An Abstract Multi-Forking Lemma

Charanjit S. Jutla

IBM T. J. Watson Research Center,
Yorktown Heights,
NY 10598, USA

Abstract. In this work we state and prove an abstract version of the multi-forking lemma of Pointcheval and Stern from EUROCRYPT'96. Earlier, Bellare and Neven had given an abstract version of forking lemma for two-collisions (CCS'06). While the original purpose of the forking lemma was to prove security of signature schemes in the random oracle methodology, the abstract forking lemma can be used to obtain security proofs for multi-signatures, group signatures, and compilation of interactive protocols under the Fiat-Shamir random-oracle methodology.

1 Forking Lemma

We start by introducing the forking lemma of Pointcheval and Stern [PS96]. Forking Lemma was introduced and proved in [PS96] to prove the unforgeability of signatures in the random-oracle model for signature schemes like DSA and EC-DSA. The main idea of the proof is to show that if an adversary can forge a signature, then the same adversary can be used to break the discrete-log problem. This is accomplished by rewinding the Adversary, so that the random-oracle can be programmed differently (yet, randomly). The adversary's responses in the different forks allow one to obtain the discrete log of the discrete log challenge. Of course, the Adversary's queries may adaptively change with each change in the random-oracle responses, and hence the forking lemma is a non-trivial probabilistic lemma.

A more abstract version of the forking lemma was stated and proved in [BN06]. However, the original lemma of [PS96] is more powerful in the sense that it can lower bound the probability of obtaining multi-collisions (in multi-forks). The [BN06] abstract version only handles a single fork and lower bounds the probability of a single collision. While the original purpose of the forking lemma was to prove security of signature schemes in the random-oracle methodology, the abstract forking lemma can be used to obtain security proofs for multi-signatures, group signatures, and compilation of interactive protocols under the Fiat-Shamir random-oracle methodolgy [FS87]. We remark that the forking lemma is not always useful, and alternate strategies are sometimes employed in compiling interactive proofs such as Valiant's extractor strategy [Val08] (see also [BCS16]). In [BDL19] a different *local* forking lemma was considered in which the random oracle is reprogrammed on just a single fork point rather than on all points past the fork.

In this work, borrowing some ideas from [BPVY00], we prove a stronger multi-collision version of the abstract forking lemma[1]. We first state the lemma as proved in [BN06], and then follow it with the more advanced lemma and its proof.

Notation. If A is a randomized algorithm, then $\mathsf{A}(x_1, x_2, ..., x_n; \rho)$ denotes the output of A on inputs $x_1, x_2, ..., x_n$ and coins ρ. We write $\sigma \xleftarrow{\$} \mathsf{A}(x_1, x_2, ...x_n)$ to indicate that σ was obtained as output of running A on inputs $x_1, x_2, ..., x_n$ with randomly chosen coins.

Lemma 1. *(Abstract Forking Lemma) Fix an integer $q \geq 1$ and a set H of size $\mathsf{h} \geq 2$. Let A be a randomized algorithm that on input $x, h_1, ...h_q$ returns a pair, the first element of which is an integer in the range $[0..q]$ and the second element of which we reer to as a side output. Let IG be a randomized algorithm with output in $\mathcal{X}$ and which we call the input generator. The accepting probability of A, denoted acc, is defined as the probability that $J \geq 1$ in the experiment*

$$x \xleftarrow{\$} \mathsf{IG};\ h_1, ...h_q \xleftarrow{\$} H;\ (J, \sigma) \xleftarrow{\$} \mathsf{A}(x, h_1, ..., h_q).$$

The forking algorithm $\mathsf{F_A}$ associated to A is the randomized algorithm that takes input x and proceeds as follows:

Algorithm $\mathsf{F_A}$
Pick coins ρ for A at random
$h_1, ..., h_q \xleftarrow{\$} H$
$(I, \sigma) \leftarrow \mathsf{A}(x, h_1, ...h_q; \rho)$
If $I == 0$ *then return* $(0, \epsilon, \epsilon)$
$h'_I, ...h'_q \xleftarrow{\$} H$
$(I', \sigma') \leftarrow \mathsf{A}(x, h_1, ..., h_{I-1}, h'_I, ..., h'_q; \rho)$
If $(I == I')$ *and* $(h_I \neq h'_I)$ *then return* $(1, \sigma, \sigma')$
Else return $(0, \epsilon, \epsilon)$.

Let

$$frk = \Pr[b = 1\ :\ x \xleftarrow{\$} \mathsf{IG}; (b, \sigma, \sigma') \xleftarrow{\$} \mathsf{F_A}].$$

Then,

$$frk\ \geq\ acc \cdot \left(\frac{acc}{q} - \frac{1}{\mathsf{h}}\right).$$

2 Abstract Multi-Forking Lemma

Lemma 2. *(Abstract Multi-Forking Lemma) Under the conditions and definitions of randomized algorithm A in the previous lemma, the accepting probability of A, denoted acc, is defined as the probability that $J \geq 1$ in the experiment*

$$x \xleftarrow{\$} \mathsf{IG}\, h_1, ...h_q \xleftarrow{\$} H;\ (J, \sigma) \xleftarrow{\$} \mathsf{A}(x, h_1, ..., h_q).$$

[1] The forking lemma in [PS96] is stated in terms of an adversary's interaction with a signature scheme.

The multi-forking algorithm F_{A} *associated to* A *is the randomized algorithm that takes a parameter* T *as input and runs in two phases as follows:*

Algorithm $\mathsf{F}_{\mathsf{A}}(T)$
Phase I:
Set $t = 0$*; Repeat*
$t = t + 1$*;*
Pick coins $\rho^{(t)}$ *for* A *at random;*
$x^{(t)} \xleftarrow{\$} \mathsf{IG}$*;* $h_1^{(t)}, ..., h_q^{(t)} \xleftarrow{\$} H$*;*
$(I^{(t)}, \sigma^{(t)}) \leftarrow \mathsf{A}(x^{(t)}, h_1^{(t)}, ...h_q^{(t)}; \rho^{(t)})$
Until $(I^{(t)} > 0)$ *or* $(t > T - 1)$*.*
Phase II:
Let $I = I^{(t)}$*. If* $I == 0$ *then return* $(0, I, \epsilon, \epsilon, \epsilon, \epsilon)$*.*
Else, choose $h'_I, ...h'_q \xleftarrow{\$} H$*.*
$(I', \sigma') \leftarrow \mathsf{A}(x^{(t)}, h_1^{(t)}, ..., h_{I-1}^{(t)}, h'_I, ..., h'_q; \rho^{(t)})$
If $(I == I')$ *then return* $(1, I, \sigma^{(t)}, \sigma', h_I^{(t)}, h'_I)$
Else return $(0, I, \epsilon, \epsilon, \epsilon, \epsilon)$*.*

1. *Let*
$$succ = \Pr[I \geq 1 : \ (b, I, \sigma, \sigma', h, h') \xleftarrow{\$} \mathsf{F}_{\mathsf{A}}(T)].$$
Then, $succ = 1 - (1 - acc)^T$*.*
2. *Define frk to be*
$$\Pr[(b = 1) \ \wedge \ (h \neq h') : \ (b, I, \sigma, \sigma', h, h') \xleftarrow{\$} \mathsf{F}_{\mathsf{A}}(T)].$$
Then,
$$frk \ \geq \ succ \cdot \left(\frac{acc}{q} - \frac{1}{|H|} \right).$$
3. *Fix a positive integer* ν*. Let* $p = (1/q) * acc - (\nu/|H|)$*. Suppose, after executing Phase I, the phase II part is repeated* $N = (\nu/p) * \log 2\nu$ *times, with outputs designated* $(b_j, I, \sigma, \sigma_j, h, h_j)$ *for* $j \in [1..N]$*. Let* $\mathcal{N}$ *be a subset of* $[1..N]$ *such that*
(a) for all $j \in \mathcal{N}$*:* $b_j = 1$ *and* $h_j \neq h$*,*
(b) and for all $j, j' \in \mathcal{N}$*,* $j \neq j'$*:* $h_j \neq h_{j'}$*.*
Then, probability that there exists an $\mathcal{N}$ *such that* $|\mathcal{N}| \geq \nu$ *is at least* $succ * e^{-1}$*.*

We state two important basic lemmas from probability theory.

Lemma 3. *Let* X *be a real-valued random variable. Then* $\mathbf{E}[X^2] \geq \mathbf{E}[X]^2$*.*

Lemma 4. *Suppose* $q \geq 1$ *is an integer, and* $x_1, ...x_q \geq 0$ *are real numbers. Then*
$$\sum_{i=1}^{q} x_i^2 \ \geq \ \frac{1}{q} \left(\sum_{i=1}^{q} x_i \right)^2.$$

Proof. (of Lemma 2)

Proof of (1): $I \geq 1$ iff at the end of phase I, $I^{(t)} \geq 1$. Thus, $I == 0$ iff Phase I ran through all T repetitions and all of them produced $I^{(t)} = 0$. The claim follows as all repetitions in Phase I are completely independent.

Proof of (2) (similar to [BN06]): Let h_I stand for $h_I^{(t)}$ at the end of Phase I. Note that the predicate in the definition of frk is equivalent to $(I' = I) \wedge (I \geq 1) \wedge (h'_I \neq h_I)$. Now,

$$\Pr[I' = I \wedge I \geq 1 \wedge h'_I \neq h_I] = \Pr[I' = I \wedge h'_I \neq h_I \mid I \geq 1] * \text{succ}$$

Also, $\Pr[I' = I \wedge h'_I \neq h_I \mid I \geq 1]$ is same as $\Pr[I' = I^{(t)} \wedge h'_I \neq h_I \mid I^{(t)} \geq 1]$, where t is the value of the variable at the end of Phase I. Now, since all repetitions of phase I are completely independent, and in particular use independent randomness, this probability is same if algorithm $\mathsf{F}_{\mathsf{A}}(T)$ was run with parameter T set to one. Thus, with $T = 1$, the above probability is same as $\Pr[I' = I^{(1)} \wedge h'_I \neq h_I \mid I^{(1)} \geq 1]$. Again, we will just refer to $I^{(1)}$ as I. This probability is then

$$\begin{aligned}
&\Pr[I' = I \wedge h'_I \neq h_I \wedge I \geq 1] * \frac{1}{\text{acc}} \\
&\geq (\Pr[I' = I \wedge I \geq 1] - \Pr[h'_I = h_I \wedge I \geq 1]) * \frac{1}{\text{acc}} \\
&\geq (\Pr[I' = I \wedge I \geq 1] - \Pr[h'_I = h_I \mid I \geq 1]) * \frac{1}{\text{acc}} \\
&\geq (\Pr[I' = I \wedge I \geq 1] - \frac{1}{|H|}) * \frac{1}{\text{acc}}.
\end{aligned}$$

We now focus on lower-bounding $\Pr[I' = I \wedge I \geq 1]$. For each $i \in [1..q]$, Define $X_i : \mathcal{X} \times \mathcal{R} \times H^{i-1} \to [0, 1]$ to be

$$\begin{aligned}
&X_i(\hat{x}, \hat{\rho}, \hat{h}_1, ..., \hat{h}_{i-1}) = \\
&\quad \Pr[\hat{I} = i; \hat{h}_i, ..., \hat{h}_q \xleftarrow{\$} H, (\hat{I}, \hat{\sigma}) \leftarrow A(\hat{x}, \hat{h}_1, ...\hat{h}_q, \hat{\rho})].
\end{aligned}$$

For each $i \in [1..q]$, we show that $\Pr[I = i \wedge I' = i] = E[X_i^2]$.

Now, for $i > 0$,

$$\begin{aligned}
\Pr[I = i \ \wedge \ I' = i] = & \sum_{x^*,\rho^*,h_1^*,...,h_{i-1}^*} \Pr[\rho = \rho^* \ \wedge \ x = x^* \wedge \\
& \vec{h}_{|i-1} = \vec{h}^*_{|i-1} \ \wedge \ (I = i) \wedge (I' = i); \ x \xleftarrow{\$} \mathcal{X}, \ \vec{h}, h'_i, ..., h'_q \xleftarrow{\$} H, \\
& (I, \sigma) \leftarrow \mathsf{A}(x^*, h_1^*, ..., h_{i-1}^*, h_i, ..., h_q; \rho^*), \\
& (I', \sigma') \leftarrow \mathsf{A}(x^*, h_1^*, ..., h_{i-1}^*, h'_1, ..., h'_q; \rho^*)] \\
= & \sum \ \Pr[\rho = \rho^* \ \wedge \ x = x^* \ \wedge \vec{h}_{|i-1} = \vec{h}^*_{|i-1}] * \\
& \Pr[(I = i) \wedge (I' = i); h_i, ..., h_q, h'_i, ..., h'_q \xleftarrow{\$} H, \\
& (I, \sigma) \leftarrow \mathsf{A}(x^*, h_1^*, ..., h_{i-1}^*, h_i, ..., h_q; \rho^*), \\
& (I', \sigma') \leftarrow \mathsf{A}(x^*, h_1^*, ..., h_{i-1}^*, h'_1, ..., h'_q; \rho^*)] \\
= & \sum \Pr[\rho = \rho^* \wedge x = x^* \wedge \vec{h}_{|i-1} = \vec{h}^*_{|i-1}] \ * \ X_i(x^*, \rho^*, \vec{h}^*_{|i-1})^2 \\
= & E[X_i^2]
\end{aligned}$$

Thus, using basic probability theory [BN06],

$$\begin{aligned}
& \sum_{i=1}^{q} \Pr[I = i \ \wedge \ I' = i] = \\
& \qquad \sum_{i=1}^{q} E[X_i^2] \geq \sum_{i=1}^{q} E[X_i]^2 \geq \frac{1}{q}(\sum_{i=1}^{q} E[X_i])^2 \geq \frac{1}{q}\mathrm{acc}^2.
\end{aligned}$$

Proof of (3): Let the experiment in the statement of the lemma (part(3)), , i.e. repetition of Phase II, be called $\mathbf{Expt}_0$. We consider an alternate experiment $\mathbf{Expt}_1$ in which before the start of Phase II a variable D, called *bad set*, is initialized to singleton set $\{h_I\}$ if $I \geq 1$. If $I = 0$, phase II terminates as before. The set D maybe updated at the end of each repetition of Phase II, to be described next. During the sampling of $h'_I, ..., h'_q$ in a repetition of Phase II, if h'_I is in set D, then it outputs $(0, I, \epsilon, \epsilon, \epsilon, \epsilon)$ instead now. At the end of j-th repetition of Phase II, if the repetition was a success, i.e. $b_j = 1$, then the value h'_I chosen in this repetition is added to the set D.
In $\mathbf{Expt}_1$, after N repetitions, denote by $p_{\nu,N}$ the probability of existence of a subset $\mathcal{N}$ of size at least ν satisfying: (a') for all $j \in \mathcal{N}$ $b_j = 1$. We now claim that the probability of existence of a subset $\mathcal{N}$ of size at least ν satisfying (a) and (b) in the original experiment $\mathbf{Expt}_0$ is at least $p_{\nu,N}$. This follows easily because the underlying probability distribution is same in both experiments, and for every choice of h' variables and subset $\mathcal{N}$ satisfying (a') in $\mathbf{Expt}_1$ there is the same choice of h' variables and subset $\mathcal{N}$ satisfying (a) and (b) in $\mathbf{Expt}_0$.
We note that the probability of frk in each repetition is slightly different now. In particular, the probability of frk in the j-th repetition of Phase II ($j \in [1..N]$) is now lower bounded by $\mathrm{succ} * (\mathrm{acc}/q - |D_j|/|H|)$, where D_j

is the set D at the start of the j-th repetition. This is most conveniently shown by considering for each $i \in [1..q]$ and $D \in 2^H$, s function $Y_{i,D} : \mathcal{X} \times \mathcal{R} \times H^{i-1} \to [0,1]$ to be

$$Y_i(\hat{x}, \hat{\rho}, \hat{h}_1, ..., \hat{h}_{i-1}) =$$
$$Pr[\hat{I} = i \wedge \hat{h}_i \notin D; \hat{h}_i, ..., \hat{h}_q \xleftarrow{\$} H, (\hat{I}, \hat{\sigma}) \leftarrow A(\hat{x}, \hat{h}_1, ...\hat{h}_q, \hat{\rho})].$$

For any D of size d, note that $Y_{i,D}$ on its arguments is at least X_i on the same arguments minus $d/|H|$. Next, For each $i \in [1..q]$, we show that $\Pr[I = i \wedge I'_D = i] = E[X_i * Y_{i,D}]$, where I'_D is I' in Phase II in $\mathbf{Expt}_1$ with the bad set initialized to D. Now, for $i > 0$,

$$\begin{aligned}
&\Pr[I = i \ \wedge \ I'_D = i] = \sum_{x^*, \rho^*, h^*_1, ..., h^*_{i-1}} \Pr[\rho = \rho^* \ \wedge \ x = x^* \wedge \\
&\vec{h}_{|i-1} = \vec{h}^*_{|i-1} \ \wedge \ h'_i \notin D \ \wedge (I = i) \wedge (I' = i); \ x \xleftarrow{\$} \mathcal{X}, \ \vec{h}, h'_i, ..., h'_q \xleftarrow{\$} H, \\
&(I, \sigma) \leftarrow \mathsf{A}(x^*, h^*_1, ..., h^*_{i-1}, h_i, ..., h_q; \rho^*), \\
&(I', \sigma') \leftarrow \mathsf{A}(x^*, h^*_1, ..., h^*_{i-1}, h'_1, ..., h'_q; \rho^*)] \\
=& \sum \Pr[\rho = \rho^* \ \wedge \ x = x^* \ \wedge \vec{h}_{|i-1} = \vec{h}^*_{|i-1}] * \\
&\Pr[h'_i \notin D \ \wedge (I = i) \wedge (I' = i); h_i, ..., h_q, h'_i, ..., h'_q \xleftarrow{\$} H, \\
&\quad (I, \sigma) \leftarrow \mathsf{A}(x^*, h^*_1, ..., h^*_{i-1}, h_i, ..., h_q; \rho^*), \\
&\quad (I', \sigma') \leftarrow \mathsf{A}(x^*, h^*_1, ..., h^*_{i-1}, h'_1, ..., h'_q; \rho^*)] \\
&= \sum \Pr[\rho = \rho^* \wedge x = x^* \wedge \vec{h}_{|i-1} = \vec{h}^*_{|i-1}] \ * \ X_i(x^*, \rho^*, \vec{h}^*_{|i-1}) * \\
&\quad Y_{i,D}(x^*, \rho^*, \vec{h}^*_{|i-1}) \\
&= E[X_i * Y_{i,D}]
\end{aligned}$$

Thus, again using basic probability theory, and denoting I' in the j-th repetition of Phase II in $\mathbf{Expt}_1$ by I'_j,

$$\begin{aligned}
&\sum_{i=1}^{q} \Pr[I = i \ \wedge \ I'_j = i] \\
&\geq \sum_{i=1}^{q} E[X_i * (X_i - |D_j|/|H|)] \\
&\geq \sum_{i=1}^{q} (E[X_i]^2 - (|D_j|/|H|) * E[X_i]) \\
&\geq (1/q) * \mathrm{acc}^2 - (|D_j|/|H|) * \mathrm{acc}.
\end{aligned}$$

Similarly to [BPVY00], we now analyze the probability of existence of set $\mathcal{N}$ of size at least ν with (a') holding in $\mathbf{Expt}_1$. Recall, $p = (1/q) * \mathrm{acc} - (\nu/|H|)$. We first focus on the first N/ν repetitions till a repetition j, $j \leq N/\nu$ reports

success, i.e. $b_j = 1$. Conditioned on $I \geq 1$, the probability that none of the first N/ν repetitions reports success is at most $(1-p)^{N/\nu}$, since all these repetitions are independent (in particular bad set D is not updated). Thus, the probability of success in the first N/ν repetitions is at least $1-(1-p)^{N/\nu} \geq 1-(1-p)^{\log 2\nu*(1/p)} \geq 1-(2\nu)^{-1}$. We next focus on the next N/ν repetitions, starting from the last success (or N/ν if none found). Again, the probability of success in this bunch is at least $1-(2\nu)^{-1}$. Thus after ν such bunched-repetitions, probability that all the ν bunches had success (and which are independent events) is at least[2] $(1-(2\nu)^{-1})^{\nu} \geq e^{-\nu/(2\nu-1)} \geq e^{-1}$.

References

[BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. Cryptology ePrint Archive, Report 2016/116, 2016. https://eprint.iacr.org/2016/116.

[BDL19] Mihir Bellare, Wei Dai, and Lucy Li. The local forking lemma and its application to deterministic encryption. In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part III*, volume 11923 of *LNCS*, pages 607–636. Springer, Heidelberg, December 2019.

[BN06] Mihir Bellare and Gregory Neven. Multi-signatures in the plain public-key model and a general forking lemma. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006*, pages 390–399. ACM Press, October / November 2006.

[BPVY00] Ernest F. Brickell, David Pointcheval, Serge Vaudenay, and Moti Yung. Design validations for discrete logarithm based signature schemes. In Hideki Imai and Yuliang Zheng, editors, *PKC 2000*, volume 1751 of *LNCS*, pages 276–292. Springer, Heidelberg, January 2000.

[FS87] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In Andrew M. Odlyzko, editor, *CRYPTO'86*, volume 263 of *LNCS*, pages 186–194. Springer, Heidelberg, August 1987.

[PS96] David Pointcheval and Jacques Stern. Security proofs for signature schemes. In Ueli M. Maurer, editor, *EUROCRYPT'96*, volume 1070 of *LNCS*, pages 387–398. Springer, Heidelberg, May 1996.

[Val08] Paul Valiant. Incrementally verifiable computation or proofs of knowledge imply time/space efficiency. In Ran Canetti, editor, *TCC 2008*, volume 4948 of *LNCS*, pages 1–18. Springer, Heidelberg, March 2008.

[2] Using the inequality $(1-p)^{1/p} \geq e^{-1/(1-p)}$ for $0 < p < 1$, which in turn follows from $e^x \geq 1+x$.