Monero's Building Blocks Part 2 of 10 – *Pointcheval & Stern's Generic Signature Scheme [1]*

Bassam El Khoury Seguias

BTC: 3FcVvBZwTUkUrcqJd16RcjR42qT2tDWHWn

ETH: 0xb79Fb9194C8Cc6221368bb70976e18609Ab9AcA8

March 3, 2018

1 Introduction

In the next few parts of this series, we look at various signature schemes and prove their security in the RO model. This part is dedicated to the analysis of a generic signature scheme introduced in [1], of which the non-interactive *Schnorr* scheme is an example. The generic scheme is built around a single (sk, pk) pair. Later parts of this series will focus on *ring signature schemes*. Ring signatures embed the actual signer in a ring of other possible signers to hide her identity. We will discuss them in parts 3, 4, 5, 6, and 7.

2 Pointcheval & Stern's generic scheme

For a given message m, our generic scheme creates a signature $\sigma(m) \equiv (r, h, \alpha)$ where r is a random element chosen from a pre-defined set, $h = \mathcal{H}(m, r)$ (i.e., RO output on query (m, r)), and α is fully determined by m, r, and h. By design, we require that the probability of selecting any particular r be upper-bounded by $\frac{1}{2^{k-1}}$ for a given security parameter k.

Schnorr's signature scheme is an example that fits this generic model. To see why, recall that $\Sigma_{Schnorr}$ chooses a random commitment $k \in \mathbb{Z}_q^*$ where q is a pre-defined prime number. It then assigns $r \equiv g^k$ where g denotes a chosen generator of $\mathbb{Z}_q^*$. Afterwards, h is set to $\mathcal{H}(m, r)$. Finally, α is calculated as $k - hx$ where x denotes the signer's private key. Note that r can be any element of $\mathbb{Z}_q^*$ and so the probability that it takes on a specific value is equal to $\frac{1}{q-1}$. By design, we choose the security parameter $k \leq log_2(q)$. This choice of k guarantees that the aforementioned probability is upper-bounded by $\frac{1}{2^k-1} \leq \frac{1}{2^{k-1}}$.

3 Security analysis - Unforgeability vis-a-vis EFACM

For unforgeability proofs, we follow the 5-step approach mentioned in part 1 of this series.

Step 1 : To prove that this generic scheme is secure against EFACM in the RO model, we proceed by contradiction and assume that there exists a PPT adversary $\mathcal{A}$ such that:

$$P_{\omega,r,\mathcal{H}}[\mathcal{A}(\omega)^{\mathcal{H},\Sigma^{\mathcal{H}}(r)} \textit{ succeeds in } EFACM] = \epsilon(k), \text{ for some } \epsilon \text{ non-negligible in k.}$$

Step 2 : Next, we build a simulator $\mathcal{S}(r')$ such that it:

- Does not have access to the private key of any signer.
- Has the same range as Σ (i.e., they output signatures taken from the same pool of potential signatures over all possible choices of RO functions and respective random tapes r and r').
- Has indistinguishable probability distribution from that of Σ over this range .

$\mathcal{S}(r')$ is specific to the particular instance of the generic scheme being used. In what follows, we build a simulator for the case of a *Schnorr* signature.

Original Signer $\Sigma(r)$

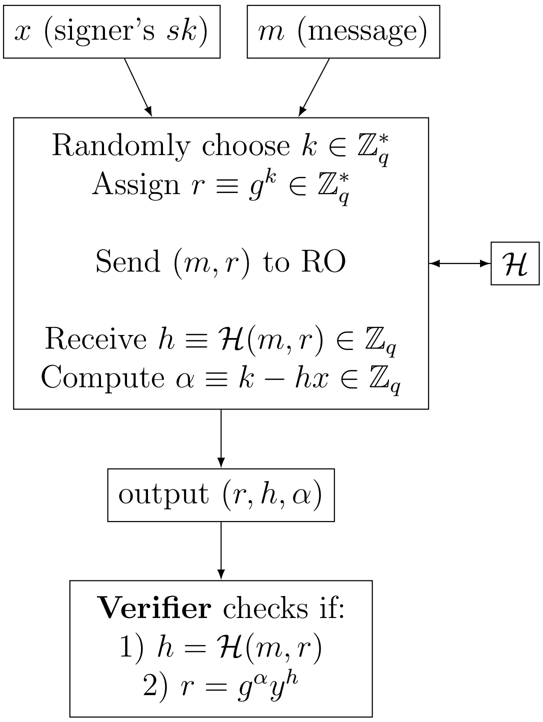
x (signer's sk)
m (message)
Randomly choose $k \in \mathbb{Z}_q^*$
Assign $r \equiv g^k \in \mathbb{Z}_q^*$
Send (m, r) to RO
$\mathcal{H}$
Receive $h \equiv \mathcal{H}(m, r) \in \mathbb{Z}_q$
Compute $\alpha \equiv k - hx \in \mathbb{Z}_q$
output (r, h, α)
Verifier checks if:
1) $h = \mathcal{H}(m, r)$
2) $r = g^{\alpha} y^{h}$

Simulator $\mathcal{S}(r')$ *bypasses RO*

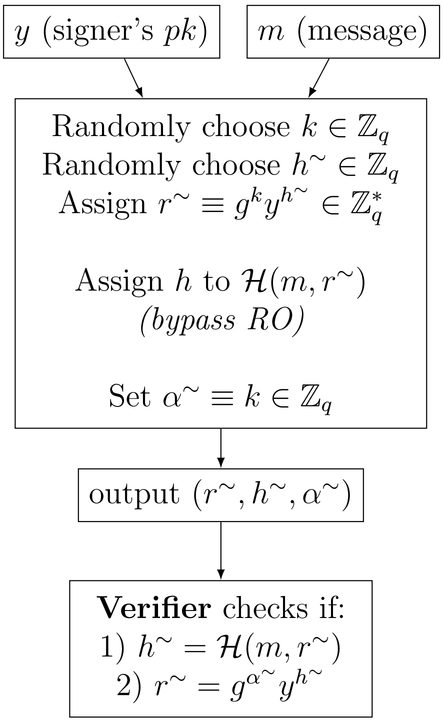
y (signer's pk)
m (message)
Randomly choose $k \in \mathbb{Z}_q$
Randomly choose $h^{\sim} \in \mathbb{Z}_q$
Assign $r^{\sim} \equiv g^k y^{h^{\sim}} \in \mathbb{Z}_q^*$
Assign h to $\mathcal{H}(m, r^{\sim})$
(bypass RO)
Set $\alpha^{\sim} \equiv k \in \mathbb{Z}_q$
output $(r^{\sim}, h^{\sim}, \alpha^{\sim})$
Verifier checks if:
1) $h^{\sim} = \mathcal{H}(m, r^{\sim})$
2) $r^{\sim} = g^{\alpha^{\sim}} y^{h^{\sim}}$

By construction, the output of $\mathcal{S}$ will satisfy the verification equations. Moreover, it assigns a random value for h and bypasses the RO in doing so. Next, note the following:

1. $\mathcal{S}$ does not use any private key.
2. Σ and $\mathcal{S}$ both have a range $R \equiv \{(\epsilon, \beta, \gamma) \in (\mathbb{Z}_q^* \times \mathbb{Z}_q \times \mathbb{Z}_q) \ s.t. \ \epsilon = g^{\gamma} \times y^{\beta}\}$.

3. Σ and $\mathcal{S}$ have the same probability distribution over R. Indeed, $\forall(\epsilon, \beta, \gamma) \in R$ we have:
   - For $\Sigma : P[(r, h, \alpha) = (\epsilon, \beta, \gamma)] = P_{k\neq 0,h}[g^k = \epsilon, h = \beta, k - hx = \gamma] = \frac{1}{(q-1).q}$.
   - For $\mathcal{S} : P[(r^{\sim}, h^{\sim}, \alpha^{\sim}) = (\epsilon, \beta, \gamma)] = P_{r^{\sim},h^{\sim}}[r^{\sim} \equiv g^k y^{h^{\sim}} = \epsilon, h^{\sim} = \beta, \alpha^{\sim} \equiv k = \gamma] = \frac{1}{(q-1).q}$.

With $\mathcal{S}$ adequately built for the Schnorr scheme, we conclude that (*refer to section 6 of part 1 of this series for a justification*):

$$P_{\omega,r,'\mathcal{H}}[\mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} \textit{ succeeds in EFACM}] = \epsilon(k), \text{ for some } \epsilon \text{ non-negligible in k.}$$

Step 3 : We now show that the probability of faulty collisions is negligible (*refer to section 6 of part 1 of this series for a description of collision types*). The 2 tyes of collisions fo the generic scheme are:

- $Col_{Type\ 1}$: A tuple (m, r) that $\mathcal{S}$ encounters – it makes its own random assignment to $\mathcal{H}(m, r)$ and bypasses RO – also appears in the list of queries that $\mathcal{A}(\omega)$ sends to RO. A conflict in the 2 values will happen with overwhelming probability and the execution will halt.
- $Col_{Type\ 2}$: A tuple (m, r) that $\mathcal{S}$ encounters – it makes its own random assignment to $\mathcal{H}(m, r)$ – is the same as another tuple (m', r') that $\mathcal{S}$ encountered at an earlier time instance – here too, $\mathcal{S}$ would have made its own random assignment to $\mathcal{H}(m', r')$. Since the 2 tuples are identical (i.e., $(m, r) = (m', r')$), it must be that the 2 random assignments match (i.e., $\mathcal{H}(m, r) = \mathcal{H}(m', r')$). With overwhelming probability, the 2 values will be different and the execution will halt.

The aforementioned collisions must be avoided. In order to do so, we first calculate the probability of their occurence. We assume that during an EFACM attack, $\mathcal{A}(\omega)$ can make a maximum of Q queries to RO and a maximum of Q_S queries to $\mathcal{S}(r')$. Q and Q_S are both assumed to be polynomial in the security parameter k, since the adversary is modeled as a PPT Turing machine.

$$P[Col_{Type\ 1}] = P[\cup_{all\ (m,r)}\{(m,r) \textit{ appeared in at least one of the } Q_S \textit{ queries to } \mathcal{S} \textit{ and } Q \textit{ queries to RO}\}]$$

$$\leq P[\cup_{all\ r}\{r \textit{ was part of at least one of the } Q_S \textit{ queries to } \mathcal{S} \textit{ and } Q \textit{ queries to RO}\}]$$

$$\leq \textstyle\sum_{all\ r\in\mathbb{Z}_q^*} P[\cup_{(j=1,..,Q_S),\ (k=1,...,Q)}\{r \textit{ was part of at least the } j^{th} \textit{ query to } \mathcal{S} \textit{ and } k^{th} \textit{ queries to RO}\}]$$

$$\leq \textstyle\sum_{all\ r\in\mathbb{Z}_q^*} \sum_{j=1}^{Q_S} \sum_{k=1}^{Q} P[r \textit{ was part of at least the } j^{th} \textit{ query to } \mathcal{S} \textit{ and } k^{th} \textit{ queries to RO}]$$

$$\leq \textstyle\sum_{all\ r\in\mathbb{Z}_q^*} \sum_{j=1}^{Q_S} \sum_{k=1}^{Q} (\frac{1}{q-1})^2 = (q-1) \times \frac{Q_S Q}{(q-1)^2} = \frac{Q_S Q}{(q-1)} \leq \frac{Q_S Q}{2^{k-1}}$$

Since Q_S and Q are polynomial in k, we conclude that $P[Col_{Type\ 1}]$ is negligible in k.

Next, we compute $P[Col_{Type\ 2}]$:

$$P[Col_{Type\ 2}] = P[\cup_{all\ \ (m,r)}\{(m,r)\ appeared\ at\ least\ twice\ in\ some\ of\ the\ queries\ to\ \mathcal{S}\}]$$

$$\leq P[\cup_{r\in\mathbb{Z}_q^*}\{r\ was\ part\ of\ at\ least\ 2\ queries\ to\ \mathcal{S}\}]$$

$$\leq \sum_{r\in\mathbb{Z}_q^*}\binom{Q_S}{2} \times \frac{1}{(q-1)^2} = \binom{Q_S}{2} \times \frac{q-1}{(q-1)^2} = \frac{Q_S(Q_S-1)}{2(q-1)} \leq \frac{Q_S^2}{2\times 2^{k-1}}$$

And so $P[Col_{Type\ 2}]$ is also negligible in k.

We then have:

$$P[Col] = P[Col_{Type\ 1} \cup Col_{Type\ 1}] \leq P[Col_{Type\ 1}] + P[Col_{Type\ 2}] \leq \frac{Q_S Q + \frac{Q_S^2}{2}}{2^{k-1}} \equiv \delta(k)$$

which is negligible in k. We can finally conclude (as was shown in section 6 of part 1), that:

$$P_{\omega,r,'\mathcal{H}}[\mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} succeeds\ in\ EFACM\ \ \cap \overline{Col}] \geq \epsilon(k) - \delta(k),\ \text{(non-negligible in } k)$$

Step 4 : In this step, our objective is to show that if $(\omega^*, r'^*, \mathcal{H}^*)$ is a successful tuple that generated a first EFACM forgery, then the following quantity is non-negligible in k:

$$P_{\mathcal{H}}[\mathcal{A}(\omega^*)^{\mathcal{H},\mathcal{S}(r'^*)}\ succeeds\ in\ EFACM\ \ \cap\ \ (\rho_\beta \neq \rho_\beta^*) \mid (\omega^*, r'^*, \mathcal{H}^*)\ is\ a\ succesfull\ first\ forgery,\ and\ (\rho_i = \rho_i^*)\ for\ i \in \{1, ...\beta - 1\}]$$

Here β is an appropriate index that we will define in the proof. And to further simplify the notation, we let $\rho_i^* \equiv \mathcal{H}^*(q_i^*)$ and $\rho_i \equiv \mathcal{H}(q_i)$. ($q_i^*$ and q_i denote respectively the i^{th} query to $\mathcal{H}^*$ and $\mathcal{H}$) for all $i \in 1, ...\beta$.

Let's take a closer look at $P_{\omega,r,'\mathcal{H}}[\mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} succeeds\ in\ EFACM \cap \overline{Col}]$.

Any successful forgery must pass the verification test. One of the verification equations is to check if $h = \mathcal{H}(m, r)$. So we distinguish between 2 scenarios (*w.l.o.g. we assume that all $\mathcal{A}$-queries sent to RO are distinct from each-other since $\mathcal{A}$ can keep a local copy of previous query results and avoid redundant calls*):

- Scenario 1: $\mathcal{A}$ was successful in its forgery, and no collisions occured, and it never queried RO on input (m, r).
- Scenario 2: $\mathcal{A}$ was successful in its forgery, and no collisions occured, and it queried RO on input (m, r) during its execution.

The probability of scenario 1 is upperbounded by the probability that $\mathcal{A}$ picks a value for h that matches the value of $\mathcal{H}(m, r)$. Here, $\mathcal{H}(m, r)$ is the value that RO returns to $\mathcal{V}$ (the verification algorithm) when verifying the validity of the forged signature. (It is upper-bounded because at the very least, the constraint $h = \mathcal{H}(m, r)$ must be observed for a valid signature). And since h can be any value in $\mathbb{Z}_q$, we get:

$$P[Scenario\ 1] \leq \frac{1}{q} \leq \frac{1}{2^k}, \text{ which is negligible in } k.$$

So we assume that a successful forgery will likely be of the Scenario 2 type. We have:

$$P[Scenario\ 2] = P_{\omega,r,'\mathcal{H}}[\mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} succeeds\ in\ EFACM\ \cap \overline{Col}] - P[Scenario\ 1]$$

$$\geq \epsilon(k) - \delta(k) - \frac{1}{2^k} \equiv \nu(k), \text{ which is non-negligible in } k$$

We then define $Ind(\omega, r', \mathcal{H})$ to be the index of the query (m, r) sent by $\mathcal{A}(\omega)$ to RO during execution. We let $Ind(\omega, r', \mathcal{H}) = \infty$ if the query (m, r) was never asked by $\mathcal{A}(\omega)$. This definition allows us to build the following sets:

- $S = \{(\omega, r', \mathcal{H}) \mid \mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} succeeds\ in\ EFACM\ \cap \overline{Col} \cap\ Ind(\omega, r', \mathcal{H}) \neq \infty\}$

  In other terms, S is the set of tuples $(\omega, r', \mathcal{H})$ that yield a successful EFACM forgery when no collisions occur, and when $\mathcal{A}(\omega)$ queried RO on input (m, r) at some point during its execution (i.e., scenario 2).

- $S_i = \{(\omega, r', \mathcal{H}) \mid \mathcal{A}(\omega)^{\mathcal{H},\mathcal{S}(r')} succeeds\ in\ EFACM\ \cap \overline{Col} \cap\ Ind(\omega, r', \mathcal{H}) = i\}$

  In other terms, S_i is the set of tuples $(\omega, r', \mathcal{H})$ that yield a successful EFACM forgery when no collisions occur, and when the index of the $\mathcal{A}(\omega)$-query on input (m, r) sent to RO is equal to i.

Recall that, $P_{\omega,r',\mathcal{H}}[(\omega, r', \mathcal{H}) \in S] = P[Scenario2] \geq \nu(k)$, which is non-negligible in k.

And clearly, $\{\cup_{i=1}^{Q} S_i\}$ partitions S. So $\sum_{i=1}^{Q} P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S] = 1$.

This implies that $\exists i \in \{1, ...Q\}\ s.t.\ P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S] \geq \frac{1}{2Q}$.

If this were not the case, then one would get the following contradiction:

$$1 = \sum_{i=1}^{Q} P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S] < Q \times \frac{1}{2Q} = \frac{1}{2} < 1.$$

So we introduce the set I consisting of all indices that meet the $\frac{1}{2Q}$ threshold, i.e.

$$I = \{i \in \{1, ...Q\} \mid P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S] \geq \frac{1}{2Q}\}$$

We claim that $P[Ind(\omega, r', \mathcal{H}) \in I \mid (\omega, r', \mathcal{H}) \in S] \geq \frac{1}{2}$.

Proof By definition of the sets S_i we have:

$$P[Ind(\omega, r', \mathcal{H}) \in I \mid (\omega, r', \mathcal{H}) \in S] = \sum_{i \in I} P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S]$$

$$= 1 - \sum_{j \notin I} P[(\omega, r', \mathcal{H}) \in S_j \mid (\omega, r', \mathcal{H}) \in S] > 1 - \sum_{j \notin I} \frac{1}{2Q} > 1 - \frac{Q}{2Q} = \frac{1}{2}$$

The next step is to apply the splitting lemma to each $S_i,\ i \in I$. First note that:

$$P_{\omega,r',\mathcal{H}}[(\omega, r', \mathcal{H}) \in S_i] = P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S] \times P_{\omega,r',\mathcal{H}}[(\omega, r', \mathcal{H}) \in S]$$

$$\geq \tfrac{1}{2Q} \times \nu(k)$$

Referring to the notation used in the splitting lemma (section 7 of part 1), we let:

$$A \equiv S_i,\ X \equiv (\omega, r', \rho_1, ..., \rho_{i-1}),\ Y \equiv (\rho_i, ..., \rho_Q),\ \epsilon \equiv \tfrac{\nu(k)}{2Q}, \text{ and } \alpha \equiv \tfrac{\nu(k)}{4Q} = \tfrac{\epsilon}{2}$$

X is defined as the space of tuples of all random tapes ω, all random tapes r', and all possibe RO answers to the first $i-1$ queries sent by $\mathcal{A}(\omega)$. Y is defined as the space of all possible RO answers to the last $(Q-i+1)$ queries sent by $\mathcal{A}(\omega)$. (Recall that $\rho_i \equiv \mathcal{H}(q_i)$). The splitting lemma guarantees the existence of a subset Ω_i of tuples $(\omega, r', \mathcal{H})$ such that:

- $P_{\omega,r',\mathcal{H}}[(\omega, r', \mathcal{H}) \in \Omega_i] \geq \frac{\nu(k)}{4Q}$
- $\forall[(\omega^\sim, r'^\sim, \mathcal{H}^\sim) \equiv (\omega^\sim, r'^\sim, \rho_1^\sim, ..., \rho_{i-1}^\sim, \rho_i^\sim ... \rho_Q^\sim]) \in \Omega_i$, we have

  $$P_{\mathcal{H}}[(\omega^\sim, r'^\sim, \rho_1^\sim, ..., \rho_{i-1}^\sim, \rho_i ... \rho_Q) \in S_i \mid (\omega^\sim, r'^\sim, \mathcal{H}^\sim) \in \Omega_i] \geq \tfrac{\nu(k)}{4Q}, \text{ and so}$$

  $$P_{\mathcal{H}}[(\omega^\sim, r'^\sim, \mathcal{H}) \in S_i \mid (\omega^\sim, r'^\sim, \mathcal{H}^\sim) \in \Omega_i,\ \rho_1 = \rho_1^\sim, ..., \rho_{i-1} = \rho_{i-1}^\sim)] \geq \tfrac{\nu(k)}{4Q}$$

- $P[(\omega, r', \mathcal{H}) \in \Omega_i \mid (\omega, r', \mathcal{H}) \in S_i] \geq (\frac{\nu(k)}{4Q})/(\frac{\nu(k)}{2Q}) = \frac{1}{2}$

We would like to compute the probability of finding a 2^{nd} successful tuple $(\omega^*, r'^*, \mathcal{H}^\sim)$ given that $(\omega^*, r'^*, \mathcal{H}^*)$ was a successful 1^{st} tuple and s.t. $\rho_j^\sim = \rho_j^*,\ j = 1, ...i-1$. That means finding the following probability:

$$P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_i \mid (\omega^*, r'^*, \mathcal{H}^*) \in S_i,\ \rho_1 = \rho_1^*, ..., \rho_{i-1} = \rho_{i-1}^*].$$

From the splitting lemma results, we have a (non-negligible in k) lower-bound on $P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_i \mid (\omega^*, r'^*, \mathcal{H}^*) \in \Omega_i,\ \rho_1 = \rho_1^*, ..., \rho_{i-1} = \rho_{i-1}^*]$.

Note however, that Ω_i and S_i are generally distinct sets. And so we **cannot** conclude that

$$P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_i \mid (\omega^*, r'^*, \mathcal{H}^*) \in S_i,\ \rho_1 = \rho_1^*, ..., \rho_{i-1} = \rho_{i-1}^*]$$

$$= P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_i \mid (\omega^*, r'^*, \mathcal{H}^*) \in \Omega_i,\ \rho_1 = \rho_1^*, ..., \rho_{i-1} = \rho_{i-1}^*]$$

and therefore we **cannot** conclude that the following quantity is non-negligible in k

$$P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_i \mid (\omega^*, r'^*, \mathcal{H}^*) \in S_i,\ \rho_1 = \rho_1^*, ..., \rho_{i-1} = \rho_{i-1}^*]$$

In order to show that the above quantity is non-negligible in k, we proceed differently. Suppose we can show that the following probability is non-negligible in k:

$$P_{(\omega,r',\mathcal{H})}[\exists \beta \in I\ s.t.\ (\omega, r', \mathcal{H}) \in (\Omega_\beta \cap S_\beta)]$$

This would imply that with non-negligible probability, we can find a tuple that belongs to S_β (and hence corresponds to a successful forgery) and at the same time belongs to Ω_β. We can then invoke the splitting lemma result just mentioned, to find a second tuple coresponding to a second forgery and that has the desired properties.

To prove the above, we proceed as follows:

$$P[\exists \beta \in I\ s.t.\ (\omega, r', \mathcal{H}) \in (\Omega_\beta \cap S_\beta) \mid (\omega, r', \mathcal{H}) \in S]$$

$$= P[\cup_{i\in I}\{(\omega, r', \mathcal{H}) \in (\Omega_i \cap S_i) \mid (\omega, r', \mathcal{H}) \in S\}]$$

$$= \textstyle\sum_{i\in I} P[(\omega, r', \mathcal{H}) \in (\Omega_i \cap S_i) \mid (\omega, r', \mathcal{H}) \in S], \text{ since the } S_i\text{'s are disjoint.}$$

$$=$$

$$\textstyle\sum_{i\in I}\{P[(\omega, r', \mathcal{H}) \in \Omega_i \mid (\omega, r', \mathcal{H}) \in S_i] \times P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S]\}$$

$$\geq \tfrac{1}{2} \textstyle\sum_{i\in I} P[(\omega, r', \mathcal{H}) \in S_i \mid (\omega, r', \mathcal{H}) \in S],\ \textit{(}3^{rd}\textit{ result of splitting lemma above)}$$

$$\geq \tfrac{1}{2} \times \tfrac{1}{2}\ \textit{(by the claim proven earlier)} = \tfrac{1}{4}.$$

And so we conclude that:

$$P_{(\omega,r',\mathcal{H})}[\exists \beta \in I\ s.t.\ (\omega, r', \mathcal{H}) \in (\Omega_\beta \cap S_\beta)]$$

$$= P[\exists \beta \in I\ s.t.\ (\omega, r', \mathcal{H}) \in (\Omega_\beta \cap S_\beta) \mid (\omega, r', \mathcal{H}) \in S] \times P_{(\omega,r',\mathcal{H})}[(\omega, r', \mathcal{H}) \in S]$$

$$\geq \tfrac{\nu(k)}{4}, \text{ which is non-negligible in } k.$$

So let β be such an index and $(\omega^*, r'^*, \mathcal{H}^*)$ such a tuple. From the result above, we know that finding such a $(\omega^*, r'^*, \mathcal{H}^*) \in (\Omega_\beta \cap S_\beta)$ can be done with non-negligible probability. And since $(\Omega_\beta \cap S_\beta) \subset \Omega_\beta$, we must have $(\omega^*, r'^*, \mathcal{H}^*) \in \Omega_\beta$. We can then invoke the 2^{nd} consequence of the splitting lemma, and write:

$$P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_\beta \mid (\omega^*, r'^*, \mathcal{H}^*) \in\ S_\beta,\ \rho_1 = \rho_1^*, ..., \rho_{\beta-1} = \rho_{\beta-1}^*)] =$$

$$P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_\beta \mid (\omega^*, r'^*, \mathcal{H}^*) \in \Omega_\beta,\ \rho_1 = \rho_1^*, ..., \rho_{\beta-1} = \rho_{\beta-1}^*)] \geq \tfrac{\nu(k)}{4Q}$$

We still have one last constraint to impose and that is that $\rho_\beta^* \neq \rho_\beta^{\sim}$. We show that the following quantity is non-negligible:

$$P_{\mathcal{H}}[((\omega^*, r'^*, \mathcal{H}) \in S_\beta) \ \cap \ (\rho_\beta \neq \rho^*_\beta)|\ (\omega^*, r'^*, \mathcal{H}^*) \in S_\beta,\ \rho_1 = \rho^*_1, ..., \rho_{\beta-1} = \rho^*_{\beta-1})]$$

To prove this, note that if B and C are independent events, then we can write:

$$P[A|C] = P[A \cap B|C] + P[A \cap \overline{B}|C] \leq P[A \cap B|C] + P[\overline{B}|C] = P[A \cap B|C] + P[\overline{B}]$$

And so we get $P[A \cap B|C] \geq P[A|C] - P[\overline{B}]$. This results allows us to write:

$$P_{\mathcal{H}}[((\omega^*, r'^*, \mathcal{H}) \in S_\beta) \ \cap \ (\rho_\beta \neq \rho^*_\beta)|\ (\omega^*, r'^*, \mathcal{H}^*) \in S_\beta,\ \rho_1 = \rho^*_1, ..., \rho_{\beta-1} = \rho^*_{\beta-1})]$$

$$\geq P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_\beta|\ (\omega^*, r'^*, \mathcal{H}^*) \in S_\beta,\ \rho_1 = \rho^*_1, ..., \rho_{\beta-1} = \rho^*_{\beta-1})] - P_{\mathcal{H}}[\rho_\beta = \rho^*_\beta]$$

$$= P_{\mathcal{H}}[(\omega^*, r'^*, \mathcal{H}) \in S_\beta|\ (\omega^*, r'^*, \mathcal{H}^*) \in \Omega_\beta,\ \rho_1 = \rho^*_1, ..., \rho_{\beta-1} = \rho^*_{\beta-1})] - P_{\mathcal{H}}[\rho_\beta = \rho^*_\beta]$$

(because we chose $(\omega^*, r'^*, \mathcal{H}^*) \in \Omega_\beta \cap S_\beta$)

$$\geq \frac{\nu(k)}{4Q} - \frac{1}{2^k}, \text{ which is non-negligible in } k.$$

Step 5 : The final step uses the 2 forgeries obtained earlier to solve an instance of the Discrete Logarithm (DL) problem. Here is a recap of Step 4 results:

- With non-negligible probability of at least $\frac{\nu(k)}{4}$ we get a successful tuple $(\omega^*, r'^*, \mathcal{H}^*)$, s.t. $(\omega^*, r'^*, \mathcal{H}^*) \in (\Omega_\beta \cap S_\beta)$ for some index $\beta \in I$. So by running $\mathcal{A}$ a number of times polynomial in k, we can confidently find such a tuple.
- Once we find such a tuple, we've also shown that with non-negligible probability of at least $\frac{\nu(k)}{4Q} - \frac{1}{2^k}$, we can find another successful tuple $(\omega^*, r'^*, \mathcal{H}^\sim)$ such that $(\omega^*, r'^*, \mathcal{H}^\sim) \in S_\beta$ and $(\rho^\sim_1 = \rho^*_1), ..., (\rho^\sim_{\beta-1} = \rho^*_{\beta-1})$, but $(\rho^\sim_\beta \neq \rho^*_\beta)$.

W.l.o.g, let $(\omega^*, r'^*, \mathcal{H}^*)$ correspond to $\sigma_{forge}(m_1) \equiv (r_1, h_1, \alpha_1)$, and $(\omega^*, r'^*, \mathcal{H}^\sim)$ correspond to $\sigma_{forge}(m_2) \equiv (r_2, h_2, \alpha_2)$.

Recall that β is the index of the query (m_1, r_1) that $\mathcal{A}$ sends to the RO. Since the 2 experiments corresponding to the 2 successful tuples have the same random tapes ω^* and r'^*, and since the 2 corresponding ROs $\mathcal{H}^*$ and $\mathcal{H}^\sim$ behave the same way on the first $\beta - 1$ queries, we can be confident that the first β queries sent to the 2 ROs are identical. In particular the two β^{th} queries are the same (i.e., $(m_1, r_1) = (m_2, r_2)$). Moreover by design, $h_1 = \mathcal{H}^*(m_1, r_1) = \rho^*_\beta \neq \rho^\sim_\beta = \mathcal{H}^\sim(m_1, r_1) = h_2$.

So we have 2 successful forgeries $\sigma_{forge}(m_1) \equiv (r_1, h_1, \alpha_1)$ and $\sigma_{forge}(m_1) \equiv (r_1, h_2, \alpha_2)$, with $h_1 \neq h_2$. Since both are valid signatures, they must satisfy the verification equations. For the particular case of a Schnorr signature scheme , they must satisfy the following 2 equations (1 verification equation per signature):

- $r_1 = g^{\alpha_1} y^{h_1}$, where y is the public key of the signer whose signature $\mathcal{A}$ is forging.
- $r_1 = g^{\alpha_2} y^{h_2}$, where y is the public key of the signer whose signature $\mathcal{A}$ is forging.

Writing $y = g^x$ (x is the secret key of the signer whose signature $\mathcal{A}$ is forging), we get:

$$g^{\alpha_1 + xh_1} = g^{\alpha_2 + xh_2} \iff \alpha_1 + xh_1 = \alpha_2 + xh_2 \longrightarrow x = \frac{\alpha_1 - \alpha_2}{h_2 - h_1}.$$

Since, $h_1 \neq h_2$, we can solve for x (the DL of y) in polynomial time. This contradicts the intractability of DL on multiplicative cyclic groups and we conclude that our signature scheme (in this case the Schnorr's scheme) is secure against EFACM in the RO model.

References

[1] D. Pointcheval and J. Stern. Security arguments for digital signatures and blind signatures. *Journal of Cryptology*, 2000.