

SECURITY POLICIES AS MEMBRANES IN SYSTEMS FOR GLOBAL COMPUTING

DANIELE GORLA^a, MATTHEW HENNESSY^b, AND VLADIMIRO SASSONE^c

^a Dip. di Informatica, Univ. di Roma “La Sapienza”
e-mail address: gorla@di.uniroma1.it

^b Dept. of Informatics, Univ. of Sussex
e-mail address: matthewh@sussex.ac.uk

^c Dept. of Informatics, Univ. of Sussex
e-mail address: vs@sussex.ac.uk

ABSTRACT. We propose a simple global computing framework, whose main concern is code migration. Systems are structured in sites, and each site is divided into two parts: a computing body, and a *membrane* which regulates the interactions between the computing body and the external environment. More precisely, membranes are filters which control access to the associated site, and they also rely on the well-established notion of *trust* between sites. We develop a basic theory to express and enforce security policies via membranes. Initially, these only control the actions incoming agents intend to perform locally. We then adapt the basic theory to encompass more sophisticated policies, where the number of actions an agent wants to perform, and also their order, are considered.

1. INTRODUCTION

Computing is increasingly characterised by the global scale of applications and the ubiquity of interactions between mobile components. Among the main features of the forthcoming “global ubiquitous computing” paradigm we list *distribution* and *location awareness*, whereby code located at specific sites acts appropriately to local parameters and circumstances, that is, it is “context-aware”; *mobility*, whereby code is dispatched from site to site to increase flexibility and expressivity; *openness*, reflecting the nature of global networks and embodying the permeating hypothesis of localised, partial knowledge of the execution environment. Such systems present enormous difficulties, both technical and conceptual, and are currently more at the stage of exciting future prospectives than that of established of engineering practice. Two concerns, however, appear to clearly have a ever-reaching

2000 ACM Subject Classification: F.1.1, F.3.1, D.3.1, D.4.6.

Key words and phrases: Process calculi, Mobile code, Language-based security, Type systems.

^a This work has been mostly carried on while the first author was at the Dept. of Informatics, Univ. of Sussex, with a Marie Curie Fellowship.
^{a,b,c} The authors would like to acknowledge the support of the EC Global Computing projects Mikado and Myths.

import: *security* and *mobility control*, arising respectively from openness and from massive code and resource migrations. They are the focus of the present paper.

We aim at classifying mobile components according to their behaviour, and at empowering sites with control capabilities which allow them to deny access to those agents whose behaviour does not conform to the site's *policy*. We see every site of a system

$$k \llbracket M \triangleright P \rrbracket$$

as an entity named k and structured in two layers: a *computing body* P , where programs run their code – possibly accessing local resources offered by the site – and a *membrane* M , which regulates the interactions between the computing body and the external environment. An agent P wishing to enter a site l must be verified by the membrane before it is given a chance to execute in l . If the preliminary check succeeds, the agent is allowed to execute, otherwise it is rejected. In other words, a membrane implements the policy each site wants to enforce locally, by ruling on the requests of access of the incoming agents. This can be easily expressed by a migration rule of the form:

$$k \llbracket M^k \triangleright \mathbf{go} l.P \mid Q \rrbracket \parallel l \llbracket M^l \triangleright R \rrbracket \rightarrow k \llbracket M^k \triangleright Q \rrbracket \parallel l \llbracket M^l \triangleright P \mid R \rrbracket \quad \text{if } M^l \vdash^k P$$

The relevant parts here are P , the agent wishing to migrate from k to l , and l , the receiving site, which needs to be satisfied that P 's behaviour complies with its policy. The latter is expressed by l 's membrane, M^l . The judgement $M^l \vdash^k P$ represents l inspecting the incoming code to verify that it upholds M^l .

Observe that in the formulation above $M^l \vdash^k P$ represents a runtime check of all incoming agents. Because of our fundamental assumption of openendedness, such kind of checks, undesirable as they might be, cannot be avoided. In order to reduce their impact on systems performance, and to make the runtime semantics as efficient as possible, we adopt a strategy which allows for efficient agent verification. Precisely, we adopt an elementary notion of *trust*, so that from the point of view of each l the set of sites is consistently partitioned between “good,” “bad,” and “unknown” sites. Then, in a situation like the one in the rule above, we assume that l will be willing to accept from a *trusted* site k a *k-certified digest* T of P 's behaviour. We then modify the primitive **go** and the judgement \vdash^k as in the refined migration rule:

$$k \llbracket M^k \triangleright \mathbf{go}_{\mathsf{T}} l.P \mid Q \rrbracket \parallel l \llbracket M^l \triangleright R \rrbracket \rightarrow k \llbracket M^k \triangleright Q \rrbracket \parallel l \llbracket M^l \triangleright P \mid R \rrbracket \quad \text{if } M^l \vdash_{\mathsf{T}}^k P$$

The notable difference is in $M^l \vdash_{\mathsf{T}}^k P$. Here, l verifies the entire code P against M^l *only if* it does not trust k , the signer of P 's certificate T . Otherwise, it suffices for l to match M^l against the digest T carried by **go** together with P from k , so effectively shifting work from l to the originator of P .

Our main concern in this paper is to put the focus on the machinery a membrane should implement to enforce *different kinds* of policies. We first distill the simplest calculus which can conceivably convey our ideas and still support a non-trivial study. It is important to remark that we are abstracting from agents' local computations. These can be expressed in any of several well-known models for concurrency, for example CCS [Mil82] or the π -calculus [Mil99]. We are concerned, instead, with agents' migration from site to site: our main language mechanism is **go** rather than intra-site (i.e. local) communication. Using this language, we examine four notions of policy and show how they can be enforced by using membranes. We start with an amusingly simple policy which only lists allowed actions. We then move to count action occurrences and then to policies expressed by *deterministic finite automata*. Note that such policies are only concerned with the behaviour of single

agents, and do *not* take into account “*coalitional*” behaviours, whereby incoming agents – apparently innocent – join clusters of resident agents – they too apparently innocent – to perform cooperatively potentially harmful actions, or at least overrule the host site’s policy. We call *resident* those policies intended to be applied to the joint, composite behaviour of the agents contained at a site. We explore resident policies as our fourth and final notion of policy. In all the cases, the theory adapts smoothly; we only need to refine the information stored in the membrane and the inspection mechanisms.

Structure of the paper. In Section 2 we define the calculus used in this paper, and start with the straightforward policy which only prescribes the actions an agent can perform when running in a site. In Section 3, we enhance the theory to control also how many (and not only which kind of) actions an agent wants to perform in a site, and their order of execution. Finally, in Section 4 we extend the theory to control the overall computation taking place at a site, and not only the behaviour of single agents. The paper concludes in Section 5 where a comparison with related work is also given. The theoretical results are proved in Appendix A. With respect to the extended abstract [GHS04], this paper contains more examples together with complete proofs.

2. A SIMPLE CALCULUS

In this section we describe a simple calculus for mobile agents, which may migrate between sites. Each site is guarded by a *membrane*, whose task is to ensure that every agent accepted at the site conforms to an *entry policy*.

2.1. The Syntax.

The syntax is given in Figure 1 and assumes two pairwise disjoint sets: basic agent actions ACT , ranged over by a, b, c, \dots , and localities LOC , ranged over by l, k, h, \dots . Agents are constructed using the standard action-prefixing, parallel composition and replication operators from process calculi, [Mil82]. The one novel operator is that for migration,

$$\mathbf{go}_T l.P$$

This agent seeks to migrate to site l in order to execute the code P ; moreover it promises to conform to the entry policy T . In practical terms this might consist of a certification that the incoming code P conforms to the policy T , which the site l has to decide whether or not to accept. In our framework, this certification is a policy that describes the (local) behaviour of the agent; thus, in $\mathbf{go}_T l.P$, T will be called the *digest* of P .

A system consists of a finite set of sites running in parallel. A site takes the form

$$l \llbracket M \bowtie P \rrbracket$$

where

- l is the site name
- P is the code currently running at l
- M is the membrane which implements the entry policy.

For convenience we assume that site names are unique in systems. Thus, in a given system we can identify the membrane associated with the site named l by M^l . We start with a very simple kind of policy, which we will then progressively enhance.

<i>Basic Actions</i>	$a, b, c, \dots \in \text{ACT}$	
<i>Localities</i>	$l, h, k, \dots \in \text{LOC}$	
<i>Agents</i>	$P, Q, R ::=$	nil nil agent
		$a.P$ basic action
		$\text{go}_T l.P$ migration
		$P \mid Q$ composition
		$!P$ replication
<i>Systems</i>	$N ::=$	0 empty system
		$l \llbracket M \triangleright P \rrbracket$ site
		$N_1 \parallel N_2$ composition

Figure 1: A Simple Calculus

Definition 2.1 (Policies). A *policy* is a finite subset of $\text{ACT} \cup \text{LOC}$. For two policies T_1 and T_2 , we write

$$T_1 \text{ enforces } T_2$$

whenever $T_1 \subseteq T_2$. ■

Intuitively an agent conforms to a policy T at a given site if

- every action it performs at the site is contained in T
- it will only migrate to sites whose names are in T .

For example, conforming to the policy $\{\text{info}, \text{req}, \text{HOME}\}$, where **info**, **req** are actions and **HOME** a location, means that the only actions that will be performed are from the set $\{\text{info}, \text{req}\}$ and migration will only occur, if at all, to the site **HOME**. With this interpretation of policies, our definition of the predicate **enforces** is also intuitive; if some code P conforms to the policy T_1 and $T_1 \text{ enforces } T_2$ then P also automatically conforms to T_2 .

The purpose of membranes is to enforce such policies on incoming agents. In other words, at a site $l \llbracket M \triangleright Q \rrbracket$ wishing to enforce a policy T_{in} , the membrane M has to decide when to allow entry to an agent such as $\text{go}_T l.P$ from another site. There are two possibilities.

- The first is to syntactically check the code P against the policy T_{in} ; an implementation would actually expect the agent to arrive with a proof of this fact, and this proof would be checked.
- The second would be to *trust* the agent that its code P conforms to the stated T and therefore only check that this conforms to the entry policy T_{in} . Assuming that checking one policy against another is more efficient than the code analysis, this would make entry formalities much easier.

Deciding on when to apply the second possibility presupposes a *trust management* framework for systems, which is the topic of much current research. To simplify matters, here we simply assume that each site contains, as part of its membrane, a record of the level of trust it has in other sites. Moreover, we assume only three possible levels: **bad**, **unknown** and **good**. Intuitively, a site is **good**/**bad** if it behaves in a reliable/unreliable way, i.e. it does/doesn't properly calculate digests. On the other hand, a site tagged as **unknown** can

$$\begin{array}{lcl}
\text{(R-ACT)} & l\llbracket M \triangleright a.P \mid Q \rrbracket & \rightarrow l\llbracket M \triangleright P \mid Q \rrbracket \\
\text{(R-PAR)} & \frac{N_1 \rightarrow N'_1}{N_1 \parallel N_2 \rightarrow N'_1 \parallel N_2} & \\
\text{(R-STRUCT)} & \frac{N \equiv N_1 \quad N_1 \rightarrow N'_1 \quad N'_1 \equiv N'}{N \rightarrow N'} & \\
\text{(R-MIG)} & k\llbracket M^k \triangleright \mathbf{go}_\top l.P \mid Q \rrbracket \parallel l\llbracket M^l \triangleright R \rrbracket & \rightarrow \\
& k\llbracket M^k \triangleright Q \rrbracket \parallel l\llbracket M^l \triangleright P \mid R \rrbracket & \text{if } M^l \vdash_{\top}^k P
\end{array}$$

Figure 2: The reduction relation

$$\begin{array}{ll}
l\llbracket M \triangleright P \mid \mathbf{nil} \rrbracket \equiv l\llbracket M \triangleright P \rrbracket & N \parallel \mathbf{0} \equiv N \\
l\llbracket M \triangleright P \mid Q \rrbracket \equiv l\llbracket M \triangleright Q \mid P \rrbracket & N_1 \parallel N_2 \equiv N_2 \parallel N_1 \\
l\llbracket M \triangleright (P \mid Q) \mid R \rrbracket \equiv l\llbracket M \triangleright P \mid (Q \mid R) \rrbracket & (N_1 \parallel N_2) \parallel N_3 \equiv N_1 \parallel (N_2 \parallel N_3) \\
l\llbracket M \triangleright !P \mid Q \rrbracket \equiv l\llbracket M \triangleright P \mid !P \mid Q \rrbracket &
\end{array}$$

Figure 3: The structural equivalence

behave in a non specified way; thus, for the sake of security, it will be considered as **bad**. In a more realistic scenario, it would be possible to refine **unknown** to either **good** or **bad**, upon collection of enough evidence to consider it reliable or not. For the sake of simplicity, we do not model this framework here.

Definition 2.2 (Membranes). A membrane M is a pair (M_t, M_p) where

- M_t is a partial function from LOC to $\{\mathbf{unknown}, \mathbf{good}, \mathbf{bad}\}$
- M_p is a policy

■

2.2. The Operational Semantics.

Having defined both *policies* and *membranes*, we now give an operational semantics for the calculus, which formalises the above discussion on how to manage agent migration. This is given as a binary relation $N \rightarrow N'$ over systems; it is defined to be the least relation which satisfies the rules in Figure 2. Rule (R-ACT) says that the agent $a.P$ running in parallel with other code in site l , such as Q , can perform the action a ; note that the semantics does not record the occurrence of a . (R-PAR) and (R-STRUCT) are standard. The first allows reductions within parallel components, while the second says that reductions are relative to a structural equivalence; the rules defining this equivalence are given in Figure 3. The interesting reduction rule is the last one, (R-MIG), governing migration; the agent $\mathbf{go}_\top l.P$ can migrate from site k to site l provided the predicate $M^l \vdash_{\top}^k P$ is true. This ‘enabling’ predicate formalises our discussion above on the role of the membrane M^l , and requires in

$$\begin{array}{c}
\text{(TC-EMPTY)} \\
\frac{}{\vdash \mathbf{nil} : \mathbb{T}}
\end{array}
\qquad
\begin{array}{c}
\text{(TC-ACT)} \\
\frac{\vdash P : \mathbb{T}}{\vdash a.P : \mathbb{T}} \quad a \in \mathbb{T}
\end{array}
\qquad
\begin{array}{c}
\text{(TC-MIG)} \\
\frac{\vdash P : \mathbb{T}'}{\vdash \mathbf{go}_{\mathbb{T}'} l.P : \mathbb{T}} \quad l \in \mathbb{T}
\end{array}$$

$$\begin{array}{c}
\text{(TC-REPL)} \\
\frac{\vdash P : \mathbb{T}}{\vdash !P : \mathbb{T}}
\end{array}
\qquad
\begin{array}{c}
\text{(TC-PAR)} \\
\frac{\vdash P : \mathbb{T} \quad \vdash Q : \mathbb{T}}{\vdash P \mid Q : \mathbb{T}}
\end{array}$$

Figure 4: Typechecking incoming agents

turn a notion of code P satisfying a policy \mathbb{T} ,

$$\vdash P : \mathbb{T}$$

With such a notion, we can then define $M^l \vdash_{\mathbb{T}}^k P$ to be:

$$\text{if } M_t^l(k) = \text{good then } (\mathbb{T} \text{ enforces } M_p^l) \text{ else } \vdash P : M_p^l \quad (2.1)$$

In other words, if the target site l trusts the source site k , it trusts that the professed policy \mathbb{T} is a faithful reflection of the behaviour of the incoming agent P , and then entry is gained provided that \mathbb{T} enforces the entry policy M_p^l (i.e., in this case, $\mathbb{T} \subseteq M_p^l$). Otherwise, if k can not be trusted, then the entire incoming code P has to be checked to ensure that it conforms to the entry policy, as expressed by the predicate $\vdash P : M_p^l$.

In Figure 4 we describe a simple inference system for checking that agents conform to policies, i.e. to infer judgements of the form $\vdash P : \mathbb{T}$. Rule (TC-EMPTY) simply says that the empty agent \mathbf{nil} satisfies all policies. (TC-ACT) is also straightforward; $a.P$ satisfies a policy \mathbb{T} and if a is allowed by \mathbb{T} , and the residual P satisfies \mathbb{T} . The rule (TC-PAR) says that to check $P \mid Q$ it is sufficient to check P and Q separately, and similarly for replicated agents. The most interesting rule is (TC-MIG), which checks $\mathbf{go}_{\mathbb{T}'} l.P$. This not only checks that migration to l is allowed by the policy, that is $l \in \mathbb{T}$, but it also checks that the code to be spawned there, P , conforms to the associated professed policy \mathbb{T}' . In some sense, if the agent $\mathbf{go}_{\mathbb{T}'} l.P$ is allowed entry into a site k , then k assumes responsibility for any promises that it makes about conformance to policies.

2.3. Safety.

We have just outlined a reduction semantics in which sites seek to enforce policies either by directly checking the code of incoming agents against entry policies, or more simply by checking the professed policy of trusted agents. The extent to which this strategy works depends, not surprisingly, on the quality of a site's trust management.

Example 2.1. Let HOME be a site name with the following trust function

$$M_t^h : \{\text{ALICE, BOB, SECURE}\} \mapsto \text{good}.$$

Consider the system

$$N \triangleq \text{HOME} \llbracket M^h \triangleright P^h \rrbracket \parallel \text{BOB} \llbracket M^b \triangleright P^b \rrbracket \parallel \text{ALICE} \llbracket M^a \triangleright P^a \rrbracket \parallel \text{SECURE} \llbracket M^s \triangleright P^s \rrbracket$$

in which the entry policy of HOME, M_p^h , is $\{\text{info}, \text{req}, \text{SECURE}\}$, and that of SECURE, M_p^s , is $\{\text{give}, \text{HOME}\}$. Since $M_t^h(\text{BOB}) = \text{good}$, agents migrating from BOB to HOME are trusted and only their digests are checked against the entry policy M_p^h . So, if P^b contains the agent

$$\text{go}_{T_1} \text{HOME} . (\text{take} . Q)$$

where T_1 enforces M_p^h , then the entry policy of HOME will be transgressed.

As another example, suppose ALICE, again trusted by HOME, contains the agent

$$\text{go}_{T_1} \text{HOME} . (\text{info} . \text{go}_{T_2} \text{SECURE} . (\text{take} . Q))$$

where T_2 is some policy which enforces the entry policy of SECURE, M_p^s . Again because T_1 enforces M_p^h , the migration is allowed from ALICE to HOME, and moreover the incoming agent conforms to the policy demanded of HOME. The second migration of the agent is also successful if SECURE trusts HOME: $M_t^s(\text{HOME}) = \text{good}$ and therefore only the digest T_2 is checked against the entry policy of SECURE. We then have the reduction

$$N \rightarrow^* \text{HOME}[\dots] \parallel \text{BOB}[\dots] \parallel \text{ALICE}[\dots] \parallel \text{SECURE}[\![M^s \triangleright \text{take} . Q \mid P^s]\!]$$

in which now the entry policy of SECURE has been foiled. ■

The problem in this example is that the trust knowledge of HOME is faulty; it trusts in sites which do not properly ensure that professed policies are enforced. Let us divide the sites into *trustworthy* and otherwise. This bipartition could be stored in an external record stating which nodes are trustworthy (i.e. typechecked) and which ones are not. However, for economy, we prefer to record this information in the membranes, by demanding that the trust knowledge at trustworthy sites is a proper reflection of this division. This is more easily defined if we assume the following ordering over trust levels:

$$\text{unknown} <: \text{bad} \quad \text{and} \quad \text{unknown} <: \text{good}$$

This reflects the intuitive idea that sites classified as **unknown** may, perhaps with further information, be subsequently classified either as **good** or **bad**. On the other hand, **good** or **bad** cannot be further refined; sites classified as either, will not be reclassified.

Definition 2.3 (Trustworthy sites and Coherent systems). In a system N , the site k is *trustworthy* if $M_t^k(k) = \text{good}$. N is *coherent* if $M_t^k(l) <: M_t^l(l)$ for every trustworthy site k . ■

Thus, if a trustworthy site k believes that a site l can be trusted (i.e., $M_t^k(l) = \text{good}$), then l is indeed trustworthy (as represented by $M_t^l(l) = \text{good}$). Similarly, if it believes l to be **bad**, then l is indeed bad. The only uncertainty is when k classifies l as **unknown**: then l may be either **good** or **bad**. Of course, in coherent systems we expect sites which have been classified as trustworthy to act in a *trustworthy manner*, which amounts to saying that code running at such a k must have at one time gained entry there by satisfying the entry policy. Note that by using policies as in Definition 2.1, if P satisfies an entry policy M_p^k , then it continues to satisfy the policy while running at k (cf. Theorem 2.2 below).

$$\begin{array}{c}
\text{(WF-EMPTY)} \\
\vdash \mathbf{0} : \mathbf{ok}
\end{array}
\qquad
\begin{array}{c}
\text{(WF-G.SITE)} \\
\frac{\vdash P : M_p}{\vdash l[M \triangleright P] : \mathbf{ok}} \quad l \text{ trustworthy}
\end{array}$$

$$\begin{array}{c}
\text{(WF-PAR)} \\
\frac{\vdash N_1 : \mathbf{ok}, \quad \vdash N_2 : \mathbf{ok}}{\vdash N_1 \parallel N_2 : \mathbf{ok}}
\end{array}
\qquad
\begin{array}{c}
\text{(WF-U.SITE)} \\
\frac{}{\vdash l[M \triangleright P] : \mathbf{ok}} \quad l \text{ not trustworthy}
\end{array}$$

Figure 5: Well-formed systems

This property of coherent systems, which we call *well-formedness*, can therefore be checked syntactically. In Figure 5, we give the set of rules for deriving the judgement

$$\vdash N : \mathbf{ok}$$

of well-formedness of N . There are only two interesting rules. Firstly, (WF-G.SITE) says that $l[M \triangleright P]$ is well-formed whenever l is trustworthy and $\vdash P : M_p$. There is a subtlety here; this not only means that P conforms to the policy M_p , but also that any digests proffered by agents in P can also be trusted. The second relevant rule is (WF-U.SITE), for typing unknown sites: here there is no need to check the resident code, as agents emigrating from such sites will not be trusted.

Example 2.2. (*Example 2.1 continued.*) Let us now re-examine the system N in Example 2.1. Suppose HOME is trustworthy, that is $M_t^h(\text{HOME}) = \text{good}$. Then, if N is to be coherent, it is necessary for each of the sites BOB, ALICE and SECURE also to be trustworthy. Consequently, N cannot be well-formed. For example, to derive $\vdash N : \mathbf{ok}$ it would be necessary to derive

$$\vdash \mathbf{go}_{T_1} \text{HOME}(\mathbf{take}.Q) : M_p^b$$

where M_p^b is the entry policy of BOB. But this requires the judgement

$$\vdash \mathbf{take}.Q : T_1$$

where T_1 enforces M_p^h . Since $\mathbf{take} \notin M_p^h$, this is not possible.

One can also check that the code running at ALICE stops the system from being well-formed. Establishing $\vdash N : \mathbf{ok}$ would also require the judgement

$$\vdash \mathbf{go}_{T_1} \text{HOME}(\mathbf{info}.\mathbf{go}_{T_2} \text{SECURE}(\mathbf{take}.Q)) : M_p^a$$

which in turn, eventually, requires

$$\vdash \mathbf{take}.Q : T_2$$

for some T_2 such that T_2 enforces M_p^s ; this is impossible, again because \mathbf{take} is not in M_p^s . ■

In well-formed systems we know that entry policies have been respected. So one way of demonstrating that our reduction strategy correctly enforces these policies is to prove that

- system well-formedness is preserved by reduction
- only legal computations take place within trustworthy sites

Figure 6: A Labelled Transition System

☐

With any collection of policies, endowed with two such relations, we can define the predicate $M \vdash_k^T P$ as in (2.1) above, and thereby get a reduction semantics for the calculus. In this section we investigate two variations on the notion of entry policies and discuss the extent to which we can prove that the reduction strategy correctly implements them.

3.1. Multisets as Entry Policies.

The policies of the previous section only express the legal actions agents may perform at a site. However in many situations more restrictive policies are desirable. To clarify this point, consider the following example.

Example 3.1. Let MAIL_SERV be the site name of a mail server with the following entry policy M_p^{ms} :

$$\{\text{list}, \text{send}, \text{retr}, \text{del}, \text{reset}, \text{quit}\}$$

The server accepts client agents performing requests for listing mail messages, sending/retrieving/deleting messages, resetting the mailbox and quitting. Now, consider the system

$$N \triangleq \text{MAIL_SERV} \llbracket M^{ms} \triangleright P^{ms} \rrbracket \parallel \text{SPAM} \llbracket M^s \triangleright \text{go}_{\top \text{MAIL_SERV}}.(!\text{send}) \rrbracket$$

where $\top = \{\text{send}\}$. According to the typechecking of Figure 4, we have that

$$\vdash !\text{send} : M_p^{ms}$$

However, the agent is a spamming virus and, in practical implementations, should be rejected by MAIL_SERV. ■

In such scenarios it would be more suitable for policies to be able to fix an upper-bound over the number of messages sent. This can be achieved in our setting by changing policies from sets of agent actions to *multisets* of actions. Consequently, predicate **enforces** is now multiset inclusion.

First let us fix some notation. We can view a multiset as a set equipped with an *occurrence function*, that associates a natural number to each element of the set. To model permanent resources, we also allow the occurrence function to associate ω to an element with an infinite number of occurrences in the multiset. Notationally, e^ω stands for an element e occurring infinitely many times in a multiset. This notation is extended to sets and multisets; for any set/multiset E , we let E^ω to denote the multiset $\{e^\omega : e \in E\}$.

Example 3.2. (*Example 3.1 continued.*) Coming back to Example 3.1, it would be sufficient to define M_p^{ms} to be $\{\dots, \text{send}^K, \dots\}$ where K is a reasonable constant. In this way, an agent can only send at most K messages in each session; if it wants to send more messages, it has to disconnect from MAIL_SERV (i.e. leave it) and then reconnect again (i.e. immigrate again later on). In practice, this would prevent major spamming attacks, because the time spent for login/logout operations would radically slow down the spam propagation. ■

The theory presented in Sections 2.2 and 2.3 can be adapted to the case where policies are multisets of actions. The judgment $\vdash P : \top$ is redefined in Figure 7, where operator \cup stands for *multiset union*. The key rules are (TC-ACT), (TC-PAR) and (TC-REPL). The first two properly decrease the type satisfied when typechecking sub-agents. The third one is needed because recursive agents can be, in general, freely unfolded; hence, the actions they intend to locally perform can be iterated arbitrarily many times. For instance, agent

$$P \triangleq !\text{send}$$

$$\begin{array}{c}
\text{(TC-EMPTY)} \\
\vdash \mathbf{nil} : \mathbb{T}
\end{array}
\qquad
\begin{array}{c}
\text{(TC-ACT)} \\
\vdash P : \mathbb{T} \\
\hline
\vdash a.P : \mathbb{T} \cup \{a\}
\end{array}
\qquad
\begin{array}{c}
\text{(TC-MIG)} \\
\vdash P : \mathbb{T}' \\
\hline
\vdash \mathbf{go}_{\mathbb{T}', l}.P : \mathbb{T} \cup \{l\}
\end{array}$$

$$\begin{array}{c}
\text{(TC-PAR)} \\
\vdash P : \mathbb{T}_1 \quad \vdash Q : \mathbb{T}_2 \\
\hline
\vdash P \mid Q : \mathbb{T}_1 \cup \mathbb{T}_2
\end{array}
\qquad
\begin{array}{c}
\text{(TC-REPL)} \\
\vdash P : \mathbb{T} \\
\hline
\vdash !P : \mathbb{T}'
\end{array}
\quad \mathbb{T}^\omega \text{ enforces } \mathbb{T}'$$

Figure 7: Typechecking with policies as Multisets

satisfies policy $\mathbb{T} \triangleq \{\mathbf{send}^\omega\}$. Notice that the new policy satisfaction judgement prevents the spamming virus of Example 3.1 from typechecking against the policy of MAIL_SERV defined in Example 3.2.

The analysis of the previous section can also be repeated here but an appropriate notion of *well-formed* system is more difficult to formulate. The basic problem stems from the difference between *entry* policies and *resident* policies. The fact that all agents who have ever entered a site l respects an entry policy \mathbb{M}_p gives no guarantees as to whether the joint effect with the code currently occupying the site l also satisfies \mathbb{M}_p . For instance, in the terms of Example 3.2, MAIL_SERV ensures that each incoming agent can only send at most K messages. Nevertheless, two such agents, having gained entry and now running concurrently at MAIL_SERV, can legally send – jointly – up to $2K$ messages. It is therefore necessary to formulate *well-formedness* in terms of the individual threads of the code currently executing at a site. Let us say P is a *thread* if it is not of the form $P_1 \mid P_2$. Note that every agent P can be written in the form of $P_1 \mid \dots \mid P_n, n \geq 1$, where each P_i is a thread. So the well-formedness judgment is modified by replacing rule (WF-G.SITE_M) in Figure 5 as below.

$$\begin{array}{c}
\text{(WF-G.SITE}_M\text{)} \\
\forall i. (P_i \text{ a thread and } \vdash P_i : \mathbb{M}_p) \\
\hline
\vdash l \llbracket M \Downarrow P_1 \mid \dots \mid P_n \rrbracket : \mathbf{ok}
\end{array}
\quad l \text{ trustworthy}$$

Theorem 3.1 (Subject Reduction for multiset policies). If $\vdash N : \mathbf{ok}$ and $N \rightarrow N'$, then $\vdash N' : \mathbf{ok}$.

Proof. Similar to that of Theorem 2.1. The necessary changes are outlined in Appendix A.2. □

The statement of safety must be changed to reflect the focus on individual threads rather than agents. Moreover, we must keep into account also multiple occurrences of actions in a trace; thus, we let $\mathbf{act}(\sigma)$ return a multiset formed by all the actions occurring in σ .

Theorem 3.2 (Safety for multiset policies). Let N be a well-formed system. Then, for every trustworthy site $l \llbracket M \Downarrow P_1 \mid \dots \mid P_n \rrbracket$ in N , where each P_i is a thread, $P_i \xrightarrow{\sigma} P'_i$ implies that $\mathbf{act}(\sigma) \text{ enforces } \mathbb{M}_p$.

Proof. See Appendix A.2. □

3.2. Finite Automata as Entry Policies.

A second limitation of the setting presented in Section 2 is that policies will sometimes need to prescribe a precise order for executing legal actions. This is very common in client/server interactions, where a precise protocol (i.e. a pattern of message exchange) must be respected. To this end we define policies as *deterministic finite automata* (DFAs, for short).

Example 3.3. Let us consider Example 3.1 again. Usually, mail servers requires a preliminary authentication phase to give access to mail services. To express this fact, we could implement the entry policy of MAIL_SERV, M_p^{ms} , to be the automaton associated to the regular expression below.

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

The server accepts client requests only upon authentication, via a username/password mechanism. Moreover, the policy imposes that each session is regularly committed by requiring that each sequence of actions is terminated by `quit`. This could be needed to save the status of the transaction and avoid inconsistencies. ■

We now give the formal definitions needed to adapt the theory developed in Section 2. We start by defining a DFA, the language associated to it, the **enforces** predicate between DFAs and a way for an agent to satisfy a DFA. As usual [HU79], a DFA is a quintuple $A \triangleq (S, \Sigma, s_0, F, \delta)$ where

- S is a finite set of *states*;
- Σ is the *input alphabet*;
- $s_0 \in S$ is a reserved state, called the *starting state*;
- $\emptyset \subset F \subseteq S$ is the set of *final states* (also called *accepting states*);
- $\delta : S \times \Sigma \rightarrow S$ is the *transition relation*.

In our framework, the alphabet of the DFAs considered is a finite subset of $\text{ACT} \cup \text{LOC}$. Moreover, for the sake of simplicity, we shall always assume that the DFAs in this paper are minimal.

Definition 3.1 (DFA Acceptance and Enforcement). Let A be a DFA. Then

- $Acp_s(A)$ contains all the $\sigma \in \Sigma^*$ such that σ leads A from state s to a final state;
- $Acp(A)$ is defined to be $Acp_{s_0}(A)$;
- $A_1 \text{ enforces } A_2$ holds true whenever $Acp(A_1) \subseteq Acp(A_2)$. ■

Notice that, as expected, there is an efficient way to establish $A_1 \text{ enforces } A_2$, once given the automata A_1 and A_2 (see Proposition A.2 in Appendix A.3). We now formally describe the language associated to an agent by exploiting the notion of *concurrent regular expressions* (CRE, for short) introduced in [GR92] to model concurrent processes. For our purposes, the following subset of CRE suffices:

$$e ::= \epsilon \mid \alpha \mid e_1.e_2 \mid e_1 \odot e_2 \mid e^\otimes$$

ϵ denotes the empty sequence of characters, α ranges over $\text{ACT} \cup \text{LOC}$, ‘.’ denotes concatenation, \odot is the interleaving (or *shuffle*) operator and $^\otimes$ is its closure. Intuitively, if e represents the language L , then e^\otimes represents $\{\epsilon\} \cup L \cup L \odot L \cup L \odot L \odot L \dots$. Given a CRE

e , the language associated to it, written $lang(e)$, can be easily defined; a formal definition is recalled in Appendix A.3. Now, given a process P , we easily define a CRE associated to it. Formally

$$\begin{aligned} \text{CRE}(\mathbf{nil}) &\triangleq \epsilon & \text{CRE}(a.P) &\triangleq a.\text{CRE}(P) \\ \text{CRE}(\mathbf{go}_A l.P) &\triangleq l & \text{CRE}(P_1 | P_2) &\triangleq \text{CRE}(P_1) \odot \text{CRE}(P_2) \\ \text{CRE}(!P) &\triangleq \text{CRE}(P)^\otimes \end{aligned}$$

Definition 3.2 (DFA Satisfaction). An agent P satisfies the DFA A , written $\vdash P : A$, if $lang(\text{CRE}(P)) \subseteq \text{Acp}(A)$, and $\vdash Q : A'$ holds for every subagent of P of the form $\mathbf{go}_{A'} l.Q$. ■

In Proposition A.2, we prove that DFA satisfaction is decidable, although extremely hard to establish. This substantiate our hypothesis that verifying digests is preferable to inspecting the full code from the point of view computational complexity. We are now ready to state the soundness of this variation. It simply consists in finding a proper notion of well-formed systems. As in Section 3.1, the entry policy can only express properties of single threads, instead of coalitions of threads hosted at a site. Thus, we modify rule (WF-G.SITE) from Figure 5 to:

$$\frac{(\text{WF-G.SITE}_A) \quad \forall i. P_i \text{ a thread and } \exists s \in S. lang(\text{CRE}(P_i)) \subseteq \text{Acp}_s(\mathbf{M}_p)}{\vdash l \llbracket M \rrbracket P_1 | \dots | P_n \rrbracket : \mathbf{ok}} \quad l \text{ trustworthy}$$

This essentially requires that the languages associated to each of the threads in l are suffixes of words accepted by \mathbf{M}_p (cf. Theorem 3.4 below). Since this may appear quite weak, it is worth remarking that the well-formedness predicate is just a ‘consistency’ check, a way to express that the agent is in a state from where it will respect the policy of l . The soundness theorems are reported below and are proved in Appendix A.3.

Theorem 3.3 (Subject Reduction for automata policies). If $\vdash N : \mathbf{ok}$ and $N \rightarrow N'$, then $\vdash N' : \mathbf{ok}$.

Theorem 3.4 (Safety for automata policies). Let N be a well-formed system. Then, for every trustworthy site $l \llbracket M \rrbracket P_1 | \dots | P_n \rrbracket$ in N , where each P_i is a thread, $\sigma \in lang(\text{CRE}(P_i))$ implies that there exists some $\sigma' \in \text{Acp}(\mathbf{M}_p)$ such that $\sigma' = \sigma''\sigma$, for some σ'' .

We conclude this section with two interesting properties enforceable by using automata.

Example 3.4. [Lock/Unlock] We have two actions, **lock** and **unlock**, with the constraint that each **lock** must be always followed by an **unlock**. Let $\Sigma_1 = \Sigma - \{\mathbf{lock}\}$ and $\Sigma_u = \Sigma - \{\mathbf{lock}, \mathbf{unlock}\}$. Thus, the desired policy (written using a regular expression formalism) is

$$(\Sigma_1^* . (\mathbf{lock} . \Sigma_u^* . \mathbf{unlock})^*)^* \quad \blacksquare$$

Example 3.5. [Secrecy] Let **secret** be a secret action; we require that, whenever an agent performs **secret**, it cannot migrate anymore (this policy enforces that agents having performed **secret** always remain co-located). Let $\Sigma_s = \Sigma - \{\mathbf{secret}\}$ and $\Sigma_g = \Sigma - \text{LOC}$; thus, the desired policy is

$$\Sigma_s^*.(\epsilon + \mathbf{secret}.\Sigma_g^*) \quad \blacksquare$$

4. RESIDENT POLICIES

Here we change the intended interpretation of policies. In the previous section a policy dictated the proposed behaviour of an agent *prior* to execution in a site, at the point of entry. This implied that safety in well-formed systems was a thread-wise property (see rules (WF-G.SITE_M) and (WF-G.SITE_A)). Here we focus on policies which are intended to describe the permitted (coalitional) behaviour of agents during execution at a site. Nevertheless these resident policies are still used to determine whether a new agent is allowed access to the site in question; entry will only be permitted if the addition of this incoming agent to the code currently executing at the site does not violate the policy.

Let us consider an example to illustrate the difference between entry and resident policies.

Example 4.1. Let LICENCE_SERV be the site name of a server that makes available K licences to download and install a software product. The distribution policy is based on a queue: the first K agents landing in the site are granted the licence, the following ones are denied. The policy of the server should be $M_p^s \triangleq \{\mathbf{get_licence}^K\}$. However if this policy is interpreted as an entry policy, applying the theory of Section 3.1, then the system grants at most K licences to *each* incoming agent. Moreover this situation continues indefinitely, effectively handing out licences to all incoming agents. \blacksquare

We wish to re-interpret the policies of the previous section as *resident policies* and here we outline two different schemes for enforcing such policies. For simplicity we confine our attention to one kind of policy, that of multisets.

4.1. Static membranes.

Our first scheme is conservative in the sense that many of the concepts developed in Section 3.1 for entry policies can be redeployed. Let us reconsider rule (R-MIG) from Figure 2. There, the membrane M^l only takes into consideration the incoming code P , and its digest T , when deciding on entry, via the predicate $M^l \vdash_T^k P$. But if the membrane is to enforce a *resident policy*, then it must also take into account the contribution of the code already running in l , namely R . To do so we need a mechanism for *joining* policies, such as those of the incoming P and the resident R in rule (R-MIG). So let us assume that the set of policies, with the relation **enforces** is a partial order in which every pair of elements T_1 and T_2 has a *least upper bound*, denoted $T_1 \sqcup T_2$. For multiset policies this is the case as \sqcup is simply multiset union. In addition we need to be able to calculate the (minimal) policy

$$\begin{array}{c}
\text{(TI-EMPTY)} \\
\frac{}{\vdash \mathbf{nil} : \emptyset}
\end{array}
\qquad
\begin{array}{c}
\text{(TI-ACT)} \\
\frac{\vdash P : \mathsf{T}}{\vdash a.P : \mathsf{T} \cup \{a\}}
\end{array}
\qquad
\begin{array}{c}
\text{(TI-MIG)} \\
\frac{\vdash P : \mathsf{T}'}{\vdash \mathbf{go}_{\mathsf{T}}^l.P : \{l\}} \quad \mathsf{T}' \text{ enforces } \mathsf{T}
\end{array}$$

$$\begin{array}{c}
\text{(TI-REPL)} \\
\frac{\vdash P : \mathsf{T}}{\vdash !P : \mathsf{T}^\omega}
\end{array}
\qquad
\begin{array}{c}
\text{(TI-PAR)} \\
\frac{\vdash P : \mathsf{T}_1 \quad \vdash Q : \mathsf{T}_2}{\vdash P \mid Q : \mathsf{T}_1 \cup \mathsf{T}_2}
\end{array}$$

Figure 8: Type inference for agents with policies as multisets

which a process R satisfies; let us denote this as $\text{pol}(R)$. For multiset policies we can adjust the rules in Figure 7, essentially by eliminating *weakening*, to perform this calculation; the resulting rules are given in Figure 8, with judgements of the form $\vdash P : \mathsf{T}$.

Lemma 4.1.

- For every P , there is at most one T such that $\vdash P : \mathsf{T}$.
- $\vdash P : \mathsf{T}$ implies that there exists some policy T' such that T' **enforces** T and $\vdash P : \mathsf{T}'$.

Proof. The first statement is proved by structural induction on P ; the second by induction on the derivation $\vdash P : \mathsf{T}$. \square

Definition 4.2. Define the partial function $\text{pol}(\cdot)$ by letting $\text{pol}(P)$ be the unique policy such that $\vdash P : \mathsf{T}$, if it exists. \blacksquare

With these extra concepts we can now change rule (R-MIG) in Figure 2 to take the current resident code into account. It is sufficient to change the side condition, from $M^l \vdash_{\mathsf{T}}^k P$ to $M^l, R \vdash_{\mathsf{T}}^k P$, where this latter is defined to be

$$\text{if } M_t^l(k) = \text{good then } (\mathsf{T} \sqcup \text{pol}(R)) \text{ enforces } M_p^l \text{ else } \vdash P \mid R : M_p^l$$

Here if only the digest needs to be checked then we compare $\mathsf{T} \sqcup \text{pol}(R)$, that is the result of *adding* the digest to the policy of the resident code R , against the resident policy M_p^l . On the other hand if the source site is untrusted we then need to analyse the incoming code in parallel with the resident code R . It should be clear that the theory developed in Section 3.1 is readily adapted to this revised reduction semantics. In particular the Subject Reduction and Safety theorems remain true; we spare the reader the details. However it should also be clear that this approach to enforcing resident policies has serious practical drawbacks. An implementation would need to:

- (1) freeze and retrieve the current content of the site, namely the agent R ;
- (2) calculate the minimal policy satisfied by R to be merged with P 's digest in order to check the predicate **enforces**, or typecheck the composed agent $P \mid R$;
- (3) reactivate R and, according to the result of the checking phase, activate P .

Even if the language were equipped with a ‘passivation’ operator, as in [SS03], the overall operation would still be computationally very intensive. Consequently we suggest below another approach.

$\frac{}{\Theta \vdash l \llbracket M \triangleright P \rrbracket : \mathbf{ok}} \quad \begin{array}{l} l \text{ trustworthy} \\ (\text{pol}(P) \sqcup M_p) \text{ enforces } \Theta(l) \end{array}$	$\frac{}{\Theta \vdash \mathbf{0} : \mathbf{ok}} \quad \text{(WF-EMPTY)}$
$\frac{}{\Theta \vdash l \llbracket M \triangleright P \rrbracket : \mathbf{ok}} \quad l \text{ not trustworthy} \quad \text{(WF-U.SITE)}$	$\frac{\Theta \vdash N_1 : \mathbf{ok}, \quad \Theta \vdash N_2 : \mathbf{ok}}{\Theta \vdash N_1 \parallel N_2 : \mathbf{ok}} \quad \text{(WF-PAR)}$

Figure 9: Well-formed systems under Θ

4.2. Dynamic membranes.

In the previous approach we have to repeatedly calculate the policy of the current resident code each time a new agent requests entry. Here we allow the policy in the membrane to “*decrease*,” in order to reflect the resources already allocated to the resident code. So at any particular moment in time the policy currently in the membrane records what resources *remain*, for any future agents who may wish to enter; with the entry of each agent there is a corresponding decrease in the membrane’s policy. Formally we need to change the migration rule rule (R-MIG) to one which not only checks incoming code, or digest, against the membrane’s policy, but also updates the membrane:

$$\begin{array}{l} \text{(R-MIG')} \quad k \llbracket M^k \triangleright \mathbf{go}_{\top} l.P \mid Q \rrbracket \parallel l \llbracket M^l \triangleright R \rrbracket \rightarrow \\ \quad k \llbracket M^k \triangleright Q \rrbracket \parallel l \llbracket \widehat{M}^l \triangleright P \mid R \rrbracket \quad \text{if } M^l \vdash_{\top}^k P \succ \widehat{M}^l \end{array}$$

where the judgement $M^l \vdash_{\top}^k P \succ \widehat{M}^l$ is defined as

$$\text{let } T' = \left\{ \begin{array}{ll} T & \text{if } M_t^l(k) = \mathbf{good} \\ \text{pol}(P) & \text{otherwise} \end{array} \right\} \text{ in } T' \text{ enforces } M_p^l \wedge M_p^l = \widehat{M}_p^l \sqcup T' \wedge M_t^l = \widehat{M}_t^l$$

First notice that if this migration occurs then the membrane at the target site changes, from M_p^l to \widehat{M}_p^l . The latter is obtained from the former by eliminating those resources allocated to the incoming code P . If the source site, k , is deemed to be **good** this is calculated via the incoming digest T ; otherwise a direct analysis of the code P is required, to calculate $\text{pol}(P)$.

This revised schema is more reasonable from an implementation point of view, but its soundness is more difficult to formalise and prove. As a computation proceeds no permanent record is kept in the system of the original resident policies at the individual sites. Therefore well-formedness can only be defined relative to an external record of what the resident policies were, when the system was initiated. For this purpose we use a function Θ , mapping trustworthy sites to policies; it is sufficient to record the original policies at these sites as we are not interested in the behaviour elsewhere.

Then we can define the notion of well-formed systems, relative to such a Θ ; this is written as $\Theta \vdash N : \mathbf{ok}$ and the formal definition is given in Table 9. The crucial rule is (WF-G.SITE), for trustworthy sites. If l is such a site then $l \llbracket M \triangleright P \rrbracket$ is well-formed relative to the original record Θ if $M_p^l \sqcup \text{pol}(P)$ guarantees the original resident policy at l , namely $\Theta(l)$.

Theorem 4.3 (Subject Reduction for resident policies). If $\Theta \vdash N : \mathbf{ok}$ and $N \rightarrow N'$, then $\Theta \vdash N' : \mathbf{ok}$.

Proof. Outlined in Appendix A.4. □

The introduction of these external records of the original resident policies also enables us to give a Safety result.

Theorem 4.4 (Safety for resident policies). Let N be a well-formed system w.r.t. Θ . Then, for every trustworthy site $l \llbracket M \triangleright P \rrbracket$ in N , $P \xrightarrow{\sigma} P'$ implies that $\mathbf{act}(\sigma)$ enforces $\Theta(l)$.

Proof. See Appendix A.4. □

5. CONCLUSION AND RELATED WORK

We have presented a framework to describe distributed computations of systems involving migrating agents. The activity of agents entering/running in ‘good’ sites is constrained by a membrane that implements the layer dedicated to the security of the site. We have described how membranes can enforce several interesting kind of policies. The basic theory presented for the simpler case has been refined and tuned throughout the paper to increase the expressiveness of the framework. Clearly, any other kind of behavioural specification of an agent can be considered a policy. For example, a promising direction could be considering logical frameworks (by exploiting model checking or proof checkers).

The calculus we have presented is very basic: it is even simpler than CCS [Mil82], as no synchronization can occur. Clearly, we did not aim at Turing-completeness, but at a very basic framework in which to focus on the rôle of membranes. We conjecture that, by suitably advancing the theory presented here, all the ideas can be lifted to more complex calculi (including, e.g., synchronization, value passing and/or name restriction).

Related Work. In the last decade, several calculi for distributed systems with code mobility have appeared in literature. In particular, structuring a system as a (flat or hierarchical) collection of named sites introduced the possibility of dealing with sophisticated concrete features. For example, sites can be considered as the unity of *failure* [FG+96, Ama00], *mobility* [FG+96, CG00] or *access control* [HR02, RH03, GP03]. The present work can be seen as a contribution to the last research line.

As in [GP03], we have presented a scenario where membranes can evolve. However, the membranes presented in Section 4 only describe ‘what is left’ in the site. On the other hand, the (dynamically evolving) type of a site in [GP03] always constrains the overall behaviour of agents in the site and it is modified upon acquisition/loss of privileges through computations.

We borrowed from [RH03] the notion of *trust* between sites, where agents coming from trusted sites are accepted without any control. Here, we relaxed this choice by examining the digest of agents coming from trusted sites. Moreover, we have a fixed net of trust; we believe that, once communication is added to our basic framework, the richer scenario of [RH03] (where the partial knowledge of a site can evolve during its computation) can be recovered.

A related paper is [IK01], where authors develop a *generic* type system that can be smoothly instantiated to enforce several properties of the π -calculus (dealing with arity

mismatch in communications, deadlock, race control and linearity). They work with one kind of type, and modify the subtyping relation in order to yield several relevant notions of safety. The main difference with our approach is that we have different kind of types (and, thus, different type checking mechanisms) for all the variations we propose. It would be nice to lift our work to a more general framework closer to theirs; we leave this for future work.

Our work is also related to [NR05]. Policies are described there as deterministic finite automata and constrain the access to critical sections in a concurrent functional language. A type and effect system is provided that guarantees adherence of systems to the policy. In particular, the sequential behaviour of each thread is guaranteed to respect the policy, and the interleavings of the threads' locks to be safe. But unlike our paper [NR05] has no code migration, and no explicit distribution; thus, only one centralised policy is used.

Membranes as filters between the computing body of a site and the external environment are also considered in [FMP04, Bou04, SS03]. There, membranes are computationally capable objects, and can be considered as a kind of process. They can evolve and communicate both with the outer and with the inner part of the associated node, in order to regulate the life of the node. This differs from our conception of membranes as simple tools for the verification of incoming agents.

To conclude, we remark that our understanding of membranes is radically different from the concept of policies in [ES99]. Indeed, in *loc. cit.*, security automata control the execution of agents running in a site by *in-lined monitoring*. This technique consists of accepting incoming code unconditionally, but blocking at runtime those actions not abiding the site policy. Clearly, in order to implement the strategy, the execution of each action must be filtered by the policy. This contrasts with our approach, where membranes are 'containers' that regulate the interactions between sites and their environments. The computation taking place within the site is out of the control of the membrane, which therefore cannot rely on in-lined monitoring.

ACKNOWLEDGEMENT

The authors wish to acknowledge the reviewers of this paper for their positive attitude and for their fruitful comments. Joanna Jedrzejowicz kindly answered some questions on regular languages with interleaving and iterated interleaving.

REFERENCES

- [Ama00] R. Amadio. On modelling mobility. *Theoretical Computer Science*, 240(1):147–176, 2000.
- [Bou04] G. Boudol. A generic membrane model. In *Proc. of Global Computing*, volume 3267 of *LNCS*, pages 208–222. Springer, 2004.
- [Bou98] Z. Bouziane. A primitive recursive algorithm for the general Petri net reachability problem. In *Proc. of FOCS'98*, pages 130–136. IEEE, 1998.
- [CG00] L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.
- [ES99] U. Erlingsson and F. Schneider. SASI Enforcement of Security Policies: A Retrospective. In *Proc. of New Security Paradigms Workshop*, pages 87–95. ACM, 1999.
- [FMP04] G. Ferrari, E. Moggi, and R. Pugliese. MetaKlaim: a type safe multi-stage language for global computing. *Mathematical Structures in Computer Science*, 14(3):367–395, 2004.
- [FG+96] C. Fournet, G. Gonthier, J. Lévy, L. Maranget, and D. Rémy. A calculus of mobile agents. In *Proc. of CONCUR'96*, volume 1119 of *LNCS*, pages 406–421. Springer, 1996.
- [GR92] V. Garg and M. Raghunath. Concurrent regular expressions and their replationship to Petri nets. *Theoretical Computer Science*, 96:285–304, 1992.

- [GHS04] D. Gorla and M. Hennessy and V. Sassone. Security Policies as Membranes in Systems for Global Computing. In *Proc. of FGUC'04*, ENTCS. Elsevier, 2004.
- [GP03] D. Gorla and R. Pugliese. Resource access and mobility control with dynamic privileges acquisition. In *Proc. of ICALP'03*, volume 2719 of *LNCS*, pages 119–132. Springer-Verlag, 2003.
- [HR02] M. Hennessy and J. Riely. Resource Access Control in Systems of Mobile Agents. *Information and Computation*, 173:82–120, 2002.
- [HU79] J. Hopcroft and J. Ullman. *Introduction to automata theory, languages and computation*. Addison-Wesley, 1979.
- [IK01] A. Igarashi and N. Kobayashi. A generic type system for the pi-calculus. In *Proceedings of POPL '01*, pages 128–141. ACM, 2001.
- [May84] E. Mayr. An algorithm for the general Petri net reachability problem. *SIAM Journal of Computing*, 13(3):441–460, 1984.
- [Mil82] R. Milner. *A Calculus for Communicating Systems*. Springer-Verlag, 1982.
- [Mil99] R. Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999.
- [NR05] N. Nguyen and J. Rathke. Typed Static Analysis for Concurrent, Policy-Based, Resource Access Control. Draft, 2005.
- [Pet81] J. Peterson. *Petri Net Theory and Modeling of Systems*. Prentice Hall, 1981.
- [RH03] J. Riely and M. Hennessy. Trust and partial typing in open systems of mobile agents. *Journal of Automated Reasoning*, 31:335–370, 2003.
- [SS03] A. Schmitt and J. Stefani. The M-calculus: a higher-order distributed process calculus. In *Proc. of POPL'03*, pages 50–61. ACM, 2003.

APPENDIX A. TECHNICAL PROOFS

We now outline the proofs of the technical results in the paper, section by section.

A.1. Proofs of Section 2.

Lemma A.1 (Subsumption). If $\vdash P : T$ and T enforces T' , then $\vdash P : T'$.

Proof. By induction on the derivation of the judgment $\vdash P : T$. □

Proof of Theorem 2.1 [Subject Reduction]: The proof is by induction over the inference of $N \rightarrow N'$. Notice that trustworthiness is invariant under reduction. Therefore coherence, which is defined in terms of the trustworthiness of sites, is also preserved by reduction.

We outline the proof when the inference is deduced using rule (R-MIG), a typical example. By hypothesis, $\vdash k \llbracket M^k \Downarrow \mathbf{go}_T l.P \mid Q \rrbracket : \mathbf{ok}$; this implies that $\vdash k \llbracket M^k \Downarrow Q \rrbracket : \mathbf{ok}$. Thus, we only need to prove that $\vdash l \llbracket M^l \Downarrow R \rrbracket : \mathbf{ok}$ and $M^l \vdash_T^k P$ imply $\vdash l \llbracket M^l \Downarrow P \mid R \rrbracket : \mathbf{ok}$. We have two possible situations:

- l **trustworthy:** Judgment $\vdash R : M_p^l$ holds by hypothesis; judgment $\vdash P : M_p^l$ is implied by $M^l \vdash_T^k P$. Indeed, because of the coherence hypothesis, $M_t^l(k) <: M_t^k(k)$. If $M_t^k(k) \neq \mathbf{good}$, then $M^l \vdash_T^k P$ is exactly the required $\vdash P : M_p^l$. Otherwise, we know that $\vdash \mathbf{go}_T l.P : M_p^k$; by rule (TC-MIG) this implies that $\vdash P : T$. Judgment $\vdash P : M_p^l$ is obtained by using Lemma A.1, since $M^l \vdash_T^k P$ is defined to be T enforces M_p^l (see (2.1) in Section 2.2). Thus, by using (TC-PAR), we obtain the desired $\vdash P \mid R : M_p^l$.
- l **not trustworthy:** This case is simple, because rule (WF-U.SITE) always allows to derive $\vdash l \llbracket M^l \Downarrow P \mid R \rrbracket : \mathbf{ok}$.

The case when (R-ACT) is used is similar, although simpler, and the case when rule (R-PAR) is used requires a simple inductive argument. Finally to prove the case when rule (R-STRUCT) is used, we need to know that *coherency* of systems is preserved by structural equivalence; the proof of this fact, which is straightforward, is left to the reader. \square

Proof of Theorem 2.2 [Safety]: Let $l\llbracket M \rrbracket P$ be a site in N such that $P \xrightarrow{\sigma} P'$. We have to prove that $\text{act}(\sigma)$ enforces M_p . The statement is proved by induction over the length of σ . The base case, when $\sigma = \epsilon$, is trivial since $\text{act}(\epsilon) = \emptyset$.

So we may assume $\sigma = \alpha\sigma'$ and $P \xrightarrow{\alpha} P'' \xrightarrow{\sigma'} P'$. Let us consider $P \xrightarrow{\alpha} P''$; by induction on $\xrightarrow{\alpha}$, we can prove that $\alpha \in M_p$ and that $\vdash l\llbracket M \rrbracket P'' : \text{ok}$. If the transition has been inferred by using rule (LTS-ACT), then $P = a.P''$ and, by rule (WF-G.SITE), we have that $\vdash a.P'' : M_p$; by definition of rule (TC-ACT), we have the desired $a \in M_p$ and $\vdash P'' : M_p$. When (LTS-MIG) is used the argument is similar, and all other cases follow in a straightforward manner by induction.

Thus, we can now apply induction on the number of actions performed in $P'' \xrightarrow{\sigma'} P'$ and obtain that $\text{act}(\sigma')$ enforces M_p . This suffices to conclude that $\text{act}(\sigma) = (\text{act}(\sigma') \cup \{\alpha\})$ enforces M_p . \square

A.2. Proofs of Section 3.1.

The proofs given in Appendix A.1 can be easily adapted to the setting in which entry policies are multisets. We outline only the main changes. First recall that **enforces** is multiset inclusion, that judgments $\vdash P : T$ must be now inferred by using the rules in Figure 7 and that rule (WF-G.SITE_M) is used for well-formedness. Then, Lemma A.1 remains true in this revised setting.

Proof of Theorem 3.1 [Subject Reduction]: A straightforward adaptation of the corresponding proof in the previous section. The only significant change is to the case when a replication is unfolded via the rule (R-STRUCT), i.e.

$$N \triangleq l\llbracket M \rrbracket !P \mid Q \equiv l\llbracket M \rrbracket P \mid !P \mid Q \rightarrow N'' \equiv N'$$

By hypothesis, $\vdash !P : M_p$; therefore, by definition of rule (TC-REPL), we have that $\vdash P : T$ for some T such that T^ω enforces M_p . Since T enforces T^ω and because of Lemma A.1, we have that $\vdash l\llbracket M \rrbracket P \mid !P \mid Q : \text{ok}$. By induction, $\vdash N'' : \text{ok}$. It is easy to prove that this suffices to obtain the desired $\vdash N' : \text{ok}$. \square

Proof of Theorem 3.2 [Safety]: From the rule (WF-G.SITE_M) we know that $\vdash P_i : M_p$, for all $i = 1, \dots, n$. We now proceed by induction over $|\sigma|$. The base case is trivial. For the inductive case, we consider $\sigma = \alpha\sigma'$ and $P_i \xrightarrow{\alpha} P_i'' \xrightarrow{\sigma'} P_i'$. By induction on $\xrightarrow{\alpha}$, we can prove that $\alpha \in M_p$ and that $\vdash l\llbracket M_i; M_p - \{\alpha\} \rrbracket P_i''$. If the transition has been inferred by using rule (LTS-ACT), then $P_i = a.P_i''$ and, by rule (WF-G.SITE_M), we have that $\vdash a.P_i'' : M_p$; by definition of rule (TC-ACT), we have the desired $M_p = T \cup \{a\}$ and $\vdash P_i'' : T$. When (LTS-MIG) is used the case is simpler, and all other cases follow in a straightforward manner by induction.

Coming back to the main claim, we use the induction and obtain that $\text{act}(\sigma')$ **enforces** $M_p - \{\alpha\}$; thus, $\text{act}(\sigma)$ **enforces** M_p . \square

A.3. Proofs of Section 3.2.

We start by recalling from [GR92] the formal definition of the language associated to a CRE, as follows.

$$\begin{aligned}
 \text{lang}(\epsilon) &\triangleq \{\epsilon\} \\
 \text{lang}(\alpha) &\triangleq \{\alpha\} \\
 \text{lang}(e_1.e_2) &\triangleq \{x_1x_2 : x_1 \in \text{lang}(e_1) \wedge x_2 \in \text{lang}(e_2)\} \\
 \text{lang}(e_1 \odot e_2) &\triangleq \{x_1y_1 \cdots x_ny_n : x_1 \cdots x_n \in \text{lang}(e_1) \wedge y_1 \cdots y_n \in \text{lang}(e_2)\} \\
 \text{lang}(e^{\otimes}) &\triangleq \bigcup_{i \geq 0} \text{lang}(e)^{\otimes i} \quad \text{where } L^{\otimes i} \triangleq \begin{cases} \{\epsilon\} & \text{if } i = 0 \\ L^{\otimes i-1} \odot L & \text{otherwise} \end{cases}
 \end{aligned}$$

Notice that the definition of the $\text{lang}(e_1 \odot e_2)$ hides a trick: the x_i s and the y_i s can also be ϵ . Thus, as expected, we can also consider for interleaving strings of different length.

We start by accounting on the complexity of predicate **enforces** and the satisfiability relation when policies are automata. This is stated by the following Proposition.

Proposition A.2.

- (1) A_1 **enforces** A_2 can be calculated in polynomial time
- (2) $\vdash P : A$ is decidable, but it is super-exponential

Proof.

- (1) Let $A_i = (S_i, \Sigma, s_0^i, F_i, \delta_i)$ and let $L_i = \text{Acp}(A_i)$. By definition, we have to check whether $L_1 \subseteq L_2$ or not. This is equivalent to check whether $L_1 \cap \overline{L_2} = \emptyset$. The following steps have been carried out by following [HU79].
 - (a) calculate the automaton associated to $\overline{L_2}$. This can be done in $O(|S_2|)$ and the resulting automaton has $|S_2|$ states.
 - (b) calculate the automaton associated to $L_1 \cap \overline{L_2}$. This can be done in $O(|S_1| \times |S_2| \times |\Sigma|)$ and creates an automaton A with $|S_1| \times |S_2|$ states.
 - (c) Checking the emptiness of $L_1 \cap \overline{L_2}$ can be done by using a breath-first search that starts from the starting state of (the graph underlying) A and stops whenever a final state is reached. If no final state is reached, $L_1 \cap \overline{L_2}$ is empty. This can be done in $O(|S_1| \times |S_2| \times |\Sigma|)$.

Thus, the overall complexity is $O(|S_1| \times |S_2| \times |\Sigma|)$.

- (2) It has been proved in [GR92] that each CRE e can be represented by a (labelled) Petri net, in that the language accepted by the Petri net is $\text{lang}(e)$. Now, we can easily construct a DFA accepting the complement of the language accepted by A (see item (a) of the previous proof). Now, we can construct the product between this DFA (that can be seen as a Petri net) and the Petri net associated to $\text{CRE}(P)$; this Petri net accepts $\text{lang}(\text{CRE}(P)) \cap \text{Acp}(A)$ (see [Pet81]). Now, the emptiness of this language can be solved with the algorithm for the reachability problem in the corresponding Petri net. This problem has been proved decidable [May84] and solvable in double-exponential time [Bou98]. \square

We now prove the subject reduction theorem in the setting where types are DFAs. To this aim, we need to adapt Lemma A.1 and we need a very simple result on the languages associated to DFAs and processes.

Lemma A.3. If $\vdash P : A$ and A enforces A' , then $\vdash P : A'$.

Proof. By transitivity of subset inclusion. \square

Lemma A.4.

- (1) $\alpha\sigma \in Acp_s(A)$ if and only if $\sigma \in Acp_{\delta(s,\alpha)}(A)$
- (2) If $\sigma \in lang(CRE(a.P))$ then $\sigma = a\sigma'$ for $\sigma' \in lang(CRE(P))$. Viceversa, if $\sigma \in lang(CRE(P))$, then $a\sigma \in lang(CRE(a.P))$.

Proof. Trivial. \square

Proof of Theorem 3.1 [Subject Reduction]: Now $\vdash N : \mathbf{ok}$ relies on rule (WF-G.SITE_A). Again, the proof is by induction on the inference of $N \rightarrow N'$. We only give the base cases, because inductive steps can be handled with in a standard way. We only consider the cases of trustworthy sites; the case for non-trustworthy sites is easier. In what follows, we write $\vdash_s P : A$ to mean that $\vdash P : A'$, where A' is the DFA obtained from A by setting s as starting state.

(R-ACT): In this case, $N = l[M \bowtie a.P \mid Q]$. By definition of rule (WF-G.SITE_A), it holds that $Q = Q_1 \mid \dots \mid Q_k$ (for Q_i threads), $\exists s. \vdash_s a.P : M_p$ and $\forall i. \exists s_i. \vdash_{s_i} Q_i : M_p$. By definition, we have that $lang(CRE(a.P)) \subseteq Acp_s(M_p)$; by Lemma A.4, we have that $lang(CRE(P)) \subseteq Acp_{\delta(s,a)}(M_p)$. This suffices to infer the well-formedness of $N' = l[M \bowtie P \mid Q]$.

(R-MIG): In this case, $N = k[M^k \bowtie \mathbf{go}_A l.P \mid Q] \parallel l[M^l \bowtie R]$ and $M^l \vdash_A^k P$. We further identify two sub-cases:

- $M^l(k) = \mathbf{good}$. In this case, because of coherence, we know that $\vdash P : A$. Moreover, by definition of $M^l \vdash_A^k P$, it holds that A enforces M_p^l . By Lemma A.3, we have that $\vdash P : M_p^l$. This suffices to conclude.
- $M^l(k) \neq \mathbf{good}$. This case is simpler because $M^l \vdash_A^k P$ is defined to be $\vdash P : M_p^l$.

\square

Proof of Theorem 3.4 [Safety]: The proof is quite easy. Indeed, by rule (WF-G.SITE_A), it holds that $\exists s_i : \vdash_{s_i} P_i : M_p^l$. By definition, this implies that every $\sigma \in lang(CRE(P_i))$ is in $Acp_{s_i}(M_p^l)$. Since the automaton M_p^l is minimal, s_i is a reachable state from the starting state s_0 , say, with a (finite) string σ'' . Then, by Definition 3.1 and by Lemma A.4.1, it holds that $\sigma''\sigma \in Acp(M_p^l)$. This proves the thesis. \square

A.4. Proofs of Section 4.

We show here the main things to modify to carry out the proofs given in Appendix A.2. Obviously, judgment $\vdash P : T$ must be now replaced everywhere with $\Vdash P : T$ and, similarly, $\vdash N : \mathbf{ok}$ becomes $\Theta \vdash N : \mathbf{ok}$.

Proof of Theorem 4.3 [Subject Reduction]: The proof is by induction over the inference of $N \rightarrow N'$. Inductive steps are simple; we only give the base steps.

(R-ACT): By hypothesis, $\Theta \vdash l \llbracket M \triangleright a.P \mid Q \rrbracket : \mathbf{ok}$. If l is not trustworthy, the case is trivial. Otherwise, we know by hypothesis that $(\text{pol}(a.P \mid Q) \sqcup M_p)$ **enforces** $\Theta(l)$. Now, by definition of judgment \Vdash (and hence of function $\text{pol}(\cdot)$) we have that $\text{pol}(a.P \mid Q) = \text{pol}(P \mid Q) \cup \{a\}$. Hence, $(\text{pol}(P \mid Q) \sqcup M_p)$ **enforces** $\Theta(l)$, as required.

(R-MIG): By hypothesis, $\Theta \vdash l \llbracket M^l \triangleright R \rrbracket : \mathbf{ok}$; we only consider the case in which l is trustworthy. Thus, we know that $(\text{pol}(R) \sqcup M_p^l)$ **enforces** $\Theta(l)$. By the premise of rule (R-MIG), it holds that $M^l \vdash_{\top}^k P \succ \widehat{M}^l$. We have two possible situations:

$M_t^l(k) = \mathbf{good}$: In this case, $M^l \vdash_{\top}^k P \succ \widehat{M}^l$ is defined to be \top **enforces** $M_p^l \wedge M_p^l = \widehat{M}_p^l \sqcup \top \wedge M_t^l = \widehat{M}_t^l$. The fact that $M_t^l = \widehat{M}_t^l$ is sufficient to preserve coherence. Moreover, by rule (TI-MIG), we know that $\Vdash P : \top'$ and \top' **enforces** \top . By rule (TI-PAR), $\text{pol}(P \mid R) = \text{pol}(R) \sqcup \top'$ and $(\text{pol}(R) \sqcup \top')$ **enforces** $(\text{pol}(R) \sqcup \top)$. Then, $\text{pol}(P \mid R) \sqcup \widehat{M}_p^l = (\text{pol}(R) \sqcup \top' \sqcup \widehat{M}_p^l)$ **enforces** $(\text{pol}(R) \sqcup \top \sqcup \widehat{M}_p^l) = (\text{pol}(R) \sqcup M_p^l)$ **enforces** $\Theta(l)$, as required.

$M_t^l(k) \neq \mathbf{good}$: In this case, the previous proof should be rephrased by using $\text{pol}(P)$ instead of the digest \top . □

Proof of Theorem 4.4 [Safety]: We prove a slightly more general result, that easily implies the claim desired.

Let N be a well-formed system w.r.t. Θ . If $l \llbracket M \triangleright P \rrbracket$ is a trustworthy site of N such that $(\text{pol}(P) \sqcup M_p^l) = \top$, then $P \xrightarrow{\sigma} P'$ implies that $\text{act}(\sigma)$ **enforces** \top .

The proof is by induction over $|\sigma|$. The base case is when $\sigma = \epsilon$ and it is trivial. In the inductive case, we consider $\sigma = \alpha\sigma'$ and $P \xrightarrow{\alpha} P'' \xrightarrow{\sigma'} P'$. To start, it is easy to prove that

$$\text{pol}(P'') \sqcup \{\alpha\} = \text{pol}(P) \tag{A.1}$$

By transitivity of multiset inclusion and by the claim (A.1) above, $(\text{pol}(P'') \sqcup M_p^l) = \top'$, where $\top = \top' \sqcup \{\alpha\}$. Thus, node $l \llbracket M^l \triangleright P'' \rrbracket$ is well-formed (and trustworthy). By induction we therefore have that $\text{act}(\sigma')$ **enforces** \top' . Hence, $\text{act}(\sigma) = \text{act}(\sigma') \sqcup \{\alpha\}$ **enforces** $\top' \sqcup \{\alpha\} = \top$, as required.

To conclude, the original claim of Theorem 4.4 is obtained from the result just proved by noticing that, because of well-formedness, \top **enforces** $\Theta(l)$. □