

Multiple-Labelled Transition Systems for nominal calculi and their logics[†]

ROCCO DE NICOLA and MICHELE LORETI

Dipartimento di Sistemi e Informatica, Viale Morgagni, 65 50134 Firenze, Italy

Received 21 May 2007

In Memory of Sauro Tulipani

Action-labelled transition systems (LTSs) have proved to be a fundamental model for describing and proving properties of concurrent systems. In this paper we introduce Multiple-Labelled Transition Systems (MLTSs) as generalisations of LTSs that enable us to deal with system features that are becoming increasingly important when considering languages and models for network-aware programming. MLTSs enable us to describe not only the actions that systems can perform but also their usage of resources and their handling (creation, revelation . . .) of names; these are essential for modelling changing evaluation environments. We also introduce MoMo, which is a logic inspired by Hennessy–Milner Logic and the μ -calculus, that enables us to consider state properties in a distributed environment and the impact of actions and movements over the different sites. MoMo operators are interpreted over MLTSs and both MLTSs and MoMo are used to provide a semantic framework to describe two basic calculi for mobile computing, namely μ KLAIM and the asynchronous π -calculus.

1. Introduction

A well-established and successful approach to modelling and verifying the properties of concurrent systems uses process algebras and temporal logics as a basis. Concurrent systems are specified as terms of a process description language, while properties are specified as temporal logic formulae. Labelled transition systems are associated with terms via a set of structural operational semantics rules, and model checking is used to determine whether the transition systems associated with terms enjoy the property specified by the temporal formulae.

A *Labelled Transition System* (LTS) consists of a set of states \mathcal{S} , a set of transition labels \mathcal{L} and a transition relation $\rightarrow \subseteq \mathcal{S} \times \mathcal{L} \times \mathcal{S}$. *States* correspond to the configurations that the system can reach. *Labels* describe the actions that the system can perform to interact with the environment. *Transition relations* describe the system's evolution as determined by the execution of specific actions. Temporal logic formulae are a mix of logical operators and modal operators. The former are the usual boolean operators, while the latter are the operators that permit reasoning about the system's evolution in time and for dealing with the LTS's dynamic aspects. The success of these two formalisms has led to interesting results and to the development of general tools and theorems for

[†] The work presented in this paper was partially supported by EU Project Software Engineering for Service-Oriented Overlay Computers (SENSORIA, contract IST-3-016004-IP-09).

minimisation and animation of LTSs and for model checking satisfaction of temporal formulae over them.

In the last ten years, and stimulated by new applications of network-aware programming, several new formalisms based on process algebras, but with new constructs for modelling network topology, name passing, resources usage and mobility, have been proposed. Here, we will just mention the Distributed Join-calculus (Fournet *et al.* 1996), the Distributed π -calculus (Hennessy and Riely 2002), the Ambient calculus (Cardelli and Gordon 2000), the Seal calculus (Castagna *et al.* 2005), Nomadic Pict (Wojciechowski and Sewell 1999) and KLAIM (De Nicola *et al.* 1998).

The new primitives of these distributed nominal calculi, are such that not only actions, but also names and resources play a central role. This makes it more difficult to use LTSs, operational semantics and classical temporal logics to deal with the new calculi. To model some of their distinguishing features, such as *name scoping* or *process and data distribution*, transition labels had to be *enriched* to carry information not only about the actions performed, but also about the state of the system. This interplay has made the development of a general approach to distributed nominal calculi more difficult, and we have witnessed the development of a number of theories tailored to specific calculi that cannot be easily generalised.

A side effect of this lack of homogeneity has been the proposal of many different temporal logics devised for reasoning about the different calculi. Two variants of HML (*Hennessy–Milner Logic* (Hennessy and Milner 1985)) for specifying and verifying properties of π -calculus processes were introduced in Milner *et al.* (1993) and in Dam (1996). Cardelli and Gordon (2006) introduced a spatial and modal logic for specifying and verifying properties of Mobile Ambients. On the other hand, *MobileUnity* (McCann and Roman 1998) and *MobAdtl* (Ferrari *et al.* 2002) are two program logics designed specifically for specifying and reasoning about mobile systems by exploiting a Unity-like proof system.

We feel that it is important to come up with a general operational model that can play the same role for distributed nominal calculi that LTSs have played for process algebras. Such an operational model should allow us to capture all main aspects of the new class of systems in a natural way and should represent the natural basis for interpreting temporal logics for describing system properties. A central role in the new set of mobile calculi is played by *names*. Names can be *public*, that is, known by the environment where the system is executing, or *private*. Systems, during execution, can *discover* some of the private names. Thus, the behaviour of terms of mobile calculi cannot be described without taking into account the names they *use*, and classical LTS do not have a natural and explicit handling of names.

In this paper we propose a variant of LTSs, which we call *Multiple-Labelled Transition Systems* (MLTSs), as a candidate general operational model. MLTSs are equipped with a number of transition relations that capture the different aspects of system behaviour, namely *action execution*, *resource creation and consumption* and *name revelation*. An MLTS consists of:

- a set of *states*, which describe system configurations;

- a set of *resources*, which model data (for example, the values exchanged over a channel), computational environments (for example, the locations where processes can be executed) and network links (which can be used for interaction);
- a set of *transition labels*;
- a *naming structure*, which enables us to associate public (known by the environment) names with each state;
- three transition relations:
 - the *computation relation* describes the interaction with the environment;
 - the *resource relation* describes resource usage;
 - the *revelation relation* describes the names revealed to the environment.

In order to model the properties of an MLTS, we introduce a temporal logic (MoMo), which consists of a number of basic operators to be used to describe specific properties/behaviours of mobile and distributed systems. Thus, together with the usual logical connectives and the operators for minimal and maximal fixed points, MoMo is equipped with operators for describing dynamic behaviours (*temporal properties*), for modelling resource management (*state properties*), for the handling of names (*nominal properties*), and for controlling mobile processes (*mobility properties*).

To show the usefulness of our proposal, we show how MLTSs can be used to describe the operational semantics of two formalisms with opposite objectives, namely μKLAIM (Bettini *et al.* 2003) and the asynchronous π -calculus ($A\pi$) (Boudol 1992; Honda and Tokoro 1991). μKLAIM is a simplified version of KLAIM, which is an experimental programming language that supports a programming paradigm where both processes and data can be moved across different computing environments, and which relies on the use of explicit localities. $A\pi$ is the generally recognised minimal common denominator of calculi for mobility. We also introduce two dialects of MoMo for reasoning about the two calculi under consideration.

The rest of the paper is organised as follows. In Section 2, we introduce our general MLTS model, while in Sections 3 and 4 we show how MLTSs can be used to describe the operational semantics of $A\pi$ and μKLAIM , respectively. The modal logic for mobility is presented in Section 5, and its dialects are discussed in Section 6, where we also specify a few properties of a simple client–server system. Section 7 contains some results about the relationships between the equivalence induced on $A\pi$ by MoMo and the classical notion of bisimulation. The final section contains some concluding remarks.

2. Multiple-Labelled Transition Systems

In this section we introduce our variant of LTSs, which will be used as a general framework for describing the semantics of mobile calculi. We will begin by introducing a few basic definitions.

Definition 2.1. Let \mathcal{N} be a set of names and X be a subset of \mathcal{N} . A *name substitution* is a function $\sigma : \mathcal{N} \rightarrow \mathcal{N}$, where:

- $\{y_1/x_1, \dots, y_n/x_n\}$ denotes the substitution that maps x_i into y_i and is the identity on the other names;

- $\sigma_1 \cdot \sigma_2$ is the composition of σ_1 and σ_2 ;
- $\sigma_1 \setminus X$ is the substitution σ_2 such that

$$\sigma_2(n) = \begin{cases} n & \text{if } n \in X \\ \sigma_1(n) & \text{otherwise;} \end{cases}$$

- \emptyset denotes the identity substitution;
- Σ denotes the set of substitutions σ .

Definition 2.2. With $X \subseteq \mathcal{N}$ and $\sigma \in \Sigma$, we define $\sigma(X)$ to be the image of X with respect to σ , namely

$$\sigma(X) = \{x' \mid \exists x \in X : \sigma(x) = x'\}.$$

Definition 2.3. A *naming structure* for a set S is a triple $\mathbb{N} = \langle \mathcal{N}, \eta, \circ \rangle$ such that:

- \mathcal{N} is a countable set of *names*
- $\eta : S \rightarrow 2^{\mathcal{N}}$ is the *naming function*
- $\circ : S \times \Sigma \rightarrow S$ is the *renaming function*

where, for each $s \in S$ and $\sigma \in \Sigma$, if we use $s \circ \sigma$ to denote $\circ(s, \sigma)$, we have

- $\eta(s)$ is finite
- $s \circ \emptyset = s$
- $\eta(s \circ \sigma) = \sigma(\eta(s))$.

Intuitively, a naming structure for S enables us to consider the elements of S as containers of names. For each element s in S , $\eta(s)$ gives the set of names that appear in s , while the application of an element s to a substitution σ ($s \circ \sigma$) yields the element of S obtained by renaming each name in s according to σ .

If one considers the set of π -calculus processes PROC where

- CH is the set of channel names
- $fn(P) \subseteq \text{CH}$ yields the free-names in P
- $P \circ \sigma$ returns the process obtained from P by replacing each free name in P as determined by substitution σ ,

then the triple $\langle \text{CH}, fn, \cdot \rangle$ is a naming structure for PROC.

A Multiple-Labelled Transition System, MLTS for short, consists of a set of *states* \mathcal{S} , a set of *resources* \mathcal{R} , a set of *transition labels* \mathcal{L} and a naming structure $\langle \mathcal{N}, \eta, \sigma \rangle$ for each of \mathcal{S} , \mathcal{R} and \mathcal{L} . *States* describe the configurations of a system modelled by an MLTS. *Naming structures* label each state by a set of names. These are the names that are public (known by the environment) when the system is in that state. *Resources* are the necessary prerequisites for system evolutions. *Transition labels* typically identify the actions a system can perform to interact with the environment. We also assume that \mathcal{L} contains a special distinct transition label τ denoting internal/silent evolution of a system.

MLTSs provide information about:

- the actions that a system can perform to interact with the environment;
- a system’s reactions to creation/deletion of new *resources*; and
- a system’s handling of private/public names.

The modelling via an MLTS is based on three different transition relations:

- the *computation relation* $\longrightarrow \subseteq \mathcal{S} \times \mathcal{L} \times \mathcal{S}$;
- the *resource relation* $\xrightarrow{\oplus} \subseteq \mathcal{S} \times (\{\oplus, \ominus\} \times \mathcal{R}) \times \mathcal{S}$;
- the *revelation relation* $\hookrightarrow \subseteq \mathcal{S} \times \mathcal{N} \times \mathcal{S}$.

As already mentioned, the three transition relations of an MLTS enable us to model a system’s behaviour with respect to three different aspects. The *computation relation* plays the same role as the transition relation in an LTS in that it describes the actions a system can perform to interact with the environment. The *resource relation* describes how a system evolves when resources are created or consumed. If r is a resource, we have that $s_1 \xrightarrow{\ominus r} s_2$ is possible if r is available at state s_1 ; after r is consumed state s_2 is reached. Similarly, we have $s_1 \xrightarrow{\oplus r} s_2$ if s_2 can be obtained from s_1 by adding resource r . The *revelation relation* describes the capability of a system to reveal a private name to the environment.

To guarantee the correct management of names, we need to require some specific properties for the three transition relations, namely that they have to be preserved by *name permutations*, that is, by bijective substitutions.

Definition 2.4. Let $\mathbb{N} = \langle \mathcal{N}, \eta, \circ \rangle$ be a naming structure for A_1, \dots, A_n . We say that a relation $\mathcal{R} \subseteq A_1 \times \dots \times A_n$ is *preserved by name permutations* if and only if for each *name permutation* σ ,

$$(a_1, \dots, a_n) \in \mathcal{R} \Leftrightarrow (a_1 \circ \sigma, \dots, a_n \circ \sigma) \in \mathcal{R}.$$

Definition 2.5. A *Multiple-Labelled Transition System* is a 7-tuple

$$\mathcal{M} = \langle \mathcal{S}, \mathcal{R}, \mathcal{L}, \mathbb{N}, \longrightarrow, \xrightarrow{\oplus}, \hookrightarrow \rangle$$

where:

- \mathcal{S} , \mathcal{R} and \mathcal{L} are countable sets of *states*, *resources* and *transition labels*;
- $\mathbb{N} = \langle \mathcal{N}, \eta, \circ \rangle$ is a naming structure for \mathcal{S} , \mathcal{R} , \mathcal{L} and \mathcal{N} ; and
- the three relations

$$\begin{aligned} \longrightarrow &\subseteq \mathcal{S} \times \mathcal{L} \times \mathcal{S}; \\ \xrightarrow{\oplus} &\subseteq \mathcal{S} \times (\{\oplus, \ominus\} \times \mathcal{R}) \times \mathcal{S} \\ \hookrightarrow &\subseteq \mathcal{S} \times \mathcal{N} \times \mathcal{S} \end{aligned}$$

are preserved by name permutations.

In the rest of this paper we will use $\mathcal{M}^{\mathcal{S}}$, $\mathcal{M}^{\mathcal{R}}$, $\mathcal{M}^{\mathcal{L}}$ and $\mathcal{M}^{\mathbb{N}}$ to denote the set of states, the set of resources, the set of transition labels and the naming structure of an MLTS \mathcal{M} , respectively.

In the next two sections we will show how MLTSs can be used for describing the behaviour of asynchronous π -calculus and μ KLAIM.

3. Asynchronous π -calculus

In this section we consider the asynchronous π -calculus ($A\pi$) (Boudol 1992; Honda and Tokoro 1991), which is a subset of the π -calculus in which there is no output prefixing.

If CH is a countable set of channels (whose elements are ranged over by a, b, \dots), we let $PROC_{A\pi}$ be the set of $A\pi$ processes (P, Q, \dots) defined by the following syntax:

$$\begin{aligned}
 P, Q &::= \bar{a}b \mid G \mid P|Q \mid \nu a.P \mid !G \\
 G &::= \mathbf{0} \mid a(b).P \mid \tau.P \mid G + G
 \end{aligned}$$

where G denotes *guarded processes*.

In $A\pi$, processes interact by exchanging messages over channels.

The process $\bar{a}b$ is used to indicate the availability of message b over channel a ; it models non-blocking outputs.

A *guarded process* G can be:

- $\mathbf{0}$, the deadlocked process;
- $a(b).P$, the process that retrieves a value over channel a and then behaves like process P , and where the name b is bound to the retrieved value;
- $\tau.P$, the process that performs an internal action and then behaves like P ;
- $G_1 + G_2$, the process that can non-deterministically behave like G_1 or G_2 .

Processes are composed via *parallel composition*: $P|Q$ describes a system composed from two components (specified by P and Q , respectively) that can proceed independently and can interact via shared channels.

Private names are defined using *restriction*: $\nu a.P$ denotes the fact that a is a private name in P .

Infinite behaviours are modelled via *process replication* ($!G$). This can be thought of as an infinite composition of G ($G|G|\dots$).

Note that both input prefixing $a(b).P$ and name restriction $\nu a.P$ act as binders for name a within P . In the rest of the paper we shall use $fn(P)$ to denote the set of names free in P . We will also write $P =_\alpha Q$ to denote the fact that P and Q are equal up-to renaming of bound names. For instance, $a(x).\bar{b}(x) =_\alpha a(y).\bar{b}(y)$ and $\nu a.\bar{b}a =_\alpha \nu c.\bar{b}c$. Let P be a process and σ a name substitution, we write $P \circ \sigma$ to mean the process obtained from P by replacing every free name a with $\sigma(a)$.

Terms that intuitively represent the same process are identified by means of standard *structural congruence* \equiv . This relation is defined as the smallest congruence relation over $A\pi$ processes induced by the laws in Table 1. The structural laws express the fact that $|$ is commutative and associative and that the empty process can always be safely removed from or added to a parallel composition. Structural equivalence also states that $!P$ is equivalent to $!P|P$, and that if a does not occur in Q , then $\nu a.(P|Q)$ is equivalent to $\nu a.P|Q$. In the rest of the paper, elements in $PROC_{A\pi}$ will be considered equal up to structural equivalence.

We will now introduce a simple $A\pi$ specification for a simple system, which we will use as a running example in the rest of the paper.

Example 3.1. A *proxy* is a system such that, given two channels a and b , if appropriate, it will emit the values read from channel a on channel b .

Table 1. PROC_{Aπ} Structural equivalence

$P \mathbf{0} \equiv P$	$P Q \equiv Q P$	$(P Q) R \equiv P (Q R)$
$!P \equiv P !P$	$\frac{a \notin fn(Q)}{\nu a.(P Q) \equiv (\nu a.P) Q}$	

Table 2. Labelled transition system for Aπ

$\tau.P \xrightarrow{\tau} P$	$a(b).P \xrightarrow{ac} P[c/b]$	$\bar{a}b \xrightarrow{\bar{a}b} \mathbf{nil}$
$\frac{P \xrightarrow{\bar{a}b} P' \quad a \neq b}{\nu b.P \xrightarrow{\bar{a}(b)} P'}$	$\frac{P \xrightarrow{\lambda} P' \quad a \notin n(\lambda)}{\nu a.P \xrightarrow{\lambda} \nu a.P'}$	$\frac{P \xrightarrow{\bar{a}b} P' \quad Q \xrightarrow{ab} Q'}{P Q \xrightarrow{\tau} P' Q'}$
$\frac{P \xrightarrow{\bar{a}(b)} P' \quad Q \xrightarrow{ab} Q' \quad b \notin fn(Q)}{P Q \xrightarrow{\tau} \nu b.P' Q'}$	$\frac{G \xrightarrow{\lambda} P}{G + G' \xrightarrow{\lambda} P}$	$\frac{G \xrightarrow{\lambda} P}{!G \xrightarrow{\lambda} P !G}$
$\frac{P \xrightarrow{\lambda} P' \quad fn(Q) \cap bn(x) = \emptyset}{P Q \xrightarrow{\lambda} P' Q'}$		

A possible Aπ specification for this system is

$$\text{PROXY}_1 = !a(x).\tau.\bar{b}x$$

where process PROXY₁ models the system as the infinite replication of a process that first reads a value x from channel a, then performs some internal actions, and, finally, emits x on channel b.

The operational semantics of Aπ processes is described in Amadio *et al.* (1998) by means of the LTS relation ($\xrightarrow{\cdot}$), which is defined in Table 2, where it is assumed that $\xrightarrow{\cdot}$ is closed with respect to the structural congruence relation \equiv . This means that

$$\frac{P \equiv Q \quad Q \xrightarrow{\lambda} Q' \quad Q' \equiv P'}{P \xrightarrow{\lambda} P'}$$

As usual, bound and free names are considered, and we have $n(\lambda) = fn(\lambda) \cup bn(\lambda)$ where

$$\begin{aligned} fn(\tau) &= \emptyset & fn(\bar{a}(b)) &= \{a\} & fn(\bar{a}b) &= fn(a(b)) = \{a, b\} \\ bn(\tau) &= \emptyset & bn(\bar{a}(b)) &= \{b\} & bn(\bar{a}b) &= bn(a(b)) = \emptyset. \end{aligned}$$

The transition relation $\xrightarrow{\cdot}$ of Amadio *et al.* (1998) makes use of labels of the form

$$\lambda ::= \tau \mid \bar{a}b \mid \bar{a}(b) \mid a(b),$$

and describes behaviours by considering different aspects at the same time. Labels τ and $a(b)$ are used to describe computations of a process: $P \xrightarrow{\tau} P'$ if P can perform an internal synchronisation and then behaves like P'; while $P \xrightarrow{a(b)} P'$ if P behaves like

Table 3. $\mathcal{M}_{A\pi}$: transition relations

$\tau.P \xrightarrow{\tau} P$	$a(b).P \bar{a}c \xrightarrow{\tau} P[c/b]$
$\frac{P \xrightarrow{\tau} P'}{P Q \xrightarrow{\tau} P' Q}$	$\frac{P \xrightarrow{\tau} P'}{v a.P \xrightarrow{\tau} v a.P'}$
$P \xrightarrow{\oplus \bar{a}b} P \bar{a}b$	$P \bar{a}b \xrightarrow{\ominus \bar{a}b} P$
$\frac{a \neq b}{v a.(P \parallel \bar{b}a) \xrightarrow{a} P \parallel \bar{b}a}$	

P' after value b is retrieved from channel a . The same relation is used to describe the state/spatial configuration of a process: $P \xrightarrow{\bar{a}b} P'$ if $P \equiv P'|\bar{a}b$. Moreover, if the value b is private in P ($P \equiv v b.(P'|\bar{a}b)$), P evolves to P' with a label $\bar{a}(b)$ denoting the fact that b is bound in P' and that b is communicated in some way to the environment, so b is no longer private in P' .

We shall now see how MLTSs can be used to describe the behaviour of $A\pi$ processes in such a way that computations, spatial configurations and name revelations are modelled separately. We will define an MLTS (named $\mathcal{M}_{A\pi}$) that can be used as a semantic model for $A\pi$ processes. The set of states in $\mathcal{M}_{A\pi}$ will be the set of all $A\pi$ processes. Since processes during their computations consume and produce values over channels, we consider our resources to be the set of terms denoting the availability of messages over channels.

The computation relation will only consider internal actions and synchronisation. For this reason, the set of transition labels in $\mathcal{M}_{A\pi}$ will contain τ labels only.

We let $\mathcal{M}_{A\pi}$ be the MLTS defined by

$$\langle \text{PROC}_{A\pi}, \text{RES}_{A\pi}, \text{LAB}_{A\pi}, \mathbb{N}_{A\pi}, \longrightarrow, \xrightarrow{\quad}, \xleftrightarrow{\quad} \rangle \tag{1}$$

where:

- $\text{RES}_{A\pi} = \{\bar{a}b \mid a, b \in \text{CH}\}$;
- $\text{LAB}_{A\pi} = \{\tau\}$;
- $\mathbb{N}_{A\pi} = \langle \text{CH}, fn, \circ \rangle$;
- $\longrightarrow, \xrightarrow{\quad}, \xleftrightarrow{\quad}$ and \xrightarrow{a} are the transition relations induced by the rules of Table 3, plus those induced by the closure of the relations under \equiv as defined in Table 1.

The computation relation describes internal computation of a system caused by internal moves ($\tau.P$) or by process synchronisation. The resource relation states that starting from a process P , a resource $\bar{a}b$ can be created leading to process $P|\bar{a}b$. Conversely, a resource $\bar{a}b$ can be consumed in $P \equiv Q|\bar{a}b$ to obtain Q . Finally, a process P reveals a name a , if this is available over a public channel b . Intuitively, this means that the environment is able to know the *private name*.

To show that the $\mathcal{M}_{A\pi}$ defined above is indeed an MLTS, we need to prove that $\mathbb{N}_{A\pi}$ is a named structures for states, resources and transition labels, and that all the transition relations considered are preserved by name permutation.




Fig. 1. A portion of the MLTS for PROXY₁ of Example 3.1

Lemma 3.2. $\mathbb{N}_{A\pi}$ is a naming structure for $\text{PROC}_{A\pi}$, $\text{RES}_{A\pi}$, $\text{LAB}_{A\pi}$ and CH .

Proof. The thesis follows easily by noting that:

- for each process P and substitution σ , we have $fn(P)$ is finite, $P \circ \emptyset = P$ and $fn(P \circ \sigma) = \sigma(fn(P))$;
- for each $a \in \text{CH}$, we have $fn(a) = \{a\}$ and $a \circ \sigma = \sigma(a)$;
- $fn(\tau) = \emptyset$ and $\tau \circ \sigma = \tau$. □

Lemma 3.3. \rightarrow , $\dots\dots\rightarrow$ and \leftrightarrow are preserved by name permutation.

Proof. We prove the claim for \rightarrow ; the proofs for $\dots\dots\rightarrow$ and \leftrightarrow are similar.

Let σ be a permutation, we have to prove that for each P and Q , if $P \xrightarrow{\tau} Q$, then $P \circ \sigma \xrightarrow{\tau} Q \circ \sigma$.

We proceed by induction on the derivation for $P \xrightarrow{\tau} Q$.

Base of induction: We can distinguish two cases:

- $P = \tau.Q$
- $P = a(b).R[\bar{a}c]$ with $Q = R[c/b]$.

In both the cases, for each permutation σ , $P \circ \sigma \xrightarrow{\tau} Q \circ \sigma$.

Induction hypothesis: For each P and Q , if $P \xrightarrow{\tau} Q$ is derivable in i steps, then for each permutation σ , $P \circ \sigma \xrightarrow{\tau} Q \circ \sigma$.

Induction step: Let $P \xrightarrow{\tau} Q$ be derivable in $i + 1$ steps. We can distinguish three cases:

- $P = P'|R$, $P' \xrightarrow{\tau} Q'$ and $Q = Q'|R$;
- $P \equiv P'$, $P' \xrightarrow{\tau} Q'$ and $Q' \equiv Q$;
- $P = \nu a.P'$ ($\sigma(a) = a$), $P' \xrightarrow{\tau} Q'$ and $Q = \nu a.Q'$;

In these cases, we use the induction hypothesis for each permutation σ to get $P' \circ \sigma \xrightarrow{\tau} Q' \circ \sigma$. Thus the thesis follows easily from the fact that for each R and S :

- $(S|R) \circ \sigma = (S \circ \sigma)|(R \circ \sigma)$;
- $\nu a.S \circ \sigma = \nu a.S \circ \sigma$ ($\sigma(a) = a$); and
- if $R \equiv S$, then for each σ , $R \circ \sigma \equiv S \circ \sigma$. □

In Figure 1, we present a portion of the MLTS describing the behaviour of PROXY₁ (see Example 3.1), where dashed arrows denote *resource* transitions. Note that only a subset

of the possible transitions is shown in the figure; because v could be any of the admissible values and resources could be added at any state.

We now consider the following lemma, which enables us to establish a strong correspondence between the LTS semantics of Amadio *et al.* (1998) and the MLTS semantics presented in this paper.

Lemma 3.4. The following assertions hold:

- 1 $P \xrightarrow{\tau} P' \Leftrightarrow P \xrightarrow{\tau} P'$.
- 2 $P \xrightarrow{\bar{a}b} P' \Leftrightarrow P \xrightarrow{\ominus \bar{a}b} P'$.
- 3 $P \xrightarrow{\bar{a}(b)} P' \Leftrightarrow \exists P''. P \xrightarrow{\ominus^b} P'' \xrightarrow{\ominus \bar{a}b} P'$.

Proof.

- 1 $P \xrightarrow{\tau} P'$ if and only if, either
 - $P \equiv v \tilde{a} . (\tau . Q_1 | Q_2)$ and $P' \equiv v \tilde{a} . Q_1 | Q_2$; or
 - $P \equiv v \tilde{a} . (\bar{a}b | a(c) . Q_1 | Q_2)$ and $P' \equiv v \tilde{a} . (Q_1 [c/b] | Q_2)$.

But the above holds if and only if $P \xrightarrow{\tau} P'$.

- 2 $P \xrightarrow{\bar{a}b} P'$ if and only if $P \equiv \bar{a}b | P'$. Moreover, $P \equiv \bar{a}b | P'$ if and only if $P \xrightarrow{\ominus \bar{a}b} P'$.
- 3 $P \xrightarrow{\bar{a}(b)} P'$ if and only if $P \equiv v b . (\bar{a}b | P')$ and $b \neq a$. However, this holds if and only if $P \xrightarrow{\ominus^b} \bar{a}b | P'$ and $\bar{a}b | P' \xrightarrow{\ominus \bar{a}b} P'$. □

4. μ KLAIM

KLAIM (De Nicola *et al.* 1998) is a formalism that can be used to model and program mobile systems. It has been designed to provide programmers with primitives for handling the physical distribution, scoping and mobility of processes. KLAIM is based on process algebras but makes use of Linda-like asynchronous communication and models distribution *via* multiple shared tuple spaces (Carriero and Gelernter 1989; Gelernter 1985; Gelernter 1989). Tuple spaces and processes are distributed over different localities and the classical Linda operations are indexed by the location of the tuple space they operate on.

For simplicity, we shall use a simplified version of KLAIM, called μ KLAIM (see, for example, Bettini *et al.* (2003)). The main differences between KLAIM and μ KLAIM is that the former allows high-level communication (processes can be used as tuple fields) while the latter only permits the remote evaluation of processes. Moreover, the simpler language does not make any distinction between physical and logical localities, and does not need allocation environments.

A μ KLAIM system, called a *net*, is a set of *nodes*, each of which is identified by a *locality*. *Localities* can be seen as the addresses of nodes. We shall use LOC to denote the set of localities l, l_1, \dots . Every node has a computational component (a set of processes running in parallel) and a data component (a tuple space). Processes interact with each other either locally or remotely, by inserting and withdrawing tuples from tuple spaces.

A tuple is a sequence of *actual* fields. Each actual field can be either a locality (l), a value v from the (finite) set of *basic values* Val (not specified here), or a *variable* x from

Table 4. μ KLAIM syntax

N	$::=$	0	NETS	P	$::=$	nil	PROCESSES
		$l :: C$	<i>empty net</i>			$act.P$	<i>Empty process</i>
		$N_1 \parallel N_2$	<i>located component</i>			$P P$	<i>Action Prefixing</i>
		$\nu l.N$	<i>net composition</i>			$P + P$	<i>Par. Composition</i>
			<i>name restriction</i>			X	<i>Nondet. Choice</i>
						$recX.P$	<i>Recursion Variable</i>
C	$::=$	$P \mid \langle t \rangle$	COMPONENTS			$\nu l.P$	<i>Recursion</i>
							<i>Name restriction</i>
act	$::=$	out (t)@ l	ACTIONS	t	$::=$	$f \mid f, t$	TUPLES
		eval (P)@ l	<i>output</i>	f	$::=$	$l \mid \nu \mid x$	FIELDS
		in (T)@ l	<i>migration</i>				
			<i>input</i>				
				T	$::=$	$F \mid F, T$	TEMPLATES
				F	$::=$	$f \mid !l \mid !x$	TEMP. FIELDS

the set of variables *Var*. Tuples are retrieved from tuple spaces via *pattern matching* using *templates* (T). Templates are sequences of *actual* and *formal* fields. The latter are *variables* that will get a value when a tuple is retrieved, and are distinguished by a ‘!’ before the variable name.

The *pattern-matching* function *match* is defined in Table 5. The meaning of the rules is straightforward: a template matches against a tuple if both have the same number of fields and the corresponding fields match; two values (localities) match only if they are identical, while formal fields match any value of the same type. A successful matching returns a substitution function σ associating the variables contained in the formal fields of the template with the values contained in the corresponding actual fields of the accessed tuple.

The syntax of μ KLAIM nets is defined in the first part of Table 4. Term **0** denotes the *empty net*, that is, the net that does not contain any node. Terms $l :: P$ (*located process*) and $l :: \langle t \rangle$ (*located tuple*) are used to describe basic μ KLAIM nodes: the former states that process P is running at l , and the latter that the tuple space located at l contains tuple $\langle t \rangle$. μ KLAIM nets are obtained by parallel composition (\parallel) of located processes and tuples. Finally, $\nu l.N$ states that l is a private name within N .

The following term denotes a net consisting of two nodes, named l_1 and l_2 .

$$l_1 :: P_1 \parallel l_1 :: P_2 \parallel l_2 :: (Q_1|Q_2) \parallel l_2 :: \langle t_1 \rangle \parallel l_2 :: \langle t_2 \rangle$$

Processes P_1 and P_2 are running at l_1 while Q_1 and Q_2 are running at l_2 . The tuple space located at l_2 contains tuples $\langle t_1 \rangle$ and $\langle t_2 \rangle$, while that located at l_1 is empty

The syntax of μ KLAIM processes is defined in the second part of Table 4, where: **nil** stands for the process that cannot perform any actions; $P_1|P_2$ stands for the parallel composition of P_1 and P_2 ; and $act.P$ stands for the process that executes action act then behaves like P . Possible actions are **out**(t)@ l , **in**(T)@ l and **eval**(P)@ l .

Action **out**(t)@ l adds t to the tuple space located at l . Action **eval**(P)@ l spawns a process P at locality l . Action **in**(T)@ l is used to retrieve tuples from tuple spaces. Unlike **out**, this is a blocking operation: the computation is blocked until a tuple matching

Table 5. Matching rules

$(M_1) \text{ match}(l, l) = []$	$(M_2) \text{ match}(!l_1, l_2) = [l_2/l_1]$
$(M_3) \frac{\text{match}(F, l) = \sigma_1 \quad \text{match}(T, t) = \sigma_2}{\text{match}(F, T), (l, t) = \sigma_1 \cdot \sigma_2}$	

Table 6. Structural congruence

$N_1 \parallel N_2 \equiv N_2 \parallel N_1$	$(N_1 \parallel N_2) \parallel N_3 \equiv N_1 \parallel (N_2 \parallel N_3)$
$l :: C \equiv l :: C \parallel l :: \mathbf{nil}$	$l :: \mathbf{rec}X.P \equiv l :: P[\mathbf{rec}X.P/X]$
$l :: (P_1 P_2) \equiv l :: P_1 \parallel l :: P_2$	$\frac{l \notin \text{fn}(N_2)}{v l.(N_1 \parallel N_2) \equiv (v l.N_1) \parallel N_2}$
$N_1 \parallel \mathbf{0} \equiv N_1$	$v l.\mathbf{0} \equiv \mathbf{0}$
$\frac{N_1 =_\alpha N_2}{N_1 \equiv N_2}$	$\frac{l_2 \neq l_1}{l_1 :: v l_2.P \equiv v l_2.l_1 :: P}$

template T is found in the tuple space located at l . When such a tuple t is found, it is removed from the tuple space and the continuation process is closed with substitution $\sigma = \text{match}(T, t)$, which replaces the formals in T with the corresponding values in t . For instance, if $T = (lu, 4)$ and $et = (l, 4)$, then $\text{match}(T, t) = [l/u]$. For this reason, $\mathbf{in}(T)@l.P$ acts as a binder for variables in the formal fields of T . Finally, $v l.P$ declares a new name l that will be used in P , so $v l.N$ and $v l.P$ act as binders for l in N and P , respectively.

In the rest of the paper, we will use \tilde{l} to denote a finite sequence of names. Moreover, if $\tilde{l} = l_1, \dots, l_n$, we will write $v \tilde{l}.N$ for $v l_1 \dots v l_n.N$, and use $\{\tilde{l}\}$ to denote the set $\{l_1, \dots, l_n\}$. We will use $\text{fn}(\gamma)$, where γ is a generic syntactic term, to denote the set of free names in γ . Finally, we will write $P_1 =_\alpha P_2$ whenever P_1 is equal to P_2 up to renaming of bound names and variables.

In the rest of this section we will use the *structural congruence* \equiv to identify terms that, intuitively, represent the same net. It is defined as the smallest congruence relation over nets that satisfies the laws in Table 6. The structural laws express the following facts:

- \parallel is commutative and associative;
- α -equivalent processes give rise to equivalent nodes;
- the null process and the empty net can always be safely removed/added;
- a process identifier can be replaced by the body of its definition; and
- it is always possible to transform a parallel of co-located processes into a parallel over nodes.

Note that the commutativity and associativity of ‘ $|$ ’ is derived from the commutativity and associativity of ‘ \parallel ’ and from the fact that $l :: (P|Q) \equiv l :: P \parallel l :: Q$.

Table 7. A simple μKLAIM system

<i>Printer</i> :: $\text{rec}X.\text{in}(!\text{from})@\text{Printer}.$
$(X \text{out}(\text{from})@\text{PrintServer}.\text{nil})$
<i>PrintServer</i> :: $\langle\text{PrintSlot}\rangle$
$\langle\text{PrintSlot}\rangle$
$\text{rec}X.\text{in}(\text{Print}, !\text{from})@\text{PrintServer}.$
$(X \text{out}(\text{from})@\text{Printer}.$
$\text{in}(\text{from})@\text{PrintServer}.$
$\text{out}(\text{PrintOk})@\text{from}.$
$\text{out}(\text{PrintSlot})@\text{PrintServer}.\text{nil})$

Example 4.1 (Proxy specification in μKLAIM). The proxy described in Example 3.1 can be specified in μKLAIM as a system consisting of two locations identified by localities l_a and l_b . When a tuple (matching $!$) is stored in the tuple space located at l_a , this is retrieved and stored in the tuple space located at l_b . Two of the possible μKLAIM implementations of such a system are:

$$\kappa\text{PROXY}_1 = l_a :: \text{rec}X.\text{in}(!l)@l_a.\text{out}(l)@l_b.X \parallel l_b :: \text{nil} \quad (2)$$

$$\kappa\text{PROXY}_2 = l_a :: \text{rec}X.\text{in}(!l)@l_a.\text{eval}(\text{out}(l)@l_b.\text{nil})@l_b.X \parallel l_b :: \text{nil} \quad (3)$$

In κPROXY_1 , located at l_a , there is a process that waits for a tuple t matching $!$ and then stores t at l_b . In κPROXY_2 , the process located at l_a withdraws a tuple matching $!$ and then spawns at l_b a process that adds the retrieved tuples locally.

Example 4.2 (A print server). In Table 7 we show how μKLAIM can be used to model a simple print server with two μKLAIM nodes: *PrintServer* and *Printer*. Located at *PrintServer* there is a process that waits for a print request. Each such request contains the locality from which it has been sent. When a request appears in the tuple space at *PrintServer*, the process sends the document to the printer ($\text{out}(\text{from})@\text{Printer}$), waits for the printing signal ($\text{in}(\text{from})@\text{PrintServer}$) and sends an ack ($\text{out}(\text{PrintOk})@\text{from}$) to the client. To send a print request, a print client has to retrieve a *PrintSlot* from the tuple space located at *PrintServer*, where two *PrintSlots* are available.

We now present the operational semantics of μKLAIM nets in terms of MLTSs. We will define an MLTS $\mathcal{M}_{\mu\text{K}}$ by singling out a set of states describing the configurations that μKLAIM nets can reach; a set of labels describing the actions nets can perform to interact with the environment; and a set of resources used/produced by μKLAIM nets in their computation. We will then define the three transition relations that describe: the actions a net can perform; how a net reacts when resources are created or consumed; and when a net reveals private names. The MLTS-based semantics proposed here is just a different presentation of the one, based on LTS, proposed in De Nicola and Loreti (2005). Indeed, it is easy to establish a correspondence between the transition relations of De Nicola and Loreti (2005) and those proposed here.

Table 8. $\mathcal{M}_{\mu K}$: computation relation

$\frac{\sigma = \text{match}(T, t)}{l_1 :: \mathbf{out}(t)@l_2.P \xrightarrow{l_1:t \triangleright l_2} l_1 :: P \parallel l_2 :: \langle t \rangle \quad l_1 :: \mathbf{in}(T)@l_2.P \xrightarrow{l_1:t \triangleleft l_2} l_1 :: P\{\sigma\}}$	
$l_1 :: \mathbf{eval}(Q)@l_2.P \xrightarrow{l_1:Q \blacktriangleright l_2} l_1 :: P \parallel l_2 :: Q$	
$\frac{N_1 \xrightarrow{l_1:t \triangleright l_2} N'_1}{N_1 \parallel l_2 :: C \xrightarrow{\tau} N'_1 \parallel l_2 :: C}$	$\frac{N_1 \xrightarrow{l_1:t \triangleleft l_2} N'_1}{N_1 \parallel l_2 :: \langle t \rangle \xrightarrow{\tau} N'_1}$
$\frac{N_1 \xrightarrow{l_1:Q \blacktriangleright l_2} N'_1}{N_1 \parallel l_2 :: C \xrightarrow{\tau} N_2 \parallel l_2 :: C}$	
$\frac{N_1 \xrightarrow{\lambda} N_2 \quad \lambda' = \lambda/l \quad \lambda' \text{ is relevant}}{v l.N_1 \xrightarrow{\lambda'} v l.N_2}$	$\frac{N_1 \xrightarrow{\lambda} N_2}{N_1 \parallel N \xrightarrow{\lambda} N_2 \parallel N}$

The set of states in $\mathcal{M}_{\mu K}$ coincides with the set NET of μ KLAIM nets. Transition labels have a complex structure and contain information about the actions performed by located processes. Moreover, the set of resources $\text{RES}_{\mu K}$ will contain locality names and tuples with an indication of their location.

Let $\mathcal{M}_{\mu K}$ be the MLTS defined by

$$\langle \text{NET}, \text{RES}_{\mu K}, \text{LAB}_{\mu K}, \mathbb{N}_{\mu K}, \longrightarrow, \dots \blacktriangleright, \leftrightarrow \rangle$$

where:

— $\text{RES}_{\mu K}$ is the set of resources r defined by the syntax

$$r ::= l \mid \langle t \rangle @ l.$$

— $\text{LAB}_{\mu K}$ is the set of transition labels λ defined using the grammar

$$\lambda ::= \tau \mid l_1 : t \triangleright l_2 \mid l_1 : t \triangleleft l_2 \mid l_1 : P \blacktriangleright l_2.$$

— $\mathbb{N}_{\mu K} = \langle \text{Loc}, fn, \circ \rangle$.

— $\longrightarrow, \dots \blacktriangleright$ and \leftrightarrow are the transition relations, closed with respect to \equiv , induced by the rules of Table 8 and Table 9.

Resource $\langle t \rangle @ l$ indicates the presence of tuple t in the tuple space located at l , and l indicates that location l does indeed exist.

Label τ denotes silent transitions while $l_1 : t \triangleright l_2$, $l_1 : t \triangleleft l_2$ and $l_1 : P \blacktriangleright l_2$ denote, respectively, the fact that a process located at l_1 :

- inserts tuple t in the tuple space located at l_2 ;
- withdraws tuple t from the tuple space located at l_2 ;
- spawns process P to be evaluated at l_2 .

The computation relation (\longrightarrow) respects the intuitive meaning of μ KLAIM operations. A net N_1 reduces to N'_1 with label $l_1 : t \triangleright l_2$ if there is a process located at l_1 that performs

Table 9. $\mathcal{M}_{\mu K}$: resource relation and revelation relation

$l :: \langle t \rangle \xrightarrow{\ominus(t)@l} l :: \mathbf{nil}$	$l :: C \xrightarrow{\ominus l} l :: C$
$N_1 \xrightarrow{\ominus r} N'_1 \quad l \notin r$	$N_1 \xrightarrow{\ominus r} N'_1$
$v l.N_1 \xrightarrow{\ominus r} v l.N'_1$	$N_1 \parallel N_2 \xrightarrow{\ominus r} N'_1 \parallel N_2$
$N_1 \xrightarrow{\oplus(t)@l} N_1 \parallel l :: \langle t \rangle$	$N_1 \xrightarrow{\oplus l} N_1 \parallel l :: \mathbf{nil}$
$\frac{l_1 \neq l_2 \quad l_2 \in t}{v l_1.(N \parallel l_2 :: \langle t \rangle) \xleftrightarrow{l_1} N \parallel l_2 :: \langle t \rangle}$	

out(t)@l₂. The same net N_1 evolves with a τ when it is composed in parallel with node l_2 . Similar considerations apply for input and eval.

The rule for restriction ($v l.N_1$) relies on name hiding: if λ is a transition label, λ/l denotes the transition label obtained from λ by *hiding* location l , namely by replacing l with a special distinct locality \star denoting the *unknown* location.

Definition 4.3.

1 Let \star be a special distinct name in Loc. We let

$$l_1/l_2 = \begin{cases} \star & \text{if } l_1 = l_2 \\ l_1 & \text{otherwise.} \end{cases}$$

2 λ/l is defined inductively as follows:

- $\tau/l = \tau$
- $(l_1 : t \triangleright l_2)/l = (l_1/l) : (t[\star/l]) \triangleright (l_2/l)$
- $(l_1 : t \triangleleft l_2)/l = (l_1/l) : (t[\star/l]) \triangleleft (l_2/l)$
- $(l_1 : P \blacktriangleright l_2)/l = (l_1/l) : P[\star/l] \blacktriangleright (l_2/l)$.

A transition label is *relevant* if it is τ or it takes effect on a visible/public location. Formally, we have the following definition.

Definition 4.4. A transition λ is *relevant* if and only if:

- $\lambda = \tau$; or
- λ is $l_1 : t \triangleright l_2$, $l_1 : t \triangleleft l_2$ or $l_1 : P \blacktriangleright l_2$ and $l_2 \neq \star$.

If N_1 reduces with λ to N'_1 , and λ/l is relevant, $v l.N_1$ reduces with λ/l to $v l.N'_1$.

The resource relation ($\xrightarrow{\dots}$) is as expected. If in N_1 a tuple space located at l contains tuple t , then N_1 evolves with $\ominus(t)@l$ by removing such a tuple. Conversely, N_1 evolves with $\oplus(t)@l$ by adding tuple t to the tuple space located at l . Reductions concerning $\ominus l$ and $\oplus l$ are similar. If N_1 contains a node named l , then $N_1 \xrightarrow{\ominus l} N_1$, while $N_1 \xrightarrow{\oplus l} N_1 \parallel l :: \mathbf{nil}$. Note that in the case of $\ominus l$ nothing is removed from N_1 : resource l is used without being consumed.

Finally, a net N_1 reveals a private name l if there exists a tuple t containing l located in a *public* tuple space, that is, a tuple space located at a non-private node.

As with the $\mathcal{A}\pi$, we now have to prove that $\mathbb{N}_{\mu K}$ is a naming structure for NET , $\text{RES}_{\mu K}$, $\text{LAB}_{\mu K}$ and LOC , and that \rightarrow , $\dots\rightarrow$ and \leftrightarrow are preserved by name (locality) permutation.

Lemma 4.5. $\mathbb{N}_{\mu K}$ is a naming structure for NET , $\text{LAB}_{\mu K}$, $\text{RES}_{\mu K}$ and LOC .

Lemma 4.6. \rightarrow , $\dots\rightarrow$ and \leftrightarrow are preserved by name (locality) permutation.

The proofs for the lemmas above follow along the same lines as the proofs of Lemmas 3.2 and 3.3.

5. MoMo: a modal logic for mobility

In this section, we present MoMo logic, which is a modal logic that enables us to specify properties of MLTSs. With MoMo we lay the basis for defining a common logical framework for specifying properties of mobile calculi whose semantics is given in terms of an MLTS. A variant of MoMo that was specifically designed for specifying and verifying properties of KLAIM systems was introduced in De Nicola and Loreti (2005).

MoMo contains four groups of logical operators:

- The *kernel fragment* contains standard first-order logic operators and *recursive formulae* that enable us to describe recursive properties.
- The remaining fragments rely on the three transition relations contained in an MLTS.
- *State formulae* are used to assert properties concerning the allocation of resources in a system; their interpretation function will be defined in terms of the relation $\dots\rightarrow$.
- *Temporal formulae*, which are interpreted by considering the computation relation \rightarrow , describe the action a system can perform to interact with the environment. These are parametrised with respect to a set \mathcal{A} of *label predicates* that will be instantiated for the different calculi.
- The properties of names, such as freshness or revelation, are specified using *nominal formulae*, whose semantics are given using the relation \leftrightarrow .

In the rest of this section we describe and motivate each fragment of the logic separately, and discuss formulae satisfaction. Models for MoMo formulae are pairs of the form $\langle \mathcal{M}, \mathbf{A} \rangle$ where \mathcal{M} is an MLTS and \mathbf{A} is a function that maps label predicates to sets of transition labels belonging to $\mathcal{M}^{\mathcal{L}}$. The formal syntax and semantics of MoMo will be defined in Section 5.5.

5.1. Kernel fragment

This fragment of the logics is standard. The propositional part of this fragment contains *True* (\mathbf{T}), *Conjunction* ($\phi_1 \wedge \phi_2$) and *Negation* ($\neg\phi$). The interpretation of this part of the logic is as expected. Every state satisfies *True*, a state satisfies $\phi_1 \wedge \phi_2$ if and only if both ϕ_1 and ϕ_2 are satisfied, and $\neg\phi$ is satisfied by each state that does not satisfy ϕ .

The kernel fragment also contains formulae for specifying recursive properties: *maximal fixed point* and *logical variables*. Intuitively, a state s satisfies $\nu\kappa.\phi$ if s satisfies $\phi[\nu\kappa.\phi/\kappa]$.

Formally, the interpretation of $\nu\kappa.\phi$ is defined as the maximal fixed point of the interpretation of ϕ .

If $\mathbb{M}[\cdot]$ is the interpretation function of the proposed logic (which is formally defined later), the set of nets satisfying $\nu\kappa.\phi$ is defined as the maximal fixed point of the function $\mathcal{F}_\phi : 2^{\mathcal{M}^S} \rightarrow 2^{\mathcal{M}^S}$ defined as follows:

$$\mathcal{F}_\phi(S) \stackrel{\text{def}}{=} \mathbb{M}[\phi][\kappa \mapsto S]$$

where $\mathbb{M}[\phi][\kappa \mapsto S]$ identifies the set of states satisfying ϕ when it is assumed that each state in S satisfies κ .

To guarantee well-definedness of the interpretation function, it is assumed that in every formula $\nu\kappa.\phi$, the logical variable κ is positive within ϕ (that is, it appears within the scope of an even number of negations). This enables us to define $\mathbb{M}[\cdot]$ as the composition of monotonic functions in $2^{\mathcal{M}^S} \rightarrow 2^{\mathcal{M}^S}$; since $2^{\mathcal{M}^S}$ is a complete lattice, Tarski's Fixed Point Theorem guarantees the existence of a unique *maximal* fixed point for $\mathbb{M}[\phi]$.

5.2. State properties

This fragment of the logic is intended to be used for describing properties concerning the availability of resources and a system's reactions to the placement of new resources in a state.

Typical properties one could specify using this fragment of the logic are:

'Two instances of resource ρ are available.' (†)

'Eventually after the insertion of resource ρ_1 , resource ρ_2 will be available.' (‡)

An analysis of the resources present in a given state can be described by means of the *consumption* operator $\rho \rightarrow \phi$. Intuitively, a state s satisfies $\rho \rightarrow \phi$ if resource ρ is available at s and, once it is removed, formula ϕ is satisfied.

By using the consumption operator, property (†) above would be rendered as

$$\rho \rightarrow \rho \rightarrow \mathbf{T}.$$

Production properties are used to state properties depending on the availability of new resources. These properties are specified by using the *production* operator $\rho \leftarrow \phi$, which is satisfied by every state that satisfies ϕ after a resource ρ is produced. Hence, a state s_1 satisfies $\rho \leftarrow \phi$ if there exists s_2 such that $s_1 \xrightarrow{\oplus \rho} s_2$ and s_2 satisfies ϕ .

Property (‡) can be formalised in MoMo as

$$\rho_1 \leftarrow \text{Eventually}(\rho_2 \rightarrow \mathbf{T})$$

where *Eventually*(ϕ) is a derivable operator, which we use as a macro in MoMo, asserting that sooner or later a state satisfying ϕ will be reached.

Note that the production operator is reminiscent of the *separating implication* of *Separation Logic* (Reynolds 2002). This operator allows one to describe the properties of *shared memories* extended with a set of values satisfying a given specification.

5.3. Temporal properties

Temporal properties are specified using *diamond* operator $\langle \cdot \rangle \phi$. This is a variant of the HML diamond operator that, instead of being indexed by a set of transition labels, is indexed by a label predicate α from a given set of *label predicates* \mathcal{A} . The interpretation of modal operator $\langle \cdot \rangle \phi$ relies on function \mathbf{A} that, for each label predicate α , yields the set of transition labels satisfying α .

Formally, a state s_1 satisfies $\langle \alpha \rangle \phi$ if and only if there exist s_2 and λ such that $\lambda \in \mathbf{A}(\alpha)$, $s_1 \xrightarrow{\lambda} s_2$ and s_2 satisfies ϕ .

5.4. Nominal properties

In this section we introduce operators for dealing with name properties. They will enable us to describe properties of the form ‘there exists a name that is used in a specific way’ or ‘a private name is used accordingly to a given protocol’. Below, we shall consider four different operators: *quantification*; *revelation*; *name freshness* and *name matching*.

Name quantification ($\exists n. \phi$) enables us to quantify properties with respect to names – a state s satisfies $\exists n. \phi$ if there exists n' such that s satisfies $\phi[n'/n]$.

Name revelation ($n \textcircled{R} \phi$) is used to describe properties concerning the disclosure of private names. A state satisfies $n \textcircled{R} \phi$ if n can be uncovered (revealed) before reaching a state satisfying ϕ . A state s_1 reveals n if there exists s_2 such that $s_1 \xrightarrow{n} s_2$. Hence, s_1 satisfies $n \textcircled{R} \phi$ if and only if there exists s_2 such that $s_1 \xrightarrow{n} s_2$ and s_2 satisfies ϕ .

Name Freshness ($\bigwedge n \phi$) acts as a quantifier over all names that do not occur free either in the formula ϕ or in the state. It is an adaptation of the Gabbay–Pitts quantifier $\bigwedge l \phi$ (Gabbay and Pitts 1999). The fresh name quantifier, when used with state formulae (and, in particular, with the production operator), enables us to assert properties related to the creation of new resources.

Given the MLTS \mathcal{M} , where $\mathcal{M}^N = \langle \mathcal{N}, \eta, \circ \rangle$, a state $s \in \mathcal{M}^S$ satisfies $\bigwedge n \phi$ if and only if there exists n' that is not in $\eta(s)$ and does not appear in ϕ such that $\phi[n'/n]$ is satisfied by s .

Name matching ($\{n_1 = n_2\}$) enables us to verify whether two names are equal. Namely, a state satisfies $\{n_1 = n_2\}$ if n_1 and n_2 refer to the same name.

5.5. Syntax and Semantics of MoMo

The actual syntax of MoMo is summarised in Table 10 and its interpretation function is formally defined in Table 11.

MoMo formulae are interpreted with respect to a pair $\langle \mathcal{M}, \mathbf{A} \rangle$, where \mathcal{M} is an MLTS and \mathbf{A} is a function that maps label predicates to sets of transition labels. Formulae interpretation function ($\mathbf{M} \llbracket \cdot \rrbracket^{\langle \mathcal{M}, \mathbf{A} \rangle}$) is parametrised with respect to such a pair. We will omit explicit reference to $\langle \mathcal{M}, \mathbf{A} \rangle$ whenever it is clear from the context.

Formulae are interpreted under a function, called the *logical environment* and denoted by ε , that associates to each logical variable a set of states. We use ε_0 to denote the logical environment such that, for each logical variable κ , $\varepsilon_0(\kappa) = \emptyset$.

Table 10. The syntax of formulae

$\phi ::=$	FORMULAE
T	<i>true</i>
$\phi \wedge \phi$	<i>conjunction</i>
$\neg\phi$	<i>negation</i>
$v\kappa.\phi$	<i>maximal fixed point</i>
κ	<i>logical variable</i>
$\rho \rightarrow \phi$	<i>consumption</i>
$\rho \leftarrow \phi$	<i>production</i>
$\langle \alpha \rangle \phi_2$	<i>diamond</i>
$\exists n.\phi$	<i>name quantification</i>
$\{n_1 = n_2\}$	<i>name matching</i>
$n \textcircled{R} \phi$	<i>revelation</i>
$\not\! n\phi$	<i>fresh name quantification</i>

Table 11. The interpretation of formulae

Kernel Fragment	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \mathbf{T} \rrbracket \varepsilon = S$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \neg\phi \rrbracket \varepsilon = S - \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi_1 \wedge \phi_2 \rrbracket \varepsilon = \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi_1 \rrbracket \varepsilon \cap \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi_2 \rrbracket \varepsilon$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket v\kappa.\phi \rrbracket \varepsilon = v\mathcal{F}$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \kappa \rrbracket \varepsilon = \varepsilon(\kappa)$	
where:	$\mathcal{F}(x) = \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon \cdot [\kappa \mapsto x]$
State Formulae	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \rho \rightarrow \phi \rrbracket \varepsilon = \{s_1 s_1 \xrightarrow{\ominus\rho} s_2 \text{ and } s_2 \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon\}$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \rho \leftarrow \phi \rrbracket \varepsilon = \{s_1 s_1 \xrightarrow{\oplus\rho} s_2 \text{ and } s_2 \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon\}$	
Temporal Formulae	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \langle \alpha \rangle \phi \rrbracket \varepsilon = \left\{ s_1 \exists s_2 \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon. s_1 \xrightarrow{\lambda} s_2, \lambda \in \mathbb{A}(\alpha) \right\}$	
Nominal Properties	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \exists n.\phi \rrbracket \varepsilon = \bigcup_{n' \in \mathcal{N}} \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi[n'/n] \rrbracket \varepsilon$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket n \textcircled{R} \phi \rrbracket \varepsilon = \left\{ s_1 \exists s_2 \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon. s_1 \hookrightarrow^n s_2 \right\}$	
$\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \not\! n\phi \rrbracket \varepsilon = \{s \exists n'. n' \notin fn(\phi) \cup \eta(s) \text{ and } s \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi[n'/n] \rrbracket \varepsilon\}$	

The interpretation function $\mathbb{M} \llbracket \cdot \rrbracket^{(\mathcal{M}, \mathbb{A})}$ takes a formula ϕ and a *logical environment* ε and yields the set of states in \mathcal{M} satisfying ϕ , that is, the set of state that are *models* of ϕ .

We shall use $\mathbb{L}_{\mathcal{A}}$ to denote the set of MoMo formulae that use label predicates belonging to \mathcal{A} .

The following lemma guarantees that names can be consistently replaced in nets and formulae while preserving the satisfaction relation.

Table 12. Derivable operators

$\mu\kappa.\phi \stackrel{\text{def}}{=} \neg v\kappa.\neg\phi[\neg\kappa/\kappa]$	$\phi_1 \vee \phi_2 \stackrel{\text{def}}{=} \neg(\neg\phi_1 \wedge \neg\phi_2)$
$\rho \Rightarrow \phi \stackrel{\text{def}}{=} \neg\rho \rightarrow \neg\phi$	$[\alpha]\phi \stackrel{\text{def}}{=} \neg\langle\alpha\rangle\neg\phi$
$\forall n.\phi \stackrel{\text{def}}{=} \neg\exists n.\neg\phi$	$\{n_1 \neq n_2\} \stackrel{\text{def}}{=} \neg\{n_1 = n_2\}$

Lemma 5.1. Let \mathcal{M} be an MLTS, and $\mathcal{M}^{\mathbb{N}} = \langle \mathcal{N}, \eta, \circ \rangle$, for every state $s \in \mathcal{M}^{\mathbb{S}}$, formula ϕ and names n_1 and n_2 in \mathcal{N} , we have $s \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon$ if and only if $s \circ \{n_2/n_1\} \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi[n_2/n_1] \rrbracket$.

Proof. The proof follows easily by induction on the structure of the syntax of formulae thanks to the fact that all the transition relations in \mathcal{M} are preserved under name permutations. □

In the rest of this paper we will use $\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket$ to denote $\mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket \varepsilon_0$, and $s \models_{(\mathcal{M}, \mathbb{A})} \phi$ will be used to indicate that $s \in \mathbb{M}^{(\mathcal{M}, \mathbb{A})} \llbracket \phi \rrbracket$.

In Table 12 we present a set of fairly standard derivable operators, which we will use as macros of the logic in the examples in the rest of the paper.

6. MoMo dialects for the analysis of mobile and distributed systems

In this section we present two different dialects of MoMo that enable us to specify and verify properties of mobile and distributed systems specified using $A\pi$ and μKLAIM . To define each of these dialects we will single out a set of label predicates (\mathbb{A}) and an interpretation function (\mathbb{A}) that identifies the set of transition labels satisfying a given label predicate.

6.1. MoMo $_{A\pi}$

Since we have only considered a single transition label (τ) in $\mathcal{M}_{A\pi}$, we will just consider a single label predicate \surd . Moreover, we shall use the trivial interpretation function $\mathbb{A}_{A\pi} \llbracket \surd \rrbracket = \{\tau\}$.

MoMo $_{A\pi}$ syntax can be summarised as follows (propositional and recursive formulae are omitted):

$$\phi ::= \dots \mid \bar{a}b \rightarrow \phi \mid \bar{a}b \leftarrow \phi \mid \langle \surd \rangle \phi \mid \exists a.\phi \mid a \textcircled{\text{R}} \phi \mid \bigvee a\phi.$$

The interpretation of these formulae, which is obtained by considering $\mathcal{M}_{A\pi}$ as a model in the definition of Table 11, can be described informally as follows:

- $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \bar{a}b \rightarrow \phi$ if and only if $P \equiv \bar{a}b|Q$ and $Q \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \phi$;
- $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \bar{a}b \leftarrow \phi$ if and only if $P|\bar{a}b \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \phi$;
- $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \langle \surd \rangle \phi$ if and only if $P \xrightarrow{\tau} Q$ and $Q \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \phi$;
- $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \exists a.\phi$ if and only if there exists $b \in \text{fn}(P)$ such that $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \phi[b/a]$;
- $P \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} a \textcircled{\text{R}} \phi$ if and only if $P \equiv v a.(Q|\bar{b}a)$ ($a \neq b$) and $Q|\bar{b}a \models_{(\mathcal{M}_{A\pi}, \mathbb{A}_{A\pi})} \phi$;

— $P \models_{\langle \mathcal{M}_{A\pi}, \mathbb{A}_{A\pi} \rangle} \bigvee a\phi$ if and only if there exists $b \notin fn(P)$ and $b \notin fn(\phi)$ such that $P \models_{\langle \mathcal{M}_{A\pi}, \mathbb{A}_{A\pi} \rangle} \phi[b/a]$.

We will now show how the proposed logic can be used to specify properties of the system described in Example 3.1. For this system, we may be interested in guaranteeing the following properties:

PROPI Every message sent over channel a will eventually be emitted over channel b .

PROP2 The order of messages is preserved by the proxy.

The first property can be rephrased as *if we let c be a new name, then if c is made available over channel a (by creating resource $\bar{a}c$), then sooner or later resource $\bar{b}c$ can be consumed*. In the logic, this can be formalised as

$$\phi_1 \stackrel{\text{def}}{=} \bigvee c \bar{a}c \leftarrow \text{Evn}(\bar{b}c \rightarrow \mathbf{true}) \tag{4}$$

where

$$\text{Evn}(\phi) \stackrel{\text{def}}{=} \nu \kappa. \phi \vee ([\sqrt{\quad}] \kappa \wedge \langle \sqrt{\quad} \rangle \mathbf{true}).$$

It easy to prove that PROXY_1 satisfies (4).

Property **PROP2** can be rephrased as *if we let c and d be two names, then it is not the case that if d is sent over a after c , then c is emitted over b after d* . This property can be rendered in $\text{MoMo}_{A\pi}$ as follows:

$$\phi_1 \stackrel{\text{def}}{=} \neg \bar{a}c \leftarrow \langle \sqrt{\quad} \rangle \bar{a}c \Rightarrow \mathbf{false} \wedge \bar{a}d \leftarrow \phi_2 \tag{5}$$

where

$$\phi_2 \stackrel{\text{def}}{=} \nu \kappa. \bar{b}c \Rightarrow \mathbf{false} \wedge (\bar{b}d \rightarrow \mathbf{true} \vee \langle \sqrt{\quad} \rangle \kappa).$$

Now we have that process PROXY_1 does not satisfy formula (5), because there is no synchronisation between two different instances of the replicated process.

Both **PROPI** and **PROP2** would be satisfied by a different implementation of the $A\pi$ proxy:

$$\text{PROXY}_2 \stackrel{\text{def}}{=} \nu c. (a(x). \tau. (\bar{b}x \bar{c}(c))) | (!c(y). a(x). \tau. (\bar{b}x \bar{c}(y)))$$

where a new instance of the replicated process can start only when the previous one is terminated. This behaviour is guaranteed by the synchronisation over the private name c .

6.2. $\text{MoMo}_{\mu K}$

Let $\mathcal{A}_{\mu K}$ be the set of label predicates α defined as follows:

- $\sqrt{\quad}$, identifying τ labels;
- $l_1 : t \triangleleft l_2$, identifying the execution of an input;
- $l_1 : t \triangleright l_2$, corresponding to an output action;
- $l_1 : \phi \blacktriangleright l_2$, characterising the remote evaluation of a process.

The interpretation of $\sqrt{\quad}$, $l_1 : t \triangleleft l_2$ and $l_1 : t \triangleright l_2$ is trivial: they are satisfied by τ , $l_1 : t \triangleleft l_2$ and $l_1 : t \triangleright l_2$, respectively. The interpretation of $l_1 : \phi \blacktriangleright l_2$ is more complicated. Intuitively, a transition label λ satisfies $l_1 : \phi \blacktriangleright l_2$ if $\lambda = l_1 : P \blacktriangleright l_2$ and $l_2 :: P$ satisfies ϕ . Hence,

to define $\mathbb{A}_{\mu K}$ (the interpretation of $\mathcal{A}_{\mu K}$), we need a mutual recursion because we make use of $\mathbb{M}^{(\mathcal{M}_{\mu K}, \mathbb{A}_{\mu K})}$ when defining $\mathbb{A}_{\mu K}$. Indeed, the interpretation function of $\mathcal{A}_{\mu K}$ is the greatest lower bound of a chain of interpretation functions that depends on the depth of temporal formulae.

Definition 6.1. Let $depth(\phi)$ and $depth(\alpha)$ be defined inductively as follows:

- $depth(\mathbf{T}) = depth(\kappa) = depth(\{n_1 = n_2\}) = 0$;
- $depth(\neg\phi) = depth(\forall\kappa.\phi) = depth(\rho \rightarrow \phi) = depth(\rho \leftarrow \phi) = depth(\exists n.\phi) = depth(n \mathbb{R} \phi) = depth(\bigvee n\phi) = depth(\phi)$
- $depth(\langle\alpha\rangle\phi) = \max\{depth(\alpha), depth(\phi)\}$
- $depth(\phi_1 \wedge \phi_2) = \max\{depth(\phi_1), depth(\phi_2)\}$
- $depth(\alpha) = \begin{cases} 1 + depth(\phi) & \text{if } \alpha = l_1 : \phi \blacktriangleright l_2 \\ 0 & \text{otherwise.} \end{cases}$

Definition 6.2. Let $\mathbb{A}[\alpha] = \bigcap_{i \in \mathbb{N}} \mathbb{A}_{\mu K}^i[\alpha]$ where for all $i \geq 0$:

- $\mathbb{A}_{\mu K}^i[\sqrt{\quad}] = \{\tau\}$;
- $\mathbb{A}_{\mu K}^i[l_1 : t \triangleleft l_2] = \{l_1 : t < l_2\}$;
- $\mathbb{A}_{\mu K}^i[l_1 : t \triangleright l_2] = \{l_1 : t > l_2\}$;
- $\mathbb{A}_{\mu K}^0[l_1 : \phi \blacktriangleright l_2] = \{l_1 : P \blacktriangleright l_2 | l_2 :: P \in \text{NET}\}$;
- $\mathbb{A}_{\mu K}^{i+1}[l_1 : \phi \blacktriangleright l_2] = \begin{cases} \{l_1 : P \blacktriangleright l_2 | l_2 :: P \in \mathbb{M}^{\mathbb{A}^i}[\phi]_{\varepsilon_0}\} & \text{if } depth(\phi) \leq i \\ \{l_1 : P \blacktriangleright l_2 | l_2 :: P \in \text{NET}\} & \text{otherwise.} \end{cases}$

To guarantee the well-definedness of \mathbb{A} , we have to prove, for each α , that $\{\mathbb{A}_{\mu K}^i[\alpha]\}_{i \in \mathbb{N}}$ is a chain in the complete lattice $2^{\text{LAB}_{\mu K}}$. Namely, that for each i , we have $\mathbb{A}_{\mu K}^{i+1}[\alpha] \subseteq \mathbb{A}_{\mu K}^i[\alpha]$.

Lemma 6.3. The following assertions hold:

- For each α and for each i and j such that $depth(\alpha) \leq i$ and $depth(\alpha) \leq j$, we have

$$\mathbb{A}_{\mu K}^i[\alpha] = \mathbb{A}_{\mu K}^j[\alpha].$$

- For each ϕ and for each i and j such that $depth(\phi) \leq i$ and $depth(\phi) \leq j$, we have

$$\mathbb{M}^{\mathbb{A}^i}[\phi] = \mathbb{M}^{\mathbb{A}^j}[\phi].$$

Proof. The proof is by induction on $depth(\alpha)$ and $depth(\phi)$.

Base of induction: It is easy to prove that if $depth(\alpha) = 0$, then for each i and j

$$\mathbb{A}_{\mu K}^i[\alpha] = \mathbb{A}_{\mu K}^j[\alpha].$$

Similarly, if $depth(\phi) = 0$, then ϕ does not contain label predicates of the form $l_1 : \psi \blacktriangleright l_2$. Hence,

$$\mathbb{M}^{\mathbb{A}^i}[\phi] = \mathbb{M}^{\mathbb{A}^j}[\phi].$$

Induction hypothesis: We assume that:

- For each α , with $depth(\alpha) \leq k$, and for each i and j such that $depth(\alpha) \leq i$ and $depth(\alpha) \leq j$, we have

$$\mathbb{A}_{\mu K}^i[\alpha] = \mathbb{A}_{\mu K}^j[\alpha].$$

- For each ϕ with $depth(\phi) \leq k$, and for each i and j such that $depth(\phi) \leq i$ and $depth(\phi) \leq j$, we have

$$\mathbb{M}^{A^i} \llbracket \phi \rrbracket = \mathbb{M}^{A^j} \llbracket \phi \rrbracket.$$

Induction step: Let α be such that $depth(\alpha) = k + 1$. Hence, $\alpha = l_1 : \psi \blacktriangleright l_2$ for some l_1, l_2 and ψ such that $depth(\psi) = k$. For each $i \geq k + 1$, we have

$$\mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket = \left\{ l_1 : P \blacktriangleright l_2 | l_2 :: P \in \mathbb{M}^{A^{i-1}} \llbracket \psi \rrbracket \right\} = (A).$$

By the induction hypothesis, we have that for each $j \geq k + 1$,

$$(A) = \left\{ l_1 : P \blacktriangleright l_2 | l_2 :: P \in \mathbb{M}^{A^{j-1}} \llbracket \psi \rrbracket \right\} = \mathbb{A}_{\mu K}^j \llbracket \alpha \rrbracket.$$

Similarly, let ϕ be such that $depth(\phi) = k + 1$. Without loss of generality we let $\phi = \langle \alpha \rangle \psi$ with $depth(\alpha) = k + 1$ and $depth(\psi) \leq k$. Let $i \geq k + 1$. We have

$$\mathbb{M}^{A^i} \llbracket \langle \alpha \rangle \psi \rrbracket = \{ N | \exists \lambda, N' : N \xrightarrow{\lambda} N', \lambda \in \mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket, \text{ and } N' \in \mathbb{M}^{A^i} \llbracket \psi \rrbracket \}.$$

Using the induction hypothesis for each $j \geq k$, we have $\mathbb{M}^{A^j} \llbracket \psi \rrbracket = \mathbb{M}^{A^j} \llbracket \psi \rrbracket$. Moreover, we have shown that, if $depth(\alpha) = k + 1$, then for each $j \geq k + 1$ $\mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket = \mathbb{A}_{\mu K}^j \llbracket \alpha \rrbracket$. Thus, $\mathbb{M}^{A^i} \llbracket \langle \alpha \rangle \psi \rrbracket = \mathbb{M}^{A^j} \llbracket \langle \alpha \rangle \psi \rrbracket$, and the statement follows. \square

Lemma 6.4. For each $i \geq 0$ and for each α , we have $\mathbb{A}_{\mu K}^{i+1} \llbracket \alpha \rrbracket \subseteq \mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket$.

Proof. We can distinguish three cases:

- If $depth(\alpha) \leq i$, then for each i and j , $\mathbb{A}_{\mu K}^{i+1} \llbracket \alpha \rrbracket = \mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket$ (Lemma 6.3).
- If $depth(\alpha) = i + 1$, then $\alpha = l_1 : \phi \blacktriangleright l_2$ for some l_1, l_2 and ϕ such that $depth(\phi) = i$. We have:

$$\mathbb{A}_{\mu K}^{i+1} \llbracket \alpha \rrbracket = \{ l_1 : P \blacktriangleright l_2 | l_2 :: P \in \mathbb{M}^{A^i} \llbracket \phi \rrbracket \varepsilon_0 \} \subseteq \{ l_1 : P \blacktriangleright l_2 | l_2 :: P \in \text{NET} \} = \mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket.$$

- If $depth(\alpha) > i + 1$, then we have

$$\mathbb{A}_{\mu K}^{i+1} \llbracket \alpha \rrbracket = \mathbb{A}_{\mu K}^i \llbracket \alpha \rrbracket = \{ l_1 : P \blacktriangleright l_2 | l_2 :: P \in \text{NET} \}. \quad \square$$

The *state fragment* of $\text{MoMo}_{\mu K}$ ($r \rightarrow \phi$ and $r \leftarrow \phi$) relates formulae satisfaction to the creation or consumption of resources and enables us to describe the *spatial* properties of a system. On the other hand, the *temporal fragment* ($\langle \alpha \rangle \phi$) enables us to describe properties concerning *interaction protocols* among system components and to *control* properties of agents migrating over the net.

We can use $\text{MoMo}_{\mu K}$ to specify properties of the proxy specifications in μKLAIM presented in Example 4.1. As with the $A\pi$ specification, we are interested in guaranteeing the following properties:

- PROPI** Every message sent over channel a will eventually be emitted over channel b .
- PROP2** The order of messages is preserved by the proxy.

These can be rendered in $\text{MoMo}_{\mu K}$ using the following formulae:

$$\bigvee l \langle l \rangle @ l_a \leftarrow \text{Evn}(\langle l \rangle @ l_b \rightarrow \text{true}) \quad (6)$$

where

$$\begin{aligned}
 \text{Evn}(\phi) &\stackrel{\text{def}}{=} v\kappa.\phi \vee ([\surd] \kappa \wedge \langle \surd \rangle \text{true}) \\
 \neg \langle l_1 \rangle @ l_a \leftarrow \langle \surd \rangle \langle l_1 \rangle @ l_a &\Rightarrow \text{false} \wedge \langle l_2 \rangle @ l_a \leftarrow \phi_2
 \end{aligned} \tag{7}$$

where

$$\phi_2 \stackrel{\text{def}}{=} v\kappa.\langle l_1 \rangle @ l_b \Rightarrow \text{false} \wedge (\langle l_2 \rangle @ l_b \rightarrow \text{true} \vee) \langle \surd \rangle \kappa.$$

It easy to prove that κPROXY_2 (see specification (2) in Section 4) satisfies both (6) and (7) while κPROXY_1 (see specification (3) in Section 4) satisfies only (6).

Note that the formalisation of properties PROP1 and PROP2 in $\text{MoMo}_{\mu K}$ (6 and 7) and $\text{MoMo}_{A\pi}$ (4 and 5) coincide apart from the resource specification. Indeed, while in the $A\pi$ specification a resource is a channel with a value ($\bar{a}c$), in μKLAIM it is a located tuple ($\langle t \rangle @ l$). We thus have that the properties of a generic proxy can be described in MoMo while ignoring specific system implementations.

$\text{MoMo}_{\mu K}$ can also be used for specifying and verifying properties of Example 4.2. For instance, a possible property to establish is that the print server handles all requests. In other words, one could require that if a client retrieves a print slot and sends a document for printing, then sooner or later the document will be printed and the client will be notified. This property can be specified by the following formula:

$$\begin{aligned}
 \phi &\stackrel{\text{def}}{=} (\text{PrintSlot}) @ \text{PrintServer} \rightarrow & (1) \\
 &l \leftarrow & (2) \\
 &(l, \text{"test"}) @ \text{PrintServer} \leftarrow & (3) \\
 &\text{Evn}(\text{PrintOk} @ l \rightarrow \mathbf{T}) & (4)
 \end{aligned}$$

This can be read as:

- 1 Let (PrintSlot) be a tuple available at PrintServer ,
- 2 and let l be the locality of a print client,
- 3 if $(l, \text{'test'})$ is inserted at PrintServer ,
- 4 then eventually tuple (PrintOK) will be at l .

7. Logical and behavioural equivalences

In the previous section we showed how dialects of MoMo logics can be used for specifying and verifying properties of mobile and distributed applications specified using calculi like $A\pi$ and μKLAIM .

An alternative method for proving properties of process calculi is the one based on behavioural equivalences that requires us to provide a concrete and an abstract specification of the behaviour of a given system and then establish that they are ‘indistinguishable’ under appropriate assumptions. Equivalences also turn out to be important because they would enable us to determine when parts of a system can be *replaced* without changing the behaviour of the whole specification. If one considers the proxy specified in Example 3.1, one could imagine that the τ action, which is performed after a message is retrieved over channel a , models an internal behaviour and the system could be refined with an

alternative implementation:

$$\text{PROXY}_3 = !v c.a(x).\bar{c}x|c(y).\bar{b}y.$$

Obviously, formulae satisfaction also induces an equivalence on the interpretation model and two system descriptions are equivalent if (and only if) they satisfy the same set of formulae. A classical result relating modal logic and behavioural equivalence was given in Hennessy and Milner (1985) and shows that the equivalence induced by a very simple modal logic coincides with the so-called bisimulation equivalence.

Correspondence results between differently characterised relations are important in many ways. They help us to gain confidence in the proposed formalisation, enable us to use minimised models when checking formulae satisfaction and help us in studying the properties of the more convenient variant of the equivalent specification.

In this section we shall provide a behavioural characterisation of the logical equivalence induced by MoMo on an MLTS, and show that it coincides with a natural adaptation of the bisimulation relations built on the three transition relations of our model.

The proposed behavioural equivalence is defined over the states of an MLTS. We first define three relations that characterise the bisimulations over the three relations associated with each MLTS.

Definition 7.1. Let \mathcal{M} be an MLTS.

1 A relation $\mathcal{R} \subseteq \mathcal{M}^S \times \mathcal{M}^S$ is a resource preserving bisimulation if and only if \mathcal{R} is symmetric and for each $(s_1, s_2) \in \mathcal{R}$ and $r \in \mathcal{M}^R$, we have:

- If $s_1 \xrightarrow{\ominus r} s'_1$, then there exists s'_2 such that $s_2 \xrightarrow{\ominus r} s'_2$ and $(s'_1, s'_2) \in \mathcal{R}$.
- If $s_1 \xrightarrow{\ominus r} s'_1$, then there exists s'_2 such that $s_2 \xrightarrow{\ominus r} s'_2$ and $(s'_1, s'_2) \in \mathcal{R}$.

We use $\sim_{\mathcal{R}}$ to denote the largest resource preserving bisimulation.

2 A relation $\mathcal{R} \subseteq \mathcal{M}^S \times \mathcal{M}^S$ is a revelation bisimulation if and only if \mathcal{R} is symmetric and for each $(s_1, s_2) \in \mathcal{R}$ and $n \in \mathcal{M}^N$, we have:

- If $s_1 \xrightarrow{\hookrightarrow^n} s'_1$, then there exists s'_2 such that $s_2 \xrightarrow{\hookrightarrow^n} s'_2$ and $(s'_1, s'_2) \in \mathcal{R}$.

We use $\sim_{\mathcal{N}}$ to denote the largest revelation bisimulation.

3 A relation $\mathcal{R} \subseteq \mathcal{M}^S \times \mathcal{M}^S$ is an \mathcal{A} -parameterised bisimulation whenever $\mathcal{A} \subseteq \mathcal{M}^L \times \mathcal{M}^L$, if and only if \mathcal{R} is symmetric and for each $(s_1, s_2) \in \mathcal{R}$ and $\lambda_1 \in \mathcal{M}^L$, we have:

- If $s_1 \xrightarrow{\lambda_1} s'_1$, then there exist λ_2 and s'_2 such that $s_2 \xrightarrow{\lambda_2} s'_2$, $(\lambda_1, \lambda_2) \in \mathcal{A}$ and $(s'_1, s'_2) \in \mathcal{R}$;

We use $\sim_{\mathcal{L}}^{\mathcal{A}}$ to denote the largest \mathcal{A} -parameterised bisimulation.

The required bisimulation for an MLTS is obtained as the intersection of the three relations we have just introduced.

Definition 7.2. Let \mathcal{M} be an MLTS and $\mathcal{A} \subseteq \mathcal{M}^L \times \mathcal{M}^L$. We define $\sim^{\mathcal{A}} = \sim_{\mathcal{R}} \cap \sim_{\mathcal{N}} \cap \sim_{\mathcal{L}}^{\mathcal{A}}$.

Let \mathcal{M} be an MLTS and \mathcal{A} be a subset of $\mathcal{M}^L \times \mathcal{M}^L$. We will now show how $\sim^{\mathcal{A}}$ can characterise the behavioural equivalence induced by formulae satisfaction $\mathbb{L}_{\mathcal{A}}$. We will first introduce the following technical lemma.

Lemma 7.3. If $\mathcal{A}_1 \subseteq \mathcal{A}_2$, then for each s_1 and s_2 , $s_1 \sim^{\mathcal{A}_1} s_2 \Rightarrow s_1 \sim^{\mathcal{A}_2} s_2$.

Proof. We prove that if $\mathcal{A}_1 \subseteq \mathcal{A}_2$, then $\sim^{\mathcal{A}_1}$ is an \mathcal{A}_2 -parameterised bisimulation.

Let $s_1 \sim^{\mathcal{A}_1} s_2$. Then, if $s_1 \xrightarrow{\lambda_1} s'_1$, there exists λ_2 and s'_2 such that $s_2 \xrightarrow{\lambda_2} s'_2$, $(\lambda_1, \lambda_2) \in \mathcal{A}_1$ and $s'_1 \sim^{\mathcal{A}_1} s'_2$. Since $\mathcal{A}_1 \subseteq \mathcal{A}_2$, we have $(\lambda_1, \lambda_2) \in \mathcal{A}_2$ and $\sim^{\mathcal{A}_1}$ is an \mathcal{A}_2 -parameterised bisimulation and $\sim^{\mathcal{A}_1} \subseteq \sim^{\mathcal{A}_2}$. \square

We now show that for each s_1 and s_2 , if $s_1 \sim^{\mathcal{A}} s_2$, then s_1 and s_2 satisfy the same set of formulae in $\mathbb{L}_{\mathcal{A}}$, under the assumption that relation \mathcal{A} is *respectful of* interpretation \mathbb{A} , where ‘respectful of’ is defined as follows.

Definition 7.4. Let $\mathcal{A} \subseteq \mathcal{M}^{\mathcal{L}} \times \mathcal{M}^{\mathcal{L}}$, \mathcal{A} be a set of label predicates and $\mathbb{A} : \mathcal{A} \rightarrow 2^{\mathcal{M}^{\mathcal{L}}}$ be the interpretation function. We say that

— \mathcal{A} is respectful of \mathbb{A} if and only if $\forall \alpha \in \mathcal{A}$,

$$(\lambda_1, \lambda_2) \in \mathcal{A} \text{ if and only if } \lambda_1 \in \mathbb{A}[\![\alpha]\!] \Leftrightarrow \lambda_2 \in \mathbb{A}[\![\alpha]\!].$$

Theorem 7.5. Let \mathcal{M} be an MLTS, \mathcal{A} be a set of label predicates, $\mathbb{A} : \mathcal{A} \rightarrow 2^{\mathcal{M}^{\mathcal{L}}}$ be the interpretation function, and $\mathcal{A} \subseteq \mathcal{M}^{\mathcal{L}} \times \mathcal{M}^{\mathcal{L}}$ be respectful of \mathbb{A} . Then

$$s_1 \sim^{\mathcal{A}} s_2 \implies \forall \phi \in \mathbb{L}_{\mathcal{A}}. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \text{ iff } s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi.$$

Proof. We prove the result for the recursion-free fragment of the logic, denoted by $\mathbb{L}_{\mathcal{A}}^-$. Indeed, it is easy to prove that

$$\forall \phi \in \mathbb{L}_{\mathcal{A}}. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$$

if and only if

$$\forall \phi \in \mathbb{L}_{\mathcal{A}}^-. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi.$$

Let:

- $\mathcal{R}_{\mathbb{L}} = \{(s_1, s_2) \mid \forall \phi \in \mathbb{L}_{\mathcal{A}}^-. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi\}$;
- $size(\phi)$ be defined inductively as follows:

- $size(\mathbf{T}) = 0$
- $size(\phi_1 \wedge \phi_2) = 1 + \max\{size(\phi_1), size(\phi_2)\}$
- $size(\neg\phi) = 1 + size(\phi)$
- $size(\rho \rightarrow \phi) = 1 + size(\phi)$
- $size(\rho \leftarrow \phi) = 1 + size(\phi)$
- $size(\langle \alpha \rangle \phi) = 1 + size(\phi)$
- $size(\exists n. \phi) = 1 + size(\phi)$
- $size(\{n_1 = n_2\}) = 0$
- $size(n \textcircled{R} \phi) = 1 + size(\phi)$
- $size(\bigvee n \phi) = 1 + size(\phi)$.

We will prove by induction on $size(\phi)$ that for each $\phi \in \mathbf{L}_{\mathcal{A}}^-$, if $s_1 \sim^{\mathcal{A}} s_2$, then $s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$ if and only if $s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$.

Base of induction: We have to prove that for each ϕ such that $size(\phi) = 0$, if $s_1 \sim^{\mathcal{A}} s_2$, then $s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$ if and only if $s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$. We can distinguish two cases:

- $\phi = \mathbf{T}$;
- $\phi = \{n_1 = n_2\}$.

In both cases the statement follows trivially.

Induction hypothesis: We assume that for each ϕ such that $size(\phi) \leq n$, if $s_1 \sim^{\mathcal{A}} s_2$, then $s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$ if and only if $s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$.

Induction step: Let ϕ be such that $size(\phi) = n + 1$. We prove that if $s_1 \sim^{\mathcal{A}} s_2$, then $s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$ if and only if $s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$. We can distinguish the following cases according to the syntax of ϕ :

- $\phi = \phi_1 \wedge \phi_2$:

$$\begin{aligned} s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_1 \wedge \phi_2 &\Leftrightarrow s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_1 \text{ and } s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_2 \\ &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_1 \text{ and } s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_2 \\ &\quad \text{for the induction hypothesis} \\ &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_1 \wedge \phi_2. \end{aligned}$$

- $\phi = \rho \rightarrow \phi$:

$$\begin{aligned} s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \rho \rightarrow \phi &\Leftrightarrow \exists s'_1, s_1 \xrightarrow{\ominus \rho} s'_1 \text{ and } s'_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\ &\Leftrightarrow \exists s'_2, s_2 \xrightarrow{\ominus \rho} s'_2 \text{ and } s'_1 \sim^{\mathcal{A}} s'_2 \\ &\quad \text{since } \sim^{\mathcal{A}} \text{ is a resource preserving bisimulation} \\ &\Leftrightarrow s'_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\ &\quad \text{for the induction hypothesis} \\ &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \rho \rightarrow \phi. \end{aligned}$$

- $\phi = \rho \leftarrow \phi$:

$$\begin{aligned} s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \rho \leftarrow \phi &\Leftrightarrow \exists s'_1, s_1 \xrightarrow{\oplus \rho} s'_1 \text{ and } s'_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\ &\Leftrightarrow \exists s'_2, s_2 \xrightarrow{\oplus \rho} s'_2 \text{ and } s'_1 \sim^{\mathcal{A}} s'_2 \\ &\quad \text{since } \sim^{\mathcal{A}} \text{ is a resource preserving bisimulation} \\ &\Leftrightarrow s'_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\ &\quad \text{for the induction hypothesis} \\ &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \rho \leftarrow \phi. \end{aligned}$$

— $\phi = \langle \alpha \rangle \phi$:

$$\begin{aligned}
 s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \langle \alpha \rangle \phi &\Leftrightarrow \exists s'_1. s_1 \xrightarrow{\lambda_1} s'_1, \lambda \in \mathbb{A}[\![\alpha]\!] \\
 &\quad \text{and } s'_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\
 &\Leftrightarrow \exists \lambda_2, s'_2. s_2 \xrightarrow{\lambda_2} s'_2, \lambda_2 \in \mathbb{A}[\![\alpha]\!] \\
 &\quad \text{and } s'_1 \sim^{\mathcal{A}} s'_2
 \end{aligned}$$

since $\sim^{\mathcal{A}}$ is an \mathcal{A} -parametrised bisimulation and \mathcal{A} is respectful of \mathbb{A}

$$\Leftrightarrow s'_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi$$

for the induction hypothesis

$$\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \langle \alpha \rangle \phi.$$

— $\phi = \exists n. \phi$:

$$\begin{aligned}
 s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \exists n. \phi &\Rightarrow s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi[n'/n] \text{ for some } n' \in \mathcal{M}^{\mathcal{N}} \\
 &\Rightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi[n'/n] \text{ for the induction hypothesis} \\
 &\Rightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \exists n. \phi.
 \end{aligned}$$

— $\phi = n \textcircled{\mathbb{R}} \phi$:

$$\begin{aligned}
 s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} n \textcircled{\mathbb{R}} \phi &\Leftrightarrow s_1 \hookrightarrow^n s'_1 \text{ and } s'_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\
 &\Leftrightarrow s_2 \hookrightarrow^n s'_2 \text{ and } s'_1 \sim^{\mathcal{A}} s'_2 \\
 &\quad \text{since } s'_1 \sim^{\mathcal{A}} s'_2 \text{ is a revelation bisimulation} \\
 &\Leftrightarrow s'_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \quad \text{for the induction hypothesis} \\
 &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} n \textcircled{\mathbb{R}} \phi.
 \end{aligned}$$

— $\phi = \bigvee n \phi$:

$$\begin{aligned}
 s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \bigvee n \phi &\Leftrightarrow s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi[n'/n] \text{ for some } n' \notin \eta_S(s_1) \\
 &\Leftrightarrow s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi[n''/n] \text{ for some } n'' \notin \eta_S(s_1) \sup \eta_S(s_2) \\
 &\quad \text{and } n'' \text{ not occurring in } \phi \text{ (see Lemma 5.1)} \\
 &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi[n''/n] \text{ for the induction hypothesis} \\
 &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \bigvee n \phi.
 \end{aligned}$$

— $\phi = \neg \phi$:

$$\begin{aligned}
 s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \neg \phi &\Leftrightarrow s_1 \not\models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \\
 &\Leftrightarrow s_2 \not\models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \quad \text{for the induction hypothesis} \\
 &\Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \neg \phi.
 \end{aligned}$$

□

Let $\mathcal{R}_{\mathbb{L}} \subseteq \mathcal{M}^S \times \mathcal{M}^S$ be the symmetric relation defined by

$$\mathcal{R}_{\mathbb{L}} = \{(s_1, s_2) \mid \forall \phi \in \mathbb{L}_{\mathcal{A}}^-. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \Leftrightarrow s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi\}.$$

We want to prove that $\mathcal{R}_{\mathbb{L}}$ is a subset of $\sim^{\mathcal{A}}$.

First, we have to guarantee that \mathcal{M} is tractable, namely that for each state in \mathcal{M}^S , only a finite set of next states has to be considered.

Definition 7.6. An MLTS \mathcal{M} is *image-finite* with respect to the interpretation function $\mathbb{A} : \mathcal{A} \rightarrow 2^{\mathcal{M}^{\mathcal{L}}}$ if and only if for each state s_1 , resource r , name n and label predicate α , the following sets are finite:

- $\{s_2 | s_1 \xrightarrow{\oplus r} s_2\}$
- $\{s_2 | s_1 \xrightarrow{\ominus r} s_2\}$
- $\{s_2 | s_1 \xrightarrow{n} s_2\}$
- $\{s_2 | \exists \lambda \in \mathbb{A}[\alpha]. s_1 \xrightarrow{\lambda} s_2\}$.

Moreover, we also need \mathbb{A} and \mathcal{A} to *discriminate* between the different transitions performed at each state.

Definition 7.7. Let \mathcal{M} be an MLTS, $\mathcal{A} \subseteq \mathcal{M}^{\mathcal{L}} \times \mathcal{M}^{\mathcal{L}}$, \mathcal{A} be a set of label predicates and $\mathbb{A} : \mathcal{A} \rightarrow 2^{\mathcal{M}^{\mathcal{L}}}$ be the interpretation function. We say that \mathcal{A} and \mathbb{A} are *transition respectful* in \mathcal{M} if and only if for each $s \in \mathcal{M}^{\mathcal{S}}$ and $\lambda \in \mathcal{M}^{\mathcal{L}}$ there exists α in \mathcal{A} such that:

- $\lambda \in \mathbb{A}[\alpha]$
 - for each λ' such that $s \xrightarrow{\lambda'}$,
- $$(\lambda, \lambda') \notin \mathcal{A} \iff \lambda' \notin \mathbb{A}[\alpha]$$

where $s \xrightarrow{\lambda'}$ if there exists s' such that $s \xrightarrow{\lambda'} s'$.

The following theorem enables us to guarantee that if s_1 and s_2 satisfy the same set of formulae, then s_1 is behaviourally equivalent to s_2 .

Theorem 7.8. Let \mathcal{M} be an MLTS, \mathcal{A} be a set of label predicates, $\mathbb{A} : \mathcal{A} \rightarrow 2^{\mathcal{M}^{\mathcal{L}}}$ be the interpretation function and $\mathcal{A} \subseteq \mathcal{M}^{\mathcal{L}} \times \mathcal{M}^{\mathcal{L}}$ such that:

- 1 \mathcal{A} is respectful of \mathbb{A} in \mathcal{M} ;
- 2 \mathcal{A} and \mathbb{A} are transition respectful in \mathcal{M} ;
- 3 \mathcal{M} is *image-finite* with \mathcal{A} and \mathbb{A} .

Then,

$$s_1 \sim^{\mathcal{A}} s_2 \iff \forall \phi \in \mathbb{L}_{\mathcal{A}}. s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi \text{ iff } s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi.$$

Proof. We will prove that $\mathcal{R}_{\mathbb{L}}$ is a resource preserving bisimulation, a revelation bisimulation and an \mathcal{A} -parameterised bisimulation.

Since $\sim_{\mathcal{R}}$, $\sim_{\mathcal{N}}$ and $\sim_{\mathcal{L}}^{\mathcal{A}}$ are the the largest resource preserving bisimulations, revelation bisimulations and \mathcal{A} -parameterised bisimulations, respectively, it follows that

$$\mathcal{R}_{\mathbb{L}} \subseteq \sim_{\mathcal{R}} \cap \sim_{\mathcal{N}} \cap \sim_{\mathcal{L}}^{\mathcal{A}} = \sim^{\mathcal{A}}.$$

We now give the formal proof that $\mathcal{R}_{\mathbb{L}}$ is an \mathcal{A} -parameterised bisimulation: we will omit the similar proofs that $\mathcal{R}_{\mathbb{L}}$ is a resource preserving bisimulation.

Let $(s_1, s_2) \in \mathcal{R}_{\mathbb{L}}$ and $s_1 \xrightarrow{\lambda_1} s'_1$ for some λ_1 . We have to prove that there exists λ_2 and s'_2 such that $(\lambda_1, \lambda_2) \in \mathcal{A}$, $s_2 \xrightarrow{\lambda_2} s'_2$ and $(s'_1, s'_2) \in \mathcal{R}_{\mathbb{L}}$.

We suppose that for each λ_2 and s'_2 such that $(\lambda_1, \lambda_2) \in \mathcal{A}$, we have $(s'_1, s'_2) \notin \mathcal{R}_{\mathbb{L}}$. Let

$$X = \left\{ s'_2 \mid \exists \lambda_2. (\lambda_1, \lambda_2) \in \mathcal{A} \text{ and } s_2 \xrightarrow{\lambda_2} s'_2 \right\}.$$

Since \mathcal{M} is image-finite, X is finite. For each $s \in X$ there exists ϕ_s such that $s'_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \phi_s$ while $s \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \neg \phi_s$. If X is non-empty, we let

$$\phi_X = \bigwedge_{s \in X} \phi_s.$$

And if X is empty, we let $\phi_X = \mathbf{T}$.

\mathcal{A} and \mathbb{A} are transition respectful in \mathcal{M} . Hence, there exists α such that:

- $\lambda_1 \in \mathbb{A}[\![\alpha]\!]$; and
- for each λ such that $s \xrightarrow{\lambda}$, if $(\lambda_1, \lambda) \notin \mathcal{A}$, then $\lambda \notin \mathbb{A}[\![\alpha]\!]$.

We have that $s_1 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \langle \alpha \rangle \phi_X$ while $s_2 \models_{\langle \mathcal{M}, \mathbb{A} \rangle} \neg \langle \alpha \rangle \phi_X$. Hence, $(s_1, s_2) \notin \mathcal{R}_{\mathbb{L}}$, which contradicts the assumptions, so the statement is proved. □

7.1. Logical characterisation of bisimulation for $A\pi$

In this section we will investigate the equivalence induced by the satisfaction of formulae in $\mathbb{L}_{A\pi}$ and their relationships with another well-studied behavioural equivalence between $A\pi$ processes, *viz.* the asynchronous bisimulation presented in Amadio *et al.* (1998). First we note the following result.

Lemma 7.9.

- $\mathcal{M}_{A\pi}$ is image-finite with respect to $\mathcal{A}_{A\pi}$ and $\mathbb{A}_{A\pi}$.
- $\mathcal{A}_{A\pi} = \{(\tau, \tau)\}$ and $\mathbb{A}_{A\pi}$ are transition respectful in $\mathcal{M}_{A\pi}$.

Proof. Image-finiteness of $\mathcal{M}_{A\pi}$ with respect to $\mathcal{A}_{A\pi}$ and $\mathbb{A}_{A\pi}$ follows easily from the following:

- For each P and $\bar{a}b$, we have $\{P' \mid P \xrightarrow{\oplus \bar{a}b} P'\} = \{P \mid \bar{a}b\}$.
- For each P and $\bar{a}b$, we have $\{P' \mid P \xrightarrow{\ominus \bar{a}b} P'\}$ is $\{Q\}$ if $P \equiv Q \parallel \bar{a}b$, otherwise it is the empty set.
- For each P and a , we have $\{P' \mid P \xrightarrow{\hookrightarrow^a} P'\} = \{P' \mid \exists P''. P \equiv \nu a.P' \text{ and } P' \equiv P'' \mid \bar{a}b\}$, which is finite since in P only a finite number of bound names can occur.
- For each P and α , we have $\{P' \mid \exists \lambda. \lambda \in \mathbb{A}[\![\alpha]\!]$ and $P \xrightarrow{\lambda} P'\}$ is equal to $\{P' \mid P \xrightarrow{\tau} P'\}$, which is finite because in P only a finite number of synchronisations can occur.

Moreover, $\mathcal{A}_{A\pi} = \{(\tau, \tau)\}$. Hence, for each P and λ_1 , we have $\lambda_1 \in \mathbb{A}_{A\pi}[\![\sqrt{\quad}]\!]$, while for each $\lambda_2 \in \text{LAB}_{A\pi}$, we have $(\lambda_1, \lambda_2) \in \mathcal{A}_{A\pi}$. Hence, $\mathcal{A}_{A\pi}$ and $\mathbb{A}_{A\pi}$ are transition respectful in $\mathcal{M}_{A\pi}$. □

Moreover, we will show that the equivalence $\sim^{\{(\tau, \tau)\}}$ defined over $A\pi$ processes coincides with the asynchronous bisimulation (\sim_a) defined in Amadio *et al.* (1998). This means that $\mathbb{L}_{\mathcal{A}_{A\pi}}$ characterises asynchronous bisimulation.

Definition 7.10. A symmetric relation \mathcal{R} on $A\pi$ terms is a strong $\sigma\tau$ -bisimulation if $P \mathcal{R} Q$, $P \xrightarrow{\alpha} P'$, α is not an input action, and $bn(\alpha) \cap fn(Q) = \emptyset$ implies $Q \xrightarrow{\alpha} Q'$ and $P' \mathcal{R} Q'$. Let $\sim_{\sigma\tau}$ be the largest $\sigma\tau$ -bisimulation.

Definition 7.11. A relation \mathcal{R} is an asynchronous bisimulation if it is an $\sigma\tau$ -bisimulation and whenever $P \mathcal{R} Q$ and $P \xrightarrow{ab} P'$, we have either:

- $Q \xrightarrow{ab} Q'$ and $P' \mathcal{R} Q'$; or
- $Q \xrightarrow{\tau} Q'$ and $P' \mathcal{R}(Q'|\bar{a}b)$.

Definition 7.12. \sim_a is the largest asynchronous bisimulation.

Definition 7.13. Let \sim_1 be the largest relation \mathcal{R} such that \mathcal{R} is an $\sigma\tau$ -bisimulation and $P \mathcal{R} Q$ implies $(P|\bar{a}b) \mathcal{R}(Q|\bar{a}b)$ for all $\bar{a}b$.

Theorem 7.14. $\sim_a = \sim_1$

Proof. See Amadio *et al.* (1998). □

Let $P = !\tau.0$ and $Q = !\tau.0|a(b).\bar{a}b$. We have that $P \sim_a Q$. Note that these processes can be distinguished using the modal logics defined in Milner *et al.* (1993), Dam (1996) and Caires (2004).

Theorem 7.15. $\sim^{\{(\tau,\tau)\}} \subseteq \sim_a$

Proof. We first prove that $\sim^{\{(\tau,\tau)\}}$ is an $\sigma\tau$ -bisimulation. We have to show that if $P \mathcal{R} Q$, then $P \xrightarrow{\alpha} P'$, α is not an input action, and $bn(\alpha) \cap fn(Q) = \emptyset$ implies $Q \xrightarrow{\alpha} Q'$ and $P' \mathcal{R} Q'$.

Thanks to Lemma 3.4, it easy to prove that if $P \sim^{\{(\tau,\tau)\}} Q$, then:

- $P \xrightarrow{\tau} P' \Rightarrow \exists Q'. Q \xrightarrow{\tau} Q'$ and $P' \sim^{\{(\tau,\tau)\}} Q'$.
- $P \xrightarrow{\ominus \bar{a}b} P' \Rightarrow \exists Q'. Q \xrightarrow{\ominus \bar{a}b} Q'$ and $P' \sim^{\{(\tau,\tau)\}} Q'$.
- $P \xrightarrow{\hookrightarrow^b} P'' \Rightarrow \exists Q', Q''. Q \xrightarrow{\hookrightarrow^b} Q'' \xrightarrow{\ominus \bar{a}b} Q'$ and $P'' \sim^{\{(\tau,\tau)\}} Q'', P' \sim^{\{(\tau,\tau)\}} Q'$.

We now have to prove that for each $\bar{a}b$, if $P \sim^{\{(\tau,\tau)\}} Q$, then $P|\bar{a}b \sim^{\{(\tau,\tau)\}} Q|\bar{a}b$. For each $\bar{a}b$, we have $P \xrightarrow{\ominus \bar{a}b} P|\bar{a}b$ and $Q \xrightarrow{\ominus \bar{a}b} Q|\bar{a}b$. Since $P \sim^{\{(\tau,\tau)\}} Q$, it follows that $P|\bar{a}b \sim^{\{(\tau,\tau)\}} Q|\bar{a}b$ and that $\sim^{\{(\tau,\tau)\}}$ is an asynchronous bisimulation.

Since \sim_a is the largest asynchronous bisimulation, this implies $\sim^{\{(\tau,\tau)\}} \subseteq \sim_a$. □

Theorem 7.16. $\sim_a \subseteq \sim^{\{(\tau,\tau)\}}$

Proof. We prove that \sim_a is a resource preserving bisimulation, a revelation bisimulation and a $\{(\tau,\tau)\}$ -parameterised bisimulation. We have to show that if $P \sim_a Q$, then:

- 1 If $P \xrightarrow{\ominus \bar{a}b} P'$, then there exists Q' such that $Q \xrightarrow{\ominus \bar{a}b} Q'$ and $P' \sim_a Q'$.
- 2 If $P \xrightarrow{\hookrightarrow^a} P'$, then there exists Q' such that $Q \xrightarrow{\hookrightarrow^a} Q'$ and $P' \sim_a Q'$.
- 3 If $P \xrightarrow{a} P'$, then there exists Q' such that $Q \xrightarrow{a} Q'$ and $P' \sim_a Q'$.
- 4 If $P \xrightarrow{\tau} P'$, then there exists Q' such that $Q \xrightarrow{\tau} Q'$ and $P' \sim_a Q'$.

The above follow directly from the definition of \sim_a by noting that:

- 1 If $P \xrightarrow{\oplus \bar{a}b} P'$, then $P' \equiv P|\bar{a}b$.
- 2 If $P \xrightarrow{\ominus \bar{a}b} P'$, then $P \xrightarrow{\bar{a}b} P'$.
- 3 If $P \hookrightarrow^b P'$, then $P' = \bar{a}b|P''$ ($a \neq b$) and $P \xrightarrow{\bar{a}(b)} P''$.
- 4 If $P \xrightarrow{\tau} P'$, then $P \xrightarrow{\tau} P'$. □

7.2. A bisimulation for $\mathcal{M}_{\mu K}$

In this section we introduce an equivalence between μK LAIM transition labels ($\mathcal{A}_{\mu K}$). This will be used to define a bisimulation for $\mathcal{M}_{\mu K}$ ($\sim^{\mathcal{A}_{\mu K}}$) that agrees with the equivalence induced by the satisfaction of formulae in $\mathbb{L}_{\mu K}$.

We will define $\mathcal{A}_{\mu K}$ as the *inf* of a chain of equivalences ($\mathcal{A}_{\mu K}^i$), each of which is respectful of the corresponding interpretation \mathbf{A}^i (see Definition 6.2).

Definition 7.17. Let $\mathcal{A}_{\mu K}^i$ be the subset of $\text{LAB}_{\mu K} \times \text{LAB}_{\mu K}$ defined inductively as follows:

- $(\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^0$ if and only if either:
 - $\lambda_1 = \lambda_2$; or
 - $\lambda_1 = l_1 : P \blacktriangleright l_2$ and $\lambda_2 = l_1 : Q \blacktriangleright l_2$ for some l_1, l_2, P and Q .
- $(\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^{i+1}$ if and only if either:
 - $\lambda_1 = \lambda_2$; or
 - $\lambda_1 = l_1 : P \blacktriangleright l_2$ and $\lambda_2 = l_1 : Q \blacktriangleright l_2$ for some l_1, l_2, P and Q , and $l_2 :: P \sim^{\mathcal{A}_{\mu K}^i} l_2 :: Q$.

We will now show that each $\mathcal{A}_{\mu K}^i$ is respectful of \mathbf{A}^i (see Definition 6.2), and that $\mathcal{A}_{\mu K}^i$ and \mathbf{A}^i are transition respectful in $\mathcal{M}_{\mu K}$. This allows us to prove (using Theorems 7.5 and 7.8) that for each i :

$$N_1 \sim^{\mathcal{A}_{\mu K}^i} N_2 \text{ if and only if } N_1 \models_{\mathbf{A}^i} \phi \Leftrightarrow N_2 \models_{\mathbf{A}^i} \phi.$$

Indeed, the following Lemma holds.

Lemma 7.18. $\mathcal{M}_{\mu K}$ is *image-finite* with respect to $\mathbf{A}_{\mu K}$.

Proof. It is easy to prove that for each N_1 :

- $\{N_2|N_1 \xrightarrow{\oplus r} N_2\}$
- $\{N_2|N_1 \xrightarrow{\ominus r} N_2\}$
- $\{N_2|N_1 \hookrightarrow^n N_2\}$
- $\{N_2|\exists \lambda \in \mathbf{A}[\![\alpha]\!] . N_1 \xrightarrow{\lambda} N_2\}$. □

Lemma 7.19. For each i :

- $\mathcal{A}_{\mu K}^i$ is respectful of \mathbf{A}^i .
- $\mathcal{A}_{\mu K}^i$ and \mathbf{A}^i are transition respectful in $\mathcal{M}_{\mu K}$.

Proof. The proof is by induction on i .

Base of induction ($i=0$): If $(\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^0$, then either $\lambda_1 = \lambda_2 \in \{\tau, l_1 : t \triangleright l_2, l_1 : t \triangleleft l_2\}$ (for some l_1, l_2 and t), or $\lambda_1 = l_1 : P \blacktriangleright l_2$ and $\lambda_2 = l_1 : Q \blacktriangleright l_2$ (for some l_1, l_2, P and Q). In the first case we let α be $\surd, l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$, respectively: $\mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket = \{\lambda_1\}$ and for each $\alpha' \neq \alpha, \lambda_1 \notin \mathbf{A}_{\mu K}^0 \llbracket \alpha' \rrbracket$.

In the second case we have that for each $\phi, \lambda_1, \lambda_2 \in \mathbf{A}_{\mu K}^0 \llbracket l_1 : \phi \blacktriangleright l_2 \rrbracket$, while neither λ_1 nor λ_2 belongs to $\mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket$ if α is $\surd, l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$.

Hence, $(\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^0$ if and only if for each $\alpha \in \mathcal{A}_{\mu K}$ we have $\lambda_1 \in \mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket \Leftrightarrow \lambda_2 \in \mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket$.

We now have to prove that for each N_1 and λ_1 there exists α such that:

— $\lambda_1 \in \mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket$; and

— for each λ such that $N_1 \xrightarrow{\lambda}$, if $(\lambda_1, \lambda) \notin \mathcal{A}_{\mu K}^0$, then $\lambda_2 \notin \mathbf{A}_{\mu K}^0 \llbracket \alpha \rrbracket$.

The statement follows easily by considering $\alpha = \surd$ if $\lambda_1 = \tau, \alpha = l_1 : t \triangleright l_2$ if $\lambda_1 = l_1 : t \triangleright l_2, \alpha = l_1 : t \triangleleft l_2$ if $\lambda_1 = l_1 : t \triangleleft l_2$ and $\alpha = l_1 : \mathbf{T} \blacktriangleright l_2$ if $\lambda_1 = l_1 : t \triangleleft l_2$.

Induction hypothesis: For each $i \leq n$:

— $\mathcal{A}_{\mu K}^i$ is respectful of \mathbf{A}^i .

— $\mathcal{A}_{\mu K}^i$ and \mathbf{A}^i are transition respectful in $\mathcal{M}_{\mu K}$.

Induction step: Let $(\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^{n+1}$. We have that either $\lambda_1 = \lambda_2$ and there exists α such that $\mathbf{A}_{\mu K}^{n+1} \llbracket \alpha \rrbracket = \{\lambda_1\}$ (and for each $\alpha' \neq \alpha, \lambda_1 \notin \mathbf{A}_{\mu K}^{n+1} \llbracket \alpha' \rrbracket$), or $\lambda_1 = l_1 : P \blacktriangleright l_2, \lambda_2 = l_1 : Q \blacktriangleright l_2$ and $l_2 :: P \sim^{\mathcal{A}_{\mu K}^n} l_2 :: Q$.

In the first case the statement follows directly from definition of \mathbf{A}^{n+1} .

In the second case, using the induction hypothesis, we have that for each $\phi, l_2 :: P \in \mathbb{M}^{\mathcal{A}_{\mu K}^n} \llbracket \phi \rrbracket$ if and only if $l_2 :: Q \in \mathbb{M}^{\mathcal{A}_{\mu K}^n} \llbracket \phi \rrbracket$. Hence, for each ϕ such that $depth(\phi) \leq n$, we have $l_1 : P \blacktriangleright l_2 \in \mathbf{A}_{\mu K}^{n+1} \llbracket \phi \rrbracket \Leftrightarrow l_1 : Q \blacktriangleright l_2$.

We now prove that for each N_1 and λ_1 there exists an α such that:

— $\lambda_1 \in \mathbf{A}_{\mu K}^{n+1} \llbracket \alpha \rrbracket$; and

— for each λ such that $N_1 \xrightarrow{\lambda}$, if $(\lambda_1, \lambda) \notin \mathcal{A}_{\mu K}^{n+1}$, then $\lambda_2 \notin \mathbf{A}_{\mu K}^{n+1} \llbracket \alpha \rrbracket$.

If λ_1 is one of $\tau, l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$ (for some l_1, t or l_2), the statement follows easily by considering $\surd, l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$.

If $\lambda_1 = l_1 : P \blacktriangleright l_2$, we let X be the finite set defined by

$$X = \{l_1 : Q \blacktriangleright l_2 \mid \exists N_2 : N_1 \xrightarrow{\lambda_2} N_2 \text{ and } l_2 :: P \sim^{\mathcal{A}_{\mu K}^n} l_2 :: Q\}.$$

From the induction hypotheses, for each $\lambda = l_1 : Q \blacktriangleright l_2 \in X$ there exists ϕ_λ such that $l_2 :: P \models_{\mathbf{A}^n} \phi$ and $l_2 :: Q \models_{\mathbf{A}^n} \neg \phi$. We let $\alpha = l_1 : \phi_X \blacktriangleright l_2$ where $\phi_X = \bigwedge_{\lambda \in X} \phi_\lambda$. Therefore:

— $\lambda_1 \in \mathbf{A}_{\mu K}^{n+1} \llbracket \alpha \rrbracket$; and

— for each λ such that $N_1 \xrightarrow{\lambda}$, if $(\lambda_1, \lambda_2) \notin \mathcal{A}_{\mu K}^{n+1}$, then $\lambda_2 \notin \mathbf{A}_{\mu K}^{n+1} \llbracket \alpha \rrbracket$. □

Let $\mathcal{A}_{\mu K}$ be the greatest lower bound of the chain $\mathcal{A}_{\mu K}^i$ in the CPO $2^{\mathbf{LAB}_{\mu K} \times \mathbf{LAB}_{\mu K}}$. Since $(2^{\mathbf{LAB}_{\mu K} \times \mathbf{LAB}_{\mu K}}, \subseteq)$ is a complete lattice, all we have to do to guarantee the well-definedness of $\mathcal{A}_{\mu K}$ is prove that for each i we have $\mathcal{A}_{\mu K}^{i+1} \subseteq \mathcal{A}_{\mu K}^i$.

Lemma 7.20. For each i we have $\mathcal{A}_{\mu K}^{i+1} \subseteq \mathcal{A}_{\mu K}^i$.

Proof. The thesis follows directly from the definition of $\mathcal{A}_{\mu K}^i$ and Lemma 6.3. □

Definition 7.21. We define $\mathcal{A}_{\mu K} = \bigcap_i \mathcal{A}_{\mu K}^i$.

The following lemmas enable us to guarantee that $\mathcal{A}_{\mu K}$ is respectful of $\mathbb{A}_{\mu K}$ and that $\mathcal{A}_{\mu K}$ and $\mathbb{A}_{\mu K}$ are transition respectful in $\mathcal{M}_{\mu K}$. Hence, $\sim^{\mathcal{A}_{\mu K}}$ characterises the equivalence induced by the satisfaction of formulae in $\mathbb{L}_{\mu K}$.

Lemma 7.22. $\mathcal{A}_{\mu K}$ is respectful of $\mathbb{A}_{\mu K}$.

Proof. We have

$$\begin{aligned} (\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K} &\Rightarrow \forall i. (\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^i \\ &\Rightarrow \forall i. \forall \alpha. \lambda_1 \in \mathbb{A}_{\mu K}^i[\alpha] \Leftrightarrow \lambda_2 \in \mathbb{A}_{\mu K}^i[\alpha] \\ &\Rightarrow \forall \alpha. (\forall i. \lambda_1 \in \mathbb{A}_{\mu K}^i[\alpha]) \Leftrightarrow (\forall i. \lambda_2 \in \mathbb{A}_{\mu K}^i[\alpha]) \\ &\Rightarrow \forall \alpha. \lambda_1 \in \mathbb{A}[\alpha] \Leftrightarrow \lambda_2 \in \mathbb{A}[\alpha]. \end{aligned}$$

Moreover,

$$\begin{aligned} (\lambda_1, \lambda_2) \notin \mathcal{A}_{\mu K} &\Rightarrow \exists k. (\lambda_1, \lambda_2) \in \mathcal{A}_{\mu K}^k \\ &\Rightarrow \exists k. \exists \alpha. \lambda_1 \in \mathbb{A}_{\mu K}^k[\alpha] \wedge \lambda_2 \notin \mathbb{A}_{\mu K}^k[\alpha] \\ &\Rightarrow \exists \alpha. \lambda_1 \in \mathbb{A}[\alpha] \wedge \lambda_2 \notin \mathbb{A}[\alpha]. \end{aligned} \quad \square$$

Lemma 7.23. If $N_1 \not\sim^{\mathcal{A}_{\mu K}^i} N_2$, there exist k and ϕ such that $depth(\phi) \leq k$ and $N_1 \models_{\mathbb{A}^k} \phi$ $N_2 \models_{\mathbb{A}^k} \neg\phi$.

Proof. If $N_1 \not\sim^{\mathcal{A}_{\mu K}^i} N_2$, there exists ϕ such that $N_1 \models_{\mathbb{A}^i} \phi$ $N_2 \models_{\mathbb{A}^i} \neg\phi$.

If $i > 0$ for each P and ψ such that $depth(\psi) \leq i - 1$, then $l_1 : P \blacktriangleright l_2 \in \mathbb{A}_{\mu K}^i[l_1 : \psi \blacktriangleright l_2]$, and it follows that $N_1 \models_{\mathbb{A}^i} \phi' N_2 \models_{\mathbb{A}^i} \neg\phi'$. Where ϕ' is obtained from ϕ by replacing each label predicate $l_1 : \psi \blacktriangleright l_2$ ($depth(\psi) \geq i$) with $l_1 : \mathbf{T} \blacktriangleright l_2$. We have that $depth(\phi') \leq i + 1$, and, moreover, $N_1 \models_{\mathbb{A}^{i+1}} \phi' N_2 \models_{\mathbb{A}^{i+1}} \neg\phi'$.

If $i = 0$, we let ϕ' be the formula obtained from ϕ by replacing each label predicate $l_1 : \psi \blacktriangleright l_2$ with $l_1 : \mathbf{T} \blacktriangleright l_2$. We have that $depth(\phi') \leq 1$, and, moreover, $N_1 \models_{\mathbb{A}^1} \phi' N_2 \models_{\mathbb{A}^1} \neg\phi'$. □

Lemma 7.24. $\mathcal{A}_{\mu K}$ and $\mathbb{A}_{\mu K}$ are transition respectful in $\mathcal{M}_{\mu K}$.

Proof. We now prove that for each N_1 and λ_1 there exists an α such that:

- $\lambda_1 \in \mathbb{A}[\alpha]$; and
- for each λ_2 and N_2 such that $N_2 \xrightarrow{\lambda_2} N'_2$ and $(\lambda_1, \lambda_2) \notin \mathcal{A}_{\mu K}$, we have $\lambda_2 \notin \mathbb{A}_{\mu K}^{n+1}[\alpha]$.

If λ_1 is one of τ , $l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$ (for some l_1, t or l_2), the statement follows easily by considering \surd , $l_1 : t \triangleright l_2$ or $l_1 : t \triangleleft l_2$.

If $\lambda_1 = l_1 : P \blacktriangleright l_2$, we let X be the finite set defined by

$$X = \{l_1 : Q \blacktriangleright l_2 \mid \exists N_2 : N_1 \xrightarrow{l_1 : Q \blacktriangleright l_2} N_2 \text{ and } l_2 :: P \not\sim^{\mathcal{A}_{\mu K}} l_2 :: Q\}.$$

There exists a k such that for each $l_1 : Q \blacktriangleright l_2 \in X$, $l_2 :: P \not\sim^{\mathcal{A}_{\mu K}} l_2 :: Q$. For each $\lambda = l_1 : Q \blacktriangleright l_2 \in X$ there exists ϕ_λ such that $l_2 :: P \models_{\mathbb{A}^k} \phi_\lambda$, while $l_2 :: Q \models_{\mathbb{A}^k} \neg\phi_\lambda$.

Let $\phi_X = \bigcap_{\lambda \in X} \phi_\lambda$. We have that $l_1 : P \blacktriangleright l_2 \in \mathbf{A}_{\mu K}^k \llbracket l_1 : \phi_X \blacktriangleright l_2 \rrbracket$, while for each $\lambda \in X$, $\lambda \notin \mathbf{A}_{\mu K}^k \llbracket l_1 : \phi_X \blacktriangleright l_2 \rrbracket$. Thus, $l_1 : P \blacktriangleright l_2 \in \mathbf{A} \llbracket l_1 : \phi_X \blacktriangleright l_2 \rrbracket$, while for each $\lambda \in X$, $\lambda \notin \mathbf{A} \llbracket l_1 : \phi_X \blacktriangleright l_2 \rrbracket$. \square

We can now state the final theorem of this section.

Theorem 7.25. $\sim^{\mathcal{A}_{\mu K}}$ characterises the equivalence induced by formulae satisfaction.

Proof. The statement follows directly from Lemmas 7.22 and 7.24 using Theorems 7.5 and 7.8. \square

8. Conclusions and future work

In this paper we have proposed a variant of LTSs, which we have called *Multiple-Labelled Transition Systems* (MLTSs), as a candidate general operational model for distributed calculi with names and mobility. To show the usefulness of our proposal, we used MLTSs to describe the operational semantics of two formalisms, namely the asynchronous π -calculus ($A\pi$) (Boudol 1992; Honda and Tokoro 1991) and μ KLAIM (Bettini *et al.* 2003), that have the opposite objectives of expressivity and usability.

For modelling the properties of MLTSs, we introduced a temporal logic (MoMo) that consists of a small set of operators to be used to describe specific properties/behaviours of mobile and distributed systems. Together with the usual logical connectives and the operators for minimal and maximal fixed points, the logic is equipped with operators for: describing dynamic behaviours (*temporal properties*); modelling resource management (*state properties*); taking into account name handling (*nominal properties*); and controlling mobile processes (*mobility properties*).

We have also studied the relationships between the equivalences induced on $A\pi$ by MoMo and by bisimulation.

The main limitation of the proposed approach is that in order to establish system properties we need to provide detailed descriptions of the whole systems under consideration. Obviously, this is a very strong assumption for wide area networks, since it is very often the case that only some components of the system are known; and one has only a limited knowledge of the overall context in which the component is operating. Nevertheless, one would like to guarantee that components behave well whenever the context guarantees specific resources or interactions.

For this reason, we plan to set up a framework for specifying contexts for distributed calculi (whose semantics is specified in term of MLTSs). By means of contexts, we will be able to provide abstract specifications of a given system and avoid the need to describe all of its components in full. Indeed, some of these components could be known or implemented only at a later stage. Then, the implemented component can be removed from the context and added to the implemented part, thus performing a *concretion* operation. We aim to set up a framework that would guarantee the preservation of formulae satisfaction at each stage of refinement if the introduced implementation *agrees* with the original specification.

We also plan to use MLTSs to describe the behaviour of other calculi with explicit notions of distribution and mobility, and, in particular, for modelling emerging calculi for Service Oriented Architecture, such as SCC (Boreale *et al.* 2006). This will enable us to specify and verify the properties of distributed service architecture using MoMo.

Acknowledgements

We would like to thank Diego Latella and Mieke Massink for suggestions, and members of the SENSORIA projects for interesting discussions.

References

- Amadio, R.M., Castellani, I. and Sangiorgi, D. (1998) On bisimulations for the asynchronous π -calculus. *Theoretical Computer Science* **195** (2) 291–324.
- Bettini, L., Bono, V., De Nicola, R., Ferrari, G., Gorla, D., Loreti, M., Moggi, E., Pugliese, R., Tuosto, E. and Venneri, B. (2003) The klaim project: Theory and practice. In: *Global Computing. Springer-Verlag Lecture Notes in Computer Science* **2874** 88–150.
- Boreale, M., Bruni, R., Caires, L., De Nicola, R., Lanese, I., Loreti, M., Martins, F., Montanari, U., Ravara, A., Sangiorgi, D., Vasconcelos, V. and Zavattaro, G. (2006) SCC: a service centered calculus. In: *Web Services and Formal Methods, Third International Workshop, WS-FM 2006. Springer-Verlag Lecture Notes in Computer Science* **4184** 38–57.
- Boudol, G. (1992) Asynchrony and the π -calculus (note). Rapport de Recherche 1702, INRIA Sophia-Antipolis.
- Caires, L. (2004) Behavioral and spatial observations in a logic for the pi-calculus. In: Walukiewicz, I. (ed.) *FoSSaCS. Springer-Verlag Lecture Notes in Computer Science* **2987** 72–89.
- Cardelli, L. and Gordon, A.D. (2000) Mobile ambients. *Theoretical Computer Science* **240** (1) 177–213.
- Cardelli, L. and Gordon, A.D. (2006) Ambient logic. *Mathematical Structures in Computer Science* (to appear).
- Carriero, N. and Gelernter, D. (1989) Linda in Context. *Communications of the ACM* **32** (10) 444–458.
- Castagna, G., Vitek, J. and Nardelli, F.Z. (2005) The Seal Calculus. *Information and Computation* **201** (1) 1–54.
- Dam, M. (1996) Model checking mobile processes. *Journal of Information and Computation* **129** (1) 35–51.
- De Nicola, R., Ferrari, G.L. and Pugliese, R. (1998) KLAIM: A kernel language for agents interaction and mobility. *IEEE Transactions on Software Engineering* **24** (5) 315–330.
- De Nicola, R. and Loreti, M. (2005) MoMo: A modal logic for reasoning about mobility. In: de Boer, F.S., Bonsangue, M.M., Graf, S. and de Roever, W.P. (eds.) *Proc. of FMCO 2004, Revised Lectures. Springer-Verlag Lecture Notes in Computer Science* **3657** 95–119.
- Ferrari, G., Montanero, C., Semini, L. and Semprini, S. (2002) Mark, a reasoning kit for mobility. *Automated Software Engineering* **9** 137–150.
- Fournet, C., Gonthier, G., Levy, J.J., Maranget, L. and Remy, D. (1996) A Calculus of Mobile Agents. In: Montanari, U. and Sassone, V. (eds.) *Proc. of 7th Int. Conf. on Concurrency Theory (CONCUR'96). Springer-Verlag Lecture Notes in Computer Science* **1119** 406–421.

- Gabbay, M. and Pitts, A. M. (1999) A new approach to abstract syntax involving binders. *Proceedings of the 14th Annual IEEE Symposium on Logic in Computer Science*, IEEE Computer Society 214–224.
- Gelernter, D. (1985) Generative Communication in Linda. *ACM Transactions on Programming Languages and Systems* **7** (1) 80–112.
- Gelernter, D. (1989) Multiple Tuple Spaces in Linda. In: Odijk, E., Rem, M. and Syre, J.-C. (eds.) Proceedings, PARLE '89. *Springer-Verlag Lecture Notes in Computer Science* **366** 20–27.
- Hennessy, M. and Milner, R. (1985) Algebraic laws for nondeterminism and concurrency. *Journal of the ACM* **32** (1) 137–161.
- Hennessy, M. and Riely, J. (2002) Resource access control in systems of mobile agents. *Information and Computation* **173** (1) 82–120.
- Honda, K. and Tokoro, M. (1991) An object calculus for asynchronous communication. In: America, P. (ed.) European Conference on Object-Oriented Programming (ECOOP'91). *Springer-Verlag Lecture Notes in Computer Science* **512** 133–147.
- McCann, P. and Roman, G.-C. (1998) Compositional programming abstraction for mobile computing. *IEEE Transactions on Software Engineering* **24** (2) 97–110.
- Milner, R., Parrow, J. and Walker, D. (1993) Modal logics for mobile processes. *Theoretical Computer Science* **114** 149–171.
- Reynolds, J. C. (2002) Separation logic: A logic for shared mutable data structures. In: *Proceedings of 17th IEEE Symposium on Logic in Computer Science*, IEEE Computer Society 55–74.
- Wojciechowski, P. T. and Sewell, P. (1999) Nomadic Pict: Language and infrastructure design for mobile agents. In: *ASA/MA*, IEEE Computer Society 2–12.