

Model Checking for Nominal Calculi^{*}

Gianluigi Ferrari, Ugo Montanari, and Emilio Tuosto

Dipartimento di Informatica, Largo Bruno Pontecorvo 3, 56127 Pisa – Italy

Abstract. Nominal calculi have been shown very effective to formally model a variety of computational phenomena. The models of nominal calculi have often infinite states, thus making model checking a difficult task. In this note we survey some of the approaches for model checking nominal calculi. Then, we focus on *History-Dependent automata*, a syntax-free automaton-based model of mobility. History-Dependent automata have provided the formal basis to design and implement some existing verification toolkits. We then introduce a novel syntax-free setting to model the symbolic semantics of a nominal calculus. Our approach relies on the notions of reactive systems and observed borrowed contexts introduced by Leifer and Milner, and further developed by Sassone, Lack and Sobocinski. We argue that the symbolic semantics model based on borrowed contexts can be conveniently applied to web service discovery and binding.

1 Summary

Model checking has been shown very effective for proving properties of system behaviour whenever a finite model of it can be constructed. The approach is convenient since it does not require formal proofs and since the same automaton-like model can accommodate system specification languages with substantially different syntax and semantics. Among the properties which can be checked, behavioural equivalence is especially important for matching specifications and implementations, for proving the system resistant to certain attacks and for replacing the system with a simpler one with the same properties.

Names have been used in process calculi for representing a variety of different informations concerning addresses, mobility links, continuations, localities, causal dependencies, security keys and session identifiers. When an unbound number of new names can be generated during execution, the models tend to be infinite even in the simplest cases, unless explicit mechanisms are introduced to allocate and garbage collect names, allowing the same states to be reused with different name meanings.

We review some existing syntax-free models for name-passing calculi and focus on *History-Dependent automata* (HD-automata), introduced by Montanari and Pistore in 1995 [62]. HD-automata [62, 63, 71] have been shown a suitable automata-based model for representing Petri nets, CCS with causality and localities and some versions of π -calculus [59, 75].

^{*} Work supported by European Union project PROFUNDIS, Contract No. IST-2001-33100.

Different versions of HD-automata have been defined. The simplest version can be easily translated to ordinary automata, but possibly with a larger number of states. In a second version, the states are equipped with name symmetries which further reduce the size of the automata. Furthermore, a theory based on coalgebras in a category of “named sets” can be developed for this kind of HD-automata, which extends the applicability of the approach to other nominal calculi and guarantees the existence of the minimal automaton within the same bisimilarity class [64, 34].

HD-automata also constitute the formal basis upon which several verification toolkits have been defined and implemented. The front end towards the π -calculus and the translation algorithm for the simplest version of HD-automata have been implemented in the HAL tool [31, 32], which relies on the JACK verification environment [7] for handling the resulting ordinary automata. The minimisation algorithm, naturally suggested by the coalgebraic framework, has been implemented in the Mihda toolkit [35, 36] within the European project PROFUNDIS. Other versions of HD-automata can be equipped with algebraic operations, and are based on a algebraic-coalgebraic theory [61].

Here we propose a further instance handling the symbolic versions of nominal calculi, where inputs are represented as variables which are instantiated only when needed. As it is the case for logic programming unification, one would like the variables to be instantiated only the least possible, still guaranteeing that all behaviours are eventually explored. The approach we follow relies on the notion of reactive system and of observable borrowed contexts introduced by Leifer and Milner [53, 52] and further developed by Sassone, Lack and Sobocinski [76, 78, 50] using G-categories and adhesive categories. The reduction semantics of reactive systems is extended in order to introduce as borrowed contexts both the variable instantiations needed in the transitions and the ordinary π -calculus actions. It is argued that the symbolic semantics model based on borrowed contexts can be conveniently applied to web service discovery and binding.

In this paper we review the main results on HD-automata setting them in the mainstream research on nominal calculi. The final part of the paper introduces a novel symbolic semantics of π -calculus based on reactive systems and observed borrowed contexts. In our approach, unification is the basic interaction mechanism. We consider this as being the first step toward the definition of a formal framework (models, proof techniques and verification toolkits) for the so-called *service oriented computing* paradigm.

2 Verification via Semantics Equivalence

In the last thirty years the application of formal methods to software engineering has generated techniques and tools to deal with the various facets of the software development process (see e.g. [19] and the references therein). One of the main advantages of exploiting formal techniques consists of the possibility of constructing *abstractions* that approximate behaviours of the system under development. Often, these abstractions are amenable to automatic verification of properties thus providing a support to the certification of software quality.

Among the different proposals, *verification via semantics equivalence* provides a well established framework to deal with the checking of behavioural properties. In

this approach, checking behavioural properties is reduced to the problem of contrasting two system abstractions in order to determine whether their behaviours coincide with respect to a suitable notion of semantics equivalence. For instance, it is possible to verify whether an abstraction of the implementation is consistent with its abstract specification. Another example is provided by the *information leak* detection; in [39] the analysis of information flow is done by verifying that the abstraction of the system P is equivalent to another abstraction obtained by suitably restricting the behaviour of P . A similar idea has been exploited in [1] for the analysis of cryptographic protocols.

Bisimilarity [69] has been proved to be an effective basis for verification based on semantics-equivalence of system abstractions described in some process calculus, i.e. Milner's Calculus of Communicating Systems (CCS) [58]. Bisimilarity is a *co-inductive* relation defined over a special class of automata called *labelled transition systems*. A generic labelled transition system (LTS) describes the evolution of a system by its interactions with the external environment. The co-inductive nature of bisimulation provides an effective proof method to establish semantics equivalence: it is sufficient to exhibit a bisimulation relating the two abstractions. Bisimulation-based proof methods have been exploited to establish properties of a variety of systems such as communication protocols, hardware designs and embedded controllers. Moreover, they have been incorporated in several toolkits for the verification of properties. Indeed, finite state verification environments have enjoyed substantial and growing use over the last years. Here, we mention the Concurrency WorkBench [21], the Meije-FC2 tools [8] and the JACK toolkit [7] to cite a few. Several systems of considerable complexity have been formalised and proved correct by exploiting these semantics-based verification environments.

The advent of mobile computing and wireless communication together with the development of applications running over the Internet (*Global Computing Systems*) have introduced software engineering scenarios that are much more dynamic than those handled with the techniques discussed above. Indeed, finite state verification of global computing systems is much more difficult: in this case, even simple systems can generate infinite state spaces. An illustrative example is provided by the π -calculus [59, 75]. The π -calculus primitives are simple but expressive: channel names can be created, communicated (thus giving the possibility of dynamically reconfiguring process acquaintances) and they are subjected to sophisticated scoping rules. The π -calculus is the archetype of name passing or nominal process calculi. Nominal calculi emphasise the principle that name mechanisms (e.g. local name generation, name exchanges, etc.) provide a suitable abstraction to formally explain a wide range of phenomena of global computing systems (see e.g. [80, 41]). Moreover, nominal calculi provide a basic programming model that has been incorporated in suitable libraries or novel programming languages [22, 4]. Finally, the usefulness of names has been also emphasised in practice. For instance, Needham [66] pointed out the role of names for the security of distributed systems. The World Wide Web provides an excellent (perhaps the most important) example of the power of names and name binding/resolution.

Nominal calculi have greater expressive power than ordinary process calculi, but the possibility of dynamically generating new names leads also to a much more complicated

theory. In particular, bisimilarity is not always a congruence even for the strong bisimilarity. Moreover, the ordinary, underlying LTSs are infinite-state and infinite branching, thus making verification via semantics equivalence a difficult task.

Bisimulation-based proof techniques for nominal calculi can be roughly divided into two main families. The first consists of the *syntax-based* approaches while the second refers to the *syntax-free* approaches. The former line of development represents the states of the LTS with their syntactic denotation, while in the latter the states are just items characterised by their properties and connections. We recall a few of the approaches of both families without the ambition of being exhaustive.

Among the syntax-based, the most efficient approaches for finite-state verification rely on symbolic semantics. Symbolic semantics [42, 6, 54], generalise standard operational semantics by keeping track of equalities among names: transitions are derived in the context of such constraints. The main advantage of the symbolic semantics is that it yields a smaller transition system. The idea of symbolic semantics has been exploited to provide a convenient characterisation of *open bisimilarity* [74] and in the design of the corresponding bisimulation checker, the *Mobility WorkBench* (MWB) [83]. The MWB adapts to the case of the π -calculus the *on-the-fly* approach of [30], where the bisimulation relation is constructed during the state space generation. The MWB checks for open bisimilarity in the case of (finite-control) π -calculus processes and has also been reworked to deal with the Fusion calculus [70]. To gain efficiency, the MWB has been extended in [44] with modules implementing certain bisimulation-preserving program transformations, the *up-to-techniques* (introduced in [73]). Symbolic semantics has been also exploited in the design of the MCC model checker for the π -calculus [84]. The key idea of the approach is to provide an encoding of π -calculus symbolic semantics as a logic programming system. It is important to emphasise that all the constructions of the symbolic semantics rely on an *external* metalanguage and on a theory to describe and reason about name equalities.

A different approach is the definition of semantic-based techniques where names have a central role and are explicitly dealt with. Basically, in these frameworks it is possible to allocate and garbage collect names, allowing the same names to be reused with different meanings. This alternative line of research explores models of name-passing calculi, regardless of their syntactic details and aims at providing uniform theories that can be used to handle a variety of calculi and semantics. A well studied approach is based on the so-called permutation model, whose ingredients are a set of names and an action of its group of permutations (renaming substitutions) on an abstract set [37, 40, 47, 64]. In this setting, transition systems for nominal calculi are constructed via suitable functors over the underlying category of names and permutations: the internal theory of names.

It is important to notice that these approaches are *syntax-free* and provide the abstract framework to capture the notions of name abstraction and fresh name that are needed to describe and reason about nominal calculi. The HD-automata [34, 64, 71] and indexed LTSs [17] are examples of syntax-free models of name process calculi developed following the permutation approach.

3 Model Checking

Probably, the most successful formal technique applied in practice in the verification of systems is *model checking* (we refer to [18] for a detailed introduction to this field). Roughly speaking, model checking is used to determine whether a system abstraction (expressed as an automata or a term of a process calculus) satisfies a property (expressed as a modal or temporal logic formula). In order to model check a system with respect to a given formula it is necessary to prove that the system is a model of the formula. Tools supporting model checking techniques have matured to be used in practice (e.g. the SPIN model checker [45, 46] and SMV [57]). Recently, these techniques have been adopted to verify properties of programs written in high level programming languages like C++ and Java (e.g. JavaPathFinder [10], BANDERA [23], SLAM [3] and BLAST [43]).

Model checking presents several advantages. It is completely automatic, provided that finiteness of the system (the model) is guaranteed. Usually, it provides counterexamples when a system does not satisfy the property. This gives information on the design choices that have lead to the implementation errors. Finally, it is possible to obtain very high efficiency by exploiting refined data structures (e.g. BDDs), or symbolic techniques.

While modal and temporal logics have been proved suitable to express many properties of interest of concurrent systems, similar logics for global computing systems are still lacking. Only recently a new class of modal logics, *spatial logics* [15, 16], has been introduced to address the characterising issues of global computing. In our opinion, this explains why traditionally model checking has been exploited on foundational models for global computing only for limited fields and has not been fully applied to the general setting.

Without the ambition of being exhaustive, we now review some of the approaches to model check properties of nominal calculi. The MWB provides a model checking functionality. This is based on the implementation of the tableau-based proof system [25, 26] for the π - μ calculus, an extension of the propositional μ -calculus in which it is possible to express name parameterisation and quantifications over names. The MCC system also provides a model checking facility for the π - μ calculus.

The *HD-automata Laboratory* (HAL) [32] supports verification by model checking of properties expressed as formulae of a suitable modal logic, a high level logic with modalities indexed by π -calculus actions. This logic, although expressive enough to describe interesting safety and liveness properties of π -calculus specifications, is less expressive than the π - μ calculus. The construction of the HAL model checker takes direct advantage of the finite representation of π -calculus specifications presented in [62]. In particular, a HAL module translates these logical formulae into classical modal logic formulae and the translation is driven by the finite state representation of the system (the π -calculus process) to be verified.

The most relevant examples of application of model checking techniques and nominal calculi are those of the verification of security protocols [56, 20]. Several prototypical tools based on nominal calculi have been in fact designed and implemented [60, 55, 27, 38]. Indeed, nominal calculi provide a solid formal context for expressing many facets of cryptographic protocols in natural way. For instance, many authentication protocols rely on *nonce-challenges* where a fresh sequence of bit must be generated; the

correctness of these protocols relies on the uniqueness of the nonces used in a given session. This can be easily modelled in nominal calculi, e.g. the π -calculus, where freshly generated names can be expressed and dealt with. An advantage of using model checking is that, when the protocol does not satisfy the security property, then the counterexample is the attack that an intruder could perform.

The main drawback of these approaches is that they require a finite state space while, in general, the generation of fresh names easily leads to infinite state spaces, if no countermeasure for garbage-collecting and reusing names is adopted. In practice, this problem has been faced by imposing strong conditions that limit the generality of the analysis. In particular, *finitary* systems, namely systems with infinite behaviour which can be finitely represented, are not considered. For instance, the analysis are performed on instances of protocols where only a limited number of participants is *a priori* fixed and in general recursion or iteration is forbidden. Hence, model checking security properties for nominal calculi can only deal with protocol sessions where a finite number of participants run in parallel and all the participants are non-recursive processes. Recently, symbolic ad-hoc model checkers have been proposed to overcome these issues e.g., [5, 82, 9, 2]. Despite the technical differences, all these approaches check a given property by generating a “symbolic” state space, where states collect constraints over the names involved in the execution. If there is a reachable state that violates the property, but whose constraints hold, then an attack is found. The symbolic techniques exploited in these approaches enforce efficiency both in the size of the generated state space and in the visit of it, but they still require finite state space.

4 History-Dependent Automata

History Dependent automata (HD-automata in brief) are one of the proposal based on the syntax-free approach. HD-automata are an operational model for history dependent formalisms, namely those formalisms accounting for systems whose behaviour at a given time might be influenced by some “historical” information which is too expensive to be included explicitly in the states. HD-automata allow for a compact representation of agent behaviour by collapsing states differing only for the renaming of local names and encompass the main characteristics of name-passing calculi, namely creation/deallocation of names. Basically, HD-automata associate a “history” to the names of the states appearing in the computation, in the sense that it is possible to reconstruct the associations which have led to the state containing the name. Clearly, if a state is reached in two different computations, different histories could be assigned to its names. Process calculi exhibiting causality, localities and mobility, and Petri nets, can be translated (preserving bisimilarity) to HD-automata [71].

Different versions of HD-automata have been defined [71, 63, 64, 34]. When handling causality, locality and the link mobility exhibited by the synchronous π -calculus without matching, the simplest version can be easily translated to ordinary automata. However, in general, a larger number of states is necessary for representing HD-automata with ordinary automata. The front-end towards the π -calculus and the translation algorithm have been implemented in the HAL toolkit, which relies on the JACK verification environment for handling the resulting ordinary automata.

In a second version, states of HD-automata are equipped with name symmetries which further reduce the size of the automata [64] and which guarantee the existence of the minimal realization. The minimal automata are computed using a partition refinement algorithm [34]. They have a very important practical fall-out: for instance, the problem of deciding bisimilarity is reduced to the problem of computing the minimal transition system [67, 29, 49]. Moreover, the minimal automaton is indistinguishable from the original one with respect to many behavioural properties (e.g., bisimilarity) and properties expressed in most modal or temporal logics. The minimisation algorithm, naturally suggested by the coalgebraic framework, has been implemented in the Mihda toolkit [36] within the European project PROFUNDIS. Other versions of HD-automata can be equipped with algebraic operations [61], thus relying on an algebraic-coalgebraic, namely bialgebraic, theory.

Similarly to ordinary automata, HD-automata consist of states and labelled transitions, their peculiarity being that states and transitions are equipped with names which are no longer dealt with as syntactic components of labels, but become an explicit part of the operational model. Noteworthy, names in states of HD-automata have *local meaning* which requires a mechanism for describing how names correspond each other along transitions.

Graphically, we can represent such correspondences using “wires” that connect names of label, source and target states of transitions as in Figure 1, where a tran-

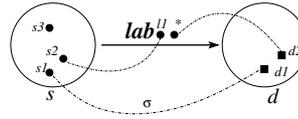


Fig. 1. A HD-automaton transition

sition from source state s to destination state d is depicted. State s has three names, $s1$, $s2$ and $s3$ while d has two names $d1$ and $d2$ which correspond to name $s1$ of s and to the new name $*$, respectively. The transition is labelled by lab and exposes two names: name $l1$ and $*$ the former corresponding to name $s2$ of s and the latter to a fresh name denoted as $*$. Notice that name $s3$ is “deallocated” along such transition.

4.1 Minimising HD-Automata: An Informal Presentation

We report the formal definitions for *named sets* and *named functions* for representing finite HD-automata. These are the basic concepts upon which the partition refinement algorithm for HD-automata has been defined. For the sake of conciseness we give here only an incomplete definition. The interested reader is referred to [36] for a full presentation.

Definition 1 (Named Sets). Let \mathcal{N} be a denumerable set of names. A named set is a pair $\langle Q, \mathbf{g} \rangle$ where Q is a totally-ordered set and $\mathbf{g} : Q \rightarrow \bigcup_{N \in \wp_{fin}(\mathcal{N})} \text{sym}(N)$ assigns a (finite) group of permutations over a finite set of names to elements in Q . For $q \in Q$, $|q|$ denotes the carrier of q defined as $\text{dom}(\rho)$, where $\rho \in \mathbf{g}(q)$.

Definition 2 (Named Functions). Let \mathcal{N}^* be $\mathcal{N} \cup \{\star\}$ where \star is an element not in \mathcal{N} . Given two named sets $\langle Q, \mathbf{g} \rangle$ and $\langle Q', \mathbf{g}' \rangle$, a named function $H : \langle Q, \mathbf{g} \rangle \rightarrow \langle Q', \mathbf{g}' \rangle$ consists of a pair of functions $\langle h, \Sigma \rangle$ where $h : Q \rightarrow Q'$ and $\Sigma : Q \rightarrow \wp_{\text{fin}}(\mathcal{N} \rightarrow \mathcal{N}^*)$ such that for all $q \in Q$ and $\sigma \in \Sigma(q)$

- σ is injective, $\sigma(|h(q)|) \subseteq |q| \cup \{\star\}$ and $\sigma|_{\mathcal{N} \setminus |h(q)|}$ is the identity;
- $\sigma; \mathbf{g}(q) \subseteq \Sigma(q)$;
- $\mathbf{g}(h(q)); \sigma = \Sigma(q)$.

Named sets and functions form a category, *NS*, since named functions can be composed and identity named functions can be easily defined (see [36] for details). Given a set of labels L , if $\wp_{\text{fin}}(\cdot)$ is the finite power set functor on category *Set*, we define the functor \wp_L on named sets as $\wp_L(\langle Q, \mathbf{g} \rangle) = \langle \wp_{\text{fin}}(L \times Q), \mathbf{g}' \rangle$ where $\mathbf{g}'(B)$ contains all those permutations ρ such that $B\rho = B$ ($B\rho$ is element-wise application of ρ to B).

Definition 3. A HD-automaton over L is a coalgebra for the functor \wp_L .

The most important operation for minimising HD-automata is the *normalisation* which removes *redundant* transitions. In nominal calculi, redundancy is strictly connected to the concept of *active names*. A name n is *inactive* for an agent P if it is not used in the future behaviour of P .

In π -calculus, if P is bisimilar to $(\nu n)P$ we say that n is inactive in P (otherwise n is *active* in P) and a transition $P \xrightarrow{xn} Q$ is *redundant* (in the early semantics of π -calculus) when n is inactive in P . Deciding whether a name is active is as difficult as deciding bisimilarity. The importance of redundancy emerges when we try to establish the equivalence of states that have different numbers of free names. For instance, consider $P \stackrel{\text{def}}{=} x(u).(\nu v)(\bar{v}z + \bar{u}y)$ and $Q \stackrel{\text{def}}{=} x(u).\bar{u}y$, which differ only for a deadlocked alternative. They are bisimilar only if, for any name substituted for u , their continuations remain bisimilar. However, the input transition $P \xrightarrow{xz}$ cannot be matched by Q when considering only the *necessary* input transitions of agents, namely those where the acquired name is either a fresh name or one of the free names of the agent (as required for a finite representation of the transition system). Thus, unless the above transition of P is recognised as redundant and removed, the automata for P and Q would not be bisimilar. Redundant transitions occur when LTSs of π -calculus processes are compiled to HD-automata and are removed during the minimisation algorithm, since it is not possible to leave them out at compiling time¹.

The minimisation algorithm relies on functor T consisting of the composition of the normalisation functor and \wp_L . Consider a T -coalgebra $\langle D, K : D \rightarrow T(D) \rangle$, the minimisation algorithm is defined by the two equations below.

¹ In general, deciding whether a free input transition is redundant or not is equivalent to decide whether a name is active or not; therefore, it is as difficult as deciding bisimilarity.

$$H_{(0)} \stackrel{\text{def}}{=} \langle q \mapsto \perp, q \mapsto \emptyset \rangle, \quad \text{where } \text{dom}(H_{(0)}) = D \quad (1)$$

$$H_{(i+1)} \stackrel{\text{def}}{=} K; T(H_i). \quad (2)$$

In words, all the states of automaton K are initially considered equivalent, indeed, the kernel of H_0 gives rise to a single equivalence class containing the whole $\text{dom}(K)$. At the generic $(i+1)$ -th iteration, the image through T of the i -th iteration is composed with K as prescribed in (2). The algorithm stops when the fixpoint \bar{H} of (2) is reached. Then \bar{H} is the unique final coalgebra morphism and states mapped together by it are bisimilar.

Theorem 1 (Convergence [36]). *The iterative algorithm described by (1) and (2) is convergent on finite state automata.*

4.2 The PROFUNDIS Web

In the last years distributed applications over the World-Wide Web, have attained wide interest. Recently, the Web is exploited as a *service distributor* and applications are no longer monolithic but rather made of components (i.e., services). Applications over the Web are developed by combining and integrating Web services. The Web service framework has emerged as the standard and natural architecture to realize the so called *Service Oriented Computing* (SOC) [24, 68]. In [33] a Web-service infrastructure was developed integrating verification toolkits for checking properties of mobile systems and related higher-level toolkits for verifying security protocols. The development of the verification infrastructure has been performed inside the PROFUNDIS project (see URL <http://www.it.uu.se/profundis>) within the Global Computing Initiative of the European Union. For this reason we called it the *PROFUNDIS WEB*, PWeb for short. The current prototype implementation of the PWeb infrastructure can be exercised on-line at the URL <http://jordie.di.unipi.it:8080/pweb>.

Beyond the current prototype implementation, we envisage the important role that will be played by PWeb service coordination. Indeed, service coordination provides several benefits:

- *Model-based verification.* The coordination rules impose constraints on the execution flow of the verification session thus enabling a *model-based* verification methodology where several descriptions are manipulated together.
- *Modularity.* The verification of the properties of a large software system can be reduced to the verification of properties over subsystems of manageable complexity: the coordination rules reflect the semantic modularity of system specifications.
- *Flexibility.* The choice of the verification toolkits involved in the verification session may depend on the specific verification requirements.

The PWeb implementation has been conceived to support reasoning about the behaviour of systems specified in some dialect of the π -calculus. It supports the dynamic integration of different verification techniques (e.g. standard bisimulation checking and symbolic techniques for cryptographic protocols). The PWeb integrates several independently-developed toolkits, e.g., Mihda [35, 36] and several tools for verifying

cryptographic protocols, like TRUST [82] and STA [5]. The PWeb has been designed by targeting also the goal of extending available verification environments (Mobility Workbench [83], HAL [31, 32]) with new facilities provided as Web services.

The core of the PWeb is a *directory service*. A PWeb directory service is a component that maps the description of the Web services into the corresponding network addresses and has two main facilities: the `publish` facility, invoked to make a toolkit available as Web service, and the `query` facility, used to discover available services. For instance, Mihda publishes the `reduce` service which accepts a (XML description of) HD-automaton describing the behaviour of a π -calculus agent. Once invoked, `reduce` performs the minimisation of the HD-automaton.

The service discovery mechanisms are exploited by the `trader` engine which manipulates pools of services distributed over several PWeb directory services. It can be used to obtain a Web service of a certain type and to bind it inside the application. The `trader` engine gives to the PWeb directory service the ability of finding and binding web services at run-time without “hard-coding” the name of the web service inside the application code. The following code describes the use of a simple trader for the PWeb directory.

```
import Trader
offers = Trader.query( "reducer" )
mihda = offers[ 0 ]
```

The code asks the trader for a `reduce` service and selects the first of them. The trader engine allows one to hide network details in the service coordination code. A further benefit is given by the possibility of replicating the services and maintaining a standard access modality to the Web services under coordination.

The fundamental technique enabling the dynamic integration of services is the separation between the service facilities (what the service provides) and the mechanisms that coordinate the way services interact (service coordination). An example of service coordination for checking whether a process `A` is a model for a formula `F` is as follows

```
hd = mihda.compile( A )
reduced_hd = mihda.reduce( hd )
reduced_hd_fc2 = mihda.Tofc2( reduced_hd )
aut = hal.unfold( reduced_hd_fc2 )
if hal.check( aut, F ):
    print 'ok'
else:
    print 'ko'
```

Variables `mihda` and `hal` have been linked by the trader engine to the required services (acquired as illustrated before). Now, the `compile` service of `mihda` is invoked yielding an HD-automaton (stored in `hd`). Next, `hd` is minimised by invoking the service `reduce` of `Mihda`; and afterward it is transformed into the FC2 format by a HAL service. Finally, the HAL service `unfold` generates an ordinary automaton from the FC2 representation of the automaton and prints a message which depends on whether the system satisfies the formula `F` or not. This is obtained by invoking the HAL model checking facility `check`.

5 A Borrowed Context Semantics for the Open π -Calculus

The version of the π -calculus implemented in the *Mihda* toolkit does not rely on a symbolic semantics. This fact makes unnecessarily large the number of states, due to the existence of different input transitions for different instantiations of the input variable. While a symbolic semantics for a syntax-based version of HD-automata for the open π -calculus has been defined in [72], it might be convenient to define a symbolic semantics for the ordinary syntax-free HD-automata. More generally, in Service Oriented Computing (SOC) [24, 68] one would like to have more sophisticated mechanisms than service call and parameter passing for modelling the phase of service discovery and binding. The SOC paradigm is the emerging technology to design and develop global computing systems: several research activities have addressed the theoretical foundations of the SOC paradigm by exploiting formal frameworks based on process calculi [12, 51, 14, 11] (see also [81] for an informal presentation on the usefulness of nominal calculi to design workflow business processes).

When looking for a generalisation of parameter passing, logic programming unification comes to mind, or rather constraint programming, when service level agreements involve nonfunctional issues. When the binding occurs, not only the callee is instantiated, but also the caller. The instantiation that must be applied to the caller is formally analogous to a missing context that must be borrowed by a process in order to undergo a reduction. In this line of thought, some recent works about systematic methods for deriving LTSs from reduction rules look relevant. In particular, the approach we follow relies on the notion of reactive system, introduced by Leifer and Milner [53, 52], used by Jensen and Milner in [48] for deriving a LTS for bigraphs and further developed by Sassone, Lack and Sobocinski [76, 78, 50] using G-categories and adhesive categories.

In this section we will consider a simplified version of open π -calculus and we will develop a semantics for it using the notion of reactive systems. While the corresponding bisimilarity semantics turns out to be finer, we think that this exercise shows the feasibility of employing context borrowing for modelling symbolic semantics. The generality of the reactive system approach gives some hope that interesting abstractions of the SOC paradigm could also be modelled that way. Note however that the transition system which can be derived from reactive rules in our development is not really suitable for a HD-automata implementation, since new names are never forgotten, thus making the transition systems infinite in all but the most trivial cases. We comment in Section 6 about possible solutions of this problem.

5.1 Open π -Calculus

One of main peculiarities of the π -calculus is the richness of its observational semantics. Initially, it came equipped with the *early* and the *late* observational semantics [59] which differ each other in the way they deal with name instantiation. Symbolic semantics [42] generalises standard operational semantics by keeping track of equalities among names: transitions are derived in the context of such constraints. The main advantage of the symbolic semantics is that it yields a smaller transition system. The idea of symbolic semantics has been exploited to provide a finitary characterisation of *open*

Table 1. Semantics of π_-

$\text{(PRE)} \alpha.p \xrightarrow{\alpha} p$	$\text{(SUM)} \frac{p \xrightarrow{\mu} p'}{p + q \xrightarrow{\mu} p'}$
$\text{(PAR)} \frac{p \xrightarrow{\mu} p'}{p \mid q \xrightarrow{\mu} p' \mid q} \text{ if } \text{bn}(\mu) \cap \text{fn}(q) = \emptyset$	$\text{(COM)} \frac{p \xrightarrow{\bar{a}b} p' \quad q \xrightarrow{a'(c)} q'}{p \mid q \xrightarrow{a=a'} p' \mid q' \{b/c\}}$
$\text{(REP)} \frac{p \mid p' \xrightarrow{\mu} q}{p' \xrightarrow{\mu} q}$	

bisimilarity [74] which, differently from the early and the late semantics, is a congruence with respect to the contexts of the π -calculus.

We consider a subset of the π -calculus without neither matching nor restriction operators. Given a numerable infinite and totally ordered set of *names* $\mathcal{N} = \{a_1, a_2, \dots\}$, the set \mathbf{P} of π_- *processes* is defined by the grammar

$$p, q ::= \mathbf{0} \mid \mu.p \mid p \mid q \mid p + q \mid p! \qquad \alpha ::= \bar{a}b \mid a(b).$$

As usual, name a is free in $\bar{a}b$ and $a(b)$, while b is free just in the former case and bound in the latter. Moreover, a is called the *subject* and b the *object* of the action. Considering $a(b).p$, the occurrences of b in p are bound, *free names* are defined as usual and $\text{fn}(p)$ indicates the set of free names of process p . Differently than in the full π -calculus, only the input prefix binds names. Processes are considered equivalent up-to α -renaming of bound names.

The operational rules for the semantics of π_- are those reported in Table 1 together with the symmetric rules for (PAR) and (SUM). The rules specify an LTS whose labels (denoted as μ) are either actions or *fusions*. The only non-standard rule is (COM) which states that an output $\bar{a}b$ and an input $a'(c)$ can synchronise provided that a and a' are fused. Notice that, if a and a' are the same, $a = a'$ is the identity fusion, denoted as ε , which corresponds to the usual silent action τ .

The transition system of π_- resulting from specification rules in Table 1 is the same as the one obtained by applying the LTS rules of [74] to π_- . The only differences between the two LTSs are in the syntax of the labels and in the rule (COM). In [74] the labels are pairs (M, μ) or (M, τ) where M are sequences of fusions. It is easy to see that our label μ corresponds to (μ, τ) if μ is a fusion label and to (\emptyset, μ) if it is an action label. The communication rule of [74] is

$$\frac{p \xrightarrow{(M, \bar{a}b)} p' \quad q \xrightarrow{(N, a'(c))} q'}{p \mid q \xrightarrow{(L, \tau)} p' \mid q' \{b/c\}} \qquad L = \begin{cases} MN[a = a'], & \text{if } a \neq a' \\ MN, & \text{if } a = a' \end{cases}$$

which resembles rule (COM) of Table 1. However it is considerably more complex since it must also collect the fusions due to matchings.

Proposition 1. *Under the label correspondende illustrate above, let $p \in \mathbf{P}$ be a π_- process, then $p \xrightarrow{\mu} q$ if, and only if, the same transition can be derived from the transition system in [74] (changing μ with the corresponding label of [74]).*

Proof. The (\Rightarrow) part trivially follows by induction on the length of the proof of $p \xrightarrow{\mu} q$. The (\Leftarrow) part follows by observing that the length of the fusions in labels of [74] is one, since π_- lacks the matching operator. \square

We recast the definition of *open bisimulation* given in [74] for π_- .

Definition 4 (Open Bisimulation). A symmetric relation $S \subseteq \mathbf{P} \times \mathbf{P}$ is an open bisimulation if whenever pSq ,

- if $p \xrightarrow{\alpha} p'$ then there is q' such that $q \xrightarrow{\alpha} q'$ and $p'Sq'$;
- if $p \xrightarrow{\varepsilon} p'$ then there is q' such that $q \xrightarrow{\varepsilon} q'$ and $p'Sq'$;
- if $p \xrightarrow{a=b} p'$ then there is q' such that $(q \xrightarrow{a=b} q' \vee q \xrightarrow{\varepsilon} q') \wedge \sigma_{a=b}(p')S\sigma_{a=b}(q')$,

where $\sigma_{a=b}$ is a substitution that maps a to b (or viceversa) and leaves the other names unchanged. Two processes p and q are open bisimilar, written $p \sim q$, when there is an open bisimulation relating them.

In order to compare the ordinary bisimilarity \sim with the one arising from the Leifer and Milner approach, it is convenient to introduce an additional bisimilarity for π_- .

Definition 5 (Syntactical Bisimilarity). The syntactical bisimilarity relation \simeq for π_- is obtained by simplifying the last condition of Definition 4 with

$$\text{if } p \xrightarrow{a=b} p' \text{ then there is } q' \text{ such that } q \xrightarrow{a=b} q' \text{ and } \sigma_{a=b}(p')S\sigma_{a=b}(q').$$

It is immediate to see that \simeq is finer than or equal to \sim . In fact its conditions for matching transition labels are more demanding than those for \sim .

Theorem 2 (Open Versus Syntactical Bisimilarity). We have $\simeq \subseteq \sim$.

An equivalence relation relating terms of an algebra is said a *congruence* if it is preserved by all the operation of the algebra, or, equivalently, if it is preserved in all the contexts of the language. In [74], \sim has been proven to be a congruence for the π -calculus.

5.2 Reactive Systems

A systematic method for deriving bisimulation congruence from reduction rules has been proposed by Leifer and Milner in [53,52], on turn inspired by [79], where the idea of interpreting $p \xrightarrow{c} q$ as “in the context c , p reacts and becomes q ” has been proposed. Also, the approach of observing contexts imposed on agents at each step has been introduced in [65], yielding the notion of *dynamic bisimilarity*. Following [28], we will call *borrowed context* the context c . The basic idea of [53,52] is to express “minimality” conditions for electing the context c among the (possibly infinite) ones that allow p to react. These conditions have been distilled by [53] in the notion of *relative push-out (RPO)* in categories of *reactive systems*. The RPO construction is reminiscent of the unification process of logic programming, which in fact can be given an interactive semantics in much the same style [13].

We want to apply this approach to a reduction semantics of π_- that reflects its LTS semantics, therefore, we collect here the main definitions and results of the RPO approach. We remark that Definitions 6, 7, 8 and Theorem 3 are borrowed from [53, 52] (aside from some minor notational conventions).

Let \mathbf{C} be an arbitrary category whose arrows are denoted by f, g, h, k and whose objects by m, n . Hereafter, $f;g$ will indicate arrow composition.

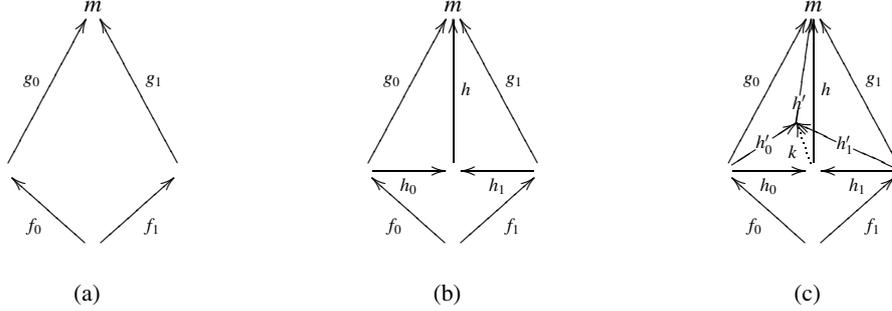


Fig. 2. Diagrams for Definitions 6

Definition 6 (Relative Push-Out and Idem Push-Out). Consider the commuting diagram in Figure 2(a) consisting of $f_0;g_0 = f_1;g_1$. A triple $\langle h_0, h_1, h \rangle$ is an RPO if diagram in Figure 2(b) commutes and for any triple $\langle h'_0, h'_1, h' \rangle$ satisfying $f_0;h'_0 = f_1;h'_1$ and $h'_i;h' = g_i$, for $i = 0, 1$ there exists a unique k such that diagram Figure 2(c) commutes. Diagram (a) is an idem push-out (IPO) if $\langle g_0, g_1, id \rangle$ is an RPO.

Definition 7 (Reactive System). A reactive system is a category \mathbf{C} with the following extra components:

- a distinguished (not necessarily initial) object \star ;
- a set of pairs of arrows $(l : \star \rightarrow m, r : \star \rightarrow m)$ called reaction rules;
- a subcategory \mathbf{D} of reactive contexts with the property that if $d; d'$ is an arrow of \mathbf{D} , then both d and d' are arrows in \mathbf{D} .

The IPO construction yields the definition of labelled transition out of a reduction semantics and the corresponding observational semantics.

Definition 8 (Labelled Transition and Bisimulation). We write $a \xrightarrow{f} a'$ iff there

exist a reaction rule (l, r) and a reactive context d such that $\begin{array}{ccc} & f & d \\ & \nearrow & \searrow \\ a & & \star \\ & \nwarrow & \nearrow \\ & l & r \end{array}$ is an IPO and

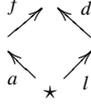
$a' = r;d$.

A symmetric binary relation $S \subseteq \bigcup_m \mathbf{C}[\star, m] \times \bigcup_m \mathbf{C}[\star, m]$, where $\mathbf{C}[x, y]$ is the set of all the arrows from x to y of category \mathbf{C} , is a bisimulation over \xrightarrow{f} iff for $(a, b) \in S$, if $a \xrightarrow{f} a'$ then there is b' such that $b \xrightarrow{f} b'$ and $(a', b') \in S$.

The central result of [53] can be stated as follows:

Theorem 3. *The largest bisimulation over \xrightarrow{f} is a congruence provided that \mathbf{C} has all redex-RPOs.*

The category \mathbf{C} has all *redex-RPOs* when for all reaction rules (l, r) , all arrow a, f and

all contexts d such that $a; f = l; d$ then the square  has an RPO.

5.3 A Reactive System for the Open π -Calculus

We shall specify a reactive system semantics for π_- taking actions and name substitutions as reactive contexts and by defining rules in such a way that the LTS will be essentially the same as the one defined in Section 5.1. However, the observational semantics resulting from the RPO approach considers labels as purely syntactical items and transitions can match only if they have identical labels. In the definition of open bisimilarity, instead, a proper fusion can be matched by an ε label. Thus it cannot be expected that the two bisimilarity relations coincide. In fact, we will show that the bisimilarity arising from the RPO approach is finer than open bisimilarity.

The reduction semantics of π_- is specified with rules of the form $P; \mu \rightarrow q$, where μ is an action or a fusion, $q \in \mathbf{P}$ and P is a *normalised process* (formally defined below). A rule $P; \mu \rightarrow q$ corresponds to a π_- transition $P \xrightarrow{\mu} q$, the only difference being that in the reactive system approach processes must be typed by (a natural number larger or equal than) the largest index of their free variables. Normalised processes can be thought of as being processes where all the occurrences of free variables are replaced by different variables $\{a_1, \dots, a_n\}$ ordered in some standard way. Normalised processes give a logic programming flavour to the reduction semantics. In fact, they are reminiscent of predicate symbols, while processes correspond to goals: as goals are instantiations of predicate symbols, any process $p \in \mathbf{P}$ can be regarded as the instantiation of a normalised process P . This amounts to say that, whenever p and $P; \mu$ (i.e., the instance and the head of the clause) unify, then a transition for p can be deduced. They unify whenever P is the normalised process of p . Moreover, the label is the borrowed context, which turns out either to be μ whenever μ is an action or to be a fusion not implied by the substitution mapping P to p , or else to be ε if it is implied.

Let $p \in \mathbf{P}$, we assume given two functions \hat{p} and σ_p such that

$$\text{fn}(\hat{p}) = \{a_1, \dots, a_n\}, \quad \hat{p} = \hat{p}, \quad p = \sigma_p(\hat{p}), \quad p = \sigma(q) \implies \hat{p} = \hat{q} \wedge \sigma_p = \sigma \circ \sigma_q,$$

where $\sigma_p : \text{fn}(\hat{p}) \rightarrow \text{fn}(p)$ and $\sigma : \text{fn}(q) \rightarrow \text{fn}(p)$ are surjective name substitutions homomorphically extended to π_- agents ($\sigma(_)$ stands for the extension of σ to agents). It is easy to show that \hat{p} is a linear process, namely each free variable occurs exactly once. Indeed, let $x \in \mathcal{N}$ occur twice in $p \in \mathbf{P}$ and assume by absurd that $\hat{p} = p$. Now, consider $p' \in \mathbf{P}$ to be the term obtained by replacing in p the first and the second occurrence of x with y and z , respectively. Then $p = \sigma(p')$, where $\sigma = \{y \mapsto x, z \mapsto x\}$, thus by definition, $\hat{p} = \hat{p}'$. But there is no $\sigma_{p'}$ such that $p' = \sigma_{p'}(\hat{p}') = \sigma_{p'}(p)$.

Notice that $\hat{_}$ and $\sigma_$ only involve syntactical aspects of agents, therefore they can be easily defined on the syntax trees of $\pi_$. For instance, \hat{p} might be defined as the agent having the same syntax tree of p where the i -th leaf is named by a_i , assuming that leaves are ordered according to a depth-first visit: substitution σ_p is defined accordingly. The order of leaves is arbitrary and different definitions might be possible, however, all of them differ only for a permutation of (the indexes of) $\text{fn}(\hat{p})$.

Definition 9 (Normalised Processes). *The processes that are fixpoints of $\hat{_}$ are the normalised processes and are ranged over by P .*

Before defining **PAC**, the category we work with, we specify its (basic) arrows where the underlying objects are elements of the set $\omega_\star = \omega \cup \{\star\}$ consisting of the natural numbers plus a distinguished element \star .

Definition 10 (Basic Arrows). *We define the following basic arrows.*

A normalised agent arrow $P_m : \star \rightarrow m$ is a pair consisting of a normalised process P and a natural number $m \in \omega$ such that, for any $a_n \in \text{fn}(P)$, $n \leq m$. We write P instead of P_m when $\text{fn}(P)$ contains exactly m names.

A fusion arrow from m to n is a surjective substitution from $\{a_1, \dots, a_m\}$ to $\{a_1, \dots, a_n\}$ written as $\sigma : m \rightarrow n$.

Action arrows are $\pi_$ actions parameterised on ω , more precisely

$$\bar{a}_i^m a_j : m \rightarrow m \quad a_i^m : m \rightarrow m + 1 \quad i, j \leq m$$

that respectively correspond to output and input transitions with the object name in the latter case being a_{m+1} .

A sequence arrow $\gamma : m_0 \rightarrow m_1$ is a tuple $\langle \mu_1, \dots, \mu_k, \sigma \rangle$ where $k \geq 0$, for each $0 < i \leq k$, $\mu_i : m_{i-1} \rightarrow m_i$ is an action arrow and $\sigma : m_k \rightarrow m'$ is a fusion arrow. In addition, we require that, if $\sigma(a_i) = \sigma(a_j)$ with $i < j$, then name a_j does not appear in actions μ_1, \dots, μ_k . Notice that for $k = 0$ we obtain fusion arrows while for $k = 1$ and $\sigma = \text{id}_m$ we obtain action arrows.

A process arrow $p : \star \rightarrow m$ is a tuple $\langle P, \mu_1, \dots, \mu_k, \sigma \rangle$ where $P : \star \rightarrow m_0$ is a normalised agent arrow and $\langle \mu_1, \dots, \mu_k, \sigma \rangle$ is a sequence arrow such that $\text{dom}(\mu_1) = m_0$. Notice that for $k = 0$, and $\sigma = \text{id}_{m_0}$ we obtain normalised agent arrows.

Definition 11 (Process-Action-Context Category). *The process-action-context category **PAC** is the category having as objects elements of ω_\star and as morphisms:*

1. the identity arrows $\text{id}_\star : \star \rightarrow \star$ and $\text{id}_m : m \rightarrow m$, the latter being the identity substitution on $\{a_1, \dots, a_m\}$;
2. the normalised agent arrows, the fusion arrows and the action arrows as generators; and
3. the arrows freely generated by 2 under the composition operation $_;$ subject to the usual associativity and identity axioms and, in addition, to the following axioms:

$$\frac{\sigma : n \rightarrow m \quad a_i^m : m \rightarrow m + 1}{\sigma; a_i^m = a_h^n; \sigma'}, \quad h = \min_l \{ \sigma(a_l) = a_i \} \quad \sigma' = \sigma[n + 1 \mapsto m + 1]$$

$$\frac{\sigma : n \rightarrow m \quad \bar{a}_i^m a_j : m \rightarrow m}{\sigma; \bar{a}_i^m a_j = \bar{a}_h^n a_k; \sigma}, \quad h = \min_l \{ \sigma(a_l) = a_i \} \quad k = \min_l \{ \sigma(a_l) = a_j \}$$

$(\sigma[n+1 \mapsto m+1])$ stands for the function that behaves as σ for any $a \in \{a_1, \dots, a_n\}$ and maps a_{n+1} to a_{m+1} .

The arrows of **PAC** can be given an intuitive standard representation that will be useful later in the proofs.

Proposition 2. *The arrows of **PAC** are exactly the process arrows, the sequence arrows and the identity arrow id_\star .*

Proof. First, observe that: (a) a normalised agent arrow is a process arrow with an empty sequence of actions and an identity substitution. (b) A fusion arrow σ is a sequence arrow with no action arrows and with σ as the fusion arrow; this also yields the identities id_m where $m \in \omega$. (c) Similarly, action arrows are sequence arrows with a single action arrow and the identity substitution. Now, we prove that the composition of a process (resp. sequence) arrow with a sequence arrow yields a process (resp. sequence) arrow. Consider $p: \star \rightarrow m$ and $\gamma: m \rightarrow n$ be the process arrow $\langle P, \mu_1, \dots, \mu_h, \sigma \rangle$ and the sequence arrow $\langle \mu'_1, \dots, \mu'_k, \sigma' \rangle$. By definition $p;\gamma = \langle P, \mu_1, \dots, \mu_h, \sigma, \mu'_1, \dots, \mu'_k, \sigma' \rangle$, and, observing that the two last axioms in 3 of Definition 11, allows to “exchange” a fusion arrow with an action arrow, we trivially conclude that $p;\gamma = \langle P, \mu_1, \dots, \mu_h, \mu'_1, \dots, \mu'_k, \sigma''; \sigma' \rangle$, for suitable μ'_1, \dots, μ'_k and σ'' . We remark that if, at any stage, two names are fused, say a_i and a_j with $i < j$, then a_j is replaced by a_i by definition and this guarantees that $\langle P, \mu_1, \dots, \mu_h, \mu'_1, \dots, \mu'_k, \sigma''; \sigma' \rangle$ is a process arrow. The prove is the same when considering composition between two sequence arrows.

The proof is concluded by showing that different arrows cannot be equated by axioms. In other words, we prove that the standard representation of an arrow is unique (up to identities). Indeed, by inspecting the initial part of the proof we see that equality between two arrow can be proved only by shifting back and forth fusion arrows or introducing/cancelling identities. In the former case, any shift uniquely determines both the action and the fusion arrow of the equated arrows (Definition 11). \square

As already mentioned, in the above definitions we have introduced typed versions (the type is a natural number m) of normalised agents and actions (substitutions are already typed), such that their names are in $\{a_1, \dots, a_m\}$. This is apparently required by the “box and wires” structure of category **PAC**. We continue defining typed versions of ordinary processes and of fusions.

Given a π_- agent p and a natural number m such that $m \geq \max\{k \mid a_k \in \text{fn}(p)\}$, we denote as $p_m: \star \rightarrow m$ the arrow $\hat{p}_n; \sigma$ where $n = |\text{fn}(\hat{p})| + m - |\text{fn}(p)|$ and $\sigma: n \rightarrow m$ is defined as:

- $\sigma(a_i) = \sigma_p(a_i)$, if $i \in \text{fn}(\hat{p})$,
- σ bijective and index monotone when restricted to $i \notin \text{fn}(\hat{p})$ (where σ is *index monotone* if $\sigma(a_i) = \sigma(a_h)$, $\sigma(a_j) = \sigma(a_k)$ and $i \leq j$ implies $h \leq k$).

Basically, p_m represents the agent p in terms of a normalised process with n variables. Given a fusion $a_i = a_j$ and $m \in \omega$, with $i < j \leq m$, the substitution $[a_i = a_j]_m: m \rightarrow m-1$ is defined as follows:

$$[a_i = a_j]_m(a_k) = \begin{cases} a_k, & k < j \\ a_i, & k = j \\ a_{k-1}, & j < k \leq m \end{cases}$$

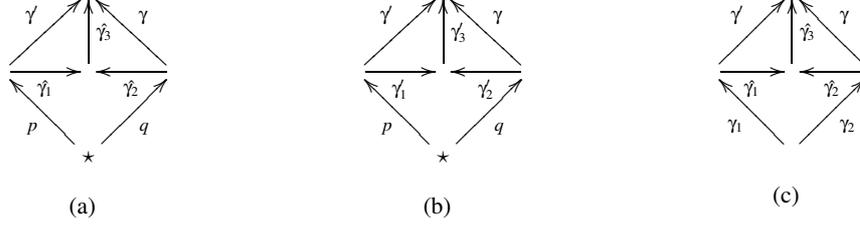


Fig. 3. Diagrams for proofs in Theorem 4

In words, $[a_i = a_j]_m$ maps the initial m names to the initial $m - 1$ by replacing a_j with a_i and mapping the names greater than a_j to their predecessors.

Definition 12 (PAC Reaction Rules). *The reaction rules are those generated by the following inference rules where $m \geq |\text{fn}(P)|$:*

$$\frac{P \xrightarrow{ab} q}{P_m; \bar{a}^m b \Longrightarrow q_m} \qquad \frac{P \xrightarrow{a(a_{m+1})} q}{P_m; \bar{a}^m \Longrightarrow q_{m+1}}$$

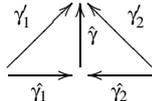
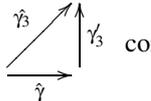
$$\frac{P \xrightarrow{a_i=a_j} q \quad i \neq j}{P_m; [a_i = a_j]_m \Longrightarrow q_m; [a_i = a_j]_m} \qquad \frac{P \xrightarrow{\varepsilon} q}{P_m \Longrightarrow q_m}$$

Definition 12 specify the reduction rules of **PAC** which rely on the LTS semantics of π_- . Take the first rule; it states that, if a normalised process P makes an output transition to q , then, in **PAC**, the corresponding arrow composed with the (output) action arrow (considered in at least $|\text{fn}(P)|$ variables m) reduces to the arrow representing q in m variables. Basically, the same can be said for the input and fusion transitions, aside that the former introduces the new variable a_{m+1} while the latter eliminates a variable. The last rule is just the special case of fusing a name with itself (i.e., $P; id$ is the lhs of the reduction).

Theorem 4. **PAC** has redex relative pushouts (RPOs).

Proof. We must prove that, given a reaction rule $q \Longrightarrow r$, for any process arrow p and any sequence arrows γ, γ' such that $p; \gamma' = q; \gamma$, there exist three sequence arrows $\hat{\gamma}_1, \hat{\gamma}_2$ and $\hat{\gamma}_3$ that satisfy the following conditions:

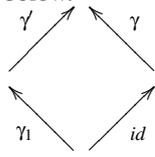
- the diagram in Figure 3(a) commutes, and
- for any sequence arrows γ'_1, γ'_2 and γ'_3 such that the diagram in Figure 3(b) com-

mutates, there is a unique $\hat{\gamma}$ such that both  and  commute.

Let us remark that the reduction contexts are all the arrows of **PAC**, however, for redex RPOs, γ and γ' can only be sequence arrows. Moreover, since $p;\gamma' = q;\gamma$, for Proposition 2, p and q are process arrows that are the composition of the *same* normalised linear arrow, say P with two sequence arrows. Hence, without loss of generality, it suffices to prove that there are arrows $\hat{\gamma}_1$, $\hat{\gamma}_2$ and $\hat{\gamma}_3$ forming an RPO for any diagram as in Figure 3(c).

The proof continues by case analysis.

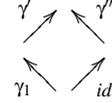
- First assume that γ_2 is an identity fusion arrow and consider the commuting diagram below.



We prove that $\hat{\gamma}_1 = id$, $\hat{\gamma}_2 = \gamma_1$ and $\hat{\gamma}_3 = \gamma'$ is an RPO. Indeed, condition a) trivially holds because the external square commutes by hypothesis. Consider three sequence arrows γ'_1 , γ'_2 and γ'_3 such that $\gamma' = \gamma'_1;\gamma'_3$, $\gamma = \gamma'_2;\gamma'_3$ and $\gamma_1;\gamma'_1 = \gamma'_2$. Then, assuming $\hat{\gamma} = \gamma'_3$ we obtain that the commutativity of the triangles corresponding to condition b) holds. Finally, uniqueness of $\hat{\gamma}$ is guaranteed by observing that $\hat{\gamma}_1$ is the identity.

- Let γ_2 is a generic fusion arrow σ . By Proposition 2, there is a sequence arrow γ''

such that $\sigma;\gamma = \gamma''$. Hence, we can equivalently prove that



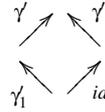
has an RPO,

which hold by the previous case.

- Finally, assume that γ_2 is an action arrow μ . By hypothesis, $\gamma_1;\gamma' = \mu;\gamma$, then, by Proposition 2,

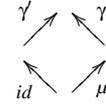
- either $\gamma_1 = \mu;\gamma'_1$
- or γ_1 is the identity and $\gamma' = \mu;\gamma''$.

In the former case, the proof reduces to show that



has an RPO, which hold

by the previous case. While, in the latter case, the redex diagram is



and,

proceeding as before, it is easy to see that μ , id and γ constitute an RPO.

□

Definition 13 (Labelled Transitions). *The diagram in Figure 3(a) is an IPO when it is an RPO and $\hat{\gamma}_1 = \gamma'$, $\hat{\gamma}_2 = \gamma$ and $\hat{\gamma}_3 = id$. We write $p \xrightarrow{\gamma'} r;\gamma$ when there is a reduction rule $q \Longrightarrow r$ and the diagram Figure 3(a) is an IPO. This defines a LTS. The corresponding bisimilarity according to Definition 8 is denoted as \simeq .*

The results in [53] and Theorem 12 guarantee the following corollary.

Corollary 1. *Bisimilarity relation \simeq is a congruence.*

The LTS of Definition 13 is essentially the same as in Section 5.1 indeed, the states are π_- processes and it is possible to show that the IPOs of **PAC** characterise the transitions of [74]. Thus bisimilarity relation \simeq essentially coincides with syntactic bisimilarity \simeq .

Theorem 5 (\simeq **Is** \simeq). *Relation \simeq , which is defined on process arrows p_m , when restricted to those p_m with $m = \max\{k \mid a_k \in \text{fn}(p)\}$, coincides with \simeq .*

Notice that, due to the missing restriction operator, two agents with different sets of free names cannot be bisimilar. Thus, observing actions or typed actions does not make a difference.

From Theorem 2 we know that \simeq is finer than or equal than \sim . It is easy to see that it is finer from this example. Consider the following processes

$$p = (\bar{a}b \mid a'(c)) + (\bar{d}e \mid d(f)) \quad q = \bar{a}b.a'(c) + a'(c).\bar{a}b + (\bar{d}e \mid d(f)),$$

then $p \sim q$ because the synchronisation between $\bar{a}b$ and $a'(c)$ in a context that identifies a and a' is matched by the (unique) synchronisation of q . On the contrary, $p \not\sim q$ because the transition $p \xrightarrow{a=a'}$ cannot be matched by q . We can thus conclude the following fact.

Theorem 6. *Relation \simeq when restricted to those p_m with $m = \max\{k \mid a_k \in \text{fn}(p)\}$, is finer than \sim .*

6 Conclusions

In the paper we surveyed some of the approaches for model checking nominal calculi, focusing on HD-automata and on the existing toolkits for handling them. We also introduced a simplified version of open π -calculus and we proposed a bisimilarity semantics for it based on a reactive system with observed borrowed contexts. This approach has been proposed by Leifer and Milner [53, 52] and further developed by Sassone, Lack and Sobocinski [76, 78, 50] using G-categories and adhesive categories. The generality of the reactive system approach gives some hope that interesting abstractions of the SOC paradigm could also be modelled that way.

However we noticed that the transition system we obtain in this manner is not really suitable for a HD-automata implementation, since new names are never forgotten. To avoid this problem, it might be necessary to take advantage of the extended theory developed by Sassone, Lack and Sobocinski [76, 78, 50]. In particular, the actions of nominal calculi which forget names could be represented as cospans of suitable adhesive categories. In fact several expressive graph-like structures can be represented by adhesive categories and the existing theory guarantees that the categories of their cospans have the *all redex-RPOs* property [77].

Acknowledgements

The authors thank Vladimiro Sassone and Pawel Sobocinski for their helpful comments on an earlier draft of this paper.

References

1. M. Abadi and A. Gordon. A Calculus for Cryptographic Protocols: The Spi Calculus. *Inf. and Comp.*, 148(1):1–70, January 1999.
2. G. Baldi, A. Bracciali, G. Ferrari, and E. Tuosto. A Coordination-based Methodology for Security Protocol Verification. In *WISP, ENTCS*, Bologna, Italy, June 2004. Elsevier. To appear.
3. T. Ball and S. Rajamani. The SLAM Toolkit. In G. Berry, H. Comon, and A. Finkel, editors, *CAV*, volume 2102 of *LNCS*, pages 260–264. Springer, 2001.
4. N. Benton, L. Cardelli, and C. Fournet. Modern Concurrency Abstractions for C#. *TOPLAS*, 26(5):269–304, Sept. 2004.
5. M. Boreale and M. Buscemi. A Framework for the Analysis of Security Protocols. In L. Brim, P. Jančar, M. Křetínský, and A. Kučera, editors, *CONCUR*, volume 2421 of *LNCS*, pages 483–498. Springer, Aug. 2002.
6. M. Boreale and R. De Nicola. A Symbolic Semantics for the π -calculus. *Inf. and Comp.*, 126(1):34–52, April 1996.
7. A. Bouali, S. Gnesi, and S. Larosa. The Integration Project for the JACK Environment. In *EATCS Bull.*, volume 54, pages 207–223. Centrum voor Wiskunde en Informatica (CWI), 1994.
8. A. Bouali, A. Ressouche, V. Roy, and R. de Simone. The FC2TOOLS Set. In R. Alur and T. Henzinger, editors, *CAV*, volume 1102 of *LNCS*, pages 441–445, New Brunswick, NJ, USA, 1996. Springer.
9. A. Bracciali, A. Brogi, G. Ferrari, and E. Tuosto. Security Issues in Component Based Design. In U. Montanari and V. Sassone, editors, *ConCoord: International Workshop on Concurrency and Coordination*, volume 54 of *ENTCS*, Lipari Island - Italy, July 2001. Elsevier.
10. G. Brat, K. Havelund, S. Park, and W. Visser. Model Checking Programs. *Automated Software Engineering*, 10(2):203–232, 2003.
11. R. Bruni, C. Laneve, and U. Montanari. Orchestrating Transactions in Join Calculus. In L. Brim, P. Jančar, M. Křetínský, and A. Kučera, editors, *CONCUR*, volume 2421 of *LNCS*, pages 321–336. Springer, 2002.
12. R. Bruni, H. Melgratti, and U. Montanari. Theoretical Foundations for Compensations in Flow Composition Languages. In *POPL*, 2005. To appear.
13. R. Bruni, U. Montanari, and F. Rossi. An Interactive Semantics of Logic Programming. *Theory and Practice of Logic Programming.*, 1(6):647–690, 2001.
14. M. Butler and C. Ferreira. An Operational Semantics for StAC, a Language for Modelling Long-Running Business Transactions. In R. De Nicola, G. Ferrari, and G. Meredith, editors, *COORDINATION*, volume 2949 of *LNCS*, pages 87–104. Springer, 2004.
15. L. Caires and L. Cardelli. A Spatial Logic for Concurrency (Part I). *Inf. and Comp.*, 186, 2003.
16. L. Caires and L. Cardelli. A Spatial Logic for Concurrency II. *TCS*, 322(3):517–565, Sept. 2004.
17. G. Cattani and P. Sewell. Models for Name-Passing Processes: Interleaving and Causal (Extended Abstract). *Inf. and Comp.*, 190(2):136–178, May 2004.
18. E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. MIT Press, 1999.
19. E. Clarke and J. Wing. Formal Methods: State of the Art and Future Directions. *ACM Computing Surveys*, 28(4):626–643, December 1996.
20. S. Clarke, Edmund M. Jha and W. Marrero. Using State Space Exploration and a Natural Deduction Style Message Derivation Engine to Verify Security Protocols. In *In Proc. IFIP Working Conference on Programming Concepts and Methods (PROCOMET)*, 1998.

21. R. Cleaveland, J. Parrow, and B. Steffen. The Concurrency Workbench: A Semantics-Based Tool for the Verification of Concurrent Systems. *TOPLAS*, 15(1):36–72, Jan. 1993.
22. S. Conchon and F. Le Fessant. Jocaml: Mobile Agents for Objective-Caml. In *International Symposium on Agent Systems and Applications*, pages 22–29, Palm Springs, California, Oct. 1999.
23. J. Corbett, M. Dwyer, J. Hatcliff, S. Laubach, S. Corina, Robby, and H. Zheng. Bandera: Extracting Finite-state Models from Java Source Code. In *International Conference on Software Engineering*, pages 439–448, Limerick, Ireland, June 2000.
24. F. Curbera, R. Khalaf, N. Mukhi, S. Tai, and S. Weerawarana. The Next Step in Web Services. *CACM*, 46(10):29–34, 2003.
25. M. Dam. Model Checking Mobile Processes. *Inf. and Comp.*, 129(1):35–51, 1996.
26. M. Dam. Proof Systems for π -Calculus Logics. *Logic for concurrency and synchronisation*, pages 145–212, 2003.
27. G. Denker and J. Millen. CAPSL Integrated Protocol Environment. Technical report, Computer Science Laboratory, SRI International, Menlo Park, CA, 1999.
28. H. Ehrig and B. König. Deriving Bisimulation Congruences in the DPO Approach to Graph Rewriting. In I. Walukiewicz, editor, *FoSSaCS*, volume 2987, pages 151–166. LNCS, 2004.
29. J. Fernandez. An Implementation of an Efficient Algorithm for Bisimulation Equivalence. *Science of Computer Programming*, 13(2–3):219–236, May 1990.
30. J. Fernandez and L. Mounier. On-the-fly Verification of Behavioural Equivalences and Preorders. In K. Larsen and A. Skou, editors, *CAV*, volume 575 of LNCS, pages 181–191. Springer, July 1991.
31. G. Ferrari, G. Ferro, S. Gnesi, U. Montanari, M. Pistore, and G. Ristori. An Automata Based Verification Environment for Mobile Processes. In E. Brinksma, editor, *TACAS*, volume 1217 of LNCS, pages 275–289. Springer, April 1997.
32. G. Ferrari, S. Gnesi, U. Montanari, and M. Pistore. A Model Checking Verification Environment for Mobile Processes. *TOPLAS*, 12(4):1–34, 2004.
33. G. Ferrari, S. Gnesi, U. Montanari, R. Raggi, G. Trentanni, and E. Tuosto. Verification on the WEB. In J. Augusto and U. Ultes-Nitsche, editors, *VVEIS*, pages 72–74, Porto, Portugal, April 2004. INSTICC Press.
34. G. Ferrari, U. Montanari, and M. Pistore. Minimizing Transition Systems for Name Passing Calculi: A Co-algebraic Formulation. In M. Nielsen and U. Engberg, editors, *FOSSACS 2002*, volume 2303 of LNCS, pages 129–143. Springer, 2002.
35. G. Ferrari, U. Montanari, and E. Tuosto. From Co-algebraic Specifications to Implementation: The Mihda toolkit. In F. de Boer, M. Bonsangue, S. Graf, and W. de Roever, editors, *FMCO*, volume 2852 of LNCS, pages 319 – 338. Springer, November 2002.
36. G. Ferrari, U. Montanari, and E. Tuosto. Coalgebraic Minimisation of HD-automata for the π -Calculus in a Polymorphic λ -Calculus. *TCS*, 2004. To appear.
37. M. Fiore, G. Plotkin, and D. Turi. Abstract Syntax and Variable Binding (Extended Abstract). In *LICS*, pages 193–202, Trento, Italy, July 1999. IEEE.
38. P. Fiore and M. Abadi. Computing Symbolic Models for Verifying Cryptographic Protocols. In *Computer Security Foundations Workshop*, CSFW, pages 160–173, Cape Breton, Nova Scotia, Canada, June 2001. IEEE.
39. R. Focardi and R. Gorrieri. A Classification of Security Properties. *J. of Computer Security*, 3(1), 1995.
40. M. Gabbay and A. Pitts. A New Approach to Abstract Syntax Involving Binders. In G. Longo, editor, *LICS*, pages 214–224, Trento, Italy, July 1999. IEEE.
41. A. Gordon. Notes on Nominal Calculi for Security and Mobility. In R. Focardi and R. Gorrieri, editors, *FOSAD*, volume 2171 of LNCS, pages 262–330. Springer, September 2002.
42. M. Hennessy and H. Lin. Symbolic Bisimulations. *TCS*, 138(2):353–389, February 1995.

43. T. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy Abstraction. In *POPL*, pages 58–70. ACM Press, 2002.
44. D. Hirschhoff. On the Benefits of Using the up-to Techniques for Bisimulation Verification. In W. Cleaveland, editor, *TACAS*, volume 1579 of *LNCS*, pages 285–299, Amsterdam, March 1999. Springer.
45. G. Holzmann. The Model Checker Spin. *TSE*, 23(5):279–295, May 1997.
46. G. Holzmann. *The Spin Model Checker: Primer and Reference Manual*. Addison-Wesley, Sept. 2003.
47. K. Honda. Elementary Structures in Process Theory (1): Sets with Renaming. *MSCS*, 10(5):617–663, 2000.
48. O. Jensen and R. Milner. Bigraphs and Transitions. In *POPL*, pages 38–49. ACM Press, 2003.
49. P. Kanellakis and S. Smolka. CCS Expressions, Finite State Processes and Three Problem of Equivalence. *Inf. and Comp.*, 86(1):272–302, 1990.
50. S. Lack and P. Sobociński. Adhesive Categories. In I. Walukiewicz, editor, *FoSSaCS*, volume 2987 of *LNCS*, pages 273–288, Barcelona, March 2004. Springer.
51. C. Laneve and G. Zavattaro. Foundations of Web Transactions. In *FoSSaCS*, LNCS, 2005. To appear.
52. J. Leifer. *Operational Congruences for Reactive Systems*. PhD thesis, Computer Laboratory, University of Cambridge, Cambridge, UK, 2001.
53. J. Leifer and R. Milner. Deriving Bisimulation Congruences for Reactive Systems. In C. Palamidessi, editor, *CONCUR*, volume 1877 of *LNCS*, pages 243–258, University Park, PA, USA, August 22–25 2000. Springer.
54. H. Lin. Complete Inference Systems for Weak Bisimulation Equivalences in the π -Calculus. *Inf. and Comp.*, 180(1):1–29, January 2003.
55. G. Lowe. Towards a Completeness Result for Model Checking of Security Protocols. In *CSFW*. IEEE, 1998.
56. W. Marrero, E. Clarke, and S. Jha. Model Checking for Security Protocols. In *Formal Verification of Security Protocols*, 1997.
57. K. McMillan. *Symbolic Model Checking*. Kluwer Academic Publisher, 1993.
58. R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
59. R. Milner, J. Parrow, and D. Walker. A Calculus of Mobile Processes, I and II. *Inf. and Comp.*, 100(1):1–40, 41–77, September 1992.
60. J. Mitchell, M. Mitchell, and U. Ster. Automated analysis of cryptographic protocols using mur ϕ . In *CSFW*, pages 141–151. IEEE, 1997.
61. U. Montanari and M. Buscemi. A First Order Coalgebraic Model of π -Calculus Early Observational Equivalence. In L. Brim, P. Jančar, M. Křetínský, and A. Kučera, editors, *CONCUR*, volume 2421 of *LNCS*, pages 449–465. Springer, Aug. 2002.
62. U. Montanari and M. Pistore. Checking Bisimilarity for Finitary π -Calculus. In I. Lee and S. Smolka, editors, *CONCUR*, volume 962 of *LNCS*, pages 42–56, Philadelphia, PA, USA, Aug. 1995. Springer.
63. U. Montanari and M. Pistore. History Dependent Automata. Technical report, Computer Science Department, Università di Pisa, 1998. TR-11-98.
64. U. Montanari and M. Pistore. π -Calculus, Structured Coalgebras, and Minimal HD-Automata. In M. Nielsen and B. Roman, editors, *MFCS*, volume 1983 of *LNCS*. Springer, 2000. An extended version will be published on Theoretical Computer Science.
65. U. Montanari and V. Sassone. Dynamic Congruence vs. Progressing Bisimulation for CCS. *Fundamenta Informaticae*, 16:171–196, 1992.
66. R. Needham. *Names*. Addison-Wesley (Mullender Ed.), 1989.
67. R. Paige and R. Tarjan. Three Partition Refinement Algorithms. *SIAM Journal on Computing*, 16(6):973–989, December 1987.

68. M. Papazoglou. Service-Oriented Computing: Concepts, Characteristics and Directions. In *Web Information Systems Engineering (WISE'03)*, LNCS, pages 3–12. Springer, 2003.
69. D. Park. Concurrency and Automata on Infinite Sequences. In *Theoretical Computer Science, 5th GI-Conf.*, volume 104 of LNCS, pages 167–183. Springer, Karlsruhe, March 1981.
70. J. Parrow and B. Victor. The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes. In *LICS*. IEEE, 1998.
71. M. Pistore. *History Dependent Automata*. PhD thesis, Computer Science Department, Università di Pisa, 1999.
72. M. Pistore and D. Sangiorgi. A Partition Refinement Algorithm for the π -Calculus. *Inf. and Comp.*, 164(2):467–509, 2001.
73. D. Sangiorgi. On the Bisimulation Proof Method (Extended Abstract). In J. Wiedermann and P. Hájek, editors, *MFCS*, volume 969 of LNCS, pages 479–488, Prague, August–September 1995. Springer.
74. D. Sangiorgi. A Theory of Bisimulation for the π -Calculus. *Acta Informatica*, 33(1):69–97, 1996.
75. D. Sangiorgi and D. Walker. *The π -Calculus: a Theory of Mobile Processes*. Cambridge University Press, 2002.
76. V. Sassone and P. Sobociński. Deriving Bisimulation Congruences using 2-categories. *Nordic J. of Computing*, 10(2), 2003.
77. V. Sassone and P. Sobociński. Congruences for Contextual Graph-Rewriting. Technical Report RS-14, BRICS, June 2004.
78. V. Sassone and P. Sobociński. Locating Reaction with 2-Categories. *TCS*, 2004. To appear.
79. P. Sewell. From Rewrite Rules to Bisimulation Congruences. *LNCS*, 1466, 1998.
80. P. Sewell. Applied π – A Brief Tutorial. Technical Report 498, Computer Laboratory, University of Cambridge, Aug. 2000.
81. H. Smith and P. Fingar. Workflow is Just a Pi process. Available at <http://www.bpm3.com/picalculus>, 2003.
82. V. Vanackère. The TRUST protocol analyser. Automatic and Efficient Verification of Cryptographic Protocols. In *VERIFY02*, 2002.
83. B. Victor and F. Moller. The Mobility Workbench — A Tool for the π -Calculus. In D. Dill, editor, *CAV*, volume 818 of LNCS, pages 428–440. Springer, 1994.
84. P. Yang, C. Ramakrishnan, and S. Smolka. A Logical Encoding of the π -Calculus: Model Checking Mobile Processes Using Tabled Resolution. *STTT*, 6(1):38–66, July 2004.