

Security Policies as Membranes in Systems for Global Computing

Vladimiro Sassone

University of Sussex, UK

GC 2004: MyThS/MIKADO/DART Meeting
Venice 15.06.04

with D. Gorla, M. Hennessy

1 Why

2 What

3 How

- Barring actions
- Counting actions
- Sequencing actions
- Controlling coalitions

4 Conclusion

Why

- Most calculi/languages for GC rely on *code mobility* to model interprocesses interactions;
- This leads to *security concerns* (malicious agents can compromise 'good' sites through viruses, spammings, denial-of-service attacks, ...);

Why

- Most calculi/languages for GC rely on *code mobility* to model interprocesses interactions;
- This leads to *security concerns* (malicious agents can compromise 'good' sites through viruses, spamming, denial-of-service attacks, ...);
- Thus, code mobility usually equipped with *security checks*:
 - 1 *static checks*: make the run-time as efficient as possible, but it may be not adequate in practice;
 - 2 *dynamic checks*: make the runtime heavier, execution slower, but are flexible.

Simple

- **Systems** are (plain) collections of sites;
- **Sites** are places for computations, divided in at least two layers:
 - a computing body
 - a *membrane*, to carry on security related issues
- **membranes** regulate the interactions between the computing body and the environment around the site
- differently from Boudol's and Stefani's: our membranes are *not* fully-fledged computing entities. They only implement higher-level (type related) verification on incoming agents.

The Objectives

Run an initial investigation into *what kind* of security policies can be implemented through membranes, and *how*.

This is related to, and aims at generalizing for the specific application

- the security types developed for $D\pi$ and $KLAIM$;
- the session types by [Honda et al](#);
- the generic types by [Igarashi, Kobayashi](#).

What

- 1 a *formal framework* to formalize processes running in a GC system, whose activities are *local computations* and *migrations*;
- 2 *membranes* to implement advanced checks on incoming agents (including notions of *trust* and *proof-carrying code*);
- 3 *tools* to enforce different kind of policies.

A Calculus for Migrations

A minimal calculus (Turing not an issue here)

BasicActions $a, b, c, \dots \in \text{ACT}$

Localities $l, h, k, \dots \in \text{Loc}$

Agents $P, Q, R ::= \mathbf{nil} \mid a.P \mid \mathbf{go}_T l.P \mid P \mid Q \mid !P$

Systems $N ::= \mathbf{0} \mid l \llbracket M \triangleright P \rrbracket \mid N_1 \parallel N_2$

where

A Calculus for Migrations

A minimal calculus (Turing not an issue here)

BasicActions $a, b, c, \dots \in \text{ACT}$

Localities $l, h, k, \dots \in \text{LOC}$

Agents $P, Q, R ::= \mathbf{nil} \mid a.P \mid \mathbf{go}_T l.P \mid P \mid Q \mid !P$

Systems $N ::= \mathbf{0} \mid l \llbracket M \triangleright P \rrbracket \mid N_1 \parallel N_2$

where

- $l \llbracket M \triangleright P \rrbracket$ is a site with address l , membrane M and hosting process P ;
- $\mathbf{go}_T l.P$ is an agent willing to migrate on l , whose body is P and exhibiting as PCC the policy T .

Dynamic Semantics – local

Local behaviours:

$$\llbracket M \triangleright a.P \mid Q \rrbracket \rightarrow \llbracket M \triangleright P \mid Q \rrbracket$$

Remark: we are not really interested in the local computations.

Dynamic Semantics – migration

Migration:

$$k[M \triangleright \mathbf{go}_T.I.P|Q] \parallel l[M' \triangleright R] \rightarrow k[M \triangleright Q] \parallel l[M' \triangleright P|R]$$

This reduction may happen only if P *complies with* M' .

Dynamic Semantics – migration

Migration:

$$k[M \triangleright \mathbf{go}_\tau I.P|Q] \parallel I[M' \triangleright R] \rightarrow k[M \triangleright Q] \parallel I[M' \triangleright P|R]$$

This reduction may happen only if *P complies with M'*.

But checking whole processes at migration can be very expensive!

Solution: PCCs. A source-generated and certified 'process outline' accepted as such at destination.

The matter with certification

When can we consider PCCs?

- They are **easy to verify** (they are usually very small, if compared to the process they refer to), **but**
- they can be **dangerous** (if they don't certify properly the process behaviour)

The matter with certification

When can we consider PCCs?

- They are **easy to verify** (they are usually very small, if compared to the process they refer to), **but**
- they can be **dangerous** (if they don't certify properly the process behaviour)

A compromise:

*we can safely consider PCCs of agents coming from **trusted sites**, i.e. sites that calculate the PCC attached to a migrating agent "properly."*

Trust

Each site store the trust it has on other sites, as part of its *membrane*.

Thus, a membrane is a couple (M_t, M_p) , where

- $M_t : \text{LOC} \rightarrow \{\text{good, bad, unknown}\};$
- M_p is an upper bound to the local actions of incoming agents.

The Migration Rule – revised

$$k[M \triangleright \mathbf{go}_T l.P|Q] \parallel l[M' \triangleright R] \rightarrow k[M \triangleright Q] \parallel l[M' \triangleright P|R] \quad \text{if } M' \vdash_T^k P$$

where $M' \vdash_T^k P$ is

if $M'_t(k) = \text{good}$ **then** (T enforces M'_p) **else** $\vdash P : M'_p$

and

- predicate **enforces** is a partial order on policies;
- \vdash is a compliance check of a process against a policy.

Policies as Constraints on Legal Actions

- a site only provides some methods (i.e. only some actions can be executed while running in it)
- a policy T is a subset of $ACT \cup Loc$ where
 - a process can only execute locally actions in T
 - a process can only migrate on sites in T

Policies as Constraints on Legal Actions

- a site only provides some methods (i.e. only some actions can be executed while running in it)
- a policy T is a subset of $\text{ACT} \cup \text{Loc}$ where
 - a process can only execute locally actions in T
 - a process can only migrate on sites in T
- T enforces T' is simply defined as $T \subseteq T'$;

Policies as Constraints on Legal Actions

- a site only provides some methods (i.e. only some actions can be executed while running in it)
- a policy T is a subset of $\text{ACT} \cup \text{Loc}$ where
 - a process can only execute locally actions in T
 - a process can only migrate on sites in T
- T enforces T' is simply defined as $T \subseteq T'$;
- judgment \vdash is simple. The key rules are

$$\frac{\vdash P : T}{\vdash a.P : T} \quad a \in T$$

$$\frac{\vdash P : T'}{\vdash \text{go}_{T'} . l.P : T} \quad l \in T$$

Policies as Constraints on Legal Actions (ctd)

- a system N is *well-formed*, written $\vdash N : \mathbf{ok}$, if “good” nodes only hosts “good” agents. Formally:

$$\frac{\vdash P : M_p}{\vdash \llbracket M \Downarrow P \rrbracket : \mathbf{ok}} \quad / \text{ good}$$

$$\frac{}{\vdash \llbracket M \Downarrow P \rrbracket : \mathbf{ok}} \quad / \text{ not good}$$

Policies as Constraints on Legal Actions (ctd)

- a system N is *well-formed*, written $\vdash N : \mathbf{ok}$, if “good” nodes only hosts “good” agents. Formally:

$$\frac{\vdash P : M_p}{\vdash \llbracket M \triangleright P \rrbracket : \mathbf{ok}} \quad / \text{ good}$$

$$\frac{}{\vdash \llbracket M \triangleright P \rrbracket : \mathbf{ok}} \quad / \text{ not good}$$

- Subject Reduction:** If $\vdash N : \mathbf{ok}$ and $N \rightarrow N'$, then $\vdash N' : \mathbf{ok}$.

Counting Legal Actions

- sometimes, legal actions can be performed only a certain number of times. E.g.:
 - a fair mail server allows its clients to send mails, **but**:
 - it should block spamming activities of malicious clients; **thus**:
 - it could allow sending at most K mails for each login of each client.

Counting Legal Actions

- sometimes, legal actions can be performed only a certain number of times. E.g.:
 - a fair mail server allows its clients to send mails, **but**:
 - it should block spamming activities of malicious clients; **thus**:
 - it could allow sending at most K mails for each login of each client.
- Policies are *multisets* containing elements from $\text{ACT} \cup \text{Loc}$;
- T enforces T' is multisets inclusion;
- \vdash adapts straightforwardly from the case of sets:

$$\frac{\vdash P : T}{\vdash a.P : T \cup \{a\}}$$

$$\frac{\vdash P : T'}{\vdash \mathbf{go}_T, l.P : T \cup \{l\}}$$

$$\frac{\vdash P : T_1 \quad \vdash Q : T_2}{\vdash P \mid Q : T_1 \cup T_2}$$

Counting Legal Actions (ctd)

This setting enforces a *thread-wise* property. Indeed,

- if two different agents P and Q individually send at most K mails,
- when they both run in the mail server, the agent $P \mid Q$ can send *more than* K mails (actually, it can send $2K$ mails)

Thus, the well-formedness predicate for good sites is changed as

$$\frac{\forall i. (P_i \text{ a thread and } \vdash P_i : M_p)}{\vdash I \llbracket M \triangleright P_1 \mid \dots \mid P_n \rrbracket : \mathbf{ok}} \quad I \text{ good}$$

Subject reduction holds for this modified judgment

Sequencing Legal Actions

- sometimes, legal actions can be performed only in a certain order. E.g.
 - before exploiting the functionalities of a mail server, you must have logged in, **and**
 - before logging out, you must have saved the status of the transaction.

This can be easily formalized by *(deterministic) finite automata*

```
usr.pwd.(list + send + retr + del + reset)*.quit
```

Sequencing Legal Actions

- sometimes, legal actions can be performed only in a certain order. E.g.
 - before exploiting the functionalities of a mail server, you must have logged in, **and**
 - before logging out, you must have saved the status of the transaction.

This can be easily formalized by *(deterministic) finite automata*

```
usr.pwd.(list + send + retr + del + reset)*.quit
```

- Policies are DFAs;
- T enforces T' is inclusion of DFAs's languages;
- $\vdash P : T$ holds if the language of P is accepted by T .

Sequencing Legal Actions (ctd)

- As well-known, inclusion of regular languages can be calculated easily, once given the associated DFAs
- What about predicate $\vdash P : T$?
 - we expect that calculating it is harder than verifying PCCs (i.e. verifying predicate **enforces**)
 - But, how harder? Is it decidable?
 - what is the language associated to an agent?

Sequencing Legal Actions (ctd2)

- an agent can be easily associated to a *concurrent regular expression*: regular exprs with *shuffle* \otimes and *shuffle closure* \odot .
- e.g., agent $!(a.b \mid c.go_l.P)$ can be represented as

$$((a \cdot b) \otimes (c \cdot l))^{\odot}$$

we are only interested in the *local behaviour* of the agent.

Sequencing Legal Actions (ctd2)

- an agent can be easily associated to a *concurrent regular expression*: regular exprs with *shuffle* \otimes and *shuffle closure* \odot .
- e.g., agent $!(a.b \mid c.go_l.P)$ can be represented as

$$((a \cdot b) \otimes (c \cdot l))^\odot$$

we are only interested in the *local behaviour* of the agent.

- we can derive the language associated to this CRE and check whether it is contained in the language accepted by the policy;
- CREs can be represented as Petri nets. Inclusion of a Petri net in a DFA is *decidable*, even if *super-exponential*;
- This is done by *static analysis algorithm*, not by a type system!

Controlling Coalitions at a Site

- policies as multisets and as DFAs can only express thread-oriented properties;
- Dealing with the overall behaviour of a site; Two options: When agent P want to migrate on I , containing agent R

Controlling Coalitions at a Site

- policies as multisets and as DFAs can only express thread-oriented properties;
- Dealing with the overall behaviour of a site; Two options: When agent P want to migrate on I , containing agent R
 - 1 freeze and retrieve the current content of the site, viz. R ;
check whether $P \mid R$ respects the policy of the site;
reactivate R and, according to the result of the checking phase, activate P .

Controlling Coalitions at a Site

- policies as multisets and as DFAs can only express thread-oriented properties;
- Dealing with the overall behaviour of a site; Two options: When agent P want to migrate on I , containing agent R
 - 1 freeze and retrieve the current content of the site, viz. R ;
check whether $P \mid R$ respects the policy of the site;
reactivate R and, according to the result of the checking phase, activate P .
 - 2 let membranes evolving at run-time: they are decreased with the privileges granted to P .

Controlling Coalitions at a Site

- policies as multisets and as DFAs can only express thread-oriented properties;
- Dealing with the overall behaviour of a site; Two options: When agent P want to migrate on I , containing agent R
 - 1 freeze and retrieve the current content of the site, viz. R ;
check whether $P \mid R$ respects the policy of the site;
reactivate R and, according to the result of the checking phase, activate P .
 - 2 let membranes evolving at run-time: they are decreased with the privileges granted to P .
- I'm sure you see that the first option is just crazy. . .

Controlling Coalitions at a Site (ctd)

A new migration rule:

$$\begin{array}{l}
 k[M \triangleright \mathbf{go}_T I.P|Q] \quad || \quad l[M' \triangleright R] \\
 \rightarrow \quad k[M \triangleright Q] \quad || \quad l[M'' \triangleright P|R] \quad \text{if } M' \vdash_T^k P \succ M''
 \end{array}$$

where $M' \vdash_T^k P \succ M''$:

- verifies whether P respects M'_p (by examining its PCC T or its code, according to the trust level in its origin, k);
- if P respects M'_p , it decrease M'_p with the privileges granted to P . This returns M''_p

Controlling Coalitions at a Site (ctd2)

Well-formed systems are now defined w.r.t. a function Θ associating each good site to a initial policy.

$$\frac{}{\Theta \vdash I \llbracket M \Downarrow P \rrbracket : \mathbf{ok}} \quad \begin{array}{l} I \text{ good} \\ (pol(P) \sqcup M_p) \text{ enforces } \Theta(I) \end{array}$$

where

- $pol(P)$ returns the minimal policy satisfied by P ;
- \sqcup merges together two policies.

Subject Reduction: If $\Theta \vdash N : \mathbf{ok}$ and $N \rightarrow N'$, then $\Theta \vdash N' : \mathbf{ok}$.

Conclusions

- a formal framework to reason on the role of membranes as security policies
- several variations expressing finer and finer policies
- to be done:
 - a richer calculus (including communications, restrictions, ...)
 - more complex policies (not expressible with DFAs)
 - ...
- the paper is available at
www.dsi.uniroma1.it/~gorla/papers/GHS-membranes.ps