MANA IV Proof of Security

Jukka Valkonen

Laboratory for Theoretical Computer Science
Helsinki University of Technology

4.12.2007

Agenda

- Introduction
- Cryptographic preliminaries
- The protocol
- Security Analysis

Introduction

- Setting up a shared key in ad-hoc network
    - No key hierarchy
    - No pre-shared secrets
    - Ordinary users without any knowledge of security protocols
- Mana IV can be used to authenticate the negotiated key

Communication Model

- Out-of-Band channels
  - Authentic, some times secret
  - Adversary can read, delay and reorder messages
  - Low bandwidth
- In-band channels
  - Routed via malicious adversary
  - Adversary can read, insert, delete and modifuy messages
  - Dolev-Yao -adversary

Key Establishment Protocols for First Connect

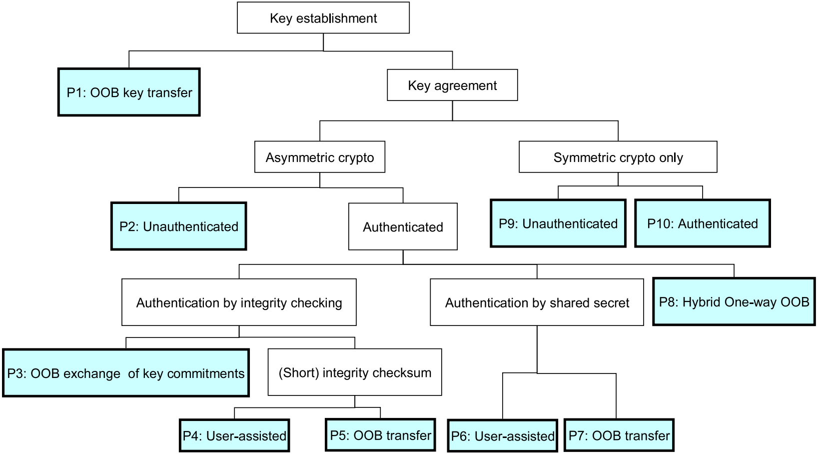

Keyed hash functions

- A hash function is ϵ_u-*almost universal* if given two inputs $x_0 \neq x_1$:

$$Pr[k \leftarrow \mathcal{K} : h(x_0, k) = h(x_1, k)] \leq \epsilon_u$$

- A hash function is ϵ_u-*almost XOR universal* if for any $x_0 \neq x_1$ and y

$$Pr[k \leftarrow \mathcal{K} : h(x_0, k) \oplus h(x_1, k) = y] \leq \epsilon_u$$

Keyed hash functions

- Special notion needed when key is divided into two sub-keys: $h : \mathcal{M} \times \mathcal{K}_a \times \mathcal{K}_b \to \mathcal{T}$
- A hash function is (ϵ_a, ϵ_b)-*almost regular* w.r.t. the sub-keys if for each data $x \in \mathcal{M}$, tag y and sub-keys $\widehat{k}_a \in \mathcal{K}$, $\widehat{k}_b \in \mathcal{K}$:

$$Pr[k_a \leftarrow \mathcal{K}_a : h(x, k_a, \widehat{k}_b) = y] \leq \epsilon_a$$

and

$$Pr[k_b \leftarrow \mathcal{K}_b : h(x, \widehat{k}_a, k_b) = y] \leq \epsilon_b$$

Keyed hash functions

- A hash function is ϵ_u-*almost universal w.r.t. the sub-key* k_a if for any two data $x_0 \neq x_1$ and $k_b, \widehat{k}_b \in \mathcal{K}_b$:

$$Pr[k_a \leftarrow \mathcal{K} : h(x_0, k_a, k_b) = h(x_1, k_a, \widehat{k}_b)] \leq \epsilon_u$$

- A hash function is *strongly* ϵ_u-*almost universal w.r.t. the sub-key* k_a if for any $(x_0, k_b) \neq (x_1, \widehat{k}_b)$ we have

$$Pr[k_a \leftarrow \mathcal{K} : h(x_0, k_a, k_b) = h(x_1, k_a, \widehat{k}_b)] \leq \epsilon_u$$

- Here $\epsilon_u, \epsilon_a, \epsilon_b \geq \frac{1}{|\mathcal{T}|}$
- If the equality holds, the word *almost* is skipped

Commitment Schemes

- Commitment scheme *Com* is specified by three algorithms:
  - `Gen` generates the public parameters `pk`
  - `Com` takes `pk` and message and transforms them into a commit value c and a decommit value d:

$$\mathcal{M} \times \mathcal{R} \rightarrow \mathcal{C} \times \mathcal{D}$$

  - `Open` opens the commitment: $\mathtt{Open}(c, d) = m$ for all $(c, d) = Com(m, r)$
- Incorrect decommit value yields to special abort value $\perp$

Commitment schemes

- A commitment scheme is (t, ϵ_1)-hiding if any t-time adversary A achieves advantage

$$\mathrm{Adv}^{\mathsf{hid}}_{\mathcal{C}om}(A) = 2 \cdot \left| \Pr\left[\begin{array}{l} \mathsf{pk} \leftarrow \mathsf{Gen}, s \leftarrow \{0,1\}, (x_0, x_1, \sigma) \leftarrow A(\mathsf{pk}) \\ (c_s, d_s) \leftarrow \mathsf{Com}_{\mathsf{pk}}(x_s) : \ A(\sigma, c_s) = s \end{array} \right] - \frac{1}{2} \right| \leq \varepsilon_1$$

- A commitment scheme is (t, ϵ_2)-binding if any t-time adversary A achieves advantage

$$\mathrm{Adv}^{\mathsf{bind}}_{\mathcal{C}om}(A) = \Pr\left[\begin{array}{l} \mathsf{pk} \leftarrow \mathsf{Gen}, (c, d_0, d_1) \leftarrow A(\mathsf{pk}) : \\ \perp \neq \mathsf{Open}_{\mathsf{pk}}(c, d_0) \neq \mathsf{Open}_{\mathsf{pk}}(c, d_1) \neq \perp \end{array} \right] \leq \varepsilon_2 \ .$$

Non-malleable commitment schemes

"Intuitively, a commitment scheme is non-malleable, if given a valid commitment c, it is infeasible to generate related commitments $c_1, \ldots, c_n$ that can be successfully opened after seeing a decommitment value d."
An adversary is a quadruple $A = (A_1, A_2, A_3, A_4)$ of algorithms, where $A_{1\ldots3}$ are active and A_4 is a distinguisher

1. The challenger draws two independent samples $x_0 \leftarrow \mathrm{MGen}, x_1 \leftarrow \mathrm{MGen}$ and computes a challenge commitment $(c, d) \leftarrow \mathcal{C}om_{\mathrm{pk}}(x_0)$
2. Challenger sends c to A_2 that computes a commitment vector $c_1, \ldots, c_n$. If some $c_i = c$ then Challenger stops A with $\perp$

Non-malleable commitment schemes

3. Challenger sends d to A_3 that must produce a *valid* decommitment vector $d_1, \ldots, d_n$ ($y_i = \mathrm{Open}_{\mathrm{pk}}(c_i, d_i)$). If some $y_i = \perp$ A is stopped with $\perp$.
4. In World_0 Challenger invokes $A_4(x_0, y_1, \ldots, y_n)$ with correct x_0 and in World_0 $A_4(x_1, y_1, \ldots, y_n)$

A commitment scheme is (t, ϵ)-non-malleable iff for any t-time adversary A the advantage of distinguishing the two worlds is

$$\mathrm{Adv}^{\mathrm{nm}}_{\mathcal{C}om}(A) = |Pr[A_4 = 0|\mathrm{World}_0] - Pr[A_4 = 0|\mathrm{World}_1]|$$

MANA IV

1. Alice computes $(c, d) \leftarrow \mathrm{Com}_{\mathrm{pk}}(k_a)$ for random $k_a \leftarrow \mathcal{K}$ and sends (m_a, c) to Bob
2. Bob chooses random $k_b \leftarrow \mathcal{K}$ and sends (m_b, k_b) to Alice
3. Alice sends d to Bob, who computes $k_a \leftarrow \mathrm{Open}_{\mathrm{pk}}(c, d)$ and halts if $k_a = \perp$. Both parties compute a test value $\mathrm{oob} = h(m_a \| m_b, k_a, k_b)$ from the received messages
4. Both parties accept (m_a, m_b) iff the local l-bit test values oob_a and oob_b coincide

h is a keyed hash function with sub-keys k_a, k_b where $\mathcal{K}_a$ is a message space of commitment scheme

Idea of the security proof

The idea is to go through all the strategies an adversary can use to attack the protocol run. These include

- Adversary attacks h by altering m_a, m_b, k_b and possible d
- Attacks based on abnormal execution paths

The attacker succeeds if Alice and Bob accept but $(m_a, \widehat{m_b}) \neq (\widehat{m_a}, m_b)$

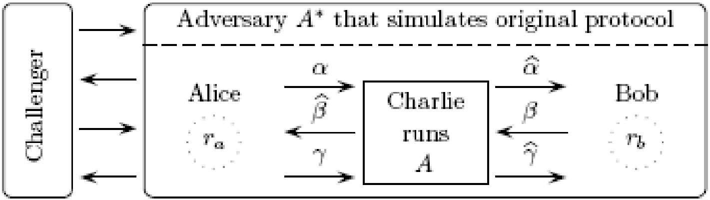

Fig. 4. Generic reduction scheme

Theorem 1: Statistically binding commitments

For any t, there exists $\tau = t + \mathcal{O}(1)$ such that if Com is (τ, ϵ_1)-hiding, ϵ_2-binding and (τ, ϵ_3)-non-malleable and h is (ϵ_a, ϵ_b)-almost regular and ϵ_u almost universal w.r.t. the sub-key k_a then the MANA IV protocol is $(2\epsilon_1 + 2\epsilon_2 + \epsilon_3 + max\{\epsilon_a, \epsilon_b, \epsilon_u\}, t)$-secure.

Theorem 2: Computationally binding commitments

For any t, there exists $\tau = 2t + \mathcal{O}(1)$ such that if Com is (τ, ϵ_1)-hiding, (τ, ϵ_2)-binding and (τ, ϵ_3)-non-malleable and h is (ϵ_a, ϵ_b)-almost regular and ϵ_u almost universal w.r.t. the sub-key k_a then the MANA IV protocol is $(2\epsilon_1 + \epsilon_2 + \sqrt{\epsilon_2} + \epsilon_3 + max\{\epsilon_a, \epsilon_b, \epsilon_u\}, t)$-secure.

Lemma 1

For any t, there exists $\tau = t + \mathcal{O}(1)$ such that if Com is τ, ϵ_1-hiding and (τ, ϵ_2)-binding and h is ϵ_u-almost universal w.r.t. the sub-key k_a, then for any t-time adversary A and input data (m_a, m_b)

$$Pr[d\text{-}forge \wedge norm \wedge c = \widehat{c}] \leq \epsilon_u \cdot Pr[norm \wedge c = \widehat{c}] + \epsilon_1 + \epsilon_2$$

Proof

Assume a t-time algorithm A which violates the previous probability

Let's construct A^* that wins the hiding game, i.e. given pk outputs (x_0, x_1, σ) and afterwards after given a commitment c_s for $s \leftarrow \{0, 1\}$ is able to correctly guess the bit s

1. Given pk, chooses $k_a, k_a^* \leftarrow \mathcal{K}_a$ as (x_0, x_1) and sends $(k_a, k_a^*, \mathrm{pk})$ to Challenger
2. When Challenger replies c_s for $(c_s, d_s) = \mathrm{Com}_{\mathrm{pk}}(x_s)$, A^* simulates a faithful execution of Mana IV with $\alpha = (m_a, c_s)$ until A queries γ. A^* stops the simulation and halts with $\perp$ if there is a protocol failure $\neg$norm or $c \neq \widehat{c}$
3. If $h(m_a \| \widehat{m}_b, k_a, \widehat{k}_b) = h(\widehat{m}_a \| m_b, k_a, k_b)$ and $(m_a, \widehat{m}_b) \neq (\widehat{m_a}, m_b)$ outputs guess $s = 0$, else $s = 1$

Proof continued

For $s = 0$ we get

$$Pr[A^* = 0|s = 0] \geq Pr[\text{d-forge} \wedge \text{norm} \wedge c = \widehat{c} \wedge k_a = \widehat{k}_a]$$

For $s = 1$,

$$Pr[A^* = 0|s = 1] \leq \epsilon_u \cdot Pr[\text{norm} \wedge c = \widehat{c}]$$

as $Pr[A^* \neq \perp |s = 1] = Pr[\text{norm} \wedge c = \widehat{c}]$ (perfect simulation until A queries γ) and c_1 and k_a are statistically independent ($Pr[A^* = 0|s = 1, A^* \neq \perp] \leq \epsilon_u$)

Proof continued

We get

$$\mathrm{Adv}^{\mathrm{hid}}(A^*) = |Pr[A^* = 0|s = 0] - Pr[A^* = 0|s = 1]| \geq$$

$$|Pr[\text{d-forge} \wedge \text{norm} \wedge c = \widehat{c} \wedge k_a = \widehat{k_a}] - \epsilon_u \cdot Pr[\text{norm} \wedge c = \widehat{c}]| > \epsilon_1$$

which contradicts the (τ, ϵ_1)-hiding property. Here $Pr[\text{d-forge} \wedge \text{norm} \wedge c = \widehat{c} \wedge k_a = \widehat{k_a}] \geq Pr[\text{d-forge} \wedge \text{norm} \wedge c = \widehat{c}] - \epsilon_2$ and the assumption that A violates the inequality

Lemma 2

For any t, there exists $\tau = t + \mathcal{O}(1)$ such that if Com is (τ, ϵ_3)-non-malleable and h is (ϵ_a, ϵ_b)-almost regular, then for any t-time adversary A and inputs (m_a, m_b)

$$Pr[d\text{-}forge \wedge norm \wedge c \neq \widehat{c}] \leq \epsilon_a \cdot Pr[norm \wedge c \neq \widehat{c}] + \epsilon_3$$

Proof

Now, A is a t-time algorithm that violates the inequality. Idea is to build an adversary $A^* = (A_1^*, A_2^*, A_3^*, A_4^*)$ that can break the non-malleability of the commitment scheme.

1. Given pk, A_1^* outputs a sampler over $\mathcal{K}_a$ and state $\sigma_1 = (\mathrm{pk}, m_a, m_b)$. Challenger computes $x_0, x_1 \leftarrow \mathcal{K}_a$ and $(c, d) \leftarrow \mathrm{Com}_{\mathrm{pk}}(x_0)$
2. Given c, σ_1, A_2^* simulates the protcol with $k_b \leftarrow \mathcal{K}_b$ and stops before A demands γ. A^* stops and halts with $\perp$ if there is a protocol failure $\neg$norm or $c = \widehat{c}$. Otherwise A_2^* outputs a commitment $\widehat{c}$ and σ_2 containing enough information to resume the simulation.

Proof continued

3 Given d, σ_2, A_3^* resumes the simulation and outputs $\widehat{d}$

4 If A_3^* was successful in opening $\widehat{c}$ then $A^*(x_s, u, \sigma_2)$ sets $k_a \leftarrow x_s$ and $\widehat{k}_a \leftarrow y$ and computes $\mathrm{oob}_a = h(m_a \| \widehat{m}_b, k_a, \widehat{k}_b)$ and $\mathrm{oob}_b = h(\widehat{m}_a \| m_b, \widehat{k}_a, k_b)$. A_4^* outputs a guess $s = 0$ if $\mathrm{oob}_a = \mathrm{oob}_b$ but $(m_a, \widehat{m}_b) \neq (\widehat{m}_a, m_b)$, else $s = 1$.

Proof continued

Now, in $\texttt{World}_0$, Step 1 provides perfect simulation and in $\texttt{World}_1$ k_a is independent of all variables computed by A. Thus

$$Pr[A_4^* = 0|\texttt{World}_0] = Pr[\text{d-forge} \land \text{norm} \land c \neq \widehat{c}]$$

and

$$Pr[A_4^* = 0|\texttt{World}_1] = \epsilon_a \cdot Pr[\text{norm} \land c \neq \widehat{c}]$$

as h is (ϵ_a, ϵ_b)-almost regular.

This results as a contradiction as

$$\text{Adv}^{\text{nm}}(A^*) = |Pr[A^* = 0|\texttt{World}_0] - Pr[A^* = 0|\texttt{World}_1| > \epsilon_3$$

Lemma 3

For any t, there exists $\tau = t + \mathcal{O}(1)$ *such that if Com is* (τ, ϵ_1)*-hiding, h is* (ϵ_a, ϵ_b)*-almost regular. Then for any t-time adversary A and input* (m_a, m_b)

$$Pr[\textit{d-forge} \wedge \widehat{\gamma} \prec \widehat{\beta}] \leq \epsilon_1 + \epsilon_a \cdot Pr[\widehat{\gamma} \prec \widehat{\beta}]$$

Proof

Again, let A be a t-time adversary that violates the previous inequlity. If $\widehat{\gamma} \prec \widehat{\beta}$, Bob'c control value oob_b is fixed before A receives γ. Now we have A^* that plays hiding game

1. Given pk, chooses $k_a, k_a^* \leftarrow \mathcal{K}_a$ as (x_0, x_1) and sends k_a, k_a^*, pk to Challenger
2. When Challenger replies c_s for $(c_s, d_s) = \text{Com}_{\text{pk}}(x_s)$, A^* simulates an execution of Mana IV with $\alpha = (m_a, c_s)$ until A outputs $\widehat{\beta}$. A^* stops the simulation and halts with $\perp$ if there is a protocol failure: $\widehat{\beta} \prec \widehat{\gamma}$ or $\text{Open}_{\text{pk}} = \perp$.
3. A^* computes $\widehat{k_a} = \text{Open}_{\text{pk}}(\hat{c}, \hat{d})$,

   $\text{ooba} = h(m_a \| \widehat{m}_b, k_a, \widehat{k}_b)$ and $\text{oob}_b = h(\widehat{m}_a \| m_b, \widehat{k}_a, k_b)$. If $\text{ooba} = \text{oob}_\text{b}$ and $(m_a, \widehat{m_b}) \neq (\widehat{m}_a, m_b)$ outputs 0 else 1

Proof continued

If $s = 0$ then $Pr[A^* = 0|s = 0] = Pr[\text{d-forge} \wedge \widehat{\gamma} \prec \widehat{\beta}]$.
If $s = 1$ then $Pr[A^* = 0|s = 1] = \epsilon_a \cdot Pr[\widehat{\gamma} \prec \widehat{\beta}]$ as

$Pr[A^* \neq \perp |s = 1] = Pr[\hat{\gamma} \prec \hat{\beta}]$ and
$Pr[A^* = 0|s = 0, A^* \neq \perp] \leq \epsilon_a$ because of (ϵ_a, ϵ_b)-almost regularity

The advantage is

$$\mathsf{Adv}^{\text{hid}}(A^*) = |Pr[A^* = 0|s = 0] - Pr[A^* = 0|s = 1]| > \epsilon_1$$

which results in a contradiction

Lemma 4

If Com is statistically ϵ_2-binding and h is (ϵ_a, ϵ_b)-almost regular, then for each adversary A and input (m_a, m_b)

$$Pr[d\text{-}forge \wedge \gamma \prec \beta] \leq \epsilon_2 + \epsilon_b \cdot Pr[\gamma \prec \beta]$$

Proof

For each $\widehat{c}$ fix a canonical $\widehat{k}_a$ such that $\widehat{k}_a = \mathrm{Open}_{\mathsf{pk}}(\widehat{c}, \widehat{d}_0)$ for some $\widehat{d}_0$. If $\gamma \prec \beta$ the oob_a is fixed before k_b. Now the probability that different k_b values lead to different openings $k'_a \neq \widehat{k}_a$ is at most ϵ_2. Otherwise, one can find valid double openings $\mathrm{Open}_{\mathsf{pk}}(\widehat{c}, \widehat{d}_0) \neq \mathrm{Open}_{\mathsf{pk}}(\widehat{c}, \widehat{d}_1)$ just by enumerating all possible protocol runs. Now $Pr[k_b \leftarrow \mathcal{K} : \mathrm{oob}_a = h(\widehat{m_a} \| m_b, \widehat{k}_a, k_b)] \leq \epsilon_b$, as k_b is independent from $\widehat{k}_a$ and oob_a and thus claim follows.

Lemma 5

For any t there exists $\tau = t + \mathcal{O}(1)$ such that if Com is (τ, ϵ_2)-binding and h is (ϵ_a, ϵ_b)-almost regular, then for any t-time adversary A and inputs m_a, m_b

$$Pr[\textit{d-forge} \wedge \gamma \prec \beta] \leq \epsilon_b \cdot Pr[\gamma \prec \beta] + \sqrt{\epsilon_2}$$

Proof omitted

Thus

by summing up the probabilities the proof is complete