


Privacy and Access Control



Rocco De Nicola
Dip. Sistemi e Informatica
Università di Firenze

denicola@dsi.unifi.it



Outline

- Motivations
- Privacy in Distributed Tuple Spaces:
 - Extending Linda with primitives for privacy in distributed and mobile applications
 - KriptoKlava
- Access Control via Typing
- μ -Klaim
- Types for μ -Klaim



Motivation

Process mobility poses to a lot of security problems

- secrecy and integrity of transmitted data and program code
- Malicious agents can attempt to access/modify private information of the nodes hosting them
- Malicious hosts can try to compromise agent's integrity/secrecy

Several programming/process languages with code mobility come equipped with security mechanisms (e.g. type systems, mechanisms for data/control flow analysis)



Attacks

- Communication channels
 - passive (e.g. traffic analysis)
 - active (e.g. message modifications/forging)
- Hosts
 - modification of host resources and data
 - denial of service
- Mobile Agents
 - modification of agent code
 - leak of sensible data

Typical defences: Cryptography, Access Control, Activity Monitoring, ... Types



Our Defenses

- Cryptography
- Types



Security Problems

- Linda provides no access protection to a tuple space
- No way to determine the issuer of an operation to the tuple space
- A process may retrieve/remove data that do not belong to it
- Shared data can be easily modified and corrupted



Our Proposal

- Extend Linda operations with cryptography:
 - Tuples can contain encrypted data
 - New Primitives for encryption/decryption
- Aims:
 - Change as little as possible the original Linda model
 - Make it Suitable for distributed application and mobile agent based application



Privacy, not Security

- Our principal aim is not to avoid that wrong data be retrieved
- Our aim is that even if data is eavesdropped or stolen, still it cannot be read
- A sort of PGP for Linda
- Smooth extension of Linda
 - The impact on the Linda model is minimal
 - Previous applications continue to work



Cryptography & Tuple Spaces

- Operations for inserting/retrieving encrypted tuples to/from tuple spaces (**ink**, **readk**, **outk**)
- Operations for encrypting/decrypting tuple contents (**encode** & **decode**)
- Operations for signing/verifying mobile agent code



ink & readk

1. look for and possibly retrieve a matching tuple,
2. attempt a decryption of the encrypted fields of the retrieved tuple
3. if the decryption fails:
 1. if the operation was an ink then put the retrieved tuple back in the tuple space,
 2. look for alternative matching tuples,
4. if all these attempts fail, then block until another matching tuple is available.



Extended Pattern Matching

- The original pattern matching has to be extended
- Two stages pattern matching:
 - In the first stage an encrypted field is seen as an ordinary field with the type “encrypted” and it can match only another “encrypted” field
 - In the second stage decryption takes place, and a further matching is performed with the decrypted clear-text fields



Keys and Mobile Agents

- symmetric and asymmetric key encryption techniques rely on the secrecy of private keys
- It is important that mobile code and mobile agents do not carry private keys when migrating to remote sites



Finer Grain Mechanisms

- Explicit operations: **enc** & **dec** acting on single tuple fields
- Mobile agents retrieve encrypted tuples with standard Linda operations (e.g. without decrypting them)
- Actual decryption will take place only at the home site (where the private key is stored) by stationary agents
- Wrong tuples retrieved by mistake have to be explicitly put back



Information Retrieval Agents

- Mobile agents can safely transport and use public keys also on remote sites
- Intermediate results can be encrypted so that they cannot be eavesdropped by other sites
- They can be decrypted only by the home site



CryptoKlava

- A subpackage of Klava providing these new modular extensions
- Based on Sun JCE (Java Cryptography Extension) providing basic interfaces and API for encryption
- Extended classes and extended operations



Types as security tools

Type systems have been successfully used to ensure *type safety* of programs since a long time to avoid *run-time errors*, and guarantee that data be used consistently with the expected operations.

Recently, work has been done on exploring and designing type systems for security:

- well-typed Java programs (and the corresponding verified bytecode) will never compromise the integrity of certain data
- Type systems for $D\pi$ -calculus (Hennessy-Riely, Yoshida-Hennessy), and (variants of) Ambient calculus (Cardelli-Ghelli-Gordon) have been proposed to control interaction



Type for access and code mobility control

- Models for Access Control
 - mechanisms to specify policies for access control
 - mechanisms to enforce such policies
- Capability-based Type System [DFPV – TCS2000]
 - Types as specification tool
 - to express control policies of nodes relative to resource access and code mobility)
 - to abstract process intentions (read, out, spawn, ...) relative to the different localities they interact with or migrate to
 - (static and dynamic) type checking as enforcement mechanism
 - only intentions that match security policies are allowed}
 - only processes whose types comply with network nodes security policies are admissible (well-typed-ness)



Types and Security Policies

- Each node, that indicates a physical machine, or a logical partition has an associated type:

$$l ::^{\delta} P$$

- Type δ describes the security policy of the node, i.e. what process P may do when running at site l .
- We have capabilities/privileges in correspondence of each process action

$$\{i, r, o, e, n\}$$



Types and Security Policies

- A type is a partial function

$$\delta : Loc \rightarrow \sqcap$$

Where Loc indicates the set of *localities* and \sqcap is a collection of non-empty set of *capabilities*.

Examples

- **LEGAL** $l ::= [l_1 \mapsto \{i, o\}, \dots] \text{ in}(\dots) @ l_1 . \text{nil}$
- **NOT LEGAL** $l ::= [l_1 \mapsto \{i, o\}, \dots] \text{ eval}(\dots) @ l_1 . \text{nil}$



The role of types

- The type of a node is set by a net coordinator and determines the access policy of the node in terms of access rights;
- Type inference permits determining *processes intentions*
- Type checking guarantees that only processes whose intentions match the rights granted by the coordinators are allowed to proceed.
- **Example:** privilege $[l' \mapsto \{e\}]$ in the type of locality l will enable processes running at l to perform an **eval** actions over l' .



The role of types

Apart from occurring in the specification of a node, type related information are introduced in two other syntactic constructs:

- in action **newloc**($u : \delta$) where δ specifies the security policy of the new node,
- in templates of formal parameter **!u: π** where π specifies the access rights corresponding to the operations that the receiving process wants to perform at u .

In both cases, the type information is not strictly necessary: it increases the flexibility of **newloc** (otherwise, some kind of 'default policy' should be assigned to the newly created node) and permits a simpler static type checking.



μ _Klaim: A core calculus for Klaim

- We take away :
 - distinction between logical and physical localities/addresses
 - allocation environments
 - higher order communication
 - types with global information
- For types, we have:
 - types with only local information
 - privilege exchanges
 - dynamic modifications of security policies
 - efficient type handling
 - simpler semantics and type systems



μ _Klaim syntax

■ Nets

$$N ::= 0 \quad | \quad l ::^{\delta} P \quad | \quad N_1 \parallel N_2$$

■ Processes

$$P ::= \text{nil} \quad | \quad a.P \quad | \quad P_1 \mid P_2 \quad | \quad A \quad (A \triangleq P)$$

■ Actions

$$a ::= \text{read}(T)@l \quad | \quad \text{in}(T)@l \quad | \quad \text{out}(t)@l \\ | \quad \text{eval}(P)@l \quad | \quad \text{newloc}(u : \delta)$$



Tuples and Templates

Templates	T	$::=$	F		F, T
Tem.Fields	F	$::=$	f		$!x$ $!u : \pi$
Tuples	t	$::=$	f		f, t
Tuple Fields	f	$::=$	e		$l : \mu$
Expressions	e	$::=$	V		x \dots



Matching Rules

$$(M_1) \quad \text{match}(V, V) = \epsilon$$

$$(M_2) \quad \text{match}(!x, V) = [V/x]$$

$$(M_3) \quad \text{match}(l, l) = \epsilon$$

$$(M_4) \quad \text{match}(!u, l) = [l/u]$$

$$(M_5) \quad \frac{\text{match}(F, f) = \sigma_1 \quad \text{match}(T, t) = \sigma_2}{\text{match}((F, T), (f, t)) = \sigma_1 \circ \sigma_2}$$



Structural Congruence

$$\text{(Com)} \quad N_1 \parallel N_2 \equiv N_2 \parallel N_1$$

$$\text{(Assoc)} \quad (N_1 \parallel N_2) \parallel N_3 \equiv N_1 \parallel (N_2 \parallel N_3)$$

$$\text{(Abs)} \quad l :: P \equiv l :: (P|\mathbf{nil})$$

$$\text{(PrInv)} \quad l :: A \equiv l :: P \quad \text{if } A \stackrel{\Delta}{=} P$$

$$\text{(Clone)} \quad l :: (P_1|P_2) \equiv l :: P_1 \parallel l :: P_2$$



Untyped μ -Klaim Semantics

$$\llbracket t \rrbracket = et$$

$$l :: \text{out}(t)@l'.P \parallel l :: P' \xrightarrow{o(l,et,l')} l :: P \parallel l :: P' \parallel l' :: \langle et \rangle$$

$$l :: \text{eval}(Q)@l'.P \parallel l' :: P' \xrightarrow{e(l, ,l')} l :: P \parallel l' :: P'|Q$$

$$\text{match}(\llbracket T \rrbracket, et) = \sigma$$

$$l :: \text{in}(T)@l'.P \parallel l' :: \langle et \rangle \xrightarrow{i(l,et,l')} l :: P\sigma \parallel l' :: \text{nil}$$

$$\text{match}(\llbracket T \rrbracket, et) = \sigma$$

$$l :: \text{read}(T)@l'.P \parallel l' :: \langle et \rangle \xrightarrow{r(l,et,l')} l :: P\sigma \parallel l' :: \langle et \rangle$$

$$l' \notin L$$

$$L \vdash l :: \text{newloc}(u).P \xrightarrow{n(l,-,l')} L \cup \{l'\} \vdash l :: P[l'/u] \parallel l' :: \text{nil}$$



Typed μ -Klaim Semantics

- We consider now a few of the previous rules by taking into account types.
- We ignore labels, these are not needed in this framework. Labelled semantics is useful for open systems and for logical specification.
- Notation $\delta \Vdash_l P$ indicates that process P , located at l complies with the restrictions imposed by type δ .
- A net is **well typed** if each node, say l , complies with $\delta \Vdash_l P$



Eval

$$\frac{\delta' \Vdash_{l'} Q}{l ::^\delta \text{eval}(Q) @ l'. P \parallel l' ::^{\delta'} P' \succ\longrightarrow l ::^\delta P \parallel l' ::^{\delta'} P' | Q}$$

Process Q must be dynamically type checked against the policy of node l' , this is necessary since no a-priori knowledge of the target node policy can be assumed, no static checking performed in l over the spawned process can be useful.



newloc

$$l' \notin L$$


$$\frac{}{L \vdash l :: \delta \text{ newloc}(u: \delta').P \rightsquigarrow L \cup \{l'\} \vdash l :: \delta[l' \mapsto \delta(l)] P[l'/u] \parallel l' :: \delta'[l'/u] \text{ nil}}$$

where $\delta_1[\delta_2]$ denotes the pointwise union of functions δ_1 and δ_2 .

It is assumed that the creating node has over the created node all the privileges it has on itself.

The check that $\delta' \preceq \delta[u \mapsto \delta(l)]$ (i.e. the specified access policy δ' is in agreement with the access policy δ of the node executing the operation extended with the ability of performing over l' all the operations allowed locally) is left to static type inference.

This check prevents a malicious node l from forging capabilities by creating a new node with more powerful privileges and then sending malicious process that takes advantage of capabilities not owned by l .



$$match_{\delta}(\llbracket T \rrbracket, t) = \sigma$$

$$l::^{\delta} \mathbf{in}(T)@l'.P \parallel l'::\langle t \rangle \succ\longrightarrow l::^{\delta} P\sigma \parallel l'::\mathbf{nil}$$

The new pattern matching function $match_{\delta}$ is defined like $match$ but it also verifies that process $P\sigma$ does not perform illegal actions w.r.t. δ . Because of the static inference, the definition of $match_{\delta}$ simply requires the following change to untyped $match$.

$$\pi \subseteq \delta(l')$$

$$match_{\delta}(!u:\pi, l') = [l'/u]$$



Type Soundness

- processes running in well-typed nets do not attempt to execute actions that are not allowed by the capabilities they own (*type safety*)
- The above property is preserved along reductions (*subject reduction*).

Dynamic security policies

- One of our goal is to enable dynamic modifications of security policies. We want to permit transition like the following where

$$\begin{aligned}
 l_1 &:: [l_2 \mapsto \{i\}] \text{ in}(!u : \{o\})@l_2.\text{out}(100)@u.\text{nil} \parallel \\
 l_2 &:: [l_2 \mapsto \{o\}, \dots] \text{ out}(l)@l_2.\text{nil}
 \end{aligned}$$

$\xrightarrow{\text{out}}$
 $\xrightarrow{\text{in}}$

$$\begin{aligned}
 l_1 &:: [l_2 \mapsto \{i\}, l \mapsto \{o\}] \text{ out}(100)@l.\text{nil} \parallel \\
 l_2 &:: [l_2 \mapsto \{o\}, \dots] \text{ nil}
 \end{aligned}$$

l_2 grants l_1 the capability of performing an **out** at l .

- Problem: How to guarantee that capabilities are not forged?**

Exchanging Privileges - 1

1. **out**: each locality is annotated with the capabilities passed along with it

$$N_1 \triangleq l_1 :: [l_2 \mapsto \{i\}] \text{ in} (!u : \{o\}) @ l_2 . \text{out}(100) @ u . \text{nil} \parallel \\ l_2 ::^\delta \text{ out}(l : [l_1 \mapsto \{o, e\}, l_3 \mapsto \{i\}]) @ l_2 . \text{nil}$$

2. When **out** is performed it is checked that the capabilities passed along with the localities be really owned by the node performing the **out**

$$N_1 \xrightarrow{\text{out}} l_1 :: [l_2 \mapsto \{i\}] \text{ in} (!u : \{o\}) @ l_2 . \text{out}(100) @ u . \text{nil} \parallel \\ l_2 ::^\delta \text{ tuple}(l : [l_1 \mapsto \{o, e\}, l_3 \mapsto \{i\}])$$

only if $\{o, e, i\} \subseteq \delta(l)$



Exchanging Privileges - 2

3. When a **read/in** is performed (communication takes place) it is verified that the accessed tuple can pass all the capabilities required in the template to the locality performing the **read/in**

$$\begin{array}{l} l_1 :: [l_2 \mapsto \{i\}] \quad \mathbf{in}(!u : \{o\})@l_2.\mathbf{out}(100)@u.\mathbf{nil} \quad || \\ l_2 ::^\delta \mathbf{tuple}(l : [l_1 \mapsto \{o, e\}, l_3 \mapsto \{i\}]) \\ \downarrow \mathbf{in} \\ l_1 :: [l_2 \mapsto \{i\}, l \mapsto \{o\}] \quad \mathbf{out}(100)@l.\mathbf{nil} \quad || \\ l_2 ::^\delta \mathbf{nil} \\ \text{because } \{o\} \subseteq \{o, e\} \end{array}$$



Additional Information

- More in a paper just presented at ICALP by D. Gorla and R. Pugliese.



<http://music.dsi.unifi.it>

- A few papers
- Current Implementation:
 - KriptoKlava
 - Type Checker for Access Control (?)