Entity Authentication

Sven Laur
swen@math.ut.ee

University of Tartu

Formal Syntax

Entity authentication

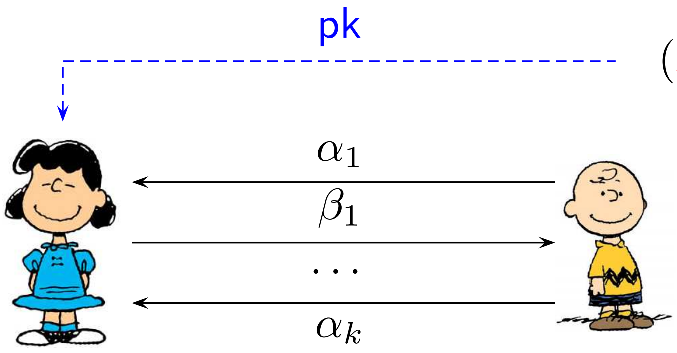

pk

$(\mathsf{sk}, \mathsf{pk}) \leftarrow \mathsf{Gen}$

α_1

β_1

$\dots$

α_k

$\beta_i \leftarrow \mathcal{V}_{\mathsf{pk}}(\alpha_1, \dots, \alpha_{i-1})$

$\alpha_i \leftarrow \mathcal{P}_{\mathsf{sk}}(\beta_1, \dots, \beta_{i-1})$

Is it Charlie?

- ▷ The communication between the prover and verifier must be authentic.
- ▷ To establish electronic identity, Charlie must generate $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}$ and convinces others that the public information pk represents him.
- ▷ The entity authentication protocol must convince the verifier that his or her opponent possesses the secret sk.
- ▷ An entity authentication protocol is functional if an honest verifier $\mathcal{V}_{\mathsf{pk}}$ always accepts an honest prover $\mathcal{P}_{\mathsf{sk}}$.

Classical impossibility results

Inherent limitations. Entity authentication is impossible if

(i) authenticated communication is unaffordable in the setup phase.

(ii) authenticated communication is unaffordable in the second phase.

Proof: Man-in-the-middle attacks. Chess-master attacks.

Conclusions

- ▷ It is impossible to establish legal identity without physical measures.
- ▷ Any bank-card is susceptible to physical attacks regardless of the cryptographic countermeasures used to authenticate transactions.
- ▷ Secure e-banking is impossible if the user does not have full control over the computing environment (secure e-banking is practically impossible).

Physical and legal identities

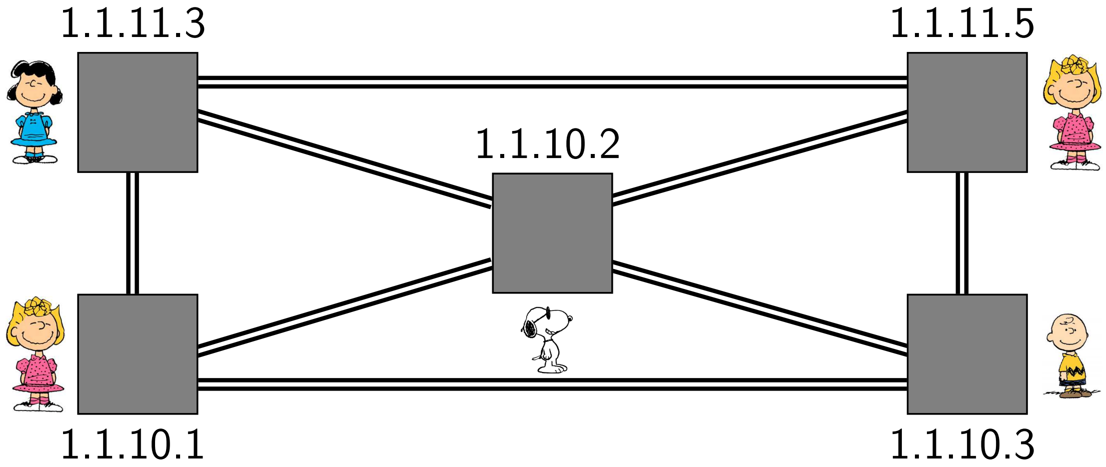

- ▷ Entity authentication is possible only if all participants have set up a network with authenticated communication links.
- ▷ A role of a entity authentication protocol is to establish a convincing bound between physical network address and legal identities.
- ▷ A same legal identity can be in many physical locations and move from one physical node to another node.

Challenge-Response Paradigm

Salted hashing

Global setup:

Authentication server $\mathcal{V}$ outputs a description of a hash function h.

Entity creation:

A party $\mathcal{P}$ chooses a password sk $\xleftarrow[u]{}\{0,1\}^\ell$ and a nonce $r \xleftarrow[u]{}\{0,1\}^k$. The public authentication information is $\mathsf{pk} = (r, c)$ where $c \leftarrow h(\mathsf{sk}, r)$.

Entity authentication:

To authenticate him- or herself, $\mathcal{P}$ releases sk to the server $\mathcal{V}$ who verifies that the hash value is correctly computed, i.e., $c = h(\mathsf{sk}, r)$.

Theorem. If h is (t, ε)-secure one-way function, then no t-time adversary $\mathcal{A}$ without sk can succeed in the protocol with probability more than ε.

- ▷ There are no secure one-way functions for practical sizes of sk.
- ▷ A malicious server can completely break the security.

RSA based entity authentication

Global setup:

Authentication server $\mathcal{V}$ fixes the minimal size of RSA keys.

Entity creation:

A party $\mathcal{P}$ runs a RSA key generation algorithm $(\mathsf{pk}, \mathsf{sk}) \leftarrow \mathsf{Gen}_{\mathrm{rsa}}$ and outputs the public key pk as the authenticating information.

Entity authentication:

1. $\mathcal{V}$ creates a challenge $c \leftarrow \mathsf{Enc}_{\mathsf{pk}}(m)$ for $m \xleftarrow[u]{} \mathcal{M}$ and sends c to $\mathcal{P}$.
2. $\mathcal{P}$ sends back $\overline{m} \leftarrow \mathsf{Dec}_{\mathsf{sk}}(c)$.
3. $\mathcal{V}$ accepts the proof if $m = \overline{m}$.

This protocol can be generalised for any public key cryptosystem.
The general form of this protocol is known as challenge-response protocol.
This mechanism provides explicit security guarantees in the SSL protocol.

The most powerful attack model

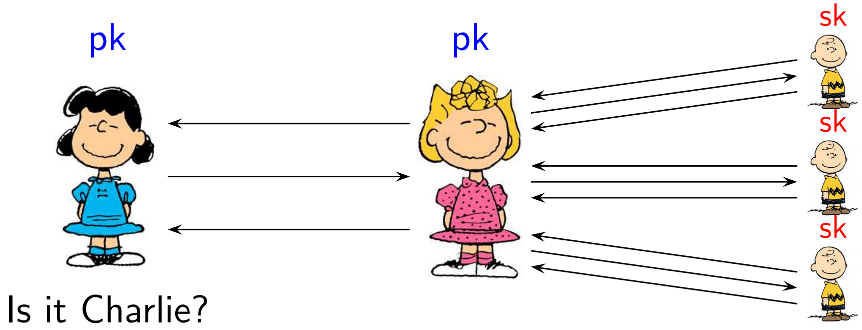

Consider a setting, where an adversary $\mathcal{A}$ can impersonate verifier $\mathcal{V}$

- ▷ The adversary $\mathcal{A}$ can execute several protocol instances with the honest prover $\mathcal{P}$ in parallel to spoof the challenge protocol.
- ▷ The adversary $\mathcal{A}$ may use protocol messages arbitrarily as long as $\mathcal{A}$ does not conduct the crossmaster attack.

Let us denote the corresponding success probability by

$$\mathsf{Adv}^{\mathsf{ea}}(\mathcal{A}) = \Pr\left[(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{Gen} : \mathcal{V}^{\mathcal{A}} = 1\right] \enspace .$$

Corresponding security guarantees

Theorem. If a cryptosystem used in the challenge-response protocol is (t, ε)-IND-CCA2 secure, then for any t-time adversary $\mathcal{A}$ the corresponding success probability $\mathsf{Adv}^{\mathsf{ea}}(\mathcal{A}) \leq \frac{1}{|\mathcal{M}|} + \varepsilon$.

Proof. A honest prover acts as a decryption oracle.

The nature of the protocol

- ▷ The protocol proves only that the prover has access to the decryption oracle and therefore the prover must possess the secret key sk.
- ▷ The possession of the secret key sk does not imply the knowledge of it. For example, the secret key sk might be hardwired into a smart card.
- ▷ Usually, the inability to decrypt is a strictly stronger security requirement than the ability to find the secret key.
- ▷ Knowledge is permanent whereas possession can be temporal.

Proofs of knowledge

Schnorr identification protocol

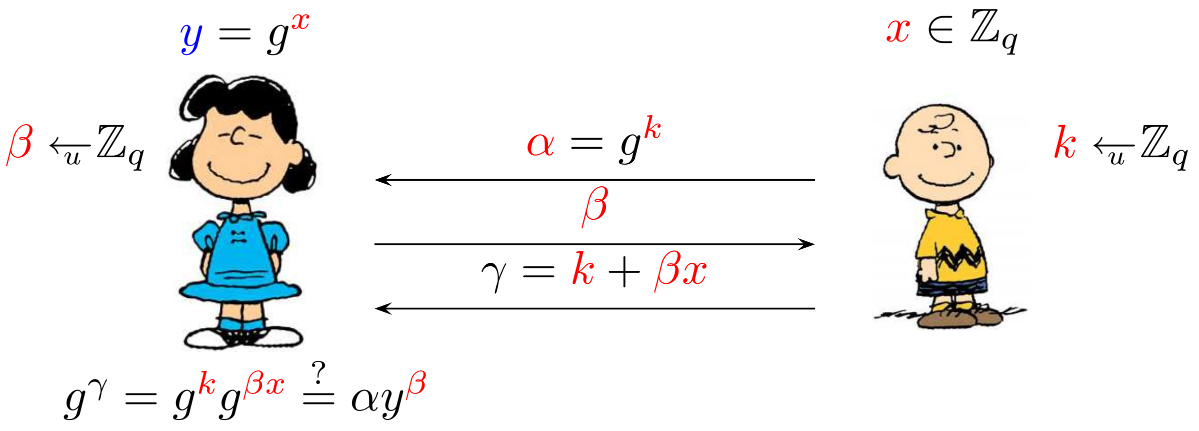

$y = g^x$

$x \in \mathbb{Z}_q$

$\beta \xleftarrow[u]{} \mathbb{Z}_q$

$k \xleftarrow[u]{} \mathbb{Z}_q$

$\xleftarrow{\quad \alpha = g^k \quad}$

$\xrightarrow{\quad \beta \quad}$

$\xleftarrow{\quad \gamma = k + \beta x \quad}$

$g^\gamma = g^k g^{\beta x} \stackrel{?}{=} \alpha y^\beta$

The group $\mathbb{G} = \langle g \rangle$ must be a DL group with a prime cardinality q.

- ▷ The secret key x is the discrete logarithm of y.
- ▷ The verifier $\mathcal{V}$ is assumed to be semi-honest.
- ▷ The prover $\mathcal{P}$ is assumed to be potentially malicious.
- ▷ We consider only security in the standalone setting.

Zero-knowledge property

Theorem. If a t-time verifier $\mathcal{V}_*$ is semi-honest in the Schnorr identification protocol, then there exists $t + \mathrm{O}(1)$-algorithm $\mathcal{V}_\circ$ that has the same output distribution as $\mathcal{V}_*$ but do not interact with the prover $\mathcal{P}$.

Proof.
Consider a code wrapper $\mathcal{S}$ that chooses $\beta \xleftarrow[u]{} \mathbb{Z}_q$ and $\gamma \xleftarrow[u]{} \mathbb{Z}_q$ and computes $\alpha \leftarrow g^\gamma \cdot y^{-\beta}$ and outputs whatever $\mathcal{V}_*$ outputs on the transcript (α, β, γ).

- ▷ If $x \neq 0$, then $\gamma = \beta + xk$ has indeed a uniform distribution.
- ▷ For fixed β and γ, there exist only a single consistent value of α.

□

Rationale: Semi-honest verifier learns nothing from the interaction with the prover. The latter is known as zero-knowledge property.

Knowledge-extraction lemma

Given two runs with a coinciding prefix α

$$\begin{tikzcd}
 & \alpha = g^{k} \arrow[dl, "\beta"'] \arrow[dr, "\beta'"] & \\
\gamma = k + \beta x & & \gamma' = k + \beta' x
\end{tikzcd}$$

We can extract the secret key $x = \frac{\gamma - \gamma'}{\beta - \beta'}$.

This property is known as special-soundness.

- ▷ If adversary $\mathcal{A}$ succeeds with probability 1, then we can extract the secret key x by rewinding $\mathcal{A}$ to get two runs with a coinciding prefix α.
- ▷ If adversary $\mathcal{A}$ succeeds with a non-zero probability ε, then we must use more advanced knowledge extraction techniques.

Find two ones in a row

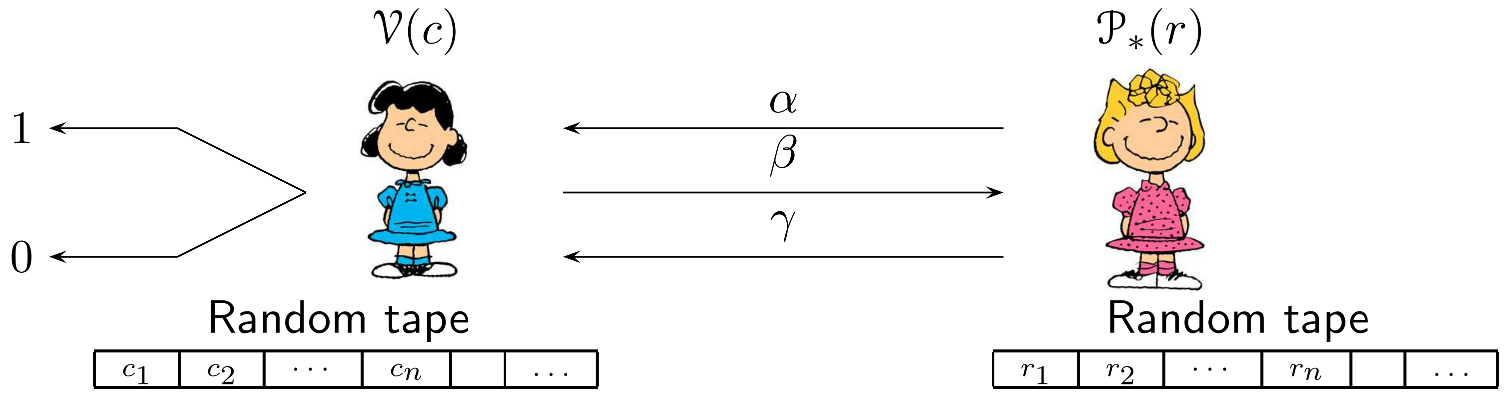

Let $A(r, c)$ be the output of the honest verifier $\mathcal{V}(c)$ that interacts with a potentially malicious prover $\mathcal{P}_*(r)$.

- ▷ Then all matrix elements in the same row $A(r, \cdot)$ lead to same α value.
- ▷ To extract the secret key sk, we must find two ones in the same row.
- ▷ We can compute the entries of the matrix on the fly.

We derive the corresponding security guarantees a bit later.

Modified Fiat-Shamir identification protocol

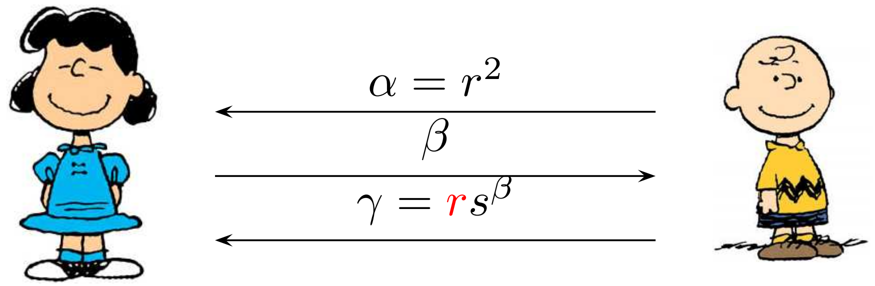

$v = s^2$

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

$s \in \mathbb{Z}_n^*$

$r \xleftarrow[u]{} \mathbb{Z}_n^*$

$\xleftarrow{\alpha = r^2}$

$\xrightarrow{\beta}$

$\xleftarrow{\gamma = rs^{\beta}}$

Halt if $\gamma \notin \mathbb{Z}_n^*$

$\gamma^2 = r^2 s^{2\beta} \overset{?}{=} \alpha v^{\beta}$

All computations are done in $\mathbb{Z}_n$, where n is an RSA modulus.

- ▷ The secret key s is a square root of v.
- ▷ The verifier $\mathcal{V}$ is assumed to be semi-honest.
- ▷ The prover $\mathcal{P}$ is assumed to be potentially malicious.
- ▷ We consider only security in the standalone setting.

Zero-knowledge property

Theorem. If a t-time verifier $\mathcal{V}_*$ is semi-honest in the modified Fiat-Shamir identification protocol, then there exists $t + \mathrm{O}(1)$-algorithm $\mathcal{V}_\circ$ that has the same output distribution as $\mathcal{V}_*$ but do not interact with the prover $\mathcal{P}$.

Proof.
Consider a code wrapper $\mathcal{S}$ that chooses $\beta \xleftarrow[u]{} \{0, 1\}$, $\gamma \xleftarrow[u]{} \mathbb{Z}_n^*$, computes $\alpha \leftarrow v^{-\beta} \cdot \gamma^2$ and outputs whatever $\mathcal{V}_*$ outputs on the transcript (α, β, γ).

▷ Since s is invertible, we can prove that $s \cdot \mathbb{Z}_n^* = \mathbb{Z}_n^*$ and $s^2 \cdot \mathbb{Z}_n^* = \mathbb{Z}_n^*$. As a result, γ is independent of β and has indeed a uniform distribution.

▷ For fixed β and γ, there exist only a single consistent value of α.

□

Knowledge-extraction lemma

Theorem. The Fiat-Shamir protocol is specially sound.

Proof. Assume that a prover $\mathcal{P}_*$ succeeds for both challenges $\beta \in \{0,1\}$:

$$\gamma_0^2 = \alpha, \quad \gamma_1^2 = \alpha v \qquad \Longrightarrow \qquad \frac{\gamma_1}{\gamma_0} = \sqrt{v}\ .$$

The corresponding extractor construction $\mathcal{K}$:

- ▷ Choose random coins r for $\mathcal{P}_*$.
- ▷ Run the protocol with $\beta = 0$ and record γ_0
- ▷ Run the protocol with $\beta = 1$ and record γ_1
- ▷ Return $\zeta = \frac{\gamma_1}{\gamma_0}$

Bound on success probability

Theorem. Let v and n be fixed. If a potentially malicious prover $\mathcal{P}_*$ succeeds in the modified Fiat-Shamir protocol with probability $\varepsilon > \frac{1}{2}$, then the knowledge extractor $\mathcal{K}^{\mathcal{P}_*}$ returns $\sqrt{v}$ with probability $2\varepsilon - 1$.

Proof. Consider the success matrix $A(r, c)$ as before. Let p_1 denote the fraction rows that contain only single one and p_2 the fraction of rows that contain two ones. Then evidently $p_1 + p_2 \leq 1$ and $\frac{p_1}{2} + p_2 \geq \varepsilon$ and thus we can establish $p_2 \geq 2\varepsilon - 1$. □

Rationale: The knowledge extraction succeeds in general only if the success probability of $\mathcal{P}_*$ is above $\frac{1}{2}$. The value $\kappa = \frac{1}{2}$ is known as knowledge error.

Matrix Games

Classical algorithm

Task: Find two ones in a same row.

Rewind:

1. Probe random entries $A(r, c)$ until $A(r, c) = 1$.
2. Store the matrix location (r, c).
3. Probe random entries $A(r, \overline{c})$ in the same row until $A(r, \overline{c}) = 1$.
4. Output the location triple $(r, c, \overline{c})$.

Rewind-Exp:

1. Repeat the procedure Rewind until $c \neq \overline{c}$.
2. Use the knowledge extraction lemma to extract sk.

Average case complexity I

Assume that the matrix contains ε-fraction of nonzero elements, i.e., $\mathcal{P}_*$ convinces $\mathcal{V}$ with probability ε. Then on average we make

$$\mathbf{E}\left[\text{probes}_1\right] = \varepsilon + 2(1-\varepsilon)\varepsilon + 3(1-\varepsilon)^2\varepsilon + \cdots = \tfrac{1}{\varepsilon}$$

matrix probes to find the first non-zero entry. Analogously, we make

$$\mathbf{E}\left[\text{probes}_2|r\right] = \tfrac{1}{\varepsilon_r}$$

probes to find the second non-zero entry. Also, note that

$$\mathbf{E}[\text{probes}_2] = \sum_r \Pr\left[r\right] \cdot \mathbf{E}[\text{probes}_2|r] = \sum_r \frac{\varepsilon_r}{\sum_{r'} \varepsilon_{r'}} \cdot \frac{1}{\varepsilon_r} = \frac{1}{\varepsilon} \enspace ,$$

where ε_r is the fraction of non-zero entries in the r^{th} row.

Average case complexity II

As a result we obtain that the Rewind algorithm does on average

$$\mathbf{E}[\text{probes}] = \tfrac{2}{\varepsilon}$$

probes. Since the Rewind algorithm fails with probability

$$\Pr\left[\text{failure}\right] = \frac{\Pr\left[\text{halting} \wedge c = \overline{c}\right]}{\Pr\left[\text{halting}\right]} \leq \frac{\kappa}{\varepsilon} \qquad \text{where} \qquad \kappa = \frac{1}{q} \enspace .$$

we make on average

$$\mathbf{E}[\text{probes}^{*}] = \frac{1}{\Pr\left[\text{success}\right]} \cdot \mathbf{E}[\text{probes}] \leq \frac{\varepsilon}{\varepsilon - \kappa} \cdot \frac{2}{\varepsilon} = \frac{2}{\varepsilon - \kappa} \enspace .$$

Strict time bounds

Markov's inequality assures that for a non-negative random variable probes

$$\Pr\left[\text{probes} \geq \alpha\right] \leq \frac{\mathbf{E}\left[\text{probes}\right]}{\alpha}$$

and thus Rewind-Exp succeeds with probability at least $\frac{1}{2}$ after $\frac{4}{\varepsilon-\kappa}$ probes.

If we repeat the experiment ℓ times, we the failure probability goes to $2^{-\ell}$.

From Soundness to Security

Soundness and subjective security

Assume that we know a constructive proof:

> If for fixed pk a potentially malicious t-time prover $\mathcal{P}_*$ succeeds with probability $\varepsilon > \kappa$, then a knowledge extractor $\mathcal{K}^{\mathcal{P}}$ that runs in time $\tau(\varepsilon) = \mathrm{O}\left(\frac{t}{\varepsilon-\kappa}\right)$ outputs sk with probability $1 - \varepsilon_2$.

and we believe:

> No human can create a $\tau(\varepsilon_1)$-time algorithm that computes sk from pk with success probability at least $1 - \varepsilon_2$.

then it is rational to assume that:

> No human without the knowledge of sk can create a algorithm $\mathcal{P}_*$ that succeeds in the proof of knowledge with probability at least ε_1.

Caveat: For each fixed pk, there exists a trivial algorithm that prints out sk. Hence, we cannot get objective security guarantees.

Soundness and objective security

Assume that we know a constructive proof:

> If for a fixed pk a potentially malicious t-time prover $\mathcal{P}_*$ succeeds with probability $\varepsilon > \kappa$, then a knowledge extractor $\mathcal{K}^{\mathcal{P}}$ that runs in time $\tau(\varepsilon) = \mathrm{O}\bigl(\frac{t}{\varepsilon-\kappa}\bigr)$ outputs sk with probability $1 - \varepsilon_2$.

and know a mathematical fact that any $\tau(2\varepsilon_1)$-time algorithm $\mathcal{A}$

$$\Pr\left[(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{Gen} : \mathcal{A}(\mathsf{pk}) = \mathsf{sk}\right] \leq \varepsilon_1(1-\varepsilon_2)$$

then we can prove an average-case security guarantee:

> For any t-time prover $\mathcal{P}_*$ that does not know the secret key

$$\mathsf{Adv}^{\mathsf{ea}}(\mathcal{A}) = \Pr\left[(\mathsf{pk},\mathsf{sk}) \leftarrow \mathsf{Gen} : \mathcal{V}^{\mathcal{P}_*(\mathsf{pk})} = 1\right] \leq 2\varepsilon_1 \ .$$

Objective security guarantees

Schnorr identification scheme

If $\mathbb{G}$ is a DL group, then the Schnorr identification scheme is secure, where the success probability is averaged over all possible runs of the setup Gen.

Fiat-Shamir identification scheme

Assume that modulus n is chosen form a distribution $\mathcal{N}$ of RSA moduli such that on average factoring is hard over $\mathcal{N}$. Then the Fiat-Shamir identification scheme is secure, where the success probability is averaged over all possible runs of the setup Gen and over all choices of modulus n.

Composability of Σ-protocols

A formal definition of sigma protocol

A sigma protocol for an efficiently computable relation $R \subseteq \{0,1\}^* \times \{0,1\}^*$ is a three move protocol that satisfies the following properties.

- **Σ-structure.** A prover first sends a commitment, next a verifier sends varying challenge, and then the prover must give a consistent response.
- **Functionality.** The protocol run between an honest prover $\mathcal{P}(\mathsf{sk})$ and verifier $\mathcal{V}(\mathsf{pk})$ is always accepting if $(\mathsf{sk}, \mathsf{pk}) \in R$.
- **Perfect simulatability.** There exists an efficient non-rewinding simulator $\mathcal{S}$ such that the output distribution of a semi-honest verifier $\mathcal{V}_*$ in the real world and the output distribution of $\mathcal{S}^{\mathcal{V}_*}$ in the ideal world coincide.
- **Special soundness.** There exists an efficient extraction algorithm Extr that, given two accepting protocol runs $(\alpha, \beta_0, \gamma_0)$ and $(\alpha, \beta_1, \gamma_1)$ with $\beta_0 \neq \beta_1$ that correspond to pk, outputs sk_* such that $(\mathsf{sk}_*, \mathsf{pk}) \in R$

AND-composition

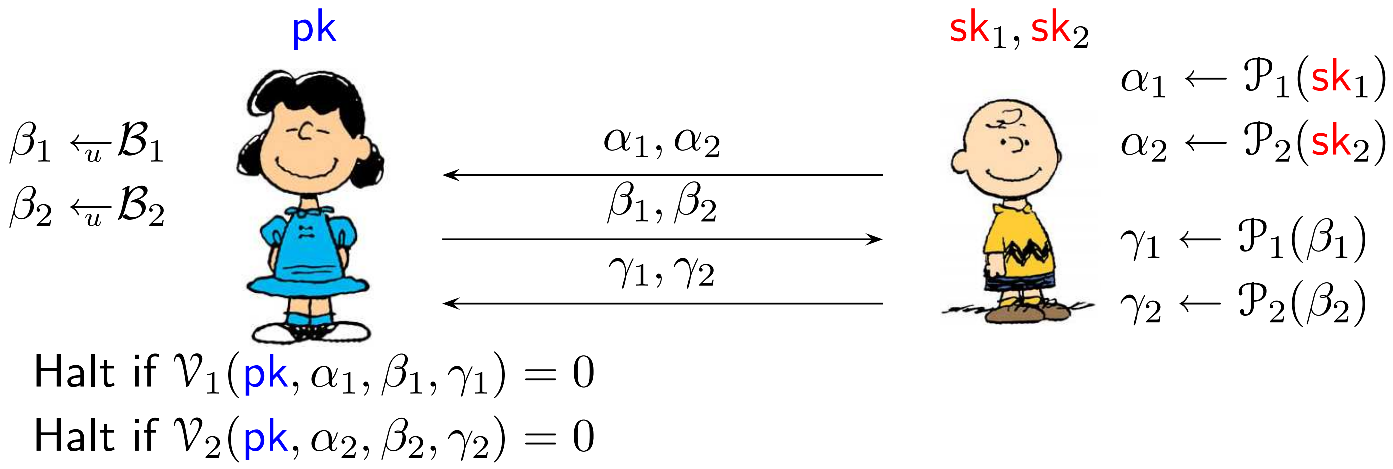

pk

$\beta_1 \xleftarrow[u]{} \mathcal{B}_1$

$\beta_2 \xleftarrow[u]{} \mathcal{B}_2$

α_1, α_2

β_1, β_2

γ_1, γ_2

$\mathsf{sk}_1, \mathsf{sk}_2$

$\alpha_1 \leftarrow \mathcal{P}_1(\mathsf{sk}_1)$

$\alpha_2 \leftarrow \mathcal{P}_2(\mathsf{sk}_2)$

$\gamma_1 \leftarrow \mathcal{P}_1(\beta_1)$

$\gamma_2 \leftarrow \mathcal{P}_2(\beta_2)$

Halt if $\mathcal{V}_1(\mathsf{pk}, \alpha_1, \beta_1, \gamma_1) = 0$

Halt if $\mathcal{V}_2(\mathsf{pk}, \alpha_2, \beta_2, \gamma_2) = 0$

If we run two sigma protocols for different relations R_1 and R_2 in parallel, we get a sigma protocol* for new relation $R_1 \wedge R_2$

$$(\mathsf{sk}_1, \mathsf{sk}_2, \mathsf{pk}) \in R_1 \wedge R_2 \quad \Leftrightarrow \quad (\mathsf{sk}_1, \mathsf{pk}) \in R_1 \wedge (\mathsf{sk}_1, \mathsf{pk}) \in R_2 \ .$$

* Modulo some minor details discussed in the next slide.

The corresponding proof

Perfect simulatability. Let $\mathcal{S}_1$ and $\mathcal{S}_2$ be canonical simulators for $\mathcal{V}_1$ and $\mathcal{V}_2$. Then $\mathcal{S}_1$ outputs a properly distributed triple $(\alpha_1, \beta_1, \gamma_1)$ and $\mathcal{S}_2$ outputs a properly distributed triple $(\alpha_2, \beta_2, \gamma_2)$. Hence, we can run $\mathcal{S}_1$ and $\mathcal{S}_2$ in parallel to create a properly distributed transcript $(\alpha_1, \alpha_2, \beta_1, \beta_2, \gamma_1, \gamma_2)$.

Special soundness*. Given two accepting transcripts

$$(\alpha_1, \alpha_2, \beta_1^0, \beta_2^0, \gamma_1^0, \gamma_2^0), (\alpha_1, \alpha_2, \beta_1^1, \beta_2^1, \gamma_1^1, \gamma_2^1), \quad \text{with} \quad \beta_1^0 \neq \beta_1^1, \beta_2^0 \neq \beta_2^1 \ ,$$

we can decompose them into original colliding transcripts

$$\begin{aligned} &(\alpha_1, \beta_1^0, \gamma_1^0), (\alpha_1, \beta_1^1, \gamma_1^1), \qquad && \beta_1^0 \neq \beta_1^1 \ , \\ &(\alpha_2, \beta_2^0, \gamma_2^0), (\alpha_2, \beta_2^1, \gamma_2^1), \qquad && \beta_2^0 \neq \beta_2^1 \ . \end{aligned}$$

OR-composition

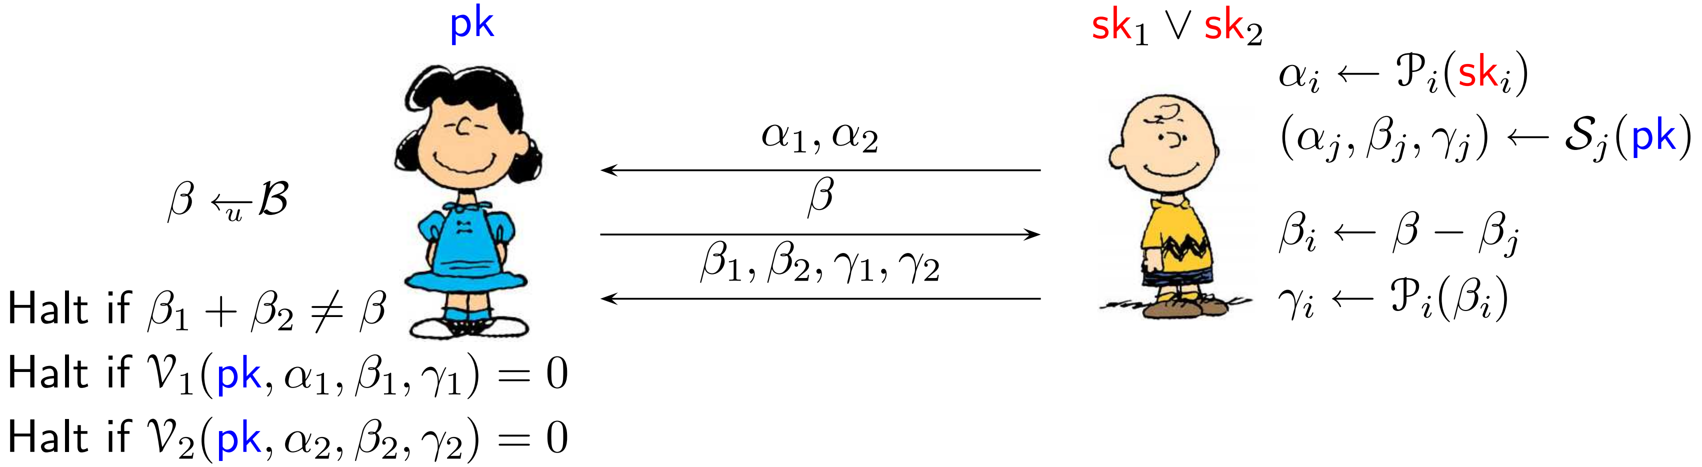

pk

$\mathsf{sk}_1 \vee \mathsf{sk}_2$

$\alpha_i \leftarrow \mathcal{P}_i(\mathsf{sk}_i)$

$(\alpha_j, \beta_j, \gamma_j) \leftarrow \mathcal{S}_j(\mathsf{pk})$

$\xleftarrow{\quad\alpha_1, \alpha_2\quad}$

$\beta \xleftarrow[u]{} \mathcal{B}$

$\xrightarrow{\quad\beta\quad}$

$\beta_i \leftarrow \beta - \beta_j$

$\gamma_i \leftarrow \mathcal{P}_i(\beta_i)$

$\xleftarrow{\quad\beta_1, \beta_2, \gamma_1, \gamma_2\quad}$

Halt if $\beta_1 + \beta_2 \neq \beta$

Halt if $\mathcal{V}_1(\mathsf{pk}, \alpha_1, \beta_1, \gamma_1) = 0$

Halt if $\mathcal{V}_2(\mathsf{pk}, \alpha_2, \beta_2, \gamma_2) = 0$

Assume that we have two sigma protocols for relations R_1 and R_2 such that the challenge is chosen uniformly from a commutative group $(\mathcal{B}; +)$.

Then a prover can use a simulator $\mathcal{S}_j$ to create the transcript for missing secret sk_j and then create response using the known secret sk_i.

OR-composition

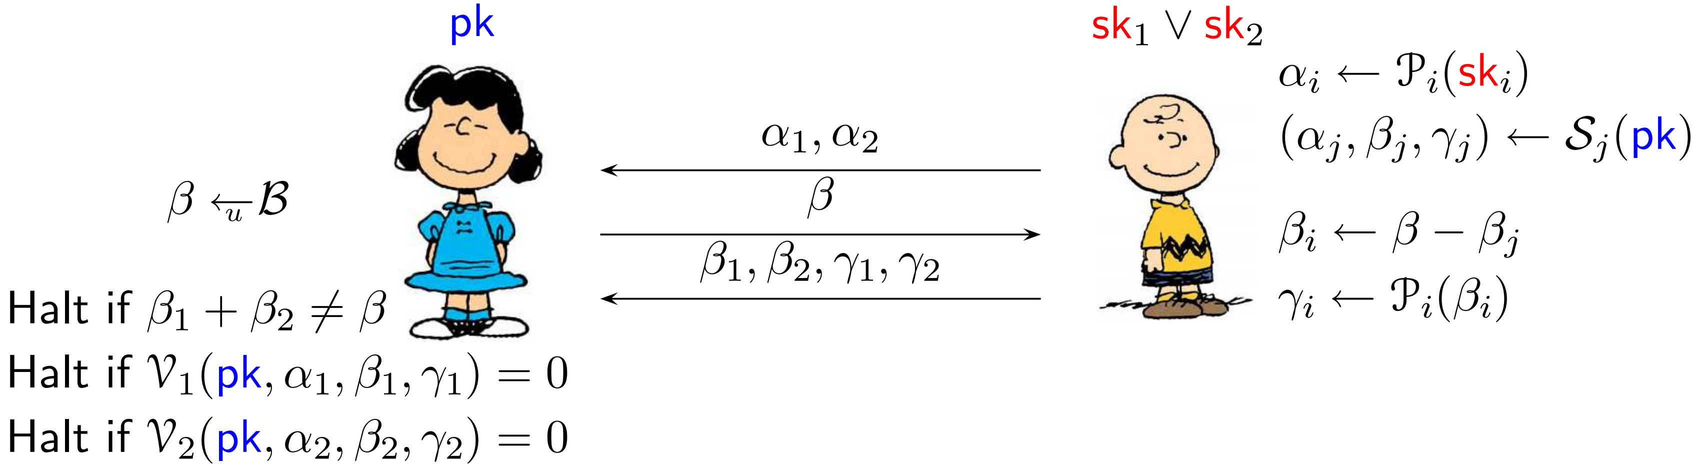

pk

$\mathsf{sk}_1 \vee \mathsf{sk}_2$

$\alpha_i \leftarrow \mathcal{P}_i(\mathsf{sk}_i)$

$(\alpha_j, \beta_j, \gamma_j) \leftarrow \mathcal{S}_j(\mathsf{pk})$

α_1, α_2

$\beta \xleftarrow[u]{} \mathcal{B}$

β

$\beta_i \leftarrow \beta - \beta_j$

$\beta_1, \beta_2, \gamma_1, \gamma_2$

$\gamma_i \leftarrow \mathcal{P}_i(\beta_i)$

Halt if $\beta_1 + \beta_2 \neq \beta$

Halt if $\mathcal{V}_1(\mathsf{pk}, \alpha_1, \beta_1, \gamma_1) = 0$

Halt if $\mathcal{V}_2(\mathsf{pk}, \alpha_2, \beta_2, \gamma_2) = 0$

As a result, we get a sigma protocol for new relation $R_1 \vee R_2$

$$(\mathsf{sk}_1, \mathsf{sk}_2, \mathsf{pk}) \in R_1 \vee R_2 \quad \Leftrightarrow \quad (\mathsf{sk}_1, \mathsf{pk}) \in R_1 \vee (\mathsf{sk}_1, \mathsf{pk}) \in R_2 \ .$$

The corresponding proof

Perfect simulatability. Note that β_1 and β_2 are independent and have a uniform distribution over $\mathcal{B}$. Consequently, we can run the canonical simulators $\mathcal{S}_1$ and $\mathcal{S}_2$ be for $\mathcal{V}_1$ and $\mathcal{V}_2$ in parallel to create the properly distributed transcript $(\alpha_1, \alpha_2, \beta_1 + \beta_2, \beta_1, \beta_2, \gamma_1, \gamma_2)$.

Special soundness. Given two transcripts

$$(\alpha_1, \alpha_2, \beta_1^0 + \beta_2^0, \beta_1^0, \beta_2^0, \gamma_1^0, \gamma_2^0), (\alpha_1, \alpha_2, \beta_1^1 + \beta_2^1, \beta_1^1, \beta_2^1, \gamma_1^1, \gamma_2^1)$$

such that $\beta_1^0 + \beta_2^0 \neq \beta_1^1 + \beta_2^1$, we can extract a colliding sub-transcript

$$\begin{cases} (\alpha_1, \beta_1^0, \gamma_1^0), (\alpha_1, \beta_1^1, \gamma_1^1), & \text{if } \beta_1^0 \neq \beta_1^1 \ , \\ (\alpha_2, \beta_2^0, \gamma_2^0), (\alpha_2, \beta_2^1, \gamma_2^1), & \text{if } \beta_2^0 \neq \beta_2^1 \ . \end{cases}$$

Monotone access structures

Let a binary properties $\pi_1, \ldots, \pi_n$ denote possible roles of participants and let $\mathsf{sk}_1, \ldots, \mathsf{sk}_n$ denote the corresponding secrets that the participant knows if the corresponding property π_i is set.

Now assume that $\psi : \{0,1\}^n \to \{0,1\}$ is a monotone predicate that maps the property vector $(\pi_1, \ldots, \pi_n)$ to a final access verdict for some object. Then there exists a sigma protocol for the corresponding relation.

As a result, we can construct identification protocols that are sound and secure and leak only the value $\psi(\pi_1, \ldots, \pi_n)$.

- ▷ Anonymous group authentication
- ▷ Anonymous verification of credentials