An Asynchronous, Distributed Implementation of Mobile Ambients

Cédric Fournet¹, Jean-Jacques Lévy², and Alan Schmitt²

¹ Microsoft Research

² INRIA Rocquencourt

Abstract We present a first distributed implementation of the Cardelli-Gordon's ambient calculus. We use Jocaml as an implementation language and we present a formal translation of Ambients into the distributed join calculus, the process calculus associated with Jocaml. We prove the correctness of the translation.

1 Introduction

We present a highly concurrent distributed implementation of the Cardelli-Gordon's calculus of Mobile Ambients [4] in Jocaml [13,6]. The ambient calculus is a simple and very esthetic model for distributed mobile computing. However, until now, it did not have a distributed implementation. Such an implementation may seem easy to build, especially with a language with distribution and strong migration (Jocaml), but we encountered several difficulties and design choices.

Ambients are nested. Their dynamics is defined by three atomic steps: an ambient may move into a sibling ambient (IN), it may move out of its parent ambient (OUT), or it may open one of its child ambients (OPEN). Each atomic migration step may involve several ambients, possibly on different sites. For instance, the source and destination ambients participate to an IN-step; similarly the source and parent ambients take part to an OUT-step; the target ambient participates to an OPEN-step. Each atomic step of the ambients calculus can be decomposed in two parts: checking whether the local structure of the ambient tree enables the step, and actually performing the migration.

The first part imposes some distributed synchronization. One may use a global synchronous primitive existing at the operating system or networking level, but such a solution is unrealistic in large-scale networks. A first range of solutions can be designed by considering locks and critical sections in order to serialize the implementation of atomic steps. For instance, the two ambients participating to a reduction step are temporary locked, and the locks are released at the end of the step. However this solution cannot be symmetric, in the same way as there is no symmetric distributed solution to the Dining Philosophers problem. Some ambients have to be distinguished, for instance, one ambient could be the synchronizer of all ambients. Naturally, the nested structure of ambients can be used, for instance each ambient controls the synchronization of its direct subambients. In both cases, one has to be careful to avoid deadlocks or

too much serialization. This solution would be similar to Cardelli’s centralized implementation of an earlier variant of the ambient calculus in Java [1,2]. One advantage of a serialized solution is the ease of the correctness proof of the implementation. On the negative side, each attempt to perform a step takes several locks higher up in the ambient hierarchy; these locks may be located at remote sites, leading to long delays before these locks are released for other local steps. Moreover, due to the mobility discipline of the ambient calculus, an ambient that migrates from one point to another in the ambient hierarchy has to travel through an ambient enclosing both the origin and the destination, thus inducing global bottlenecks.

A different set of solutions is fully asynchronous. Atomic steps of ambients are decomposed into several elementary steps, each involving only local synchronization. In this approach, each ambient step is implemented as a run of a protocol involving several messages. Concurrency is higher, as only the moving ambient might not be available for other reduction steps. For instance, our solution never blocks steps involving parents of a moving ambient. The implementation of migration towards a mobile target may be problematic, but can be handled independently of the implementation of ambient synchronization, e.g., using a forwarding mechanism. In our case, we simply rely on the strong migration primitive of Jocaml. On the negative side, the correctness proof is more involved.

In this paper, we present an asynchronous distributed algorithm for implementing ambients, we make it precise as a translation into the join calculus—the process calculus that is a model of Jocaml [9], and we refine this translation into a distributed implementation of ambients written in Jocaml. The algorithm provides an insight into the implementability of ambients. The Jocaml prototype is a first, lightweight, distributed implementation of ambients. The translation is proved correct in two stages: first we use barbed coupled simulations for the correctness of the algorithm, then we use an hybrid barbed bisimulation for the actual translation into the join calculus. Technically, the first stage is a first application of coupled-simulations [17] in a reduction-based, asynchronous setting; it relies on the introduction of an auxiliary ambient calculus extended with transient states; it does not depend of the target language. The second stage is a challenging application of the decreasing diagram technique [16]. In combination, these results imply that the translation preserves and reflects a variety of global observation predicates.

The paper is organized as follows. In section 2, we present the asynchronous algorithm and we show a formal translation from ambient processes to join calculus processes. In section 3, we discuss the correctness of the translation in terms of observations. In section 4, we focus on the operational correspondence between a process and its translation; to this end, we refine the ambient calculus to express additional transient states induced by the translation. In section 5, we state our main technical results and give an idea of their proofs. In section 6, we describe more practical aspects of the Jocaml implementation. We conclude in section 7. In an appendix, we recall the operational semantics of the distributed

join calculus and of the calculus of mobile ambients, and we give an overview of both calculi. Additional discussions, technical details, and all the proofs appear in the full version of this paper [10].

2 From Ambients to the Join Calculus

We describe the asynchronous algorithm, then we specify it as a translation from ambients to the join calculus. We begin with a fragment of the ambient calculus given by the grammar $P ::= a[P] \mid \text{in } a.P \mid \text{out } a.P \mid P \mid P' \mid \mathbf{0}$. In a second stage, we incorporate OPEN steps and other ambient constructs.

2.1 An Asynchronous Algorithm

The dynamic tree structure of ambients is represented by a doubly linked tree. Each node in the tree implements an ambient: each node contains non-ambient processes such as $\text{in } b.P$ or $\text{out } a.c[Q]$ running in parallel; each node also hosts an *ambient manager* that controls the steps performed in this ambient and in its direct subambients. Different nodes may be running at different physical sites, so their ambient managers should use only asynchronous messages to communicate with one another. Since several ambients may have the same name, each node is also associated with a unique identifier. (Informally, we still refer to ambients by name, rather than unique identifier.)

Each ambient points to its subambients and to its parent ambient. The down links are used for controlling subambients, the up link is used for proposing new actions. The parent of the moving ambient for an IN-step knows the destination ambient; the parent also knows the destination ambient—its own parent—for an OUT-step; it controls the opened ambient for an OPEN-step. Hence, the decision to perform a step will always be taken by the parent of the affected ambient.

Moves of ambient a in and out of ambient b correspond to three successive steps, depicted below. Single arrows represent current links; double arrows represent messages in transit.

We detail the dynamics of an IN-step, e.g., $c[a[\text{in } b.Q] \mid b[\mathbf{0}]] \rightarrow c[b[a[Q]]]$.

- 0-step:** initially, a delegates the migration request $\text{IN } b$ to its current parent (here c); to this end, it uses its current up link to send a message to c saying that a is willing to move into an ambient named b .
- 1-step:** the enclosing ambient c matches a 's request with a 's and b 's down links. Atomically, a 's request and the down link to a are erased, and a relocation message is sent to a ; this message contains the address of b , so that a will be able to relocate to b , and also a descriptor of a 's successful action, so that a can complete this step by triggering its guarded process Q .
- 2-step:** the moving ambient a receives c 's relocation message, relocates to b 's site, and updates its up link to point to b . It also sends a message to b that eventually registers a as a subambient of b , establishing the new downlink.

The 1-step may preempt other actions delegated by a to its former parent c . Such actions should now be delegated to its new parent b . For that purpose, a 's ambient manager keeps a log of the pending actions delegated in 0-steps, and, as it completes one of these actions in a 2-step, it re-delegates all other actions towards its new parent. The log cannot be maintained by the parent, because delegation messages may arrive long after a 's departure. Moreover, in the case an ambient moves back into a former parent, former delegation messages may still arrive, and should not be confused with fresh ones. Such stale messages must be deleted. This is not directly possible in an asynchronous world, but equivalently each migration results in a modification of the unique identifier of the moving ambient, each delegation message is tagged with the current identifier, and the parent discards every message with an old identifier.

An OUT-step of a out of b corresponds to the same series of three steps. The main different is in step 1, as the enclosing ambient b matches a 's request with a 's down link and its own name b , and passes its own up link in the relocation message sent back to a .

2.2 A Simple Translation

The compositional translation $\llbracket \cdot \rrbracket_e$ appears in Figure 1. Overall, the tree of nested ambients is mapped to an isomorphic tree of nested locations. Each ambient is mapped to a join calculus location containing the definition D of the channel names that form the ambient interface, and containing processes that represent the ambient state. The definition D is composed of three groups of rules D_0 , D_1 , and D_2 that respectively implement 0, 1, and 2-steps of the algorithm.

To represent the distributed data structure used in the algorithm of section 2.1, an ambient is represented by an interface e , which is a record that contains fields *here*, *amb*, *sub_{in}*, *sub_{out}*, *reloc*, *in*, and *out*. The *here*-field is the name of the location hosting the ambient, whereas the other fields are channel names used to interact with this ambient. The translation is parameterized by the interface e of the current enclosing ambient. A down link to a subambient named b with interface e_b and unique identifier (uid) j is represented as a message $\text{amb}(j, b, e_b)$. For every ambient, the up link to its parent ambient is represented by the parent interface e , which is stored in the state message $s(a, i, e, l)$. In

$$\begin{aligned}
\llbracket a[P] \rrbracket_e &= \mathbf{def} \ AM_{a,e}(P) \ \mathbf{in} \ \mathbf{0} \\
\llbracket \mathbf{in} \ a.P \rrbracket_e &= \mathbf{def} \ \kappa() \triangleright \llbracket P \rrbracket_e \ \mathbf{in} \ e.in(a, \kappa) \\
\llbracket \mathbf{out} \ a.P \rrbracket_e &= \mathbf{def} \ \kappa() \triangleright \llbracket P \rrbracket_e \ \mathbf{in} \ e.out(a, \kappa) \\
\llbracket P \mid Q \rrbracket_e &= \llbracket P \rrbracket_e \mid \llbracket Q \rrbracket_e \\
\llbracket \mathbf{0} \rrbracket_e &= \mathbf{0}
\end{aligned}$$

where the ambient manager $AM_{a,e}(P)$ is defined as:

$$\begin{aligned}
D_0 &\stackrel{\text{def}}{=} s(a, i, e, l) \mid in(b, \kappa) \triangleright s(a, i, e, l \cup \{\text{In } b \ \kappa\}) \mid e.sub_{in}(i, b, \kappa) \\
&\quad \wedge s(a, i, e, l) \mid out(b, \kappa) \triangleright s(a, i, e, l \cup \{\text{Out } b \ \kappa\}) \mid e.sub_{out}(i, b, \kappa) \\
D_1 &\stackrel{\text{def}}{=} s(a, i, e, l) \mid amb(j, b, e_b) \mid amb(k, c, e_c) \mid sub_{in}(k, b, \kappa) \triangleright \\
&\quad s(a, i, e, l) \mid amb(j, b, e_b) \mid e_c.reloc(e_b, \kappa) \\
&\quad \wedge s(a, i, e, l) \mid amb(j, b, e_b) \mid sub_{out}(j, a, \kappa) \triangleright s(a, i, e, l) \mid e_b.reloc(e, \kappa) \\
D_2 &\stackrel{\text{def}}{=} s(a, i, e, l) \mid reloc(e', \kappa) \triangleright \mathbf{go}(e'.here); (I_{a,e_h,e'} \mid \kappa() \mid Flush(l, in, out, \kappa)) \\
D &\stackrel{\text{def}}{=} D_0 \wedge D_1 \wedge D_2 \\
I_{a,e_h,e} &\stackrel{\text{def}}{=} \mathbf{def} \ \mathbf{uid} \ i \ \mathbf{in} \ s(a, i, e, \emptyset) \mid e.amb(i, a, e_h) \\
AM_{a,e}(P) &\stackrel{\text{def}}{=} \mathbf{here} \left[D : I_{a,e_h,e} \mid \llbracket P \rrbracket_{e_h} \right]
\end{aligned}$$

with the record notation $e_h \stackrel{\text{def}}{=} \left\{ \begin{array}{l} here = here, amb = amb, sub_{in} = sub_{in}, \\ sub_{out} = sub_{out}, reloc = reloc, in = in, out = out \end{array} \right\}$.

Figure 1. Translation from IN/OUT ambient processes to the join calculus

addition, the state message contains the name a and the current uid i of the ambient, and the log l of IN and OUT actions that have been delegated to the parent ambient using e .

We resume our study of an IN-action, considering the role of each message in the translation of $c[a[\mathbf{in} \ b.Q] \mid b[\mathbf{0}]]$. Initially, the translation defines a continuation κ for Q and issues a message $in(b, \kappa)$ in a , which is a subjective migration request into an ambient named b .

The 0-step consists of delegating the request to the parent ambient. Using the first rule of a 's D_0 , the messages $s(a, i, e, l)$ and $in(b, \kappa)$ are consumed, the request is recorded in a 's log as an entry $\text{In } b \ \kappa$, and the request is forwarded to the enclosing ambient c described by the interface e . The ambient a remains active, with new state $s(a, i, e, l \cup \{\text{In } b \ \kappa\})$. In parallel, the message $e.sub_{in}(i, b, \kappa)$ is a subambient move request sent to c , with the explicit identifier i of the requester a .

The 1-step is performed by c 's ambient manager. The message $sub_{in}(i, b, \kappa)$ may be consumed using the first rule of D_1 . The rule also requires that both the ambient that issued the request and another destination ambient with name b be actually present. This step removes the down link for the moving ambient—the message $amb(i, a, e_a)$ —, thus blocking other actions competing for the same message, whereas the destination ambient remains available for concurrent steps. A

$$\begin{aligned}
[\text{open } a.P]_e &= \text{def } \kappa() \triangleright [P]_e \text{ in } e.\text{open}(a, \kappa) \\
[\langle n \rangle]_e &= e.\text{send}(n) \\
[(n).P]_e &= \text{def } \kappa(n) \triangleright [P]_e \text{ in } e.\text{recv}(\kappa) \\
[!P]_e &= \text{def } \kappa() \triangleright [P]_e \mid \kappa() \text{ in } \kappa() \\
[\nu a.P]_e &= \text{def fresh } a \text{ in } [P]_e
\end{aligned}$$

with additional rules in the definition of $AM_{a,e}(P)$:

$$\begin{aligned}
D'_1 &\stackrel{\text{def}}{=} s(a, i, e, l) \mid \text{amb}(j, b, e_b) \mid \text{open}(b, \kappa) \triangleright s(a, i, e, l) \mid e_b.\text{opening}(\kappa) \\
D'_2 &\stackrel{\text{def}}{=} s(a, i, e, l) \mid \text{opening}(\kappa) \triangleright f(e) \mid \kappa() \mid \text{Flush}(l, e.\text{in}, e.\text{out}, \kappa) \\
D_C &\stackrel{\text{def}}{=} s(a, i, e, l) \mid \text{recv}(\kappa) \mid \text{send}(n) \triangleright s(a, i, e, l) \mid \kappa(n) \\
D_F &\stackrel{\text{def}}{=} f(e) \mid \text{in}(b, \kappa) \triangleright f(e) \mid e.\text{in}(b, \kappa) \\
&\quad \wedge f(e) \mid \text{out}(b, \kappa) \triangleright f(e) \mid e.\text{out}(b, \kappa) \\
&\quad \wedge f(e) \mid \text{open}(b, \kappa) \triangleright f(e) \mid e.\text{open}(b, \kappa) \\
&\quad \wedge f(e) \mid \text{amb}(j, b, e_b) \triangleright f(e) \mid e.\text{amb}(j, b, e_b) \\
&\quad \wedge f(e) \mid \text{sub}_{\text{in}}(k, b, \kappa) \triangleright f(e) \mid e.\text{sub}_{\text{in}}(k, b, \kappa) \\
&\quad \wedge f(e) \mid \text{sub}_{\text{out}}(k, b, \kappa) \triangleright f(e) \mid e.\text{sub}_{\text{out}}(k, b, \kappa) \\
&\quad \wedge f(e) \mid \text{recv}(\kappa) \triangleright f(e) \mid e.\text{recv}(\kappa) \\
&\quad \wedge f(e) \mid \text{send}(b) \triangleright f(e) \mid e.\text{send}(b)
\end{aligned}$$

$$D \stackrel{\text{def}}{=} D_0 \wedge D_1 \wedge D'_1 \wedge D_2 \wedge D'_2 \wedge D_C \wedge D_F$$

with the extended record notation

$$e_h \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \text{here} = \text{here}, \text{amb} = \text{amb}, \text{sub}_{\text{in}} = \text{sub}_{\text{in}}, \text{sub}_{\text{out}} = \text{sub}_{\text{out}}, \text{open} = \text{open}, \\ \text{reloc} = \text{reloc}, \text{in} = \text{in}, \text{out} = \text{out}, \text{opening} = \text{opening}, \text{recv} = \text{recv}, \text{send} = \text{send} \end{array} \right\}$$

Figure 2. Additional clauses for the full translation

relocation message $e_a.\text{reloc}(e_b, \kappa)$ is emitted, signalling to the requesting ambient a that it must migrate to the ambient with interface e_b , with continuation κ .

The 2-step, using a 's rule D_2 , consumes the message on reloc and the current state message, performs a join calculus migration to the location of the destination ambient, then resumes the activities of a with parent interface e_b . To this end, the process I_{a,e_a,e_b} restores an active state: it generates a fresh uid i' , issues a local state message $s(a, i', e_b, \emptyset)$ representing the up link, and sends to the new parent a message $e_b.\text{amb}(i', a, e_a)$ representing the down link. (Since no down link will ever mention the previous uid i , previous delegation messages tagged with i will never match a rule of D_1 . In our implementation, these stale messages are actually discarded.) In addition, the message $\kappa()$ triggers the continuation. Finally, the process $\text{Flush}(l \cup \{\text{In } b \ \kappa\}), \text{in}, \text{out}, \kappa)$ restarts any preempted actions appearing in the log. As defined in the appendix, this process emits a message $\text{in}(d, \kappa')$ or $\text{out}(d, \kappa')$ for every entry $\text{In } d \ \kappa'$ or $\text{Out } d \ \kappa'$ appearing in the log l and such that $\kappa' \neq \kappa$. These entries correspond to actions preempted by the

migration; they will be delegated to a new parent through other iterations of 0-steps.

Similarly, an OUT-step is performed according to the algorithm by using the second rule of D_0 of the moving ambient, the second rule of D_1 of the enclosing ambient, and finally rule D_2 of the moving ambient.

2.3 Dealing with Other Ambient Constructs

The translation of Figure 2 generalizes the translation above to the full ambient calculus. For each additional construct, we add a clause to the compositional translation $\llbracket \cdot \rrbracket_e$. We also upgrade $AM_{a,e}(\cdot)$, and use a larger environment e with extra fields for *open*, *opening*, *recv*, and *send*.

Values and Scopes. Ambient names are mapped to identical names in the join calculus. The two calculi rely on similar lexical scope disciplines, with scope extrusion performed by structural equivalence (rule SCOPE in join, rules R1 and R2 in ambients). Thus, it suffices to translate the creation of local ambient names $\nu a.P$ into binders **fresh** a with the same scope in the join calculus.

Communication. Ambient communication is implemented by supplementing every ambient manager with a rule D_C that binds two channels *send* and *recv* and synchronizes message outputs and message requests. This encoding is much like the encoding of pi-calculus channels into the join calculus (see [8]).

Replication. Each replicated process $!P$ is coded using a standard recursive encoding of infinite loops in the join calculus.

Open. Ambient processes may dissolve ambient boundaries using the **open** capabilities. In contrast, join calculus names are statically attached to their defining location, and location boundaries never disappear. We thus lose the one-to-one mapping from ambients to locations, and distinguish two states for each location: either the ambient is still running and the message s is present, or it has been opened and henceforth messages sent to its interface are forwarded to the enclosing ambient. The indirection is achieved by using a persistent message sent on f defined in D_F . This leads to complications in the proofs, as one must prove that these opened locations do not interfere with the rest of the translation.

3 Correctness of the Translation

The distributed synchronization algorithm seems to depart from the operational semantics of ambients, and the translation of nested ambients yields arbitrarily large terms with numerous instances of the algorithm running in parallel. This makes the correctness of the distributed implementation problematic. Technically, both calculi have a reduction-based semantics, which can be equipped with standard notions of observation. This provides a precise setting for establishing correctness on the translation, rather than on an abstraction of the algorithm.

(Of course, there are still minor discrepancies between the translation and the actual code in Jocaml; see section 6.)

We first define a syntactic notion of observation. For each calculus, we use a family of predicates on processes P indexed by names b , written $P \Downarrow_b$.

- In the ambient calculus, $P \Downarrow_b$ when b is free in P and $P \equiv \nu \tilde{v}.(b[Q] \mid R)$.
- In the join calculus, $P \Downarrow_b$ when b is free in P and $P \equiv \alpha[D' : b(\tilde{v}) \mid P'] \wedge D$.

Next, we express the correctness of the translation in terms of the following predicates, for both ambient and join processes:

- A process P has a *weak barb* on b (written $P \Downarrow_b$) when $P \rightarrow^* P' \Downarrow_b$.
- A process P *diverges* (written $P \Uparrow$) when P has an infinite series of steps.
- A process P has a *fair-must barb* on b (written $\Box P \Downarrow_b$) when for all P' such that $P \rightarrow^* P'$, we have $P' \Downarrow_b$.

In combination, these predicates give a precise content to the informal notion of correctness: “the translation should neither suppress existing behaviors, nor introduce additional behaviors.” The minimal notion of correctness for an implementation is the reflection of weak barbs, which rules out spurious behaviors; the converse direction states that the implementation does not discard potential behaviors; for instance, it rules out new deadlocks, or even an empty translation. The preservation of convergence is of pragmatic importance. In addition, correctness for fair-must tests relates infinite computations [7], and rules out implementations with restrictive scheduling policies.

Since top-level ambients are not translated into messages on free names, the observation of translated ambients requires some special care. To this end, we supplement the translation at top-level with a definition D_t that reveals the presence of a particular barb $\Downarrow_b$ in the source process. Using D and e_h as defined in figure 2, and for a given interface e , we define the top-level translation

$$\begin{aligned} \llbracket P \rrbracket^b &\stackrel{\text{def}}{=} \text{here}[D \wedge D_t \wedge \text{uid } i : s(a, i, e, \emptyset) \mid t(b) \mid \llbracket P \rrbracket_{e_h}] \\ D_t &\stackrel{\text{def}}{=} s(a, i, e, l) \mid \text{amb}(j, b, e_b) \mid t(b) \triangleright s(a, i, e, l) \mid \text{amb}(j, b, e_b) \mid \text{yes}() \end{aligned}$$

Without loss of generality, we always assume that the names in a, i, t, e, yes do not clash with any name free in P , and that the location and channel names introduced by the translation do not clash with any name of P .

We are now ready to state that all the derived observations discussed above are preserved and reflected by the translation:

Theorem 1. *For every ambient process P and name b , we have $P \Downarrow_b$ if and only if $\llbracket P \rrbracket^b \Downarrow_{\text{yes}}$; $P \Uparrow$ if and only if $\llbracket P \rrbracket^b \Uparrow$; and $\Box P \Downarrow_b$ if and only if $\Box \llbracket P \rrbracket^b \Downarrow_{\text{yes}}$.*

While correctness is naturally expressed in terms of observations along the reduction traces of processes, its proof is challenging. In particular, a direct approach clearly leads to intractable inductions on both the syntax of the source process and on the series of reductions.

4 A Calculus of Ambients Extended with Transient States

In order to prove theorem 1, we introduce an intermediate calculus of ambients with constructs that materialize the key transient states of the algorithm of section 2.1, and we equip this calculus with a reduction semantics in direct correspondence with the algorithm. For instance, atomic IN steps are decomposed into series of 1- and 2-steps. (However, 0-steps are not represented, inasmuch as requests are always eventually delegated to the current parent.) In the next section, we rely on the extended calculus to establish correctness as the composition of two equivalences. First, we use *coupled simulations* [17] to relate the two semantics for ambients; then, we use bisimulations to relate ambients equipped with the extended semantics to their join calculus translations.

The grammar for the extended calculus appears in Figure 3. It has new processes representing ambients that are committed to move or to be opened, as the result of their father's 1-step—we call such transient ambients *stubs*—and also new processes marking the future position of migrating ambients—we call such precursors *scions*. Pairs of stubs and scions are syntactically connected by a marker i . The extended operational semantics appears in Figure 4. It is a reduction-based semantics with auxiliary labels for stubs and scions. Each of the reduction steps IN, OUT, and OPEN is decomposed in two steps. Initial steps $\rightarrow_1$ introduce stubs and scions; completion steps $\rightarrow_2$ consume them. The reduction rules for RECV and REPL are those of the original semantics, except that we write $\rightarrow_C$ instead of $\rightarrow$. Overall, we obtain a reduction system for extended ambient processes with steps $\rightarrow_{12C} \stackrel{\text{def}}{=} \rightarrow_1 \cup \rightarrow_2 \cup \rightarrow_C$. For the sake of comparison, we also extend the original reduction semantics and the observation predicates from ambients to extended ambients.

Initially, stubs and scions are neighbors, but they may drift apart as the result of other steps before performing the matching completion step, so an auxiliary labeled transition system is used to match stubs and scions. The completion of a deferred migration is thus rendered as a global communication step between processes residing in two different ambient contexts, and syntactically linked by the name i . This may seem difficult to implement, but actually scions have no operational contents; they represent the passive target for a strong migration. To illustrate the extended calculus, we give below the reductions for a process with a critical pair. Processes P_i are regular ambient processes; processes Q_i are transient processes reflecting intermediate states of our algorithm.

$P ::=$	extended ambient process
$\dots$	all the constructors of Figure 5
$ \quad Xn[P]$	stub
$ \quad i$	scion
$ \quad \nu i.P$	marker restriction
$X ::=$	state extension
$\bar{i}\{P\}-$	the stub is committed to move to i
$ \quad o\{P\}-$	the stub is being opened

Well-formed conditions: stubs and scions may occur only in extended evaluation contexts; restricted markers i must be used linearly (exactly one stub and one scion). In the following, we write $X \equiv n[P]$ for either $Xn[P]$ or $n[P]$.

Figure 3. Syntax for an ambient calculus extended with transient states

Extended evaluation contexts $E(\cdot)$ are defined by the grammar

$$E(\cdot) ::= \cdot \mid P \mid E(\cdot) \mid E(\cdot) \mid P \mid X \equiv n[E(\cdot)] \mid \nu n.E(\cdot) \mid \nu i.E(\cdot)$$

Structural equivalence $\equiv$ is the smallest equivalence relation on processes that is closed by application of extended evaluation contexts and by α -conversion, and that satisfies the axioms P0, P1, P2, C0, C1, R1, R2 of figure 6 and:

$$\text{R2X} \quad \frac{m \neq n \quad m \text{ is not free in } X}{Xn[\nu m.P] \equiv \nu m.Xn[P]}$$

Labeled transitions $\xrightarrow{\alpha}$ are the smallest families of relations closed by application of restriction-free extended evaluation contexts and such that

$$\text{STUB} \quad \bar{i}\{Q\}-n[R] \xrightarrow{\bar{i}.n[Q|R]} \mathbf{0} \quad \text{SCION} \quad i \xrightarrow{i.P} P$$

Original reduction steps $\rightarrow$ are defined as in Figure 6 with extended evaluation contexts. *Initial steps* $\rightarrow_1$, *Completion steps* $\rightarrow_2$, and *Other steps* $\rightarrow_C$ are the smallest relations closed by structural equivalence, by application of extended evaluation contexts, and such that:

$$\begin{array}{ll} \text{IN 1} \quad \frac{m[P] \mid n[\text{in } m.Q \mid R]}{\rightarrow_1 \nu i. m[i \mid P] \mid \bar{i}\{Q\}-n[R]} & \text{OUT 1} \quad \frac{X \equiv m[P \mid n[\text{out } m.Q \mid R]]}{\rightarrow_1 \nu i. i \mid X \equiv m[P \mid \bar{i}\{Q\}-n[R]]} \\ \text{MOVE 2} \quad \frac{P \xrightarrow{\bar{i}.S} P' \quad Q \xrightarrow{i.S} Q'}{\nu i.(P \mid Q) \rightarrow_2 P' \mid Q'} & \\ \text{OPEN 1} \quad \text{open } n.Q \mid n[R] \rightarrow_1 o\{Q\}-n[R] & \text{OPEN 2} \quad o\{Q\}-n[R] \rightarrow_2 Q \mid R \\ \text{RECV} \quad \langle n \rangle \mid (x).P \rightarrow_C P\{n/x\} & \text{REPL} \quad !P \rightarrow_C P \mid P \end{array}$$

Figure 4. Semantics for an ambient calculus extended with transient states

5 Coupled Simulations and Operational Correspondence

We continue our study of correctness in terms of equivalences based upon weak barbs. These equivalences are essential to obtain a modular proof. As a side benefit, they also provide a finer account of correctness. (See [7] for a discussion of equivalences and encodings in process calculi.) Instead of equivalences, we actually often use relations ranging over different domains, equipped with different notions of reduction steps $\rightarrow_a$, $\rightarrow_b$ and families of observations $\Downarrow_{a,x}$ and $\Downarrow_{b,x}$.

Definition 1 (Barbed bisimulations). *A relation $\mathcal{R} \in \mathcal{P}_a \times \mathcal{P}_b$ is a weak barbed simulation when, for all $P \mathcal{R} Q$, we have (1) if $P \rightarrow_a^* P'$, then there exists Q' such that $Q \rightarrow_b^* Q'$ and $P' \mathcal{R} Q'$; (2) for all x , if $P \Downarrow_{a,x}$, then $Q \Downarrow_{b,x}$. $\mathcal{R}$ is a barbed bisimulation when $\mathcal{R}$ and its converse $\mathcal{R}^{-1}$ are barbed simulations.*

Bisimulations come with effective proof techniques that consider only a few steps at a time, rather than whole execution traces. Unfortunately, barbed bisimilarity $\approx$ —the largest barbed bisimulation closed by application of evaluation contexts—is too discriminating for our protocol. Transient processes such as Q_1 in the example above account for a partially-committed internal choice: Q_1 may reduce to P_1 and P_3 , but not to P_2 . In some contexts, they are not bisimilar to any derivative of P in the original ambient semantics. To address this issue of gradual commitment, Parrow and Sjödin proposed coarser relations called coupled simulations [17,15]. We liberally adapt their definition to ambients:

Definition 2 (Coupled simulations). *The relations $\leq \in \mathcal{P}_a \times \mathcal{P}_b$ and $\leq \in \mathcal{P}_b \times \mathcal{P}_a$ form a pair of barbed coupled simulations when $\leq$ and $\leq$ are barbed simulations that meet the coupling conditions: (1) if $P \leq Q$, then $Q \rightarrow_b^* \leq P$; (2) if $Q \leq P$, then $P \rightarrow_a^* \leq Q$.*

The discrepancy between $\leq$ and $\geq$ is most useful for handling transient states such as Q_1 . The coupling conditions guarantee that every transient state can be mapped both to a less advanced state and to a more advanced one. In our case, we would have $Q_1 \leq P$ and $P_i \leq Q_1$ for $i = 1, 3$.

Correctness of the Asynchronous Algorithm The first stage of our correctness argument is expressed as coupled simulations between ambient processes equipped with the original and the extended semantics. In the statement below, related processes have the same syntax, but live in different calculi, equipped with different reduction semantics.

Theorem 2. *Let $\leq$ be the union of $\leq \cap \geq$ for all barbed coupled simulations between ambient and extended ambient processes that are closed by application of evaluation contexts. For all ambient processes P , we have $P \leq P$.*

The proof appears in [10]; it makes apparent some subtleties of our algorithm due to additional concurrency. After a first series of results on partial commutation properties for extended steps, the key lemmas establish that, for any ambient process P and extended ambient process Q , if $P \rightarrow_{12C}^* Q$, then (1) for some ambient process P' we have $Q \rightarrow_{12C}^* P'$ and (2) for any such process P' , we also have $P \rightarrow^* P'$ in the original semantics.

Operational Correspondence The second stage of the proof relates ambients equipped with the extended semantics to their join calculus translations. It is simpler than the first one, in principle, but its proof is complicated because the translation makes explicit many details of the implementation that are inessential to the algorithm.

In order to express the correspondence of observations across the translation, we supplement the top-level translation $\llbracket \cdot \rrbracket^b$ of theorem 1 with an external choice of the ambient barb to be tested. With the same notations, we write $\llbracket \cdot \rrbracket^t$ for the translation that maps every process P to the process $[D \wedge D_t \wedge \text{uid } i : s(a, i, e, \emptyset) \mid p(t) \mid \llbracket P \rrbracket_{e_0}]$. As before, we assume that names in a, i, e, p, t , and yes do not clash with names free in P . At any point, we can use the evaluation context $T_b(\cdot) \stackrel{\text{def}}{=} h_T[p(t) \triangleright t(b) : \mathbf{0}] \wedge (\cdot)$ to test a translated ambient barb on b by testing the plain join calculus barb $T_b(\cdot) \Downarrow_{yes}$. We have $\llbracket P \rrbracket^b \approx T_b(\llbracket P \rrbracket^t)$ in the join calculus.

Theorem 3 (Correctness of the translation). *Let $\approx^{aj}$ be the largest bisimulation between extended ambient processes with reductions $\rightarrow_{12C}$ and join processes such that $Q \approx^{aj} R$ implies $Q \Downarrow_b$ iff $T_b(R) \Downarrow_{yes}$. For all ambient processes P , we have $P \approx^{aj} \llbracket P \rrbracket^t$.*

In the long version of the paper [10], a more precise *strong bisimulation up to bookkeeping* result is proved for the translations of all reachable extended ambient processes. Since every significant transient state induced by the translation has been lifted to the extended ambient calculus, its proof essentially amounts to an operational correspondence between the two calculi. We partition reductions in the join calculus according to the static rule of the translation being used. For instance, we let $\rightarrow_1$ steps in the join calculus be the steps using a rule of a definition $D_1 \wedge D'_1$ of figure 2; these steps create a *reloc* or an *opening* message, and are in operational correspondence with source $\rightarrow_1$ steps. We obtain two main classes of join calculus steps: steps $\rightarrow_{12C}$ that can be traced back to extended ambient steps, and “bookkeeping” steps $\rightarrow_B$, which are auxiliary steps used to trigger continuations, migrate, manage the logs, or unfold new ambient managers.

The main lemmas describe dynamic simplifications of derivatives of the translation, which are required to obtain translations of derivatives in the extended source calculus. These lemmas are expressed as elementary commutation diagrams between simplification relations and families of reduction steps. For instance, one lemma states that “stale messages” can be discarded; another, more complex lemma states that locations and ambient managers representing opened ambients can be eliminated, effectively merging the contents of opened ambients with the contents of their previously-enclosing ambient. To conclude, we exhibit a bisimulation relation between extended ambients equipped with steps $\rightarrow_{12C}$ and global translations of these extended ambients equipped with steps $\rightarrow_B^* \rightarrow_{12C} \rightarrow_B^*$. The proof is structured using the decreasing diagram technique of [16], whose conditions guarantee that every weak simulation diagram in the final proof can be obtained by gluing previously-established diagrams.

6 Distributed Implementation

In this section, we briefly describe the actual implementation in Jocasml. The initialization and the dynamics of distribution for ambient processes among several Jocasml runtimes lead to some design choice, discussed in the long version of the paper [10]. We also refer to [11] for the source code, setup instructions, and programming examples.

Our implementation closely follows the translation given in figures 1 and 2. Since Jocasml already provides support for mobility, local synchronization, and run-time distribution, our code is very compact—less than 400 lines for the interpreter, less than 40k in bytecode for the object files. The main differences between the formal translation and the code are given below:

- Messages in the implementation may pass names, but also arbitrary chains of capabilities [4], as for instance in the ambient process $\langle \text{in } a.\text{out } b \rangle \mid !(x).x.\langle x \rangle$.
- The implementation is an interpreter that delays the translation of guarded processes. The implementation maintains an environment for local variables, and performs dynamic type checking whenever a value is used either as a name or as a capability.
- The translation relies on non-linear join-patterns, which are not available in Jocasml. More explicitly, the implementation relies on hash tables to cache some messages: when a message arrives on sub_{in} , sub_{out} , amb , or open , either it is immediately used in combination with the message cache, or it is added to the cache.
- The formal translation of replication always yields a diverging computation (as in the source ambient process). More reasonably, the interpreter unfolds replication on demand: since every ambient reduction involves at most two copies of a replicated process, it suffices to initially unfold two copies, then to unfold an additional copy whenever a fresh copy is used or modified. Hence, the process $!a[\]$ does not diverge, while $!a[\text{in } a]$ still does.

7 Conclusions

We translated Mobile Ambients into the join calculus, and gave a first, asynchronous, distributed implementation of the ambient calculus in Jocasml, with a high level of concurrency. The synchronization mechanisms of Ambients turned out to be challenging first to implement, then to prove correct. This provides an insight into the ambient calculus as a model of concurrency. At the same time, this shows how Jocasml and its formal model can be used to tackle distributed and mobile implementations. For instance, the translation takes full advantage of join patterns to describe complex local synchronization steps, while a more traditional language would decompose these steps into explicit series of reads and updates protected by locks.

In order to get a safer and more efficient implementation, one should care about typing information for passed values and mobility capabilities [5,3]. Our implementation insures dynamic type-checking on values, whereas it would be

preferable to use the static type-checking discipline of Jocaml. Similarly, static knowledge of the actions that never appear in a given ambient can lead to more efficient, specialized ambient managers.

Finally, little is known about actual programming with Ambients, or the relevant abstractions to build a high-level language on top of the ambient calculus. While we did not consider changing the source language, we believe that our implementation provides an adequate platform for experimenting with ambient-based language design. For instance, our translation would easily accommodate the co-capabilities proposed in [14].

Acknowledgments. This work benefited from discussions with Luca Cardelli, Fabrice Le Fessant, and Luc Maranget.

References

1. L. Cardelli. *Ambit*, 1997. Available from <http://www.luca.demon.co.uk/Ambit/Ambit.html>.
2. L. Cardelli. Mobile ambient synchronization. Technical note 1997-013, Digital Systems Research Center, July 1997.
3. L. Cardelli, G. Ghelli, and A. D. Gordon. Mobility types for mobile ambients. In *ICALP'99*, volume 1644 of *LNCS*, pages 230–239, 1999.
4. L. Cardelli and A. Gordon. Mobile ambients. In *FoSSaCS'98*, volume 1378 of *LNCS*, pages 140–155, 1998.
5. L. Cardelli and A. D. Gordon. Types for mobile ambients. In *POPL'99*, pages 79–92. ACM, Jan. 1999.
6. S. Conchon and F. Le Fessant. Jocaml: Mobile agents for objective-caml. In *ASA/MA'99*, pages 22–29. IEEE Computer Society, Oct. 1999.
7. C. Fournet. *The Join-Calculus: a Calculus for Distributed Mobile Programming*. PhD thesis, Ecole Polytechnique, Palaiseau, Nov. 1998. INRIA, TU-0556.
8. C. Fournet and G. Gonthier. The re exive chemical abstract machine and the join-calculus. In *POPL'96*, pages 372–385. ACM, Jan. 1996.
9. C. Fournet, G. Gonthier, J.-J. Lévy, L. Maranget, and D. Rémy. A calculus of mobile agents. In *CONCUR'96*, volume 1119 of *LNCS*, pages 406–421, Aug. 1996.
10. C. Fournet, J.-J. Lévy, and A. Schmitt. A distributed implementation of Ambients. Long version of this paper, available from <http://join.inria.fr/ambients.html>, 1999.
11. C. Fournet and A. Schmitt. An implementation of Ambients in JoCAML. Software available from <http://join.inria.fr/ambients.html>, 1999.
12. A. D. Gordon and L. Cardelli. Equational properties of mobile ambients. In *FoSSaCS'99*, volume 1578 of *LNCS*, pages 212–226, 1999.
13. F. Le Fessant. The JoCAML system prototype. Software and documentation available from <http://pauillac.inria.fr/jocaml>, 1998.
14. F. Levi and D. Sangiorgi. Controlling interference in ambients. In *POPL'00*, pages 352–364. ACM, Jan. 2000.
15. U. Nestmann and B. C. Pierce. Decoding choice encodings. In *CONCUR'96*, volume 1119 of *LNCS*, pages 179–194, Aug. 1996. Revised full version as report ERCIM-10/97-R051, 1997.
16. V. Oostrom. Con uence by decreasing diagrams. *Theoretical Computer Science*, 126:259–280, 1994.
17. J. Parrow and P. Sjödin. Multiway synchronization verified with coupled simulation. In *CONCUR'92*, volume 630 of *LNCS*, pages 518–533, 1992.

Appendix: Two Notions of Mobile Computations

We define our syntax and semantics for the calculus of Mobile Ambients and for the distributed join calculus. We refer to the long version of the paper [10] for an overview of the two calculi.

Operational Semantics for Ambients Our syntax and semantics of the calculus of ambients are given in figures 5 and 6.

This presentation slightly differs from Cardelli and Gordon’s [4] on several counts. In spirit, it is actually closer to the harness semantics of [12]. Our structural equivalence is more restrictive; it does not introduce or remove νx binders; it operates only in evaluation contexts. Our operational semantics represents the unfolding of replication as a silent reduction step rather than a structural law. Also, communicated values are just ambient names, rather than both names and chains of capabilities. (Chains of capabilities are fully supported in the Jocaml implementation, but their encoding is heavy.)

Operational Semantics for the Join Calculus Our syntax and semantics for the join calculus are given in figures 7 and 1.

A *path* α is a string of location names $a, b, \dots$. *Active locations* are locations not under a **def**; they can be nested. The path of an active sublocation $a[D : P]$ is $\alpha.a$, where α is the path of its enclosing location. A *configuration* is a conjunction of top-level locations such that every location has a unique name, such that the set of paths for all active locations is prefix-closed except for the empty prefix (i.e. active locations form a tree whose nodes are indexed by location names), and such that every channel is defined in at most one location. In a configuration, a location with path $\alpha.a$ is *folded* when it is the only top-level location whose path contains a .

Names can be bound either as parameters $\tilde{y}$ in a message pattern $x(\tilde{y})$ or as names defined in D by **def** D **in** P . The definition $a[D' : P]$ defines a and names defined in D' . A definition containing a rule with message pattern $x(\tilde{y})$ also defines x .

To simplify the translation of section 2, we supplement the join calculus with some convenient extensions, which are easily encoded in the plain join calculus. We supplement definitions with new constructs **uid** i and **fresh** a that bind names i and a , which we use to generate unique identifiers—we could use instead dummy rules such as $i() \triangleright 0$. We use a record notation $e = \{l_1 = x_1; \dots l_n = x_n\}$ as a shortcut for a tuple of names $x_1, \dots, x_n$ passed in a consistent order, and write $e.l_i$ for x_i . We use an algebraic notation **In** $b \ \kappa$, **Out** $b \ \kappa$ for log entries with tags **IN**, **OUT** and names b, κ . We use finite sets of log entries, interpreted in the standard mathematical sense. Rather than making explicit a standard encoding of set iterators for implementing the process $\text{Flush}(l, \text{in}, \text{out}, \kappa)$, we supplement the operational semantics of the join calculus with a rule for flushing logs of messages:

$$\text{FLUSH} \quad \text{Flush}(l, \text{in}, \text{out}, \kappa) \rightarrow \prod_{\text{In } d \ \kappa' \in l \mid \kappa \neq \kappa'} \text{in}(d, \kappa') \mid \prod_{\text{Out } d \ \kappa' \in l \mid \kappa \neq \kappa'} \text{out}(d, \kappa')$$

$P ::=$	ambient process
$n[P]$	ambient
$P \mid P'$	parallel composition
$C.P$	guarded process
$\nu n.P$	name restriction
$\langle n \rangle$	asynchronous message
$(x).P$	message reception
$!P$	replication
$\mathbf{0}$	inert process
$C ::=$	capability
$\text{in } n$	ingoing migration
$\text{out } n$	outgoing migration
$\text{open } n$	ambient dissolution

Figure 5. Syntax for the ambient calculus

Evaluation contexts $E(\cdot)$ are defined by the grammar

$$E(\cdot) ::= \cdot \mid P \mid E(\cdot) \mid E(\cdot) \mid P \mid n[E(\cdot)] \mid \nu n.E(\cdot)$$

Structural equivalence $\equiv$ is the smallest equivalence relation closed by application of evaluation contexts, by α -conversion, and such that

$$\begin{array}{ll}
\text{P0} & P \mid \mathbf{0} \equiv P \\
\text{P1} & P \mid P' \equiv P' \mid P \\
\text{P2} & (P \mid P') \mid P'' \equiv P \mid (P' \mid P'') \\
\text{R1} & \frac{n \text{ is not free in } P}{P \mid \nu n.Q \equiv \nu n.(P \mid Q)} \\
\text{R2} & \frac{m \neq n}{m[\nu n.P] \equiv \nu n.m[P]}
\end{array}$$

Ambient reduction $\rightarrow$ is the smallest relation closed by structural equivalence, by application of evaluation contexts, and such that

$$\begin{array}{ll}
\text{IN} & \frac{m[P] \mid n[\text{in } m.Q \mid R]}{\rightarrow m[P \mid n[Q \mid R]]} \quad \text{OUT} \quad \frac{m[P \mid n[\text{out } m.Q \mid R]]}{\rightarrow m[P] \mid n[Q \mid R]} \\
\text{OPEN} & \text{open } n.Q \mid n[R] \rightarrow Q \mid R \\
\text{RECV} & \langle n \rangle \mid (x).P \rightarrow P\{n/x\} \quad \text{REPL} \quad !P \rightarrow P \mid P
\end{array}$$

Figure 6. Operational semantics for the ambient calculus

$P ::=$	join calculus process
$\mathbf{0}$	inert process
$ P \mid P'$	parallel composition
$ x(\tilde{y})$	asynchronous message
$ \mathbf{go}(a); P$	migration request
$ \mathbf{def } D \mathbf{ in } P$	local definition
$D ::=$	join calculus definition
$\top$	void definition
$ D \wedge D'$	composition
$ J \triangleright P$	reaction rule
$ a[D : P]$	sub-location (named a , running D and P)
$ \alpha[D : P]$	top-level location (with path α , running D and P)
$J ::=$	join pattern
$x(\tilde{y})$	message pattern
$ J \mid J'$	synchronization

Figure 7. Syntax for the distributed join calculus

Structural equivalence $\equiv$ (on both processes and definitions) is the smallest equivalence relation closed by application of contexts $\cdot \wedge \cdot$, $\cdot \mid \cdot$ and $\alpha[\cdot : \cdot]$, by α -conversion on bound names, and such that:

P0	$P \mid \mathbf{0} \equiv P$	D0	$D \wedge \top \equiv D$
P1	$P \mid P' \equiv P' \mid P$	D1	$D \wedge D' \equiv D' \wedge D$
P2	$(P \mid P') \mid P'' \equiv P \mid (P' \mid P'')$	D2	$(D \wedge D') \wedge D'' \equiv D \wedge (D' \wedge D'')$
TREE	$\alpha[a[D' : P'] \wedge D : P]$ $\equiv \alpha.a[D' : P'] \wedge \alpha[D : P]$	SCOPE	$\frac{\text{names defined in } D' \text{ are fresh}}{\alpha[D : P \mid \mathbf{def } D' \mathbf{ in } P'] \equiv \alpha[D \wedge D' : P \mid P']}$

Join calculus reduction $\rightarrow$ is the smallest relation on configurations that is closed by structural equivalence and such that:

COMM	$\frac{x \text{ is defined in } D'}{\alpha[D : x(\tilde{v}) \mid P] \wedge \beta[D' : P'] \wedge E \rightarrow \alpha[D : P] \wedge \beta[D' : x(\tilde{v}) \mid P'] \wedge E}$	JOIN	$\frac{\sigma \text{ operates on message contents of } J}{\alpha[D \wedge J \triangleright Q : J\sigma \mid P] \wedge E \rightarrow \alpha[D \wedge J \triangleright Q : Q\sigma \mid P] \wedge E}$
Go	$\frac{a \text{ folded}}{\alpha.a[D : P \mid \mathbf{go}(b); Q] \wedge \beta.b[D' : P'] \wedge E \rightarrow \beta.b.a[D : P \mid Q] \wedge \beta.b[D' : P'] \wedge E}$		

Figure 8. Operational semantics for the distributed join calculus