A Crash Course to Coin Flipping

Sven Laur
swen@math.ut.ee

University of Tartu

Coin flipping by telephone

	Gen	
$b_1 \xleftarrow[u]{} \{0,1\}$	pk $\qquad\qquad$ pk	$b_2 \xleftarrow[u]{} \{0,1\}$
$(c,d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(b_1)$	$\xrightarrow{\ c\ }$	
	$\xleftarrow{\ b_2\ }$	
	$\xrightarrow{\ d\ }$	$b_1 \leftarrow \mathsf{Open}_{\mathsf{pk}}(c,d)$
return $b_1 \oplus b_2$	 	return $b_1 \oplus b_2$

The protocol above assures that participants output a uniformly distributed bit even if one of the participants is malicious.

- ▷ If the commitment scheme is perfectly binding, then Lucy can also generate public parameters for the commitment scheme.
- ▷ If the commitment scheme is perfectly hiding, then Charlie can also generate public parameters for the commitment scheme.

Weak security guarantee

Theorem. If we consider only such adversarial strategies that do not cause premature halting and additionally assume that the commitment scheme is (t, ε_1)-hiding and (t, ε_2)-binding, then

$$\frac{1}{2} - \max\{\varepsilon_1, \varepsilon_2\} \leq \Pr[b_1 \oplus b_2 = 1] \leq \frac{1}{2} + \max\{\varepsilon_1, \varepsilon_2\}$$

provided that at least one participant is honest.

Proof

- ▷ Lucy cannot cheat unless it double opens the commitment.
- ▷ As commitment is hiding the Charlie cannot guess b_1.

Real and Ideal World

Real versus ideal world approach

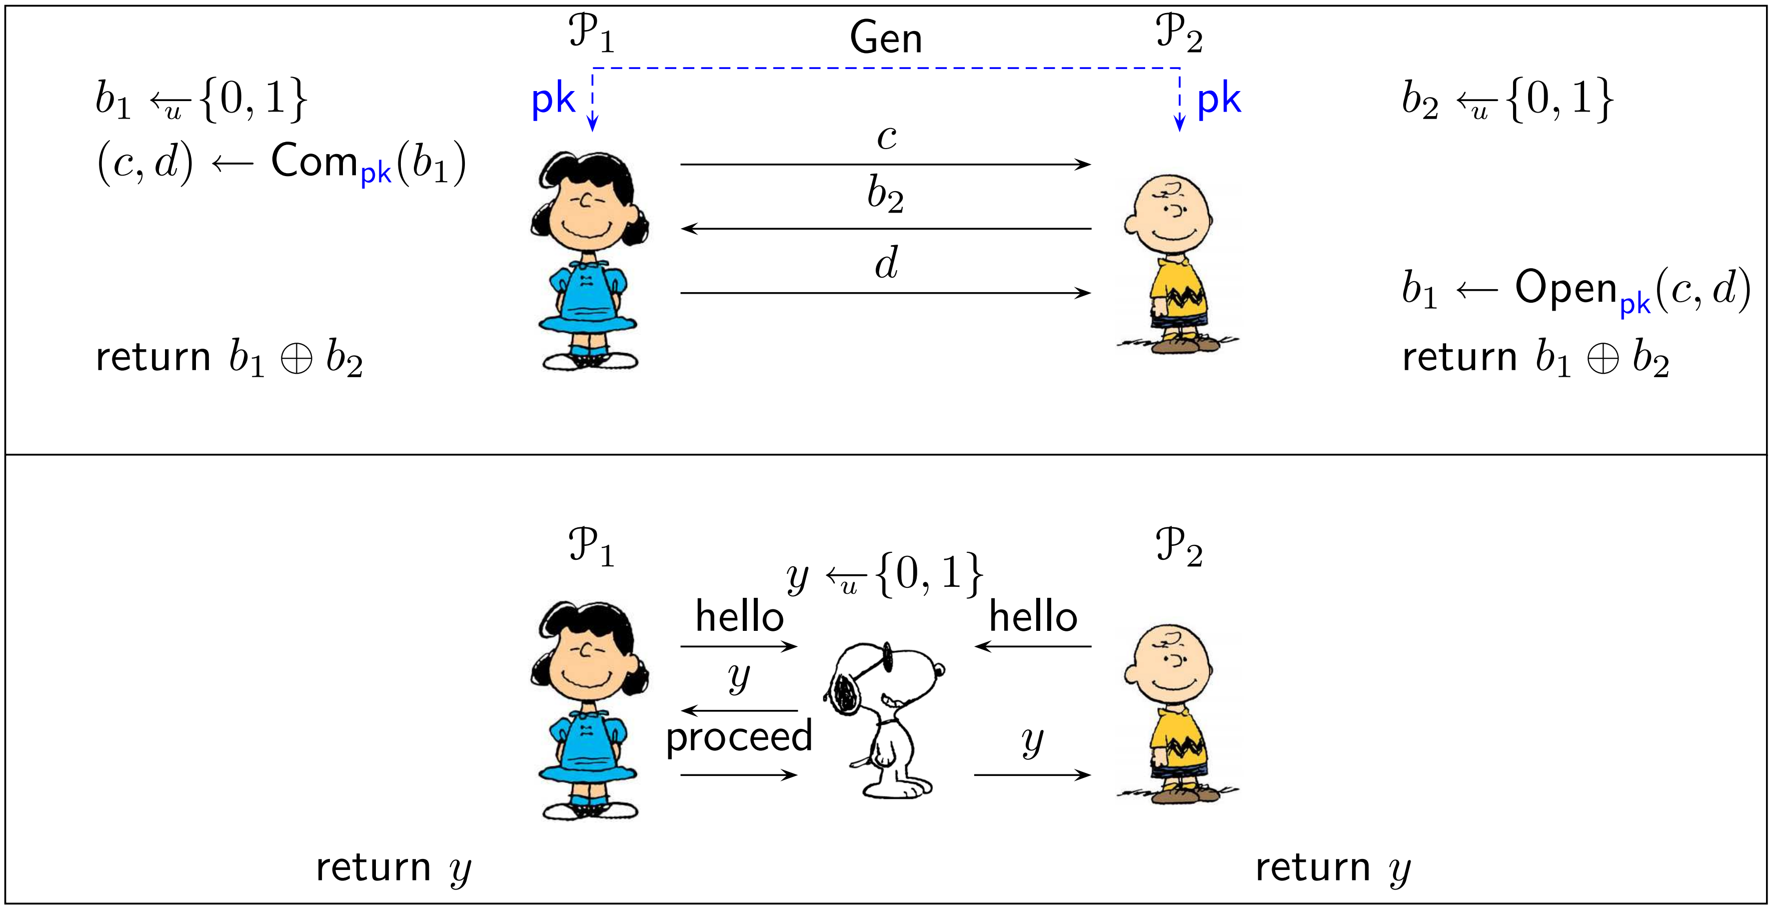

Formal definition

Let $\boldsymbol{\phi} = (\phi_1, \phi_2, \phi_{\mathrm{a}})$ be the set of input states of protocol participants $\mathcal{P}_1$ and $\mathcal{P}_2$, and the adversary $\mathcal{A}$ before the protocol. Let $\boldsymbol{\psi} = (\psi_1, \psi_2, \psi_{\mathrm{a}})$ be the set of output states after the execution of the protocol.

Similarly, let $\boldsymbol{\phi}^\circ = (\phi_1^\circ, \phi_2^\circ, \phi_{\mathrm{a}}^\circ)$ and $\boldsymbol{\psi}^\circ = (\psi_1^\circ, \psi_2^\circ, \psi_{\mathrm{a}}^\circ)$ denote the input and output states in the ideal world. Normally, one assumes that $\boldsymbol{\phi}^\circ \equiv \boldsymbol{\phi}$.

A protocol is $(t_{\mathrm{re}}, t_{\mathrm{id}}, \varepsilon)$-secure if for any t_{re}-time real world adversary $\mathcal{A}$ there exists a t_{id}-time ideal world adversary $\mathcal{A}^\circ$ such that for any input distribution $\mathfrak{D}$ the output distributions $\boldsymbol{\psi}$ and $\boldsymbol{\psi}^\circ$ are statistically ε-close.

The exact nature of the definition depends on the details

- ▷ What kind of malicious behaviour is allowed...
- ▷ What kind of ideal world model we use...
- ▷ In which contexts the protocol is executed...

Canonical constructive correspondence

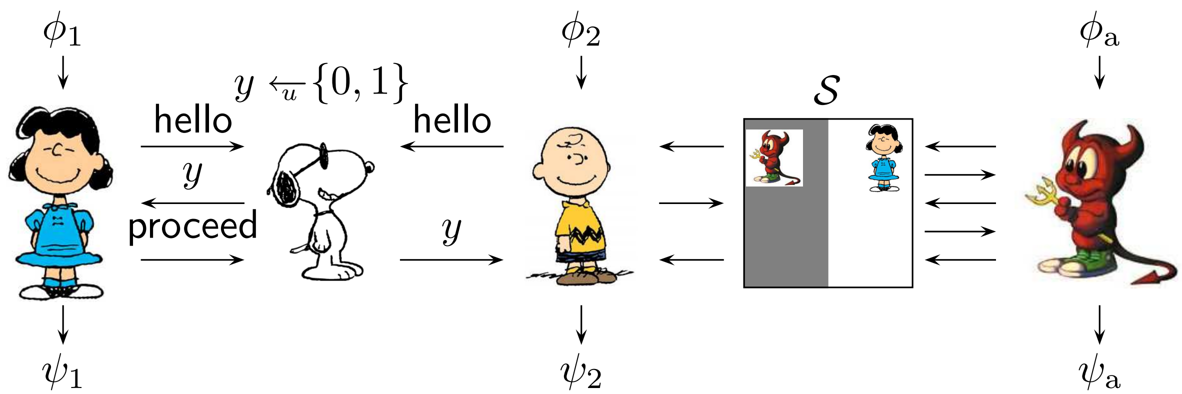

The desired mapping $\mathcal{A} \mapsto \mathcal{A}^\circ$ is defined through a code wrapper $\mathcal{S}$.

- ▷ The simulator $\mathcal{S}$ controls corrupted parties:
  - ◇ it submits their inputs to the trusted party $\mathcal{T}$,
  - ◇ it learns the response of $\mathcal{T}$.
- ▷ The simulator $\mathcal{S}$ controls the adversary $\mathcal{A}$:
  - ◇ it must mimic the real protocol execution,
  - ◇ it can rewind adversary if something goes wrong.

Simulator for the second party

$\mathcal{S}_2^{\mathcal{P}_2^*}(y)$

$\omega_2 \leftarrow \Omega_2$, $\boxed{\mathsf{pk} \leftarrow \mathsf{Gen}}$

For $i = 1, \ldots k$ do

$b_1 \xleftarrow[u]{} \{0, 1\}$

$(c, d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(b_1)$

$b_2 \leftarrow \mathcal{P}_2^*(\mathsf{pk}, c; \omega_2)$

if $b_1 \oplus b_2 = y$ then

Send d to $\mathcal{P}_2^*$ and output whatever $\mathcal{P}_2^*$ outputs.

return Failure

Failure probability

$\mathcal{S}_2^{\mathcal{P}_2^*}(y)$

- $\omega_2 \leftarrow \Omega_2$, $\mathsf{pk} \leftarrow \mathsf{Gen}$
- For $i = 1, \ldots k$ do
  - $b_1 \xleftarrow[u]{} \{0,1\}$
  - $\boxed{(c,d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(b_1)}$
  - $b_2 \leftarrow \mathcal{P}_2^*(\mathsf{pk}, c; \omega_2)$
  - if $b_1 \oplus b_2 = y$ then
    - return Success
- return Failure

$\mathcal{S}_4^{\mathcal{P}_2^*}(y)$

- $\omega_2 \leftarrow \Omega_2$, $\mathsf{pk} \leftarrow \mathsf{Gen}$
- For $i = 1, \ldots k$ do
  - $b_1 \xleftarrow[u]{} \{0,1\}$
  - $\boxed{(c,d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(0)}$
  - $b_2 \leftarrow \mathcal{P}_2^*(\mathsf{pk}, c; \omega_2)$
  - if $b_1 \oplus b_2 = y$ then
    - return Success
- return Failure

$\mathcal{S}_6^{\mathcal{P}_2^*}(y)$

- $\omega_2 \leftarrow \Omega_2$, $\mathsf{pk} \leftarrow \mathsf{Gen}$
- For $i = 1, \ldots k$ do
  - $(c,d) \leftarrow \mathsf{Com}_{\mathsf{pk}}(0)$
  - $b_2 \leftarrow \mathcal{P}_2^*(\mathsf{pk}, c; \omega_2)$
  - $b_1 \xleftarrow[u]{} \{0,1\}$
  - if $b_1 \oplus b_2 = y$ then
    - return Success
- return Failure

If commitment scheme is $(k \cdot t, \varepsilon_1)$-hiding, then for any t-time adversary $\mathcal{P}_2^*$ the failure probability

$$\Pr\left[\mathsf{Failure}\right] \leq \Pr\left[\mathcal{S}_6^{\mathcal{P}_2^*}(y) = \mathsf{Failure}\right] + k \cdot \varepsilon_1 \leq 2^{-k} + k \cdot \varepsilon_1 \enspace .$$

The corresponding security guarantee

If the output y is chosen uniformly over $\{0,1\}$, then the last effective value of b_1 has also an almost uniform distribution: $\left|\Pr\left[b_1 = 1 \middle| \neg\mathsf{Failure}\right] - \frac{1}{2}\right| \leq k \cdot \varepsilon_1$. Hence, the outputs of games

$$\mathcal{G}_{\mathrm{ideal}}^{\mathcal{S}_2^{\mathcal{P}_2^*}} \left[\begin{array}{l} (\phi_1, \phi_2) \leftarrow \mathfrak{D} \\ y \xleftarrow[u]{} \{0,1\} \\ \psi_1 \leftarrow (\phi_1, y) \\ \psi_2 \leftarrow \mathcal{S}_2^{\mathcal{P}_2^*(\phi_2)} \\ \text{return } (\psi_1, \psi_2) \end{array}\right. \qquad \mathcal{G}_{\mathrm{real}}^{\mathcal{P}_2^*} \left[\begin{array}{l} (\phi_1, \phi_2) \leftarrow \mathfrak{D} \\ \mathcal{P}_1 \text{ and } \mathcal{P}_2^* \text{ run the protocol.} \\ \psi_1 \leftarrow \mathcal{P}_1 \\ \psi_2 \leftarrow \mathcal{P}_2^* \\ \text{return } (\psi_1, \psi_2) \end{array}\right.$$

are at most $k \cdot \varepsilon_2$ apart if the run of $\mathcal{S}_2^{\mathcal{P}_2^*}$ is successful. Consequently, the statistical distance between output distributions is at most $2^{-k} + 2k \cdot \varepsilon_1$.

Simulator for the first party

$\mathcal{S}_1^{\mathcal{P}_1^*}(y)$

$\omega_1 \xleftarrow[u]{} \Omega_1,\ \boxed{\mathsf{pk} \leftarrow \mathsf{Gen}},\ c \leftarrow \mathcal{P}_1^*(\mathsf{pk};\omega_1)$

$d_0 \leftarrow \mathcal{P}_1^*(0;\omega_1),\ d_1 \leftarrow \mathcal{P}_1^*(1;\omega_1)$

$b_1^0 \leftarrow \mathsf{Open}_{\mathsf{pk}}(c, d_0),\ b_1^1 \leftarrow \mathsf{Open}_{\mathsf{pk}}(c, d_1)$

if $\perp \neq b_1^0 \neq b_1^1 \neq \perp$ then Failure

if $b_1^0 = \perp = b_1^1$ then

- Send the Halt command to $\mathcal{T}$.
- Choose $b_2 \xleftarrow[u]{} \{0,1\}$ and re-run the protocol with ω_1 and b_2.
- Return whatever $\mathcal{P}_1^*$ returns.

if $b_1^0 = \perp$ then $b_1 \leftarrow b_1^1$ else $b_1 \leftarrow b_1^0$

$b_2 \leftarrow b_1 \oplus y$

Re-run the protocol with ω_1 and b_2

if $b_1^{b_2} = \perp$ then Send the Halt command to $\mathcal{T}$.

Return whatever $\mathcal{P}_1^*$ returns.

Further analysis

If the commitment scheme is (t, ε_2)-binding, then the failure probability is less than ε_2. If the output y is chosen uniformly over $\{0, 1\}$, then the value of b_2 seen by $\mathcal{P}_1^*$ is uniformly distributed.

Consequently, the output distributions of $\mathcal{S}_1^{\mathcal{P}_1^*}$ and $\mathcal{P}_2$ in the ideal world coincide with the real world outputs if $\mathcal{S}_1$ does not fail.

Strong security guarantee

Theorem. If a commitment scheme is $(k \cdot t, \varepsilon_1)$-hiding and (t, ε_2)-binding, then for any plausible t-time real world adversary there exists $\mathrm{O}(k \cdot t)$-time ideal world adversary such that the output distributions in the real and ideal world are $\max\left\{2^{-k} + 2k \cdot \varepsilon_1, \varepsilon_2\right\}$-close.

Corollary. (Weak security guarantee) If we consider only such adversarial strategies that do not cause premature halting and additionally assume that the commitment scheme is $(k \cdot t, \varepsilon_1)$-hiding and (t, ε_2)-binding, then

$$\frac{1}{2} - \max\left\{2^{-k} + 2k \cdot \varepsilon_1, \varepsilon_2\right\} \leq \Pr\left[b_1 \oplus b_2 = 1\right] \leq \frac{1}{2} + \max\left\{2^{-k} + 2k \cdot \varepsilon_1, \varepsilon_2\right\}$$

provided that at least one participant is honest.

Sequential composition

If we execute the Blum protocol π sequentially ℓ times, then we can also stack simulators sequentially to get the ideal world adversary.

$$\mathcal{G}_{\mathrm{real}}^{\mathcal{P}_1^*}\left[\begin{array}{l}(\phi_1,\phi_2)\leftarrow\mathfrak{D}\\ \text{Run } \pi \text{ to get } (\psi_1,\psi_2)\\ (\phi_1,\phi_2)\leftarrow(\psi_1,\psi_2)\\ \text{Run } \pi \text{ to get } (\psi_1,\psi_2)\\ \ldots\\ \text{return } (\psi_1,\psi_2)\end{array}\right.\qquad\mathcal{G}_{\mathrm{ideal}}^{(\mathcal{S}_1^*)^{\mathcal{P}_1^*}}\left[\begin{array}{l}(\phi_1,\phi_2)\leftarrow\mathfrak{D}\\ \text{Use } \mathcal{S}_1 \text{ to get } (\psi_1,\psi_2)\\ (\phi_1,\phi_2)\leftarrow(\psi_1,\psi_2)\\ \text{Use } \mathcal{S}_1 \text{ to get } (\psi_1,\psi_2)\\ \ldots\\ \text{return } (\psi_1,\psi_2)\end{array}\right.$$

The final difference is a sum of individual differences.

Parallel composition

$\mathcal{P}_1$ Gen $\mathcal{P}_2$

$b_1^1, \ldots, b_1^\ell \xleftarrow[u]{} \{0,1\}$ pk pk $b_2^1, \ldots b_2^\ell \xleftarrow[u]{} \{0,1\}$

$(c_i, d_i) \leftarrow \mathsf{Com}_{\mathsf{pk}}(b_1^i)$

$\xrightarrow{c_1, \ldots, c_\ell}$

$\xleftarrow{b_2^1, \ldots, b_2^\ell}$

$\xrightarrow{d_1, \ldots, d_\ell}$

$b_1^i \leftarrow \mathsf{Open}_{\mathsf{pk}}(c_1, d_i)$

return $\boldsymbol{b}_1 \oplus \boldsymbol{b}_2$ return $\boldsymbol{b}_1 \oplus \boldsymbol{b}_2$

The simulation of this protocol is significantly more complex

- ▷ The number of potential replies $b_2^1, \ldots b_2^\ell$ grows exponentially wrt ℓ.
- ▷ We cannot sequentially alter values $c_1, \ldots, c_\ell$ to get the correct output.

Classical simulation strategies have exponential time-complexity wrt ℓ.

Non-rewinding simulators

▷ If the commitment scheme is extractable, then the simulator $\mathcal{S}_1$ can create $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}$ and choose b_2 according to $\mathsf{Extr}_{\mathsf{sk}}(c)$.

▷ If the commitment scheme is equivocable, then the simulator $\mathcal{S}_2$ can create $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}$ and then send a fake commitment to $\mathcal{P}_2^*$ and later open it with $\mathsf{Equiv}_{\mathsf{sk}}$ according to the reply b_2 to get the desired output.

▷ If the commitment scheme is both extractable and equivocable, then simulators $\mathcal{S}_1$ and $\mathcal{S}_2$ are non-rewinding and it is easy to construct simulators also for the parallel composition of several protocols.