MTAT.07.003 Cryptology II
Spring 2012 / Exercise session ?? / Example Solution

Exercise (Security of MAP-1 protocol). *Bellare and Rogaway proposed a two-party entity authentication protocol* MAP-1 *where parties* $\mathcal{P}_1$ *and* $\mathcal{P}_2$ *share the secret key* $k \xleftarrow{u} \mathcal{K}$*. This secret key is used a way to select a function instance from a pseudorandom function family* $f : \{0,1\}^* \times \mathcal{K} \to \mathcal{T}$*. To establish mutual authentication, parties execute the following protocol:*

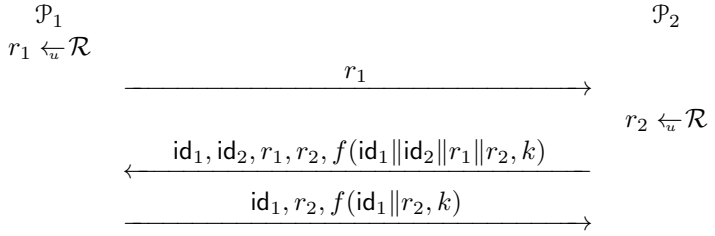

where $\mathcal{P}_1$ *halts in the second step if the authentication tag* $t_1 = f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| r_2, k)$ *and* $\mathcal{P}_2$ *halts in the third step if the authentication tag* $t_2 = f(\mathsf{id}_1 \| r_2, k)$ *is not is not consistent with their knowledge.*

Analyse the security of MAP-1 *protocol in the standalone setting, where* $\mathcal{P}_1$ *and* $\mathcal{P}_2$ *run a single instance of the protocol by sending messages through the adversary* $\mathcal{A}$ *who can alter, drop or insert messages into the conversation. An adversary* $\mathcal{A}$ *succeeds in deception if both parties reach the accepting state while some message in the protocol is altered or injected. For the security analysis, construct a game* $\mathcal{G}_0$ *where the adversary* $\mathcal{A}$ *wins if and only if the* $\mathcal{A}$ *would be successful against the protocol instance. After that show that the success probability can be bounded provided that* f *is a* (t, ε)*-secure pseudorandom function.*

Solution. A proper solution to this exercise consists of a formalisation of the security goal and the proof that the goal is achievable with negligible probability.

FORMALISATION OF THE SECURITY GOAL. We firstly formalise what constitutes a successful attack on the protocol instance. First of all, we can assume that the adversary only conducts attacks where the messages are modified as it is trivial for the parties to check if they have received some additional messages or some message is missing. Second, we can assume that the adversary does not modify messages id_1 and id_2, since parties would notice if these identities are changed – we explicitly assume that both of them knows to whom it wants to authenticate. Given these observations and the fact that the adversary has ability to intercept any of the messages between $\mathcal{P}_1$ and $\mathcal{P}_2$, we can construct the following game

$\mathcal{G}_0$
$k \xleftarrow{u} \mathcal{K}$
$r_1 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \xleftarrow{u} \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1 || \mathsf{id}_2 || \hat{r}_1 || r_2, k)$
$\hat{\hat{r}}_1, \hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $r_1 \neq \hat{\hat{r}}_1 \vee f(\mathsf{id}_1 || \mathsf{id}_2 || r_1 || \hat{r}_2, k) \neq \hat{t}_1$ then **return** 0
$t_2 \leftarrow f(\mathsf{id}_1 || \hat{r}_2, k)$
$\hat{\hat{r}}_2, \hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $r_2 \neq \hat{\hat{r}}_2 \vee f(\mathsf{id}_1 || r_2, k) \neq \hat{t}_2$ then **return** 0
return $[r_1 \neq \hat{r}_1] \vee [r_2 \neq \hat{r}_2] \vee [t_1 \neq \hat{t}_1] \vee [t_2 \neq \hat{t}_2]$

which models the protocol execution and ends with one only if the adversary achieves its goal. Instantly, we can apply some simplifications to the game. First of all, the adversary is guaranteed to loose the game whenever $r_1 \neq \hat{\hat{r}}_1$ or $r_2 \neq \hat{\hat{r}}_2$. Hence, we can assume that $\mathcal{A}$ outputs $\hat{\hat{r}}_1 = r_1$ and $\hat{\hat{r}}_2 = r_2$. If $\mathcal{A}$ does not do it

then we can write a wrapper around $\mathcal{A}$ that enforces this requirement. The resulting adversary $\mathcal{A}_*$ is at least as successful as $\mathcal{A}$. Now that we know that $\mathcal{A}$ always outputs outputs $\hat{r}_1 = r_1$ and $\hat{r}_2 = r_2$, we can further simplify the interface between the adversary and the game:

$$
\begin{array}{l}
\mathcal{G}_1 \\
\left\lceil
\begin{array}{l}
k \xleftarrow{u} \mathcal{K} \\
r_1 \xleftarrow{u} \mathcal{R} \\
\hat{r}_1 \leftarrow \mathcal{A}(r_1) \\
r_2 \xleftarrow{u} \mathcal{R} \\
t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2, k) \\
\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\mathsf{if}\ f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2, k) \neq \hat{t}_1\ \mathsf{then}\ \mathbf{return}\ 0 \\
t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2, k) \\
\hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\mathsf{if}\ f(\mathsf{id}_1||r_2, k) \neq \hat{t}_2\ \mathsf{then}\ \mathbf{return}\ 0 \\
\mathbf{return}\ [r_1 \neq \hat{r}_1] \vee [r_2 \neq \hat{r}_2] \vee [t_1 \neq \hat{t}_1] \vee [t_2 \neq \hat{t}_2]
\end{array}
\right.
\end{array}
$$

Note that games $\mathcal{G}_0$ and $\mathcal{G}_1$ are formally incompatible – we cannot use the adversary $\mathcal{A}_0$ against the game $\mathcal{G}_0$ in the game $\mathcal{G}_1$ and vice versa. However, it is trivial two write a wrapper around $\mathcal{A}_0$ that drops irrelevant outputs $\hat{r}_1$ and $\hat{r}_2$ to make it compatible with $\mathcal{G}_1$. As the success probability does not decease and the overhead in the running time is constant games $\mathcal{G}_0$ and $\mathcal{G}_1$ are equivalent.

Now note that if $r_1 = \hat{r}_1$ and $r_2 = \hat{r}_2$, then the adversary passes two verification checks in the middle of the game only if $t_1 = \hat{t}_1$ and $t_2 = \hat{t}_2$. The latter follows form the fact that f is deterministic function. Consequently, the adversary cannot win the game when $r_1 = \hat{r}_1$ and $r_2 = \hat{r}_2$. This allows to omit second half of the last check. As a result, we must analyse the security of the following game:

$$
\begin{array}{l}
\mathcal{G}_2 \\
\left\lceil
\begin{array}{l}
k \xleftarrow{u} \mathcal{K} \\
r_1, r_2 \xleftarrow{u} \mathcal{R} \\
\hat{r}_1 \leftarrow \mathcal{A}(r_1) \\
t_1 \leftarrow f(\mathsf{id}_1\|\mathsf{id}_2\|\hat{r}_1\|r_2, k) \\
\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\mathsf{if}\ f(\mathsf{id}_1\|\mathsf{id}_2\|r_1\|\hat{r}_2, k) \neq \hat{t}_1\ \mathsf{then}\ \mathbf{return}\ 0 \\
t_2 \leftarrow f(\mathsf{id}_1\|\hat{r}_2, k) \\
\hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\mathsf{if}\ f(\mathsf{id}_1\|r_2, k) \neq \hat{t}_2\ \mathsf{then}\ \mathbf{return}\ 0 \\
\mathbf{return}\ [r_1 \neq \hat{r}_1] \vee [r_2 \neq \hat{r}_2]\ .
\end{array}
\right.
\end{array}
$$

In other words, we must bound the success probability

$$\mathsf{Adv}^{\mathsf{win}}_{\mathcal{G}_2}(\mathcal{A}) = \Pr\left[\mathcal{G}_2^{\mathcal{A}} = 1\right]$$

for all t-time adversaries to prove the desired claim about the MAP-1 protocol.

Security analysis. Here, we consider three proof schemes for bounding the success probability $\mathsf{Adv}^{\mathsf{win}}_{\mathcal{G}_2}(\mathcal{A})$. They all rely on the same basic idea that without knowing the key k, the function f is indistinguishable from a random function, even if we know the input to the function. Different proofs use this equivalence in different places. The first proof splits the game into two sub-games checking conditions $[r_1 \neq \hat{r}_1]$ and $[r_2 \neq \hat{r}_2]$ separately and then replaces f with $f \leftarrow \mathcal{F}_{\text{all}}$. The second proof first replaces f with $f \leftarrow \mathcal{F}_{\text{all}}$ and

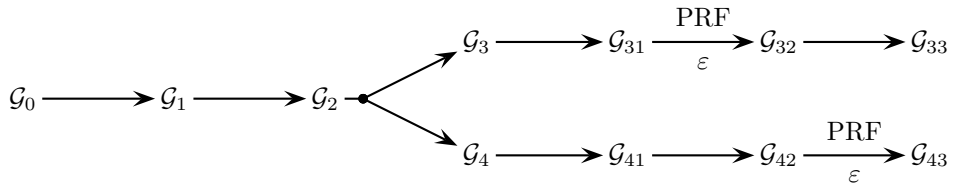

Figure 1: Game tree corresponding to the initial solution

then bounds the probability of events $[r_1 \neq \hat{r}_1]$ and $[r_2 \neq \hat{r}_2]$ separately. In the third proof, we replace f with $f \leftarrow \mathcal{F}_{\text{all}}$ and test for condition $[r_1 \neq \hat{r}_1] \vee [r_2 \neq \hat{r}_2]$ directly. The pictorial illustration of all three proofs are given in Figures 1, 2 and 3.

Initial solution. As $x \vee y = x \vee (\neg x \wedge y)$, then we can define two subgames:

$\mathcal{G}_3$

$k \xleftarrow{u} \mathcal{K}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| \hat{r}_1 \| r_2, k)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2, k) \neq \hat{t}_1$ then **return** 0
$t_2 \leftarrow f(\mathsf{id}_1 \| \hat{r}_2, k)$
$\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $f(\mathsf{id}_1 \| r_2, k) \neq \hat{t}_2$ **return** 0
return $[r_1 \neq \hat{r}_1]$

$\mathcal{G}_4$

$k \xleftarrow{u} \mathcal{K}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| \hat{r}_1 \| r_2, k)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2, k) \neq \hat{t}_1$ then **return** 0
$t_2 \leftarrow f(\mathsf{id}_1 \| \hat{r}_2, k)$
$\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $f(\mathsf{id}_1 \| r_2, k) \neq \hat{t}_2$ **return** 0
return $[r_1 \stackrel{?}{=} \hat{r}_1] \wedge [r_2 \neq \hat{r}_2]$

such that

$$\Pr\left[\mathcal{G}_2^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] + \Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] \ .$$

Next, we will bound the success probability of an adversary against these two games. For the game $\mathcal{G}_3$, we can upperbound the success probability of an adversary by converting $\mathcal{G}_3$ into an easier game

$\mathcal{G}_{31}$

$k \xleftarrow{u} \mathcal{K}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| \hat{r}_1 \| r_2, k)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2, k) \neq \hat{t}_1$ then **return** 0
return $[r_1 \neq \hat{r}_1]$.

Moreover, as f is a (t, ε)-pseudorandom function, there is at most ε distance between the games $\mathcal{G}_{31}$ and

the game

$\mathcal{G}_{32}$
$f \leftarrow \mathcal{F}_{\text{all}}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| \hat{r}_1 \| r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2) \neq \hat{t}_1$ then **return** 0
return $[r_1 \neq \hat{r}_1]$.

in which the function is changed into a random function. Note that adversary $\mathcal{A}$ can win only if it outputs $\hat{r}_1 \neq r_1$. We can easily modify $\mathcal{A}$ to always output $\hat{r}_1 \neq r_1$ by increasing the running time by a small constant. This change can only increase the success probability. Thus, we can further only observe adversaries $\mathcal{A}$ that always output $\hat{r}_1 \neq r_1$. If $\hat{r}_1 \neq r_1$ then also $\mathsf{id}_1 \| \mathsf{id}_2 \| \hat{r}_1 \| r_2 \neq \mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2$ and thus the random function f is evaluated at two different arguments. Consequently, we can replace function calls by random sampling. This leads to the following game:

$\mathcal{G}_{33}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$t_1 \xleftarrow{u} \mathcal{T}$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
$\hat{\hat{t}}_1 \xleftarrow{u} \mathcal{T}$
if $\hat{\hat{t}}_1 \neq \hat{t}_1$ then **return** 0
return $[r_1 \neq \hat{r}_1]$.

As the adversary chooses $\hat{t}_1$ before $\hat{\hat{t}}_1$ is sampled, we get

$$\Pr\left[\mathcal{G}_{33}^{\mathcal{A}} = 1\right] = \frac{1}{|T|} \ ,$$

which itself implies

$$\Pr\left[\mathcal{G}_{3}^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{31}^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{32}^{\mathcal{A}} = 1\right] + \varepsilon \leq \Pr\left[\mathcal{G}_{23}^{\mathcal{A}} = 1\right] + \varepsilon \leq \frac{1}{|T|} + \varepsilon \ .$$

Now consider the game $\mathcal{G}_4$. In this game $\mathcal{A}$ must create $r_1 = \hat{r}_1$ to win and thus we can always modify the adversary to output $r_1 = \hat{r}_1$ without decreasing the success probability. However, then the first interaction $\hat{r}_1 \leftarrow \mathcal{A}(r_1)$ is unnecessary. Therefore, we can further simplify the security game as follows:

$\mathcal{G}_{41}$
$k \xleftarrow{u} \mathcal{K}$
$r_1, r_2 \xleftarrow{u} \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| r_2, k)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| \hat{r}_2, k) \neq \hat{t}_1$ **return** 0
$t_2 \leftarrow f(\mathsf{id}_1 \| \hat{r}_2, k)$
$\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $f(\mathsf{id}_1 \| r_2, k) \neq \hat{t}_2$ **return** 0
return $[r_2 \neq \hat{r}_2]$.

To simplify the analysis, we can relax the game $\mathcal{G}_{41}$ by dropping the first check:

$$
\begin{array}{l}
\mathcal{G}_{42} \\
\lceil\, k \xleftarrow{u} \mathcal{K} \\
\quad r_1, r_2 \xleftarrow{u} \mathcal{R} \\
\quad t_1 \leftarrow f(\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| r_2, k) \\
\quad \hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\quad t_2 \leftarrow f(\mathsf{id}_1 \| \hat{r}_2, k) \\
\quad \hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\quad \mathsf{if}\ f(\mathsf{id}_1 \| r_2, k) \neq \hat{t}_2\ \mathbf{return}\ 0 \\
\quad \mathbf{return}\ [r_2 \neq \hat{r}_2]\ . \\
\rfloor
\end{array}
$$

Again, this can only increase the success probability. The game $\mathcal{G}_{42}$ has distance at most ε from the game where f is replaced by a random function. Again the adversary $\mathcal{A}$ must output $r_2 \neq \hat{r}_2$ to succeed and by using the analogous reasoning as above we can consider only such adversaries $\mathcal{A}$ that always output $r_2 \neq \hat{r}_2$. Again, if $r_2 \neq \hat{r}_2$ then also $\mathsf{id}_1 \| r_2 \neq \mathsf{id}_1 \| \hat{r}_2$. Since the third function argument $\mathsf{id}_1 \| \mathsf{id}_2 \| r_1 \| r_2$ has different length than the others, the random function f is evaluated on distinct inputs and we can replace function calls by random sampling without changing the semantics of the game. This leads to the following game:

$$
\begin{array}{l}
\mathcal{G}_{43} \\
\lceil\, k \xleftarrow{u} \mathcal{K} \\
\quad r_1, r_2 \xleftarrow{u} \mathcal{R} \\
\quad t_1 \xleftarrow{u} \mathcal{T} \\
\quad \hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\quad t_2 \xleftarrow{u} \mathcal{T} \\
\quad \hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\quad \hat{\hat{t}}_2 \xleftarrow{u} \mathcal{T} \\
\quad \mathsf{if}\ \hat{\hat{t}}_2 \neq \hat{t}_2\ \mathbf{return}\ 0 \\
\quad \mathbf{return}\ [r_2 \neq \hat{r}_2]\ . \\
\rfloor
\end{array}
$$

As before, we see that in $\mathcal{G}_{43}$, the adversary chooses $\hat{t}_2$ before $\hat{\hat{t}}_2$ is sampled and thus we can establish

$$\Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{41}^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{42}^{\mathcal{A}} = 1\right] + \varepsilon \leq \Pr\left[\mathcal{G}_{43}^{\mathcal{A}} = 1\right] + \varepsilon \leq \frac{1}{|T|} + \varepsilon\ .$$

Now by combining the results, we obtain a reasonable but loose bound

$$\mathsf{Adv}^{\mathsf{win}}_{\mathcal{G}_0}(\mathcal{A}) \leq \Pr\left[\mathcal{G}_2^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] + \Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] \leq \frac{2}{|\mathcal{T}|} + 2\varepsilon\ .$$

First improvement. The first solution is sub-optimal, as we had to replace function f with a random

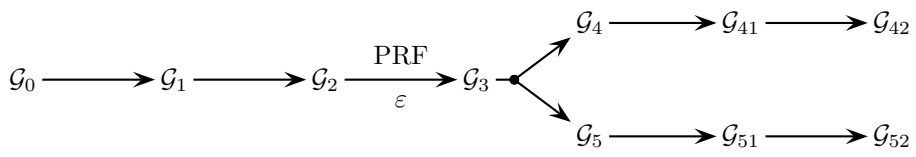

Figure 2: Game tree corresponding to the first improvement

function in both branches of the proof. We could do it directly by defining the game

$$
\begin{array}{l}
\mathcal{G}_3 \\
\left[\begin{array}{l}
f \xleftarrow{u} \mathcal{F}_{\text{all}} \\
r_1 \xleftarrow{u} \mathcal{R} \\
\hat{r}_1 \leftarrow \mathcal{A}(r_1) \\
r_2 \xleftarrow{u} \mathcal{R} \\
t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2) \\
\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\mathsf{if}\ f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2) \neq \hat{t}_1\ \mathsf{then}\ \mathbf{return}\ 0 \\
t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2) \\
\hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\mathsf{if}\ f(\mathsf{id}_1||r_2) \neq \hat{t}_2\ \mathsf{then}\ \mathbf{return}\ 0 \\
\mathbf{return}\ [r_1 \neq \hat{r}_1] \vee [r_2 \neq \hat{r}_2]
\end{array}\right.
\end{array}
$$

Since f is a (t, ε)-pseudorandom function and the key k is not leaked to the adversary, the change in the success probability is bounded:

$$\left|\Pr\left[\mathcal{G}_2^{\mathcal{A}} = 1\right] - \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right]\right| \leq \varepsilon \ .$$

Now for any $\mathcal{A}$, we can decompose the success probability into a sum

$$\Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] = \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1 \wedge r_1 \neq \hat{r}_1\right] + \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1 \wedge r_1 = \hat{r}_1 \wedge r_2 \neq \hat{r}_2\right] \ .$$

For both terms on the right, we can define corresponding games:

$\mathcal{G}_4$

$\lceil f \leftarrow_u \mathcal{F}_{\text{all}}$
$r_1 \leftarrow_u \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \leftarrow_u \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2) \neq \hat{t}_1$ then **return** 0
$t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2)$
$\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $f(\mathsf{id}_1||r_2) \neq \hat{t}_2$ then **return** 0
return $[r_1 \neq \hat{r}_1]$ $\rfloor$

$\mathcal{G}_5$

$\lceil f \leftarrow_u \mathcal{F}_{\text{all}}$
$r_1 \leftarrow_u \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \leftarrow_u \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2) \neq \hat{t}_1$ then **return** 0
$t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2)$
$\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
if $f(\mathsf{id}_1||r_2) \neq \hat{t}_2$ then **return** 0
return $[r_1 \stackrel{?}{=} \hat{r}_1] \wedge [r_2 \neq \hat{r}_2]$. $\rfloor$

Note that in the game $\mathcal{G}_4$ it is straightforward to pass the second check by submitting $\hat{r}_2 = r_2$ and $\hat{t}_2 = t_2$. Thus the second check does not make the game harder for the adversary and we can drop it. As a result, we obtain the following game:

$\mathcal{G}_{41}$

$\lceil f \leftarrow_u \mathcal{F}_{\text{all}}$
$r_1 \leftarrow_u \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \leftarrow_u \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2) \neq \hat{t}_1$ then **return** 0
return $[r_1 \neq \hat{r}_1]$ $\rfloor$

The adversary can win this game only if $r_1 \neq \hat{r}_1$ and thus we can assume without loss of generality that $\mathcal{A}$ always outputs $\hat{r}_1 \neq r_1$. As result, the function f is evaluated on distinct arguments and both evaluations can be replaced with random sampling:

$\mathcal{G}_{42}$

$\lceil f \leftarrow_u \mathcal{F}_{\text{all}}$
$r_1 \leftarrow_u \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \leftarrow_u \mathcal{R}$
$t_1 \leftarrow_u \mathcal{T}$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
$\hat{\hat{t}}_1 \leftarrow_u \mathcal{T}$
if $\hat{\hat{t}}_1 \neq \hat{t}_1$ then **return** 0
return $[r_1 \neq \hat{r}_1]$ $\rfloor$

which itself implies

$$\Pr\left[\mathcal{G}_{41}^{\mathcal{A}} = 1\right] = \frac{1}{|\mathcal{T}|} \ .$$

To analyse the game $\mathcal{G}_5$, note that for the analogous reasoning we can assume that $\mathcal{A}$ always outputs $\hat{r}_1 = r_1$ and $\hat{r}_2 \neq r_2$. As a result, we can simplify the game:

$$
\begin{array}{l}
\mathcal{G}_{51} \\
\left\lceil
\begin{array}{l}
f \xleftarrow{u} \mathcal{F}_{\text{all}} \\
r_1 \xleftarrow{u} \mathcal{R} \\
r_2 \xleftarrow{u} \mathcal{R} \\
t_1 \leftarrow f(\mathsf{id}_1 || \mathsf{id}_2 || r_1 || r_2) \\
\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\textsf{if } f(\mathsf{id}_1 || \mathsf{id}_2 || r_1 || \hat{r}_2) \neq \hat{t}_1 \textsf{ then } \textbf{return } 0 \\
t_2 \leftarrow f(\mathsf{id}_1 || \hat{r}_2) \\
\hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\textsf{if } f(\mathsf{id}_1 || r_2) \neq \hat{t}_2 \textsf{ then } \textbf{return } 0 \\
\textbf{return } [r_1 \stackrel{?}{=} \hat{r}_1] \wedge [r_2 \neq \hat{r}_2] \ .
\end{array}
\right.
\end{array}
$$

Again, it is easy to see that under these assumption all four evaluations of f are done on distinct inputs and we can replace them all with random sampling:

$$
\begin{array}{l}
\mathcal{G}_{52} \\
\left\lceil
\begin{array}{l}
r_1 \xleftarrow{u} \mathcal{R} \\
r_2 \xleftarrow{u} \mathcal{R} \\
t_1 \xleftarrow{u} \mathcal{T} \\
\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1) \\
\hat{\hat{t}}_1 \xleftarrow{u} \mathcal{T} \\
\textsf{if } \hat{\hat{t}}_1 \neq \hat{t}_1 \textsf{ then } \textbf{return } 0 \\
t_2 \leftarrow f(\mathsf{id}_1 || \hat{r}_2) \\
\hat{t}_2 \leftarrow \mathcal{A}(t_2) \\
\hat{\hat{t}}_2 \xleftarrow{u} \mathcal{T} \\
\textsf{if } \hat{\hat{t}}_2 \neq \hat{t}_2 \textsf{ then } \textbf{return } 0 \\
\textbf{return } [r_1 \stackrel{?}{=} \hat{r}_1] \wedge [r_2 \neq \hat{r}_2] \ .
\end{array}
\right.
\end{array}
$$

It is straightforward to see that the adversary passes the first check with the probability $\frac{1}{|\mathcal{T}|}$ and both checks with the probability

$$\Pr\left[\mathcal{G}_{52}^{\mathcal{A}} = 1\right] = \frac{1}{|\mathcal{T}|^2} \ .$$

By combining all results, we get a much tighter bound:

$$\Pr\left[\mathcal{G}_0^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] + \varepsilon \leq \Pr\left[\mathcal{G}_{42}^{\mathcal{A}} = 1\right] + \Pr\left[\mathcal{G}_{52}^{\mathcal{A}} = 1\right] + \varepsilon \leq \frac{1}{|\mathcal{T}|} + \frac{1}{|\mathcal{T}|^2} + \varepsilon \ .$$

Final proof. The result obtained in the last proof is much better but still sub-optimal. Intuitively, the adversary must guess either the value of $f(\mathsf{id}_1 || \mathsf{id}_2 || r_1 || \hat{r}_2)$ or $f(\mathsf{id}_1 || r_2)$ in game $\mathcal{G}_3$ depending whether it outputs $\hat{r}_1 \neq r_1$ or $\hat{r}_2 \neq r_2$. Thus, we should be able to prove $\mathsf{Adv}^{\mathsf{win}}_{\mathcal{G}_3}(\mathcal{A}) \leq \frac{1}{|\mathcal{T}|}$, although we somehow obtained the slightly looser bound:

$$\Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] \leq \frac{1}{|\mathcal{T}|} + \frac{1}{|\mathcal{T}|^2} \ .$$

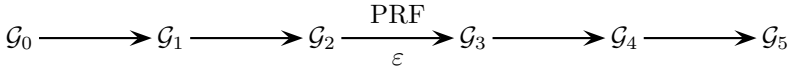

Figure 3: Game tree corresponding to the final proof

The reason why the bound is sub-optimal is subtle. First of all note that the claim

$$\Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] = \Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] + \Pr\left[\mathcal{G}_5^{\mathcal{A}} = 1\right]$$

is tight. Also, the bounds

$$\Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{41}^{\mathcal{A}} = 1\right]$$
$$\Pr\left[\mathcal{G}_5^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_{51}^{\mathcal{A}} = 1\right]$$

are separately taken optimal. However, we maximise probabilities in both branches separately and this leads to sub-optimality:

$$\max_{\mathcal{A}} \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] < \max_{\mathcal{A}} \Pr\left[\mathcal{G}_4^{\mathcal{A}} = 1\right] + \max_{\mathcal{A}} \Pr\left[\mathcal{G}_5^{\mathcal{A}} = 1\right] \quad .$$

To getter better result, we have to rewrite the game $\mathcal{G}_3$ in a more revealing way:

$\mathcal{G}_4$
$f \xleftarrow{u} \mathcal{F}_{\text{all}}$
$r_1 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \xleftarrow{u} \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $r_1 \neq \hat{r}_1$ then
  if $f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2) \neq \hat{t}_1$ then **return** 0
  $t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2)$
  $\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
  if $f(\mathsf{id}_1||r_2) \neq \hat{t}_2$ then **return** 0
  return 1
else if $r_2 \neq \hat{r}_2$ then
  $t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2)$
  $\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
  if $f(\mathsf{id}_1||r_2) \neq \hat{t}_2$ then **return** 0
  return 1
else **return** 0 .

It is easy to see that the adversary can win the branch $r_1 \neq \hat{r}_1$ if it manages to pass the first test by choosing $\hat{r}_2 = r_2$ and $\hat{t}_2 = t_2$. Hence, we can drop the second test. As $r_1 \neq \hat{r}_1$ in this branch, we can conclude that $\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2 \neq \mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2$ and thus we can replace the evaluation of $f(\mathsf{id}_1||\mathsf{id}_2||r_1||\hat{r}_2)$ by random sampling. In the branch $r_2 \neq \hat{r}_2$, we query the function f on an argument that has different length than

the previous argument $\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2$. Hence, we can replace this function evaluation also with the random sampling. This leads to a modified game:

$\mathcal{G}_5$
$f \xleftarrow{u} \mathcal{F}_{\text{all}}$
$r_1 \xleftarrow{u} \mathcal{R}$
$\hat{r}_1 \leftarrow \mathcal{A}(r_1)$
$r_2 \xleftarrow{u} \mathcal{R}$
$t_1 \leftarrow f(\mathsf{id}_1||\mathsf{id}_2||\hat{r}_1||r_2)$
$\hat{r}_2, \hat{t}_1 \leftarrow \mathcal{A}(r_2, t_1)$
if $r_1 \neq \hat{r}_1$ then
  $\hat{\hat{t}}_1 \xleftarrow{u} \mathcal{T}$
  if $\hat{\hat{t}}_1 \neq \hat{t}_1$ then **return** 0
  return 1
else if $r_2 \neq \hat{r}_2$ then
  $t_2 \leftarrow f(\mathsf{id}_1||\hat{r}_2)$
  $\hat{t}_2 \leftarrow \mathcal{A}(t_2)$
  $\hat{\hat{t}}_2 \xleftarrow{u} \mathcal{T}$
  if $\hat{\hat{t}}_2 \neq \hat{t}_2$ then **return** 0
  return 1
else **return** 0 .

From this construction it is evident that

$$\Pr\left[\mathcal{G}_5^{\mathcal{A}} = 1\right] = \Pr\left[r_1 \neq \hat{r}_1\right] \cdot \Pr\left[\hat{\hat{t}}_1 = \hat{t}_1\right] + \Pr\left[r_1 = \hat{r}_1 \wedge r_2 \neq \hat{r}_2\right] \cdot \Pr\left[\hat{\hat{t}}_2 = \hat{t}_2\right] \leq \frac{1}{|\mathcal{T}|}$$

which gives the optimal success bound

$$\Pr\left[\mathcal{G}_0^{\mathcal{A}} = 1\right] \leq \Pr\left[\mathcal{G}_3^{\mathcal{A}} = 1\right] + \varepsilon \leq \Pr\left[\mathcal{G}_5^{\mathcal{A}} = 1\right] + \varepsilon \leq \frac{1}{|\mathcal{T}|} + \varepsilon \ .$$

Conclusion. The example proofs show that splitting the proof into several branches is a powerful technique. However, there is a cost to pay. First of all, you have to use cryptographic assumptions separately in each branch and this leads to more relaxed upper bounds compared to the case when you use cryptographic assumptions before branching. Secondly, even the game simplification through branching can lead to sub-optimal result. The latter does not mean you should avoid branching. Branching is a valuable technique, as it provides a fast way to get the initial bound that can be later reduced by more advanced analysis.