

# Topics in Concurrency

## Lecture 10

Jonathan Hayman

6 March 2015

# Public-key cryptography

Public key cryptography:

- for each entity/participant/agent  $A$ , there is a key  $Pub(A)$  and a key  $Priv(A)$ .
- $Pub(A)$  is intended to be known by everybody: it is **public**
- $Priv(A)$  is intended to be known only by  $A$ : it is **private**
- Any agent can encrypt using a key that it knows
- To decrypt a message encrypted under  $Pub(A)$  it is necessary to know  $Priv(A)$
- To decrypt a message encrypted under  $Priv(A)$  it is necessary to know  $Pub(A)$

Will also allow symmetric keys e.g.  $Key(A, B)$ .

# The Needham-Schröder-Lowe Protocol

The goal of the NSL protocol: two agents use public-key cryptography to ensure

- **authentication**: For A as the initiator: upon completion of the protocol, A can be sure that B generated the messages that A received following the protocol in response to A's request
- **shared secret**: if two entities complete the protocol with each other, at the end they both know a value not known to any potential attacker (e.g. to be used in more efficient symmetric-key cryptographic operations)

Formally, the correctness properties are subtle (e.g. what if B chose to release its private key?)

# The original protocol

- (1)  $A \rightarrow B: \{m, A\}_{Pub(B)}$
- (2)  $B \rightarrow A: \{m, n\}_{Pub(A)}$
- (3)  $A \rightarrow B: \{n\}_{Pub(B)}$

$m$  and  $n$  are nonces: fresh randomly-generated (very) long integers

# The original protocol

- (1)  $A \longrightarrow B: \{m, A\}_{Pub(B)}$
- (2)  $B \longrightarrow A: \{m, n\}_{Pub(A)}$
- (3)  $A \longrightarrow B: \{n\}_{Pub(B)}$

$m$  and  $n$  are nonces: fresh randomly-generated (very) long integers

Original protocol introduced by Needham and Schröder in 1978  
contained a flaw revealed (and fixed) by Lowe in 1995 [using CSP].

# An attack against the original protocol

*Man-in-the-middle attacker E convinces A to start communication with E and uses the messages generated by A to follow the protocol with B, posing as A.*

A

E

B

$A \rightarrow B : \{m, A\}_{Pub(B)}$

$B \rightarrow A : \{m, n\}_{Pub(A)}$

$A \rightarrow B : \{n\}_{Pub(B)}$

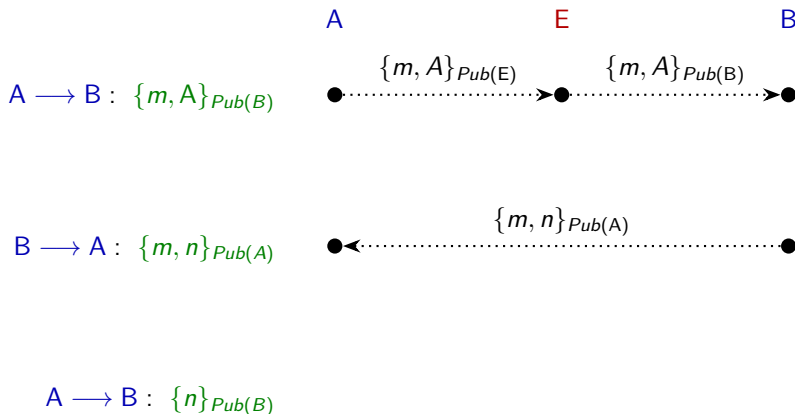






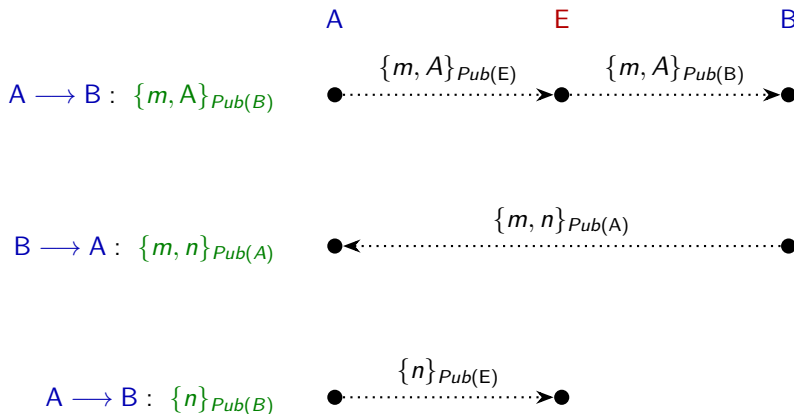
# An attack against the original protocol

Man-in-the-middle attacker **E** convinces **A** to start communication with **E** and uses the messages generated by **A** to follow the protocol with **B**, posing as **A**.



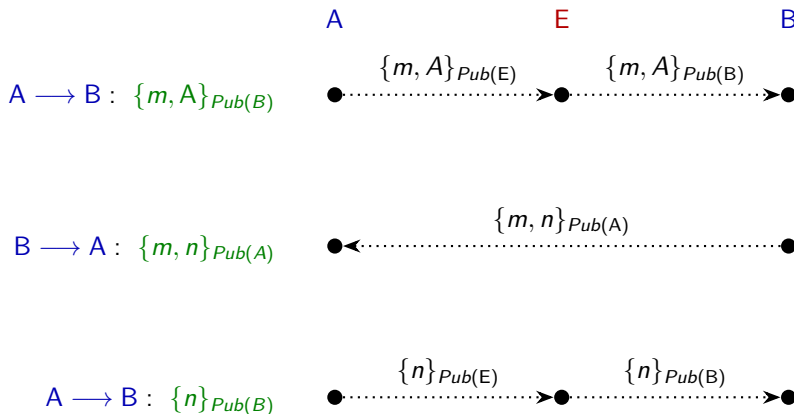
# An attack against the original protocol

Man-in-the-middle attacker **E** convinces **A** to start communication with **E** and uses the messages generated by **A** to follow the protocol with **B**, posing as **A**.



# An attack against the original protocol

Man-in-the-middle attacker **E** convinces **A** to start communication with **E** and uses the messages generated by **A** to follow the protocol with **B**, posing as **A**.



# The fixed protocol

- (1)  $A \longrightarrow B: \{m, A\}_{Pub(B)}$
- (2)  $B \longrightarrow A: \{m, n, B\}_{Pub(A)}$
- (3)  $A \longrightarrow B: \{n\}_{Pub(B)}$

- Only  $B$  can decrypt the message sent in (1)
- $A$  knows that only  $B$  can have sent the message in (2)
- $B$  knows that only  $A$  can have sent the message in (1)
- the nonces  $m$  and  $n$  are shared secrets

But these properties are informal and approximate, and we've only described what's *supposed* to happen . . .

## Security Protocol Language

- One of a range of languages and models for analyzing crypto-protocols
- Others include Spi calculus, strand spaces
- Supports reasoning based on events (vs transitions)
  
- Asynchronous communication
- Messages persist on network
- New-name generation on output
- Input pattern-matches messages on network

# Syntax

- We take an infinite set of **names**

$$\mathbf{Names} = \{m, n, \dots, A, B, \dots\}$$

$m, n, \dots$  stand for nonces,  $A, B$  agent identifiers

- with **name variables** inside the language

$$x, y, \dots, X, Y$$

$X, Y$  tend to range over agent identifiers,  $x, y$  over nonces

- The language also contains **message variables**

$$\psi, \psi', \psi_1, \dots$$

- Indices shall be used to identify components of parallel compositions

$$i \in \mathbf{Indices}$$

# SPL syntax

Name expressions  $v ::= n \mid A \mid \dots \mid x \mid X$

Key expressions  $K ::= Pub(v) \mid Priv(v) \mid Key(v, v')$

Messages  $M ::= \psi \mid v \mid k \mid M_1, M_2 \mid \{M\}_k$

Processes  $p ::=$   
    out new  $\vec{x} M.p$   
    | in pat  $\vec{x}, \vec{\psi} M.p$   
    |  $\parallel_{i \in I} p_i$

- out new  $\vec{x} M.p$  generates a new vector of nonce values  $\vec{n}$ , using them for  $\vec{x}$  to output  $M[\vec{n}/\vec{x}]$  before resuming as  $p[\vec{n}/\vec{x}]$
- input uses pattern matching...

# Messages as patterns

Messages can contain variables.

$$\psi \quad n, x \quad \{m, y, B\}_{Pub(A)}$$

These are used to perform matching.

Examples:

- match  $\{A, B\}_{Pub(A)}$  against the pattern  $\psi$
- match  $\{m, n, B\}_{Pub(A)}$  against the pattern  $\{m, x, Y\}_{Pub(A)}$
- match  $m, (n, A)$  against the pattern  $n, x$  where  $m \neq n$



# Messages as patterns

Messages can contain variables.

$$\psi \quad n, x \quad \{m, y, B\}_{Pub(A)}$$

These are used to perform matching.

Examples:

- match  $\{A, B\}_{Pub(A)}$  against the pattern  $\psi \quad \psi \mapsto \{A, B\}_{Pub(A)}$
- match  $\{m, n, B\}_{Pub(A)}$  against the pattern  $\{m, x, Y\}_{Pub(A)}$
- match  $m, (n, A)$  against the pattern  $n, x$  where  $m \neq n$

# Messages as patterns

Messages can contain variables.

$$\psi \quad n, x \quad \{m, y, B\}_{Pub(A)}$$

These are used to perform matching.

Examples:

- match  $\{A, B\}_{Pub(A)}$  against the pattern  $\psi \quad \psi \mapsto \{A, B\}_{Pub(A)}$
- match  $\{m, n, B\}_{Pub(A)}$  against the pattern  $\{m, x, Y\}_{Pub(A)}$   
 $x \mapsto n, Y \mapsto B$
- match  $m, (n, A)$  against the pattern  $n, x$  where  $m \neq n$

# Messages as patterns

Messages can contain variables.

$$\psi \quad n, x \quad \{m, y, B\}_{Pub(A)}$$

These are used to perform matching.

Examples:

- match  $\{A, B\}_{Pub(A)}$  against the pattern  $\psi \quad \psi \mapsto \{A, B\}_{Pub(A)}$
- match  $\{m, n, B\}_{Pub(A)}$  against the pattern  $\{m, x, Y\}_{Pub(A)}$   
 $x \mapsto n, Y \mapsto B$
- match  $m, (n, A)$  against the pattern  $n, x$  where  $m \neq n$       **no match**

# Conventions

- Messages with no free variables are **closed**.
- The input **in pat  $\vec{x}, \vec{\psi}$   $M.p$**  binds the variables  $\vec{x}, \vec{\psi}$  in  $M$  and  $p$ , attempting to match the pattern  $M$  against any closed message that has been output to the network
- The output **out new  $\vec{x}$   $M.p$**  binds the variables  $\vec{x}$  in  $M$  and  $p$
- A process with no unbound (free) variables is closed.

We write:

- **out  $M.p$**  where the list of **new** variables is empty
- **in  $M.p$**  where the lists of name and message variables are precisely the free name and message variables in  $M$
- **nil** is the empty parallel composition, which may be freely omitted
- use infix notation for finite parallel composition:  $p_1 \parallel p_2$  is  $\parallel_{i \in \{1,2\}} p_i$
- replication of a process **!p** is  $\parallel_{i \in \omega} p$

# Names and variables

- A closed process can contain names: nonce values  $n$  or real agent identifiers  $A$
- Variables are **not** names
- The set of all names in a process term is  $\text{names}(p)$  and in a message is  $\text{names}(M)$
- Example:

$$\text{names}(\text{out new } x\{n, x\}_{\text{Pub}(A)}) = \{n, A\}$$

# The NSL protocol in SPL

The initiator initiator of the protocol is parameterized by the identity of the initiator and their intended participant:

$$\begin{aligned} \mathit{Init}(A, B) \quad \equiv \quad & \text{out new } x \{x, A\}_{\text{Pub}(B)}. \\ & \text{in } \{x, y, B\}_{\text{Pub}(A)}. \\ & \text{out } \{y\}_{\text{Pub}(B)} \end{aligned}$$

The responder:

$$\begin{aligned} \mathit{Resp}(B) \quad \equiv \quad & \text{in } \{x, Z\}_{\text{Pub}(B)}. \\ & \text{out new } y \{x, y, B\}_{\text{Pub}(Z)}. \\ & \text{in } \{y\}_{\text{Pub}(B)} \end{aligned}$$

# Dolev-Yao assumptions

We can program various forms of attacker process. Viewing messages as **persisting** once output to the network, they output new messages built from existing ones.

$$Spy_1 \equiv \text{in } \psi_1.\text{in } \psi_2.\text{out } (\psi_1, \psi_2)$$

$$Spy_2 \equiv \text{in } (\psi_1, \psi_2).\text{out } \psi_1.\text{out } \psi_2$$

$$Spy_3 \equiv \text{in } X.\text{in } \psi.\text{out } \{\psi\}_{Pub(X)}$$

$$Spy_4 \equiv \text{in } Priv(X).\text{in } \{\psi\}_{Pub(X)}.\text{out } \psi$$

$$Spy \equiv \parallel_{i \in \{1,2,3,4\}} Spy_i$$

# The NSL system

We reason about concurrent runs of the protocol in parallel with  $\omega$ -copies of the attacker.

$$\begin{aligned} P_{spy} &\equiv !Spy \\ P_{init} &\equiv \prod_{A, B \in \mathbf{Agents}} !Init(A, B) \\ P_{resp} &\equiv \prod_{A \in \mathbf{Agents}} !Resp(A) \end{aligned}$$

Messages from one run of the protocol can be used by the attacker against another run of the protocol.

$$NSL \equiv \prod_{i \in \{resp, init, spy\}} P_i$$



# Operational semantics

- A **configuration** is a tuple

$$\langle p, s, t \rangle$$

- $p$  is a **closed** process term
- $s$  is a finite subset of names: **the names already in use**
- $t$  is a subset of closed messages: **the messages that have been output to the network**
- **Proper** configurations:
  - 1  $\text{names}(p) \subseteq s$
  - 2  $A \in s$  for every agent identifier  $A$
  - 3  $\bigcup \{\text{names}(M) \mid M \in t\} \subseteq s$
- Transitions are labelled with **actions**

$$\alpha ::= \text{out new } \vec{n} M \mid \text{in } M \mid i : \alpha$$

# Operational semantics

- **Output:** if  $\vec{n}$  all distinct and not in  $s$

$$\langle \text{out new } \vec{x} M.p, s, t \rangle \xrightarrow{\text{out new } \vec{n} M[\vec{n}/\vec{x}]} \langle p[\vec{n}/\vec{x}], s \cup \{\vec{n}\}, t \cup \{M[\vec{n}/\vec{x}]\} \rangle$$

- **Input:** if  $M[\vec{n}/\vec{x}][\vec{N}/\vec{\psi}] \in t$

$$\langle \text{in pat } \vec{x}, \vec{\psi} M.p, s, t \rangle \xrightarrow{\text{in } M[\vec{n}/\vec{x}][\vec{N}/\vec{\psi}]} \langle p[\vec{n}/\vec{x}][\vec{N}/\vec{\psi}], s, t \rangle$$

- **Parallel:**

$$\frac{\langle p_j, s, t \rangle \xrightarrow{\alpha} \langle p'_j, s', t' \rangle \quad j \in I}{\langle \parallel_{i \in I} p_i, s, t \rangle \xrightarrow{j:\alpha} \langle \parallel_{i \in I} p'_i, s', t' \rangle}$$

where  $p'_j = p_i$  for  $j \neq i$

# Reasoning from the transition semantics

## Secrecy of the responder's nonce:

Suppose  $Priv(A)$  and  $Priv(B)$  do not occur as the contents of any message in  $t_0$ . For all runs

$$\langle NSL, s_0, t_0 \rangle \xrightarrow{\alpha_1} \dots \langle p_{r-1}, s_{r-1}, t_{r-1} \rangle \xrightarrow{\alpha_r} \dots$$

where  $\langle NSL, s_0, t_0 \rangle$  is proper, if  $\alpha_r$  has the form  $resp : B : j : out\ new\ n\{m, n, B\}_{Pub(A)}$ , then  $n \notin t_l$  for any  $l \in \omega$ .

Proof idea: strengthen hypothesis, prove by induction / assume earliest violation.

The model obscures the key reasoning technique: that a violation must be by an event that causally depends (either through input/output or control) on an earlier event that violates the invariant.

↪ a Petri net semantics for SPL

# Petri net semantics of SPL

A net with persistent conditions representing all of SPL (not just particular processes at first).

Conditions viewed as being: **control**, **network** and **name**

- **Control** conditions form a set **C** of capacity-1 conditions

$$b ::= \text{out new } \vec{x} M.p \mid \text{in pat } \vec{x}, \vec{\psi} M.p \mid i : b$$

the control state of each thread

- **Network** conditions: form a set **O** of persistent conditions

$$\mathbf{O} = \{\text{closed messages}\}$$

the messages already output

- **Name** conditions: form a set **S** of capacity-1 conditions

$$\mathbf{S} = \mathbf{Names}$$

the names in use

# Control conditions

For a process  $p$ , the subset of control conditions

$$lc(p)$$

is called its **initial conditions**.

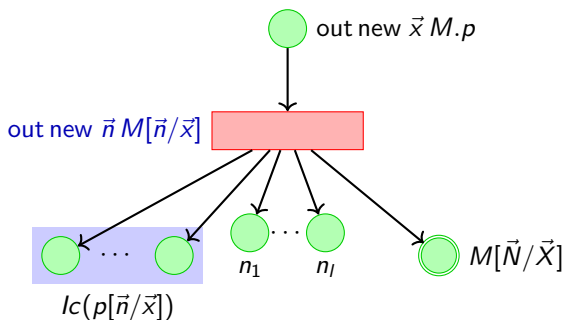
$$\begin{aligned}lc(\text{out new } \vec{x} M.p) &= \text{out new } \vec{x} M.p \\lc(\text{in pat } \vec{x}, \vec{\psi} M.p) &= \text{in pat } \vec{x}, \vec{\psi} M.p \\lc(\parallel_{i \in I} p_i) &= \bigcup_{i \in I} i : lc(p)\end{aligned}$$

where  $i : C = \{i : b \mid b \in C\}$  for  $C \subseteq \mathbf{C}$ .

# The events of SPL: output

The set **Events** includes:

if  $\text{out new } \vec{x} M.p$  is a closed term and  $\vec{n} = n_1, \dots, n_l$  are distinct names to match  $\vec{x} = x_1, \dots, x_l$




Events are labelled with an action.

# The events of SPL: input


The set **Events** includes:

if  $\text{in pat } \vec{x}, \vec{\psi} M.p$  is a closed term and  $\vec{n} = n_1, \dots, n_l$  are distinct names to match  $\vec{x} = x_1, \dots, x_l$  and  $\vec{L} = L_1, \dots, L_k$  are messages to match  $\vec{\psi} = \psi_1, \dots, \psi_k$



# The events of SPL: tags

If e.g. there is an event



then there is an event

