

Static semantics of secret channel abstractions

Marco Giunti

University of Porto and University of Beira Interior, Portugal

NordSec, October 16 2014

Motivation

- ▶ The pi calculus and its variants based on cryptographic operations are often used for protocol analysis
- ▶ E.g. googling - "pi calculus" protocol - returns 50k hits
- ▶ All pi calculus variants make use of the **new** (restriction) operator
- ▶ The **new** operator allows to
 1. create a channel name and limit its use within a certain scope
 2. enlarge the channel's scope by communicating the channel to others

Security problems

- ▶ The scope extrusion mechanism allows the mobility of the communication structure (and the great expressiveness of the pi calculus), but comports security problems
- ▶ Restricted channels cannot be implemented as dedicated channels, and open channels are not secure by default
- ▶ The spi calculus and the applied pi calculus do not rely on restriction for secure communication and use cryptographic encryption

Motivating example

A simple protocol to exchange a **confidential** information

$$P = (\text{new } c)((\text{new } s)(\bar{c}\langle s \rangle . \bar{s}\langle \text{pwd} \rangle) \mid c(x).x(y).\bar{p}\langle x \rangle)$$

- ▶ Two parallel threads communicating over restricted channel c
- ▶ The left thread generates a (secure) channel s to send the password, and forwards s over c
- ▶ The right thread receives a channel x from c , uses x to retrieve some data, and releases x over a public (free) channel p
- ▶ How to **implement** this protocol in an open network ?

Example: naive implementation

To avoid dedicated channels we use public key cryptography.

– (new s) mapped into generation of keys (new s^+, s^-)

– **Aim: to encrypt the password:** $\{\text{pwd}\}_{s^+}$

$$\mathbf{pi:} \quad c(x).x(y).\bar{p}\langle x \rangle \quad (1)$$

$$\mathbf{spi:} \quad \text{net}(z).\text{decrypt } z \text{ as } \{x^+, x^-\}_{c^-} \text{ in} \quad (2)$$

$$\text{net}(w).\text{decrypt } w \text{ as } \{y\}_{x^-} \text{ in } \bar{p}\langle x^+, x^- \rangle$$

- ▶ (2) is the (spi calculus) code of the receiver in (1)
- ▶ Keys sent on the network through the packet $\{s^+, s^-\}_{c^+}$
- ▶ To retrieve s^+, s^- processes must use the decryption key c^-

Example: naive implementation

Lack of forward secrecy

The implementation above suffers from a number of problems.

- ▶ The most serious is the lack of forward secrecy
- ▶ Informally: password in $\{\text{pwd}\}_{s^+}$ can be retrieved by buffering the message and subsequently using the key s^-
- ▶ **Formally**: the behavioral equation of pi calculus below is not preserved by the spi calculus translation

$$\begin{aligned}
 P &= (\text{new } c)((\text{new } s)(\bar{c}\langle s \rangle . \bar{s}\langle \text{pwd} \rangle) \mid c(x).x(y).\bar{p}\langle x \rangle) \\
 P &\cong (\text{new } s)(\bar{p}\langle s \rangle) \quad (p \in \text{fv}(P))
 \end{aligned}$$

- ▶ The equation ensures a well-known fact: in the pi calculus restricted communications are invisible

A secret pi calculus ($S\pi$)

- ▶ To avoid this problem in EXPRESS/SOS'12 we introduced a pi calculus featuring both a **new** and a **hide** operator
- ▶ The **new** operator does not ensure any secrecy: that is, in secret pi:

$$P \not\cong (\text{new } s)(\bar{p}\langle s \rangle)$$

- ▶ To recover the equation programmers must use the **hide** operator:

$$H = (\text{new } c)([\text{hide } s][\bar{c}\langle s \rangle.\bar{s}\langle \text{pwd} \rangle \mid c(x).x(y).\bar{p}\langle x \rangle])$$

$$H \cong_{S\pi} [\text{hide } s][\bar{p}\langle s \rangle]$$

- ▶ The brackets delimit the **static** scope of **hide**, which includes the receiver. Note: s **cannot be extruded** (e.g. by $\bar{p}\langle s \rangle$)

A type system to control the scope of channels

- ▶ In the secret pi calculus the scope of channels protected by **hide** is managed by the reduction system
- ▶ The runtime system can be interpreted as a specialized middleware for secure communications featuring local channels
- ▶ **This talk:** a type system for a standard pi calculus that disallows the extrusion of channels "declared" as *static*
- ▶ Our construction can be seen as an **API** for secure programming:
 - channels protected by **hide** are translated into typed channels with static scope
 - processes trying to leak secret (static) channels are rejected

Syntax of pi calculus types and processes

$T ::=$	Types:	$P ::=$	Processes:
$m \text{ chan} \langle T \rangle$	channel	$x(y \div B).P$	input
\top	top	\dots	pi
$m ::=$	Modalities	$B ::=$	Blocked entry:
s	static	\emptyset	empty
d	dynamic	$B \cup \{T\}$	type

- ▶ I/o types are decorated with a scope modality
- ▶ Input processes decorated with blocked types to instruct the type checker: semantics unaffected
- ▶ When B is empty: $x(y).P \stackrel{\text{def}}{=} x(y \div \emptyset).P$

Example, typed syntax

- ▶ We rewrite the secret pi calculus process

$$H = (\text{new } c)([\text{hide } s][H'])$$

$$H' = \bar{c}\langle s \rangle . \bar{s}\langle \text{pwd} \rangle \mid c(x).x(y).\bar{p}\langle x \rangle$$

- ▶ Typed syntax:

$$P = (\text{new } c : d \text{ chan} \langle T_2 \rangle)((\text{new } s : s \text{ chan} \langle T \rangle)(H'))$$

$$T_2 = d \text{ chan} \langle T \rangle$$

- ▶ **Note:** An upcast mechanism allows to send s over c by changing the type of c to $d \text{ chan} \langle s \text{ chan} \langle T \rangle \rangle$

An (abstract) API for secure programming

- ▶ We let programmers write code with the secret pi syntax

$$H = (\text{new } c)([\text{hide } s][H'])$$

$$H' = \bar{c}\langle s \rangle . \bar{s}\langle \text{pwd} \rangle \mid c(x).x(y).\bar{p}\langle x \rangle$$

- ▶ Code translated by guessing payload types of channels, scope modalities inferred automatically
- ▶ E.g. `pwd` has top type, `s` brings values of top type, ...

$$\llbracket H \rrbracket = (\text{new } c : d \text{ chan} \langle T_2 \rangle)((\text{new } s : s \text{ chan} \langle \top \rangle)(H'))$$

- ▶ Payload types different from top have a dynamic modality, e.g. $T_2 = d \text{ chan} \langle \top \rangle$

Static type checking

- ▶ Given the expected (dynamic) type T for p , we have

$$p: T \vdash \llbracket H \rrbracket$$

$$\llbracket H \rrbracket = (\text{new } c: d \text{ chan} \langle T_2 \rangle)(\text{new } s: s \text{ chan} \langle T \rangle)$$

$$(\bar{c} \langle s \rangle. \bar{s} \langle \text{pwd} \rangle) \mid c(x).x(y).\bar{p} \langle x \rangle$$

- ▶ More interestingly, the type system **rejects** attempts to **leak** channel s from p
- ▶ Specifically: the composition $\llbracket H \rrbracket \mid p(x)$ is ill-typed
- ▶ This is mandatory, as the reduction semantics of the pi calculus would allow the interaction of the two threads
- ▶ How we obtain this?

Downcasting to the rescue

- ▶ To type check $\llbracket H \rrbracket$ the payload type T_2 of c in the left thread must be upcasted to the type $s \text{ chan}\langle T \rangle$ (*)

$$\llbracket H \rrbracket = (\text{new } c : d \text{ chan}\langle T_2 \rangle)(\text{new } s : s \text{ chan}\langle T \rangle) \\ (\bar{c}\langle s \rangle . \bar{s}\langle \text{pwd} \rangle) \mid c(x) . x(y) . \bar{p}\langle x \rangle$$

- ▶ The right thread must assign T_2 as payload type of c as well, since channel c is used in i/o (specifically, it is used in input)
- ▶ In turn, the variable x gains type (*), and the “final” type of p is downcasted to the special type \bullet (void) to disallow extrusion
- ▶ The void type is not accessible to the programmer and is used in *return* environments to forbid the leak of static channels

Tracking the usage of channels

- ▶ We use return environments to keep track of the effective usage of channels
- ▶ Our judgements have the form

$$\Gamma \vdash P \triangleright \Delta$$

where Δ is a type environment with codomain = $Types \cup \{\bullet\}$

- ▶ The technique is reminding of those for algorithmic type checking of linear systems
- ▶ The typing rule for parallel crucially asks that return environments can be composed

$$\frac{\Gamma \vdash P_1 \triangleright \Delta_1 \quad \Gamma \vdash P_2 \triangleright \Delta_2}{\Gamma \vdash P_1 \triangleright \Delta_1 \otimes \Delta_2}$$

- ▶ Otimes: a void type can only be composed with top

Running example

- ▶ Given a suitable type T , we have that

$$p: T \not\vdash \llbracket H \rrbracket \mid p(x) \triangleright \Delta'$$

$$\llbracket H \rrbracket = (\text{new } c: d \text{ chan} \langle T_2 \rangle)(\text{new } s: s \text{ chan} \langle T \rangle) \\ (\bar{c} \langle s \rangle. \bar{s} \langle \text{pwd} \rangle) \mid c(x). x(y). \bar{p} \langle x \rangle$$

for any Δ' since:

- ▶ $p: T \vdash \llbracket H \rrbracket \triangleright p: \bullet$
- ▶ $p: T \vdash p(x) \triangleright \Delta$ with $\Delta(p) \neq \top$
- ▶ In contrast:
 - $p: T \vdash \llbracket H \rrbracket \mid (\text{new } p': T') p'(x) \triangleright p: \bullet$ since
 - $p: T \vdash (\text{new } p': T') p'(x) \triangleright p: \top$

Blocked types in input

- ▶ Following standard lines, we consider a pi calculus with reduction semantics and structural congruence (\equiv)
- ▶ Blocked types in input inserted in \equiv scope extrusion rule
- ▶ Example:

$$\llbracket H \rrbracket \mid p(z \div \emptyset) \equiv (\text{new } c : d \text{ chan} \langle T_2 \rangle)(\text{new } s : s \text{ chan} \langle T \rangle) \\ (\bar{c}\langle s \rangle.\bar{s}\langle \text{pwd} \rangle) \mid c(x).x(y).\bar{p}\langle x \rangle \mid p(z \div \{s \text{ chan} \langle T \rangle\})$$

- ▶ Process $p(z \div \{s \text{ chan} \langle T \rangle\})$ cannot upcast the required payload type since it is blocked
- ▶ In detail: types must have identifiers in order to avoid clashes: $(\text{new } c : d \text{ chan} \langle T_2 \rangle_{\vee})(\text{new } s : s \text{ chan} \langle T \rangle_n)(\dots)$ n perfect id

Soundness and expressiveness

- ▶ Typed processes reduce to typed processes (SR)
- ▶ Operational correspondence among (a fragment of) secret π -calculus processes and their typed translation

Assume Γ, Δ such that $\Gamma \vdash \llbracket H \rrbracket \triangleright \Delta$.

1. $H \rightarrow H'$ implies $\llbracket H \rrbracket \rightarrow \llbracket H' \rrbracket$
 2. $\llbracket H \rrbracket \rightarrow Q$ implies $H \rightarrow H'$ with $\llbracket H' \rrbracket \equiv Q$
- ▶ **Note** the typability assumption, essential to switch from middleware to software support of secret channels

Applications: protection against 3rd party code

Example: malicious list handler

$$\{ () \}_z = \bar{z}\langle \perp, \perp, \perp \rangle$$

$$\{ (\langle a_0, b_0 \rangle, \dots, \langle a_n, b_n \rangle) \}_z = (\text{new } z')(\bar{z}\langle a_0, b_0, z' \rangle \mid \\ \{ (\langle a_1, b_1 \rangle, \dots, \langle a_n, b_n \rangle) \}_{z'})$$

$$\text{ADD}(x, y, z) = z(h_1, h_2, z').((\text{new } z'')(\bar{z}\langle x, y, z'' \rangle \mid \bar{z}''\langle h_1, h_2, z'' \rangle) \mid \\ \overline{\text{port888}}\langle h_1, h_2 \rangle) \quad \text{\% \% Suspicious}$$

Fix: re-program the list, compile and ...

$$\text{STORESECCH}(H, y) = [\text{hide } x][H \mid (\text{new } z)(\{ () \}_z \mid \text{ADD}(x, y, z))]$$

ask to the type-checker! $\Gamma \vdash \llbracket \text{STORESECCH}(H, y) \rrbracket \mid Q$

Applications: Mandatory access control

DBUS is an *IPC* system using private and public bus for communication

–Previous versions: bug allows users to listen private bus

```
[marco]# echo $DBUS_SESSION_BUS_ADDRESS > Public/address  
[guest]# dbus-monitor --address /home/marco/Public/address
```

- ▶ We interpret this issue as MAC problem
- ▶ The private session bus address cannot be disclosed by its owner
- ▶ **Fix:** program the bus with `hide`. All users trying to leak the channel will be rejected

Limitations

- ▶ We just deal with direct information flows
 - We need protection against indirect flows, covert channels...
- ▶ Typed analysis does not scale
 - $\Gamma \vdash P$ and $\Gamma \vdash Q$ does not imply $\Gamma \vdash P \mid Q$
- ▶ Static typing is too demanding
 - We would need lightweight (dynamic) typing integrated with advanced functionalities
 - E.g. contracts, certificates, functions, cryptographic operations, ...

Extensions

- ▶ To understand better the static semantics of programs we need typed behavioural equivalences, typed bisimulation, ...
- ▶ The system has been designed to be easily integrated with other type systems
 - E.g. linear types, affine types, session types, ...
- ▶ Further design choice: keep the system algorithmic as possible
 - Algorithmic type checking and inference obtainable easily (by extending code of previous tools)

Thanks!

Questions?

Recent related work

- ▶ Myself: Algorithmic type checking for a pi-calculus with name matching and session types. J. Log. Algebr. Program. 82(8)
- ▶ Myself, Antonio Ravara: Towards Static Deadlock Resolution in the pi-Calculus. TGC 2013: 136-155
- ▶ Myself, Catuscia Palamidessi, Frank D. Valencia: Hide and New in the Pi-Calculus. EXPRESS/SOS 2012