A Note on Various Forking Lemmas

Chelsea Komlo

September 22, 2022

1 Introduction

First introduced by Pointcheval and Stern [3], the forking lemma is commonly used in proofs of security that require *rewinding* an adversary in order to demonstrate a reduction to breaking some known-to-be-hard mathematical problem. In this note, we review the original forking lemma and several variations thereof. Each variation is made in the effort to provide a tighter proof of security for the context that it is used, or better abstraction as to allow for more general applications.

The intuition for the forking lemma is as follows. We begin with an adversary that is modeled as a probabilistic Turing machine $\mathcal{A}$ that is initialized with a random tape and access to a hash function modeled as a random oracle. While the behavior of the adversary is generally not defined (making the adversary "black box," the adversary outputs some value that will either satisfy some pre-defined conditions (thus winning the security game), or not satisfy these conditions. Note this setup requires the assumption that a hash function can be simulated by a truly random function, which is known as the "random oracle model" [4].

If $\mathcal{A}$ completes its attack successfully, we can lower bound the probability that $\mathcal{A}$ again completes successfully in a second execution with the same random tape but *different* outputs from the random oracle. Determining this lower bound on the success probability of $\mathcal{A}$ across two executions is important for many proofs of security, as the two outputs are then employed to demonstrate the reduction (i.e, solve for the discrete logarithm of a challenge).

In this note, we begin by first reviewing the original forking lemma by Pointcheval and Stern. We then look at several subsequent variants, and how each variant allows for better abstraction (and hence applications beyond signature schemes), tighter bounds in the security proof, or tailoring to the context of a specific security proof.

2 Forking Lemma by Pointcheval and Stern

The first variant of the forking lemma was introduced by Pointcheval and Stern in their proof of security for Schnorr signatures [3]. The lemma is with respect

to a digital signature scheme in the random oracle model. More specifically, the lemma assumes an adversary modeled as a probabilistic Turing machine given access to a random tape (source of randomness) and a hash function modeled as a (programmable) random oracle. Then, the lemma demonstrates that if the adversary outputs a valid forgery σ for random oracle output h in its first execution, then with non-negligible probability, the adversary will output a valid second forgery with respect to a different different random oracle output h' with non-negligible success. Here, non-negligible specifically means the probability of success must be greater than $1/f(n)$, where $f(n)$ is any polynomial function.

Lemma 2.1 Forking Lemma. *Let $\mathcal{A}$ be a probabilistic polynomial time Turing machine given only public data as input. If $\mathcal{A}$ can find with non-negligible probability a valid signature σ with respect to a message m and random oracle output h, then, with non-negligible probability, if $\mathcal{A}$ is executed a second time with the same source of randomness but a different random oracle, then $\mathcal{A}$ will output a second valid forgery σ' for the same message m but with respect to a random oracle output h', such that $h \neq h'$.*

Pointcheval and Stern later give a different variant of the forking lemma [4] with respect to the adversary's probability of success ϵ and time bound T.

2.1 Applications

The forking lemma allows for demonstrating a reduction to discrete log for Schnorr signatures. Recall that a Schnorr signature is the tuple $\sigma = (R, z)$, such that $R = g^r$ and $z = r + c \cdot x$, where x is the secret key and $Y = g^x$ is the public key.

In the random oracle model, the simulator receives Y as the challenge and can simulate signing with respect to Y by programming the random oracle. After the simulator runs $\mathcal{A}$ and receives two valid forgeries $\sigma = (R, z)$ and $\sigma' = (R, z')$ with respect to Y, the simulator can output x by deriving:

$$\frac{z - z'}{(c - c')}$$

where $c = H(R, m)$ in the first execution of the adversary, and $c' = H(R, m)$ in the second execution of the adversary, but critically, such that $c \neq c'$. The fact that $c \neq c'$ cannot be detected by the adversary though, as it is executed twice, without keeping state between each execution.

3 General Forking Lemma

Unlike the forking lemma by Pointcheval and Stern, the general forking lemma introduced by Bellare and Neven [2] abstracts away the details of signature schemes and random oracles. Instead, the lemma simply asserts on the probability of some output of an algorithm when run on two related inputs.

Let $q \geq 1$ be an integer, and H be a set of size q_h, where $q_h \geq 2$. Let $X \in \mathbb{G}$, such that $X \xleftarrow{\$} \mathsf{IG}$, where IG is the instance generator. As before, let $\mathcal{A}$ be a randomized algorithm, that outputs $i \in \{\perp\} \cup \{1, \ldots, q\}$, and auxiliary output aux. We denote the general forking algorithm as $\mathsf{F_g}$, and show it below.

Algorithm $\mathsf{F_g}(X)$

Pick coins ρ for $\mathcal{A}$ at random.
$\{h_1, \ldots, h_q\} \xleftarrow{\$} H^q$
$\mathsf{Q} \leftarrow \{h_1, \ldots, h_q\}$
$(i, \mathsf{aux}) \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}; \rho)$
return $\perp$ **if** $i = \perp$
$\{h'_i, \ldots, h'_q\} \xleftarrow{\$} H$
$\mathsf{Q}' \leftarrow \{h_1, \ldots, h_{i-1}\} \cup \{h'_i, \ldots, h'_q\}$
// Q *and* Q' *differ at* $(q - i + 1)$ *points*
$(i', \mathsf{aux}') \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}'; \rho)$
return $\perp$ **if** $i = \perp$
if $i = i'$ **and** $h_i \neq h'_i$
  return $(1, \mathsf{aux}, \mathsf{aux}')$
return $\perp$

Let $\mathsf{acc}(X)$ denotes the probability that $\mathcal{A}$ completes successfully. I.e, acc denotes that $\mathcal{A}$ outputs a value that is accepted in its first execution in $\mathsf{F_g}$ when given the input X. We show this probability in Equation 1.

$$\mathsf{acc}(X) = \Pr\left[\, i \geq 1 \;:\; \{h_1, \ldots, h_1\} \xleftarrow{\$} H \;;\; (i, \mathsf{aux}) \xleftarrow{\$} \mathcal{A}(X, \{h_1, \ldots, h_1\}) \,\right] \quad (1)$$

Let fork denote the probability that $\mathsf{F_g}$ returns successfully given some $X \xleftarrow{\$} \mathsf{IG}$, as in Equation 2:

$$\mathsf{fork} = \Pr\left[\, b = 1 \;:\; X \xleftarrow{\$} \mathsf{IG} \;;\; (b, \mathsf{aux}, \mathsf{aux}') \xleftarrow{\$} \mathsf{F_g}(X) \,\right] \quad (2)$$

Bellare and Neven then demonstrate that Equations 1 generalize to any $X \xleftarrow{\$} \mathsf{IG}$. So Lemma 3.1 can be proved with respect to simply acc, as follows:

Lemma 3.1 *Given* acc *and* fork*, as defined, then* fork *is lower bounded by* acc *as in Equation 3.*

$$\mathsf{fork} \geq \mathsf{acc} \cdot \left(\frac{\mathsf{acc}}{q} - \frac{1}{q_h} \right) \quad (3)$$

Alternatively:

$$\mathsf{acc} \leq \frac{q}{q_h} + \sqrt{q \cdot \mathsf{fork}} \quad (4)$$

3.1 Applications

While the general forking lemma models the adversary $\mathcal{A}$ as what is executed by the forking algorithm directly, when employed in a proof, there is need for an *intermediate* adversary which simulates the environment to the actual adversary which attacks the scheme. So in practice, $\mathsf{F}_{\mathsf{g}}(X)$ will execute some adversary $\mathcal{B}$, which will set up and simulate the environment (generally with respect to some challenge), such as a signing protocol. $\mathcal{B}$ will then execute $\mathcal{A}$ which actually attacks the scheme, interacting with $\mathcal{B}$ in order to do so.

Why the need for all of this indirection? We need to argue that $\mathcal{A}$ interacts with the scheme in a way that is *indistinguishable* from a real run of the protocol. The job of $\mathcal{B}$ is to ensure that, using the inputs of $\mathsf{F}_{\mathsf{g}}(X)$ to do so.

4 Local Forking Lemma

Unlike the generalized forking algorithm F_{g} where Q and Q' differ by $q - i$ elements, the local forking lemma [1] is such that the sets employed as input to $\mathcal{A}$ differ by only a *single* element. Intuitively, this translates to the random oracle being re-programmed at only a single point after the fork, as opposed to every point after the fork. In particular, the oracle is re-programmed only at the index i output by $\mathcal{A}$ in its first execution. Why is this useful? The bounds in the proof of security can then be tighter, which allows for a more accurate representation of the security of the scheme.

We show the local forking algorithm below:

Algorithm $\mathsf{F}_{\mathsf{l}}(X)$

Pick coins ρ for $\mathcal{A}$ at random.
$\{h_1, \ldots, h_q\} \xleftarrow{\$} H^q$
$\mathsf{Q} \leftarrow \{h_1, \ldots, h_q\}$
$(i, \mathsf{aux}) \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}; \rho)$
return $\perp$ **if** $i = \perp$
$h'_i \xleftarrow{\$} H$
$\mathsf{Q}' \leftarrow \{h_1, \ldots, h_{i-1}\} \cup \{h'_i\} \cup \{h_{i+1} \ldots, h_q\}$
// Q *and* Q' *differ at exactly one point; at index* i
$(i', \mathsf{aux}') \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}'; \rho)$
return $\perp$ **if** $i = \perp$
if $i = i'$
  return $(i, \mathsf{aux}, \mathsf{aux}')$
return $\perp$

Let's now see how the local forking algorithm gives tighter bounds than the general forking lemma.

$$\mathsf{fork} \geq \frac{\mathsf{acc}^2}{q} \tag{5}$$

Lemma 4.1 *Given* acc *and* fork*, as defined, then* fork *is lower bounded by* acc *as in Equation 5.*

5 Extension of Generalized Forking Lemma

In recent analysis of the FROST signature scheme, Bellare, Tessaro, and Zhu introduced an extension of the generalized forking lemma as well as an extension to the local forking lemma. We review the generalized forking lemma extension here, and the extension of the local forking lemma in Section 6.

Let $S \subseteq \{1, \ldots, q\}$. Let $\mathcal{A}$ be a randomized algorithm that on input $(X, \{h_1, \ldots, h_q\})$, outputs an index $i \in \{\perp\} \cup S$, as well as auxiliary output aux.

The probability that the adversary will complete successfully acc and the probability that fork will output success is then as in Equation 6.

$$\mathsf{fork} \geq \frac{\mathsf{acc}^2}{|S|} \tag{6}$$

Intuitively, Equation 6 gives tighter bounds than the plain generalized forking lemma, as it bounds the acceptance probability to a subset $S \subseteq \{1, \ldots, q\}$. If it can be guaranteed that the adversary's output is strictly within S, then this variant allows for tighter bounds. Doing so is possible in the proof for FROST, as the adversary must query two random oracles in strict succession (as the output from the first random oracle is input into the second), which ensures that the environment can program the second random oracle at the time that the first is queried. Hence, the bounds can be scoped to the number of queries that the adversary has actually made to the first random oracle, as opposed to the range of allowed queries to the second random oracle.

6 Extension of Local Forking Lemma

We now review a "looser" local forking lemma again given by Bellare, Tessaro, and Zhu, this time in their analysis of the unforgeability for FROST.

The difference in this forking lemma is as follows. Here, forking occurs on two indices (i, j) in the first execution and (i', j') in the second execution. In this setting, i, i' are indices with respect to a random oracle. j is an auxiliary index that is guaranteed to be in the set J, and is used in the context of the wider proof. However, this forking lemma variant additionally enforces that the indices j, j are the same both before and after the fork, which results in a slightly looser bound than the original local forking lemma.

Why is this useful? This adaptation to the forking lemma is able to qualify over *more* information than simply some index i, as opposed to prior variants.

In other words, additional auxiliary information can be reflected in the analysis of the adversary's likelihood of success.

This "looser" local forking lemma is given in the following algorithm.

Algorithm $\mathsf{F}_{\mathsf{le}}(X)$

Pick coins ρ for $\mathcal{A}$ at random.
$\{h_1, \ldots, h_q\} \xleftarrow{\$} H^q$
$\mathsf{Q} \leftarrow \{h_1, \ldots, h_q\}$
$(i, j, \mathsf{aux}) \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}; \rho)$
  // *j is guaranteed to be strictly in some set J*
return $\perp$ **if** $i = \perp$
$h'_i \xleftarrow{\$} H$
$\mathsf{Q}' \leftarrow \{h_1, \ldots, h_{i-1}\} \cup \{h'_i\} \cup \{h_{i+1} \ldots, h_q\}$
$(i', j', \mathsf{aux}') \xleftarrow{\$} \mathcal{A}(X, \mathsf{Q}'; \rho)$
return $\perp$ **if** $i = \perp$
if $i = i'$ **and** $j = j'$
  return $(i, j, \mathsf{aux}, \mathsf{aux}')$
return $\perp$

The probability that fork will complete successfully is then reflected in Equation 7.

$$\mathsf{fork} \geq \frac{\mathsf{acc}^2}{q \cdot |J|} \tag{7}$$

7 Conclusion

In this note, we review the first variant of the forking lemma presented by Pointcheval and Stern. We then review its generalization beyond signature schemes, and several subsequent variants. Each variant is tailored to the specific proof of security which it is employed, or allows for tighter bounds or generalization beyond signature schemes.

References

[1] M. Bellare, W. Dai, and L. Li. The local forking lemma and its application to deterministic encryption. In S. D. Galbraith and S. Moriai, editors, *Advances in Cryptology - ASIACRYPT 2019 - 25th International Conference on the Theory and Application of Cryptology and Information Security, Kobe, Japan, December 8-12, 2019, Proceedings, Part III*, volume 11923 of *Lecture Notes in Computer Science*, pages 607–636. Springer, 2019.

[2] M. Bellare and G. Neven. Multi-signatures in the plain public-key model and a general forking lemma. In A. Juels, R. N. Wright, and S. D. C. di Vimercati, editors, *Proceedings of the 13th ACM Conference on Computer and Communications Security, CCS 2006, Alexandria, VA, USA, October 30 - November 3, 2006*, pages 390–399. ACM, 2006.

[3] D. Pointcheval and J. Stern. Security proofs for signature schemes. In U. M. Maurer, editor, *Advances in Cryptology - EUROCRYPT '96, International Conference on the Theory and Application of Cryptographic Techniques, Saragossa, Spain, May 12-16, 1996, Proceeding*, volume 1070 of *Lecture Notes in Computer Science*, pages 387–398. Springer, 1996.

[4] D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *J. Cryptol.*, 13(3):361–396, 2000.