algorithm that takes as input the punctured key $K_x \in \mathcal{K}_p$, and an input $x' \in \mathcal{X}$. Let $K \in \mathcal{K}, x \in \mathcal{X}$ and $K_x \leftarrow \mathsf{pPRF.puncture}(K, x)$. The correctness guarantee stipulates that:

$$
\mathsf{pPRF.eval}(K_x, x') = \begin{cases} \mathsf{pPRF}(K, x) & \text{if } x \neq x' \\ \bot & \text{otherwise} \end{cases}
$$

Security of Puncturable PRFs: The security of puncturable PRFs is depicted by a game between a challenger and an adversary. The security game consists of the following four stages:

1. **Setup Phase:** The challenger chooses uniformly at random a (master) PRF key $K \in \mathcal{K}$ and a bit $b \in \{0, 1\}$.

2. **Evaluation Query Phase:** In this phase the adversary $\mathcal{A}$ queries a point $x \in \mathcal{X}$. The challenger sends back the evaluation $\mathsf{pPRF}(K, x)$ to $\mathcal{A}$. These queries can be made arbitrarily and adaptively by $\mathcal{A}$ polynomially many times. Let $E \subset \mathcal{X}$ be the set of evaluation queries.

3. **Challenge Phase:** In this phase, the adversary $\mathcal{A}$ chooses a challenge point $x^* \in \mathcal{X}$. The challenger computes $K_{x^*} \leftarrow \mathsf{pPRF.puncture}(K, x^*)$. If bit $b = 0$, $\mathcal{C}$ sends back $(K_{x^*}, \mathsf{pPRF}(K, x^*))$. Else, the challenger samples a uniformly random $y^* \leftarrow \mathcal{Y}$, and sends back (K_{x^*}, y^*) to $\mathcal{A}$.

4. **Guess Phase:** $\mathcal{A}$ outputs a guess b' for the bit b chosen by the challenger.

The advantage of $\mathcal{A}$ in the above game is defined by:

$$
\mathsf{Adv}_{\mathcal{A}}^{\mathsf{pPRF}}(\kappa) = \Pr\big[b' = b \mid x^* \leftarrow \mathcal{A}^{\mathsf{pPRF.eval}(K, \cdot)}(\kappa, K_x^*) \ \wedge \ x^* \notin E\big].
$$

Definition 4. *The punctured PRF pPRF is said to be secure if for all PPT adversaries $\mathcal{A}$ participating in the above game, $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{pPRF}}(\kappa)$ is negligible in κ.*

2.4 Indistinguishability Obfuscation

A uniform PPT machine $i\mathcal{O}$ is called an indistinguishability obfuscator for a circuit class $\{\mathcal{C}_\kappa\}_{\kappa \in \mathbb{N}}$ if it satisfies the following conditions:

- (*Functionality Preserving*). For all security parameters $\kappa \in \mathbb{N}$, for all $C \in \mathcal{C}_\kappa$, for all inputs x, we have that:

$$\Pr\big[C'(x) = C(x) : C' \leftarrow i\mathcal{O}(\kappa, C)\big] = 1$$

- (*Indistinguishability of Obfuscation*). For any (not necessarily uniform) PPT adversaries $Samp$, $\mathcal{D}$, there exists a negligible function $\mathsf{negl}(\cdot)$ such that the following holds: if for all security parameters $\kappa \in \mathbb{N}$, $\Pr\big[\forall x, C_0(x) = C_1(x) : (C_0, C_1, \mathsf{st}) \leftarrow Samp(\kappa)\big] > 1 - \mathsf{negl}(\kappa)$, then we have:

$$\Big| \Pr\big[\mathcal{D}(\mathsf{st}, i\mathcal{O}(\kappa, C_0)) = 1 : (C_0, C_1, \mathsf{st}) \leftarrow Samp(\kappa)\big]$$
$$- \Pr\big[\mathcal{D}(\mathsf{st}, i\mathcal{O}(\kappa, C_1)) = 1 : (C_0, C_1, \mathsf{st}) \leftarrow Samp(\kappa)\big] \Big| \leq \mathsf{negl}(\kappa).$$

  where the probability is over the coins of $\mathcal{D}$ and $i\mathcal{O}$.

We remark that the algorithms $Samp$ and $\mathcal{D}$ pass state st, which can equivalently be viewed as a single stateful algorithm $\mathcal{B} = (Samp, \mathcal{D})$.

2.5 Entropic Leakage-resilient OWF

In this section, we recall the definition of leakage-resilient one-way functions (LR-OWF) from [7]. Informally, a one-way function (OWF) $g : \{0,1\}^n \rightarrow \{0,1\}^m$ is leakage-resilient if it remains one-way, even in the presence of some leakage about pre-image. In entropy-bounded leakage model, instead of bounding the length of the output of leakage functions (as in bounded leakage model), we bound the entropy loss that happens due to seeing the output of the leakage functions. We follow the definition of [16] to consider the entropy loss over the uniform distribution as a measure of leakiness. We follow this definition since it has nice composability properties as stated below.

Definition 5. [16]. A (probabilistic) function $h : \{0,1\}^* \rightarrow \{0,1\}^*$ is ℓ-leaky, if for all $n \in \mathbb{N}$, we have $\widetilde{H}_\infty(U_n | h(U_n)) \geq n - \ell$, where U_n denote the uniform distribution over $\{0,1\}^n$.

As observed in [16], if a function is ℓ-leaky, i.e, it decreases the entropy of uniform distribution by at most ℓ bits, then it decreases the entropy of every distribution by at most ℓ bits. Moreover, this definition composes nicely in the sense that, if the adversary adaptively chooses different ℓ_i-leaky functions, it learns only $\sum_i \ell_i$ bits of information. We now define the security model for weak PRFs in this entropy-bounded leakage model.

Definition 6. (Entropic leakage-resilient one-wayness). *Let $\mathcal{A}$ be an adversary against* $g : \{0,1\}^n \rightarrow \{0,1\}^m$. *We define the advantage of the adversary $\mathcal{A}$ as* $\mathsf{Adv}_{\mathcal{A}}^{\mathsf{LR\text{-}OWF}}(\kappa) = \Pr[g(x) = y \,|\, x^* \xleftarrow{\$} \{0,1\}^n, y^* = g(x^*); x \leftarrow \mathcal{A}^{\mathcal{O}_{\mathsf{Leak}}(\cdot)}(y^*)]$.

Here $\mathcal{O}_{\mathsf{Leak}}$ is an oracle that on input $h : \{0,1\}^n \rightarrow \{0,1\}^*$ returns $f(x^*)$, subject to the restriction that h is λ-entropy leaky. We say that g is λ-entropic leakage-resilient one-way function (λ-ELR-OWF) if not any PPT adversary $\mathcal{A}$ its advantage defined as above is negligible in κ.

As shown in [17], a second-preimage resistant (SPR) function with $n(\kappa)$ bits input and $m(\kappa)$ bits output is also a $\lambda(\kappa)$-entropy leaky OWF for $\lambda(\kappa) = n(\kappa) - m(\kappa) - \omega(\log \kappa)$.

2.6 Σ-Protocols

A Σ-protocol for an NP relation $\mathcal{R}$ is a special form of *3-move honest verifier zero knowledge* protocol, that enables a prover to prove knowledge to a witness w associated to a statement $x \in \mathcal{L}_{\mathcal{R}}$, without revealing any additional information about the witness W.

Let $\mathcal{R} = \{(x, w)\}$ be an efficiently computable binary relation, with $|w| \leq p(|x|)$, for some polynomial $p = p(\kappa)$. Also let $L_{\mathcal{R}} = \{x : (w, x) \in \mathcal{R}\}$ be the language corresponding to the relation $\mathcal{R}$. A Sigma protocol $\Pi = (P, V)$ is a 3-round interactive protocol between a prover P (holding a private input w corresponding to a statement $x \in L_{\mathcal{R}}$) and a verifier V (holding the statement x). The protocol proceeds as follows: (i) the prover's first message (also called *commitment*) is denoted by $a = P(x, w)$; (ii) the verifier then chooses a random *challenge* $c \in_R \mathsf{Ch}$, where Ch denotes the challenge space, (iii) the prover's second message (also called *response*) is denoted by $z = P(x, w, a, c)$. The tuple (a, c, z) is known as a transcript/proof. We write $V(x, a, c, z) = 1$ *iff* verifier V accepts, and we say that the transcript (a, c, z) is *accepting* for x. A Σ-protocol is *public-coin*, which means that the challenge c is chosen uniformly at random by the verifier V, without having to store any private information about it. We now formally define the properties to be satisfied by a Σ-protocol.

Definition 7. A protocol $\Pi = (P, V)$ is a three-round public-coin protocol satisfying the following properties:

- **Completeness:** *If P and V follow the protocol on input x and private input w to P where $(x, w) \in \mathcal{R}$, then V always accepts.*

- **Special soundness:** *There exists a polynomial time algorithm (or knowledge extractor) K, that on input x and a pair of accepting transcripts $(a, c, z), (a, c', z')$ for x, where $c' \neq c$, outputs w such that $(x, w) \in \mathcal{R}$.*

- **Perfect honest verifier zero knowledge (HVZK):** *There exists a probabilistic polynomial time simulator* Sim *such that the following two distributions are identically distributed:*

$$\left\{ \mathsf{Sim}(x, c) \right\}_{x \in L_{\mathcal{R}}, c \in_R \mathsf{Ch}} \equiv \left\{ \langle P(x, w), V(x, c) \rangle \right\}_{x \in L_{\mathcal{R}}, c \in_R \mathsf{Ch}}$$

13

where $\mathsf{Sim}(x,c)$ *denotes the output of the simulator upon input x and c, and $\langle P(x,w), V(x,c)\rangle$ denotes the output transcript of an execution between P and V, where P has input (x,w), V has input x, and V's random tape (challenge) equals c.*

As shown in [13], HVZK implies witness indistinguishability. The authors in [4] rephrased this property in a slightly different way. In particular, they showed that the oracle access to a prover $P(x,w)$ does not decrease the min-entropy of w in any experiment in which x is given to the predictor. We also follow this formulation in our work.

Lemma 2. [4]. *Let (P,V) be a HVZK protocol for the relation $\mathcal{R}$, (X,W) be random variables over $\mathcal{R}$, and let $\mathcal{A}$ be a PPT adversary. Let $\mathcal{E}_1$ be an arbitrary experiment in which $\mathcal{A}$ is given X at the start of the experiment, and let $\mathcal{E}_2$ be the same as $\mathcal{E}_1$, except that $\mathcal{A}$ is also given oracle access to $P(x,w)$ throughout the experiment. Then $\widetilde{\mathrm{H}}_\infty(W|\mathcal{E}_2) = \widetilde{\mathrm{H}}_\infty(W|\mathcal{E}_1)$.*

2.7 Representation problem.

Let $\mathcal{L} \in \mathbf{NP}$ be a language with an efficiently computable binary relation $\mathcal{R} \subset \{0,1\}^* \times \{0,1\}^*$, i.e, $x \in \mathcal{L}$ if and only if there exists w such that $(x,w) \in \mathcal{R}$ and $|w| = \mathsf{poly}(|x|)$. The value w is called a witness for the statement/theorem $x \in \mathcal{L}$. We say that a representation problem for the relation $\mathcal{R}$ is (t,ε)-hard if for all PPT adversaries running in time t, we have:

$$\Pr[w \neq w' \wedge (x,w), (x,w') \in \mathcal{R} : (x,w,w') \leftarrow \mathcal{A}(1^\kappa)] \leq \varepsilon$$

In most of the cases, the hardness of the representation problem for $\mathcal{R}$ is equivalent to the hardness of the underlying relation $\mathcal{R}$. We will require the hardness of the RSA representation problem as defined below:

RSA Representation problem. Let $N = p \cdot q$, where p and q are prime numbers. Let (e,d) be chosen in such a way that $e \cdot d = 1 \mod \phi(N)$, and e is a prime. Consider the language $\mathcal{L}_{\mathsf{RSA}} = \{(g_1, \cdots, g_\ell, h) : \exists (\rho, \omega_1, \cdots, \omega_\ell) \in \mathbb{Z}_N^* \times \mathbb{Z}_e^\ell \text{ s.t. } h = \prod_{i=1}^\ell g_i^{\omega_i} \cdot \rho^e \mod N\}$, where $(g_1, \cdots, g_\ell)$ are generators of a prime-order cyclic subgroup of $\mathbb{Z}_N^*$. We say that the tuple $w = (\rho, \omega_1, \cdots, \omega_\ell)$ is a representation of h. The ℓ-representation problem asks to compute two different representations w, w' for some $x = (g_1, \cdots, g_\ell, h) \in \mathcal{L}_{\mathsf{RSA}}$. As shown in [40], the ℓ-representation problem is hard for $\mathcal{L}_{\mathsf{RSA}}$ if and only if the RSA problem[3] is hard in $\mathbb{Z}_N^*$.

3 The Entropic Continual Memory Leakage (ECL) Model

As discussed before, no cryptographic primitive can be secure against continual leakage if it has *deterministic* update/ key-evolution procedure. To this end, we introduce a variant of continual memory leakage (CL) model, which we call the *entropic continual memory leakage* (ECL) model. Similar to the CL model, the ECL model also follows a key-evolving paradigm where the secret key is periodically updated, keeping the public key unchanged. We assume that when the secret key evolves to the next time period,

---

[3] The RSA problem is: Given (N, u, e), such that $u = \rho^e \mod N$, compute ρ.

the prior keys are securely erased from the system. Indeed, given the choice of different security models, it might be prudent to look at them through the lens of adversarial classes. In particular, we first present a unified security model that guarantees different levels of security by considering different classes of adversaries. Our unified security framework is powerful enough to capture the existing leakage security models, for e.g., the *bounded* leakage, *continual* leakage and entropy bounded or *noisy* leakage models. In addition, this gives rise to some natural intermediate notions, one of which we call the entropic continual leakage (ECL) model. Looking ahead, we will present this approach for the primitives of Key Evolving Signature Schemes (KES) and Key Evolving Encryption (KEE) schemes.

We will now proceed to describe classes of adversaries that capture not only existing notions but allow us to define new intermediate ones as well. These classes are defined via the types of leakage functions. These are classified as *local* and *non-local* leakage functions. Informally, a leakage function is local when the adversary does not compute the update function as a part of the function. We find this to be a natural requirement in practice and it is critical in the scenario of deterministic update. This is formally defined in Definition 8.

Before we proceed with the definition, we will introduce some additional notations. We will define the random variable representing the secret key at time i as SK_i. Clearly, $SK_1, \ldots, SK_T$ are the random variables representing the secret keys at time periods $1, \ldots, T$ respectively defined by a key update scheme, i.e., for all $1 \leq i \leq T - 1$, SK_{i+1} is obtained from SK_i by using the key update algorithm. Let L_i be the leakage query issued by the adversary at time i. Let $\mathbb{L}_i$ denote the random variable which represents the leakage $L_i(SK_i)$. Also, let R_{L_i} denote the random coins of the adversary corresponding to the leakage function L_i. Further, it is helpful to define Σ_i as the random variable representing the state of the adversary. This captures the cumulative knowledge of the adversary based on all leakage queries and responses, except the query at time period i, and also includes the random coins of the adversary. In other words, $\Sigma_i = \{\{\mathbb{L}_j\}_{j=1}^T \setminus \mathbb{L}_i\} \cup \{\{R_{L_j}\}_{j=1}^T \setminus R_{L_i}\}$.

Definition 8 (Local Leakage). *A leakage query L_t acting on a key sk_t, at time period t is a local leakage query if any key $SK_j \neq SK_t$ does not lose its entropy due to $\mathbb{L}_t = L_t(sk_t)$. In other words,*

$$\mathrm{H}(SK_j | \Sigma_t, \mathbb{L}_t = L_t(sk_t), R_{L_t}) = \mathrm{H}(SK_j | \Sigma_t)$$

Here, H is a suitable notion of computational *entropy.*

In other words, the above definition says that: If the entropy of a key at a time period j is the *same* with and without conditioning on the leakage from time $t \neq j$, then the key and the leakage are *independent*. A local leakage query is one which is independent of every other key.

Remark 1. Before we proceed, we would like to point out that while the above definition of "local leakage" is defined for computational entropy, it can be generalized to include *information-theoretic* notions of entropy. However, for our application of *deterministic*

15

update, it is not hard to see that the notion of entropy needs to be *computational* to avoid a trivial reduction to the bounded leakage model (as explained later).

Now, we present a unified definition of leakage security models based on the allowed adversarial classes.

Admissible adversaries: It is prudent to define different classes of adversaries based on the nature of their leakage queries - based on the length of the output of the leakage functions or the residual entropy of the key(s), given the output of the leakage functions. More formally, let us denote the maximum number of time periods to be T. Let κ denote the security parameter, and let $\alpha = \alpha(\kappa)$, and $\lambda = \lambda(\kappa)$ be two parameters. We define the following classes of adversaries:

- $\mathcal{A}^\alpha$: It is the class of α-*entropic adversaries* where *all* the leakage queries made by $\mathcal{B}$ satisfy the following condition:

$$\forall j \in [T], \quad \mathrm{H}\big(SK_j \,|\, \mathbb{L}_1, \cdots, \mathbb{L}_T, R_{L_1}, \cdots, R_{L_T}\big) \geq \alpha,$$

  for some parameter α, and H being a suitable notion of entropy (either information theoretic or computational). Recall that, the random variables $\mathbb{L}_i$ represents the leakage from the i^{th} secret key sk_i, i.e., $\mathbb{L}_i = L_i(sk_i)$, and R_{L_i} denotes the random coins of the adversary. Next, we look at conditioning on a subset of these queries such as the set of non-local leakage queries, as defined below.

- $\mathcal{A}_N^\alpha$: The class α-*entropic non-local* adversaries is the class of PPT adversaries $\mathcal{B}$ such that the non-local leakage queries made by $\mathcal{B}$ satisfy the following definition:

$$\forall j \in [T], \mathrm{H}\big(SK_j \,|\, \mathbb{L}_{a_1}, \cdots, \mathbb{L}_{a_i}, R_{L_{a_1}}, \cdots, R_{L_{a_i}}\big) \geq \alpha \tag{1}$$

  where $(R_{L_{a_1}}, \cdots, R_{L_{a_i}})$ are random coins of the adversary corresponding to the subset $(L_{a_1}, \ldots, L_{a_i}) \subseteq (L_1, \ldots, L_T)$ of non-local leakage queries. Note that this *does not* restrict the local leakage queries in any way.

- $\mathcal{A}^\lambda$: The class of λ-*length adversaries* is the set of PPT adversaries $\mathcal{B}$ making leakage queries such that $\forall j \in [T], |L_i(sk_i)| < \lambda$.
  We now look at imposing this length bound condition only on a subset of the leakage queries. We get,

- $\mathcal{A}_L^\lambda$: The class of λ-*length local adversaries* is the set of all PPT adversaries $\mathcal{B}$ making leakage queries such that, for all local leakage queries L_i we have that, $|L_i(sk_i)| < \lambda$. This does not restrict the output length of the non-local leakage queries.

In the above definitions, the L and N in subscript indicate local and non-local respectively. In addition, the presence of a 1 or 0 in the superscript indicates the number of queries the adversary can make. By default, the adversary can make a polynomial number of queries. From the above notations, it is it is easy to recover the class of adversaries for the continual leakage Model, bounded leakage model and entropy bounded/noisy leakage model as follows.

16

- *Continual Leakage (CL) Model*: The adversary is not constrained by the number or the type of queries. The only constraint is on the length of the output of the leakage functions in each time period. Let this be denoted by $\lambda < |sk|$. Therefore, $\mathcal{A}_{\mathsf{CL}} = \mathcal{A}^{\lambda}$. Note that, in this context the right notion of entropy that must be considered in the definition of local leakage (see Def. 8) is *average conditional min-entropy*, since the update function is randomized.

- *Bounded Leakage (BL) Model*: Let us recall that the BL model is envisioned for the setting where there is a single secret key and there is no key update procedure. Cast in our setting, this can be thought of as a scheme over multiple time periods where the key evolves to itself, i.e the update is the *identity* function. Note that, the only constraint is that the length of the leakage cannot exceed λ. Therefore, $\mathcal{A}_{\mathsf{BL}} = \mathcal{A}_L^{\lambda,1} \cap \mathcal{A}_N^0$.

- *Entropic Bounded Leakage (EBL) Model*: This is a generalization of the BL model where the constraint is on residual entropy of the secret key conditioned on the leakage. Thus, $\mathcal{A}_{\mathsf{EBL}} = \mathcal{A}_L^{\alpha,1} \cap \mathcal{A}_N^{\alpha,0}$, and the right notion of entropy to be considered in Def. 8 is *average conditional min-entropy*.

The Entropic Continual Leakage Model. This brings us to the definition of the Entropic Continual Leakage (ECL) model.

An initial attempt. Recall that in the CL model the update function is randomized. In contrast, in the ECL model, we want to consider *deterministic* update functions. As a first attempt, suppose we impose the following restriction: there are a maximum of T time periods and the keys need to have residual entropy of α conditioned on the leakage information, at each time period. In other words, this model considers the class of adversary $\mathcal{A}_{\mathsf{ECL}'} = \mathcal{A}^{\alpha}$.

However, it is not hard to see that this model is actually the *one-time* entropy-bounded leakage (EBL) model (without any key updates). To see this, consider a *trivial* scheme where the update function is simply the identity function. This, along with the requirement that each of the keys be entropic implies that any key-evolving scheme that is secure against *one-time* entropic leakage (i.e, in the EBL model) is also secure in this model. We thus seek something stronger.

Defining the ECL model. Our approach is as follows:

1. We allow the adversary to make *local* leakage queries, provided it leaks only a maximum of λ bits per time period. This rules out the existence of the trivial scheme (where the update function is an identity function), because there exists an adversary which leaks λ physical bits from the keys in each time period and thereby recovering the entire key sk in $|sk|/\lambda$ time periods.

2. In addition, we allow the adversary to make *non-local* leakage queries, provided that the keys have minimum residual entropy of α conditioned on the leakage information at each time period. Since the update function is deterministic, the notion of entropy we consider is computational, namely *average conditional HILL entropy*.

17

Thus, from the above we get that, $\mathcal{A}_{\mathsf{ECL}} = \mathcal{A}_L^\lambda \cap \mathcal{A}_N^\alpha$. The notion of entropy considered while defining the classes $\mathcal{A}_L^\lambda$ and $\mathcal{A}_N^\alpha$ is $\mathrm{H}_{\epsilon,s}^{\mathsf{HILL}}$. Here $\mathcal{A}_{\mathsf{ECL}}$ denote the class of adversary for the ECL model. When a scheme is secure against $\mathcal{A}_{\mathsf{ECL}}$, we call that scheme to be (α, λ)-ECL-secure.

3.1 Forward Security under Entropic Continual Memory Leakage (FS + ECL) Model

One of the key advantages of considering continuous leakage-resilient cryptographic primitives with deterministic update function is that, they can be seamlessly amalgamated with the notion of forward security (FS). To this end, we provide a unified model that captures both these aspects simultaneously– forward security and resilience to entropic continual leakage. We call this model as *forward-security under entropic continual leakage* ("FS+ECL"), motivated by the FS+CL model of [6]. In simpler terms, we also require FS to hold under the ECL model.

The key difference from the ECL model is that an adversary also has access to an oracle Exp which allows it to break into the system at some point in time (say time period $t^* \in [T]$) and obtain the secret key sk_{t^*} in full. The access to oracle Lk allows it to obtain leakage from the secret keys of all the time periods prior to the break-in period t^*[4] subject to the attacker being legitimate (more on this below).

The security under this model is guaranteed against the class of "*valid/ admissible adversaries*" $\mathcal{A}_{\mathsf{FS+ECL}}$, which we define next. Recall that, for the ECL model the class of admissible adversaries was defined as $\mathcal{A}_{\mathsf{ECL}} = \mathcal{A}_L^\lambda \cap \mathcal{A}_N^\alpha$. However, before formally defining the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of adversaries, we will need to modify the definitions of "local leakage" (see Def. 8) and the α-entropic condition of the ECL model (see Eq. 1) for the FS+ECL model as follows:

Definition 9 (Local Leakage - FS+ECL Model). *We will define the random variable representing the secret key at time i as SK_i, then $SK_1, \ldots, SK_T$ are the random variables representing the secret keys at time periods $1, \ldots, T$ respectively defined by an appropriate Key Update function. Let $t^* \in [T]$ denote the period of break-in. Let L_i be the leakage query issued by the adversary at time i. Let $\mathbb{L}_i$ denote the random variable which represents the leakage $L_i(SK_i)$. Also, let R_{L_i} denote the random coins of the adversary corresponding to the function L_i We can then define the leakage query L_t acting on a key sk_t, at time period t as a local leakage query if any key $SK_j \neq SK_t$ does not lose its entropy due to $\mathbb{L}_t = L_t(sk_t)$. In other words,*

$$\mathrm{H}(SK_j \mid \Sigma_t, \mathbb{L}_t = L_t(sk_t), R_{L_t}) = \mathrm{H}(SK_j \mid \Sigma_t)$$

where we let Σ_t represent the other queries and the respective responses, i.e, $\Sigma_t = \{\mathbb{L}_j\}_{j=1}^{t^*-1} \setminus \mathbb{L}_t \cup \{R_{L_j}\}_{j=1}^{t^*-1} \setminus R_{L_t}$.

---

[4] It is vacuous to look at leakage queries starting from time t^* as the adversary has the secret key sk_{t^*} in full and can compute future keys under the regime of deterministic updates.

Note that, as remarked before, we only include the leakage queries until time period $t^* - 1$ in this definition. This is the subtle difference from the definition of local leakage in Definition 8.

Definition 10 (λ-length, local adversary - FS+ECL Model). *The class of λ-length, local adversary in the* FS+ECL *model ($\mathcal{A}_{L,t^*}^\lambda$) is the class of all PPT adversaries such that: for all local leakage queries L_i where $i \in [t^* - 1]$, $|L_i(sk_i)| < \lambda$. This adversary does not have any constraints on the non-local queries, by definition.*

Note that, this is similar to the definition $\mathcal{A}_L^\lambda$ defined in the ECL model.

Definition 11 (α-entropic, non-local adversary - FS+ECL Model). *The class of α-entropic non-local adversaries in the* FS+ECL *model ($\mathcal{A}_{N,t^*}^\alpha$) is the class of PPT adversaries such that the leakage functions satisfy the following condition:*

$$\forall j \in [t^* - 1], \mathrm{H}\big(SK_j \,|\, \mathbb{L}_{a_1}, \cdots, \mathbb{L}_{a_i}, R_{L_{a_1}}, \cdots, R_{L_{a_i}}\big) \geq \alpha \qquad (2)$$

where where $(R_{L_{a_1}}, \cdots, R_{L_{a_i}})$ are random coins of the adversary corresponding to the subset $(\mathbb{L}_{a_1}, \ldots, \mathbb{L}_{a_i}) \subseteq (\mathbb{L}_1, \ldots, \mathbb{L}_{(t^*-1)})$ of non-local leakage queries made by the adversary before the time period t^*. This adversary does not have any constraints on the local queries, by definition.

Remark 2. Note that, both the definitions above use computational notions of entropy, namely *average conditional HILL entropy*. This is needed, since the update function is deterministic.

Defining the FS+ECL model. Combining Definitions 10 and 11 we define the class of admissible adversaries $\mathcal{A}_{\mathsf{FS+ECL}}$ in the FS+ECL model as: $\mathcal{A}_{\mathsf{FS+ECL}} = \mathcal{A}_{L,t^*}^\lambda \cap \mathcal{A}_{N,t^*}^\alpha$.

Remark 3. Finally, we note that, we model the adversary by classifying types of leakage functions rather than providing different oracles (corresponding to local and non-local queries) because it sidesteps the issues of entropy estimation which is currently not known to be doable in polynomial time.

3.2 Modelling Signature Schemes in the FS+ECL Model.

In this section, we consider key-evolving signature schemes in the FS+ECL setting. Unfortunately, it is not difficult to see that one *cannot* achieve the standard notion of *existential unforgeability* in the FS+ECL setting. This is because the adversary can use the secret signing keys $sk_1, \cdots, sk_i$ (corresponding to time periods $1, \cdots, i$ respectively) to leak bits of the signature under the key sk_{i+1} in small chunks (for an arbitrary hard-coded choice of random coins used by the signing algorithm), then concatenate the recovered chunks and claim the resulting signature as a forgery in time period $i + 1$. Therefore, to construct meaningful signature schemes in the FS+ECL setting, we consider the weaker notion of *entropic unforgeability*, as defined by [4] in the context of bounded retrieval

Game $\mathrm{FUFECL}_{\mathsf{KES}}^{\mathcal{A}_{\mathsf{sig}}=(\mathcal{A}_1,\mathcal{A}_2)}(\kappa,\alpha,\lambda,T)$	Game $\mathrm{FINDECL}_{\mathsf{KEE}}^{\mathcal{A}_{\mathsf{pke}}}(\kappa,\alpha,\lambda,T)$
$S \leftarrow \emptyset$; $t \leftarrow 1$; $t^* \leftarrow T+1$	$t \leftarrow 1$; $t^* \leftarrow T+1$
$(vk, sk_1) \leftarrow_\$ \mathsf{KES.Kg}(1^\kappa)$	$(pk, sk_1) \leftarrow_\$ \mathsf{KEE.Kg}(1^\kappa)$
$v \leftarrow \mathcal{A}_1^{\mathsf{Sign},\mathsf{Lk},\mathsf{Up},\mathsf{Exp}}(1^\kappa, vk)$	$(i, m_0, m_1, state) \leftarrow_\$ \mathcal{A}_1^{\mathsf{Up},\mathsf{Lk},\mathsf{Exp},\mathsf{Dec}}(1^\kappa, pk)$
$(i, m, \sigma) \leftarrow \mathcal{A}_2^{\mathsf{Sign},\mathsf{Up},\mathsf{Exp}}(1^\kappa, vk, v)$	$b \leftarrow_\$ \{0,1\}$; $(i, c) \leftarrow_\$ \mathsf{KEE.Enc}(1^\kappa, pk, i, m_b)$
$\mathsf{win}_1 \leftarrow (1 \le i < t^*) \wedge ((i, m, \sigma) \notin S)$	$b' \leftarrow_\$ \mathcal{A}_2^{\mathsf{Up},\mathsf{Exp},\mathsf{Dec}}(1^\kappa, state, (i, c))$
$\mathsf{win}_2 \leftarrow \mathsf{KES.Vfy}(1^\kappa, vk, m, (i, \sigma))$	If not $(1 \le i < t^*)$ then return false
Return $(\mathsf{win}_1 \wedge \mathsf{win}_2)$	If $|m_0| \ne |m_1|$ then return false
	Return $(b' = b)$
$\underline{\mathsf{Sign}(t, m)}$	
$(t, \sigma) \leftarrow_\$ \mathsf{KES.Sign}(1^\kappa, vk, t, sk_t, m)$	$\underline{\mathsf{Up}()}$
$S \leftarrow S \cup \{(t, m, \sigma)\}$; Return (t, σ)	If $t < T$ then
	$sk_{t+1} := \mathsf{KEE.Upd}(1^\kappa, pk, t, sk_t)$
$\underline{\mathsf{Lk}(L)}$	$\quad t \leftarrow t+1$
Return $L(sk_t)$	
	$\underline{\mathsf{Lk}(L)}$
$\underline{\mathsf{Up}()}$	Return $L(sk_t)$
If $t < T$ then	
$\quad sk_{t+1} := \mathsf{KES.Upd}(1^\kappa, vk, t, sk_t)$	$\underline{\mathsf{Exp}()}$
$\quad t \leftarrow t+1$	$t^* \leftarrow t$; Return sk_t
$\underline{\mathsf{Exp}()}$	$\underline{\mathsf{Dec}(t, c_t)}$
$t^* \leftarrow t$; Return sk_t	If $(t, c_t) \ne (i, c)$
	$\quad$ Return $\mathsf{KEE.Dec}(1^\kappa, pk, (t, sk_t), c_t)$
	Else, return $\bot$.

Fig. 1. Games defining forward unforgeability of key-evolving signature scheme KES under entropic continual leakage, and forward indistinguishability of key-evolving encryption scheme KEE under entropic continual leakage.

model (BRM)[5]. An attack against entropic unforgeability in the FS+ECL model is valid, if the adversary manages to forge a message (for some time period t) which is chosen from some distribution of significant entropy, *after* the adversary finishes interacting with the leakage oracle. Similar to [4], we allow the forger to select this distribution.

A key-evolving signature (KES) scheme consists of the following algorithms (KES.Kg, KES.Upd, KES.Sign, KES.Vfy). The key generation algorithm KES.Kg takes as input the security parameter 1^κ (in unary), and outputs a key pair (vk, sk_1), where sk_1 is the base signing key (corresponding to time period 1). The *deterministic* key update algorithm KES.Upd takes as input $(1^\kappa; vk, t, sk_t)$ and outputs sk_{t+1}, the secret key for the next time period $t + 1 \in [T]$. Here T denotes the total number of time periods supported by the scheme. We stress that the update algorithm is *deterministic*. The signing algorithm KES.Sign takes as input $(1^\kappa, vk, t, sk_t)$ and a message $m \in \mathcal{M}$ (where $\mathcal{M}$ is the message space) and outputs a signature σ_t for the t^{th} time period. The verification algorithm

---

[5] A similar impossibility result for constructing existentially unforgeable signatures arises in the BRM setting, since the adversary can simply leak a signature, whose size may be much less than the size of the signing key.

KES.Vfy takes as input $(1^\kappa, vk, m, (t, \sigma_t))$ to return a decision in $\{\mathsf{true}, \mathsf{false}\}$ regarding whether σ_t is a valid signature of m relative to vk in time period t.

Correctness: The correctness requirement for a KES scheme states that: for all $\kappa \in \mathbb{N}$, all $(vk, sk_1,) \leftarrow \mathsf{KES.Kg}(1^\kappa)$, all $i \in [T]$, all $m \in M$, all $sk_2, \cdots, sk_i$ satisfying $sk_j \leftarrow \mathsf{KES.Upd}(1^\kappa, vk, j-1, sk_{j-1})$, $\forall\, 2 \leq j \leq i$, and all $\sigma_i \leftarrow \mathsf{KES.Sign}(1^\kappa, vk, i, sk_i, m)$, it should hold that $\mathsf{KES.Vfy}(1^\kappa, vk, m, (i, \sigma_i)) = \mathsf{true}$.

Forward unforgeability under entropic continual leakage (FUFECL).

Given a key evolving signature scheme KES, consider the game $\mathrm{FUFECL}_{\mathsf{KES}}^{\mathcal{A}_{\mathsf{sig}}=(\mathcal{A}_1, \mathcal{A}_2)}(\kappa, \alpha, \lambda, T)$ as defined in Figure 1 running between a challenger $\mathcal{C}$ and a PPT adversary $\mathcal{A}_{\mathsf{sig}} = (\mathcal{A}_1, \mathcal{A}_2)$. The game is parametrized by the security parameter κ, entropic leakage parameters (α, λ), and the total number of time periods T that can be supported by the scheme. We separate the attacker $\mathcal{A}_{\mathsf{sig}}$ into two parts $(\mathcal{A}_1, \mathcal{A}_2)$, where $\mathcal{A}_1$ runs for the first stage of the attack with access to the signing, leakage, update and expose oracles. Once this stage is over, $\mathcal{A}_1$ outputs an arbitrary hint for v for $\mathcal{A}_2$, who then attempts to forge the signature of some message while having access to the all the oracles, *except* the leakage oracle. Below we provide the definition of entropic adversaries.

Definition 12 (Entropic Adversaries). *For any adversary $\mathcal{A}_{\mathsf{sig}} = (\mathcal{A}_1, \mathcal{A}_2)$, let $\mathsf{View}_{\mathcal{A}_1}$ be a random variable describing the view of $\mathcal{A}_1$ including its random coins and the responses of all the oracles* (Sign, Lk, Exp, Up). *Let $\mathsf{MSG}_{\mathcal{A}_2}$ be a random variable describing the message output by $\mathcal{A}_2$ in the game* FUFECL *(see Figure 1). We say that an adversary $\mathcal{A}_{\mathsf{sig}} = (\mathcal{A}_1, \mathcal{A}_2)$ is "entropic" if $\widetilde{\mathrm{H}}_\infty(\mathsf{MSG}_{\mathcal{A}_2} | \mathsf{View}_{\mathcal{A}_1}) \geq \omega(\log \kappa)$ for security parameter κ.*

To prevent trivial attacks in the FUFECL Game, we define the class of "valid" adversary in Definition 13.

Definition 13 (Valid Adversary against the FUFECL Game). *We call an adversary $\mathcal{A}_{\mathsf{sig}}$ "valid" if the following holds true: (i) it makes at most* one *query to its* Exp *oracle and this is the last query, (ii) the adversary $\mathcal{A}_{\mathsf{sig}}$ belongs to the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of* admissible *adversaries (see Section 3.1 for the definition of the class $\mathcal{A}_{\mathsf{FS+ECL}}$), and (iii) the adversary $\mathcal{A}_{\mathsf{sig}} = (\mathcal{A}_1, \mathcal{A}_2)$ is entropic.*

Definition 14. (Forward unforgeability under entropic continual leakage). *We say that* KES = (KES.Kg, KES.Upd, KES.Sign, KES.Vfy) *is (λ, α)-forward unforgeable under entropic continual leakage $((\lambda, \alpha)$-FUFECL) for all "valid" PPT adversaries $\mathcal{A}_{\mathsf{sig}}$ (as defined in Definition 13), if the advantage of $\mathcal{A}_{\mathsf{sig}}$ defined as*

$\mathsf{Adv}_{\mathsf{KES}, \mathcal{A}_{\mathsf{sig}}}^{\mathsf{fufecl}}(\kappa) = \Pr\left[\mathrm{FUFECL}_{\mathsf{KES}}^{\mathcal{A}_{\mathsf{sig}}=(\mathcal{A}_1, \mathcal{A}_2)}(\kappa, \alpha, \lambda, T) = 1\right]$ *in the game* FUFECL *(see Figure 1) is negligible in the security parameter κ.*

3.3 Modelling Public Key Encryption in the FS+ECL Model.

A key-evolving encryption (KEE) scheme consists of the following algorithms (KEE.Kg, KEE.Upd, KEE.Enc, KEE.Dec). The key generation algorithm KEE.Kg takes as input the security parameter 1^κ (in unary), and output a public-secret key pair (pk, sk_1),

where sk_1 is secret key corresponding to the base time period. The key update algorithm KEE.Upd takes as input $(1^\kappa, pk, i)$ and secret key sk_i for time period i, and outputs a secret key sk_{i+1} for the next time period. Assume that the total number of time period supported by the scheme is $T \in \mathbb{N}$. The encryption algorithm KEE.Enc takes as input $(1^\kappa, pk, i)$ and a message m to return c_i, where c_i is the ciphertext encrypting m under pk for time period $i \in [T]$. The decryption algorithm KEE.Dec takes $(1^\kappa, pk, i, sk_i, c_i)$ as input to return an output in $m \cup \{\bot\}$.

Correctness: The correctness requirement for a KEE scheme states that: for all $\kappa \in \mathbb{N}$, for all $(pk, sk_1) \leftarrow$ KEE.Kg$(1^\kappa, T)$, all $m \in \mathcal{M}$, all $sk_2, \cdots, sk_i$ satisfying
$sk_j \leftarrow$ KEE.Upd$(1^\kappa, pk, j - 1, sk_{j-1}), \forall\, 2 \le j \le i$, and all $c_i \leftarrow$ KEE.Enc$(1^\kappa, pk, i, m)$, it should hold that KEE.Dec$(1^\kappa, pk, i, sk_i, c_i) = 1$.

Forward Indistinguishability under entropic continual leakage (FINDECL**).**
Given a key-evolving encryption scheme KEE, consider the experiment FINDECL$_{\mathsf{KEE}}^{\mathcal{A}}(\kappa, \alpha, \lambda, T)$ as defined in Figure 1, running between a challenger $\mathcal{C}$ and a PPT adversary $\mathcal{A}_{\mathsf{pke}} = (\mathcal{A}_1, \mathcal{A}_2)$, parametrized by the security parameter κ, entropic leakage parameters (α, λ), and the total number of time periods T that can be supported by the scheme. To achieve a definition devoid of trivial attacks, we define the validity of an adversary in Definition 15.

Definition 15 (Valid Adversary against the FINDECL Game). *We call an adversary* $\mathcal{A}_{\mathsf{pke}} = (\mathcal{A}_1, \mathcal{A}_2)$ *"valid" if the following holds true: (i) it makes at most one query to its* Exp *oracle and this is the last query, (ii) the adversary* $\mathcal{A}_{\mathsf{pke}}$ *belongs to the class* $\mathcal{A}_{\mathsf{FS+ECL}}$ *of admissible adversaries (see Section 3.1 for the definition of the class* $\mathcal{A}_{\mathsf{FS+ECL}}$*).*

Definition 16. (Forward indistinguishability under entropic continual leakage). *We say that* KEE $=$ (KEE.Kg, KEE.Upd, KEE.Enc, KEE.Dec) *is* (λ, α)-*forward indistinguishable under entropic continual leakage (*(λ, α)-FINDECL*) with respect to the class* $\mathcal{A}_{\mathsf{FS+ECL}}$ *of adversaries, if the advantage defined as*
Adv$_{\mathsf{KEE}, \mathcal{A}_{\mathsf{pke}}}^{\mathsf{findecl}}(\kappa) = \Pr[\text{FINDECL}_{\mathsf{KEE}}^{\mathcal{A}_{\mathsf{pke}}}(\kappa, \alpha, \lambda, T) = 1] - \frac{1}{2}$ *(see Figure 1) is negligible for all "valid" PPT adversaries* $\mathcal{A}_{\mathsf{pke}} = (\mathcal{A}_1, \mathcal{A}_2)$ *(as defined above).*

Remark 4. Note that, in our FINDECL$_{\mathsf{KEE}}^{\mathcal{A}_{\mathsf{pke}}}(\kappa, \alpha, \lambda, T)$ game, the adversary $\mathcal{A}_2$ does *not* get access to the leakage oracle. This is related to the impossibility of achieving security against *after-the-fact* leakage, where the adversary, after receiving the challenge ciphertext c, can use the leakage oracle to specifically leak the challenge bit b hidden by c [26, 39]. To bypass this, one can consider weaker leakage models, like the split-state leakage model [26]. However, in this case the machinery would become significantly more complex and take away from our core focus.

4 Construction of FS+ECL Encryption Scheme.

In this section, we present our construction of key-evolving encryption (KEE) scheme in the FS+ECL model. To this end, we first introduce and formalize a notion of *forward-secure entropic leakage-resilient non-interactive key exchange* (FS-ECLR-NIKE) protocol (see

Section 4.1). In Section 4.4, we show how to construct a FS-ECLR-NIKE protocol from indistinguishability obfuscation and one-way functions. Finally, in Section 4.6, we show a *generic transformation* from a FS-ECLR-NIKE protocol to a FINDECL secure encryption scheme. We view our security model and construction of FS-ECLR-NIKE as a result of independent interest, that may find more useful applications beyond the one shown in this work.

4.1   NIKE in FS+ECL Model.

Non-interactive key exchange (NIKE) protocols allow two (or more) parties to establish a shared key between them, without any interaction. It is assumed that the public keys of all the parties are pre-distributed and known to each other. In this work, we consider two-party NIKE protocols and extend them to the setting of forward-security and entropic continual leakage model. We provide the definition of forward secure (FS) non-interactive key exchange (NIKE) protocol in the entropic continual leakage model (dubbed as "FS-ECLR-NIKE"). To bypass the black-box impossibility result of constructing leakage-resilient NIKE protocol in the plain model [36], we consider the NIKE protocols in the *common reference string* (CRS) model, where we rely on leak-free randomness to generate the CRS. Our security model for FS-ECLR-NIKE can be seen as a leakage-resilient adaptation of the model of forward-secure NIKE (FS-NIKE) of Pointcheval and Sanders ($\mathcal{PS}$ model) [41]. We call our model of NIKE as the $\mathcal{ECL}$-$\mathcal{PS}$ model.

4.2   Syntax of FS-ECLR NIKE.

A FS-ECLR-NIKE scheme FS-ECLR-NIKE, consists of the tuple of algorithms (NIKE.Setup, NIKE.Gen, NIKE.Upd, NIKE.Key). We associate to FS-ECLR-NIKE a public parameter space $\mathcal{PP}$, public key space $\mathcal{PK}$, secret key space $\mathcal{SK}$, shared key space $\mathcal{SHK}$, and an identity space $\mathcal{IDS}$. Identities are used to track which public keys are associated with which users; we are *not* in the identity-based setting.

- NIKE.Setup$(1^\kappa, (\alpha, \lambda), T)$: This is a randomized algorithm that takes as input the security parameter κ (expressed in unary), leakage parameters (α, λ) of the ECL model, and a maximum number of time period T supported by the system and outputs public parameters $params \in \mathcal{PP}$. The current time period t is initially set to 1.
- NIKE.Gen$(1^\kappa, ID)$: On input an identity $ID \in \mathcal{IDS}$, the key generation outputs a public-secret key pair (pk, sk_t) for the current time period t. We assume that the secret keys implicitly contain the time periods.
- NIKE.Upd(sk_t) : The update algorithm takes as input the secret key sk_t at time period t and outputs the updated secret key sk_{t+1} for the next time period $t + 1$, if $t < T$. They key sk_t is then securely erased from memory. If $t = T$, then the secret key is erased and there is no new key

23

$$\underline{\text{Game } \mathcal{ECL}\text{-}\mathcal{PS}^{\mathcal{A}_{\text{nike}}}_{\text{NIKE}}(\kappa, \alpha, \lambda, T)}$$

$params \leftarrow \mathsf{NIKE.Setup}(1^\kappa, (\alpha, \lambda), T) \; ; \; S, C, Q \leftarrow \emptyset \; // \; S, C$ and Q maintains the list of *honest*, *corrupt* and *exposed* users respectively.

$(ID_A, ID_B, \tilde{t}) \leftarrow \mathcal{A}_1^{\mathsf{RegHon},\mathsf{RegCor},\mathsf{CorrReveal},\mathsf{Lk},\mathsf{Exp}}(params)$
$b \leftarrow_\$ \{0, 1\}$
If $b = 0$ then Return $shk_{\tilde{t}}^{AB} \leftarrow \mathsf{NIKE.Key}(ID_A, pk_A, ID_B, sk_{\tilde{t}}^B)$
Else Return $shk_{\tilde{t}}^{AB} \leftarrow_\$ \mathcal{SHK}$
$b' \leftarrow_\$ \mathcal{A}_2^{\mathsf{RegHon},\mathsf{RegCor},\mathsf{Reveal},\mathsf{Lk},\mathsf{Exp}}(shk_{\tilde{t}}^{AB})$

If $(ID_A, -, -, \mathsf{corrupt}) \in C$ or $(ID_B, -, -, , \mathsf{corrupt}) \in C$, then return $\perp$
If $(ID_A, t^*) \in Q$ and $t^* \leq \tilde{t}$, then return $\perp$
If $(ID_B, t^*) \in Q$ and $t^* \leq \tilde{t}$, then return $\perp$
Return $(b' = b)$

Fig. 2. Game defining ECL-PS of NIKE scheme NIKE under entropic continual leakage.

- $\mathsf{NIKE.Key}(ID_A, pk^A, ID_B, sk_t^B)$: On input an identity $ID_A \in \mathcal{IDS}$ associated with public key pk_A, and another identity $ID_B \in \mathcal{IDS}$ with secret key sk_t^B corresponding to the current time period t, output the shared key $shk_t^{AB} \in \mathcal{SHK}$ or a failure symbol $\perp$. If $ID_A = ID_B$, the algorithm outputs $\perp$. Since the secret key sk_t^B is associated with time period t, the shared key shk_t^{AB} between the two users ID_A and ID_B also corresponds to the same time period t.

Correctness: The correctness requirement states that the shared keys computed by any two users ID_A and ID_B in the *same* time period are *identical*. In other words, for any time period $t \geq 1$, and any pair (ID_A, ID_B) of users having secret keys sk_t^A and sk_t^B respectively, it holds that:

$$\mathsf{NIKE.Key}(ID_A, pk^A, ID_B, sk_t^B) = \mathsf{NIKE.Key}(ID_B, pk^B, ID_A, sk_t^A).$$

4.3 Security Model for FS+ECL NIKE.

Our security model for NIKE generalizes the model of forward-secure NIKE of [41] (henceforth refereed to as the $\mathcal{PS}$ model) in the ECL setting. We refer to our model as the $\mathcal{ECL}\text{-}\mathcal{PS}$ model. Security of a NIKE protocol NIKE in the $\mathcal{ECL}\text{-}\mathcal{PS}$ model is defined by a game ECL-PS between an adversary $\mathcal{A}_{\text{nike}} = (\mathcal{A}_1, \mathcal{A}_2)$ and a challenger $\mathcal{C}$ (see Figure 2). Before the beginning of the game, the challenger $\mathcal{C}$ also initializes threee sets S, C and Q to be empty sets. Our security model makes use of the following oracles:

1. $\mathsf{RegHon}(ID)$: This oracle is used by $\mathcal{A}_{\text{nike}}$ to register a new honest user ID at the initial time period. The challenger runs the NIKE.Gen algorithm with the current time period as 1, returns the public key pk to $\mathcal{A}_{\text{nike}}$ and records $(ID, sk_1, pk, \mathsf{honest})$. This implicitly defines all the future keys $sk_2, \cdots, sk_T$. This query may be asked at most *twice* by $\mathcal{A}_{\text{nike}}$. The challenger adds the tuple $(ID, sk_1, pk, \mathsf{honest})$ to the set S. Users registered by this query are called "honest".

2. RegCor(ID, pk) : This oracle allows the adversary to register a new corrupted user ID with public key pk. The challenger adds the tuple $(ID, --, pk, \mathsf{corrupt})$ to the set C. We call the users registered by this query as "corrupt".

3. CorrReveal(ID_A, ID_B, t) : $\mathcal{A}_{\mathsf{nike}}$ supplies two indices where ID_A was registered as *corrupt* and ID_B was registered as *honest*. The challenger looks up the secret key sk_1^B (corresponding to ID_B) and computes the updated key sk_t^B corresponding to time period t. Then it runs NIKE.Key$(ID_A, pk^A, ID_B, sk_t^B)$ to get the shared key shk_t^{AB} for time period t. The challenger then returns shk_t^{AB} to $\mathcal{A}_{\mathsf{nike}}$.

4. Lk(L, ID, t) : The adversary $\mathcal{A}_{\mathsf{nike}}$ submits a leakage function $L : \mathcal{PP} \times \mathcal{SK} \to \{0,1\}^*$ to leak on the secret key of user ID for time period t, provided that $\mathcal{A}_{\mathsf{nike}}$ belongs to the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of admissible adversaries (see Section 3.1 for the definition of the class $\mathcal{A}_{\mathsf{FS+ECL}}$.).

5. Exp(ID, t^*) : This query is used by $\mathcal{A}_{\mathsf{nike}}$ to get user $ID's$ secret key at the time period t. The challenger looks for a tuple $(ID, sk_1, pk, \mathsf{honest})$. If there is a match, it computes sk_t corresponding to time period t and returns sk_t to $\mathcal{A}_{\mathsf{nike}}$. Else, it returns $\bot$. The challenger adds (ID, t^*) to the set Q.

The formal details of our ECL-PS game is given in Figure 2.

Definition 17 (Forward Secure Entropic Continual Leakage Resistant NIKE). *A NIKE protocol* NIKE *is* (λ, α)-*forward-secure under entropic continual leakage* $((\lambda, \alpha)$-FS-ECLR$)$ *with respect to any adversary* $\mathcal{A}_{\mathsf{nike}}$ *playing the above* $\mathcal{ECL}$-$\mathcal{PS}$ *game (see Figure 2), if the advantage defined as* $\mathsf{Adv}_{\mathcal{A}_{\mathsf{nike}}}^{\mathsf{fs\text{-}ecl}}(\kappa) = |\Pr[\mathcal{ECL}$-$\mathcal{PS}_{\mathsf{ECL}}^{\mathcal{A}}(\kappa, \alpha, \lambda, T) = 1] - 1/2|$ *is negligible in* κ.

In other words, the adversary $\mathcal{A}_{\mathsf{nike}}$ succeeds in the above experiment if it is able to distinguish a valid shared key between two users from a random session key. To avoid trivial win, some restrictions are enforced, namely: (i) both the targeted users needs to be honestly registered (ii) the adversary $\mathcal{A}_{\mathsf{nike}}$ is not allowed to obtain the secret keys corresponding to any of the targeted users prior to the challenge time period $\tilde{t}$, (iii) $\mathcal{A}_{\mathsf{nike}}$ is allowed to leak on the secret keys of both the target users ID_A and ID_B, as long as it belongs to the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of admissible adversaries in the FS+ECL model (see Section 3.1). We emphasize that the adversary can still obtain the secret keys of the target users ID_A and ID_B for time periods $t^* > \tilde{t}$, which models *forward security*.

Variants of NIKE. Similar to [36], we consider different variants of NIKE depending on whether the setup algorithm just outputs a uniformly random coins or sample from some structured distributions. In particular, we say a NIKE scheme is:

- a *plain* NIKE, if NIKE.Setup(1^κ) just outputs (some specified number of) uniform random coins. In particular, NIKE.Setup$(1^\kappa; r) = r$.
- a NIKE in the *common reference string* model, if NIKE.Setup(1^κ) can be arbitrary (i.e., sample from an arbitrary distribution). In this case, we rely on *leak-free randomness* to run the setup algorithm.

Remark 5. We note that, in the original $\mathcal{PS}$ model of forward secure NIKE, there can be multiple honest users, and the adversary is allowed to obtain the secret keys of the

honest users other than the target users (even prior to the challenge time period $\tilde{t}$). In this work, we consider a simplified version of the $\mathcal{PS}$ model where there are only *two* honest users, and extend it to consider the ECL setting. The above simplified model can be shown to be polynomially equivalent to the full-fledged $\mathcal{PS}$ model by following the same reduction strategy as in [21][Theorem 8, Appendix B], where they show that the CKS-light model (with two honest users) is polynomially equivalent to the CKS-heavy model (where they can be multiple honest users). We emphasize that, in our application of constructing FS+ECLR PKE scheme from FS+ECLR NIKE, we only need the above simplified model.

4.4 Construction of FS + ECL NIKE scheme

In this section, we present our construction of forward-secure NIKE protocol resilient to entropic continual leakage in the common reference string model.

Let $i\mathcal{O}$ be an indistinguishability obfuscator for circuits, $\mathsf{pPRF} = (\mathsf{pPRF.keygen}, \mathsf{pPRF.puncture}, \mathsf{pPRF.eval})$ be a puncturable PRF with image space $\mathcal{Y} = \{0,1\}^y$, $\mathsf{LF} = (\mathsf{Inj}, \mathsf{Lossy}, f)$ be a (κ, k, m)-lossy function, and $\mathsf{LF}' = (\mathsf{Inj}', \mathsf{Lossy}', f')$ be a (κ', k', m')-lossy function, where $\kappa' \geq m$.

- $\mathsf{NIKE.Setup}(1^\kappa, T)$: Choose a random key $K \leftarrow \mathsf{pPRF.keygen}(1^\kappa)$. Sample two injective evaluation keys $\mathsf{ek} \leftarrow \mathsf{Inj}(1^\kappa)$, $\mathsf{ek}' \leftarrow \mathsf{Inj}'(1^\kappa)$. Consider the circuit $C(r, X_i, X_j)$ that has the key K hard-coded as defined below. Set $params = (\widehat{C} = i\mathcal{O}(C), \mathsf{ek}, \mathsf{ek}')$.

Circuit $C(r, X_i, X_j)$
Constant: K
Inputs: r, X_i, X_j.
  If $f_{\mathsf{ek}}(r) = X_i$ or $f_{\mathsf{ek}}(r) = X_j$, output $\mathsf{pPRF.eval}(K, (f'_{\mathsf{ek}'}(X_i), f'_{\mathsf{ek}'}(X_j))$.
  Else output $\bot$.

Fig. 3. Circuit $C(r, X_i, X_j)$

- $\mathsf{NIKE.Gen}(1^\kappa, params, ID_i)$: To compute the key pair of an user ID_i, sample $sk_1^i \xleftarrow{\$} \{0,1\}^\kappa$. Consider the circuit $C_i(sk_t, t)$ that has the keys $\mathsf{ek}, \mathsf{ek}'$, the base secret key sk_1^i, and the obfuscated circuit $\widehat{C}$ (which is part of $params$) hard-coded and defined below. Set the public key as $pk^i = \widehat{C_i} = i\mathcal{O}(C_i)$.

- $\mathsf{NIKE.Upd}(1^\kappa, sk_t^i)$: On input of the user $ID_i's$ secret key sk_t^i at time period t, computes sk_{t+1}^i, the secret key for the next time period $t + 1$. The instantiation of the update function is mentioned below.

- $\mathsf{NIKE.Key}(ID_i, pk^i = \widehat{C_i}, ID_j, sk_t^j)$: The user ID_j runs the obfuscated circuit $\widehat{C_i} = i\mathcal{O}(C_i)$ on inputs the secret key sk_t^j and the time period t to obtain the shared key shk_t^{ij} at time period t.

Circuit $C_i(sk_t, t)$
Constants: sk_1^i, ek, ek', $\hat{C}$, T.
Inputs: sk_t, t.

1. Check if $t \leq T$. If not, output $\perp$.
2. Update $sk_t^i = \text{NIKE.Upd}^{t-1}(sk_1^i)$.
3. Compute $X_t^i = f_{\text{ek}}(sk_t^i)$ and $X_t^j = f_{\text{ek}}(sk_t)$.

Output the shared key $shk_t^{ij} = \hat{C}(sk_t, X_t^i, X_t^j)$.

Fig. 4. Circuit $C_i(sk_t, t)$

Note on Update Function: The update function NIKE.Upd is one which takes a secret key of the current period and produces a new secret key. As defined in the security model, the adversary can issue leakage queries provided the keys are α-entropic conditioned on the set of non-local leakage queries. It is not hard to see that the update function should necessarily satisfy the one-wayness property, essentially guaranteeing the non-invertibility of the earlier keys once the secret key is exposed. Interestingly, for the above construction, we can abstract away the update function to any *entropic leakage resilient one-way function*, i.e., $\text{NIKE.Upd}(\cdot) = g(\cdot)$, where $g : \{0,1\}^\kappa \to \{0,1\}^\kappa$ be a ℓ-entropic leakage-resilient one-way function (ℓ-ELR-OWF). The definition of entropic leakage-resilient OWF is given in Section 2.5.

Correctness Analysis. It is not hard to see that both the parties ID_i and ID_j end up with the *same* shared key. The correctness of the computation of the shared key shk_t^{ij} by both the parties is shown below:

Shared key computation by P_i: P_i computes the shared key as:

$$
\begin{aligned}
shk_t^{ij} &= \text{NIKE.Key}(ID_j, pk^j = \widehat{C_j}, ID_i, sk_t^i) \\
&= \widehat{C_j}(sk_t^i, (X_t^i, X_t^j)) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(X_t^i), f'_{\text{ek}'}(X_t^j)\big) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(f_{\text{ek}}(sk_t^i)), f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^j)))\big) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^i))), f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^j)))\big)
\end{aligned}
$$

Shared key computation by P_j: P_j computes the shared key as:

$$
\begin{aligned}
shk_t^{ij'} &= \text{NIKE.Key}(ID_i, pk^i = \widehat{C_i}, ID_j, sk_t^j) \\
&= \widehat{C_i}(sk_t^j, (X_t^i, X_t^j)) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(X_t^i), f'_{\text{ek}'}(X_t^j)\big) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^i))), f'_{\text{ek}'}(f_{\text{ek}}(sk_t^j))\big) \\
&= \text{pPRF.eval}\big(K, f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^i))), f'_{\text{ek}'}(f_{\text{ek}}(\text{NIKE.Upd}^{t-1}(sk_1^j)))\big)
\end{aligned}
$$

Hence, we can see that shared keys computed by both P_i and P_j corresponding to time period t are *same*, i.e., $shk_t^{ij} = shk_t^{ij'}$.

Instantiations. Our FS+ECL NIKE construction can be instantiated based on a variety of assumptions, thanks to the recent constructions of $i\mathcal{O}$ from well-founded assumptions [8, 23, 30, 43]. Besides one can construct lossy functions from DDH or LWE [36]. In particular we obtain FS+ECL NIKE from either DDH or LWE along with any of these assumptions: (i) sub-exponential SXDH on asymmetric bilinear groups, sub-exponential LPN, Boolean PRGs in NC^0, (ii) sub-exponential circular-secure LWE, (iii) oblivious LWE sampleability (the construction from this assumption is still heuristic). Assumption (ii) combined with the instantiation of lossy functions from LWE gives us the *first* post-quantum secure FS-NIKE construction.

4.5 Security Proof of our **FS-ECLR-NIKE** construction.

In this section, we prove the following theorem statement.

Theorem 5. *Let κ be the security parameter. Assume that $i\mathcal{O}$ is an indistinguishability obfuscator for circuits, LF is an (κ, k, m)-lossy function, LF' is an (κ', k', m')-lossy function where $\kappa' \geq m$, pPRF is a family of puncturable PRFs with image size $\mathcal{Y} = \{0, 1\}^y$. Assume that $y \geq \alpha - \lambda - kT - k' + 2\kappa$, where T denote the total number of time period supported by the scheme. Then, Construction 4.4 is a (α, λ)-forward-secure entropic leakage-resilient NIKE in the $\mathcal{ECL}\text{-}\mathcal{PS}$ model, for $(\alpha, \kappa) = (\mathsf{poly}(\kappa), \mathsf{poly}'(\kappa))$, for arbitrary polynomials $\mathsf{poly}(\cdot)$ and $\mathsf{poly}'(\cdot)$ in κ.*

Our proof in is similar to the proof template of in [36], with some crucial differences. Let $\mathsf{Ext} : \{0, 1\}^n \times \{0, 1\}^d \to \mathcal{Y}$ be an average-case strong randomness extractor to be determined later.

We will now present the detailed proof of Theorem 5.

Proof. We prove the security of Theorem 5 via a sequence of hybrid experiments. We will use S_i to refer to the event that $\mathcal{A}_{\mathsf{ECL}}$ wins in $\mathbf{Game}_i$.

$\mathbf{Game}_0$. This corresponds to the original security game in the experiment $\mathsf{Exp}_{\mathcal{A}_{\mathsf{ECL}}}^{\mathsf{fs\text{-}ecl}}(\kappa)$, except that the challenger guesses that challenge time period $\tilde{t}$. The challenger does the following:

1. It chooses a time period $\tilde{t}$ as a guess for the challenge time period of the adversary $\mathcal{A}_{\mathsf{ECL}}$. If $\mathcal{A}_{\mathsf{ECL}}$ chooses any period other than $\tilde{t}$, the challenger aborts. Note that the guess is correct with probability $1/T$.
2. The challenger then runs NIKE.Setup to get $params = (\widehat{C} = i\mathcal{O}(C), \mathsf{ek}, \mathsf{ek}')$. The definition of the circuit C can be found in figure 3.
3. It also chooses two secret keys sk_1^i and sk_1^j. It sets $pk^i = \widehat{C_i} = i\mathcal{O}(C_i)$ where C_i is defined in figure 4.
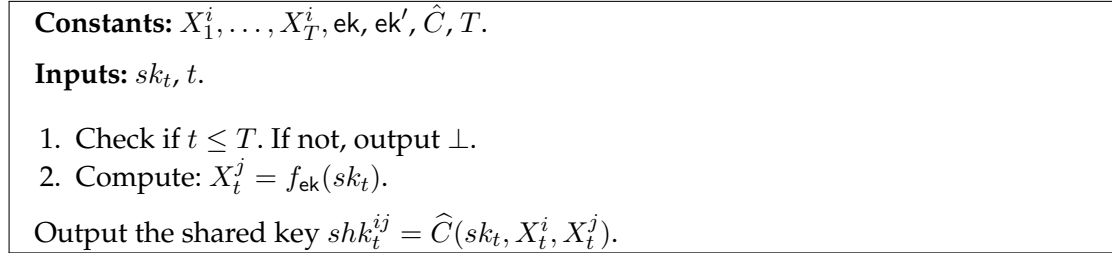4. It computes the challenge appropriately.

5. All the leakage information that $\mathcal{A}_{\mathsf{ECL}}$ asks can also be answered by the challenger with the knowledge of the secret keys sk_1^i and sk_1^j.

Let $\mathcal{E}$ be the event that the challenger guesses the correct challenge time period. Then,

$$\Pr[S_0] \le \Pr[S_0|\mathcal{E}] + \Pr[\bar{\mathcal{E}}] = \mathsf{Adv}^{\mathsf{fs\text{-}ecl}}_{\mathcal{A}_{\mathsf{nike}}}(\kappa) + \left(1 - \frac{1}{T}\right)$$

Game$_1$. This is similar to **Game$_0$**, except that the challenger sets $pk^i = \widehat{C_i^*}$ (resp. $pk^j = \widehat{C_j^*}$) where C_i^* (resp. C_j^*) is a program closely related to the original C_i (resp. C_j). The definition of C_i^* is defined in Figure 5. The difference in the programs are that the challenger, with knowledge of sk_1^i (resp. sk_1^j), computes the values of $X_1^i, \ldots, X_T^i$ (resp. $X_1^j, \ldots, X_T^j$) and embeds it into the program C_i^* instead of embedding sk_1^i. Note that, $X_t^i = f_{\mathsf{ek}}(\mathsf{NIKE.Upd}^{t-1}(sk_1^i))$. Hence, the two circuits C_i and C_i^* are functionally equivalent, and hence the distributions in **Game$_1$** and **Game$_0$** are computationally indistinguishable by security of $i\mathcal{O}$. In other words,

$$|\Pr[S_1] - \Pr[S_0]| \le \mathsf{negl}(\kappa).$$

Constants: $X_1^i, \ldots, X_T^i, \mathsf{ek}, \mathsf{ek}', \hat{C}, T$.

Inputs: sk_t, t.

1. Check if $t \le T$. If not, output $\bot$.
2. Compute: $X_t^j = f_{\mathsf{ek}}(sk_t)$.

Output the shared key $shk_t^{ij} = \widehat{C}(sk_t, X_t^i, X_t^j)$.

Fig. 5. Circuit $C_i^*(sk_t, t)$.

Game$_2$. This is similar to **Game$_1$** with the exception that there is a change in the generation of $params$. Instead of computing $\widehat{C} = i\mathcal{O}(C)$, the challenger computes $\widehat{C^{(1)}} = i\mathcal{O}(C^{(1)})$ where the program $C^{(1)}$ is closely related to C and is defined in Figure 6.

The difference in the programs are as follows: the challenger uses its guess of the challenge period $\tilde{t}$ to embed the values of $f'_{\mathsf{ek}'}(X_{\tilde{t}}^i)$ and $f'_{\mathsf{ek}'}(X_{\tilde{t}}^j)$ in the program $C^{(1)}$ using its knowledge of the secret keys sk_1^i, sk_1^j. Further, it also embeds the value of $\mathsf{pPRF.eval}\left(K, f'_{\mathsf{ek}'}(X_{\tilde{t}}^i), f'_{\mathsf{ek}'}(X_{\tilde{t}}^j)\right)$. Set $params = (\widehat{C^{(1)}}, \mathsf{ek}, \mathsf{ek}')$.

Note that the programs C and $C^{(1)}$ are functionally equivalent. This follows from the definitions of functions f and f', i.e., they are injective. Therefore, we have that the

distributions in $\mathbf{Game}_1$ and $\mathbf{Game}_2$ are computationally indistinguishable by security of $i\mathcal{O}$. In other words,

$$|\Pr[S_2] - \Pr[S_1]| \leq \mathsf{negl}(\kappa)$$

Constant: $K, \mathsf{ek}, \mathsf{ek}', f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}}), \mathsf{pPRF.eval}\left(K, f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}})\right)$

Inputs: r, X_i, X_j.

If $f'_{\mathsf{ek}'}(X_i) = f'_{\mathsf{ek}'}(X^i_{\hat{t}})$ and $f'_{\mathsf{ek}'}(X_j) = f'_{\mathsf{ek}'}(X^j_{\hat{t}})$, and:

   If $f'_{\mathsf{ek}'}(f_{\mathsf{ek}}(r)) = f'_{\mathsf{ek}'}(X^i_{\hat{t}})$ or $f'_{\mathsf{ek}'}(f_{\mathsf{ek}}(r)) = f'_{\mathsf{ek}'}(X^j_{\hat{t}})$:

      output $\mathsf{pPRF.eval}(K, (f'_{\mathsf{ek}'}(X_i), f'_{\mathsf{ek}'}(X_j))$.

If $f_{\mathsf{ek}}(r) = X_i$ or $f_{\mathsf{ek}}(r) = X_j$, output $\mathsf{pPRF.eval}(K, (f'_{\mathsf{ek}'}(X_i), f'_{\mathsf{ek}'}(X_j))$.

Else output $\bot$.

Fig. 6. Circuit $C^{(1)}(r, X_i, X_j)$.

Game$_3$. This is similar to **Game$_2$** with the exception that now we will puncture the pPRF at the points $f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}})$ while generating $params$, i.e.

$$K' \leftarrow \mathsf{pPRF.puncture}(K, f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}})) \ .$$

Instead of computing $\widehat{C^{(1)}} = i\mathcal{O}(C^{(1)})$, the challenger computes $\widehat{C^{(2)}} = i\mathcal{O}(C^{(2)})$ where the program $C^{(2)}$ is defined in Figure 7. The circuit $C^{(2)}$ has hard-coded in it the values the punctured key $K', \mathsf{ek}, \mathsf{ek}', f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}})$ and some *random* value $\gamma \leftarrow \mathcal{Y}$ (instead of $\mathsf{pPRF.eval}\left(K', f'_{\mathsf{ek}'}(X^i_{\hat{t}}), f'_{\mathsf{ek}'}(X^j_{\hat{t}})\right)$). Set $params = (\widehat{C^{(2)}}, \mathsf{ek}, \mathsf{ek}')$. Therefore, we have that the distributions in **Game$_3$** and **Game$_2$** are computationally indistinguishable by the security of punctured PRF. In other words,

$$|\Pr[S_3] - \Pr[S_2]| \leq \mathsf{negl}(\kappa).$$

Game$_4$. We again change how we generate the parameters $params$. The challenger first samples two seeds $s_i, s'_i \leftarrow \{0,1\}^d$.

- We first sample γ from $\mathcal{Y}$.
- We compute $z_i = \mathsf{Ext}(sk^i_{\hat{t}}, s_i) \oplus \gamma$ and $z_j = \mathsf{Ext}(sk^j_{\hat{t}}, s'_i) \oplus \gamma$.
- We now replace $i\mathcal{O}(C^{(2)})$ with $i\mathcal{O}(C^{(3)})$ where $C^{(3)}$ is a closely related program defined in Figure 8.

Constant: $K', \text{ek}, \text{ek}', f'_{\text{ek}'}(X^i_{\tilde{t}}), f'_{\text{ek}'}(X^j_{\tilde{t}}), \gamma$

Inputs: r, X_i, X_j.

If $f'_{\text{ek}'}(X_i) = f'_{\text{ek}'}(X^i_{\tilde{t}})$ and $f'_{\text{ek}'}(X_j) = f'_{\text{ek}'}(X^j_{\tilde{t}})$, and:

   If $f'_{\text{ek}'}(f_{\text{ek}}(r)) = f'_{\text{ek}'}(X^i_{\tilde{t}})$ or $f'_{\text{ek}'}(f_{\text{ek}}(r)) = f'_{\text{ek}'}(X^j_{\tilde{t}})$ output γ.

If $f_{\text{ek}}(r) = X_i$ or $f_{\text{ek}}(r) = X_j$, output $\text{pPRF.eval}(K', (f'_{\text{ek}'}(X_i), f'_{\text{ek}'}(X_j)))$.

Else output $\perp$.

Fig. 7. Circuit $C^{(2)}(r, X_i, X_j)$.

Note that the programs $C^{(2)}$ and $C^{(3)}$ are functionally equivalent. This follows from the definitions of functions f and f', i.e., they are injective, and consequently both $z_i \oplus \text{Ext}(r, s_i)$ and $z_j \oplus \text{Ext}(r, s'_i)$ evaluate to γ. Therefore, we have that the distributions in **Game$_4$** and **Game$_3$** are computationally indistinguishable by security of $i\mathcal{O}$. In other words,

$$|\Pr[S_4] - \Pr[S_3]| \leq \text{negl}(\kappa).$$

Constant: $K', \text{ek}, \text{ek}', f'_{\text{ek}'}(X^i_{\tilde{t}}), f'_{\text{ek}'}(X^j_{\tilde{t}}), z_i, z_j, s_i, s'_i$

Inputs: r, X_i, X_j.

If $f'_{\text{ek}'}(X_i) = f'_{\text{ek}'}(X^i_{\tilde{t}})$ and $f'_{\text{ek}'}(X_j) = f'_{\text{ek}'}(X^j_{\tilde{t}})$, and:

   If $f'_{\text{ek}'}(f_{\text{ek}}(r)) = f'_{\text{ek}'}(X^i_{\tilde{t}})$ output $z_i \oplus \text{Ext}(r, s_i)$

   If $f'_{\text{ek}'}(f_{\text{ek}}(r)) = f'_{\text{ek}'}(X^j_{\tilde{t}})$ output $z_j \oplus \text{Ext}(r, s'_i)$

If $f_{\text{ek}}(r) = X_i$ or $f_{\text{ek}}(r) = X_j$, output $\text{pPRF.eval}(K', (f'_{\text{ek}'}(X_i), f'_{\text{ek}'}(X_j)))$.

Else output $\perp$.

Fig. 8. Circuit $C^{(3)}(r, X_i, X_j)$.

Game$_5$ This is similar to **Game$_4$** with the only difference being in the way we generate $params$. In this game, the challenger samples ek and ek' from Lossy instead of Inj. This is now embedded in the new program $C^{(4)}$. The resulting distribution is indistinguishable from that of **Game$_4$** from the security of lossy functions. In other words,

$$|\Pr[S_5] - \Pr[S_4]| \leq \text{negl}(\kappa).$$

Game$_6$ This is similar to **Game$_5$** with the only difference being that we will rely on the definition of HILL entropy to switch to a distribution with min-entropy. Recall that $\mathcal{A}_{\mathsf{ECL}}$ adversary guarantees that each secret key conditioned on the non-local leakage queries have at least α computational entropy. Let $\mathcal{L}_{\mathsf{non\text{-}local}}$ denote the view of the adversary generated by all non-local queries, then we have that $\mathrm{H}^{\mathsf{HILL}}_{\epsilon,s}(SK_p|\mathcal{L}_{\mathsf{non\text{-}local}}) \geq \alpha$. By Definition 3, this means that there exists distributions Y_p over $\{0,1\}^\kappa$ such that $\widetilde{\mathrm{H}}_\infty(Y_p|\mathcal{L}_{\mathsf{non\text{-}local}}) \geq \alpha)$ and $(SK_p, \mathcal{L}_{\mathsf{non\text{-}local}}) \approx_c (Y_p, \mathcal{L}_{\mathsf{non\text{-}local}})$. Therefore, in this hybrid we switch to Y_p. By definition of HILL entropy we have that

$$|\mathsf{Pr}[S_6] - \mathsf{Pr}[S_5]| \leq \mathsf{negl}(\kappa).$$

Game$_7$ This is similar to **Game$_6$** except that the value returned as the shared key is entirely random. We will argue that **Game$_6$** and **Game$_7$** are statistically indistinguishable by relying on the security of Ext. Specifically, we prove the following claim.

Claim. The distribution of γ conditioned on $E := (pk^{(i)} = i\mathcal{O}(C_i^*), pk^{(j)} = i\mathcal{O}(C_j^*), params, \mathcal{L})$ is *statistically* indistinguishable from random. Here $\mathcal{L}$ denotes the view of the adversary based on the leakage queries it makes, and the responses it receives.

Proof. To prove the above claim, it is sufficient to prove that the values $z_i = \mathsf{Ext}(Y_{\bar{t}}^i, s_i) \oplus \gamma$ and $z_j = \mathsf{Ext}(Y_{\bar{t}}^j, s_i') \oplus \gamma$ look uniformly random. Notice that all the elements of E can be generated from

$$\left(K', \mathsf{ek}, \mathsf{ek}', \mathcal{L}, X_1^i, \ldots, X_T^i, X_1^j, \ldots, X_T^j, T, s_i, s_i', f'_{\mathsf{ek}'}(X_{\bar{t}}^i), f'_{\mathsf{ek}'}(X_{\bar{t}}^j)\right)$$

Further, note that the local leakage on the key would have at most length λ. Therefore, we now use the property of chain rule to conclude that

$$\widetilde{\mathrm{H}}_\infty(Y_{\bar{t}}^i|\mathcal{L})) \geq \alpha - \lambda$$

Also, note that we ensured that C_i^*, C_j^* did not have the secret key embedded in them. However, these do have the values of $X_1^i, \ldots, X_T^i$ and $X_1^j, \ldots, X_T^j$ which are dependent on the secret key chain. We need to estimate the loss of entropy due to this value. However, observe that we have switched the key $\mathsf{ek} \leftarrow \mathsf{Inj}$ to lossy function $\mathsf{ek} \leftarrow \mathsf{Lossy}$. By definition of lossy function, this means that

$$\widetilde{\mathrm{H}}_\infty(Y_{\bar{t}}^i|\mathcal{L}, X_1^i, \ldots, X_T^i, f'_{\mathsf{ek}'}(X_{\bar{t}}^i))) \geq \alpha - \lambda - kT - k'$$

To complete the proof, we will make use of the following Theorem.

Theorem 6 (Generalized Leftover Hash Lemma [19]). *Assume $\{H_x : \{0,1\}^\kappa \to \{0,1\}^y\}_{x\in X}$ be a family of universal hash functions. Then for any random variables W and I, we have:*

$$\delta\left((H_X(W), W, I), (U_y, W, I)\right) \leq \epsilon,$$

where $\epsilon \leq \frac{1}{2}\sqrt{2^{\widetilde{H}_\infty(W|I)}2^y}$ and $\delta(X, Y)$ denotes the statistical distance between two random variables X and Y.

In particular, universal hash functions are average-case (κ, n, y, ϵ)-*strong extractors whenever* $y \leq n - 2\log(\frac{1}{\epsilon}) + 2$.

Combining the above and the fact that $y \leq \alpha - \lambda - kT - k' + 2\kappa$, we obtain that z_i is statistically close to uniform in $\{0, 1\}^y$. A similar analysis can be done for the case of z_j. As a result, the distribution of γ is statistically close to uniform in $\{0, 1\}^y$. This concludes the proof of the claim. ∎

Finally, note that the adversary has no advantage in $\mathbf{Game}_7$. □

Parameters. We can set the value of κ appropriately, such that $\alpha = \mathsf{poly}(\kappa)$ and $\lambda = \mathsf{poly}'(\kappa)$ for arbitrary polynomials $\mathsf{poly}(\cdot)$ and $\mathsf{poly}'(\cdot)$ and any leakage rate $\rho \leq 1 - o(1)$.

4.6 FINDECL secure PKE from FS-ECLR-NIKE scheme

As a central application of forward-secure entropic leakage-resilient NIKE (FS-ECLR-NIKE), we show how to construct a IND-CCA secure key-evolving key encapsulation mechanism (KEM) resilient to ECL attacks, generically starting from any FS-ECLR-NIKE. From such a FINDCL-secure KEM it is easy to construct a FINDECL-secure PKE scheme using standard techniques. Hence, we focus on the construction of the FINDECL-secure KEM scheme. Our generic transformation essentially adapts the ideas of Freire et al. [21] to deal with forward security and entropic continual leakage.

The main idea of our transformation is as follows: the base public-secret key pair (pk, sk_1) of the FINDECL-secure KEM scheme is sampled using the key generation algorithm of the underlying FS-ECLR-NIKE scheme. The secret keys of the KEM scheme is updated using the key update algorithm of the FS-ECLR-NIKE scheme. To encrypt a message m for time period t, independently sample another base key pair (pk', sk_1') of the FS-ECLR-NIKE scheme and the secret key sk_1' is updated (using the key-update algorithm of FS-ECLR-NIKE) to the time period t, resulting in the key sk_t'. The public key pk' is also signed using a one-time signature scheme that binds the public key to its identity. The encapsulation key K is then generated by running the shared key generation algorithm of the FS-ECLR-NIKE with input (pk, sk_t'). The ciphertext consists of the randomly sampled public key pk' and the signature σ. The receiver can compute the same encapsulated key by running the shared key generation algorithm of the FS-ECLR-NIKE with the inputs (pk', sk_t) (where sk_t is the updated version of sk_1 to time period t), assuming the one-time signature verifies.

We now show the detailed construction of our key-evolving KEM scheme $\mathsf{KEM} = (\mathsf{KEM.Kg}, \mathsf{KEM.Upd}, \mathsf{KEM.Enc}, \mathsf{KEM.Dec})$. Before proceeding with the construction we present the syntax and security model for a forward-secure entropic leakage-resilient KEM scheme.

Modelling KEM in the FS+ECL Model. A key-evolving KEM scheme consists of the following algorithms $\mathsf{KEM} = (\mathsf{KEM.Kg}, \mathsf{KEM.Upd}, \mathsf{KEM.Enc}, \mathsf{KEM.Dec})$. The key generation algorithm $\mathsf{KEM.Kg}$ takes as input the security parameter 1^κ (in unary) and op-

tionally the parameters of the scheme (e.g., the maximum number of time periods supported by the scheme), and output a public-secret key pair (pk, sk_1), where sk_1 is secret key corresponding to the base time period. The key update algorithm KEM.Upd takes as input $(1^\kappa, pk, i)$ and secret key sk_i for time period i, and outputs a secret key sk_{i+1} for the next time period. Assume that the total number of time period supported by the scheme is $T \in \mathbb{N}$. The encapsulation algorithm KEM.Enc takes as input $(1^\kappa, pk, i)$ to return (c_i, K_i), where K_i is the encapsulated key and c_i is the ciphertext encrypting K_i under pk for time period $i \in [T]$. The decapsulation algorithm KEM.Dec takes $(1^\kappa, pk, i, sk_i, c_i)$ as input to return an output in $K_i \cup \{\bot\}$. Let us denote the encpasulated key sapce by $\mathcal{K}$.

Correctness: The correctness requirement for KEM states that: for all $\kappa \in \mathbb{N}$, for all $(pk, sk_1) \leftarrow$ KEM.Kg$(1^\kappa, T)$, all $sk_2, \cdots, sk_i$ satisfying $sk_j \leftarrow$ KEM.Upd$(1^\kappa, pk, j-1, sk_{j-1})$, $\forall\, 2 \le j \le i$, and all $c_i, K_i \leftarrow$ KEM.Enc$(1^\kappa, pk, i)$, it holds that KEM.Dec$(1^\kappa, pk, i, sk_i, c_i) = K_i$.

Forward Indistinguishability under entropic continual leakage (FINDECL).
Given a KEM scheme, consider the experiment $\mathrm{FINDECL}_{\mathsf{KEM}}^{\mathcal{A}_{\mathsf{kem}}}(\kappa, \alpha, \lambda, T)$ as defined in Figure 9, running between a challenger $\mathcal{C}$ and a PPT adversary $\mathcal{A}_{\mathsf{kem}} = (\mathcal{A}_1, \mathcal{A}_2)$, parametrized by the security parameter κ, the maximum leakage bound λ in a single time-period, an entropy parameter α, and the total number of time periods T that can be supported by the scheme. Similar to Definition 15, we define what it means for an adversary $\mathcal{A}_{\mathsf{kem}}$ to be "valid".

Definition 18 (Valid Adversary against the FINDECL Game of KEM). *We call an adversary $\mathcal{A}_{\mathsf{kem}} = (\mathcal{A}_1, \mathcal{A}_2)$ "valid" if the following holds true: (i) it makes at most one query to its Exp oracle and this is the last query, (ii) the adversary $\mathcal{A}_{\mathsf{kem}}$ belongs to the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of admissible adversaries (see Section 3.1 for the definition of the class $\mathcal{A}_{\mathsf{FS+ECL}}$.).*

Definition 19. (Forward indistinguishability under entropic continual leakage). *We say that* KEM = (KEM.Kg, KEM.Upd, KEM.Enc, KEM.Dec) *is (λ, α)-forward indistinguishable under entropic continual leakage ((λ, α)-FINDECL) with respect to the class $\mathcal{A}_{\mathsf{FS+ECL}}$ of adversaries, if the advantage defined as*
$\mathsf{Adv}_{\mathsf{KEM}, \mathcal{A}_{\mathsf{kem}}}^{\mathsf{findecl}}(\kappa) = \Pr[\mathrm{FINDECL}_{\mathsf{KEM}}^{\mathcal{A}_{\mathsf{pke}}}(\kappa, \alpha, \lambda, T) = 1]$ *(see Figure 9) is negligible for all "valid" PPT adversaries $\mathcal{A}_{\mathsf{kem}} = (\mathcal{A}_1, \mathcal{A}_2)$ (as defined above).*

The Construction. Let FS-ECLR-NIKE = (NIKE.Setup, NIKE.Gen, NIKE.Upd, NIKE.Key) be a *forward-secure entropic leakage-resilient* NIKE secure in the $\mathcal{ECL}\text{-}\mathcal{PS}$ model, with shared key space $\mathcal{SHK}$, and let OTS = (OTS.Gen, OTS.Sign, OTS.Vfy) be a *strong* existentially unforgeable *one-time* signature scheme. The construction of FINDECL KEM is shown in Figure 10.

Game FINDECL$_{\mathsf{KEM}}^{\mathcal{A}_{\mathsf{kem}}}(\kappa, \alpha, \lambda, T)$	$\underline{\mathsf{Up}()}$
$t \leftarrow 1 \,;\; t^* \leftarrow T + 1$	If $t < T$ then
$(pk, sk_1) \leftarrow\!\!\text{\$}\ \mathsf{KEM.Kg}(1^\kappa)$	$sk_{t+1} := \mathsf{KEM.Upd}(1^\kappa, pk, t, sk_t)$
$(i, state) \leftarrow\!\!\text{\$}\ \mathcal{A}_1^{\mathsf{Up,Lk,Exp,Dec}}(1^\kappa, pk)$	$\quad t \leftarrow t + 1$
$b \leftarrow\!\!\text{\$}\ \{0,1\} \,;\; (i, c, K_0) \leftarrow\!\!\text{\$}\ \mathsf{KEM.Enc}(1^\kappa, pk, i) \,;$	$\underline{\mathsf{Lk}(L)}$
$K_1 \leftarrow\!\!\text{\$}\ \mathcal{K} \,;$	Return $L(sk_t)$
$b' \leftarrow\!\!\text{\$}\ \mathcal{A}_2^{\mathsf{Up,Exp,Dec}}(1^\kappa, state, (i, c, K_b))$	
If not $(1 \leq i < t^*)$ then return false	$\underline{\mathsf{Exp}()}$
Return $(b' = b)$	$t^* \leftarrow t \,;$ Return sk_t
	$\underline{\mathsf{Dec}(t, c_t)}$
	If $(t, c_t) \neq (i, c)$
	$\quad$ Return $\mathsf{KEM.Dec}(1^\kappa, pk, (t, sk_t), c_t)$

Fig. 9. Game defining forward indistinguishability of key-evolving KEM scheme KEM under entropic continual leakage.

4.7 Security proof.

Theorem 7. *Let* FS-ECLR-NIKE *be a* (α, λ)-*entropic continual leakage-resilient forward-secure NIKE* ((α, λ)-FS-ECLR-NIKE) *secure in the* $\mathcal{ECL\text{-}PS}$ *model, and* OTS *be a strong existentially unforgeable one-time signature scheme. Then* KEM $=$ (KEM.Kg, KEM.Upd, KEM.Enc, KEM.Dec) *is a* (α, λ)-FINDECL-secure KEM *scheme.*

Proof. Let $\mathcal{A}_{\mathsf{kem}}$ be an adversary against the FINDECL-secure KEM scheme KEM. We now show how to use $\mathcal{A}_{\mathsf{kem}}$ to construct another adversary $\mathcal{A}_{\mathsf{nike}}$ against FS-ECLR-NIKE, thereby contradicting its $\mathcal{ECL\text{-}PS}$ security. $\mathcal{A}_{\mathsf{nike}}$ simulates the environment to $\mathcal{A}_{\mathsf{kem}}$ in the following way:

$\mathcal{A}_{\mathsf{nike}}$ on input $params$ picks one identity ID_i uniformly at random and runs OTS.Gen to obtain a key pair $(signk, vk)$. It then sets $ID_j = vk$ and makes two queries to its RegHon oracle, namely RegHon(ID_i) and RegHon(ID_j) queries to receive the public keys pk^i and pk^j respectively. $\mathcal{A}_{\mathsf{nike}}$ then returns $pk_{\mathsf{KEM}} = (params, ID_i, pk^i)$ to $\mathcal{A}_{\mathsf{kem}}$. When $\mathcal{A}_{\mathsf{kem}}$ makes an update query, $\mathcal{A}_{\mathsf{nike}}$ forwards the query to its own challenger and answers $\mathcal{A}_{\mathsf{kem}}$. On input a leakage query $f(\cdot, t)$ on ID_1 from $\mathcal{A}_{\mathsf{kem}}$, the adversary $\mathcal{A}_{\mathsf{nike}}$ queries it leakage oracle Lk(f, ID_i, t). It then forwards the answer to $\mathcal{A}_{\mathsf{kem}}$. When the adversary $\mathcal{A}_{\mathsf{kem}}$ enters into the the challenge phase for some time period $\tilde{t}$ (before the break-in period t^*), the adversary $\mathcal{A}_{\mathsf{nike}}$ also enters its own challenge phase and queries the tuple $(ID_i, ID_j, \tilde{t})$ to its challenger to receive a shared key $shk_{\tilde{t}}^{ij}$ for time period $\tilde{t}$. The key $shk_{\tilde{t}}^{ij}$ can either be a real key or a random key. $\mathcal{A}_{\mathsf{nike}}$ then sets the encapsulated key $K^* = shk_{\tilde{t}}^{ij}$, and computes the signature $\sigma^* \leftarrow$ OTS.Sign($signk, pk_j$). Finally, $\mathcal{A}_{\mathsf{nike}}$ sets $C_{\mathsf{KEM}} = (ID_j, pk_j, \sigma^*)$, and returns the challenge tuple (K^*, C_{KEM}) to $\mathcal{A}_{\mathsf{kem}}$. When the adversary $\mathcal{A}_{\mathsf{kem}}$ decides to break-in to some time period say t^*, $\mathcal{A}_{\mathsf{nike}}$ makes a query to its oracle Exp for the same time period t^* and returns the answer to $\mathcal{A}_{\mathsf{kem}}$.

1. KEM.Kg$(1^\kappa, (\lambda, \alpha), T)$: Run NIKE.Setup$(1^\kappa, (\lambda, \alpha), T)$ algorithm to obtain the system parameters $params$. It then chooses $ID \in \mathcal{IDS}$ uniformly at random, and runs NIKE.Gen$(1^\kappa, ID)$ to obtain a (base) key pair (pk, sk_1). It then sets $pk_{\mathsf{KEM}} = (params, ID, pk)$, and $sk_{\mathsf{KEM}} = (ID, sk_1)$.

2. KEM.Upd$(1^\kappa, pk, t, sk_t)$: The key update algorithm of KEM runs NIKE.Upd(sk_t) (where sk_t is the secret key for the current time period t) to obtain the updated key sk_{t+1} for the next time period $t + 1$.

3. KEM.Enc$(1^\kappa, pk_{\mathsf{KEM}}, t)$: This algorithm does the following:
    - Parses pk_{KEM} as $(params, ID, pk)$.
    - Runs $(signk, vk) \leftarrow$ OTS.Gen(1^κ), and repeats this until $vk \neq ID$.
    - Runs NIKE.Gen$(1^\kappa, vk = ID')$ to obtain another key pair (pk', sk_1'), and computes $\sigma \leftarrow$ OTS.Sign$(signk, pk')$ to obtain a signature σ on pk'.
    - Runs $sk_t' =$ NIKE.Upd$^{t-1}(sk_1')$ to obtain the updated key corresponding to time period t.
    - It then runs NIKE.Key(ID, pk, ID', sk_t') to obtain a shared key $K \in \mathcal{SHK}$.

    The output is the tuple $\big(K, C_{\mathsf{KEM}} = (pk', vk, \sigma)\big)$.

4. KEM.Dec$(1^\kappa, pk_{\mathsf{KEM}}, t, sk_1, C_{\mathsf{KEM}})$: This algorithm does the following:
    - Parse pk_{KEM} as $(params, ID, pk)$ and C_{KEM} as (pk', vk, σ).
    - Run OTS.Vfy(vk, pk', σ). If the signature does not verify, output $\perp$. Also, if $vk = ID$, output $\perp$.
    - Run $sk_t =$ NIKE.Upd$^{t-1}(sk_1)$ to the time period t.
    - Finally, run NIKE.Key$(vk = ID', pk', ID, sk_t)$ to obtain a shared key $K' \in \mathcal{SHK} \cup \{\perp\}$. (note that K' can also be $\perp$)

Fig. 10. Construction of FINDECL-secure KEM scheme.

$\mathcal{A}_{\mathsf{kem}}$ also makes Dec queries which can are handled by $\mathcal{A}_{\mathsf{nike}}$ as follows: For each decryption query of the form (t, C_t), parse C_t as (ID', pk', σ'). If $ID' = ID_j$, and $(pk', \sigma') \neq (pk_j, \sigma^*)$, then it is easy to see that we can build another adversary $\mathcal{A}_{\mathsf{OTS}}$ that breaks the strong existential unforgeability of the OTS scheme. If $ID' = ID_i$, it returns $\perp$. If $ID' \notin \{ID_i, ID_j\}$, $\mathcal{A}_{\mathsf{nike}}$ makes a call to its oracle RegCor(ID', pk') and then makes a call to the oracle Reveal(ID_i, ID', t) to get a shared key $K \in \mathcal{SHK}$ or $\perp$. It then returns the key K to $\mathcal{A}_{\mathsf{kem}}$.

This completes the description of $\mathcal{A}_{\mathsf{nike}}$'s simulation. Note that, the view of $\mathcal{A}_{\mathsf{kem}}$ is identical when playing either against $\mathcal{A}_{\mathsf{nike}}$ in this simulation or against a real IND-CCA secure FINDECL KEM challenger. When $\mathcal{A}_{\mathsf{kem}}$ outputs a guess b', $\mathcal{A}_{\mathsf{nike}}$ also outputs the same bit b' as guess for the shared key (whether it is real or random). Also, note that, the leakage tolerated by the KEM scheme is exactly the *same* as the amount of leakage

tolerated by the underlying FS-ECLR-NIKE NIKE scheme. This concludes the proof.

□

5 Signatures Schemes in the FS+ECL Model

In this section, we present our construction of a key-evolving signature scheme in the FS+ECL model. To this end, we first define and construct a new notion of *forward-secure entropic continual leakage-resilient identification* (FS-ECLR-ID) scheme in Sections 5.1 and 5.2 respectively. Our notion of FS-ECLR-ID schemes generalizes the prior definitions of ID schemes, which were either leakage-resilient or forward-secure, but not both. Later, in Section 5.3, we show how to transform such a FS-ECLR-ID scheme into a FUFECL-secure signature scheme using (generalized) Fiat-Shamir (FS) transform. This shows the applicability of FS transform even in the FS+ECL setting. The FUFECL signature scheme obtained via the FS-transform is secure in the RO model and can tolerate a leakage rate of $1/2 - o(1)$. However, one drawback of our construction is that, the resulting signature scheme can support an a-priori bounded (but arbitrary polynomial) number of time periods.

5.1 Forward-secure Entropic Leakage-resilient Identification schemes

An identification scheme is an interactive protocol that enables a prover (who holds a public-secret key pair) to prove its identity to a verifier. In a forward-secure identification scheme, the time is divided into discrete time periods, such that the secret key for period $i + 1$ can be computed from the secret key of period i. The public key remains the same in every time period. More formally, a forward-secure identification scheme consists of the five algorithms (ParamGen$_{ID}$, Gen$_{ID}$, Update$_{ID}$, $\mathcal{P}$, $\mathcal{V}$) as described below:

1. ParamGen$_{ID}(1^\kappa)$: The parameter generation algorithm takes as input the security parameter κ (in unary) and outputs a set of system parameters params, and the maximum number of time periods T supported by the system. The parameters params are taken as implicit input by all the algorithms.
2. Gen$_{ID}(1^\kappa, T)$: The key generation algorithm outputs a pair (pk, sk_1) containing the public key pk and a base secret key sk_1 for the first time period.
3. Update$_{ID}(pk, sk_t)$: The *deterministic* key update algorithm takes the secret key sk_t of the current time period t and outputs the secret key sk_{t+1} for the next time period $t + 1$, if sk_t is a secret key for time period $t < T$. We assume that the secret keys implicitly contain the information about the time periods they are associated with.
4. $\mathcal{P}(pk, sk_t)$: The prover algorithm takes as input the secret key sk_t for the current time period t, the current conversation transcript and the associated state and outputs the next message (if any) to be sent to the verifier.
5. $\mathcal{V}(pk, t)$: The deterministic verification algorithm takes the public key and the period t and outputs a decision in $\{accept, reject\}$ at the end of the protocol execution. We denote the interaction between $\mathcal{P}$ and $\mathcal{V}$ corresponding to time period t as $\{\mathcal{P}(pk, sk_t) \rightleftharpoons \mathcal{V}(pk, t)\}$.
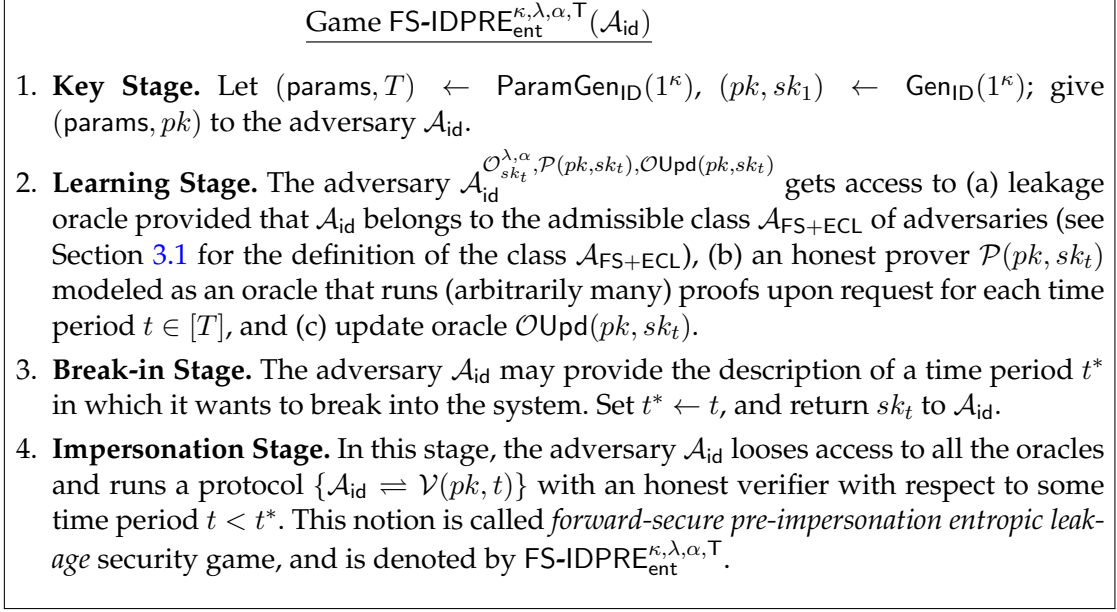
<div style="border:1px solid black; padding:10px;">

$$\underline{\text{Game FS-IDPRE}_{\mathsf{ent}}^{\kappa,\lambda,\alpha,\mathsf{T}}(\mathcal{A}_{\mathsf{id}})}$$

1. **Key Stage.** Let $(\text{params}, T) \leftarrow \mathsf{ParamGen}_{\mathsf{ID}}(1^\kappa)$, $(pk, sk_1) \leftarrow \mathsf{Gen}_{\mathsf{ID}}(1^\kappa)$; give (params, pk) to the adversary $\mathcal{A}_{\mathsf{id}}$.

2. **Learning Stage.** The adversary $\mathcal{A}_{\mathsf{id}}^{\mathcal{O}_{sk_t}^{\lambda,\alpha}, \mathcal{P}(pk,sk_t), \mathcal{O}\mathsf{Upd}(pk,sk_t)}$ gets access to (a) leakage oracle provided that $\mathcal{A}_{\mathsf{id}}$ belongs to the admissible class $\mathcal{A}_{\mathsf{FS+ECL}}$ of adversaries (see Section 3.1 for the definition of the class $\mathcal{A}_{\mathsf{FS+ECL}}$), (b) an honest prover $\mathcal{P}(pk, sk_t)$ modeled as an oracle that runs (arbitrarily many) proofs upon request for each time period $t \in [T]$, and (c) update oracle $\mathcal{O}\mathsf{Upd}(pk, sk_t)$.

3. **Break-in Stage.** The adversary $\mathcal{A}_{\mathsf{id}}$ may provide the description of a time period t^* in which it wants to break into the system. Set $t^* \leftarrow t$, and return sk_t to $\mathcal{A}_{\mathsf{id}}$.

4. **Impersonation Stage.** In this stage, the adversary $\mathcal{A}_{\mathsf{id}}$ looses access to all the oracles and runs a protocol $\{\mathcal{A}_{\mathsf{id}} \rightleftharpoons \mathcal{V}(pk, t)\}$ with an honest verifier with respect to some time period $t < t^*$. This notion is called *forward-secure pre-impersonation entropic leakage* security game, and is denoted by $\mathsf{FS\text{-}IDPRE}_{\mathsf{ent}}^{\kappa,\lambda,\alpha,\mathsf{T}}$.

</div>

Fig. 11. Attack game for FS-ECLR-ID scheme.

An ID scheme should satisfy the standard *completeness* property, i.e., if the prover is honest the (honest) verifier will accept the transcript generated by the interaction $\{\mathcal{P}(pk, sk_t) \rightleftharpoons \mathcal{V}(pk, t)\}$. We now define the security of forward-secure entropic leakage-resilient ID schemes. Informally, in the *learning stage* the adversary is given access to polynomially many "copies of the prover", and also access to the leakage and update oracles. The adversary can obtain leakage on each of the secret keys corresponding to each time period $i \in [T]$, provided the leakage functions satisfy the constraints of the FS+ECL model. The adversary may also break into the system and obtain the secret key sk_{t^*} for any time period t^*. After the learning stage, the adversary enters into an *impersonation stage* in which it either looses access to all the oracles (called *forward-secure pre-impersonation entropic leakage* security). In this stage, the adversary tries to impersonate the prover to the honest verifier with respect to some time period t prior to the break-in period t^*, and wins the game if it succeeds. The attack game is defined in Figure 11. The advantage of an adversary $\mathcal{A}_{\mathsf{id}}$ in the game $\mathsf{FS\text{-}IDPRE}_{\mathsf{ent}}^{\kappa,\lambda,\mathsf{T}}(\mathcal{A}_{\mathsf{id}})$ is the probability that the verifier $\mathcal{V}$ accepts in the impersonation stage.

Definition 20. *Let* $\mathsf{FS\text{-}ECLR\text{-}ID} = (\mathsf{ParamGen}_{\mathsf{ID}}, \mathsf{Gen}_{\mathsf{ID}}, \mathsf{Update}_{\mathsf{ID}}, \mathcal{P}, \mathcal{V})$ *be a forward secure identification scheme parametrized with the security parameter* κ, *leakage parameter* (α, λ), *number of time periods* T, *and satisfying perfect completeness. We say that the scheme is secure with* forward-secure pre-impersonation entropic leakage-resilient *with respect to the class* $\mathcal{A}_{\mathsf{FS+ECL}}$ *of adversaries, if the advantage of any PPT adversary* $\mathcal{A}_{\mathsf{id}}$ *in the game* $\mathsf{FS\text{-}IDPRE}_{\mathsf{ent}}^{\kappa,\lambda,\alpha,\mathsf{T}}(\mathcal{A}_{\mathsf{id}})$ *is negligible in* κ.

5.2 Construction of FS-ECLR-ID scheme.

In this section, we present our construction of forward-secure entropic continual leakage-resilient ID scheme.

THE SCHEME. We show that a forward-secure version of the generalized GQ identification scheme is secure against entropic continual key leakage attacks (see Figure 12). To analyze the scheme, we use the relation $\mathcal{R} = \{(pk, sk) : sk = (\rho, \omega_1, \cdots, \omega_\ell), pk = (g_1, \cdots, g_\ell, h),$ s.t. $h = \prod_{i=1}^{\ell} g_i^{\omega_i} \cdot \rho^e \mod N\}$, where $N = p \cdot q$ (p and q are prime numbers), e and d are chosen such that $e \cdot d = 1 \mod \phi(N)$, and e is a prime number.

- $\underline{\mathsf{ParamGen}_{\mathsf{ID}}(1^\kappa, 1^T)}$: Let $N = p \cdot q$, where p and q are two distinct ℓ_N-bit primes, and let $e_1, \cdots, e_T$ be distinct ℓ_e-bit primes, co-prime to $\phi(N) = (p-1)(q-1)$, chosen uniformly at random. Here, T denote the number of time periods supported by the system. Also, let $g_1, \cdots, g_\ell \xleftarrow{\$} (\mathbb{Z}_N^*)^\ell$ be generators of a prime-order cylic subgroup of $\mathbb{Z}_N^*$. Set params $:= (N, (e_1, \cdots, e_T), (g_1, \cdots, g_\ell), T)$. Also, let us denote $f_i = e_{i+1} \cdots e_T$, $f_T = 1$.

- $\underline{\mathsf{Gen}_{\mathsf{ID}}(\mathsf{params}, T)}$: This is the initial key generation algorithm. Choose $\rho, \omega_1, \cdots, \omega_\ell \xleftarrow{\$} \mathbb{Z}_N^* \times \mathbb{Z}_{e_1}^\ell$, and set $pk = (N, (e_1, \cdots, e_T), (g_1, \cdots, g_\ell), h)$, where $h = \prod_{j=1}^{\ell} g_j^{\omega_j} \cdot \rho^E \mod N$, and $E = \prod_{j=1}^{T} e_j$. Let $\rho_1 = \rho^{E/e_1}$ and $\rho_1' = \rho^{E/f_1}$. The secret key for the base time period is $sk_1 = (N, e_1, \rho_1, \rho_1')$.

- $\underline{\mathsf{Update}_{\mathsf{ID}}(pk, sk_i)}$: Parse $pk = (N, (e_1, \cdots, e_T), (g_1, \cdots, g_\ell), h)$ and $sk_i = (N, e_i, \rho_i, \rho_i')$, where $\rho_i = \rho^{E/e_i}$ and $\rho_i' = \rho^{E/f_i}$. Compute, $\rho_{i+1} = \rho_i'^{f_{i+1}}$ and $\rho_{i+1}' = \rho_i'^{e_{i+1}}$. Set $sk_{i+1} = (N, e_{i+1}, \rho_{i+1}, \rho_{i+1}')$.

- $\underline{\mathcal{P}(pk, sk_i), \mathcal{V}(pk, i)}$: The prover and the verifier run the following protocol corresponding to some time period i. The public key is $pk = (N, e_i, (g_1, \cdots, g_\ell), h)$, and the secret key $sk_i = (N, e_i, \rho_i, \rho_i')$. The goal of the prover $\mathcal{P}$ is to prove that h is a e_i-residue. Observe that, $\prod_{j=1}^{\ell} g_j^{\omega_j} \cdot \rho_i^{e_i} \mod N = h$.

    1. $\mathcal{P}$: Randomly chooses $\vec{a_i} = (a_i^{(1)} \cdots, a_i^{(\ell)}) \leftarrow \mathbb{Z}_{e_i}^\ell$, and $s_i \leftarrow \mathbb{Z}_N^*$. Then compute $\mathsf{com}_i := \prod_{j=1}^{\ell} g_j^{a_i^{(j)}} \cdot s_i^{e_i} \mod N$. Output $(\vec{a}, \mathsf{com}_i)$, and send com_i to $\mathcal{V}$.

    2. $\mathcal{V}$: Chooses a random $\mathsf{ch}_i \leftarrow \mathbb{Z}_{e_i}$ and send ch_i to $\mathcal{P}$.

    3. $\mathcal{P}$: Compute $\vec{z_i} = (\mathsf{ch}_i \cdot w_1 + a_i^{(1)}, \cdots, \mathsf{ch}_i \cdot w_\ell + a_i^{(\ell)})$, and $u_i = (s_i \cdot \rho_i^{\mathsf{ch}_i}) \mod N$. Output $\mathsf{resp}_i = (\vec{z_i} = (z_i^{(1)}, \cdots, z_i^{(\ell)}), u_i)$ and send resp_i to $\mathcal{V}$.

    The verifier $\mathcal{V}$ *accepts* if and only if $u_i^{e_i} \cdot \prod_{j=1}^{\ell} g_j^{z_i^{(j)}} = h^{\mathsf{ch}_i} \cdot \mathsf{com}_i \mod N$.

Fig. 12. Construction of forward-secure entropic continuous leakage-resilient ID scheme FS-ECLR-ID.

Remark 6. Note that, in our construction the sizes of the keys are *linear* in the number of time periods. This is because we store the exponents $e_1, \cdots, e_T$ in the public and the secret key. However, it is possible to have *constant* sized keys by computing the exponents using a *random oracle*, using techniques similar to [2, Sec. 5.1]. We refer the reader to [2] for the details.

Theorem 8. *Assuming the hardness of the* RSA representation problem, *the construction of our identification scheme* FS-ECLR-ID *as shown in* Figure 12 *is* $\big((\ell \log e_i + \log \phi(N) - \lambda), \lambda\big)$-forward-secure pre-impersonation entropic leakage-resilient, *where* $\lambda = \ell \log e - \kappa$.

Proof. To prove the above theorem, we first prove the following lemma:

Lemma 3. *The following three properties hold for our* FS-ECLR-ID *construction:*

1. *It is difficult to find a public key* pk *and two different secret keys* sk_i *and* sk_i' *for* pk *for any time period* $i \in [T]$. *In particular,*

$$\Pr[sk_i' \neq sk_i \ and \ (pk, sk_i), (pk, sk_i') \in \mathcal{R} \,|\, (pk, sk_i, sk_i') \leftarrow \mathcal{A}(\mathsf{params});$$
$$\mathsf{params} \leftarrow \mathsf{ParamGen}_{\mathsf{ID}}(1^\kappa))] \leq \mathsf{negl}(\kappa).$$

2. *The protocol* $\mathcal{P}, \mathcal{V}$ *is a* Σ *protocol for* $\mathcal{R}$.
3. *Let* $\mathsf{PK}, \mathsf{SK}_i$ *are random variables defined over the key pairs* (pk, sk_i) *for any time period* i. *Then it holds that* $\widetilde{\mathrm{H}}_\infty(\mathsf{SK}_i|\mathsf{PK}) \geq \ell \log \min\{e_1, \cdots, e_T\}$.

Using the properties of Lemma 3, we will complete the proof of Theorem 8.

Proof of Lemma 3. We now prove the three properties stated in Lemma 3.

`Proof of property 1`: The proof of property 1 follows in a straightforward manner from the RSA assumption. The secret key in our construction of the FS-ECLR-ID scheme for a particular time period is actually a RSA ℓ-representation of the public key pk. Property 1 then follows immediately from the hardness of finding two distinct representations for the same public key.

`Proof of property 2`: We now show that FS-ECLR-ID is a Σ protocol. In particular, we need to show that FS-ECLR-ID satisfies completeness, special-soundness and HVZK properties. The completeness is trivial to see.

Special soundness: We will show that given two accepting transcripts $(\mathsf{com}_i, \mathsf{ch}_i, \mathsf{resp}_i)$ and $(\mathsf{com}_i, \mathsf{ch}_i', \mathsf{resp}_i')$ with $\mathsf{ch}_i \neq \mathsf{ch}_i'$ for some time period $i \in [T]$, we can find another representation sk_i' corresponding to pk as follows.

Parse $\mathsf{com}_i = \prod_{j=1}^\ell g_j^{a_i^{(j)}} \cdot s_i^{e_i} \mod N$, $\mathsf{resp}_i = \big(\vec{z_i} = (z_i^{(1)}, \cdots, z_i^{(\ell)}), u_i\big)$ and $\mathsf{resp}_i' = \big(\vec{z_i'} = (z_i'^{(1)}, \cdots, z_i'^{(\ell)}), u_i'\big)$, where $z_i^{(j)} = \mathsf{ch}_i \cdot w_i + a_i^{(j)}$, $z_i'^{(j)} = \mathsf{ch}_i' \cdot w_i + a_i^{(j)}$, $u_i = (s_i \cdot \rho_i^{\mathsf{ch}_i})$ $\mod N$ and $u_i' = (s_i \cdot \rho_i^{\mathsf{ch}_i'}) \mod N$. Also, assume that $\mathsf{ch}_i < \mathsf{ch}_i'$. Then it is possible to extract all the values $(w_1, \cdots, w_\ell)$ by solving the above systems of linear equations,

namely by computing $w_j = \frac{(z_i'^{(j)} - z_i^{(j)})}{\Delta \mathsf{ch}_i}$, where $\Delta \mathsf{ch}_i = (\mathsf{ch}_i' - \mathsf{ch}_i)$. Similarly, another e_i-th residue can be extracted by computing $\rho_i' = \left(\frac{u_i'}{u_i}\right)^{1/(\Delta \mathsf{ch}_i)}$. Note that, the inverse exists since $\Delta \mathsf{ch}_i > 0$.

Honest Verifier Zero Knowledge: To prove this, we need to design a simulator Sim who is only given the public key and the challenges corresponding to each time period and has to produce transcripts that are identically distributed to the original transcripts for all the time periods. Without loss of generality, let us show the simulation for any time period $i \in [T]$. The simulator Sim is first given the public key $pk = \left(N, (e_1, \cdots, e_T), (g_1, \cdots, g_\ell), h\right)$, and has to produce a simulated transcript identical to $(\mathsf{com}_i, \mathsf{ch}_i, \mathsf{resp}_i)$. Sim then samples $\vec{z}_i = (z_i^{(1)}, \cdots, z_i^{(\ell)}) \xleftarrow{\$} \mathbb{Z}_{e_i}^\ell$ and $u_i \xleftarrow{\$} \mathbb{Z}_N^*$. The simulator then receives the challenge $\mathsf{ch}_i \in \mathbb{Z}_{e_i}$, and computes $\mathsf{com}_i = \frac{u_i^{e_i} \cdot \prod_{j=1}^\ell g_j^{z_i^{(j)}}}{h^{\mathsf{ch}_i}} \mod N$. It is trivial to see that the simulated transcript is identically distributed to the original transcript.

`Proof of property 3:` The length of a secret key sk_i is $|sk_i| = \ell \log e_i + \log \phi(N)$, and the public key $pk \in \mathbb{Z}_N^*$. Hence we have:

$$\widetilde{\mathrm{H}}_\infty(\mathsf{sk}_i|\mathsf{pk}) \geq \widetilde{\mathrm{H}}_\infty(\mathsf{sk}_i) - \log \phi(N) = \ell \log e_i \geq \ell \log \min\{e_1, \cdots, e_T\}.$$

The proofs of these properties proves Lemma 3. □

We now continue with the proof of Theorem 8. The proof of this part is similar to the proof of the leakage-resilient ID scheme, as shown in [4]. Suppose there is an adversary $\mathcal{A}_{\mathsf{id}}$ running in time t and having advantage ε in the game $\mathsf{FS\text{-}IDPRE}_{\mathsf{ent}}^{\kappa,\lambda,\mathsf{T}}(\mathcal{A})$. Then, we can construct another adversary $\mathcal{B}$ that runs in time $\approx 2t$ and

$$\Pr[sk_i' \neq sk_i \text{ and } (pk, sk_i), (pk, sk_i') \in \mathcal{R} \,|\, (pk, sk_i, sk_i') \leftarrow \mathcal{B}(\mathsf{params});$$
$$\mathsf{params} \leftarrow \mathsf{ParamGen}_{\mathsf{ID}}(1^\kappa))] \leq \varepsilon^2 - \frac{1}{\log \phi(N)} - 2^{-\kappa}$$

The adversary $\mathcal{B}$ chooses a random base key pair $(pk, sk_1) \leftarrow \mathsf{Gen}_{\mathsf{ID}}(\mathsf{params})$ and simulates the environment for the adversary $\mathcal{A}_{\mathsf{id}}$. $\mathcal{B}$ then gives pk to $\mathcal{A}_{\mathsf{id}}$ and uses sk_1 to simulate the leakage queries from the secret keys corresponding to each time period and also the prover oracle $\mathcal{P}(pk, sk_i)$. When $\mathcal{A}_{\mathsf{id}}$ reaches the impersonation stage corresponding to some time period t, $\mathcal{B}$ choose a fresh random challenge $\mathsf{ch}_t \leftarrow \mathbb{Z}_{e_t}$, and receives the transcript $(\mathsf{com}_t, \mathsf{ch}_t, \mathsf{resp}_t)$. Then $\mathcal{B}$ rewinds $\mathcal{A}_{\mathsf{id}}$ and sends a fresh random challenge $\mathsf{ch}_t' \leftarrow \mathbb{Z}_{e_t}$, which results in the transcript $(\mathsf{com}_t, \mathsf{ch}_t', \mathsf{resp}_t')$. If both the conversations are accepting and $\mathsf{ch}_t' \neq \mathsf{ch}_t$, then the special soundness property guarantees the existence of an extractor which can find a secret key sk_t' such that $(pk, sk_t') \in \mathcal{R}$. Let E_1^t be the event that the above happens and E_2^t denote the event that $\mathsf{ch}_t' = \mathsf{ch}_t$.

Claim. $\Pr[E_1^t] \geq \varepsilon^2 - \frac{1}{\log \phi(N)}$

Proof. This follows from a rather straightforward probabilistic argument.

Claim. $\Pr[E_2^t] \leq 2^{-\kappa}$

Proof. Let us think of an experiment $\mathcal{E}_0$ where the adversary $\mathcal{A}_{\mathsf{id}}$ gets access to all the oracles, namely $\mathcal{O}_{sk_t}^{\lambda(\kappa)}$, $\mathcal{P}(pk, sk_t)$ and $\mathcal{O}\mathsf{Upd}(pk, sk_t)$. Let $\mathcal{E}_1$ denote the same experiment as $\mathcal{E}_0$, except that the adversary $\mathcal{A}_{\mathsf{id}}$ is not given access to the prover oracle $\mathcal{P}(pk, sk_t)$, and let $\mathcal{E}_2$ be the experiment in which $\mathcal{A}_{\mathsf{id}}$ has access to only the public key pk. so, we have:

$$\widetilde{\mathrm{H}}_\infty(\mathsf{SK}_t | \mathcal{E}_0) \geq \widetilde{\mathrm{H}}_\infty(\mathsf{SK}_t | \mathcal{E}_1) - \lambda \geq \widetilde{\mathrm{H}}_\infty(\mathsf{SK}_t | \mathcal{E}_2) - \lambda = \widetilde{\mathrm{H}}_\infty(\mathsf{SK}_t | \mathsf{PK}) - \lambda$$
$$\geq \ell \log e_t - \lambda \geq \kappa.$$

where the first inequality follows from chain rule for min-entropy, the second inequality follows from Lemma 2, and the last inequality follows from property (3) of Lemma 3. The final inequality holds since $\lambda \leq \ell \log \min\{e_1, \cdots, e_T\} - \kappa$.

Note that, the secret keys for each of the time periods retain enough *min-entropy* in them, even given the leakage in that particular time period. Now, by the entropic continual leakage assumption, all the keys are unpredictable, even given the leakage across all the time periods.

Finally, observe that $\Pr[E_2^t] \leq 2^{\widetilde{\mathrm{H}}_\infty(\mathsf{SK} | \mathcal{E}_0)}$

Combining the above two claims, the first part of the theorem follows as:

$$\Pr[E_1^t \wedge \neg E_2^t] \geq \varepsilon^2 - \frac{1}{\log \phi(N)} - 2^{-\kappa} \geq \varepsilon^2 + \mathsf{negl}(\kappa.)$$

The proof for forward-secure anytime entropic leakage security is similar to above, with a subtle difference. In this setting, the leakage functions of the adversary $\mathcal{A}_{\mathsf{id}}$ may also depend on the challenge ch_t. Hence, while rewinding the adversary to fresh random challenge ch_t' the leakage functions may also depend on the new challenge ch_t'. Hence, for anytime leakage security the allowed leakage is half of the original leakage tolerated for pre-impersonation security. □

5.3 FUFECL signatures from FS-ECLR-ID schemes

In this section, we show how to transform any public-coin FS-ECLR-ID scheme into a FUFECL signature scheme using a generalized version of Fiat-Shamir (FS) transform [2]. More precisely, the signature in period i is just the signature obtained via FS transform using the secret key of the i^{th} period $sk_i = \mathsf{Update}_{\mathsf{ID}}^{i-1}(pk, sk_1)$ (with period i included in the random oracle input). The amount of leakage tolerated by the signature scheme is exactly the *same* as the leakage tolerated by the underlying identification scheme. We now present the details of the construction.

THE CONSTRUCTION. Let $\mathsf{FS\text{-}ECLR\text{-}ID} = (\mathsf{ParamGen}_{\mathsf{ID}}, \mathsf{Gen}_{\mathsf{ID}}, \mathsf{Update}_{\mathsf{ID}}, \mathcal{P}, \mathcal{V})$ be a forward-secure entropic leakage-resilient identification scheme. Let H be a hash function modeled as a random oracle, and let $\mathsf{KES} = (\mathsf{KES.Kg}, \mathsf{KES.Upd}, \mathsf{KES}.$

Sign, KES.Vfy) be the signature scheme obtained via Fiat-Shamir transform applied to FS-ECLR-ID as shown in Fig. 13.

$\mathsf{KES.Kg}(1^\kappa, 1^T)$: Run params $\leftarrow \mathsf{ParamGen_{ID}}(1^\kappa, 1^T)$; return params.

$\mathsf{KES.Sign}(1^\kappa, vk, sk_i, m)$: Compute: (1) $\mathsf{com}_i \xleftarrow{\$} \mathcal{P}(vk, sk_i)$; (2) $\mathsf{ch}_i \leftarrow \mathsf{H}(\mathsf{com}_i, m, i)$; and (3) $\mathsf{resp}_i \leftarrow \mathcal{P}(vk, sk_i, \mathsf{com}_i, \mathsf{ch}_i)$. Finally, return the signature $\sigma_i \leftarrow (\mathsf{com}_i, \mathsf{resp}_i, i)$

$\mathsf{KES.Upd}(1^\kappa, vk, i, sk_i)$: Run $sk_{i+1} \leftarrow \mathsf{Update_{ID}}(vk, sk_i)$; return sk_{i+1}.

$\mathsf{KES.Vfy}(1^\kappa, vk, m, i, \sigma_i)$: Parse $(\mathsf{com}_i, \mathsf{resp}_i, i) \leftarrow \sigma_i$; compute $\mathsf{ch}_i \leftarrow \mathsf{H}(\mathsf{com}_i, m, i)$, and $d \leftarrow \mathcal{V}(vk, \mathsf{com}_i, \mathsf{ch}_i, \mathsf{resp}_i)$. Return d.

Fig. 13. Generalized Fiat-Shamir transform for forward-secure entropic leakage-resilient signature

Theorem 9. *Let* FS-ECLR-ID $= (\mathsf{ParamGen_{ID}}, \mathsf{Gen_{ID}}, \mathsf{Update_{ID}}, \mathcal{P}, \mathcal{V})$ *be a* three-round public-coin (α, λ)-forward-secure anytime entropic leakage-resilient *identification scheme. Then the signature scheme* KES $= (\mathsf{KES.Kg}, \mathsf{KES.Upd}, \mathsf{KES.Sign}, \mathsf{KES.Vfy})$ *obtained by the generalized Fiat-Shamir transform applied to* FS-ECLR-ID *is* $(\alpha, \frac{\lambda}{2})$-FUFECL-secure.

Proof Sketch. The main idea of the proof follows from [4]. In our case, however, we need to guess the time period i of the signature output by the adversary, in order to embed the challenge correctly, hence resulting in a loss of factor T in the security reduction. The leakage queries of the FUFECL adversary can be easily handled by the adversary of the FS-ECLR-ID scheme by querying its own leakage oracle. This concludes the proof.

6 Encryption scheme in the FS+CL model

In the section, we consider the problem of constructing forward-secure encryption scheme in the continual leakage model (dubbed as FS+CL as in [6]). In the FS+CL model, the update process is randomized and the leakage happens according to the CL model (i.e., bounded leakage between two successive invocations). We refer to the reader to [6] for the details of the FS+CL model. To construct such a forward-indistinguishable encryption scheme secure in the CL model (FINDCL-secure PKE), we first introduce a notion of *continual leakage-resilient binary tree encryption* (CLR-BTE) (see below), which can be seen as a restricted version of CLR hierarchical identity-based encryption (CLR-HIBE). The construction of CLR-BTE follows in a straightforward manner from the CLR-HIBE scheme of Lewko et al. [35] (which is based on static assumptions over composite-order bilinear groups). For appropriate choice of parameters, the CLR-HIBE scheme achieves the optimal leakage rate of $1 - o(1)$. Hence, our CLR-BTE scheme also achieves the same leakage rate. Finally, we show how to transform such a CLR-BTE scheme to a FINDCL-secure encryption scheme using the Canetti-Halevi-Katz (CHK) transform. This approach of constructing FINDCL-secure encryption scheme was already suggested in [6]. However, it was intuitively claimed in [6] that the above approach does

not work, due to the following reason: "The problem is that FS+CL security of the resulting scheme requires that multiple nodes of the BTE construction can be leaked on jointly, whereas the CL security of HIBE only buys us leakage on each such node individually." Surprisingly, we prove the contrary and show that, indeed, it is possible to simulate the joint leakage by leaking on a single node. This requires a careful analysis of the CHK transform in the FS+CL setting.

6.1 Continual Leakage-resilient Binary Tree Encryption

In this section, we introduce the notion of continual leakage-resilient binary tree encryption (CLR-BTE). Our security model of CLR-BTE generalizes the definition of binary tree encryption (BTE) (as proposed by Canetti et al. [10]) in the setting of continual leakage. A BTE can be seen as a restricted version of HIBE, where the identity tree is represented as a *binary* tree.[6] In particular, as in HIBE, a BTE is also associated with a "master" public key MPK corresponding to a tree, and each node in the tree has their respective secret keys. To encrypt a message for a node, one specifies the identity of the node and the public key MPK. The resulting ciphertext can be decrypted using the secret key of the target node. Moreover, the secret key of any node can be used to derive the secret keys of its children.

Definition 21. *(Continual leakage-resilient BTE).* A continual leakage-resilient binary tree encryption (CLR-BTE) is a tuple of the PPT algorithms (Gen, Der, Enc, Dec) such that:

1. The key generation algorithm Gen takes as input the security parameter κ and a value ℓ for the depth of the tree. It returns a master public key MPK and an initial (root) secret key SK_ε.
2. The key derivation algorithm Der takes as input MPK, the identity of a node $w \in \{0,1\}^{\leq \ell}$, and its secret key SK_w. It returns secret keys SK_{w_0}, SK_{w_1} for the two children of w.
3. The encryption algorithm Enc takes as input MPK, the identity of a node $w \in \{0,1\}^{\leq \ell}$ and a message M to return a ciphertext C.
4. The decryption algorithm Dec takes as input MPK, the identity of a node $w \in \{0,1\}^{\leq \ell}$, its secret key SK_w, and a ciphertext C. It returns a message M or $\perp$ (to denote decryption failure).

We require the standard correctness requirement, i.e., for all (MPK, SK_ε) output by Gen, any node $w \in \{0,1\}^{\leq \ell}$, and secret key SK_w correctly generated for this node, and any message M, we have $M = \mathsf{Dec}(MPK, w, SK_w, \mathsf{Enc}(MPK, w, M))$.

We now present our security model for CLR-BTE. Our model generalizes the notion of *selection-node chosen-plaintext attacks* (SN-CPA) put forward by Canetti et al. [10] to define the security of BTE. In our model, the adversary first specifies the identity of the

---

[6] Recall that, in HIBE the tree can have arbitrary degree.

target node[7] $w^* \in \{0, 1\}^{\leq \ell}$. The adversary receives the public key MPK and the secret keys of all the nodes that do not trivially allow him/her to derive the secret key of w^*[8]. Besides, the adversary is also allowed to leak *continuously* from the secret keys of all the nodes that lie on the path from the root node and w^* (including both). The goal of the adversary is then to win the indistinguishability game with respect to the target node w^*.

Definition 22. A CLR-BTE scheme is secure against *continual leakage selective-node, chosen-plaintext attacks* ($\lambda(\kappa)$-CLR-SN-CPA) if for all polynomially-bounded functions $\ell(\cdot)$, and leakage bound $\lambda(\kappa)$, the advantage of any PPT adversary $\mathcal{A}$ in the following game is negligible in the security parameter κ:

1. The adversary $\mathcal{A}(1^\kappa, \ell)$ outputs the name of a node $w^* \in \{0, 1\}^{\leq \ell}$. We will denote the path from the root node to the target node w^* by P_{w^*}.

2. The challenger runs the algorithm $\mathsf{Gen}(1^\kappa, \ell)$ and outputs (MPK, SK_ε). In addition, it runs $\mathsf{Der}(\cdot, \cdot, \cdot)$ to generate the secret keys of all the nodes on the path P_{w^*}, and also the secret keys for the two children w_0^* and w_1^*. The adversary is given MPK and the secret keys $\{SK_w\}$ for all nodes w of the following form:

   – $w = w'\bar{b}$, where $w'b$ is a prefix of w^* and $b \in \{0, 1\}$ (i.e., w is a sibling of some node in P_{w^*}).

   – $w = w_0^*$ or $w = w_1^*$ (i.e., w is a child of w^*; this is only when $|w^*| < \ell$).

   The challenger also creates a set $\mathcal{T}$ that holds tuples of all the (node) identities, secret keys and the number of leaked bits from each key so far.

3. The adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ may also ask leakage queries. The adversary $\mathcal{A}$ provides the description of a probabilistic leakage function h with constant output size acting on the set of keys, and an identity of a node w in the path P_{w^*} (that may also include both the root note and the target node w^*). The challenger scans $\mathcal{T}$ to find the tuple with identity w. It should be of the form (w, SK_w, L). The challenger then checks if $L + |h(SK_w)| \leq \lambda(\kappa)$. If this is true, it responds with $h(SK_w)$ and updates the L in the tuple with $L = L + |h(SK_w)|$. If the check fails, it returns $\perp$ to the adversary.

4. The adversary $\mathcal{A}$ then sends two messages M_0 and M_1 to the challenger such that $|M_0| = |M_1|$. The challenger samples a random bit $b \xleftarrow{\$} \{0, 1\}$, and computes $C^* \leftarrow \mathsf{Enc}(MPK, w^*, M_b)$. It then returns C^* to the adversary $\mathcal{A}$. The adversary is not allowed to ask any further leakage queries after receiving the challenge ciphertext C^*.[9]

---

[7] Note that, this model where the adversary specifies the target node w^* ahead of time is weaker than the model where the adversary may choose the target *adaptively* (analogous to the adaptive security of HIBE schemes). However, as we will show, this model already suffices to construct of a forward-secure CLR encryption scheme.

[8] In particular, the adversary receives the secret keys of all the nodes that are siblings of all the nodes that are on the path from the root node to the target node w^*.

[9] If the adversary is allowed to ask leakage queries after receiving the challenge ciphertext, it can encode the entire decryption algorithm of C^* as a function on a secret key, and thus win the game trivially.

At the end of this game, the adversary outputs a bit $b' \in \{0, 1\}$; it succeeds if $b' = b$. The advantage of the adversary is the absolute value of the difference between its success probability and $1/2$.

Construction of CLR-BTE scheme. Our construction of the CLR-BTE scheme essentially follows in a straightforward manner from the continuous leakage-resilient HIBE (CLR-HIBE) construction of Lewko et al. [35], tuned to the setting of a binary tree. The resulting CLR-BTE is *adaptively secure*, since the CLR-HIBE of [35] enjoys security against adaptive adversaries employing the dual-system encryption technique. The security of the CLR-BTE scheme can be proven under static assumptions over composite-order bilinear groups. We refer the readers to [35] for the details of the CLR-HIBE construction and its proof. For appropriate choice of parameters, the CLR-HIBE scheme achieves the optimal leakage rate of $1 - o(1)$.

6.2 FINDCL encryption from CLR-BTE scheme

In this section, we show that a generic construction of a FINDCL-secure encryption scheme from any CLR-BTE scheme. The main idea of our construction is very simple: use the Canetti-Halevi-Katz (CHK) transform [10] to the underlying CLR-BTE scheme to construct a FINDCL encryption scheme. In particular, we show the applicability of the CHK transform[10] even in the setting of continuous leakage. However, as we show later, the analysis of the CHK transform in the setting of leakage turns out to be quite tricky.

Let $(\mathsf{Gen}, \mathsf{Der}, \mathsf{Update}, \mathsf{Enc}, \mathsf{Dec})$ be a CLR-BTE scheme. We construct our FINDCL PKE scheme $(\mathsf{KEE.Kg}, \mathsf{KEE.Upd}, \mathsf{KEE.Enc}, \mathsf{KEE.Dec})$ as shown below. The construction is identical to the CHK transform, with the underlying building blocks appropriately changed.

Some additional notations: To obtain a FINDCL-secure encryption scheme with $T = 2^\ell - 1$, time periods (labeled through 1 to T), we use a CLR-BTE of depth ℓ. We associate the time periods with all nodes of the tree according to a pre-order traversal. The node associated with time period i is denoted by w^i. In a pre-order traversal, $w^1 = \varepsilon$ (the root node), if w^i is an internal node then $w^{i+1} = w^i 0$ (i.e., left child of w^i). If w^i is a leaf node and $i < T - 1$ then $w^{i+1} = w'1$, where w' is the longest string such that $w'0$ is a prefix of w^i.

1. $\mathsf{KEE.Kg}(1^\kappa, T)$: Run $\mathsf{Gen}(1^\kappa, \ell)$, where $T \leq 2^\ell - 1$, and obtain (MPK, SK_ε). Set $pk = (MPK, T)$, and $sk_1 = SK_\varepsilon$.

2. $\mathsf{KEE.Upd}(1^\kappa, pk, i, sk_i)$: The secret key sk_i organized as a stack of node keys, with the secret key SK_{w^i} on top. We first pop this key off the stack. If w^i is a leaf node, the next node on top of the stack is $SK_{w^{i+1}}$. If w^i is an internal node, compute

---

[10] The original CHK transform [10] is used to construct a forward-secure PKE scheme starting from a BTE scheme.

$(SK_{w^i0}, SK_{w^i1}) \leftarrow \mathsf{Der}(pk, w^i, SK_{w^i})$ and push SK_{w^i1} and then SK_{w^i0} onto the stack. In either case, the node SK_{w^i} is erased.

3. $\mathsf{KEE.Enc}(pk, i, m)$: Run $\mathsf{Enc}(pk, w^i, m)$. Note that w^i is publicly computable given i and T.

4. $\mathsf{KEE.Dec}(1^\kappa, pk_i, sk_i, c_i)$: Run $\mathsf{Dec}(pk, w, SK_{w^i}, c_i)$. Note that, SK_{w^i} is stored as part of sk_i.

Theorem 10. *Let* $\lambda : \mathbb{N} \to [0, 1]$. *Let* $\Pi = (\mathsf{Gen}, \mathsf{Der}, \mathsf{Update}, \mathsf{Enc}, \mathsf{Dec})$ *be a* $\lambda(\kappa)$-CLR-SN-CPA *continual leakage-resilient binary-tree encryption* (CLR-BTE) *scheme. Let* $\ell : \mathbb{N} \to \mathbb{N}$ *be a polynomial such that* $T \leq 2^{\ell-1}$. *Then* $\Pi' = (\mathsf{KEE.Kg}, \mathsf{KEE.Upd}, \mathsf{KEE.Enc}, \mathsf{KEE.Dec})$ *is a* $\lambda(\kappa)$-FINDCL *secure encryption scheme supporting up to* T *time periods.*

Proof Sketch. Our proof follows the template of the CHK transformation for converting a BTE scheme to forward-secure encryption scheme, with the crucial difference in simulating the leakage queries. In particular, we show that if there exists a $\lambda(\kappa)$-bounded valid adversary $\mathcal{A}_{\mathsf{kee}}$ that breaks the $\lambda(\kappa)$-FINDCL security of Π', we can build another $\lambda(\kappa)$-bounded valid adversary $\mathcal{A}_{\mathsf{clr-bte}}$ breaking the $\lambda(\kappa)$-CLR-SN-CPA security of Π. The adversary $\mathcal{A}_{\mathsf{clr-bte}}$ uses $\mathcal{A}_{\mathsf{kee}}$ in a black-box manner. It is very easy for $\mathcal{A}_{\mathsf{clr-bte}}$ to simulate key generation, update and encryption queries asked by $\mathcal{A}_{\mathsf{kee}}$. The adversary $\mathcal{A}_{\mathsf{clr-bte}}$ knows the secret keys of all the nodes that are right siblings of the nodes that lie in the path P_{w^*} from the root node to w^* (the target node). Besides, it also knows the secret keys of both the children of w^{*}[11]. Hence, $\mathcal{A}_{\mathsf{clr-bte}}$ can itself simulate the update queries asked by $\mathcal{A}_{\mathsf{kee}}$.

However, $\mathcal{A}_{\mathsf{kee}}$ may also ask leakage queries. We partition the nodes of the binary tree into two disjoint sets and simulate the leakage queries on each of these two sets differently. In particular, $\mathcal{A}_{\mathsf{kee}}$ may ask a leakage query on a node w such that w does not lie in the path P_{w^*}, i.e, $w \notin P_{w^*}$, or it may ask leakage query on a node w that lie on the path P_{w^*} (including the root node and w^*), i.e, $w \in P_{w^*}$. Simulation of leakage queries on all the nodes $w \notin P_{w^*}$ is trivial, since $\mathcal{A}_{\mathsf{clr-bte}}$ already knows the secret keys of all such nodes. Hence, for these nodes it can simulate the leakage queries by itself. However, for all nodes $w \in P_{w^*}$, $\mathcal{A}_{\mathsf{clr-bte}}$ does not know their secret keys. Let us denote the path from the root to node w as P_w, which is certainly a prefix of the path P_{w^*}. One may think that the leakage on such nodes $w \in P_{w^*}$ can be simulated by $\mathcal{A}_{\mathsf{clr-bte}}$ by simply querying the leakage oracle of the challenger of the CLR-BTE scheme. However, this is not true, since the challenger of the CLR-BTE scheme expects a function that leaks on each node of the tree *individually*, rather than leaking *jointly* on multiple nodes of the tree. In particular, the secret key sk_w of any node w in the PKE scheme is a tuple of key components, namely $sk_w = (SK_w, \{SK_{rs(P_w)}\})$, rather than a single key component[12], where $\{SK_{rs(P_w)}\}$ denote the secret keys corresponding to all the right siblings of the nodes that lie on the path P_w. However, at this point, the key observation is that the

---

[11] Recall, in the CLR-SN-CPA security game (please see Def. 22) the adversary gets all these secret keys.

[12] Recall that, the secret key of any node w in our construction contains the secret key of w, i.e., SK_w, and also the keys corresponding to all right siblings of the nodes on the path P_w.

adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ already knows all the secret key components $\{SK_{rs(P_w)}\}$, except the key component SK_w. This is because it knows the secret keys of all the right siblings of the nodes that lie in the path P_{w^*}, and hence also the secret keys of all the right siblings of the nodes that lie in the path P_w. To simulate the leakage f on sk_w, the adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ now modifies the function f into a related leakage function f' that acts only on the secret key component SK_w of sk_w, and at the same time is consistent with the output of f. Thus, the joint leakage on all the nodes is transformed to a leakage on the single node w. The way that we accomplish this is that: when $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ receives the leakage function f from $\mathcal{A}_{\mathsf{kee}}$, it hardwires the secret keys $\{SK_{rs(P_w)}\}$ into the function f.[13] $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ then sends this modified function f' to its own challenger. Hence, with this leakage information and the full knowledge of the other keys of sk_w, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ can consistently simulate the joint leakage by just leaking on a single node. The formal proof follows.

Proof. Assume that we have an adversary $\mathcal{A}_{\mathsf{kee}}$ with advantage $\epsilon(\kappa)$ in an $\lambda(\kappa)$-FINDCL security game of $\Pi' = (\mathsf{KEE.Kg}, \mathsf{KEE.Upd}, \mathsf{KEE.Enc}, \mathsf{KEE.Dec})$. We construct an adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ that obtains an advantage $\epsilon(\kappa)/T$ in the corresponding attack against the underlying the CLR-BTE scheme $\Pi = (\mathsf{Gen}, \mathsf{Der}, \mathsf{Update}, \mathsf{Enc}, \mathsf{Dec})$. The leakage rate tolerated by Π is exactly the same as Π'. We now describe how $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ simulates the environment for $\mathcal{A}_{\mathsf{kee}}$:

1. $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ chooses uniformly at random a time period $i^* \in [T]$ and outputs w^{i^*} (the identity of the node corresponding to i^*). $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ then obtains MPK and $\{SK_w\}$ for all the appropriate nodes w[14] from its challenger. $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ then sets $pk = (MPK, T)$, and forwards the public key pk to the adversary $\mathcal{A}_{\mathsf{kee}}$.

2. When $\mathcal{A}_{\mathsf{kee}}$ decides to break into the system, it provides the time period, say j. If $j \leq i^*$, then $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ outputs a random bit and halts. Otherwise, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ computes the appropriate secret key sk_j and gives it to $\mathcal{A}_{\mathsf{kee}}$. Note that, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ can efficiently compute the secret keys sk_j for any $j > i^*$ from the knowledge of $\{SK_w\}$ (the set of secret keys received in Step 1).

3. $\mathcal{A}_{\mathsf{kee}}$ may ask leakage queries on the secret key of any node w in the tree.[15] This node can either be of any one of the following types: (1) $w \notin P_{w^{i^*}}$ or (2) $w \in P_{w^{i^*}}$, where $P_{w^{i^*}}$ is the path containing the nodes from the root node to the target node w^{i^*} (including both).
   For the first case, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ repeatedly runs the algorithm Der with the appropriate secret keys in the set $\{SK_w\}$ to derive the secret key sk_i (corresponding to node w^i). Hence, any leakage query asked on sk_i for $i > i^*$ can be simulated by $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ by simply computing the corresponding secret key and answering the leakage function.

---

[13] Note that this is possible to do since $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ has full knowledge of the secret key component $\{SK_{rs(P_w)}\}$ of sk_w.

[14] Recall that $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ receives the secret keys of all the nodes that are right siblings of the nodes that lie on the path P from the root node to w^{i^*}.

[15] Practically, it will only ask for leakage on sk_i for any $i < j$, where j is the period of break-in. This is because, for any $i > j$, the adversary can itself compute the secret key. However, we consider the general case, where all the node are prone to leakage.

For the second case, i.e, when the leakage function is asked on a node $w^i \in P_{w^{i*}}$ [16], the adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ does not know sk_i or the secret key of any of the ancestors of w^i. The secret key sk_i can be seen a stack of node keys (derived using the underlying CLR-BTE scheme) with the key SK_{w^i} on top of the stack. The other node keys in the stack are secret keys corresponding to the right siblings of all the nodes in the path P_{w^i} from the root node to w^i. Note that, the adversary $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ already knows all of these node keys, since the path P_{w^i} is a prefix of $P_{w^{i*}}$. Let us denote $sk_i = (SK_{w^i}, \{SK\}_{rs(P_{w^i})})$, where $\{SK\}_{rs(P_{w^i})}$ denote the secret keys of the right siblings of all nodes in path P_{w^i}. The adversary now does the following:

- Receive as input the leakage function f from $\mathcal{A}_{\mathsf{kee}}$. Modify the description of the function as $h = f_{\{SK\}_{rs(P_{w^i})}}(\cdot) = f(\cdot, \{SK\}_{rs(P_{w^i})})$. In other words, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ hardwires the secret keys $\{SK\}_{rs(P_{w^i})}$ in the function f, and forwards h as the leakage function to its challenger.
- On input the answer $h(SK_{w^i}, \{SK\}_{rs(P_{w^i})})$ from its challenger, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ forwards this answer as the output of the leakage function f to $\mathcal{A}_{\mathsf{kee}}$.

It is clear that $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ perfectly simulates the answers to the leakage queries of the adversary $\mathcal{A}_{\mathsf{kee}}$, regardless of which node w the leakage query is asked.

4. When $\mathcal{A}_{\mathsf{kee}}$ asks a challenge query with input (i, m_0, m_1), if $i \neq i^*$ then $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ outputs a random bit and halts. Otherwise, it forwards the tuple (m_0, m_1) to its challenger and obtains the challenge ciphertext C^*. It then gives C^* to $\mathcal{A}_{\mathsf{kee}}$.

5. When $\mathcal{A}_{\mathsf{kee}}$ outputs b', $\mathcal{A}$ outputs b' and halts.

It is easy to see that, if $i = i^*$, the above simulation by $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ is perfect. Since, $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ guesses i^* with probability $1/T$, we have that $\mathcal{A}_{\mathsf{clr\text{-}bte}}$ correctly predicts the bit b with advantage $\epsilon(\kappa)/T$. $\qquad\square$

7 Conclusion

In this work, we propose the ECL and FS+ECL models as means to construct cryptographic primitives in the continual leakage model with *deterministic* key update procedures. Some of the key open problems left by our work are:

- Construct FS+ECL-secure NIKE and ID schemes that can support *unbounded* number of time periods (currently our constructions can handle only bounded (but an arbitrary polynomial) number of time periods).
- Construct efficient FS+ECL-secure NIKE schemes from standard assumptions. Currently, our FS+ECL NIKE scheme relies on $i\mathcal{O}$ and one-way functions. The reliance on $i\mathcal{O}$ does not seem to be inherent; yet there does not seem to be straightforward way to get a construction without using it.
- Finally, one could also propose alternative security models that capture deterministic key updates and at the same time enable secure and efficient constructions of different cryptographic primitives in these models.

---

[16] Note that, in this case $i < i^*$, since we follow a pre-order traversal.

References

1. M. Abdalla, J. H. An, M. Bellare, and C. Namprempre. From identification to signatures via the Fiat-Shamir transform: Minimizing assumptions for security and forward-security. In L. R. Knudsen, editor, *EUROCRYPT 2002*, volume 2332 of *LNCS*, pages 418–433. Springer, Heidelberg, Apr. / May 2002.

2. M. Abdalla, F. Benhamouda, and D. Pointcheval. On the tightness of forward-secure signature reductions. *Journal of Cryptology*, pages 1–67, 2018.

3. A. Akavia, S. Goldwasser, and V. Vaikuntanathan. Simultaneous hardcore bits and cryptography against memory attacks. In O. Reingold, editor, *TCC 2009*, volume 5444 of *LNCS*, pages 474–495. Springer, Heidelberg, Mar. 2009.

4. J. Alwen, Y. Dodis, and D. Wichs. Leakage-resilient public-key cryptography in the bounded-retrieval model. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 36–54. Springer, Heidelberg, Aug. 2009.

5. M. Bellare and S. K. Miner. A forward-secure digital signature scheme. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 431–448. Springer, Heidelberg, Aug. 1999.

6. M. Bellare, A. O'Neill, and I. Stepanovs. Forward-security under continual leakage. *In 16th International Conference on Cryptology And Network Security*, 2017.

7. E. Boyle, G. Segev, and D. Wichs. Fully leakage-resilient signatures. In K. G. Paterson, editor, *EUROCRYPT 2011*, volume 6632 of *LNCS*, pages 89–108. Springer, Heidelberg, May 2011.

8. Z. Brakerski, N. Döttling, S. Garg, and G. Malavolta. Factoring and pairings are not necessary for io: Circular-secure lwe suffices. Cryptology ePrint Archive, Report 2020/1024, 2020. https://eprint.iacr.org/2020/1024.

9. Z. Brakerski, Y. T. Kalai, J. Katz, and V. Vaikuntanathan. Overcoming the hole in the bucket: Public-key cryptography resilient to continual memory leakage. In *51st FOCS*, pages 501–510. IEEE Computer Society Press, Oct. 2010.

10. R. Canetti, S. Halevi, and J. Katz. A forward-secure public-key encryption scheme. In E. Biham, editor, *EUROCRYPT 2003*, volume 2656 of *LNCS*, pages 255–271. Springer, Heidelberg, May 2003.

11. J.-S. Coron, C. Gentry, S. Halevi, T. Lepoint, H. K. Maji, E. Miles, M. Raykova, A. Sahai, and M. Tibouchi. Zeroizing without low-level zeroes: New MMAP attacks and their limitations. In R. Gennaro and M. J. B. Robshaw, editors, *CRYPTO 2015, Part I*, volume 9215 of *LNCS*, pages 247–266. Springer, Heidelberg, Aug. 2015.

12. J.-S. Coron, M. S. Lee, T. Lepoint, and M. Tibouchi. Cryptanalysis of GGH15 multilinear maps. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 607–628. Springer, Heidelberg, Aug. 2016.

13. R. Cramer, I. Damgård, and B. Schoenmakers. Proofs of partial knowledge and simplified design of witness hiding protocols. In Y. Desmedt, editor, *CRYPTO'94*, volume 839 of *LNCS*, pages 174–187. Springer, Heidelberg, Aug. 1994.

14. D. Dachman-Soled, S. D. Gordon, F.-H. Liu, A. O'Neill, and H.-S. Zhou. Leakage-resilient public-key encryption from obfuscation. In C.-M. Cheng, K.-M. Chung, G. Persiano, and B.-Y. Yang, editors, *PKC 2016, Part II*, volume 9615 of *LNCS*, pages 101–128. Springer, Heidelberg, Mar. 2016.

15. Ö. Dagdelen and D. Venturi. A second look at Fischlin's transformation. In D. Pointcheval and D. Vergnaud, editors, *AFRICACRYPT 14*, volume 8469 of *LNCS*, pages 356–376. Springer, Heidelberg, May 2014.

16. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Cryptography against continuous memory attacks. In *51st FOCS*, pages 511–520. IEEE Computer Society Press, Oct. 2010.

17. Y. Dodis, K. Haralambiev, A. López-Alt, and D. Wichs. Efficient public-key cryptography in the presence of key leakage. In M. Abe, editor, *ASIACRYPT 2010*, volume 6477 of *LNCS*, pages 613–631. Springer, Heidelberg, Dec. 2010.

18. Y. Dodis, A. B. Lewko, B. Waters, and D. Wichs. Storing secrets on continually leaky devices. In R. Ostrovsky, editor, *52nd FOCS*, pages 688–697. IEEE Computer Society Press, Oct. 2011.

19. Y. Dodis, L. Reyzin, and A. Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. In C. Cachin and J. Camenisch, editors, *EUROCRYPT 2004*, volume 3027 of *LNCS*, pages 523–540. Springer, Heidelberg, May 2004.

20. S. Dziembowski, T. Kazana, and D. Wichs. Key-evolution schemes resilient to space-bounded leakage. In P. Rogaway, editor, *CRYPTO 2011*, volume 6841 of *LNCS*, pages 335–353. Springer, Heidelberg, Aug. 2011.

21. E. S. V. Freire, D. Hofheinz, E. Kiltz, and K. G. Paterson. Non-interactive key exchange. In K. Kurosawa and G. Hanaoka, editors, *PKC 2013*, volume 7778 of *LNCS*, pages 254–271. Springer, Heidelberg, Feb. / Mar. 2013.

22. S. Garg, C. Gentry, S. Halevi, and D. Wichs. On the implausibility of differing-inputs obfuscation and extractable witness encryption with auxiliary input. In J. A. Garay and R. Gennaro, editors, *CRYPTO 2014, Part I*, volume 8616 of *LNCS*, pages 518–535. Springer, Heidelberg, Aug. 2014.

23. R. Gay and R. Pass. Indistinguishability obfuscation from circular security. Cryptology ePrint Archive, Report 2020/1010, 2020. https://eprint.iacr.org/2020/1010.

24. L. C. Guillou and J.-J. Quisquater. A "paradoxical" indentity-based signature scheme resulting from zero-knowledge. In S. Goldwasser, editor, *CRYPTO'88*, volume 403 of *LNCS*, pages 216–231. Springer, Heidelberg, Aug. 1990.

25. J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten. Lest we remember: cold-boot attacks on encryption keys. *Commun. ACM*, 52(5):91–98, 2009.

26. S. Halevi and H. Lin. After-the-fact leakage in public-key encryption. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 107–124. Springer, Heidelberg, Mar. 2011.

27. J. Håstad, R. Impagliazzo, L. A. Levin, and M. Luby. A pseudorandom generator from any one-way function. *SIAM Journal on Computing*, 28(4):1364–1396, 1999.

28. S. Hohenberger, V. Koppula, and B. Waters. Adaptively secure puncturable pseudorandom functions in the standard model. In T. Iwata and J. H. Cheon, editors, *ASIACRYPT 2015, Part I*, volume 9452 of *LNCS*, pages 79–102. Springer, Heidelberg, Nov. / Dec. 2015.

29. C.-Y. Hsiao, C.-J. Lu, and L. Reyzin. Conditional computational entropy, or toward separating pseudoentropy from compressibility. In M. Naor, editor, *EUROCRYPT 2007*, volume 4515 of *LNCS*, pages 169–186. Springer, Heidelberg, May 2007.

30. A. Jain, H. Lin, and A. Sahai. Indistinguishability obfuscation from well-founded assumptions. Cryptology ePrint Archive, Report 2020/1003, 2020. https://eprint.iacr.org/2020/1003.

31. J. Katz and V. Vaikuntanathan. Signature schemes with bounded leakage resilience. In M. Matsui, editor, *ASIACRYPT 2009*, volume 5912 of *LNCS*, pages 703–720. Springer, Heidelberg, Dec. 2009.

32. P. C. Kocher. Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In N. Koblitz, editor, *CRYPTO'96*, volume 1109 of *LNCS*, pages 104–113. Springer, Heidelberg, Aug. 1996.

33. P. C. Kocher, J. Jaffe, and B. Jun. Differential power analysis. In M. J. Wiener, editor, *CRYPTO'99*, volume 1666 of *LNCS*, pages 388–397. Springer, Heidelberg, Aug. 1999.

34. A. B. Lewko, M. Lewko, and B. Waters. How to leak on key updates. In L. Fortnow and S. P. Vadhan, editors, *43rd ACM STOC*, pages 725–734. ACM Press, June 2011.

35. A. B. Lewko, Y. Rouselakis, and B. Waters. Achieving leakage resilience through dual system encryption. In Y. Ishai, editor, *TCC 2011*, volume 6597 of *LNCS*, pages 70–88. Springer, Heidelberg, Mar. 2011.

36. X. Li, F. Ma, W. Quach, and D. Wichs. Leakage-resilient key exchange and two-seed extractors. In D. Micciancio and T. Ristenpart, editors, *Advances in Cryptology - CRYPTO 2020 - 40th Annual International Cryptology Conference, CRYPTO 2020, Santa Barbara, CA, USA, August 17-21, 2020, Proceedings, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 401–429. Springer, 2020.

37. E. Miles, A. Sahai, and M. Zhandry. Annihilation attacks for multilinear maps: Cryptanalysis of indistinguishability obfuscation over GGH13. In M. Robshaw and J. Katz, editors, *CRYPTO 2016, Part II*, volume 9815 of *LNCS*, pages 629–658. Springer, Heidelberg, Aug. 2016.

38. M. Naor. On cryptographic assumptions and challenges (invited talk). In D. Boneh, editor, *CRYPTO 2003*, volume 2729 of *LNCS*, pages 96–109. Springer, Heidelberg, Aug. 2003.

39. M. Naor and G. Segev. Public-key cryptosystems resilient to key leakage. In S. Halevi, editor, *CRYPTO 2009*, volume 5677 of *LNCS*, pages 18–35. Springer, Heidelberg, Aug. 2009.

40. T. Okamoto. Provably secure and practical identification schemes and corresponding signature schemes. In E. F. Brickell, editor, *CRYPTO'92*, volume 740 of *LNCS*, pages 31–53. Springer, Heidelberg, Aug. 1993.

41. D. Pointcheval and O. Sanders. Forward secure non-interactive key exchange. In M. Abdalla and R. D. Prisco, editors, *SCN 14*, volume 8642 of *LNCS*, pages 21–39. Springer, Heidelberg, Sept. 2014.
42. A. Sahai and B. Waters. How to use indistinguishability obfuscation: deniable encryption, and more. In D. B. Shmoys, editor, *46th ACM STOC*, pages 475–484. ACM Press, May / June 2014.
43. H. Wee and D. Wichs. Candidate obfuscation via oblivious lwe sampling. Cryptology ePrint Archive, Report 2020/1042, 2020. https://eprint.iacr.org/2020/1042.