

A Calculus of Mobile Resources^{*}

Jens Chr. Godskesen¹, Thomas Hildebrandt¹, and Vladimiro Sassone²

¹ IT University of Copenhagen
{jcg,hilde}@itu.dk
² University of Sussex
vs@susx.ac.uk

Abstract. We introduce a calculus of *Mobile Resources* (MR) tailored for the design and analysis of systems containing mobile, possibly nested, computing devices that may have resource and access constraints, and which are not copyable nor modifiable per se. We provide a reduction as well as a labelled transition semantics and prove a correspondence between barbed bisimulation congruence and a higher-order bisimulation. We provide examples of the expressiveness of the calculus, and apply the theory to prove one of its characteristic properties.

Introduction

Mobile computing resources moving in and out of other computing resources abound in our daily life. Prime examples are smart cards [12] used e.g. in Subscriber Identity Module (SIM) cards or next generation credit cards, moving from card issuers to card holders and in and out of mobile phones or automatic teller machines (ATMs). Accordingly, the ability to reason about correctness of the behavior of concurrent systems containing such resources, as well as the need of design and implementation tools, will raise to an increasingly prominent role. We propose a calculus of *mobile resources* (MR) aimed at designing and analysing systems containing nested, mobile computing resources residing in named locations that have *capacity* constraints. Our goals include to devise a formal framework to express and prove properties that may depend on the assumption that such resources are neither *copyable* nor arbitrarily *modifiable* per se. These assumptions are crucial for the security of systems based on smart cards as trusted computing bases, such as e-cash and SIMs.

The calculus MR is inspired by the Mobile Ambient calculus [5, 16], bears relationships to Boxed Ambients [3] and the Seal calculus [23], and to distributed process algebras [13, 22, 10], but differs from all these in important ways, motivated by our specific goals. Building upon a CCS-like calculus [18] with prefix, restriction, parallel composition, replication, no summation nor recursion, we introduce *named slots*, i.e., if p is a process, then $n[p]$ represents a resource p in a slot named by n . In general, we allow slot aliasing, that is slots to be named

* The third author is supported by ‘MyThS: Models and Types for Security in Mobile Distributed Systems’, EU FET-GC IST-2001-32617.


by more than one name, writing $\tilde{n}[p]$ for a resource p in a slot named by a set of names \tilde{n} . We omit brackets around singleton sets.

We postulate that a resource can move from a location to another only if an empty slot can be found at the target location. This makes $n[\bullet]$ very different from a slot containing a terminated process, and allows us to model locations that can only contain a bounded number of resources, thus capturing a very relevant aspect of real-world devices carrying embedded processors. To abstract away from this, replication in the style of the π calculus can be used to recover the usual semantics of locations by generating unboundedly many slots at a location, as, e.g., in $!n[\bullet]$. Since resources are processes, they might themselves contain slots, giving rise to a nested spatial structure. By allowing restriction of location names, we can represent restricted access to a location.

To help focusing our ideas, let us consider the processes

$$\begin{aligned} \text{Alice} &\triangleq (m)(a[C] \parallel A) \\ \text{Bob} &\triangleq (b)(n)(b[\bullet] \parallel B \parallel \{d, n\}[\bullet]) \\ C &\triangleq c[c'] \parallel m \\ P &\triangleq \text{Alice} \parallel \text{Bob} \end{aligned}$$


consisting of a process Alice with a resource C in a public slot named a and a process Bob having an empty, private slot named b and an empty slot with a public name d and a private name n . For the sake of this discussion, the spatial structure of P can be depicted as in the labelled tree below, where edges represent slots, labels slot names, and nodes processes other than slots.



Mobility of resources in MR is ‘*objective*,’ as opposed to ‘*subjective*,’ i.e. the migration of a resource is initiated and controlled not by the resource itself, but by an external process. More precisely, a resource is controlled by a process *outside* the slot where the resource is placed, that is a process residing at a super location. We introduce this notion by means of *move actions* of the form $n \triangleright \bar{m}$, a capability that should be read ‘move a resource from a slot at the location n to a slot at the location m .’ We use a notation reminiscent of action/co-action pairs to stress the dual roles of n and m that, respectively, give and take a resource, and we will adopt consistent conventions throughout the paper. If for instance $A \triangleq a \triangleright \bar{d}. A'$, we would have

$$P \searrow (m)(b)(n)(a[\bullet] \parallel A' \parallel b[\bullet] \parallel B \parallel \{d, n\}[C])$$

whose spatial structure can be drawn as follows.



Observe that the movement of C from a to d causes a *scope extension* for m .

Carrying on with our example, supposing $B \triangleq d \triangleright \bar{b}. B'$ we have the reduction

$$\begin{aligned}
 (m)(b)(n)(a [\bullet] \parallel A' \parallel b [\bullet] \parallel d \triangleright \bar{b}. B' \parallel \{d, n\} [C]) &\searrow \\
 (m)(b)(n)(a [\bullet] \parallel A' \parallel b [C] \parallel B' \parallel \{d, n\} [\bullet]) & .
 \end{aligned}$$

Observe that the last two reductions illustrate the passage of resource C from Alice to the private slot of Bob, without Alice’s knowing the name of Bob’s private slot.

In order to allow the number of slots to *decrease*, slots may be removed. We denote slot remotion with actions of the form $\mathfrak{h}\bar{n}$. Thus, if $B' \triangleq \mathfrak{h}\{d, n\}. D$, the following reduction is possible.

$$\begin{aligned}
 (m)(b)(n)(a [\bullet] \parallel A' \parallel b [C] \parallel \mathfrak{h}\{d, n\}. D \parallel \{d, n\} [\bullet]) &\searrow \\
 (m)(b)(n)(a [\bullet] \parallel A' \parallel b [C] \parallel D) & .
 \end{aligned}$$

The remarkable features here are that a slot can only be removed by processes knowing all its names. In particular, holding a private name to a slot will prevent the slot’s deletion. In our example, the slot named $\{d, n\}$ cannot be deleted by Alice or by another process in the environment, even though it is globally accessible by name d .

While explicit mobility captures asynchronous communication via resource passing, *synchronous communication* is the second central concept of MR, covering several different aspects of process interaction in our application domain. As in CCS, co-located parallel processes can communicate synchronously by performing respectively an a -action and a \bar{a} -action. In addition, we allow a process to communicate with any of its descendants, by performing a *directed action* of the form δa , where δ is a sequence of slot names. For example, if $D = bc\bar{c}'. D'$ in our running example, we have the reduction

$$\begin{aligned}
 (m)(b)(n)(a [\bullet] \parallel A' \parallel b [c [c'] \parallel m] \parallel bc\bar{c}'. D') &\searrow \\
 (a')(b')(b')(a [\bullet] \parallel A' \parallel b [c [0] \parallel m] \parallel D') & ,
 \end{aligned}$$

where the co-action from D synchronises with the corresponding action from slot c inside slot b . In this way, the actions of the resource C (and its sub-resources) are *dynamically bound* to the directed actions of Bob. Unlike e.g. the Seal calculus, we do not distinguish between undirected actions and actions that may synchronise with ascendants.

Using sequences of names in move actions as for the synchronisation, we can move a resource (subtree) from a slot at an arbitrarily deep sub-location to an empty slot (a black leaf) at another arbitrarily deep sub-location. For instance, if $D' = bc \triangleright \bar{a}. D''$ we have the reduction

$$(m)(b)(n)(a[\bullet] \parallel A' \parallel b[c[0] \parallel m] \parallel bc \triangleright \bar{a}. D'') \searrow (m)(b)(n)(a[0] \parallel A' \parallel b[c'[\bullet] \parallel m] \parallel D'').$$

The reductions presented above constitute the primary mechanisms of MR.

Structure of the paper & Results. After introducing the syntax of MR in §1, in §2 we lay the foundations of its semantic theory by giving a reduction semantics formalising the different ways of interaction discussed above; §3 discusses several small examples aimed at illustrating some particularities of MR. We then proceed in §4 to give a labelled transition semantics to MR equivalent to the reduction one. This is well known to be a non-trivial task for calculi allowing (higher-order) process mobility and scope extension, as in MR when resources containing restricted names are moved. In §5 we provide a characterisation of the barbed congruence in terms of a higher-order labelled transition bisimulation. Predictably, the main difficulty in proving the transition bisimulation to be a congruence is the insertion of processes into slots. One of the examples in §3 will point out one of the reasons for that. The detailed proofs of our results can be found in [11]. As usual with higher-order bisimulations, the characterisation here uses a selected set of contexts that play the role of destructors for the higher-order values, namely *receiving contexts* dealing with the reception of resources into slots. We will return on this later on. In §6 we give an application of the characterisation, proving a linearity property of the calculus by giving a bisimulation of two processes.

Design issues & Related work. As already mentioned, MR shares ideas with the Mobile Ambients (MA) [5]. In both calculi, in fact, processes are equipped with nested, named locations – the ambients – containing processes, and the spatial structure can be dynamically extended or change due to movement. However, likewise the Seal calculus [23], it is the *anonymous contents* of locations to be moved in MR, and it is moved by a process *external* to the location. On the contrary, in MA it is the *named location* to be moved by a process *within* it. Another departure point with MA, where ambients communicate only asynchronously, processes in MR may communicate both *synchronously*, as in CCS and the π -calculus, and *asynchronously*, by exchanging resources. Resource movement is a three-party interaction in both MR and the Seal calculus. However, slots in Seal are pure references that disappear after interaction, while in MR they remain as empty slots until explicitly removed. Moreover, the reception of a resource is via a pair action/co-action in Seal.

To the best of our knowledge, the boundedness of resources is unique, among process algebras, to the calculus proposed in this paper. Similar ideas may of course be found in related area, most notably bounded places in Petri nets, but,

besides the obvious analogies, there seem to be no formal relationships with our notion here.

Our calculus shares with Safe Ambients (SA) [16] – and with several other proposals that space does not allow us to survey upon – the wish to put a stricter control on mobility and access to locations, taking the objective mobility viewpoint. While this is realised by an action/co-action synchronisation between mover and movee in the SA approach, MR relies on move actions performed by the mover. Also, the idea of direct actions across location boundaries is reminiscent of the forms of communications found in the Seal calculus, Boxed Ambients, in $D\pi$ [13] and in the distributed Join calculus [10], though in the latter communication is asynchronous and locations distributed.

Observe that, differently from all these, MR does not allow explicit communication of names. This design choice seems consistent with our application, in that it confines information inside resources and allows network topology evolution only by means of extensions and replacement of substructures, maintaining a strictly hierarchical network structure. We leave to future work the investigation of a capability-passing version of MR, as well as the impact of asynchronous communication between remote, non directly nested sites, and the expressiveness of movements that refer to sibling slots.

Concerning the choice of moves and communication that span multiple slot boundaries our hypothesis is that, slots do not necessarily represent physical location boundaries that enforce a notion of communication distance. Distance may be enforced by use of restricted slot names akin to private fields in Java. For this reason we prefer to develop the theory in full generality. After all, that a location may not be accessed from “grand parent nodes” is an issue that can easily be demanded to the control of a type system. Of course, this choice makes the calculus more complex; its price is a more complex semantic theory, yet – we believe – still manageable. For future reference, let us call MR_2 the calculus restricted to paths of length at most two, that is with only directed communication across at most one boundary and short moves of the form $a \triangleright \bar{c}$ (flat), $ab \triangleright \bar{c}$ (up), and $a \triangleright \bar{c}b$ (down). All the results in the paper carry naturally over for this sub-calculus.

1 The Calculus

We assume an infinite set of *names* \mathcal{N} ranged over by n and m . Let \tilde{n} range over sets of names. Let $\bar{\mathcal{N}} = \{\bar{n} \mid n \in \mathcal{N}\}$ be the set of *co-names*. We let α range over $\mathcal{A} = \mathcal{N} \cup \bar{\mathcal{N}}$ and γ over the set \mathcal{N}^* of sequences of names, referred to as *direction paths*, with ϵ denoting the empty sequence. We use δ to denote elements of \mathcal{N}^+ , the set of non-empty direction paths. The set \mathcal{L} of *prefix labels* is then defined by:

$$\lambda ::= \gamma\alpha \mid \delta \triangleright \bar{\delta}' \mid \mathfrak{h}\tilde{n}.$$

The actions α play the same role as in CCS. However, as explained in the introduction, we allow actions to be extended with a sequence of slot names, so

that $n\alpha$ is an action directed to a resource in a slot identified by n . An action $n\alpha$ synchronises with the corresponding co-action $\bar{\alpha}$ performed by a resource in a slot at n .

The sets \mathcal{P} of *process expressions* is defined by:

$$\begin{aligned} p, q &::= \mathbf{0} \mid \lambda . p \mid p \parallel q \mid !p \mid (n)p \mid \tilde{n}[r] & (\mathcal{P}) \\ r &::= \bullet \mid p \end{aligned}$$

Processes $\mathbf{0}$, $\lambda . p$, and $p \parallel q$ are the ordinary CCS-like constructs, representing respectively the inactive process, the prefixed process, and the parallel composition of processes. The replicated process $!p$ provides as many parallel instances of p as required and adds to the calculus the power of recursive definitions. The restriction $(n)p$ makes name n local to p . The novelty of the calculus resides in the slot processes already described in the introduction: $\tilde{n}[\bullet]$, an empty *slot*, identified by the names in \tilde{n} , and $\tilde{n}[p]$ a slot identified by the names \tilde{n} containing a process p . We will write $n[r]$ for a slot $\{n\}[r]$ possessing only one name. We refer to processes within slots as *resources*.

The restriction operator (n) is the only binding construct; the set $fn(p)$ of free names of p is defined accordingly as usual. By convenience, we omit trailing $\mathbf{0}$ s and hence write λ instead of $\lambda . \mathbf{0}$. As usual, we let prefixing, replication, and restriction be right associative and bind stronger than parallel composition hence writing e.g. $!(n)n . p \parallel q$ instead of $!(n)(n . p) \parallel q$. For a set of names $\tilde{n} = \{n_1, \dots, n_k\}$ we let $(\tilde{n})p$ denote $(n_1) \cdots (n_k)p$.

2 Reduction Semantics

Contexts C are, as usual, terms with a hole $(-)$. We write $C(p)$ for the insertion of p in the hole of context C . An equivalence relation S on \mathcal{P} is a *congruence* if it is preserved by all contexts. As our calculus allows actions involving terms at depths arbitrarily far apart, in order to express its notions with formal precision, yet in succinct terms, we need to make an essential use of a particular kind of contexts throughout the paper. We define an \mathcal{N}^* -indexed family of *path contexts*, C_γ , inductively as:

$$C_\epsilon ::= (-) \quad C_{n\gamma} ::= \tilde{n}[C_\gamma \parallel p] \ , \quad n \in \tilde{n}.$$

Observe that the direction path γ for a context C_γ indicates a path under which the context's 'hole' is found. We extend $fn()$ to path contexts by $fn(C_\gamma) = fn(C_\gamma(\mathbf{0}))$. We also define a family of resource contexts

$$D_{\gamma n} ::= C_\gamma(\tilde{n}[-]) \ , \quad n \in \tilde{n},$$

for the special case where the hole is the only content of a slot.

The structural congruence relation \equiv is the least congruence on \mathcal{P} satisfying alpha-conversion and the rules in Table 1. The equations express that $(P, \parallel, \mathbf{0})$ is a commutative monoid (E_1 – E_3) and enforce the usual rules for scope (E_4 – E_6) and replication (E_7). We write $p \equiv_\alpha q$ if p and q are alpha-convertible.

Table 1. Structural equivalence

$E_1. p \parallel \mathbf{0} \equiv p$	$E_4. (n)\mathbf{0} \equiv \mathbf{0}$	$E_7. !p \equiv p \parallel !p$
$E_2. p \parallel q \equiv q \parallel p$	$E_5. (n)p \parallel p' \equiv (n)(p \parallel p')$, if $n \notin \text{fn}(p')$	
$E_3. (p \parallel p') \parallel p'' \equiv p \parallel (p' \parallel p'')$	$E_6. (n)\tilde{n}[p] \equiv \tilde{n}[(n)p]$, if $n \notin \tilde{n}$	

Table 2. Reduction rules

$\gamma\alpha. p \parallel C_\gamma(\bar{\alpha}. q) \searrow p \parallel C_\gamma(q)$
$\gamma\delta_1 \triangleright \gamma\bar{\delta}_2. p \parallel C_\gamma(D_{\delta_1}(q) \parallel D_{\delta_2}(\bullet)) \searrow p \parallel C_\gamma(D_{\delta_1}(\bullet) \parallel D_{\delta_2}(q))$
$\sharp\tilde{n}. p \parallel \tilde{n}[r] \searrow p$

Evaluation contexts E are contexts whose hole does not appear under prefix or replication, i.e.

$$E ::= (-) \mid \tilde{n}[E] \mid E \parallel p \mid (n)E$$

Define \searrow as the least binary relation on \mathcal{P} satisfying the rules in Table 2 and closed under \equiv and under all evaluation contexts E .

The first rule captures both the standard CCS synchronous communication and a synchronisation reminiscent of the one found in the Seal calculus and in Boxed ambients, in that communication in $n\alpha$ is directed downward and may synchronize with a local communication on α inside n . The purpose of context C_γ there is to express that $\gamma\alpha$ synchronises with an $\bar{\alpha}$ found under path γ . The second rule defines movement of resources. This movement is ‘*objective*,’ meaning that resources are moved from the outside and not by the resource itself, as in the Ambient calculus. The third rule defines deletion of slots. The process that performs it must hold all the names of the slot.

As already remarked, moves across multiple boundaries make the calculus formally more complex. The price we pay is a pervasive use of contexts, starting here in the reduction rules and with effects reaching – as we will see – our bisimulation congruence. We remark that in MR_2 , where paths have length at most two, the movement rules above specialise to

$$\begin{aligned} a \triangleright \bar{c}b. p \parallel a[s] \parallel c[q \parallel b[\bullet]] &\searrow p \parallel a[\bullet] \parallel c[q \parallel b[s]], \\ ab \triangleright \bar{c}. p \parallel a[b[s] \parallel q] \parallel c[\bullet] &\searrow p \parallel a[b[\bullet] \parallel q] \parallel c[s], \\ a \triangleright \bar{b}. p \parallel a[s] \parallel b[\bullet] &\searrow p \parallel a[\bullet] \parallel b[s], \end{aligned}$$

and similarly for the communication rules.

We move now to study the semantic theory of MR. We start by discussing the notion of observation. It seems fair to observe the communication actions processes offer to the environment. We then define barbs as:

$$p \downarrow n \quad \text{if} \quad p \equiv (\tilde{n})(\alpha. p' \parallel q), \quad \alpha \in \{n, \bar{n}\},$$

where $n \notin \tilde{n}$. This excludes observing restricted actions, as well as directed actions and move actions. Several alternative choices of observation appear natural, as e.g. observing at top level – i.e., not inside any slots – free slot names, empty slots, path actions, or movements. Our choice is robust, as none of these alternatives would actually give rise to a different semantic theory from the one we develop below.

Definition 1. A *barbed bisimulation* is a symmetric relation S on \mathcal{P} such that whenever $p S q$

$$\begin{aligned} p \downarrow n \text{ implies } q \downarrow n \\ p \searrow p', \text{ then } \exists q' \searrow q' \text{ with } p' S q' \end{aligned}$$

Barbed bisimulation congruence \sim_b is the largest congruence that is a barbed bisimulation.

The definition above is in principle stricter than the classical notion of barbed congruence, defined as the largest congruence *contained* in the barbed bisimulation. It is however gaining credit in process algebra theory for its good properties (cf. e.g [9, 1]).

3 Examples

Slots with multiple names. As already remarked in the introduction, slots may be given multiple names and some of the names may be hidden. As a further example, consider

$$(a)(a')(\{a, b\}[\bullet] \parallel \{a', b\}[\bullet] \parallel a \triangleright \bar{a} . \mathbf{!}\{a, b\} \parallel a[\bullet]).$$

Here the environment can insert resources non-deterministically into one of two slots using the name b . If a resource is inserted into the buffer slot named $\{a, b\}$ then it is moved to an internal slot and the buffer is removed.

Linearity. It is a fundamental property of MR that resources cannot be copied. This interacts with the usual scoping rules enforced by restriction to yield an interesting ‘*linearity*’ property. Consider the term $P \triangleq (b)(a[b] \parallel \mathbf{!}\bar{a}b . c)$. Because of the restriction (b) only the resource inside slot a will ever be able to use b to interact with the replicated term on its side. However, it may in principle be possible that exporting it to some external context, such a resource may be able to ‘copy’ itself, replicating the reference to b . This would be possible in several (higher-order) calculi. Yet, it is not possible in MR. In particular, if a resource containing a reference to a local name is given out, only (a residual of) *that* resource may in future refer that name. This is stated by the following equation, that we prove in §6.

$$(b)(a[b] \parallel \mathbf{!}\bar{a}b . c) \sim_b (b)(a[b] \parallel \bar{a}b . c)$$

Suggestively, the process $(b)(a[b] \parallel \mathbf{!}\bar{a}b . c)$ can be regarded as a model of a pre-paid cash card (the resource b) in the slot a of a vending machine $a[\bullet] \parallel \mathbf{!}\bar{a}b . c$

that delivers a cup of coffee (action c) for each cash card of the right kind, b , inserted in a . The \sim_b -equation above then states that if there exists only *one* card of the ‘right’ type, then there will ever be only *one* cup of coffee; in other words, the cash card cannot be copied. This is the kind of properties relevant to our intended application area.

Scope extension and mobility The interplay of upward and downward moves and scope extension gives rise to interactions unexpected at first, and is a major challenge for the theory of the observational congruence presented in the next section. Here, we exemplify it as follows. Consider the contexts

$$C_1 \triangleq c[(-)] \parallel a, \quad C_2 \triangleq d[(-)] \parallel \overline{d\bar{a}}.b,$$

and the process $p \triangleq (a)C_1(C_2(\bullet)) = (a)(c[d[\bullet] \parallel \overline{d\bar{a}}.b] \parallel a)$. Since name a is private, it would appear that no resource inserted into the empty slot d can synchronise with the $\overline{d\bar{a}}$ -action and, similarly, that no process can ever synchronise with the a action at top-level. It would then follow that the b action can never be revealed and, in particular, that p behaves like $q \triangleq (a)(c[d[\bullet] \parallel a])$, no matter the context. Yet, this is not the case. Under a suitable context, it is possible for the process a to change its role from being the parent of $\overline{d\bar{a}}.b$ to being its child in the slot named d . Suppose in fact that p and q are inserted into the context

$$C = x[(-)] \parallel y[\bullet] \parallel xc \triangleright \bar{y}. x \triangleright \overline{y\bar{d}}.$$

Then $C(p)$ reduces (in two steps) to

$$(a)(x[\bullet] \parallel y[C_2(C_1(\bullet))]) = (a)(x[\bullet] \parallel y[d[c[\bullet] \parallel a] \parallel \overline{d\bar{a}}.b]),$$

where C_1 and C_2 have swapped place. Now the b -action may be unleashed upon synchronisation on a . Since $b \notin \text{fn}(C(q))$, clearly $C(q)$ cannot reduce to a process with a b -action, so $p \not\sim_b q$.

Digital Signature Card. The following example models a digital signature card. For readability we use the names reg , in and out for slots representing respectively a *register*, an *in-buffer* and a *out-buffer*. We then give a model of a process Enc_k parametrized by a name k (the key) that (repeatedly) encrypts resources received in its *in-buffer* with key k , and returns the encrypted resource via its *out-buffer*.

$$Enc_k \triangleq !(reg)(in \triangleright \overline{reg\bar{k}}. reg \triangleright \overline{out} \parallel reg[k[\bullet] \parallel a]) \parallel in[\bullet] \parallel out[\bullet]$$

Dually, we can define a process Dec_k that (repeatedly) decrypts resources received in its *in-buffer* with key k and returns the decrypted resource via its *out-buffer*.

$$Dec_k \triangleq !(reg)(in \triangleright \overline{reg}. reg\ k \triangleright \overline{out}. \parallel reg[\bullet] \parallel in[\bullet] \parallel out[\bullet])$$

If the name k is globally known, anyone can perform encryption and decryption. On the other hand, if k is a shared secret between two processes, e.g. Alice and Bob, and Alice (resp. Bob) possesses an encryption (resp. decryption) process as a private resource, then Alice can send messages secretly to Bob.

$$\begin{aligned} Alice_{k,M} &\triangleq (m)(a)(a[Enc_k] \parallel m[M] \parallel m \triangleright \overline{a \text{ in. } a \text{ out}} \triangleright \overline{network}) \\ Bob_k &\triangleq (m)(b)(b[Dec_k] \parallel m[\bullet] \parallel network \triangleright \overline{b \text{ in. } b \text{ out}} \triangleright \overline{m}) \\ SecretCom_M &\triangleq (k)(Alice_{k,M} \parallel Bob_k) \parallel network[\bullet] \end{aligned}$$

We may prove the encryption property by showing that for any processes (messages) of the form $M = a_1 . a_2 . \dots . a_i$ and $M' = a'_1 . a'_2 . \dots . a'_j$, we have

$$SecretCom_M \sim_b SecretCom_{M'}.$$

We then may model a digital signature card which generates the key and exports the decryption resource (as many times as needed) but keeps the encryption resource private.

$$SignatureCard \triangleq (k)(!export[Dec_k] \parallel Enc_k)$$

4 Transition Semantics

In this section we set out to provide MR with a labelled transition semantics. The interplay of mobility and local names as illustrated by the two examples in the previous section has interesting consequences in this respect. The first example, in fact, shows that a certain amount of information must be retained about resources given out to the context. This is not similar to the transition semantics for the π calculus (cf.

[18, 20]). In the π calculus the relevant information concerns the extruded names; in MR things may in principle be more complex, since interaction involves passing around resources, i.e. higher-order, evolving entities. The second example also points out that we must consider that exported resources may be received in arbitrarily complex contexts.

We focus on explaining the transition rules for the characteristic features of our calculus, i.e. slots, directed communication and objective mobility, which are shown in Table 4. Also, we explain the interplay between labels.

To capture directed communication, we introduce labels of the form $\overline{\delta}a$ that may synchronise with the directed actions of the form δa that appear as prefixes in the calculus. We do this by defining $\overline{\overline{\delta}a} = \overline{\delta}a$. For example, we have

$$n[a.p] \xrightarrow{\overline{\delta}a} n[p] \quad (\textit{nesting})$$

and the usual synchronisation rule yields

$$\overline{\delta}a . p \parallel n[a.q] \xrightarrow{\tau} p \parallel n[q] ,$$

The three-party interaction required for the movement of resources is modelled by means of higher-order labels. We introduce

$$\bar{\delta} \triangleright \langle p \rangle \quad (p \text{ exits from } \delta) \quad \text{and} \quad (p) \triangleright \delta \quad (p \text{ enters in } \delta).$$

The corresponding co-labels will be indicated by $\delta \triangleright \langle p \rangle$ and $\langle p \rangle \triangleright \bar{\delta}$, respectively.

The move action $\delta_1 \triangleright \bar{\delta}_2$ – whose co-action we denote by $\bar{\delta}_1 \triangleright \delta_2$ – and the two higher-order labels will partially match each other in pairs, so as to give rise to *co-label* corresponding to the third party, the one missing to perfection the three-way synchronisation. That is,

$$\begin{array}{lll} \delta_1 \triangleright \bar{\delta}_2 & \text{coalesces with} & \bar{\delta}_1 \triangleright \langle p \rangle & \text{yielding} & \langle p \rangle \triangleright \bar{\delta}_2, \\ \delta_1 \triangleright \bar{\delta}_2 & \text{coalesces with} & (p) \triangleright \delta_2 & \text{yielding} & \delta_1 \triangleright (p), \\ \bar{\delta}_1 \triangleright \delta_2 & \text{coalesces with} & (p) \triangleright \delta_2 & \text{yielding} & \bar{\delta}_1 \triangleright \delta_2. \end{array}$$

Hence, the three labels match in any order and annihilate their matching action/co-action particles to yield, at last, a τ . For instance, a resource that exits a slot produces a transition

$$n[p] \xrightarrow{\bar{m} \triangleright \langle p \rangle} n[\bullet] \quad (\textit{exit}),$$

and similarly, an empty slot that receives a resource, gives rise to a higher-order transition

$$m[\bullet] \xrightarrow{(p) \triangleright m} m[p] \quad (\textit{enter}).$$

These ‘exit’ and ‘enter’ transitions may synchronise, yielding a *co-move* transition

$$n[p] \parallel m[\bullet] \xrightarrow{\bar{m} \triangleright m} n[\bullet] \parallel m[p] \quad (\textit{co-move})$$

which, in turn, can synchronise with the corresponding move action to yield a τ -action that represent the completed interaction:

$$n[p] \parallel m[\bullet] \parallel n \triangleright \bar{m} . q \xrightarrow{\tau} n[\bullet] \parallel m[p] \parallel q.$$

Symmetrically, ‘exit’ and ‘move’ transitions may synchronise, resulting in a ‘give’ transition

$$n[p] \parallel n \triangleright \bar{m} . q \xrightarrow{\langle p \rangle \triangleright \bar{m}} n[\bullet] \parallel q \quad (\textit{give})$$

which may synchronise with the dual ‘enter’ transition. Dually, ‘enter’ and ‘move’ transitions may synchronise, resulting in a ‘take’ transition

$$m[\bullet] \parallel n \triangleright \bar{m} . q \xrightarrow{m \triangleright \langle p \rangle} m[p] \parallel q \quad (\textit{take}),$$

which is ready to synchronise with the dual ‘exit’ transition.

Rules (*exit*) and, in particular, (*enter*) may at first appear to be ‘spontaneous’ rules. A closer analysis though reveals that they are akin to ‘output’ and ‘input’ transitions in the ‘early’ labelled transitions semantics of the (higher-order) π calculus, rather than transition that may fire autonomously an unbounded number of times.

Table 3. Transition rules, standard

$(prefix) \frac{}{\lambda \cdot p \xrightarrow{\lambda} p}$	$(rest) \frac{p \xrightarrow{\pi} p'}{(n)p \xrightarrow{\pi} (n)p'}, n \notin fn(\pi) \cup bn(\pi)$
$(rep) \frac{p \parallel !p \xrightarrow{\pi} p'}{!p \xrightarrow{\pi} p'}$	$(sync) \frac{p_1 \xrightarrow{(\tilde{n})\bar{\pi}} p'_1 \quad p_2 \xrightarrow{\pi} p'_2}{p_1 \parallel p_2 \xrightarrow{\tau} (\tilde{n})(p'_1 \parallel p'_2)}, fn(p_2) \cap \tilde{n} = \emptyset$
$(par) \frac{p \xrightarrow{\pi} p'}{p \parallel q \xrightarrow{\pi} p' \parallel q}, fn(q) \cap bn(\pi) = \emptyset$	$(sym) \frac{p \parallel q \xrightarrow{\pi} p' \parallel q}{q \parallel p \xrightarrow{\pi} q \parallel p'}$

The last issue involved in the movement of resources is the treatment of scope extension when resources are moved. The phenomenon is totally analogous to that in the (higher-order) pi-calculus, and we handle it as usual (cf. [20]) by restrictions on the labels of (*exit*) and (*give*) transitions, e.g.,

$$(m)n[m] \xrightarrow{(m)\bar{n}\triangleright\langle m \rangle} n[\bullet] \quad (O1)$$

and by explicit scope extension in the synchronisation rule.

The directed communication and movement actions are generalised to actions spanning several levels by the (*nesting*) rule. This uses an operation $n \cdot (-)$ to prepend n to labels coming from processes enclosed into slot n defined as follows:

$$\begin{aligned} n \cdot (\tau) &= \tau, & n \cdot (\bar{\gamma}\alpha) &= \bar{n}\gamma\alpha, & n \cdot (\bar{\delta}_1 \triangleright \delta_2) &= \bar{n}\delta_1 \triangleright n\delta_2, \\ n \cdot ((p) \triangleright \delta) &= (p) \triangleright n\delta, & n \cdot ((\tilde{n})\bar{\delta} \triangleright \langle p \rangle) &= (\tilde{n})\bar{n}\delta \triangleright \langle p \rangle, \end{aligned}$$

where $n \notin \tilde{n}$. By using $n \cdot (\pi)$, we implicitly assume that π is a label of one of the kinds above. Finally, the rule (*delete*) allows a slot to be deleted if all its names are free.

We let π range over the complete set Π of labels used in our transition semantics, defined formally as follows.

$$\pi ::= \beta \mid (\tilde{n})\bar{\delta} \triangleright \langle p \rangle \mid (\tilde{n})\langle p \rangle \triangleright \bar{\delta} \quad \text{where} \quad \beta ::= \tau \mid \lambda \mid \bar{\delta}\alpha \mid \bar{\delta} \triangleright \delta \mid (p) \triangleright \delta \mid \delta \triangleright (p).$$

The set of *bound names* in a label π , $bn(\pi)$, is \tilde{n} if π is $(\tilde{n})\bar{\delta} \triangleright \langle p \rangle$ or $(\tilde{n})\langle p \rangle \triangleright \bar{\delta}$, and \emptyset if π is a β -label.

The rules in Table 3 and 4 then define a labelled transition system

$$(\mathcal{P}, \longrightarrow \subseteq \mathcal{P} \times \Pi \times \mathcal{P}).$$

The following two propositions state the correspondence between the reduction semantics and the transition semantics.

Proposition 1. $p \xrightarrow{\tau} p'$ if and only if $p \searrow p'$.

Proposition 2. $p \xrightarrow{n}$ or $p \xrightarrow{\bar{n}}$ if and only if $p \downarrow n$.

Table 4. Transition rules for resources and mobility

$(exit) \frac{}{\tilde{n}[p] \xrightarrow{\bar{n}\triangleright\langle p \rangle} \tilde{n}[\bullet]}, n \in \tilde{n} \quad (enter) \frac{}{\tilde{n}[\bullet] \xrightarrow{\langle p \rangle\triangleright n} \tilde{n}[p]}, n \in \tilde{n}}$
$(give) \frac{p_1 \xrightarrow{(\tilde{n})\bar{\delta}_1\triangleright\langle q \rangle} p'_1 \quad p_2 \xrightarrow{\delta_1\triangleright\bar{\delta}_2} p'_2}{p_1 \parallel p_2 \xrightarrow{(\tilde{n})\langle q \rangle\triangleright\bar{\delta}_2} p'_1 \parallel p'_2}, fn(p_2) \cap \tilde{n} = \emptyset \quad (take) \frac{p_2 \xrightarrow{\delta_1\triangleright\bar{\delta}_2} p'_2 \quad p_1 \xrightarrow{\langle q \rangle\triangleright\delta_2} p'_1}{p_1 \parallel p_2 \xrightarrow{\delta_1\triangleright\langle q \rangle} p'_1 \parallel p'_2}$
$(co-move) \frac{p_1 \xrightarrow{(\tilde{n})\bar{\delta}_1\triangleright\langle q \rangle} p'_1 \quad p_2 \xrightarrow{\langle q \rangle\triangleright\delta_2} p'_2}{p_1 \parallel p_2 \xrightarrow{\bar{\delta}_1\triangleright\delta_2} (\tilde{n})(p'_1 \parallel p'_2)}, fn(p_2) \cap \tilde{n} = \emptyset$
$(nesting) \frac{p \xrightarrow{\pi} p'}{\tilde{n}[p] \xrightarrow{n \cdot (\pi)} \tilde{n}[p']}, n \in \tilde{n} \quad (delete) \frac{}{\tilde{n}[r] \xrightarrow{\bar{h}\tilde{n}} \mathbf{0}} \quad (alpha) \frac{p \equiv_{\alpha} p' \quad p' \xrightarrow{\pi} q}{p \xrightarrow{\pi} q}$
$(O1) \frac{p \xrightarrow{(\tilde{n})\bar{\delta}\triangleright\langle q \rangle} p'}{(n)p \xrightarrow{(n\tilde{n})\bar{\delta}\triangleright\langle q \rangle} p'}, n \in fn(q) \setminus (fn(\delta) \cup \tilde{n}) \quad (O2) \frac{p \xrightarrow{(\tilde{n})\langle q \rangle\triangleright\bar{\delta}} p'}{(n)p \xrightarrow{(n\tilde{n})\langle q \rangle\triangleright\bar{\delta}} p'}, n \in fn(q) \setminus (fn(\delta) \cup \tilde{n})$

5 Bisimulation Congruence

In this section we provide a labelled transition bisimulation, and prove that it coincides with the barbed bisimulation congruence.

As remarked in §4, we need to take into account that resources may be moved at arbitrary depth. As often happens in higher-order bisimulations (cf. e.g. [19, 8, 17]), we need to use an appropriate selection of destructors in order to test and assess the higher-order values exchanged by interaction. Analogously to what is done in [17] for the ambient calculus, we embody such contexts – resource receptors in our case – in the label. That is, we replace the higher-order actions $(\tilde{n})\langle p \rangle\triangleright\bar{\delta}$ and $(\tilde{n})\bar{\delta}\triangleright\langle p \rangle$ with the family of actions $\triangleright\bar{\delta}(D_{\delta})$ and $(C_{\gamma})\bar{\delta}'\triangleright(D_{\delta})$, respectively. The path contexts D_{δ} and C_{γ} represent the surrounding slots that the resource crosses during its movement.

Definition 2. For D_{δ} and C_{γ} path contexts, we define:

- $p \xrightarrow{\triangleright\bar{\delta}(D_{\delta})} (\tilde{n})(p' \parallel D_{\delta}(q))$ if $p \xrightarrow{(\tilde{n})\langle q \rangle\triangleright\bar{\delta}} p'$, and $fn(D_{\delta}) \cap \tilde{n} = \emptyset$.
- $p \xrightarrow{(C_{\gamma})\bar{\delta}'\triangleright(D_{\delta})} (\tilde{n})(C_{\gamma}(p') \parallel D_{\delta}(q))$ if $p \xrightarrow{(\tilde{n})\bar{\delta}'\triangleright\langle q \rangle} p'$, and $(fn(C_{\gamma}) \cup fn(D_{\delta})) \cap \tilde{n} = \emptyset$.

The set of actions considered in the bisimulation game below is thus:

$$\psi ::= \beta \mid \triangleright\bar{\delta}(D_{\delta}) \mid (C_{\gamma})\bar{\delta}'\triangleright(D_{\delta}).$$

Definition 3. A *simulation* is a binary relation \mathcal{S} over \mathcal{P} such that whenever $p \mathcal{S} q$

$$\text{if } p \xrightarrow{\psi} p' \text{ then } \exists q' \xrightarrow{\psi} q' \text{ such that } p' \mathcal{S} q'$$

\mathcal{S} is a *bisimulation* if \mathcal{S} and \mathcal{S}^{-1} are simulations. We write $p \sim q$ if there exists a bisimulation \mathcal{S} such that $p \mathcal{S} q$.

It follows immediately from the definition of bisimulation that it is an *equivalence relation*. Also, \sim can be proved to be a congruence.

Theorem 1. \sim is a congruence.

From the correspondence between τ transitions and reductions (Prop. 1 and 2), and from the fact that \sim is a congruence (Thm. 1), it follows easily that \sim is sound with respect to the barbed bisimulation congruence. On the other hand, the proof that \sim_b is a bisimulation can be found in [11].

Theorem 2. $\sim_b = \sim$

It can be argued that the use of ‘receptor embodying’ labels and their employment in the bisimulation make the latter a form of contextual equivalence, so that proving processes bisimilarity becomes overly hard. We claim that establishing \sim is still much easier than proving barbed congruence, since the needed contexts have a very simple structure. The next section aims at supporting this claim by analysing an example.

In [11] we show path contexts with paths at length most two are enough for MR *altogether*.

6 An Application

In this section we prove the \sim_b -equivalence illustrated in §3 about the vending machine and the ‘linear’ behaviour of resources, taking up the opportunity to put \sim at work. Exploiting Thm. 2, we prove that

$$(b)(a[b] \parallel !\bar{a}b.c) \sim_b (b)(a[b] \parallel \bar{a}b.c)$$

by showing that the two processes are \sim -bisimilar that, by co-induction, can be proved by proving that $\mathcal{S} = \mathcal{S}_1 \cup \mathcal{S}_2 \cup \{(p, p)\}$ is a bisimulation where

$$\begin{aligned} \mathcal{S}_1 &= \left\{ \langle (b)(C_\gamma(C_{\gamma'}(!\bar{a}b.c) \parallel D_\delta(b)) \parallel p), (b)(C_\gamma(C_{\gamma'}(\bar{a}b.c) \parallel D_\delta(b)) \parallel p) \rangle \right. \\ &\quad \left. | b \notin \text{fn}(C_\gamma) \cup \text{fn}(C_{\gamma'}) \cup \text{fn}(D_\delta) \cup \text{fn}(p) \right\} \\ \mathcal{S}_2 &= \left\{ \langle (b)(C_\gamma(!\bar{a}b.c) \parallel p), (b)(C_\gamma(q) \parallel p) \rangle | b \notin \text{fn}(C_\gamma) \cup \text{fn}(p) \wedge q \in \{\mathbf{0}, \bar{a}b.c\} \right\} \end{aligned}$$

It is then relatively easy to verify that \mathcal{S} is a bisimulation that contains the pair of processes under analysis. Note that it would have been considerably more difficult to prove barbed congruence directly, since that would have required considering *all* contexts, in particular arbitrary replication.

Conclusions and Further Work

We have presented MR, a calculus of nested mobile resources designed to provide fine control on the migration and duplication of resources, as relevant for application in the analysis of mobile embedded devices. Its key properties are: the enforcement of bounded capacity for locations; the synchronous communication between co-located processes and toward children location; and the objective mobility provided by move actions. We have studied a semantic theory for MR based on a reduction semantics and a labelled transition systems, culminating in a bisimulation congruence that coincides with the barbed bisimulation. We provided examples of the expressiveness of MR and put the theory at work by exploiting the correspondence between semantics to prove a characteristic ‘linearity’ property of the calculus.

Among the open issues we plan to address in future work, we mention the study of spatial logics in the style of [7], the provision of suitable type theories, as e.g. [6, 4], to enforce communication and migration safety as well as access control. We plan to extend MR with name-passing, while maintaining the orthogonality of communication by mobility and by exchange of messages. We also plan to explore expressiveness issues by considering alternative design choice and reduced versions of MR, including the absence of communication primitives, move actions that only span single slot boundaries, disallowing scope extension through slot-boundaries, asynchronous messaging, and the cohabitation of slots with ‘soft’ slots that allow copying resources. Also, we are working on an encoding of (a form of linear) capability-passing in MR, and studying the formal relationships with MR₂. We think the theory we have developed carries on smoothly to weak bisimulation; the details are under investigation.

We have applied Sewell’s [21] to derive a transition semantics for a finitary fragment of MR₂ and proved that the bisimulation that so arises is included in ours. We conjecture that they coincide. It would then be interesting to carry on and recast (a larger fragment of) MR in a framework where to understand the relationship with Leifer-Milner’s RPOs based bisimulation [15].

References

- [1] Martin Abadi and Cédric Fournet. Mobile values, new names, and secure communication. In *Proceedings of POPL ’01*. ACM, January 2001. 279
- [2] ACM. *27th Annual Symposium on Principles of Programming Languages (POPL)* (Boston, MA), January 2000.
- [3] Michele Bugliesi, Giuseppe Castagna, and Silvia Crafa. Boxed ambients. In Benjamin Pierce, editor, *TACS’01*, volume 2215 of *Lecture Notes in Computer Science*, pages 38–63. Springer-Verlag, 2001. 272
- [4] Luca Cardelli, Giorgio Ghelli, and Andrew D. Gordon. Ambient groups and mobility types. In J. van Leeuwen, O. Watanabe, M. Hagiya, P. D. Mosses, and T. Ito, editors, *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics, Proceedings of the International IFIP Conference TCS 2000 (Sendai, Japan)*, volume 1872 of *LNCS*, pages 333–347. IFIP, Springer, August 2000. 286

- [5] Luca Cardelli and Andrew D. Gordon. Mobile ambients. In Maurice Nivat, editor, *Proceedings of FoSSaCS'98*, volume 1378 of *LNCS*, pages 140-155. Springer, 1998. [272](#), [275](#)
- [6] Luca Cardelli and Andrew D. Gordon. Types for mobile ambients. In *Proceedings of POPL'99*, pages 79-92. ACM, 1999. [286](#)
- [7] Luca Cardelli and Andrew D. Gordon. Anytime, anywhere: Modal logics for mobile ambients. In *Proceedings of POPL'00 [2]*, pages 365-377. [286](#)
- [8] William Ferreira, Matthew Hennessy, and Alan Jeffrey. A theory of weak bisimulation for core CML. *Journal of Functional Programming*, 8(5):447-491, 1998. [284](#)
- [9] Cédric Fournet and Georges Gonthier. A hierarchy of equivalences for asynchronous calculi. In Larsen et al. [14], pages 844-855. [279](#)
- [10] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In Ugo Montanari and Vladimiro Sassone, editors, *Proceedings of CONCUR'96*, volume 1119 of *LNCS*, pages 406-421. Springer, 1996. [272](#), [276](#)
- [11] Jens Chr. Godskesen, Thomas Hildebrandt, and Vladimiro Sassone. A calculus of mobile resources. Technical Report TR-2002-16, The IT University of Copenhagen, 2002. [275](#), [285](#)
- [12] Uwe Hansmann, Martin S. Nicklous, Thomas Schäck, and Frank Seliger. *Smart Card Application Development Using Java*. Springer, 200. [272](#)
- [13] Matthew Hennessy and James Riely. Resource access control in systems of mobile agents. In Uwe Nestmann and Benjamin C. Pierce, editors, *Proceedings of HLCL'98*, volume 16.3 of *ENTCS*, pages 3-17. Elsevier Science Publishers, 1998. [272](#), [276](#)
- [14] Kim G. Larsen, Sven Skyum, and Glynn Winskel, editors. *25th Colloquium on Automata, Languages and Programming (ICALP) (Aalborg, Denmark)*, volume 1443 of *LNCS*. Springer, July 1998.
- [15] James J. Leifer and Robin Milner. Deriving bisimulation congruences for reactive systems. In *Proceedings of CONCUR 2000*, volume 1877 of *LNCS*, pages 243-258, 2000. [286](#)
- [16] Francesca Levi and Davide Sangiorgi. Controlling interference in ambients. In *Proceedings of POPL'00 [2]*, pages 352-364. [272](#), [276](#)
- [17] Massimo Merro and Matthew Hennessy. Bisimulation congruences in safe ambients. *ACM SIGPLAN Notices*, 31(1):71-80, January 2002. [284](#)
- [18] Robin Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, May 1999. [272](#), [281](#)
- [19] Davide Sangiorgi. Bisimulation for Higher-Order Process Calculi. *Information and Computation*, 131(2):141-178, 1996. [284](#)
- [20] Davide Sangiorgi and David Walker. *The pi-calculus: a Theory of Mobile Processes*. Cambridge University Press, 2001. [281](#), [283](#)
- [21] Peter Sewell. From rewrite rules to bisimulation congruences. In D. Sangiorgi and R. de Simone, editors, *Proceedings CONCUR'98*, volume 1466, pages 269-284, 1998. [286](#)
- [22] Peter Sewell. Global/local subtyping and capability inference for a distributed pi-calculus. In Larsen et al. [14], pages 695-706. [272](#)
- [23] Jan Vitek and Giuseppe Castagna. Seal: A framework for secure mobile computations. In *Internet Programming Languages*, 1999. [272](#), [275](#)