

Finite-Control Mobile Ambients

Witold Charatonik^{1,2}, Andrew D. Gordon³, and Jean-Marc Talbot⁴

¹ Max-Planck-Institut für Informatik, Germany.

² University of Wrocław, Poland.

³ Microsoft Research, United Kingdom.

⁴ Laboratoire d'Informatique Fondamentale de Lille, France.

Abstract. We define a finite-control fragment of the ambient calculus, a formalism for describing distributed and mobile computations. A series of examples demonstrates the expressiveness of our fragment. In particular, we encode the choice-free, finite-control, synchronous π -calculus. We present an algorithm for model checking this fragment against the ambient logic (without composition adjunct). This is the first proposal of a model checking algorithm for ambients to deal with recursively-defined, possibly nonterminating, processes. Moreover, we show that the problem is PSPACE-complete, like other fragments considered in the literature. Finite-control versions of other process calculi are obtained via various syntactic restrictions. Instead, we rely on a novel type system that bounds the number of active ambients and outputs in a process; any typable process has only a finite number of derivatives.

1 Introduction

The ambient calculus [6] is a formalism for describing distributed and mobile computation in terms of *ambients*, named collections of running processes and nested sub-ambients. A state of computation has a tree structure induced by ambient nesting. Mobility is represented by re-arrangement of this tree (an ambient may move inside or outside other ambients) or by deletion of a part of this tree (a process may dissolve the boundary of some ambient, revealing its content).

There are proposals for analysing systems expressed in the ambient calculus and its variants [2,14] via several techniques, such as equational reasoning [3], type systems [8], control flow analysis [16], and abstract interpretation [13]. Still, the ambient calculus is Turing-complete, and little attention has been paid to finding expressive finite-state fragments that admit automatic verification via state-space exploration. The goal of this work is to identify such a fragment, and to develop a model checking algorithm for verifying properties expressible in the ambient logic [5,7]. The long term intention is that automatic verification tools for a finite-state ambient calculus will be useful either by themselves or in conjunction with methods for obtaining finite-state abstractions of infinite-state systems. Similar abstractions [9] are being developed for the π -calculus [15], the formalism from which the ambient calculus derives.

A finite-state version of π exists [12]. It is described as a finite-control calculus because its control structure is finite. Starting in any state, the number of states reachable via internal reduction steps is finite. However, if we allow inputs of external data, the number of reachable states may be infinite. We define in this paper a

finite-control ambient calculus. It is a substantial extension of the replication-free fragment [5,10,11] (sometimes referred to as the “finite-state ambient calculus”). In particular, in the replication-free fragment every process can make only a finite number of computation steps—no recursion or iteration is possible.

We begin, in Section 2, by presenting a variant of the ambient calculus in which recursion is defined by means of an explicit recursive definition instead of replication. We specify standard spatial rearrangements via a structural congruence relation on processes, and specify the operational semantics as a reduction relation on processes. This variant is easily seen to simulate the original one. Then, in Section 3, we design a type system for ambient processes and show that typability of a process guarantees finitary behaviour. The basic idea of the type system is to count the number of active outputs and ambients in a process. Theorem 1 asserts that the number of processes, up to structural congruence, reachable from any typable process is finite. We define the finite-control fragment as those processes that are typable. In contrast, finite-control fragments of the π -calculus are defined via simple syntactic restrictions. In Section 4, we explore the expressivity of our calculus by presenting and developing some standard examples, including an encoding of a finite-control π -calculus.

Turning to the verification problem, Section 5 reviews the syntax and semantics of the ambient logic we use to specify process properties. We prove that the verification problem—model checking against the ambient logic without composition adjunct—is decidable for the finite-control fragment. To achieve this, we adapt the model checking algorithm from [11]. Theorem 2 states that the algorithm is correct with respect to the semantics of the logic. Moreover, our final result is Theorem 3, that the verification problem remains PSPACE-complete, which is the same complexity as verifying the replication-free fragment against the same logic.

A difficulty in designing a finite-control fragment of a process calculus is striking a balance between the expressivity of the fragment and the complexity of the verification problem. The general goal is to make the calculus as expressive as possible while keeping the verification problem decidable. The methods we use differ substantially from the methods used to define finite-control π -calculi. Therefore, for the sake of a simple exposition we omit several possible features from our finite-control ambient calculus while including enough to model interesting iterative computations. Section 6 discusses some of these additional features. Section 7 concludes the paper.

2 An Ambient Calculus with Recursion

We present in this section an ambient calculus with recursive definitions instead of replication. We give examples of the calculus in Section 4; see [6] for more elementary examples.

The following table defines the syntax of *capabilities* and *processes* of our calculus. We assume countably many *names* ranging over by n, m, a, b, c, \dots and countably many *identifiers* ranging over by A, B, C, \dots . For the sake of a simple presentation we allow only names to be communicated whereas the original calculus allows also the transmission of sequences of capabilities.

Processes and Capabilities:

$\alpha ::=$	capabilities		
$\text{in } n$	can enter n	$\text{out } n$	can exit n
$\text{open } n$	can open n		
$P, Q, R ::=$	processes		
$\mathbf{0}$	inactivity	$P \mid Q$	composition
$n[P]$	ambient	$\alpha.P$	action prefix
$(n).P$	input	$\langle n \rangle$	output
$(\nu n)P$	name restriction	A	identifier
$(\text{fix } A=P)$	recursion		

We consider input (n) , name restriction (νn) to be binders for the name n and $\text{fix } A$ to be a binder for the identifier A . A name n or an identifier A occurring in the scope of respectively $(n), (\nu n)$ and $\text{fix } A$ is *bound*. Otherwise it is *free*. We write $fn(P)$ for the set of free names in P . We say that a process is *closed* if it contains no free identifier. We identify processes up-to capture-avoiding α -renaming of both bound names and bound identifiers. For instance, $(\text{fix } A=(\nu n)\text{open } n.A)$ and $(\text{fix } B=(\nu m)\text{open } m.B)$ are identical processes. Slightly abusing the notation, we write $bn(P)$ for the bound names in an implicitly given syntactic representation of P . We write $P\{m \leftarrow n\}$ for the outcome of substituting n for each free occurrence of m in P . Similarly, $P\{A \leftarrow Q\}$ is the outcome of substituting Q for each free occurrence of A in P . We will assume without loss of generality that two distinct bound identifiers are different as well as being distinct from any free identifier.

The semantics of our calculus is given by two relations. The *reduction relation* $P \rightarrow Q$ describes the evolution of ambient processes over time. We write \rightarrow^* for the reflexive and transitive closure of \rightarrow . The *structural congruence* relation $P \equiv Q$ relates different syntactic representations of the same process; it is used to define the reduction relation.

Structural Congruence $P \equiv Q$:

$P \equiv P$	(Str Refl)	$P \equiv Q \Rightarrow (\nu n)P \equiv (\nu n)Q$	(Str Res)
$P \equiv Q \Rightarrow Q \equiv P$	(Str Symm)	$P \equiv Q \Rightarrow P \mid R \equiv Q \mid R$	(Str Par)
$P \equiv Q, Q \equiv R \Rightarrow P \equiv R$	(Str Trans)	$P \equiv Q \Rightarrow n[P] \equiv n[Q]$	(Str Amb)
$P \mid \mathbf{0} \equiv P$	(Str Par Zero)	$P \equiv Q \Rightarrow \alpha.P \equiv \alpha.Q$	(Str Action)
$P \mid Q \equiv Q \mid P$	(Str Par Comm)	$P \equiv Q \Rightarrow (n).P \equiv (n).Q$	(Str Input)
$(P \mid Q) \mid R \equiv P \mid (Q \mid R)$	(Str Par Assoc)	$P \equiv Q \Rightarrow (\text{fix } A=P) \equiv (\text{fix } A=Q)$	(Str Fix)
$(\nu n)\mathbf{0} \equiv \mathbf{0}$	(Str Res Zero)	$(\text{fix } A=A) \equiv \mathbf{0}$	(Str Fix Id)
$(\nu n)(\nu m)P \equiv (\nu m)(\nu n)P$	(Str Res Res)	$(\text{fix } A=P) \equiv P\{A \leftarrow (\text{fix } A=P)\}$	(Str Fix Rec)
$(\nu n)(P \mid Q) \equiv P \mid (\nu n)Q$ if $n \notin fn(P)$	(Str Res Par)		
$(\nu n)(m[P]) \equiv m[(\nu n)P]$ if $n \neq m$	(Str Res Amb)		

Reduction: $P \rightarrow Q$

$n[\text{in } m.P \mid Q] \mid m[R] \rightarrow m[n[P \mid Q] \mid R]$	(Red In)
$m[n[\text{out } m.P \mid Q] \mid R] \rightarrow n[P \mid Q] \mid m[R]$	(Red Out)
$\text{open } n.P \mid n[Q] \rightarrow P \mid Q$	(Red Open)
$\langle m \rangle \mid (n).P \rightarrow P\{n \leftarrow m\}$	(Red I/O)

$P \rightarrow Q \Rightarrow P \mid R \rightarrow Q \mid R$	(Red Par)
$P \rightarrow Q \Rightarrow n[P] \rightarrow n[Q]$	(Red Amb)
$P \rightarrow Q \Rightarrow (\nu n)P \rightarrow (\nu n)Q$	(Red Res)
$P' \equiv P, P \rightarrow Q, Q \equiv Q' \Rightarrow P' \rightarrow Q'$	(Red \equiv)

As in other process calculi with recursion, it is convenient to regard certain unwanted recursive processes as ill-formed, and to disregard them. An example is $(\text{fix } A=A \mid A)$. We define a well-formed process as follows.

Definition 1. *A process P is said to be well-formed if every recursive subprocess $(\text{fix } A=Q)$ of P satisfies the following two requirements: (i) A is the only free identifier in Q , and (ii) A occurs at most once in Q .*

From now on, we only consider well-formed processes. These processes are stable with respect to structural congruence and reduction: if P is well-formed and either $P \equiv P'$ or $P \rightarrow P'$ then P' is well-formed.

As in the π -calculus [15], we can easily simulate replication with recursion. To simulate $!P$, which behaves like an unbounded number of replicas of P running in parallel, we introduce a new identifier A_P and replace $!P$ by $(\text{fix } A_P=P \mid A_P)$. The resulting process is well-formed. This encoding of replication fulfils the axioms of structural congruence for replication $!P \equiv P \mid !P$ and $!0 \equiv 0$ given in [6]. It does not obey the two additional axioms $!P \equiv !!P$ and $!(P \mid Q) \equiv !P \mid !Q$ from [5], but these axioms are unnecessary for computing reduction steps.

Sangiorgi [17] also considers an ambient calculus with recursion. A difference is that our formulation allows recursively defined ambient structures. In a recursion $(\text{fix } A=P)$, the identifier A can appear in P within an ambient construct; for example, the processes $(\text{fix } A=m[A])$ and $(\text{fix } A=\text{open } n.m[A])$ are well-formed. The latter but not the former belongs to the finite-control fragment defined next.

3 A Finite-Control Ambient Calculus

The finite-control (synchronous) π -calculus [12] is obtained by disallowing parallel composition through recursion. So, a finite-control π -calculus process is a finite parallel composition of threads each of which is a recursive process without parallel composition. This ensures that there is only finitely many pairwise non-congruent configurations reachable from such a process. In the ambient calculus this restriction is both too strong and too weak. It is too strong because it limits the admissible computation too much. In particular, due to the asynchronous communication mechanism in the ambient calculus, it completely excludes communication in recursive programs. On the other hand, it is too weak, as exemplified below, because it does not ensure finitary behaviour of processes. The example shows that in the ambient calculus a bound on the number of parallel threads gives neither a bound on the size of reachable processes nor a bound on the number of possible interactions between threads. This is in contrast to the situation in the π -calculus, and seems to arise from the spatial characteristics of the ambient calculus.

Example 1. Consider the two ambient processes P_A and P_B defined respectively by $(\text{fix } A = n[\text{open } m.A])$ and $(\text{fix } B = m[\text{in } n.B])$. Neither process contains parallel composition. However, the process $P_A \mid P_B$ reduces in $2k$ steps to $\underbrace{n[\cdots n[P_A \mid P_B]\cdots]}_k$,

denoted $(P_A \mid P_B)_k$. Since for $k \neq k'$, $(P_A \mid P_B)_k \not\equiv (P_A \mid P_B)_{k'}$, there are infinitely many non-congruent processes reachable from $P_A \mid P_B$.

Now, if we place another process $m[(\text{fix } C = \text{in } n.C) \mid (\text{fix } D = \text{out } n.D)]$ in parallel with $P_A \mid P_B$, this process can traverse the structure of $(P_A \mid P_B)_k$ in an arbitrary way. Thus, although there are only four recursively defined processes and none of them contains parallel composition, they may create an arbitrary number of locations and interact in any of these locations.

This example shows that directly adopting the syntactic restriction from the finite-control synchronous π -calculus is problematic, but it does not show undecidability of the verification problem. However, if we adopt a more liberal condition from the finite-control asynchronous π -calculus [1] (which ensures that there is only a bounded number of active threads and seems more appropriate here due to the asynchronous communication used in ambient calculus), we obtain undecidability even for the reachability problem. One can adapt the encoding of the Post correspondence problem from [11], in which only finitely many active threads are used.

Thus, one may consider a more severe syntactic restriction, to forbid both parallel composition and ambient construction within recursion. We will see later on that this restriction indeed ensures finite-control. Still, we can see at once that it is too drastic. In such a restricted process only sequences of action prefixes, each invoking a capability, could be defined recursively. Moreover, those sequences would operate on a process whose spatial structure has a bounded size. Inspecting the individual effect of such sequences, one sees that the open capability is somehow the more powerful as it changes the spatial structure of the process by deleting part of it whereas both capabilities in and out only re-arrange this structure. As a consequence, only finitely many occurrences of the powerful capability open can be executed by such a restricted process.

Instead of defining finite-control by means of syntactic restrictions over processes, we adopt a semantic point of view based on a type system. Intuitively, a type of a process P is a natural number that bounds the number of active outputs and ambients in any process reachable from P . We present the type system in Section 3.1. In Section 3.2 we show that typability ensures finitary computation.

3.1 The Type System FC

A *type environment* Γ is a finite set of pairs $\{(A_1, \tau_1), \dots, (A_n, \tau_n)\}$ such that each A_i is an identifier, τ_i is a natural number and for any two pairs (A_i, τ_i) and (A_j, τ_j) , $i \neq j$ implies $A_i \neq A_j$. We say that an environment Γ is defined for A if Γ contains a pair (A, τ) . Whenever Γ is defined for no identifier, we simply write \emptyset .

Definition 2. Given a type environment Γ , a type judgment $\Gamma \vdash Q : \tau$ holds for a process Q and a natural number τ if there exists a finite proof tree built with the inference rules from the table below such that its root is labelled by $\Gamma \vdash Q : \tau$ and none of its leaves contains a type judgment.

Process Typing : $\Gamma \vdash P : \tau$

(Identifier)	(Zero)	(Par)	(Res)
$\frac{A \text{ is identifier, } (A, \tau) \in \Gamma}{\Gamma \vdash A : \tau}$	$\frac{}{\Gamma \vdash \mathbf{0} : 0}$	$\frac{\Gamma \vdash P : \tau, \Gamma \vdash Q : \theta}{\Gamma \vdash P \mid Q : \tau + \theta}$	$\frac{\Gamma \vdash P : \tau}{\Gamma \vdash (\nu n)P : \tau}$
(Output)	(Input)	(In/Out)	
$\frac{}{\Gamma \vdash \langle n \rangle : 1}$	$\frac{\Gamma \vdash P : \tau}{\Gamma \vdash (n).P : \max(\tau - 1, 1)}$	$\frac{\Gamma \vdash P : \tau, \text{cap} \in \{\text{in}, \text{out}\}}{\Gamma \vdash \text{cap } n.P : \max(\tau, 1)}$	
(Amb)	(Open)	(Fix)	
$\frac{\Gamma \vdash P : \tau}{\Gamma \vdash n[P] : \tau + 1}$	$\frac{\Gamma \vdash P : \tau}{\Gamma \vdash \text{open } n.P : \max(\tau - 1, 1)}$	$\frac{\Gamma \cup \{(A, \tau)\} \vdash P : \theta, \theta \leq \tau}{\Gamma \vdash (\text{fix } A=P) : \tau}$	

The basic idea of the type system is to bound the number of active outputs and ambients in all processes reachable from a given one. In the rules (Input) and (Open) the process P is guarded and thus not active. These rules express that P may become active only after dissolving some active output or ambient. The function $\max(\cdot, 1)$ is used to avoid negative types. Without it some processes of unbounded (or even infinite) size like $(\text{fix } A=\text{open } n.\mathbf{0} \mid m[] \mid A)$ could be still typable. We take maximum with 1 and not 0 to obtain a property—used in some proofs—that $\mathbf{0}$ is the only process of type 0.

Example 2. The proof tree stating that the type judgment $\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) \mid n[\mathbf{0}] : 2$ holds is given below. In a similar way we can build a proof tree for $\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) \mid n[\mathbf{0}] : 3$.

$\frac{(A, 1) \in \{(A, 1)\}}{\{(A, 1)\} \vdash A : 1}$
$\frac{\{(A, 1)\} \vdash A : 1}{\{(A, 1)\} \vdash m[A] : 2}$
$\frac{\{(A, 1)\} \vdash m[A] : 2}{\{(A, 1)\} \vdash \text{open } n.m[A] : 1, \quad 1 \leq 1}$
$\frac{\{(A, 1)\} \vdash \text{open } n.m[A] : 1, \quad 1 \leq 1}{\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) : 1}$
$\frac{\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) : 1 \quad \emptyset \vdash \mathbf{0} : 0}{\emptyset \vdash n[\mathbf{0}] : 1}$
$\frac{\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) : 1 \quad \emptyset \vdash n[\mathbf{0}] : 1}{\emptyset \vdash (\text{fix } A=\text{open } n.m[A]) \mid n[\mathbf{0}] : 2}$

Recall Example 1. The process $(\text{fix } A=n[\text{open } m.A]) \mid (\text{fix } B=m[\text{in } n.B])$ is not typable because $(\text{fix } B=m[\text{in } n.B])$ is not typable.

Definition 3. A type environment Γ well-types a process P if there exists some natural number τ_P such that the type judgment $\Gamma \vdash P : \tau_P$ holds. A process P is typable if there exists a type environment Γ that well-types P .

If P is typable and does not contain free identifiers, then \emptyset well-types P .

We say that a process P is *balanced* if the number of occurrences of ambients and outputs in P is equal to the number of occurrences of open capabilities and inputs in P . We say that a recursive process $(\text{fix } A=P)$ has a balanced type if for every type environment $\Gamma \cup \{(A, \tau)\}$ such that $\Gamma \cup \{(A, \tau)\} \vdash P : \theta$ and $\theta \leq \tau$ we have $\theta = \tau$. In most natural examples, like in $(\text{fix } A=n[\text{open } n.A])$, if P is balanced then $(\text{fix } A=P)$ is typable. But there are exceptions like $(\text{fix } A=A \mid \text{open } n.n[\mathbf{0}])$ which is balanced but not typable. If P is not balanced because it contains more outputs and ambients than inputs and opens (like in $(\text{fix } A=n[A])$) then $(\text{fix } A=P)$ is not typable. Finally, note that not all

typable processes have balanced types. For example $(\text{fix } A = \text{open } n.A)$ is typable but it does not have a balanced type (the environment $\{(A, 2)\}$ being a counter-example).

In most examples of typable recursive processes that are considered in this paper we will want the types to be balanced. This is because if such a process does not have a balanced type then in each execution it consumes some (strictly more than it creates) messages or ambients in the global context in which it is placed, and thus it can be executed only finitely many times.

For a given type environment Γ and a given process P there may be many natural numbers τ such that the type judgment $\Gamma \vdash P : \tau$ holds. For example, $\emptyset \vdash (\text{fix } A = A) : \tau$ holds for any natural τ . However, since every set of natural numbers has a least element, we may define a least type.

Definition 4 (Least Type). *For any process P and any type environment Γ that well-types P , the least type of P with respect to Γ , denoted $\mathcal{L}^{\text{FC}}(P, \Gamma)$, is the least natural number τ such that $\Gamma \vdash P : \tau$ holds.*

Proposition 1 (Type Stability). *Let P and P' be typable processes and let Γ be a type environment that well-types P and P' . Then $P \equiv P'$ implies $\mathcal{L}^{\text{FC}}(P, \Gamma) = \mathcal{L}^{\text{FC}}(P', \Gamma)$.*

Proof. The proof goes by induction over the proof tree for $P \equiv P'$. The only difficult case is for the axiom $(\text{fix } A = Q) \equiv Q\{A \leftarrow (\text{fix } A = Q)\}$ (Str Fix Rec) which requires an induction over the structure of Q . \square

Proposition 2 (Type Checking - Type Inference). *Type checking (that is, deciding, given Γ, P and τ , whether the type judgment $\Gamma \vdash P : \tau$ holds) is decidable. For any process P and type environment Γ , we can decide whether Γ well-types P and compute $\mathcal{L}^{\text{FC}}(P, \Gamma)$.*

Proof. Both type checking and type inference amount to solving easy systems of inequalities with addition, subtraction of a constant, and max as the only arithmetic operations. \square

Proposition 3 (Subject Reduction). *Let P be a process and Γ an environment that well-types P . Then for all processes P' such that $P \rightarrow P'$, Γ well-types P' and $\mathcal{L}^{\text{FC}}(P', \Gamma) \leq \mathcal{L}^{\text{FC}}(P, \Gamma)$.*

Sketch of proof. The reductions (Red In) and (Red Out) do not change the type of a process at all. The reductions (Red Open) and (Red I/O) reducing processes of the form $\text{open } n.Q$ or $(n).Q$, respectively, do not change the type if $\mathcal{L}^{\text{FC}}(Q, \Gamma) > 1$. Otherwise, they strictly decrease the type by removing the ambient n or consuming a message. For the other reductions (Red Par), (Red Amb), (Red Res), and (Red \equiv), it follows from induction over P , using Proposition 1 in case of (Red \equiv). \square

Due to Proposition 1, the least type of any process congruent to $\mathbf{0}$ is 0. Conversely, we have:

Proposition 4. *For all closed and typable processes P , if $\mathcal{L}^{\text{FC}}(P, \emptyset) = 0$ then $P \equiv \mathbf{0}$.*

Proof. It is easy to see that if a closed process P contains either an ambient construct, a capability, an input or an output, then its least type is greater or equal to one. Therefore, it is enough to show that closed and well-formed processes built up with identifiers, $\mathbf{0}$, parallel composition, name restriction, and fix are congruent to $\mathbf{0}$. The proof goes by induction on the structure of P . \square

Additionally, we can prove some other properties. All recursion-free processes are typable. The encoding of replication $!P$ given earlier is not typable for any P non-congruent to $\mathbf{0}$. Processes built without parallel composition and ambient construct are typable. This last property implies that processes are typable if they satisfy the syntactic restriction—to forbid both composition and ambients within recursion—considered in Section 3. As we see in the next section, it follows that processes obeying this syntactic restriction are finite-state.

3.2 Typability and Finite-Control

The goal of this section is to prove that for a typable and closed process, there exist finitely many \equiv -congruence classes $\mathcal{K}_1, \dots, \mathcal{K}_n$ such that for all processes P' with $P \rightarrow^* P'$, there exists i for which $P' \in \mathcal{K}_i$. Instead of proving this directly, we show that for any typable and closed process P and any process P' reachable from P , there exists a *representative* P'' of P' (that is, $P' \equiv P''$) such that:

- the size of P'' is bounded and depends only on P ;
- the set of free names of P'' is a subset of the free names of P .

Here, by the size $|P|$ of a process P we mean the number of nodes in the tree representation of P . The two statements above imply that there exist only finitely many pairwise non-congruent processes reachable from P . Simply showing the size is bounded is insufficient as there are infinitely many different names. For example, processes from the set $\{n[0] \mid n \text{ being a name}\}$ have a bounded size, but being non-congruent with each other, they represent infinitely many \equiv -congruence classes.

The second requirement about free names is straightforward and actually does not rely on typability.

Proposition 5. *For all processes P, P' , if $P \equiv P'$ or $P \rightarrow P'$ then $fn(P') \subseteq fn(P)$.*

The first requirement is much more involved for various reasons. We need to characterize representatives of structural congruence classes of reachable processes; this requires to consider a process split into several parts.

First, we define pre-normalized processes. Let a process P be *pre-normalized* if it takes the form $(\nu n_1) \dots (\nu n_k)Q$ and, (i) every n_j occurs free in Q , (ii) n_1, \dots, n_k are pairwise distinct, and (iii) any other name restriction occurring in Q appears in the scope of some input or of some action prefix. Intuitively, pre-normalization is rewriting a process using the scope extrusion rules (Str Res Par) and (Str Res Amb) to a kind of prenex normal form.

The second part of processes consists in *outermost guarded* subprocesses. A process P is *guarded* if either $P \equiv \mathbf{0}$, $P \equiv \langle M \rangle$, $P \equiv \alpha.P'$ for some P' and some α , $P \equiv (x).P'$ for some P' , or recursively $P \equiv (\nu n)Q$ for some guarded Q . This property is clearly stable with respect to structural congruence, that is, if P is guarded and $P \equiv P'$ then P' is guarded as well. Let a subprocess P' be *outermost guarded* in a process P if P' is guarded and for any subprocess P'' of P enclosing P' , P'' is not guarded. For instance, $\text{out } m.\mathbf{0}$ is outermost guarded in $n[\text{out } m.\mathbf{0}]$ and as a consequence, $\mathbf{0}$ is guarded but not outermost guarded. In $n[\text{out } m.\mathbf{0} \mid \mathbf{0}]$, $\text{out } m.\mathbf{0} \mid \mathbf{0}$ is outermost guarded

(because $\text{out } m.0 \mid 0$ is congruent to a process of the form $\alpha.P$, namely $\text{out } m.0$) and thus, $\text{out } m.0$ is not outermost guarded. This last example shows that outermost guardedness is a pure syntactic condition and is not stable with respect to structural congruence.

Finally, the remaining part of the process is captured by a context. A *context* \mathcal{C} with l holes (or, for short, an l -context) is a process where exactly l subprocesses have been replaced by a hole \star_i occurring exactly once in \mathcal{C} . We write $\mathcal{C}[P_1, \dots, P_l]$ for the process obtained by filling each hole \star_i in \mathcal{C} with P_i .

A context is *active* if it consists only of holes, ambients, parallel compositions, and void processes 0 and furthermore if each process 0 occurs as a child node of an ambient in the tree representation of the context.

A process is *normalized* if this process is either 0 or a pre-normalized process of the form $(\nu n_1) \dots (\nu n_k)Q$, where Q is of the form $\mathcal{C}[P_1, \dots, P_l]$ such that:

- \mathcal{C} is an active l -context,
- P_1, \dots, P_l are the outermost guarded subprocesses from Q that are not congruent to 0 .

By the *one-step unfolding* of a process P we mean the process obtained from P by replacing every subprocess of the form $(\text{fix } A=Q)$ by $Q\{A \leftarrow (\text{fix } A=Q)\}$. If Q is obtained by one-step unfolding from P then $|Q| \leq |P|^2$.

Lemma 1. *Any typable and closed process P admits a congruent normalized process Q such that $|Q| \leq |P|^2$.*

Sketch of proof. First, by structural induction we prove that in any recursive process $(\text{fix } A=Q)$ either A is guarded in Q or $(\text{fix } A=Q)$ is congruent to 0 . We obtain a normalized version of a pre-normalized process by replacing all recursive definitions congruent to 0 by 0 , applying a one-step unfolding to the result, and then removing all 0 's from the context that are not child nodes of an ambient. \square

We say that a process Q is a *subprocess up to renaming* of a process P if Q can be obtained from some subprocess of P by renaming its free names.

Proposition 6. *Let P be a closed, typable, normalized and non-congruent to 0 process and P' be its one-step unfolding. Then for all processes Q reachable from P (that is, such that $P \rightarrow^* Q$), there exists a normalized process $(\nu n_1) \dots (\nu n_k)\mathcal{C}[P_1, \dots, P_l]$ structurally congruent to Q and such that*

- k is bounded by the size of \mathcal{C} ,
- the size of \mathcal{C} is bounded by $3 \cdot \mathcal{L}^{\text{FC}}(P, \emptyset)$,
- each Q_j is a subprocess up to renaming of some outermost guarded part from P' .

Sketch of proof. Since $(\nu n_1) \dots (\nu n_k)\mathcal{C}[P_1, \dots, P_l]$ is pre-normalized, every restricted name from the set $\{n_1, \dots, n_k\}$ must occur freely in \mathcal{C} . Thus k is bounded by the size of \mathcal{C} .

Since the process is normalized, the subprocesses P_1, \dots, P_l are not congruent to 0 and thus have strictly positive types. \mathcal{C} is an active context, so its tree representation consists of four kinds of nodes:

- leaves representing a hole, whose number is smaller than $\mathcal{L}^{\text{FC}}(P_1, \emptyset) + \dots + \mathcal{L}^{\text{FC}}(P_l, \emptyset)$,
- leaves representing $\mathbf{0}$, whose number is smaller than the number of unary nodes,
- binary nodes representing parallel compositions whose number is smaller than the number of leaves,
- unary nodes representing ambients; the number of such nodes summed with $\mathcal{L}^{\text{FC}}(P_1, \emptyset) + \dots + \mathcal{L}^{\text{FC}}(P_l, \emptyset)$ gives $\mathcal{L}^{\text{FC}}(Q, \emptyset)$.

This together with the subject reduction theorem (Proposition 3) gives that the size of \mathcal{C} is bounded by $3 \cdot \mathcal{L}^{\text{FC}}(P, \emptyset)$.

Finally, the processes Q_j are either directly subprocesses up to renaming of the initial process P or of unfoldings of the recursive definitions, which are already unfolded in P' . This is because the only possibility (apart from using the structural congruence) to modify a process below a guard is to substitute some of its free names with other names coming from communication. \square

The following theorem is a direct corollary from Propositions 5 and 6.

Theorem 1 (Finite-State). *For any closed and typable process P , there exist only finitely many pairwise non-congruent processes reachable from P .*

4 Examples

The model checking algorithm from [5] is limited to replication-free processes. We want to have at least some restricted version of recursion that would help us in modelling mobile computations while keeping model checking decidable. This section gives examples of programs that are typable and that therefore, by Theorem 1, are finite-state.

4.1 Simple Examples with Infinite Behaviour

Probably the simplest possible example with infinite behaviour is $n[P_A] \mid m[]$ where P_A is the process $(\text{fix } A = \text{in } m. \text{out } m.A)$. It is typable with the type of P_A equal to 1 and the type of the whole process being 3. We have $n[P_A] \mid m[] \rightarrow m[n[\text{out } m.P_A]] \rightarrow n[P_A] \mid m[]$, which creates an infinite loop.

Another simple example is $P_A \mid P_B$ where P_A is $(\text{fix } A = a[\text{open } b.A])$ and P_B is $(\text{fix } B = \text{open } a.b[B])$. Here the least type of P_A is 2 and the least type of P_B is 1. One can see it as a simple synchronization mechanism—we will use such a mechanism later in the encoding of the (synchronous) finite-control π -calculus. We have $P_A \mid P_B \equiv a[\text{open } b.P_A] \mid \text{open } a.b[P_B] \rightarrow \text{open } b.P_A \mid b[P_B] \rightarrow P_A \mid P_B$.

A similar behaviour can be obtained from a simpler process $(\text{fix } A = \text{open } a.A) \mid (\text{fix } B = a[B])$, but we cannot use it since it is not typable.

Our last example in this section shows that we can obtain not only infinite computation paths, but also infinitely many syntactically different processes along these paths. Consider the process $P_A \mid P_B$ where P_A is $(\text{fix } A = (\nu a) \text{open } n. \text{open } m. (\langle a \rangle \mid a[A]))$ and P_B is $(\text{fix } B = n[m[(x). \text{open } x.B]])$. The process is typable with the least types of P_A and P_B being respectively 1 and 3. Here, in every iteration, the process P_A creates a new fresh name and sends it to P_B .

4.2 Objective Moves

The only iterative definition in the encoding of objective moves in [6] is $\text{allow } n \triangleq !\text{open } n$. This can be directly translated to $(\text{fix } A = \text{open } n.A)$, but such a translation leads to a definition of $\text{mv in } n.P$ where the type of $\text{mv in } n.P$ is one greater than the type of P , and so does not allow the use of objective moves inside recursion. Therefore we propose an alternative definition.

$$\begin{aligned}
 \text{allow } n &\triangleq (\text{fix } A = \text{open } n.n[A]) \\
 n^\downarrow[P] &\triangleq n[P \mid \text{allow in}] \\
 n^\downarrow[P] &\triangleq n[P] \mid \text{allow out} \\
 n^\downarrow \downarrow[P] &\triangleq n[P \mid \text{allow in}] \mid \text{allow out} \\
 \text{mv in } n.P &\triangleq (\nu k)k[\text{in } n.\text{in}[\text{out } k.\text{open } k.\text{open in}.P]] \\
 \text{mv out } n.P &\triangleq (\nu k)k[\text{out } n.\text{out}[\text{out } k.\text{open } k.\text{open out}.P]]
 \end{aligned}$$

It is easy to see that all these processes are typable and are balanced; the least type of $\text{allow } n$ is 1, the least type of $\text{mv in } n.P$ and $\text{mv out } n.P$ is the maximum of the type of P and 2. One can check that $n^\downarrow[Q] \mid \text{mv in } n.P \rightarrow^* n^\downarrow[P \mid Q]$ and $n^\downarrow[\text{mv out } n.P \mid Q] \rightarrow^* n^\downarrow[Q] \mid P$.

4.3 Firewalls

Consider the firewall from [6]. This is a replication-free process

$$\text{firewall} = (\nu w)k[\text{in } k.\text{in } w] \mid w[\text{open } k.P],$$

but it allows only one agent to enter the firewall. Let us first extend this example to allow for more agents. To avoid some confusion we replace one of the two occurrences of the name k with k' : $\text{firewall} = (\nu w)!k[\text{in } k'.\text{in } w] \mid w[!\text{open } k' \mid P]$. Then we have $k'[\text{open } k.Q] \mid \text{firewall} \rightarrow^* (\nu w)!k[\text{in } k'.\text{in } w] \mid w[!\text{open } k' \mid P \mid Q]$ and the firewall is still ready to allow more agents that are aware of the password (k, k') .

We still have a little problem with modelling this firewall as a typable program. The process $k[\text{in } k'.\text{in } w]$ is at the beginning outside the ambient w , but at the end (after the agent enters the firewall) it is inside w . In typable programs we need to always start a recursion in the same place that we end it. Therefore we first modify the firewall: $\text{firewall} = (\nu w)w[!k[\text{out } w.\text{in } k'.\text{in } w] \mid !\text{open } k' \mid P]$.

Now it is easy to see that this process behaves in the same way as the following program *firewall* where

$$\begin{aligned}
 \text{firewall} &= (\nu w)w[\text{hook} \mid \text{initiator} \mid P] \\
 \text{hook} &= (\text{fix } A = k[\text{out } w.\text{in } k.\text{in } w.\text{open } b.A]) \\
 \text{initiator} &= (\text{fix } B = \text{open } k.b[B])
 \end{aligned}$$

We use the additional ambient b to balance the bodies of the procedures *hook* and *initiator*.

4.4 Routable Packets

Following [6] we define *packet* pkt as an empty packet named pkt that can be routed repeatedly to various destinations. Contrary to [6], we do not model routing as communicating the path to be followed (we restricted the calculus not to contain communication of compound messages), but by sending it another ambient containing the path.

$$\begin{aligned} packet\ pkt &\triangleq pkt[(\text{fix } R = \text{open } route.route[R])] \\ route\ pkt\ \text{with } P\ \text{to } M &\triangleq route[\text{in } pkt.\text{open } route.M \mid P] \\ forward\ pkt\ \text{to } M &\triangleq route\ pkt\ \text{with } \mathbf{0}\ \text{to } M \end{aligned}$$

Then there is an execution $packet\ pkt \mid route\ pkt\ \text{with } P\ \text{to } M \rightarrow^* pkt[M \mid P \mid P_R]$ where P_R is the process $(\text{fix } R = \text{open } route.route[R])$. Similarly, $pkt[P \mid P_R] \mid forward\ pkt\ \text{to } M \rightarrow^* pkt[M \mid P \mid P_R]$.

4.5 A Finite-Control π -Calculus

Here we encode a version of the finite-control π -calculus [12] without name passing in recursive procedures (that is, with parameterless recursive definitions) and without non-deterministic choice. An encoding of the full finite-control π -calculus seems possible using the extensions of our calculus discussed in Section 6.

Processes of the finite-control π -calculus

$P ::=$	process		
$(\nu n)P$	name restriction		
$T_1 \mid \dots \mid T_k$	parallel threads		
$T ::=$	thread		
$\mathbf{0}$	inactivity	$(\nu n)T$	name restriction
$\alpha.T$	action	A	identifier
$(\text{fix } A = T)$	recursion		
$\alpha ::=$	action		
$n(x)$	input on channel n		
$n\langle x \rangle$	output on channel n		

The encoding of the (asynchronous) π -calculus given in [6] cannot be used here for at least two reasons. First, the finite-control π -calculus uses synchronous communication while the communication in the ambient calculus is asynchronous. In order to simulate synchronous communication we have to run a synchronisation protocol. Second, dynamic generation of new channels strictly increases the size of the encoding and thus cannot be typable. Instead of this, we create new channels for every single communication and we destroy this channel immediately after the communication is finished.

To synchronize the communication, for every thread T_i we introduce an ambient $sync_i[]$ that avoids mixing the order of actions taken by this thread: every thread can send or receive at most one message at a time. Additionally we introduce one ambient $lock[]$ that allows processing only one communication at a time. These ambients are

present at the beginning, but they disappear (that is, they get opened) when the respective action starts and they reappear when the action is finished.

In the encoding given below we use an ambient named ch as a place where communication happens. The idea of this encoding is as follows. If two processes $n\langle M \rangle.P$ and $n(x).Q$ in threads T_i and T_j are willing to communicate, they start by opening the respective ambients $sync_i$ and $sync_j$ (if some of these ambients are not present, it means that the thread is busy with some other action, and the process has to wait). Then the output process leaves an ambient $n[]$ (this is the information that there is a message sent over the channel n) and moves inside the ambient ch . There it sends the message $\langle M \rangle$ within another ambient n . The input process opens the ambient $n[]$ (if there is no such ambient, it means that there is no message sent over channel n and the process has to wait), then it opens $lock[]$ (again, if there is no $lock[]$ ambient, it means that there is another communication just taking place and the process has to wait until it is finished) and goes inside ch and inside n where it reads the message M . The rest of the encoding is just to clean up afterwards: both processes go out of ch and together with the auxiliary processes $Sync_i$ and $Sync_j$ they synchronize the two threads and release the lock on communication, and remove all auxiliary ambients used in the meantime (more precisely, after the communication the ambient b gets opened, $done_i$ goes out of n and n goes inside $done_i$ where it gets opened; then $done_i$ moves outside ch , c gets opened, and the two $Sync$ processes open $done_i$ and $done_j$; at this moment the ambients $sync_i[], sync_j[],$ and $lock[]$ appear again at the top level). The ambients a, b, c are used to balance the process and to move inside ch .

Formally, the encoding is defined by the function $\llbracket \cdot \rrbracket$ from processes of the finite-control π -calculus to the finite-control ambient calculus. Except from communication, the encoding is quite straightforward: we have $\llbracket (\nu n)P \rrbracket \triangleq (\nu n)\llbracket P \rrbracket$,

$$\begin{aligned} \llbracket T_1 \mid \dots \mid T_k \rrbracket &\triangleq \llbracket T_1 \rrbracket_1 \mid Sync_1 \mid \dots \mid \llbracket T_k \rrbracket_k \mid Sync_k \mid ch^{\uparrow}[] \mid lock[] \\ \llbracket 0 \rrbracket_i &\triangleq 0, \llbracket (\nu n)T \rrbracket_i \triangleq (\nu n)\llbracket T \rrbracket_i, \llbracket A \rrbracket_i \triangleq A, \llbracket (\text{fix } A=T) \rrbracket_i \triangleq (\text{fix } A=\llbracket T \rrbracket_i), \\ \llbracket n\langle M \rangle.P \rrbracket_i &\triangleq \text{open } sync_i.(n[] \mid \text{mv in } ch.n[\langle M \rangle \mid \text{open } a.PostOut_i(P)]) \\ \llbracket n(x).Q \rrbracket_j &\triangleq \text{open } sync_j.\text{open } n.\text{open } lock.\text{mv in } ch.a[\text{in } n.(x).PostIn_j(Q)] \end{aligned}$$

where

$$\begin{aligned} Sync_i &\triangleq sync_i[] \mid (\text{fix } S_i=\text{open } done_i.(sync_i[] \mid S_i)) \\ PostOut_i(P) &\triangleq \text{open } b.(\text{in } done_i.lock[] \mid done_i[\text{out } n.\text{open } n.\text{out } ch.\text{open } c.(\llbracket P \rrbracket_i)]) \\ PostIn_j(Q) &\triangleq b[c[done_j[\llbracket Q \rrbracket_j]]] \end{aligned}$$

5 Ambient Logic and Model Checking

To reason about distributed and mobile computations programmed in the ambient calculus, Cardelli and Gordon [5] introduce a modal logic that apart from standard temporal modalities for describing the evolution of processes includes novel spatial modalities for describing the tree structure of ambient processes. In a recent paper, Cardelli and Gordon extend the logic with the constructs for describing private names [7].

The *model checking* problem is to decide whether a given object (in our case, an ambient process) satisfies (that is, is a model of) a given formula. Cardelli and Gordon [5]

give a model checking algorithm for the fragment of the calculus in which the processes contain no replications and no dynamic name generation against a fragment of the logic in which formulas contain no composition adjunct. It was then proved in [10] that model checking this fragment of the calculus against this fragment of the logic is PSPACE-complete. Recently, in [11] it has been shown that on the one hand, extending the calculus with name restriction and the logic with corresponding logical operators is harmless for the complexity of model checking—it remains PSPACE—and on the other hand that either considering replication in the calculus or composition adjunct in the logic makes the model checking problem undecidable.

5.1 Ambient Logic

We recall in this section definitions concerning this logic (omitting the composition adjunct).

In addition to the reduction relation and the structural congruence, we introduce an additional relation called *location* and denoted \downarrow to reason about the shape of ambients (that is, space). The location relation is defined as $P \downarrow Q$ if there exists Q', n such that $P \equiv n[Q] \mid Q'$. We write \downarrow^* for the reflexive and transitive closure of \downarrow .

We describe the syntax of the ambient logic and its satisfaction relation in the following tables.

Logical Formulas:

η	a name n or a variable x		
$\mathcal{A}, \mathcal{B} ::=$	formula	$\eta[\mathcal{A}]$	location
T	true	$\mathcal{A}@\eta$	location adjunct
$\neg \mathcal{A}$	negation	$\eta\textcircled{\mathcal{A}}$	revelation
$\mathcal{A} \vee \mathcal{B}$	disjunction	$\mathcal{A} \odot \eta$	revelation adjunct
0	void	$\Diamond \mathcal{A}$	sometime modality
$\mathcal{A} \mid \mathcal{B}$	composition match	$\Downarrow \mathcal{A}$	somewhere modality
		$\exists x.\mathcal{A}$	existential quantification

We assume that names and variables belong to two disjoint vocabularies. We write $\mathcal{A}\{x \leftarrow m\}$ for the outcome of substituting each free occurrence of the variable x in the formula \mathcal{A} with the name m . We say a formula \mathcal{A} is closed if and only if it has no free variables (though it may contain free names).

The *satisfaction relation* $P \models \mathcal{A}$ provides the semantics of our logic. It is stable with respect to structural congruence, that is, if $P \models \mathcal{A}$ and $P \equiv P'$ then $P' \models \mathcal{A}$.

Satisfaction $P \models \mathcal{A}$ (for \mathcal{A} closed):

$P \models \mathbf{T}$	$P \models n[\mathcal{A}] \triangleq \exists P'. P \equiv n[P'] \wedge P' \models \mathcal{A}$
$P \models \neg \mathcal{A} \triangleq \neg(P \models \mathcal{A})$	$P \models \mathcal{A}@\eta \triangleq n[P] \models \mathcal{A}$
$P \models \mathcal{A} \vee \mathcal{B} \triangleq P \models \mathcal{A} \vee P \models \mathcal{B}$	$P \models n\textcircled{\mathcal{A}} \triangleq \exists P'. P \equiv (\nu n)P' \wedge P' \models \mathcal{A}$
$P \models \mathbf{0} \triangleq P \equiv \mathbf{0}$	$P \models \mathcal{A} \odot n \triangleq (\nu n)P \models \mathcal{A}$
$P \models \mathcal{A} \mid \mathcal{B} \triangleq \exists P', P''. P \equiv P' \mid P'' \wedge P' \models \mathcal{A} \wedge P'' \models \mathcal{B}$	$P \models \Diamond \mathcal{A} \triangleq \exists P'. P \rightarrow^* P' \wedge P' \models \mathcal{A}$
	$P \models \Downarrow \mathcal{A} \triangleq \exists P'. P \downarrow^* P' \wedge P' \models \mathcal{A}$
	$P \models \exists x.\mathcal{A} \triangleq \exists m. P \models \mathcal{A}\{x \leftarrow m\}$

5.2 Model Checking Finite-Control Mobile Ambients

In this section we show how closed and typable processes can be model checked against formulas of the ambient logic. We assume that in any process bound identifiers are pairwise distinct and that bound names are pairwise distinct and different from free names (not only free names from the process itself, but also free names occurring in formulas).

We consider here normalized processes as introduced in Section 3.2. To single out name restrictions, we write a normalized process $(\nu n_1) \dots (\nu n_k) \mathcal{C}[P_1, \dots, P_l]$ as a pair $\langle \{n_1, \dots, n_k\}, \mathcal{C}[P_1, \dots, P_l] \rangle$, separating name restriction prefix from the rest of the process and considering these name restrictions as a set of names.¹

In a normalized process, only the active-context part is addressed by spatial modalities from the logic, that is, the modalities $\mathcal{A} \mid \mathcal{B}$, $n[\mathcal{A}]$, $\diamond \mathcal{A}$ and $\mathbf{0}$. This allows us to control the size and the number of normalized processes considered for model checking these spatial modalities.

The following propositions (Propositions 7–10) express that in polynomial space we can test whether a process is congruent to $\mathbf{0}$, we can decompose it in all possible ways to a parallel composition of two other processes, we can remove the given leading ambient (if it exists), and we can compute all sublocations of the process. This is possible because it requires examining only the active context of the given process. The proof of these propositions is based on the following lemma.

Lemma 2 (Inversion). *Let P, Q, Q' be normalized processes.*

1. $(\nu n)P \equiv \mathbf{0}$ if and only if $P \equiv \mathbf{0}$.
2. If n and m are different names, then $(\nu n)P \equiv m[Q]$ if and only if there exists a normalized process R such that $P \equiv m[R]$ and $Q \equiv (\nu n)R$.
3. $(\nu n)P \equiv Q \mid Q'$ if and only if there exist normalized processes R, R' such that $P \equiv R \mid R'$ and either $Q \equiv (\nu n)R$ and $Q' \equiv R'$ and $n \notin \text{fn}(Q')$ or $Q \equiv R$ and $Q' \equiv (\nu n)R'$ and $n \notin \text{fn}(Q)$.

Proof. The proof is the same as for the replication-free fragment of the ambient calculus (Proposition 5.1 in [11]), observing that we have to examine only the active context of a normalized process. \square

Proposition 7. *For any normalized process $\langle N, \tilde{P} \rangle$, $\langle N, \tilde{P} \rangle \equiv \mathbf{0}$ if and only if $\tilde{P} \equiv \mathbf{0}$. Furthermore, we can test whether \tilde{P} is congruent to $\mathbf{0}$ in polynomial-time.*

Proposition 8. *For any normalized process $\langle N, \tilde{P} \rangle$, we can compute in polynomial space a finite set of pairs of normalized processes that we denote $\text{Decomp}(\langle N, \tilde{P} \rangle)$ and defined as $\{(\langle N_1, \tilde{Q}_1 \rangle, \langle N'_1, \tilde{R}_1 \rangle), \dots, (\langle N_p, \tilde{Q}_p \rangle, \langle N'_p, \tilde{R}_p \rangle)\}$ satisfying:*

- for all Q, R satisfying $\langle N, \tilde{P} \rangle \equiv Q \mid R$, there exists i such that $\langle N_i, \tilde{Q}_i \rangle \equiv Q$ and $\langle N'_i, \tilde{R}_i \rangle \equiv R$.
- for all i in $1 \dots p$, $N_i \cup N'_i = N$ and $N_i \cap N'_i = \emptyset$, $\text{fn}(\tilde{Q}_i) \cap N'_i = \emptyset$ and $\text{fn}(\tilde{R}_i) \cap N_i = \emptyset$ and $|\tilde{Q}_i|, |\tilde{R}_i| \leq |\tilde{P}|$.

¹ Whenever two different normalized processes P and P' have the same pair-representation, then $P \equiv P'$ by the axiom (Str Res Res) from the structural congruence.

Proposition 9. *For any normalized process $\langle N, \tilde{P} \rangle$ and any name n , we can test in polynomial time if there exists Q such that $\langle N, \tilde{P} \rangle \equiv n[Q]$. Moreover, if such a Q exists, then $n \notin N$ and we can compute in polynomial time the normalized version $\langle N, \tilde{Q} \rangle$ of Q such that $|\tilde{Q}| \leq |\tilde{P}|$.*

Combining the two previous propositions, we obtain:

Proposition 10. *For any normalized process $\langle N, \tilde{P} \rangle$, we can compute a finite set of normalized processes $\text{Sublocations}(\langle N, \tilde{P} \rangle) = \{\langle N_1, \tilde{Q}_1 \rangle, \dots, \langle N_p, \tilde{Q}_p \rangle\}$ such that (i) for all Q such that $\langle N, \tilde{P} \rangle \downarrow^* Q$, there exists i satisfying that $Q \equiv \langle N_i, \tilde{Q}_i \rangle$ and (ii) for all i in $1 \dots p$, $N_i \subseteq N$ and $|\tilde{Q}_i| \leq |\tilde{P}|$.*

Moreover, using results from Section 3.2, we obtain:

Proposition 11. *For any typable and normalized process $\langle N, \tilde{P} \rangle$, we can compute a finite set of normalized processes $\text{Reachable}(\langle N, \tilde{P} \rangle) = \{\langle N_1, \tilde{Q}_1 \rangle, \dots, \langle N_p, \tilde{Q}_p \rangle\}$ such that (i) for all Q such that $\langle N, \tilde{P} \rangle \rightarrow^* Q$, there exists i satisfying that $Q \equiv \langle N_i, \tilde{Q}_i \rangle$ and (ii) for all i in $1 \dots p$, $N_i \subseteq N$ and $|\tilde{Q}_i| \leq |\tilde{P}|^2$.*

The algorithm presented here is very close to the one given in [11].

Model Checking Algorithm: *Check($\langle N, \tilde{P} \rangle, \mathcal{A}$) where $N \cap \text{fn}(\mathcal{A}) = \emptyset$, by convention*

$$\begin{aligned}
 \text{Check}(\langle N, \tilde{P} \rangle, \mathbf{T}) &\triangleq \mathbf{T} \\
 \text{Check}(\langle N, \tilde{P} \rangle, \neg \mathcal{A}) &\triangleq \neg \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A} \vee \mathcal{B}) &\triangleq \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A}) \vee \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{B}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \mathbf{0}) &\triangleq \begin{cases} \mathbf{T} & \text{if } \tilde{P} \equiv \mathbf{0} \\ \mathbf{F} & \text{otherwise} \end{cases} \\
 \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A} \mid \mathcal{B}) &\triangleq \bigvee_{(P_1, P_2) \in \text{Decomp}(\langle N, \tilde{P} \rangle)} \text{Check}(P_1, \mathcal{A}) \wedge \text{Check}(P_2, \mathcal{B}) \\
 &\quad P_1, P_2 \text{ being respectively } \langle N_1, \tilde{P}_1 \rangle \text{ and } \langle N_2, \tilde{P}_2 \rangle \\
 \text{Check}(\langle N, \tilde{P} \rangle, n[\mathcal{A}]) &\triangleq \tilde{P} \equiv n[\tilde{Q}] \wedge \text{Check}(\langle N, \tilde{Q} \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A} @ n) &\triangleq \text{Check}(\langle N, n[\tilde{P}] \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, n @ \mathcal{A}) &\triangleq \bigvee_{m \in N} \text{Check}(\langle N - \{m\}, \tilde{P} \{m \leftarrow n\} \rangle, \mathcal{A}) \\
 &\quad \vee (n \notin \text{fn}(\tilde{P}) \wedge \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A})) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A} \odot n) &\triangleq \text{Check}(\langle N \cup \{n\}, \tilde{P} \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \diamond \mathcal{A}) &\triangleq \bigvee_{\langle N', \tilde{P}' \rangle \in \text{Reachable}(\langle N, \tilde{P} \rangle)} \text{Check}(\langle N', \tilde{P}' \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \spadesuit \mathcal{A}) &\triangleq \bigvee_{\langle N, \tilde{P}' \rangle \in \text{Sublocations}(\langle N, \tilde{P} \rangle)} \text{Check}(\langle N, \tilde{P}' \rangle, \mathcal{A}) \\
 \text{Check}(\langle N, \tilde{P} \rangle, \exists x. \mathcal{A}) &\triangleq \text{let } n_0 \notin N \cup \text{fn}(\tilde{P}) \cup \text{bn}(\tilde{P}) \cup \text{fn}(\mathcal{A}) \text{ be a fresh name in} \\
 &\quad \bigvee_{n \in \text{fn}(\langle N, \tilde{P} \rangle) \cup \text{fn}(\mathcal{A}) \cup \{n_0\}} \text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A} \{x \leftarrow n\})
 \end{aligned}$$

Theorem 2 (Correctness). *For all normalized and typable processes $\langle N, \tilde{P} \rangle$ and all closed formulas \mathcal{A} , we have $\langle N, \tilde{P} \rangle \models \mathcal{A}$ if and only if $\text{Check}(\langle N, \tilde{P} \rangle, \mathcal{A}) = \mathbf{T}$.*

Sketch of proof. The proof follows the lines of the proof of Theorem 5.1 in [11] and goes by induction on the formula \mathcal{A} . In the cases of \mathbf{T} , $\neg \mathcal{A}$, $\mathcal{A} \vee \mathcal{B}$, $\mathcal{A} @ n$, $\mathcal{A} \odot n$ the result follows directly from the definition of the satisfaction relation. In the case of $\mathbf{0}$ and $n[\mathcal{A}]$ it follows from Propositions 7 and 9, and in the case of $\mathcal{A} \mid \mathcal{B}$ from Proposition 8. The cases of $\spadesuit \mathcal{A}$ and $\diamond \mathcal{A}$ follow from Propositions 10 and 11. The case of $\exists x. \mathcal{A}$ follows

the lines of the proof of Proposition 4.11 in [4]. Finally, the case of $n\textcircled{\mathcal{A}}$ reflects the two possibilities that either n is one of the bounded names occurring in the process or it does not occur there (in the latter case observe that for all processes Q , $n \notin \text{fn}(Q)$ implies $(\nu n)Q \equiv Q$). \square

Theorem 3 (PSPACE-complete). *The model checking problem for finite-control processes against the ambient logic is decidable. Moreover, it is PSPACE-complete.*

Sketch of proof. Decidability follows from Theorem 2. One obtains the PSPACE upper bound by implementing disjunction in polynomial space, as is done in [10]. The PSPACE lower bound is proved in [10]. \square

6 Extensions of the System

In this section we discuss some extensions that are possible to the calculus without affecting decidability or complexity of the model checking problem. We did not introduce these extensions before because we want to keep our formal presentation of a finite-control ambient calculus as simple as possible.

Parameters in recursive definitions. In the system we defined the identifiers used in recursive definitions do not carry any name parameters. It is however quite straightforward to allow definitions of the form $(\text{fix } A(x) = P[A(y)])$ — one has to clearly distinguish between definitions (λ -abstractions) and calls (λ -application) of such functions and then respectively handle the renaming of parameters.

Nondeterministic choice. In the ambient calculus one may encode an internal nondeterministic choice $P + Q$ (see [6] for an encoding of an external choice) as the process $(\nu n)(n[0] \mid \text{open } n.P \mid \text{open } n.Q)$.

Then reducing $P + Q$ leads to either $P \mid (\nu n)\text{open } n.Q$ or to $Q \mid (\nu n)\text{open } n.P$. As $(\nu n)\text{open } n.R$ is bisimilar to $\mathbf{0}$, this is a good approximation of nondeterministic choice. However, even this simple encoding is no longer possible in recursive processes of the finite-control fragment, since it goes beyond well-formed processes. But even if we ignore the well-formedness restriction, such an encoding does not work because $(\nu n)\text{open } n.R$ is not congruent (it is only bisimilar) to $\mathbf{0}$, which means that its type must be strictly positive, so it is not possible to balance the type of $P + Q$ if P and Q are balanced (and thus $P + Q$ is not typable). In the encoding of the finite-control π -calculus in Section 4.5 all recursive processes have balanced types and thus to extend the encoding to accommodate nondeterministic choice we need a balanced encoding of choice.

A possible solution is to add nondeterministic choice as a primitive construct in the calculus. To do so, we need to relax the definition of well-formed processes from one occurrence of identifier in a recursive definition to one occurrence per option of a nondeterministic choice. The reduction rules for processes can be then extended in a straightforward way, and an appropriate typing rule is:

$$\begin{array}{c} \text{(Choice)} \\ \frac{\Gamma \vdash P : \tau, \Gamma \vdash Q : \theta}{\Gamma \vdash P + Q : \max(\tau, \theta)} \end{array}$$

Replication-free fragment of the ambient calculus. Our initial motivation was to find a fragment of the ambient calculus that extends the replication-free fragment (for which the decidability and complexity of the model checking problem was known [5,10]) to allow some infinite computation, while retaining a decidable model checking problem. The calculus of this paper does not extend the replication-free fragment because it does not allow for sending capabilities inside messages. It is however quite obvious that a typable finite-control process can be put in a replication-free context without any change to the model checking algorithm. The only subtle point is that if one wants to achieve a PSPACE algorithm one should apply the data structure from [10] only to the replication-free context; otherwise storing an explicit substitution for every communication might lead (in the case of recursive communications) to infinitely growing substitutions.

Sending capabilities in communication. In the current version of the calculus we allow for sending only names. The extension to sending single capabilities is however not difficult. The problem with sending single capabilities is probably best shown in the following process: $\langle \text{in } n \rangle \mid (\text{fix } A = (x).(\langle \text{in } x \rangle \mid A))$. After the first iteration the process sends the message in $(\text{in } n)$, then in $(\text{in } (\text{in } n))$, and so on, growing infinitely. Probably the simplest solution is to observe that in x cannot be executed if x is not a name, so it is enough to introduce a special deadlock capability and replace these complex capabilities with the deadlock capability.

An intriguing alternative possibility to solve the problem would be to combine the calculus with the type system of [8], where in $(\text{in } n)$ cannot be well-typed.

Sending sequences of capabilities. In the original definition of the ambient calculus [6] it is possible to send not only single capabilities but also sequences of capabilities. We do not see a very easy solution to this problem. Consider as an example the following process: $\langle \text{in } n \rangle \mid (\text{fix } A = (x).(\langle x.x \rangle \mid A))$. Here, after k iterations we obtain a sequence of capabilities in $n \dots \text{in } n$ of length 2^k . Thus the process grows infinitely. A possible solution is to distinguish between simple communications (sending a name or a single capability) and complex communications (sending a sequence of length ≥ 2), and to give a simple output type 1 (as in the current version of the system) while a complex output is typed 2 (that is, to introduce an additional typing rule $\Gamma \vdash \langle M \rangle : 2$ for complex M). Then the above process is not typable, but for example the process $(\text{fix } A = (x).(y).(\langle x.y \rangle \mid A))$ is typable. The decidability of model checking relies then on the observation that such a process cannot be executed infinitely many times (roughly, it can be executed as many times as many outputs are present in the context around).

Again, to achieve PSPACE complexity one has to be careful about substitutions—one should apply the substitution in the case of simple communication but one should store the substitution in the data structure (as is done in [10]) in the case of complex communication.

7 Conclusion

Previous work on model checking the spatial and temporal logic of the ambient calculus is limited to processes lacking any form of recursion or iteration. This work shows the possibility of model checking a richer, more expressive class of mobile behaviours. We hope it will lead to the discovery of further applications of the ambient logic.

References

1. R.M. Amadio and Ch. Meyssonnier. On the decidability of fragments of the asynchronous pi-calculus. In *Electronic Notes in Theoretical Computer Science, Proceedings EXPRESS 2001*, 2001. 299
2. M. Bugliesi and G. Castagna. Secure safe ambients. In *28th ACM Symposium on Principles of Programming Languages (POPL'01)*, pages 222–235, 2001. 295
3. L. Cardelli and A. D. Gordon. Equational properties of mobile ambients. In *Proceedings FoSSaCS'99*, volume 1578 of *LNCS*, pages 212–226. Springer, 1999. 295
4. L. Cardelli and A. D. Gordon. Modal logics for mobile ambients: Semantic reasoning. Unpublished annex to [5], 1999. 311
5. L. Cardelli and A. D. Gordon. Anytime, anywhere: Modal logics for mobile ambients. In *Proceedings POPL'00*, pages 365–377. ACM, January 2000. 295, 296, 298, 304, 307, 307, 312, 313
6. L. Cardelli and A. D. Gordon. Mobile ambients. *Theoretical Computer Science*, 240(1):177–213, 2000. 295, 296, 298, 305, 305, 306, 306, 311, 312
7. L. Cardelli and A. D. Gordon. Logical properties of name restriction. In *Proceedings of the 5th International Conference on Typed Lambda Calculi and Applications (TLCA'01)*, volume 2044 of *LNCS*, pages 46–60. Springer, 2001. 295, 307
8. L. Cardelli and A.D. Gordon. Types for mobile ambients. In *26th ACM Symposium on Principles of Programming Languages (POPL'99)*, pages 79–92, 1999. 295, 312
9. S. Chaki, S.K. Rajamani, and J.Rehof. Types as models: Model checking message-passing programs. In *29th ACM Symposium on Principles of Programming Languages (POPL'02)*, 2002. To appear. 295
10. W. Charatonik, S. Dal Zilio, A. D. Gordon, S. Mukhopadhyay, and J.-M. Talbot. The complexity of model checking mobile ambients. In *Proceedings FoSSaCS'01*, volume 2030 of *LNCS*, pages 152–167. Springer, 2001. An extended version appears as Technical Report MSR–TR–2001–03, Microsoft Research, 2001. 296, 308, 311, 311, 312, 312, 312
11. W. Charatonik and J.-M. Talbot. The decidability of model checking mobile ambients. In *Proceedings of the 15th Annual Conference of the European Association for Computer Science Logic*, volume 2142 of *LNCS*, pages 339–354. Springer, 2001. 296, 296, 299, 308, 309, 310, 310
12. M. Dam. Model checking mobile processes. *Information and Computation*, 121(1):35–51, 1996. 295, 298, 306
13. R.R. Hansen, J.G. Jensen, F. Nielson, and H. Riis Nielson. Abstract interpretation of mobile ambients. In *Static Analysis (SAS'99)*, volume 1694 of *Lecture Notes in Computer Science*, pages 134–148. Springer, 1999. 295
14. F. Levi and D. Sangiorgi. Controlling interference in ambients. In *27th ACM Symposium on Principles of Programming Languages (POPL'00)*, pages 352–364, 2000. 295
15. R. Milner. *Communicating and Mobile Systems: the π -Calculus*. Cambridge University Press, 1999. 295, 298
16. F. Nielson, H. Riis Nielson, R.R. Hansen, and J.G. Jensen. Validating firewalls in mobile ambients. In *Concurrency Theory (Concur'99)*, volume 1664 of *Lecture Notes in Computer Science*, pages 463–477. Springer, 1999. 295
17. D. Sangiorgi. Extensionality and intensionality of the ambient logics. In *Proceedings POPL'01*, pages 4–13. ACM, January 2001. 298