

Enforcing Security Policies via Types^{*}

Daniele Gorla Rosario Pugliese

Dipartimento di Sistemi e Informatica, Università di Firenze

e-mail: {gorla,pugliese}@dsi.unifi.it

Abstract. Security is a key issue for distributed systems/applications with code mobility, like, e.g., e-commerce and on-line bank transactions. In a scenario with code mobility, traditional solutions based on cryptography cannot deal with all security issues and additional mechanisms are necessary. In this paper, we present a flexible and expressive type system for security for a calculus of distributed and mobile processes. The type system has been designed to supply real systems security features, like the assignment of different privileges to users over different data/resources. Type soundness is guaranteed by using a combination of static and dynamic checks, thus enforcing specific security policies on the use of resources. The usefulness of our approach is shown by modeling the simplified behaviour of a bank account management system.

1 Introduction

Code mobility is a fundamental aspect of global computing; however it gives rise to a lot of relevant security problems like, e.g., *secrecy* and *integrity* of data and program code. Indeed, in mobile distributed systems/applications, other than attacks to inter-process communication over the communication channels (e.g. traffic analysis, message modifications/forging), several other kinds of attacks could take place. For instance, malicious mobile processes can attempt to access private information, or modify private data of the nodes hosting them. Hence, a server receiving a mobile process for execution needs to impose strong requirements to ensure that the incoming process does not violate the secrecy and jeopardize the integrity of the information. Similarly, mobile processes need tools to ensure that their execution at the server node does not compromise their integrity (e.g. modification of process code) or secrecy (e.g. leak of sensible data). Such problems have increasingly importance due to the spreading of security critical applications, like, e.g., electronic commerce and on-line bank transactions. Moreover, global computing environments, like e.g. the Internet, are highly dynamic and open systems. In these environments static information could be partial, inaccurate or missing, therefore for ensuring security properties a certain amount of dynamic checks is needed (e.g. mobile agents should be dynamically checked at run-time when they migrate).

^{*} This work has been partially supported by EU within the FET - Global Computing initiative, project MIKADO IST-2001-32222, and by MIUR project NAPOLI. The funding bodies are not responsible for any use that might be made of the results presented here.

Code mobility strongly restricts a safe use of cryptography, that is one of the most used techniques for ensuring security in distributed systems. In fact, because of attacks like those mentioned before, we can hardly imagine to use mobile processes carrying confidential data (e.g. private keys) with them, or host nodes with classified information accessible to all incoming processes (whatever their source node be). Hence, the use of security mechanisms that back up and supplement cryptographic mechanisms becomes a major issue when developing systems of distributed and mobile processes where the compliance with some security policies must be guaranteed.

Several alternative approaches have been exploited to enforce security policies in distributed computing systems. The approaches may differ in the level of trust required, the flexibility of the enforced security policy and their costs to components producers and users. A comprehensive security framework could result from the combination of complementary features. Approaches like *code signing* and *sand-boxing* (for instance, consider the Java implementation of these concepts [23, 20]) have low costs but cannot enforce flexible security policies (signed components may behave in arbitrary ways and the user must trust the component producer, while sand-boxed components are isolated and cannot interact with each other).

Type systems can be sensible and flexible language-based security techniques, like [32] shows. Recently, a number of process/programming languages supporting process distribution and mobility have been designed that come equipped with type systems that guarantee some kind of security properties, see, e.g., [24, 15, 16, 25, 8]. However, to the best of our knowledge, the type system we present in this paper is the first that exploits the source of mobile processes for granting them different privileges over different kinds of data (thus, e.g., preventing dangerous operations over specific sensible data). These desirable features can be found in real systems like, e.g., UNIX, where different users can have different privileges and different files can be manipulated with different allowed operations.

Our type system permits expressing and enforcing security policies for controlling the access of host resources by possibly malicious mobile processes. It is expressly designed for the process calculus μ KLAIM [21] that puts forward a programming paradigm where there is a clear separation between the programmer level and the net coordinator/administrator level. Programmers write processes, while coordinators write nets, hence manage the initial distribution of processes and set the security policies for accessing the resources. The policies are specified by assigning each node of a net a type expressing the operations a process is allowed to perform once spawned at it. Hence types are part of the language for configuring the underlying net architecture and must be taken into account in the language operational semantics. Other than to express security policies, types are used to record processes intended operations, but programmers are relieved from typing processes because this task is carried on by a static type inference system. By using a combination of static and dynamic type checking, our system guarantees the absence of run-time errors due to lack of privileges. As an application of our approach we model the simplified behaviour of a bank

$N ::= \mathbf{0}$	(empty net)	$a ::= \mathbf{read}(T)@l$	(process actions)
$l ::=^\Delta P$	(single node)	$\mathbf{in}(T)@l$	
$N_1 \parallel N_2$	(net composition)	$\mathbf{out}(t)@l$	
$P ::= \mathbf{nil}$	(null process)	$\mathbf{eval}(P)@l$	
$a.P$	(action prefixing)	$\mathbf{newloc}(u : \Delta)$	
$P_1 \mid P_2$	(parallel composition)	$T ::= F \quad F, T$	(templates)
A	(process invocation)	$F ::= f \quad !x \quad !u : \pi$	(template fields)
$e ::= V \quad x \quad \dots$	(expressions)	$t ::= f \quad f, t$	(tuples)
		$f ::= e \quad l$	(tuple fields)

Table 1. μ KLAIM Syntax

account management system where the compliance with the bank security policy must be enforced.

The rest of the paper is organized as follows. We present the syntax of μ KLAIM in Section 2, its type system in Section 3, and its operational semantics in Section 4, that also contains the type soundness results. In Section 5, we illustrate an application of our approach to model a bank account management system. Finally, in Section 6 we point out a few concluding remarks and comment on related work. Due to lack of space, in this extended abstract we omit some technical details and all proofs; they can be found in the full paper [22].

2 The Process Language μ KLAIM

In this section we briefly present the syntax and informally describe the semantics of μ KLAIM [21], a calculus to program distributed and mobile processes communicating asynchronously via shared data. Due to lack of space, some formal aspects (akin to those presented in [21]) are omitted.

The syntax of μ KLAIM is reported in Table 1. We assume the existence of the following countable sets: \mathcal{A} , *process identifiers*, ranged over by A, B, \dots ; \mathcal{L} , *localities*, ranged over by l ; \mathcal{U} , *locality variables*, ranged over by u ; \mathcal{V} , *basic values*, ranged over by V . We let ℓ to range over $\mathcal{L} \cup \mathcal{U}$, x over value variables, π over sets of *capabilities* and Δ over *types* (capabilities and types are formally defined in Section 3).

The syntax of *expressions*, ranged over by e , is deliberately not specified; we just assume that expressions contain, at least, basic values and variables. *Localities* l are the addresses of nodes. *Tuples* t are sequences of actual fields f , that contain information items (expressions, localities or locality variables). Tuples are collected into multisets called *tuple spaces* (TSs, for short). *Templates* T are used to select tuples in a TS; they are sequences of actual and formal fields F . The latter are used to bind variables to values and are written $!x$ or $!u : \pi$ (the set of capabilities π constraints the use of the address dynamically bound to u and is crucial for the type checking).

Processes are built up from the inactive process **nil** and from the basic operations by using prefixing, parallel composition and process invocation. For the

sake of simplicity, we assume that each process identifier A has a *single* defining equation $A \triangleq P$ and all these equations are available at any locality of a net. Recursive behaviours can be modelled via process definitions.

μ KLAIM supplies five different basic operations, also called *actions*. $\mathbf{out}(t)@l$ adds the tuple et resulting from the evaluation¹ of t to the TS located at l . The presence of the evaluated tuple et in the TS at l is represented by putting in parallel with the process located at l the auxiliary process $\mathbf{out}(et)$. Operation $\mathbf{eval}(Q)@l$ sends process Q for execution to l , where a run-time typechecking of the incoming code will take place: if Q does not comply with l 's security policy the operation is blocked. Operation $\mathbf{in}(T)@l$ evaluates T and looks for a matching² tuple et in the TS located at l ; if et is found, it is withdrawn and the values it contains are used to replace the corresponding variables of T within the continuation process, otherwise the operation is suspended until a matching et is available. Operation \mathbf{read} behaves similarly but leaves the accessed tuple et in the tuple space. Operation $\mathbf{newloc}(u : \Delta)$ dynamically creates a new net node with a fresh address whose security policy is specified by type Δ . The last operation is not indexed with an address because it always acts locally; all the other operations explicitly indicate the (possibly remote) address where they will take place.

Nets are finite collections of nodes where processes and tuple spaces can be allocated. A *node* is a triple $l ::^\Delta P$, where locality l is the address (i.e. network reference) of the node, P is the (parallel) process located at l and Δ is the type of the node, i.e. the specification of its access control policy. The nodes of a net can be thought of both as physically distributed machines and as logical partitions of the same machine. As we already said, the TS located at l is part of P because evaluated tuples are semantically represented as special processes.

3 A Capability-Based Type System

In this section we introduce a type system for μ KLAIM that permits granting different privileges to processes coming from different nodes and constraining the operations allowed over different kinds of data. Thus, for example, if l trusts l' , then l security policy could accept processes coming from l' (that will be called l' -processes) and let them accessing any tuple in l 's TS. If l' is not totally trusted, then l 's security policy could grant l' -processes the capabilities for executing $\mathbf{in/read}$ only over tuples that do not contain classified data.

3.1 Capabilities

Capabilities are used to specify the allowed process operations and are formally defined as

¹ Tuple/template evaluation consists in replacing each expression with the value resulting from its evaluation.

² An evaluated tuple matches against an evaluated template if both have the same number of fields and corresponding fields match; formal fields of a given type match values of the same type and two values match only if identical.

$\{\langle i, p \rangle\} \sqsubseteq_{\Pi} \{\langle r, p \rangle\}$	$\frac{\mathcal{T}(p') \subseteq \mathcal{T}(p)}{\{\langle c, p \rangle\} \sqsubseteq_{\Pi} \{\langle c, p' \rangle\}}$	$\frac{\pi_1 \subseteq \pi_2}{\pi_2 \sqsubseteq_{\Pi} \pi_1}$	$\frac{\pi_1 \sqsubseteq_{\Pi} \pi'_1 \quad \pi_2 \sqsubseteq_{\Pi} \pi'_2}{\pi_1 \cup \pi_2 \sqsubseteq_{\Pi} \pi'_1 \cup \pi'_2}$

Table 2. Capability Ordering Rules

$$\mathcal{C} \triangleq \{e, n\} \cup \{ \langle c, p \rangle : c \in \{i, r, o\} \wedge p \subseteq_{\text{fin}} \mathcal{P} \}$$

where $\mathcal{P} \triangleq (\mathcal{L} \cup \mathcal{V} \cup \{\mathbf{from}, -\})^+$ is the set of all *patterns*. Capabilities e and n enable process migration and node creation (i.e. operations **eval** and **newloc**, resp.). A capability of the form $\langle c, p \rangle$ enables the operation whose name's first character is c (i.e. **in** if c is i , and so on); operation arguments must comply with the finite set of patterns p if $p \neq \emptyset$, and are not restricted otherwise (in this case, we write c instead of $\langle c, \emptyset \rangle$). Like tuples and templates, *patterns* are finite, not empty sequences of fields; *pattern fields* may be localities, basic values, the reserved word **from** (denoting the last locality visited by a mobile process) and the ‘don't care’ symbol $-$ (denoting any template field). Thus, for instance, the capability $\langle i, \{(\text{“public”}, -), (3, -, \mathbf{from})\} \rangle$ enables the operations **in**(“public”, !x)@... and **in**(3, !u : π, l)@... for an l -process, while disables operation **in**(“private”, !x)@... .

We use π to denote a non-empty subset of \mathcal{C} such that, if $\langle c, p \rangle \in \pi$ and $\langle c', p' \rangle \in \pi$, then $c \neq c'$. Π will denote the set of all these π .

We say that a template *complies with* a pattern if the template is obtained by replacing in the pattern all occurrences of **from** with a locality, and any occurrence of ‘ $-$ ’ with any template field allowed by the syntax. Given a non-empty set of patterns p , we write $\mathcal{T}(p)$ to denote the set of all templates complying with patterns in p . By definition, $\mathcal{T}(\emptyset)$ denotes the set of all templates. Since tuples are also templates (see Table 1), the previous definitions also apply to tuples.

Notice that the definition of pattern fields affects, via the relation ‘complies with’, the ability of our types to control the tuples accessed by process operations. However, our framework is largely independent of the choice of a specific set of fields. For instance, we could also permit fields of the form $-\delta$, for any type δ of legal values, with the idea that, when defining the relation ‘complies with’, an occurrence of $-\delta$ could be replaced by any value/variable of type δ . In μKLAIM , this corresponds to adding only fields $-\mathcal{L}$ and $-\mathcal{V}$; in this way, a finer control could be exercised on the tuples accessed by processes because we could distinguish between a tuple field containing a locality from one containing a basic value.

We now introduce an ordering between capabilities, \sqsubseteq_{Π} ; formally, it is the least reflexive and transitive relation induced by the rules in Table 2. The chosen ordering relies on the following assumptions: (i) if a process is allowed to perform an **in** then it is also allowed to perform a **read** over the same arguments, (ii) if a process is allowed to perform a **read/in/out** over arguments complying with patterns in p then it is allowed to perform the same operation over arguments complying with any set of patterns p' that has at most the same ‘complying templates’ as p , and (iii) if a process owns a set of capabilities π_2 then it also owns any subset π_1 .

3.2 Types

Types, ranged over by Δ , are functions of the form

$$\Delta : \mathcal{L} \cup \mathcal{U} \cup \{\mathbf{any}\} \rightarrow_{\text{fin}} \left((\mathcal{L} \cup \mathcal{U} \cup \{\mathbf{any}, \mathbf{from}\} \rightarrow_{\text{fin}} \Pi \cup \{\emptyset\}) \cup \perp \right)$$

where \rightarrow_{fin} means that the function maps only a finite subset of its domain to meaningful values (i.e. values different from \perp and \emptyset). With abuse of notation, we use \perp to also denote the empty type, i.e. the function mapping all its domain to \perp . Moreover, by letting λ to range over $\mathcal{L} \cup \mathcal{U} \cup \{\mathbf{any}, \mathbf{from}\}$, we shall write a Δ different from \perp as a non-empty list $[\lambda_i \mapsto [\lambda_{i,j} \mapsto \pi_{i,j}]_{j=1,\dots,k_i}]_{i=1,\dots,n}$. Types are used to express the security policies of nodes. Intuitively, if the type Δ of a node with address l contains the element $[l' \mapsto l'' \mapsto \pi]$, then l' -processes located at l are allowed to perform over l'' only the operations enabled by π . The reserved word **any** is used to refer any node of the net. If it occurs in the domain of Δ then it collects the privileges granted to processes coming from any node of the net (i.e. $[\mathbf{any} \mapsto l'' \mapsto \pi]$ grants all processes the privileges π over l''). If **any** is contained in the domain of $\Delta(l')$, for some l' , then it is used for denoting the operations that l' -processes located at l are allowed to perform over any node of the net (i.e. $[l' \mapsto \mathbf{any} \mapsto \pi]$ grants l' -processes the privileges π over all net nodes). The reserved word **from** stands for the last node visited by a process and is used to grant privileges over this node whatever it is; thus, for instance, $[\mathbf{any} \mapsto \mathbf{from} \mapsto \pi]$ grants l' -process spawned at l the privileges π over l' . The type \perp expresses total absence of privileges.

For the type Δ of a locality l to be *l-well-formed* the following conditions must hold:

1. The keyword **from** can occur only in the function $\Delta(\mathbf{any})$.
2. For each $\ell \in \text{dom}(\Delta)$, it holds that $\Delta(\ell) \preceq \Delta(l)$, where relation \preceq holds true if and only if for all $\lambda \in \text{dom}(\Delta(\ell))$ it holds that $\Delta(l)(\lambda) \cup \Delta(l)(\mathbf{any}) \sqsubseteq_{\Pi} \Delta(\ell)(\lambda)$.
3. For each $\lambda \in \text{dom}(\Delta(\mathbf{any})) - \{\mathbf{from}\}$, it holds that $\Delta(l)(\lambda) \sqsubseteq_{\Pi} \Delta(\mathbf{any})(\lambda)$, and that $\Delta(l)(\mathbf{any}) \sqsubseteq_{\Pi} \Delta(\mathbf{any})(\mathbf{from})$.

The first condition is not too restrictive, because the use of **from** is really necessary only when no knowledge of the last node visited by processes is available (i.e. when using **any**). The second condition says that l grants to ℓ -processes (for $\ell \in \text{dom}(\Delta)$) no more privileges than those granted to its local processes, i.e. those processes statically allocated at l . Finally, the last condition is similar to the previous one, but applies to processes coming from any node; in this case, it is also required that processes coming from any node own over the source node no more privileges than those owned by local processes over any node.

Notice that the syntax of types allows locality variables to occur within types. Basically, they are used when specifying the type of a node dynamically created for referring localities that will be dynamically determined. By exploiting this feature, we can write processes like the following: $\mathbf{in}(!u : \dots)@ \dots \mathbf{newloc}(v : [u \mapsto v \mapsto \{r\}])$.

$\Gamma \upharpoonright_T \mathbf{nil}$		
$\frac{\Gamma(\ell) \cup \Gamma(\mathbf{any}) \sqsubseteq_{\Pi} \langle o, p \rangle \quad t \in \mathcal{T}(p) \quad \Gamma \upharpoonright_T P}{\Gamma \upharpoonright_T \mathbf{out}(t) @ \ell . P}$		
$\frac{\Gamma(\ell) \cup \Gamma(\mathbf{any}) \sqsubseteq_{\Pi} \langle i, p \rangle \quad T \in \mathcal{T}(p) \quad \mathit{upd}(\Gamma, T) \upharpoonright_T P}{\Gamma \upharpoonright_T \mathbf{in}(T) @ \ell . P}$		
$\frac{\Gamma(\ell) \cup \Gamma(\mathbf{any}) \sqsubseteq_{\Pi} \langle r, p \rangle \quad T \in \mathcal{T}(p) \quad \mathit{upd}(\Gamma, T) \upharpoonright_T P}{\Gamma \upharpoonright_T \mathbf{read}(T) @ \ell . P}$		
$\frac{\Gamma(\ell) \cup \Gamma(\mathbf{any}) \sqsubseteq_{\Pi} \{e\} \quad \Gamma \upharpoonright_T P}{\Gamma \upharpoonright_T \mathbf{eval}(Q) @ \ell . P}$		
$\frac{\Gamma(l) \sqsubseteq_{\Pi} \{n\} \quad \Delta \text{ is } u\text{-well-formed} \quad \Gamma \uplus [u \mapsto (\Gamma(l) - \{n\})] \upharpoonright_T P}{\Gamma \upharpoonright_T \mathbf{newloc}(u : \Delta) . P}$		
$\frac{\Gamma \upharpoonright_T P \quad \Gamma \upharpoonright_T Q}{\Gamma \upharpoonright_T P \mid Q}$		$\frac{\Gamma \upharpoonright_T P}{\Gamma \upharpoonright_T A} \text{ if } A \triangleq P$

Table 3. Type Inference Rules

3.3 Static Type Checking

For each node of a net, say $l ::^{\Delta} P$, the static type checker analyzes the operations that P intends to perform when running at l and determines whether they are enabled by the access policy Δ or not (in fact, it is enough to consider $\Delta(l)$). To this aim, a *type context* Γ is a function of the form $\mathcal{L} \cup \mathcal{U} \cup \{\mathbf{any}\} \rightarrow_{\text{fin}} (\Pi \cup \emptyset)$. To update a type context with the type annotations specified within a template, we use the auxiliary function upd that behaves like the identity function for all fields but for formal fields binding locality variables. Formally, it is defined by:

$$\mathit{upd}(\Gamma, T) = \begin{cases} \mathit{upd}(\mathit{upd}(\Gamma, F), T') & \text{if } T = F, T' \\ \Gamma \uplus [u \mapsto \pi] & \text{if } T = !u : \pi, \\ \Gamma & \text{otherwise} \end{cases}$$

where \uplus denotes the pointwise union of functions.

The type judgments for processes take the form $\Gamma \upharpoonright_T P$, where the domain of Γ includes all the localities and all the free locality variables in P . The set of bindings for the localities in Γ implements the access policy of l for the processes statically located at l , while the remaining bindings record the type annotations for the locality variables that are free in P . Intuitively, the judgment $\Gamma \upharpoonright_T P$ states that, within the context Γ , P can be safely executed once located at l .

Type judgments are inferred by using the rules in Table 3 that should be quite explicative. For operations **out**, **in**, **read** and **eval**, the inference requires the capability associated to the operation to be enabled by the capabilities owned over the target ℓ or over all the net sites. Instead, for operation **newloc**, the

capability n must be owned by the site l executing the operation. Moreover, in this case, it is assumed that the creating node owns over the created one all the privileges it owns on itself (except, obviously, for the n capability).

We conclude this section by introducing the notion of *well-typed* net.

Definition 1. A net N is well-typed if for each node $l ::^\Delta P$ in N it holds that Δ is l -well-formed and $\Delta(l) \vdash_{\Gamma} P$.

4 Operational Semantics and Type Soundness

The operational behaviour of μ KLAIM nets can be formalized via a structural congruence and a reduction relation, see [22] for details. Here we just point out some crucial points.

The *structural congruence* gives a convenient way of rearranging the nodes of a net without affecting the behaviour of the net. It says that ‘ \parallel ’ is commutative and associative, that **nil** and **0** are the identities for ‘ $|$ ’ and ‘ \parallel ’ resp., and that process identifiers can be replaced by the processes in the body of their definitions.

The *reduction relation*, \succrightarrow , specifies the basic computational steps and formalizes the informal behaviours sketched in Section 2. Because of the highly dynamic nature of our calculus, the operational semantics uses types to perform some dynamic checks, e.g., when processes migrate (to block migration of processes that do not comply with the security policy of the target node) and when node addresses are retrieved from the TS (to ensure that the local security policy enables correct usability of these addresses). In both cases, the check occurs in the premises of an inference rule thus, if it fails, the rule cannot be used in the inference (i.e. the corresponding net reduction step is blocked). In the rest of this section, we present details on these two specific points.

As regards the first check, the rule for process migration takes the form

$$\frac{\Delta'(l) \uplus (\Delta'(\mathbf{any})[l/\mathbf{from}]) \vdash_{\Gamma'} Q}{l ::^\Delta \mathbf{eval}(Q)@l'.P \parallel l' ::^{\Delta'} P' \succrightarrow l ::^\Delta P \parallel l' ::^{\Delta'} P'|Q}$$

where $\Delta'(\mathbf{any})[l/\mathbf{from}]$ denotes syntactic substitution of **from** with l in function $\Delta'(\mathbf{any})$. Hence, the premise of the rule says that the migrating process Q must be checked against the union of the privileges that the security policy Δ' of the target node l' assigns to processes coming from l and to processes coming from any node (in this last case, occurrences of **from** must be interpreted as l).

The second run-time check is invoked for establishing matching of a formal field $!u : \pi$ against a locality l' when performing **read/in** operations. The reduction rule for **in** is

$$\frac{\mathit{match}_{\Delta(l)}(\mathcal{E}\llbracket T \rrbracket, et)}{l ::^\Delta \mathbf{in}(T)@l'.P \parallel l' ::^{\Delta'} \mathbf{out}(et) \succrightarrow l ::^\Delta P[et/T] \parallel l' ::^{\Delta'} \mathbf{nil}}$$

where $\mathcal{E}\llbracket \cdot \rrbracket$ evaluates the actual fields of T by replacing each expression with the value corresponding to its evaluation (the rule for **read** is similar but leaves

the tuple et in the TS of l'). If the match between the evaluation of the template and the chosen tuple succeeds, all the formal fields of T are replaced with the corresponding values of et in the continuation process P (written $P[et/T]$). In particular, the matching succeeds if for each formal field $!u : \pi$ the corresponding value l'' that will replace u is such that the security policy Δ of the node l where the **in/read** operation is performed allows local processes to perform all the operations enabled by π over l'' , using if needed also the capabilities owned by l 's static code over all the net. This control is implemented by the following matching rule (the remaining matching rules are standard and are omitted)

$$\frac{\Delta(l)(l'') \cup \Delta(l)(\mathbf{any}) \sqsubseteq_{\pi} \pi}{\text{match}_{\Delta(l)}(!u : \pi, l'')}$$

Other than for these two checks, the operational semantics must take types into account for updating the security policy Δ of a node l when it creates a new (fresh) node l' . The semantics prescribes that l -processes can perform over l' all the operations that they can perform locally, except for **newloc**, and hence Δ is extended with $[l \mapsto l' \mapsto (\Delta(l)(l) - \{n\})]$.

Type Soundness. We can now state two standard results for type systems, namely, *subject reduction* and *type safety*. The former means that well-typedness is an invariant of the operational semantics; the latter means that well-typed nets are free from immediate run-time errors. In our framework, such errors would arise when processes attempt to execute operations that are disabled by the security policy of the node where they are running. We use predicate $N \uparrow l$ to express the presence in N of a node l with an illegal behaviour. The two properties together amount to saying that well-typed nets never give rise to run-time errors due to misuse of access privileges. Function $\mathbf{loc}(N)$ returns the set of localities occurring in N and can be easily defined inductively on the syntax of terms, while \succrightarrow^* denotes the reflexive and transitive closure of \succrightarrow .

Theorem 1 (Subject Reduction). *If N is well-typed and $N \succrightarrow N'$ then N' is well-typed.*

Theorem 2 (Type Safety). *If N is well-typed then $N \uparrow l$ for no $l \in \mathbf{loc}(N)$.*

Corollary 1 (Global Type Soundness). *If N is well-typed and $N \succrightarrow^* N'$ then $N \uparrow l$ for no $l \in \mathbf{loc}(N)$.*

Type soundness is one of the main goal of a type system. However, in our framework it is formulated in terms of a property requiring the typing of whole nets. When dealing with larger nets, it is certainly more realistic to reason in terms of parts of the whole net. Hence, we put forward a more *local* formulation of our properties and results. To this aim, we define the *restriction* of a net N to a set of localities D , written N_D , as the subnet obtained from N by deleting all nodes whose addresses are not in D . The wanted local type soundness result can be formulated as follows.

Theorem 3 (Local Type Soundness). *Let N be a net and $D \subseteq \mathbf{loc}(N)$. If N_D is well-typed and $N \succrightarrow^* N'$ then for no $l \in D$ it holds that $N' \uparrow l$.*

5 A Bank Account Management System

In this section, we use our approach to model the simplified behaviour of a bank account management system. For ensuring compliance with the security policy of the bank some aspects of our setting, such as the possibility of granting different privileges to processes coming from different source nodes and the dynamic type checking of mobile processes when they migrate, have proved to be crucial.

We suppose that a bank is located at a node with address l_B and can receive and manage requests coming from many users located at nodes with addresses $l_U, l_{U'}, \dots$. The bank must provide the users with typical account managing operations: opening/closing accounts, putting/getting money in/from accounts, and making statements of accounts. For simplicity, we shall omit some details and technical operations that in reality take place, like, e.g., the charge of taxes, dealing with improper operations like the attempt of getting more money than that really available, \dots .

For the sake of readability, in the rest of this section we will omit trailing occurrences of process **nil**, and use parameterized process definitions (that can be easily implemented in our setting using **out/in** operations to pass/recover the parameters), integer values (to denote, e.g., amounts of money) and strings (to identify the various operations).

For permitting the bank to check the operations that users intend to perform, we assume that users cannot perform remote operations over l_B except for sending processes. Hence, if a user U wants to require an operation to the bank, it has to send a process to l_B (thus virtually moving to the bank) which will interact locally with the proper operation handler. The user process, once it has been accepted (i.e. after its compliance with the bank security policy has been checked), can require the operation by locally producing a tuple whose first field contains the name of the operation and whose second field contains the address of the user node (used to identify the user that made the request). Depending on the operation, the tuple could have other fields containing the amount of money involved in the operation and the account receiving the money.

The node implementing the bank is illustrated in Table 4. First, the bank creates a new node that will contain its clients accounts, stored as tuples of the form $(userAddress, amount)$. This node acts just as a repository for tuples and will not be used for spawning processes, thus it has assigned the empty type \perp . Then, five different handler processes, one for each kind of operation, are concurrently spawned. Each handler continuously waits for a request. When such a request arrives, the proper handler executes its task by remotely accessing the reserved locality and then reports locally a confirmation of action completion. The client process performing the request waits for such a confirmation and then brings it back to its original locality. This last operation is performed by means of a migration thus providing the user node with the chance of controlling the operation.

Notice that, by taking advantage of the semantics of μ KLAIM operations, the simple handlers of Table 4 implement the mutual exclusion needed to ensure the correctness of concurrent operations over shared data. Indeed, once a handler

$l_B ::^{\Delta_B} \text{newloc}(u : \perp). (OpenH(u) \mid PutH(u) \mid GetH(u) \mid ReadH(u) \mid CloseH(u))$ <p>where:</p> $OpenH(u) \triangleq \text{in}(\text{"open"}, !x, !y)@l_B.$ $(OpenH(u) \mid \text{out}(x, y)@u.\text{out}(\text{"OKopen"}, x, y)@l_B)$ $PutH(u) \triangleq \text{in}(\text{"put"}, !x, !y, !w)@l_B.$ $(PutH(u) \mid \text{in}(w, !z)@u.\text{out}(w, z + y)@u.\text{out}(\text{"OKput"}, x, y, w)@l_B)$ $GetH(u) \triangleq \text{in}(\text{"get"}, !x, !y)@l_B.$ $(GetH(u) \mid \text{in}(x, !z)@u.\text{out}(x, z - y)@u.\text{out}(\text{"OKget"}, x, y)@l_B)$ $ReadH(u) \triangleq \text{in}(\text{"read"}, !x)@l_B.$ $(ReadH(u) \mid \text{read}(x, !y)@u.\text{out}(\text{"OKread"}, x, y)@l_B)$ $CloseH(u) \triangleq \text{in}(\text{"close"}, !x)@l_B.$ $(CloseH(u) \mid \text{in}(x, !y)@u.\text{out}(\text{"OKclose"}, x, y)@l_B)$ $\Delta_B \triangleq [l_B \mapsto [l_B \mapsto \{i, o, r, n\},$ $\quad \text{any} \mapsto \{e\}],$ $\quad \text{any} \mapsto [\text{from} \mapsto \{e\},$ $\quad \quad l_B \mapsto \{ \langle o, \{ (\text{"open"}, \text{from}, -),$ $\quad \quad \quad (\text{"put"}, \text{from}, -, -),$ $\quad \quad \quad (\text{"get"}, \text{from}, -),$ $\quad \quad \quad (\text{"read"}, \text{from}),$ $\quad \quad \quad (\text{"close"}, \text{from})$ $\quad \quad \quad \} \rangle,$ $\quad \langle i, \{ (\text{"OKopen"}, \text{from}, -),$ $\quad \quad (\text{"OKput"}, \text{from}, -, -),$ $\quad \quad (\text{"OKget"}, \text{from}, -),$ $\quad \quad (\text{"OKread"}, \text{from}, -),$ $\quad \quad (\text{"OKclose"}, \text{from}, -)$ $\quad \quad \} \rangle$ $\quad \quad \quad]$ $\quad \quad \quad]$
--

Table 4. The node implementing the bank

H has withdrawn the tuple representing an account (i.e. once H has locked the account), in order to proceed in their tasks, all the other handlers have to wait for H to write the updated tuple (i.e. for H to release the lock).

The security policy Δ_B is so defined that ‘sensible’ operations over the accounts of a user U (like getting some money and reading/closing the account) can only be requested by l_U -processes, while operations like putting some money can be requested by processes coming from any node. Moreover, the only remote operation processes are allowed to perform is to come back to their source site. Therefore, a l_U -process can request to the bank sensible operations only over U ’s accounts and can deliver the confirmations only to l_U . Typical processes acting on behalf of a user U are illustrated in Table 5, where the parameter s denotes an amount of money and the parameter l_U ’ denotes an account.

$\begin{aligned} \text{OpenR}(s) &\triangleq \mathbf{eval}(\mathbf{out}(\text{"open"}, l_U, s)@l_B.\mathbf{in}(\text{"OKopen"}, l_U, s)@l_B. \\ &\quad \mathbf{eval}(\mathbf{out}(\text{"OKopen"}, s)@l_U)@l_U)@l_B \\ \text{PutR}(s, l_{U'}) &\triangleq \mathbf{eval}(\mathbf{out}(\text{"put"}, l_U, s, l_{U'})@l_B.\mathbf{in}(\text{"OKput"}, l_U, s, l_{U'})@l_B. \\ &\quad \mathbf{eval}(\mathbf{out}(\text{"OKput"}, s, l_{U'})@l_U)@l_U)@l_B \\ \text{GetR}(s) &\triangleq \mathbf{eval}(\mathbf{out}(\text{"get"}, l_U, s)@l_B.\mathbf{in}(\text{"OKget"}, l_U, s)@l_B. \\ &\quad \mathbf{eval}(\mathbf{out}(\text{"OKget"}, s)@l_U)@l_U)@l_B \\ \text{ReadR} &\triangleq \mathbf{eval}(\mathbf{out}(\text{"read"}, l_U)@l_B.\mathbf{in}(\text{"OKread"}, l_U, !x)@l_B. \\ &\quad \mathbf{eval}(\mathbf{out}(\text{"OKread"}, x)@l_U)@l_U)@l_B \\ \text{CloseR} &\triangleq \mathbf{eval}(\mathbf{out}(\text{"close"}, l_U)@l_B.\mathbf{in}(\text{"OKclose"}, l_U, !x)@l_B. \\ &\quad \mathbf{eval}(\mathbf{out}(\text{"OKclose"}, x)@l_U)@l_U)@l_B \end{aligned}$
--

Table 5. Processes of a user U requesting bank operations

The only possibility for a malicious node to illegally access U 's accounts is to pass through l_U , using a process like $\mathbf{eval}(\mathbf{eval}(\text{MaliciousReq})@l_B)@l_U$. Hence, U has to protect itself from these attacks by granting an e capability over l_B only to processes coming from totally trusted nodes: the security policy of l_U must contain the element $[l \mapsto l_B \mapsto \{e\}]$ only if U trusts the user located at l . However, U can trust l only if U trusts all l' trusted by l (in fact, a node trusted by l can send to l a process that is then allowed to spawn a process at U containing requests on U 's accounts).

Finally, notice that only the handler processes can access the node dynamically created whose address, say l_S , is bound to u . Indeed, when such node is created, the operational semantics dynamically extends Δ_B with $[l_B \mapsto l_S \mapsto \{i, r, o\}]$ thus enabling all the processes initially allocated at l_B to perform **in/read/out** operations over l_S .

6 Concluding Remarks

We presented a new capability based type system for the calculus μKLAIM [21] which controls data/resource access and process mobility in a flexible and expressive way. It has been designed to supply real systems security features, e.g. granting different privileges to processes coming from different nodes and constraining the operations allowed over different kinds of data/resources. Due to the highly dynamic nature of distributed and mobile systems/applications, our framework uses a combination of static and dynamic type checking to guarantee compliance with net security policies. As a future work we plan to integrate in μKLAIM other security mechanisms, like e.g. those based on cryptographic techniques, both for the establishment of secure channels, and for process code security and authentication.

The choice of the process calculus μKLAIM [21], that is at the core of the programming language KLAIM [14] and hence is based on the Linda [19, 11] coordination model, is motivated by the fact that μKLAIM has a number of features that make it appealing also for network computing environments where, in general, connections are not stable and host machines are heterogenous. Indeed, it

permits *time uncoupling* (tuples life time is independent of the producer process life time), *destination uncoupling* (the producer of a tuple does not need to know the future use or the destination of that tuple) and *space uncoupling* (communicating processes need to know a single interface, i.e. the operations over the tuple space). As shown in [17], where several messaging models for mobile processes are examined, the *blackboard* approach, of which tuple space based models are variants, is one of the most appreciated, also because of its flexibility. Evidence of the success gained by the tuple space paradigm is given by the many tuple space based run-time systems, both from industries, e.g. JavaSpaces [33, 2] and TSpaces [36], and from universities, e.g. PageSpace [13], WCL [31], Lime [29] and TuCSon [28].

Many type systems for guaranteeing security properties have been proposed for process calculi with distribution and mobility, but, as far as we know, ours is the first one implementing such fine grained policies. Among those type systems more strictly related to security, we mention those disciplining the types of the values exchanged in communications [9, 3, 25], those for controlling Ambients [10] mobility and ability to be opened [6, 7, 27, 18, 12], that for controlling resource access via policies for mandatory access control [4], that for checking that all processes that intend to perform inputs at a given channel are co-located [37], that for controlling the effect of transmitted process abstractions over local channels [38], and that for restricting the mobility of values/processes only to some part of a distributed system [26]. We also applied the latter approach to μ KLAIM, defining a type system that enforces security policies by confining mobility of processes/values; we left its presentation and the comparisons with the present setting for a full paper.

The research line closest to ours is that on the $D\pi$ -calculus [25], a distributed version of the π -calculus equipped with a type system to control access rights of mobile processes over located resources (i.e. communication channels). Like μ KLAIM, the $D\pi$ -calculus relies on a flat net architecture; however, differently from μ KLAIM, communication is local and channel-based, types describe permissions to use channels, and the net architecture is not independent from the processes involved. [24, 30] present two improved type systems for the $D\pi$ -calculus that permit establishing well-typedness of part of a net. This is similar to our local type soundness result that, however, has been obtained by using only local type information.

[37] presents $D\pi\lambda$, a process calculus that results from the integration of the call-by-value λ -calculus and the π -calculus, together with primitives for process distribution and remote process creation. Apart from the higher order and channel-based communication, the main difference with μ KLAIM is that $D\pi\lambda$ localities are anonymous (i.e. not explicitly referrable by processes) and simply used to express process distribution. In [38], a fine-grained type system for $D\pi\lambda$ is defined that permits controlling the effect of transmitted process abstractions (parameterized with respect to channel names) over local channels. Processes are assigned fine-grained types that, like interfaces, record the channels to which processes have access together with the corresponding capabilities, and process

abstractions are assigned dependent functional types that abstract from channel names and types. This use of types is similar to that of μ KLAIM.

Finally, a number of process calculi base their security policies on transmission of encrypted data over communication channels so that only those processes knowing the proper keys can access these information. [1, 34, 5] present this approach in various settings, but none of them consider process distribution and mobility.

References

1. M. Abadi and A. Gordon. A calculus for cryptographic protocols: The spi calculus. *Information and Computation*, 148(1):1–70, 1999.
2. K. Arnold, E. Freeman, and S. Hupfer. *JavaSpaces Principles, Patterns and Practice*. Addison-Wesley, 1999.
3. M. Bugliesi, G. Castagna, and S. Crafa. Boxed ambients. In *TACS 2001*, number 2215 in LNCS, pages 38–63. Springer, 2001.
4. M. Bugliesi, G. Castagna, and S. Crafa. Reasoning about security in mobile ambients. In *Concur 2001*, number 2154 in LNCS, pages 102–120. Springer, 2001.
5. N. Busi, R. Gorrieri, R. Lucchi, and G. Zavattaro. Secspaces: a data-driven coordination model for environments open to untrusted agents. In *To appear in FOCLASA '02*, ENTCS, 2002.
6. L. Cardelli, G. Ghelli, and A. D. Gordon. Mobility types for mobile ambients. In J. Wiederman, P. van Emde Boas, and M. Nielsen, editors, *Proceedings of ICALP '99*, volume 1644 of LNCS, pages 230–239. Springer, 1999.
7. L. Cardelli, G. Ghelli, and A. D. Gordon. Ambient groups and mobility types. In J. van Leeuwen, O. Watanabe, M. Hagiya, P. Mosses, and T. Ito, editors, *Proceedings of TCS 2000*, volume 1872 of LNCS, pages 333–347. IFIP, Springer, 2000.
8. L. Cardelli, G. Ghelli, and A. D. Gordon. Types for the ambient calculus. *Journal of Information and Computation*, 2002.
9. L. Cardelli and A. D. Gordon. Types for mobile ambients. In *Proceedings of POPL '99*, pages 79–92.
10. L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000.
11. N. Carriero and D. Gelernter. Linda in context. *Communications of the ACM*, 32(4):444–458, 1989.
12. G. Castagna, G. Ghelli, and F. Z. Nardelli. Typing mobility in the seal calculus. In *Concur 2001*, number 2154 in LNCS, pages 82–101. Springer, 2001.
13. P. Ciancarini, R. Tolksdorf, F. Vitali, D. Rossi, and A. Knoche. Coordinating multi-agent applications on the WWW: A reference architecture. *IEEE Transactions on Software Engineering*, 24(5):362–366, 1998.
14. R. De Nicola, G. Ferrari, and R. Pugliese. KLAIM: a Kernel Language for Agents Interaction and Mobility. *IEEE Transactions on Software Engineering*, 24(5):315–330, 1998.
15. R. De Nicola, G. Ferrari, and R. Pugliese. Types as Specifications of Access Policies. In Vitek and Jensen [35], pages 117–146.
16. R. De Nicola, G. Ferrari, R. Pugliese, and B. Venneri. Types for Access Control. *Theoretical Computer Science*, 240(1):215–254, 2000.
17. D. Deugo. Choosing a Mobile Agent Messaging Model. In *Proc. of ISADS 2001*, pages 278–286. IEEE, 2001.

18. M. Dezani-Ciancaglini and I. Salvo. Security types for mobile safe ambients. In *ASIAN Computing Science Conference - ASIAN'00*, volume 1961 of *LNCS*, pages 215–236. Springer, 2000.
19. D. Gelernter. Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems*, 7(1):80–112, 1985.
20. L. Gong. *Inside Java 2 platform security: architecture, API design, and implementation*. Addison-Wesley, Reading, MA, USA, 1999.
21. D. Gorla and R. Pugliese. Resource access and mobility control with dynamic privileges acquisition. Research report, Dipartimento di Sistemi e Informatica, Università di Firenze, 2002. Available at <http://rap.dsi.unifi.it/~pugliese/DOWNLOAD/muklaim-full.pdf>.
22. D. Gorla and R. Pugliese. Enforcing Security Policies via Types. Research report, Dipartimento di Sistemi e Informatica, Università di Firenze, 2003. Available at <http://rap.dsi.unifi.it/~pugliese/DOWNLOAD/spc-full.pdf>.
23. G. M. Graw and E. Felten. *Securing Java*. John Wiley and Son, 1999.
24. M. Hennessy and J. Riely. Type-Safe Execution of Mobile Agents in Anonymous Networks. In Vitek and Jensen [35], pages 95–115.
25. M. Hennessy and J. Riely. Resource Access Control in Systems of Mobile Agents. *Information and Computation*, 173:82–120, 2002.
26. D. Kirli. Confined mobile functions. In *Proc. of the 14th IEEE Computer Security Foundations Workshop*. IEEE Computer Society, 2001.
27. F. Levi and D. Sangiorgi. Controlling interference in ambients. In *Proceedings of POPL '00*, pages 352–364. ACM, Jan. 2000.
28. A. Omicini and F. Zambonelli. Coordination for internet application development. *Autonomous Agents and Multi-agent Systems*, 2(3):251–269, 1999. Special Issue on Coordination Mechanisms and Patterns for Web Agents.
29. G. Picco, A. Murphy, and G.-C. Roman. LIME: Linda Meets Mobility. In D. Garlan, editor, *Proc. of the 21st Int. Conference on Software Engineering (ICSE'99)*, pages 368–377. ACM Press, 1999.
30. J. Riely and M. Hennessy. Trust and partial typing in open systems of mobile agents. In *Proceedings of POPL '99*, pages 93–104.
31. A. Rowstron. WCL: A web co-ordination language. *World Wide Web Journal*, 1(3):167–179, 1998.
32. F. B. Schneider, G. Morrisett, and R. Harper. A language-based approach to security. In *Informatics: 10 Years Ahead, 10 Years Back. Conference on the Occasion of Dagstuhl's 10th Anniversary*, number 2000 in *LNCS*, pages 86–101. Springer, 2000.
33. Sun Microsystems. Javaspaces specification. available at: <http://java.sun.com/>, 1999.
34. J. Vitek, C. Bryce, and M. Oriol. Coordinating processes with secure spaces. *Science of Computer Programming*, 2002.
35. J. Vitek and C. Jensen, editors. *Secure Internet Programming: Security Issues for Mobile and Distributed Objects*, number 1603 in *LNCS*. Springer-Verlag, 1999.
36. P. Wyckoff, S. McLaughry, T. Lehman, and D. Ford. TSpaces. *IBM Systems Journal*, 37(3):454–474, 1998.
37. N. Yoshida and M. Hennessy. Subtyping and locality in distributed higher order processes. In *CONCUR '99*, volume 1664 of *LNCS*, pages 557–572. Springer-Verlag, 1999.
38. N. Yoshida and M. Hennessy. Assigning types to processes. In *Proceedings of LICS'00*, pages 334–348. IEEE, Computer Society Press, June 2000.