

Nomadic Pict: Programming Languages, Communication Infrastructure Overlays, and Semantics for Mobile Computation

PETER SEWELL

Computer Laboratory, University of Cambridge

and

PAWEŁ T. WOJCIECHOWSKI

Institute of Computing Science, Poznań University of Technology

and

ASIS UNYAPOTH

Government Information Technology Services, Thailand

Mobile computation, in which executing computations can move from one physical computing device to another, is a recurring theme: from OS process migration, to language-level mobility, to virtual machine migration. This paper reports on the design, implementation, and verification of overlay networks to support reliable communication between migrating computations, in the Nomadic Pict project. We define two levels of abstraction as calculi with precise semantics: a low-level Nomadic π calculus with migration and location-dependent communication, and a high-level calculus that adds location-independent communication. Implementations of location-independent communication, as overlay networks that track migrations and forward messages, can be expressed as translations of the high-level calculus into the low. We discuss the design space of such overlay network algorithms and define three precisely, as such translations. Based on the calculi, we design and implement the Nomadic Pict distributed programming language, to let such algorithms (and simple applications above them) to be quickly prototyped. We go on to develop the semantic theory of the Nomadic π calculi, proving correctness of one example overlay network. This requires novel equivalences and congruence results that take migration into account, and reasoning principles for agents that are temporarily immobile (e.g. waiting on a lock elsewhere in the system). The whole stands as a demonstration of the use of principled semantics to address challenging system design problems.

Categories and Subject Descriptors: C.2.2 [Network Protocols]: ; C.2.4 [Distributed Systems]: ; D.3.3 [Language Constructs and Features]: ; F.3.1 [Specifying and Verifying and Reasoning about Programs]: ; F.3.2 [Semantics of Programming Languages]:

General Terms: Algorithms, Design, Languages, Theory, Verification

1. INTRODUCTION

Authors' addresses: P. Sewell, Computer Laboratory, University of Cambridge, J. J. Thomson Avenue, Cambridge CB3 0FD, UK. Peter.Sewell@cl.cam.ac.uk

P. T. Wojciechowski, Institute of Computing Science, Poznań University of Technology, Piotrowo 2, Poznań 60-965, Poland. Pawel.T.Wojciechowski@cs.put.poznan.pl

Asis Unyapoth, Government Information Technology Services, National Science and Technology Development Agency, Ministry of Science and Technology, Thailand. asis@gits.net.th

Mobile computation, in which executing computations can move (or be moved) from one physical computing device to another, has been a recurring focus of research, spanning disparate communities. The late 1970s and the 1980s saw extensive work on process migration, largely in the setting of operating system support for local-area distributed computation, using migration for load-balancing, checkpointing, etc. This was followed in the late 1990s by work on programming language support for mobility, largely in the mobile agent community, aiming at novel wide-area distributed applications. The late 1990s also saw work on semantics, using the tools of process calculi and operational semantics. In parallel, there has been a great deal of interest in the related areas of *mobile code*, popularised by Java applets, in which executable (but not yet executing) code can be moved, and in *mobile devices*, such as smartphones, PDAs, and the other devices envisaged in ubiquitous computing, which provide applications for both mobile computation and mobile code. Recently, the late 2000s have seen a renewed interest in mobile computation, now driven by the rise of virtualisation systems, such as VMWare and Xen, which support migration of client OS images. These are finally realising the prospect of commercial *commodity computation*, in which management of services and applications can be decoupled from physical machines in a datacentre, and in which flexible markets for computational resources can emerge.

Building systems with mobile computation, whether it be at the hypervisor, OS process, or programming language level, raises challenging problems, ranging from security concerns to interaction between changing versions of the infrastructure. In this paper we focus on one of these problems: that of providing reliable communication between migrating computations, with messages being delivered correctly even if the sending and receiving computation migrate. Such high-level *location independent* communication may greatly simplify the development of mobile applications, allowing movement and interaction to be treated as separate concerns. To provide reliable communication in the face of migration, above the low-level *location dependent* communication primitives of the existing Internet Protocol (IP) network, one essentially has to build an overlay network, to track migrations and route application messages to migrating computations. This infrastructure must address fundamental network issues such as failures, network latency, locality, and concurrency; the algorithms involved are thus inherently rather delicate, and cannot provide perfect location independence. Moreover, applications may be distributed on widely different scales (from local to wide-area networks), may exhibit different patterns of communication and migration, and may demand different levels of performance and robustness; these varying demands will lead to a multiplicity of infrastructures, based on a variety of algorithms. Lastly, these infrastructure algorithms will be to some extent exposed, via their performance and behaviour under failure, to the application programmer — some understanding of an algorithm will be required for the programmer to understand its robustness properties under, for example, failure of a site.

The need for clear understanding and easy experimentation with infrastructure algorithms, as well as the desire to simultaneously support multiple infrastructures on the same network, suggests a two-level architecture—a low-level consisting of a single set of well-understood, location-dependent primitives, in terms of which

a variety of high-level, location-independent communication abstractions may be expressed. This two-level approach enables one to have a standardized low-level runtime that is common to many machines, with divergent high-level facilities chosen and installed at run time.

For this approach to be realistic, it is essential that the low-level primitives should be directly implementable above standard network protocols. The IP network supports asynchronous, unordered, point-to-point, unreliable packet delivery; it abstracts from routing. We choose primitives that are directly implementable using asynchronous, unordered, point-to-point, reliable messages. This abstracts away from a multitude of additional details—error correction, retransmission, packet fragmentation, etc.—while still retaining a clear relationship to the well-understood IP level. It is also well suited to the process calculus presentation that we use below. More substantially, we also include migration of running computations among the low-level primitives. This requires substantial runtime support in individual network sites, but not sophisticated distributed algorithms—only one message need be sent per migration. By treating it as a low-level primitive we focus attention more sharply on the distributed algorithms supporting location-independent communication. We also provide low-level primitives for creation of running computations, for sending messages between computations at the same site, for generating globally unique names, and for local computation.

Many forms of high-level communication can be implemented in terms of these low-level primitives, for example synchronous and asynchronous message passing, remote procedure calls, multicasting to agent groups, etc. For this paper we consider only a single representative form: an asynchronous message-passing primitive, similar to the low-level primitive for communication between co-located computations, but independent of their locations, and transparent to migrations.

This two-level framework can be formulated cleanly using techniques from the theory of process calculi. We precisely define the low and high levels of abstraction as process calculi, the *Nomadic π calculi*, equipped with operational semantics and type systems. The overlay networks implementing the high level in terms of the low can then be treated rigorously as translations between these calculi. The semantics of the calculi provides a precise and clear understanding of the algorithms' behaviour, aiding design, and supporting proofs of correctness. Our calculi draw on ideas first developed in Milner, Parrow, and Walker's π calculus [Milner et al. 1992; Milner 1992] and extended in the *Pict* language of Pierce and Turner [Pierce and Turner 2000; Turner 1996], the distributed join-calculus of Fournet et al. [1996], and the JoCaml programming language [Conchon and Le Fessant 1999].

To facilitate experimentation, we designed and implemented a *Nomadic Pict* programming language based on our calculi. The low-level language extends the compiler and run-time system of Pict, a concurrent language based on the π calculus, to support our primitives for computation creation, migration, and location-dependent communication. High-level languages, with particular infrastructures for location-independent communication, can then be obtained by applying user-supplied translations into the low-level language. In both cases, the full language available to the user remains very close to the process calculus presentation, and can be given rigorous semantics in a similar style.

We begin in §2 by introducing the Nomadic π calculi, discussing their primitives and semantics, and giving examples of common programming idioms.

In §3 we present a first example overlay network, expressed as a semantics-preserving translation of the high-level Nomadic π calculus into the low-level calculus. This is a central forwarding server, relatively simple but still requiring subtle locking to ensure correctness.

In §4 we give a brief overview of the design space for such overlay networks, presenting a range of basic techniques and distributed algorithms informally, and discussing their scalability and fault tolerance properties with respect to possible applications.

For two of these more elaborate overlay algorithms, one using forwarding-pointer chains (broadly similar to that used in the JoCaml implementation) and one using query servers with caching, we give detailed definitions as Nomadic π calculus translations, in §5 and §6 (and in on-line Appendix C) respectively.

In §7 (with further details in on-line Appendices D, E, and F) we describe the design and implementation of the Nomadic Pict programming language, which lets us build executable distributed prototypes of these and many other overlay algorithms, together with simple example applications that make use of them.

We then return to semantics, to prove correctness of such overlay networks. In §8 we flesh out the semantic definition of the Nomadic π calculi and their basic metatheory: type preservation, safety, and correspondence between reduction and labelled transition semantics, and in §9 we develop operational reasoning techniques for stating and proving correctness. We:

- (1) extend the standard π calculus reduction and labelled transition semantics to deal with computation mobility, location-dependent communication, and a rich type system;
- (2) consider *translocating* versions of behavioural relations (bisimulation [Milner et al. 1992] and expansion [Sangiorgi and Milner 1992] relations) that are preserved by certain spontaneous migrations;
- (3) prove congruence properties of some of these, to allow compositional reasoning;
- (4) deal with partially committed choices, and hence state the main correctness result in terms of *coupled simulation* [Parrow and Sjödin 1992];
- (5) identify properties of agents that are *temporarily immobile*, waiting on a lock somewhere in the system; and,
- (6) as we are proving correctness of an encoding, we must analyse the possible reachable states of the encoding applied to an arbitrary high-level source program – introducing an intermediate language for expressing the key states, and factoring out as many ‘house-keeping’ reduction steps as possible.

We apply these to the Central Forwarding Server overlay of §3, describing a full correctness proof in Section 10. Finally, we discuss related work in §11 and conclude in §12.

This paper thus gives a synoptic view of the results of the Nomadic Pict project, covering calculi, semantics, overlay network design, programming language design and implementation, proof techniques, and overlay network verification. Elements

of this have previously appeared in conferences: the initial calculi of Sewell, Wojciechowski and Pierce [Sewell et al. 1998; 1999]; the programming language implementation and example algorithms by Wojciechowski and Sewell [Wojciechowski and Sewell 1999; Wojciechowski 2001; 2006]; and an outline of the metatheory and algorithm verification of Unyapoth and Sewell [2001]. Further details of the implementation and algorithms, and of the semantics and proof, can be found in the PhD theses of Wojciechowski and Unyapoth respectively [Wojciechowski 2000b; Unyapoth 2001]. The implementation is available on-line [Wojciechowski 2010].

Nomadic Pict was originally thought of in terms of computation mobility at the programming-language level, and the terminology of the body of the paper is chosen with that in mind (we speak of mobile agents and language threads). Later work on the Acute programming language [Sewell et al. 2007] developed this point of view: Acute has slightly lower-level constructs than low-level Nomadic Pict for checkpointing running multi-threaded computations, using which we built a small Acute library providing the low-level Nomadic Pict primitives; overlay-network implementations of the high-level Nomadic Pict abstraction could be expressed as ML-style modules above that. The underlying ideas may also be equally applicable to mobility at the virtual-machine OS image level, as we argued in a position paper [Sewell and Wojciechowski 2008] in the Joint HP-MSR Research Workshop on *The Rise and Rise of the Declarative Datacentre*.

2. THE NOMADIC π CALCULI

In this section we introduce the abstractions of the low- and high-level Nomadic π calculi.

The main entities are *sites* s and *agents* a . Sites represent physical machines or, more accurately, instantiations of the Nomadic Pict runtime system on physical machines; each site has a unique name.

Agents are units of running computation. Each agent has a unique name and a body consisting of some Nomadic Pict concurrent process P ; at any moment it is located at a particular site. An agent can *migrate*, at any point, to any other site (identified by name), new agents can be *created* (with the system synthesising a new unique name, bound to a lexically scoped identifier) and agents can *interact* by sending messages to each other.

A key point in the design of the low-level calculus is to make it easy to understand the behaviour of the system in the presence of partial failure. To do so, we choose interaction primitives that can be directly implemented above the real-world network (the Sockets API and TCP or UDP), without requiring a sophisticated distributed infrastructure. Our guiding principle is that each reduction step of the low-level calculus should be implementable using at most one inter-site asynchronous communication.

To provide an expressive language for local computation within each agent body, but keep the calculus concise, we include the constructs of a standard asynchronous π calculus. The Nomadic Pict concurrent process of an agent body can involve parallel composition, new channel creation, and asynchronous messaging on those channels within the agent.

In the rest of this section we give the syntax of processes, with accompanying

definitions of values, patterns, and types, and the key points of their reduction semantics. The full semantics is defined in Section 8 and on-line Appendices A and B.

2.1 Processes of the Low-Level Calculus

The syntax of the low-level calculus is as follows, grouped into the three agent primitives, two useful communication forms that are expressible as syntactic sugar, and the local asynchronous π calculus.

$P, Q ::=$	create ^{Z} $a = P$ in Q	spawn agent a with body P , on local site
	migrate to $s \rightarrow P$	migrate current agent to site s
	iflocal $\langle a \rangle c!v$ then P else Q	send $c!v$ to agent a if it is co-located here, and run P , otherwise run Q
.....		
	$\langle a \rangle c!v$	(sugar) send $c!v$ to agent a if it is co-located here
	$\langle a@s \rangle c!v$	(sugar) send $c!v$ to agent a if it is at site s
.....		
	0	empty process
	$P Q$	parallel composition of processes P and Q
	new $c : \sim^I T$ in P	declare a new channel c
	$c!v$	output of v on channel c in current agent
	$c?p \rightarrow P$	input on channel c in current agent
	*c?p $\rightarrow P$	replicated input
	if v then P else Q	conditional
	let $p = ev$ in P	local declaration

Executing the construct **create** ^{Z} $b = P$ **in** Q spawns a new agent, with body P , on the current site. After the creation, Q commences execution, in parallel with the rest of the body of the spawning agent. The new agent has a unique name which may be referred to with b , both in its body and in the spawning agent (b is binding in P and Q). The Z is a mobility capability, either **s**, requiring this agent to be static, or **m**, allowing it to be mobile. We return to this when we discuss the type system.

Agents can migrate to named sites: the execution of **migrate to** $s \rightarrow P$ as part of an agent results in the whole agent migrating to site s . After the migration, P commences execution in parallel with the rest of the body of the agent.

There is a single primitive for interaction between agents, allowing an atomic delivery of an asynchronous message between two agents that are co-located on the same site. The execution of **iflocal** $\langle a \rangle c!v$ **then** P **else** Q in the body of agent b has two possible outcomes. If the agent a is on the same site as agent b then the message $c!v$ will be delivered to a (where it may later interact with an input) and P will commence execution in parallel with the rest of the body of b ; otherwise the message will not be delivered and Q will execute as part of b . This is analogous to test-and-set operations in shared memory systems—delivering the message and starting P , or discarding it and starting Q , atomically. It can greatly simplify algorithms that involve communication with agents that may migrate away at any

time, yet is still implementable locally, by the runtime systems on a single site.

Two other useful constructs can be expressed as sugar: $\langle a \rangle c!v$ and $\langle a@s \rangle c!v$ attempt to deliver $c!v$ (an output of v on channel c), to agent a , on the current site and on s , respectively. They fail silently if a is not where it is expected to be, and so are usually used only in a context where a is predictable. The first is implementable simply as **iflocal** $\langle a \rangle c!v$ **then** **0** **else** **0**; the second as **create**^m $b =$ **migrate to** $s \rightarrow \langle a \rangle c!v$ **in** **0**, for a fresh name b that does not occur in s , a , c , or v .

Turning to the π calculus constructs, the body of an agent may be empty (**0**) or a parallel composition $P|Q$ of processes.

Execution of **new** $c : \tilde{^I}T$ **in** P creates a new unique channel name for carrying values of type T ; c is binding in P . The I is a capability: as in Pierce and Sangiorgi [1996], channels can be used for input only **r**, output only **w**, or both **rw**; these induce a subtype order.

An output $c!v$ (of value v on channel c) and an input $c?p \rightarrow P$ in the same agent may synchronise, resulting in P with the appropriate parts of the value v bound to the formal parameters in the pattern p . Note that, as in other asynchronous π calculi, outputs do not have continuation processes. A replicated input $*c?p \rightarrow P$ behaves similarly except that it persists after the synchronisation, and so might receive another value.

Finally, we have conditionals **if** v **then** P **else** Q , and local declarations **let** $p = ev$ **in** P , assigning the result of evaluating a simple value expression ev to a pattern p . In $c?p \rightarrow P$, $*c?p \rightarrow P$ and **let** $p = ev$ **in** P the names in pattern p are binding in P . If **if** v **then** P is sugar for **if** v **then** P **else** **0**.

For a simple example program in the low-level calculus, consider the following applet server.

$$\begin{aligned} & *getApplet?[a\ s] \rightarrow \\ & \quad \mathbf{create}^m\ b = \\ & \quad \quad \mathbf{migrate\ to}\ s \rightarrow \\ & \quad \quad \quad (\langle a@s' \rangle ack!b \mid B) \\ & \quad \mathbf{in}\ \mathbf{0} \end{aligned}$$

It can receive (on the channel named *getApplet*) requests for an applet. This is a replicated input ($*getApplet?[a\ s] \rightarrow \dots$) so the server persists and can repeatedly grant requests. The requests contain a pair (bound to the tuple $[a\ s]$ of a and s) consisting of the name of the requesting agent and the name of the site for the applet to go to. When a request is received the server creates an applet agent with a new name bound to b . This agent immediately migrates to site s . It then sends an acknowledgement to the requesting agent a (which here is assumed to be on site s') containing its name. In parallel, the body B of the applet commences execution.

2.2 Processes of the High-Level Calculus

The high-level calculus is obtained by extending the low-level language with a single location-independent communication primitive.

$$\begin{aligned} P ::= & \dots \\ & \mid \langle a@? \rangle c!v \quad \text{send } c!v \text{ to agent } a \text{ wherever it is} \end{aligned}$$

The intended semantics is that this will reliably deliver the message $c!v$ to agent a , irrespective of the current site of a and of any migrations. The high-level calculus includes all the low-level constructs, so those low-level communication primitives are also available for interaction with application agents whose locations are predictable. We write π_{LD} for the processes of the low-level calculus, with location-dependent communication only, and $\pi_{\text{LD,LI}}$ for the processes of the high-level calculus, with location-dependent and location-independent communication.

2.3 Values and Patterns

Channels allow the communication of first-order values: constants t , names x (including channel names c , agent names a , and site names s), tuples, and packages $\{T\}v$ of existential types, containing a witness type T and a value v . Patterns p are of similar shapes as values, but are subject to the condition that the names x and type variables X that they bind are all distinct.

$$\begin{aligned} v &::= t \mid x \mid [v_1 \dots v_n] \mid \{T\}v \\ p &::= _ \mid x \mid [p_1 \dots p_n] \mid \{X\}p \end{aligned}$$

The value grammar is extended with some basic functions, including equality tests, to give *expressions*, ranged over by ev .

2.4 Types

Typing infrastructure algorithms requires a moderately expressive type system. We take types

$T ::= B$	base type
$[T_1 \dots T_n]$	tuple
$\sim^I T$	channel name
$\{X\}T$	existential
X	type variable
Site	site name
Agent ^{Z}	agent name

where B might be `int`, `bool` etc., taken from a set \mathcal{T} of base types, and X is taken from a set \mathcal{TV} of type variables. Existentials are needed as an infrastructure must be able to forward messages of any type (see the `message` and `deliver` channels in Figure 2 later). For more precise typing, and to support the proof techniques we develop in §9, channel and agent types are refined by annotating them with *capabilities*, ranged over by I and Z respectively. Channel capabilities were described in §2.2: channels can be used for input only `r`, output only `w`, or both `rw`. In addition, agents are either static `s`, or mobile `m` [Sewell 1998; Cardelli et al. 1999].

2.5 Outline of the Reduction Semantics

Located Processes and Located Type Contexts The basic process terms given above only allow the source code of the body of a single agent to be expressed. During computation, this agent may evolve into a system of many agents, distributed over

many sites. To denote such systems, we define *located processes*

$$LP, LQ ::= @_a P \mid LP \mid LQ \mid \mathbf{new} \ x : \mathbf{Agent}^Z @_s \ \mathbf{in} \ LP \mid \mathbf{new} \ x : \sim^I T \ \mathbf{in} \ LP$$

Here the body of an agent a may be split into many parts, for example written $@_a P_1 \mid \dots \mid @_a P_n$. The construct $\mathbf{new} \ x : \mathbf{Agent}^Z @_s \ \mathbf{in} \ LP$ declares a new agent name x (binding in LP); since this is an agent name, we have an annotation $@_s$ giving the name s of the site where the agent is currently located. Channels, on the other hand, are not located – the construct $\mathbf{new} \ x : \sim^I T \ \mathbf{in} \ LP$ declares a new channel name (binding in LP) and the annotation is omitted.

Correspondingly, we add location information to type contexts. *Located type contexts* Γ include data specifying the site where each declared agent is located; the operational semantics updates this when agents move.

$$\Gamma, \Delta, \Phi ::= \bullet \mid \Gamma, X \mid \Gamma, x : \mathbf{Agent}^Z @_s \mid \Gamma, x : T \quad T \neq \mathbf{Agent}^Z$$

For example, the located type context below declares two sites, s and s' , and a channel c , which can be used for sending or receiving integers. It also declares a mobile agent a , located at s , and a static agent b , located at s' .

$$s : \mathbf{Site}, s' : \mathbf{Site}, c : \sim^{rw} \mathbf{Int}, a : \mathbf{Agent}^m @_s, b : \mathbf{Agent}^s @_{s'}$$

Pattern Matching When an input process receives a value v along a channel, it needs to deconstruct v , producing a substitution to be applied to its continuation process. As usual, this is done with an auxiliary partial function for *matching*, mapping pairs of patterns and values to name substitutions, whenever they are of the same shape. Its formal definition is given below.

$$\begin{aligned} \mathbf{match}(_, v) &\stackrel{\text{def}}{=} \{\} \\ \mathbf{match}(x, v) &\stackrel{\text{def}}{=} \{v/x\} \\ \mathbf{match}([p_1 \dots p_n], [v_1 \dots v_n]) &\stackrel{\text{def}}{=} \mathbf{match}(p_1, v_1) \cup \dots \cup \mathbf{match}(p_n, v_n) \\ \mathbf{match}(\{X\} p, \{T\} v) &\stackrel{\text{def}}{=} \{T/X\} \cup \mathbf{match}(p, v) \\ \mathbf{match}(p, v) &\stackrel{\text{def}}{=} \perp \text{ (undefined) } \quad \text{otherwise} \end{aligned}$$

Reductions To capture our informal understanding of the calculus in as lightweight a way as possible, we give a reduction semantics. It is defined with a structural congruence and reduction axioms, extending that for the π calculus [Milner 1993]. Reductions are over *configurations*, which are pairs $\Gamma \Vdash LP$ of a located type context Γ and a located process LP . We use a judgement $\Gamma \vdash a @_s$, meaning that an agent a is located at s in the located type context Γ . We shall give some examples of reductions, illustrating the new primitives, before giving the formal definition of reduction later, in Section 8 and Appendix B. The most interesting axioms for the low-level calculus are given in Figure 1.

An agent a can spawn a new mobile agent b , with body P , and continues with Q . The new agent is located at the same site as a (say s , with $\Gamma \vdash a @_s$). The agent b is initially bound and the scope is over the process Q in a and the whole of the new agent.

$$\begin{aligned} &\Gamma \Vdash @_a (R \mid \mathbf{create}^m b = P \ \mathbf{in} \ Q) \\ \rightarrow &\Gamma \Vdash @_a R \mid \mathbf{new} \ b : \mathbf{Agent}^m @_s \ \mathbf{in} \ (@_a Q \mid @_b P) \end{aligned}$$

$\Gamma \Vdash @_a \mathbf{create}^Z b = P \text{ in } Q$	$\rightarrow \Gamma \Vdash \mathbf{new } b : \mathbf{Agent}^Z @_s \text{ in } (@_b P \mid @_a Q)$	if $\Gamma \vdash a @_s$
$\Gamma \Vdash @_a \mathbf{migrate to } s \rightarrow P$	$\rightarrow (\Gamma \oplus a \mapsto s) \Vdash @_a P$	
$\Gamma \Vdash @_a (c!v \mid c?p \rightarrow P)$	$\rightarrow \Gamma \Vdash @_a \mathbf{match}(p, v) P$	
$\Gamma \Vdash @_a \mathbf{iflocal } \langle b \rangle c!v \text{ then } P \text{ else } Q$	$\rightarrow \Gamma \Vdash @_a P \mid @_b c!v$	if $\Gamma \vdash a @_s \wedge \Gamma \vdash b @_s$
$\Gamma \Vdash @_a \mathbf{iflocal } \langle b \rangle c!v \text{ then } P \text{ else } Q$	$\rightarrow \Gamma \Vdash @_a Q$	if $\Gamma \vdash a @_s \wedge \Gamma \vdash b @_{s'} \wedge s \neq s'$

Fig. 1. Selected Reduction Rules

When an agent a migrates to a new site s , we simply update the located type context.

$$\begin{aligned} & \Gamma \Vdash @_a (R \mid \mathbf{migrate to } s \rightarrow Q) \\ & \rightarrow \Gamma \oplus a \mapsto s \Vdash @_a (R \mid Q) \end{aligned}$$

A **new**-bound agent may also migrate; in this case, we simply update the location annotation.

$$\begin{aligned} & \Gamma \Vdash @_a R \mid \mathbf{new } b : \mathbf{Agent}^m @_{s'} \text{ in } @_b \mathbf{migrate to } s \rightarrow Q \\ & \rightarrow \Gamma \Vdash @_a R \mid \mathbf{new } b : \mathbf{Agent}^m @_s \text{ in } @_b Q \end{aligned}$$

An agent a may send a location-dependent message to an agent b if they are on the same site. The message, once delivered may then react with an input in b . Assuming that $\Gamma \vdash a @_s$ and $\Gamma \vdash b @_s$.

$$\begin{aligned} & \Gamma \Vdash @_a (\mathbf{iflocal } \langle b \rangle c![] \text{ then } P \text{ else } Q) \mid @_b (c?[] \rightarrow R) \\ & \rightarrow \Gamma \Vdash @_a P \mid @_b (c![] \mid c?[] \rightarrow R) \\ & \rightarrow \Gamma \Vdash @_a P \mid @_b R \end{aligned}$$

If a and b are at different sites, say if $\Gamma \vdash a @_s$ and $\Gamma \vdash b @_{s'}$ for $s \neq s'$, then the message will get lost.

$$\begin{aligned} & \Gamma \Vdash @_a (\mathbf{iflocal } \langle b \rangle c![] \text{ then } P \text{ else } Q) \mid @_b (c?[] \rightarrow R) \\ & \rightarrow \Gamma \Vdash @_a Q \mid @_b (c?[] \rightarrow R) \end{aligned}$$

Synchronisation of a local output $c!v$ and an input $c?x \rightarrow P$ only occurs within an agent, but in the execution of **iflocal** a new channel name can escape the agent where it was created, to be used elsewhere for output and/or input. Consider for example the process below, executing as the body of an agent a .

```

createm b =
  c?x → (x!3|x?n → 0)
in
  new d : rwint in
    iflocal ⟨b⟩ c!d then 0 else 0
    | d!7

```

It has a reduction for the creation of agent b , a reduction for the **iflocal** that delivers the output $c!d$ to b , and then a local synchronisation of this output with the input on c . Agent a then has body $d!7$ and agent b has body $d!3 \mid d?n \rightarrow 0$. Only the latter output on d can synchronise with b 's input $d?n \rightarrow 0$. For each channel name there is therefore effectively a π calculus-style channel in each agent. The channels are distinct, in that outputs and inputs can only interact if they are in the

same agent. This provides a limited form of dynamic binding, with the semantics of a channel name (i.e., the set of partners that a communication on that channel might synchronise with) dependent on the agent in which it is used; it proves very useful in the infrastructure algorithms that we develop.

The high-level calculus has one additional axiom, allowing location-independent communication between agents.

$$\Gamma \Vdash @_a \langle b@? \rangle c!v \rightarrow \Gamma \Vdash @_b c!v$$

This delivers the message $c!v$ to agent b irrespective of where b (and the sender a) are located. For example, below an empty-tuple message on channel c is delivered to an agent b with a waiting input on c .

$$\begin{aligned} & \Gamma \Vdash @_a (P \mid \langle b@? \rangle c![] \mid @_b (c?[] \rightarrow R)) \\ \rightarrow & \Gamma \Vdash @_a P \mid @_b (c![] \mid c?[] \rightarrow R) \end{aligned}$$

2.6 Discussion of Design Choices

The only inter-site communication required in an implementation of the low-level language is for the **migrate to** reduction, in which the body of the migrating agent a must be sent from its current site to site s . (For performance, one might also implement the location-dependent output $\langle a@s \rangle c!v$ directly, with a single inter-site message, rather than via the syntax desugaring into an agent creation and migration.)

This makes it easy to understand the behaviour of the implementation in the presence of fail-stop site failure: if a site crashes, all agents are lost; and a migration from one site to another is guaranteed to succeed if those two sites do not fail. Elsewhere we develop distributed infrastructure algorithms that address site failure and/or disconnection [Wojciechowski 2000b; 2001]. They use an additional primitive for timeouts, which we do not include in the semantics in this paper — our focus here is on the failure mode of message loss for location-dependent messages to agents that are not in the specified location.

One could also envisage extending the semantics with network topology information, so that link failure and network partitions could be modelled. As far as the operational semantics goes, that would be straightforward, but developing reasoning principles above the extended semantics would be a substantial task.

The inter-site messages that must be sent in an implementation (representations of migrating agents, and tuple-structured location-dependent messages) should be reliable in the face of intermittent network packet loss — our low-level semantics does not allow messages to be spontaneously discarded. They are also of unbounded size, and could often exceed the approx. 1500 bytes that can be sent in a UDP datagram over Ethernet without IP fragmentation. Hence, an implementation would send messages via TCP, not via UDP. This raises the question of whether the low-level calculus should guarantee that inter-site messages are received in the same order as they are sent. In favour, it would be easy to implement ordering guarantees, if all messages from one site to another are multiplexed on a single underlying TCP connection, and such guarantees may be useful for some distributed algorithms. Against this, the operational semantics would be much more complex, with queues of messages in the network, and reasoning principles above it would

be correspondingly more complex. Moreover, if the low-level calculus guaranteed message ordering, it would be natural for the high-level calculus to also guarantee it. Implementing that, as agents migrate, would require more complex algorithms. Accordingly, we choose simple unordered asynchronous messages, in both the low- and high-level calculus.

A similar argument applies to the question of whether inter-site messages should be asynchronous or synchronous. If they are implemented above TCP, the implementation could conceivably acknowledge when each message is delivered to the destination Nomadic Pict runtime. This would add a non-trivial but modest communication cost (especially if messages are often relatively large, involving multiple TCP segments). However, the added semantic complexity would be large, and efficient implementations of synchronous messaging in the high-level calculus, between migrating agents, would be yet more complex. Accordingly, we stay with the asynchronous choice.

Another design choice is whether one allows agents to be nested. This might be desirable for a full-scale programming language design, but again would complicate reasoning, and would introduce many further choices as to how inter-agent communication happens across the nesting structure. We therefore stay with the simple choice described above, in which new agents are created as siblings, on the same site as their creator.

3. EXAMPLE INFRASTRUCTURE: CENTRAL FORWARDING SERVER ALGORITHM

In this section we present our first example distributed infrastructure, the *Central Forwarding Server (CFS)* algorithm. In subsequent sections we survey the algorithm design space and present two more algorithms in detail: a forwarding-pointers algorithm and a query server algorithm. In the last part of the paper we develop semantic techniques and prove correctness of the CFS algorithm.

The problem that these algorithms solve is to implement the high-level calculus using the low-level primitives — specifically, to implement the high-level location-independent semantics

$$\Gamma \Vdash @_a \langle b@? \rangle c!v \rightarrow \Gamma \Vdash @_b c!v$$

that delivers a message to agent b irrespective of any migrations of agents a and b . To do so, they also use non-trivial implementations of the other high-level agent primitives, e.g. adding some synchronisations around agent migrations and creations. The algorithms are expressed as translations of the high-level calculus into the low-level calculus.

The CFS algorithm translation is based on that in Sewell et al. [1998]. It involves a central daemon that keeps track of the current sites of all agents and forwards any location-independent messages to them. The daemon itself is implemented as an agent which never migrates; the translation of a program then consists roughly of the daemon agent in parallel with a compositional translation of the program. When a new agent is created, it has to register with the daemon, telling its site. Before an agent can migrate, it has to inform the daemon about its intent, and wait for an acknowledgement. After the migration, the agent tells the daemon it has finished moving and continues. Locks are used to ensure that an agent




Fig. 2. The Central Server Daemon and the Interface Context

does not migrate away while a message forwarded by the daemon is on its way; this ensures that all messages forwarded from the daemon are delivered before the agent migrates away.

This is a relatively simple algorithm, rather sequential and with a centralized server daemon, but it still requires delicate synchronization that is easy to get wrong. Expressing it as a translation between well-defined low- and high-level languages provides a solid basis for discussion about design choices, and enables correctness proofs; the Nomadic Pict language implementation makes it possible to execute and use the algorithm in practice.

The daemon is implemented as a static agent; the translation $\mathcal{C}_\Phi \llbracket LP \rrbracket$ of a located process $LP = \mathbf{new} \Delta \mathbf{in} @_{a_1} P_1 \mid \dots \mid @_{a_n} P_n$ (well-typed with respect to a type context Φ) then consists roughly of the daemon agent in parallel with a compositional translation $\llbracket P_i \rrbracket_{a_i}$ of each source agent:

$$\mathcal{C}_\Phi \llbracket LP \rrbracket \stackrel{\text{def}}{=} \mathbf{new} \Delta, \Phi_{aux} \mathbf{in} \\
 @_D(\dots | \mathit{Daemon}) \\
 | \prod_{i \in \{1..n\}} @_{a_i}(\dots | \llbracket P_i \rrbracket_{a_i})$$

(we omit various initialisation code, and will often elide type contexts Φ). For each

$\llbracket \langle b@? \rangle c!v \rrbracket_a$	$\stackrel{\text{def}}{=} \langle D@SD \rangle \text{message!} \{T\} [b \ c \ v]$
$\llbracket \text{create}^Z b = P \text{ in } Q \rrbracket_a$	$\stackrel{\text{def}}{=} \text{currentloc?}s \rightarrow \text{new pack} : \text{ }^{\text{rw}}[], \text{ rack} : \text{ }^{\text{rw}}[] \text{ in}$ $\text{create}^Z b =$ $\langle D@SD \rangle \text{register!} [b \ s \ \text{rack}]$ $ \text{rack?}[] \rightarrow \text{iflocal} \langle a \rangle \text{pack!}[] \text{ then}$ $(\text{currentloc!}s \mid \llbracket P \rrbracket_b \mid \text{Deliverer})$ in $\text{pack?}[] \rightarrow (\text{currentloc!}s \mid \llbracket Q \rrbracket_a)$
	where $\text{Deliverer} \stackrel{\text{def}}{=} \text{*deliver?} \{X\} [c \ v \ \text{dack}] \rightarrow (\langle D@SD \rangle \text{dack!}[] \mid c!v)$
$\llbracket \text{migrate to } s \rightarrow P \rrbracket_a$	$\stackrel{\text{def}}{=} \text{currentloc?} _ \rightarrow \text{new mack} : \text{ }^{\text{rw}}[\text{ }^{\text{w}}[\text{Site } \text{ }^{\text{w}}[]]] \text{ in}$ $\langle D@SD \rangle \text{migrating!} [a \ \text{mack}]$ $ \text{mack?}[\text{migrated}] \rightarrow$ $\text{migrate to } s \rightarrow \text{new ack} : \text{ }^{\text{rw}}[] \text{ in}$ $(\langle D@SD \rangle \text{migrated!} [s \ \text{ack}]$ $ \text{ack?}[] \rightarrow \text{currentloc!}s \mid \llbracket P \rrbracket_a)$
$\llbracket 0 \rrbracket_a$	$\stackrel{\text{def}}{=} 0$
$\llbracket P \mid Q \rrbracket_a$	$\stackrel{\text{def}}{=} \llbracket P \rrbracket_a \mid \llbracket Q \rrbracket_a$
$\llbracket c?p \rightarrow P \rrbracket_a$	$\stackrel{\text{def}}{=} c?p \rightarrow \llbracket P \rrbracket_a$
$\llbracket *c?p \rightarrow P \rrbracket_a$	$\stackrel{\text{def}}{=} *c?p \rightarrow \llbracket P \rrbracket_a$
$\llbracket c!v \rrbracket_a$	$\stackrel{\text{def}}{=} c!v$
$\llbracket \text{iflocal} \langle b \rangle c!v \text{ then } P \text{ else } Q \rrbracket_a$	$\stackrel{\text{def}}{=} \text{iflocal} \langle b \rangle c!v \text{ then } \llbracket P \rrbracket_a \text{ then } \llbracket Q \rrbracket_a$
$\llbracket \text{new } x : \text{ }^I T \text{ in } P \rrbracket_a$	$\stackrel{\text{def}}{=} \text{new } x : \text{ }^I T \text{ in } \llbracket P \rrbracket_a$
$\llbracket \text{if } v \text{ then } P \text{ else } Q \rrbracket_a$	$\stackrel{\text{def}}{=} \text{if } v \text{ then } \llbracket P \rrbracket_a \text{ then } \llbracket Q \rrbracket_a$
$\llbracket \text{let } p = ev \text{ in } P \rrbracket_a$	$\stackrel{\text{def}}{=} \text{let } p = ev \text{ in } \llbracket P \rrbracket_a$

Fig. 3. The Compositional Encoding (selected clauses)

term P_i of the source language $n\pi_{\text{LD,LI}}$, considered as the body of an agent named a_i , the result $\llbracket P_i \rrbracket_{a_i}$ of the translation is a term of the target language $n\pi_{\text{LD}}$. The body of the daemon and selected clauses of the compositional translation are shown in Figures 2 and 3. They interact using channels of an *interface context* Φ_{aux} , also defined in Figure 2, which in addition declares lock channels and the daemon name D . It uses a map type constructor, which (together with the map operations) can be translated into the core language.

The original algorithm in Sewell et al. [1998] has been modified in the following ways to simplify the correctness proof.

- Type annotations have been added and checked with the Nomadic Pict type checker [Wojciechowski 2000b] (although this does not check the static/mobile subtyping).
- Fresh channels are used for transmitting acknowledgements, making such channels linear [Kobayashi et al. 1996]. This simplifies the proof of correctness, since

communication along a linear channel yields an expansion.

—We consider programs with many agents initiated separately on different sites, rather than only programs that are initiated as single agents (this more general translation is needed to make our co-inductive proof techniques go through, analogous to strengthening of an induction hypothesis).

The daemon consists of three replicated inputs, on the `message`, `register`, and `migrating` channels, ready to receive messages from the encodings of agents. It is at a fixed site SD . Part of the initialisation code places *Daemon* in parallel with an output on `lock` which carries a reference to a *site map*: a finite map from agent names to site names, recording the current site of every agent. Finite maps, with lookup operation

$$\begin{aligned} &\mathbf{lookup}[T_1 \ T_2] \ a \ \mathbf{in} \ m \ \mathbf{with} \\ &\quad \mathbf{found}(v) \rightarrow P \\ &\quad \mathbf{notfound} \rightarrow Q \end{aligned}$$

and update operation ($m \ \mathbf{with} \ a \mapsto v$), are expressed with a standard *pi* calculus encoding [Unyapoth 2001, §6.5], so they do not need to be added as a primitive.

The single-threaded nature of the daemon is ensured by using `lock` to enforce mutual exclusion between the three replicated inputs — each of them begins with an input on `lock`, thereby acquiring both the lock and the site map, and does not relinquish the lock until the daemon finishes with the request. The code preserves the invariant that at any time there is at most one output on `lock`.

Turning to the compositional translation $\llbracket \cdot \rrbracket$, it is defined by induction on type derivations. Only three clauses are non-trivial — for the location-independent output, agent creation and agent migration primitives. We discuss each one in turn, together with their interactions with the daemon. For the rest, $\llbracket \cdot \rrbracket$ is homomorphic.

Location-Independent Output A location-independent output $\langle b@? \rangle c!v$ in an agent a (of message $c!v$ to agent b) is implemented simply by requesting the central server daemon to deliver it; the request is sent to the daemon D , at its site SD , on its channel `message`, using a location-dependent output:


$$\llbracket \langle b@? \rangle c!v \rrbracket_a \stackrel{\text{def}}{=} \langle D@SD \rangle \mathbf{message!} \{T\} [b \ c \ v]$$

The corresponding replicated input on channel `message` in the daemon

$$\begin{aligned} &\mathbf{*message?} \{X\} [a \ c \ v] \rightarrow \\ &\quad \mathbf{lock?} m \rightarrow \\ &\quad \mathbf{lookup}[\mathbf{Agent}^s \ \mathbf{Site}] \ a \ \mathbf{in} \ m \ \mathbf{with} \\ &\quad \quad \mathbf{found}(s) \rightarrow \mathbf{new} \ \mathbf{dack} : \ ^{rw}[] \ \mathbf{in} \\ &\quad \quad \quad \langle a@s \rangle \mathbf{deliver!} \{X\} [c \ v \ \mathbf{dack}] \\ &\quad \quad \quad | \ \mathbf{dack?} [] \rightarrow \mathbf{lock!} m \\ &\quad \mathbf{notfound} \rightarrow 0 \end{aligned}$$

first acquires the lock and current site map m , then looks up the target agent's site in the map and sends a location-dependent message to the `deliver` channel of that agent; the message also carries the name of a freshly created channel `dack`. It then waits to receive an acknowledgment (on the `dack` channel) from the agent before

relinquishing the lock (with `lock!m`). This prevents the agent from migrating before the `deliver` message arrives, as the migration translation (below) also requires the lock. Note that the `notfound` branch of the lookup will never be taken, as the algorithm ensures that all agents register before messages can be sent to them. In each agent the `deliver` message is handled by a `Deliverer` process (see Figure 3), which reacts to `deliver` messages by emitting a local `c!v` message in parallel with sending the `dack` message to the daemon. The inter-agent communications involved in delivery of a single location-independent output are illustrated below.



Creation In order for the daemon’s site map to be kept up to date, agents must register with the daemon, telling it their site, both when they are created and when they migrate. Each agent records its current site internally as an output on its `currentloc` channel. This channel is also used as a lock, to enforce mutual exclusion between the encodings of all agent creation and migration commands within the body of the agent. The encoding of an agent creation in an agent a (in Figure 3)

$$\begin{aligned}
 \llbracket \text{create}^Z b = P \text{ in } Q \rrbracket_a &\stackrel{\text{def}}{=} \\
 &\text{currentloc?}s \rightarrow \text{new pack} : \text{^rw}[], \text{rack} : \text{^rw}[] \text{ in} \\
 &\quad \text{create}^Z b = \\
 &\quad \langle D@SD \rangle \text{register!}[b s \text{rack}] \\
 &\quad | \text{rack?}[] \rightarrow \text{iflocal} \langle a \rangle \text{pack!}[] \text{ then} \\
 &\quad \quad (\text{currentloc!}s | \llbracket P \rrbracket_b | \text{Deliverer}) \\
 &\text{in} \\
 &\quad \text{pack?}[] \rightarrow (\text{currentloc!}s | \llbracket Q \rrbracket_a) \\
 \\
 &\text{where Deliverer} \stackrel{\text{def}}{=} * \text{deliver?} \{X\} [c v \text{dack}] \rightarrow (\langle D@SD \rangle \text{dack!}[] | c!v)
 \end{aligned}$$

first acquires the local lock and current site s and then creates the new agent b , as well as channels `pack` and `rack`. The body of b sends a `register` message to the daemon, supplying `rack`; the daemon uses `rack` to acknowledge that it has updated its site map. After the acknowledgement is received from the daemon, b sends an acknowledgement to a using `pack`, initialises the local lock of b with s , installs a `Deliverer`, and allows the encoding of the body P of b to proceed. Meanwhile, the local lock of a and the encoding of the continuation process Q are blocked until the acknowledgement via `pack` is received.

The body of b is put in parallel with the replicated input


$$\mathbf{*deliver?} \{X\} [c \ v \ \mathit{dack}] \rightarrow (\langle D@SD \rangle \mathit{dack}! [] \mid c!v)$$

which will receive forwarded messages for channels in b from the daemon, send an acknowledgment back, and deliver the value locally to the appropriate channel.

The replicated input on **register** in the daemon

$$\begin{aligned} & \mid \mathbf{*register?} [b \ s \ \mathit{rack}] \rightarrow \\ & \quad \mathit{lock?} m \rightarrow \\ & \quad \quad \mathbf{let} [\mathit{Agent}^s \ \mathit{Site}] \ m' = (m \ \mathbf{with} \ b \mapsto s) \ \mathbf{in} \\ & \quad \quad (\mathit{lock}! m' \mid \langle b@s \rangle \mathit{rack}! []) \end{aligned}$$

first acquires the lock and current site map, replaces the site map with an updated map, thereby relinquishing the lock, and sends an acknowledgment to the registering agent; the updated map records that a new agent b is located at site s . The inter-agent communications involved in a single agent creation are illustrated below.



Migration The encoding of a **migrate to** in agent a

$$\begin{aligned} \llbracket \mathbf{migrate\ to} \ s \rightarrow P \rrbracket_a & \stackrel{\text{def}}{=} \\ & \mathit{currentloc?} _ \rightarrow \mathbf{new} \ \mathit{mack} : \hat{\text{rw}} [\hat{\text{w}} [\mathit{Site} \ \hat{\text{w}} []]] \ \mathbf{in} \\ & \quad \langle D@SD \rangle \mathit{migrating}! [a \ \mathit{mack}] \\ & \quad \mid \mathit{mack?} [\mathit{migrated}] \rightarrow \\ & \quad \quad \mathbf{migrate\ to} \ s \rightarrow \mathbf{new} \ \mathit{ack} : \hat{\text{rw}} [] \ \mathbf{in} \\ & \quad \quad (\langle D@SD \rangle \mathit{migrated}! [s \ \mathit{ack}] \\ & \quad \quad \mid \mathit{ack?} [] \rightarrow \mathit{currentloc}! s \mid \llbracket P \rrbracket_a) \end{aligned}$$

first acquires the output on **currentloc** at a (discarding the current site data). It then creates a fresh channel **mack**, sends a **migrating** message to the daemon with a tuple $[a \ \mathit{mack}]$, and waits for an acknowledgement on **mack**.

Reacting to the message on migrating message, the daemon

```


| *migrating?[a mack] →
  lock?m →
    lookup[Agents Site] a in m with
      found(s) → new migrated : ^rw[Site ^w[]] in
        ⟨a@s⟩mack![migrated]
        | migrated?[s' ack]
          let m' = (m with a ↦ s') in
            (lock!m' | ⟨a@s'⟩ack![])
          notfound → 0

```

acquires its lock, looks up the current site of a in the acquired map m , creates a fresh channel `migrated`, and sends it (using an LD primitive) to a along channel `mack`. The daemon then waits to receive a message from `migrated`.

Once the waiting agent a receives a message from `mack`, it migrates to the new site s , then creates a fresh channel `ack` and sends a tuple $[s \text{ ack}]$ to the daemon via channel `migrated` (using an LD primitive). Meanwhile, the local lock and the encoding of the continuation process P is kept until the acknowledgement via `ack` is received from the daemon.

When the blocked daemon receives a message on `migrated`, it updates the site map, then relinquishes the lock and then sends an acknowledgement to a at its new site. The inter-agent communications involved in the migration of a single agent are illustrated below.



4. ALGORITHM DESIGN SPACE

Prospective applications may use some form of mobility for many different purposes, for example: to improve locality of computation; to support disconnected operation on mobile devices; to avoid transferring large volumes of data; to facilitate fault-tolerance by moving computation from partially faulty machines; or to adapt to changes in the network characteristics and in the user environment. The different applications may have very different patterns of agent migration and communication, and require different performance and robustness properties. Agent migration would often be limited, e.g. to cases where agents migrate only once or twice, where migration is within a local-area network or between a few sites which

are known in advance, where agents can only migrate to or from a central site, and between a mobile computer and the network, and so on.

In this section, we characterise some basic techniques and algorithms that can be useful for building such application-specific infrastructures, and assess their usefulness. We do not attempt to specify all the algorithms formally, so we use natural language descriptions. However, almost all algorithms have been implemented in Nomadic Pict, and the code is available with the language distribution [Wojciechowski 2010]. We also discuss informally the scalability and fault-tolerance properties of the algorithms. We do not attempt to give quantitative theoretical or empirical characterisations of the algorithms, because it would be too hard to take under consideration all the factors which exist in real systems — the range of possible migration and communication patterns is too great.

In the following sections, we describe two algorithms in more detail, presenting complete executable descriptions of the infrastructure in Nomadic Pict. They eliminate some of the drawbacks of the CFS algorithm in Section 3.

4.1 Background

We first discuss the space of all (deterministic) algorithms for location-independent message delivery to migrating entities. Awerbuch and Peleg [1995] (see also Mullender and Vitányi [1988]) stated the analogous problem of keeping track of mobile users in a distributed network. They consider two operations: “move”, for a move of a user to a new destination, and “find”, enabling one to contact a specified user at its current address. The problems of minimizing the communication overhead of these two operations appear to be in conflict. They examined two extreme strategies: full information and no information.

The *full-information* strategy requires every site in the network to maintain complete up-to-date information on the whereabouts of every user. This makes the “find” operation cheap. On the other hand, “move” operations are very expensive, since it is necessary to update information at every site. In contrast, the *no-information* approach does not assume any updates while migrating, thus the “move” operation has got a null cost. On the other hand, the “find” operation is very expensive because it requires global searching over the whole network. However, if a network is small and migrations frequent, the strategy can be useful. In contrary, the *full-information* strategy is appropriate for a near-static setting, where agents migrate relatively rarely, but frequently communicate with each other. Between these two extreme cases, there is space for designing intermediate strategies, that will perform well for any or some specific communication to migration pattern, making the costs of both “find” and “move” operations relatively cheap.

Awerbuch and Peleg [1995] describe a distributed directory infrastructure for online tracking of mobile users. They introduced the graph-theoretic concept of *regional matching*, and demonstrated how finding a regional matching with certain parameters enables efficient tracking of mobile users in a distributed network. The communication overhead of maintaining the distributed directory is within a polylogarithmic factor of the lower bound. This result is important in the case of mobile telephony and infrastructures which support mobile devices, where the infrastructure should perform well, considering *all* mobile users and their potential communication to migration patterns. These patterns can vary, depending on peo-

ple, and can only be estimated probabilistically. The infrastructure should therefore support all migration and communication scenarios, and optimise those scenarios which are likely to happen more often (preferably it should adapt to any changes in behaviour of mobile users dynamically). In mobile agent applications, however, the communication to migration pattern of mobile agents usually can be predicted precisely [Wojciechowski 2000b]. Therefore we can design algorithms which are optimal for these special cases and simpler than the directory infrastructure mentioned above.

4.2 Central Server Algorithms

Central Forwarding Server Algorithm The server records the current site of every agent. Before migration an agent A informs the server and waits for ACK (containing the number of messages sent from the server to A). It then waits for all the messages due to arrive. After migration it tells the server it has finished moving. If B wants to send a message to A , B sends the message to the server, which forwards it. During migrations (after sending the ACK) the server suspends forwarding. A variant of this algorithm was described in Section 3.

Central Query Server Algorithm The server records the current site of every agent. If B wants to send a message to A , B sends a query (containing the message ID) to the server asking for the current site of A , gets the current site s of A and sends the message to s . The name s can be used again for direct communication with A . If a message arrives at a site that does not have the recipient then a message is returned saying ‘you have to ask the name server again’. Migration support is similar as above.

Home Server Algorithm Each site s has a server (one of the above) that records the current site of some agents — usually the agents which were created on s . Agent names contain an address of the server which maintains their locations. On every migration agent A synchronises with the server whose name is part of A ’s name. If B wants to send a message to A , B resolves A ’s name and contacts A ’s server. Other details are as above.

Discussion If migrations are rare, and also in the case of stream communication or large messages, the Query Server seems the better choice. However, the Central Forwarding and Query Server algorithms do not scale well. If the number of agents is growing and communication and migration are frequent, the server can be a bottleneck. Home Servers can improve the situation. The infrastructure can work fine for small-to-medium systems, where the number of agents is small.

These algorithms do not support locality of agent migration and communication, i.e. migration and communication involve the cost of contacting the server, which might be far away. If agents are close to the server, the cost of migration, search, and update is relatively low. Our algorithms clearly explore only a part of the design space — one can envisage e.g. splitting the servers into many parts (e.g. one dealing with agents created for each user). An exhaustive discussion is beyond the scope of this paper.

In all three, the server is a single point of failure. In these and other algorithms, we can use some of the classical techniques of fault-tolerance, e.g. based

on state checkpointing, message logging and recovery. We can also replicate the server on different sites to enhance system availability and fault-tolerance, e.g. using the *primary-backup* or *active* replication techniques (see, e.g., the tutorial [Guerraoui and Schiper 1996]). However, the implementation of a replicated system with replica crashes and unpredictable communication delay is a difficult task. The difficulty can be formally explained by theoretical impossibility results, such as the impossibility of solving consensus in an asynchronous system when processes can crash [Fischer et al. 1985]. These impossibility results can be overcome by strengthening the system model slightly [Chandra and Toueg 1996].

Mechanisms similar to Home Servers have been used in many systems which support *process migration*, such as Sprite [Douglis and Ousterhout 1991]. Caching has been used, e.g. in LOCUS [Popek and Walker 1986], and V [Cheriton 1988], allowing operations to be sent directly to a remote process without passing through another site. If the cached address is wrong a home site of the process is contacted (LOCUS) or multicasting is performed (V). A variant of the Central Query Server algorithm, combined with Central Forwarding Server and data caching, will be described in detail in Section 6 and Appendix C; it also appeared in Wojciechowski and Sewell [2000].

4.3 Forwarding Pointers

Algorithm There is a forwarding daemon on each site. The daemon on site s maintains a current guess about the site of agents which migrated from s . Every agent knows the initial home site of every agent (the address is part of an agent's name). If A wants to migrate from s_1 to s_2 it leaves a forwarding pointer at the local daemon. Communications follow all the forwarding pointers. If there is no pointer to agent A , A 's home site is contacted. Forwarding pointers are preserved forever. This algorithm will be described in detail in Section 5.

Discussion There is no synchronisation between migration and communication as there was in centralised algorithms. A message may follow an agent which frequently migrates, leading to a race condition. The Forwarding Pointers algorithm is not practical if agents perform a large number of migrations to distinct sites (the chain of pointers grows, increasing the cost of search). Some “compaction” methods can be used to collapse the chain, e.g. movement-based and search-based. In the former case, an agent would send backward a location update after performing a number of migrations; in the latter case, after receiving a number of messages (i.e. after a fixed number of “find” operations occurred).

Some heuristics can be further used such as search-update. A plausible algorithm can be as follows. On each site there is a daemon which maintains forwarding addresses (additionally to forwarding pointers) for all agents which ever visited this site. A *forwarding address* is a tuple (*timestamp, site*) in which the site is the last known location of the agent and timestamp specifies the age of the forwarding address. Every message sent from agent B to A along the chain of forwarding pointers contains the latest available forwarding address of A . The receiving site may then update its forwarding address (and/or forwarding pointer) for the referenced agent, if required. Given conflicting guesses for the same agent, it is simple to determine which one is most recent using timestamps. When the message is eventually deliv-

ered to the current site of the agent, the daemon on this site will send an ACK to the daemon on the sender site, containing the current forwarding address. The address received replaces any older forwarding address but *not* the forwarding pointer (to allow updating the chain of pointers during any subsequent communication). A similar algorithm has been used in Emerald [Jul et al. 1988], where the new forwarding address is piggybacked onto the reply message in the object invocation. It is sufficient to maintain the timestamp as a counter, incremented every time the object moves.

A single site fail-stop in a chain of forwarding pointers breaks the chain. A solution is to replicate the location information in the chain on k consecutive sites, so that the algorithm is tolerant of a failure of up to $k - 1$ adjoint sites. Stale pointers should be eventually removed, either after waiting a sufficiently long time, or purged as a result of a distributed garbage collection. Distributed garbage collection would require detecting global termination of all agents that might ever use the pointer, therefore the technique may not always be practically useful. Alternatively, some weaker assumptions could be made and the agents decide arbitrarily about termination, purging the pointers beforehand.

4.4 Broadcast Algorithms

Data Broadcast Algorithm Sites know about the agents that are currently present. An agent notifies a site on leaving and a forwarding pointer is left over until agent migration is finished. If agent B wants to send a message to A , B sends the message to all sites in a network. A site s discards or forwards the message if A is not at s (we omit details).

Query Broadcast Algorithm As above but if agent B wants to send a message to A , B sends a query to all sites in a network asking for the current location of A . If site s receives the query and A is present at site s , then s suspends any migration of A until A receives the message from B . A site s discards or forwards the query if A is not at s .

Notification Broadcast Algorithm Every site in a network maintains a current guess about agent locations. After migration an agent distributes in the network information about its new location. Location information is time-stamped. Messages with stale location information are discarded. If site s receives a message whose recipient is not at s (because it has already migrated or the initial guess was wrong), it waits for information about the agent's new location. Then s forwards the message.

Discussion The cost of communication in Query and Data Broadcasts is high (packets are broadcast in the network) but the cost of migration is low. Query Broadcast saves bandwidth if messages are large or in the case of stream communication. Notification Broadcast has a high cost of migration (the location message is broadcast to all sites) but the communication cost is low and similar to forwarding pointers with pointer chain compaction. In Data and Notification Broadcasts, migration can be fast because there is no synchronisation involved (in Query Broadcast migration is synchronised with communication); the drawback is a potential for race conditions if migrations are frequent. Site failures do not disturb the algorithms. The simplest (partially) fault-tolerant algorithm could involve Data

Broadcast with buffering of broadcast messages at target sites; however, two conditions should hold: buffers need to be infinite, and the broadcasting server needs to use *reliable broadcast* [Chandra and Toueg 1996].

Although we usually assume that the number of sites is too large to broadcast anything, we may allow occasional broadcasts within, e.g. a local Internet domain, or local Ethernet. Broadcasts can be accomplished efficiently in bus-based multiprocessor systems. They are also used in radio networks. A realistic variant is to broadcast within a group of sites which belong to the itinerary of mobile agents that is known in advance. Broadcast has also been used in Emerald to find an object, if a node specified by a forwarding pointer is unreachable or has stale data. To reduce message traffic, only a site which has the specified object responds to the broadcast. If the searching daemon receives no response within a time limit, it sends a second broadcast requesting a positive or negative reply from all other sites. All sites not responding within a short time are sent a reliable, point-to-point message with the request. The Jini lookup and connection infrastructure [Arnold et al. 1999] uses multicast in the discovery protocol. A client wishing to find a Lookup Service sends out a known packet via multicast. Any Lookup Service receiving this packet will reply (to an address contained in the packet) with an implementation of the interface to the Lookup Service itself.

4.5 Agent-based Broadcast

Algorithm Agents are *grouped*, with the agents forming a group maintaining a current record about the site of every agent in the group. Agent names form a totally ordered set. We assume communication which takes place within a group only.

Before migration an agent A informs the other agents in the group about its intention and waits for ACKs (containing the number of messages sent to A). It then waits for all the messages due to arrive and migrates. After migration it tells the agents it has finished moving. Multicast messages to each agent within a group are reliably delivered in the order sent (using *first-in-first-out broadcast*). If B wants to send a message to A , B sends the message to site s which is A 's current location. During A 's migrations (i.e. after sending the ACK to A) B suspends sending any messages to A (in particular any migration requests). If two (or more) agents want to migrate at the same time there is a conflict which can be resolved as follows. Suppose A and C want to migrate. If B receives migration requests from A and C , it sends ACKs to both of them and suspends sending any messages to agents A and C (in particular any migration requests). If A receives a migration request from C after it has sent its own migration request it can either grant ACK to C (and C can migrate) or postpone the ACK until it has finished moving to a new site. The choice is made possible by ordering agent names. Thus, there is an invariant that at any time at most one agent can migrate in a given group.

Discussion The advantage of this algorithm is that sites can be stateless (the location data are part of agent's state). The algorithm is suitable for frequent messages (or stream communication) between mobile agents and when migrations are rare.

Agents can be organised into *dynamic groups*, using the primitives of *group communication systems* (designed for non-movable groups of distributed processes).

The membership of a group can change over time, as agents *join* or *leave* the group, or as crashed (or suspected as crashed) agents are collectively *removed* from the group. The current set of agents that are members of a group is called the *group view*. Agents are added to and deleted from the group view via *view changes*, handled by a *membership* service.

Mobile agents forming a group can dynamically change sites. This creates a problem how to implement the join operation so that the agents joining a group will be able to localize the group. One solution is that migrating group agents could leave forwarding pointers that would be followed by agents joining the group to “catch up” with at least one group member. Another solution is to have one agent within a group—a *group coordinator*, which never migrates and can be used to contact the group. The inter-group communication algorithm could use either the pointers or coordination agents for delivering messages that cross group boundaries.

Other variants are also possible. For example, if agent migration would be limited to a fixed set of target sites that are known in advance, then the algorithms could broadcast only to such sites; the names of these sites could be encoded as part agent’s name.

4.6 Hierarchical Location Directory

Algorithm A tree-like hierarchy of servers forms a location directory (similar to DNS). Each server in the directory maintains a current guess about the site of some agents. Sites belong to regions, each region corresponds to a sub-tree in the directory (in the extreme cases the sub-tree is simply a leaf-server for the smallest region, or the whole tree for the entire network). The algorithm maintains an invariant that for each agent there is a unique path of forwarding pointers which forms a single branch in the directory; the branch starts from the root and finishes at the server which knows the actual site of the agent (we call this server the “nearest”). Before migration an agent A informs the “nearest” server X_1 and waits for ACK. After migration it registers at a new “nearest” server X_2 , tells X_1 it has finished moving and waits for ACK. When it gets the ACK there is already a new path installed in the tree (this may require installing new and purging old pointers within the smallest sub-tree which contains X_1 and X_2). Messages to agents are forwarded along the tree branches. If B wants to send a message to A , B sends the message to the B ’s “nearest” server, which forwards it in the directory. If there is no pointer the server will send the message to its parent.

Discussion Certain optimisations are plausible, e.g. if an agent migrates very often within some sub-tree, only the root of the sub-tree would contain the current location of the agent (the cost of a “move” operation would be cheaper). Moreau [2002] describes an algorithm for routing messages to migrating agents which is also based on distributed directory service. A proposal of Globe uses a hierarchical location service for worldwide distributed objects [van Steen et al. 1998]. The Hierarchical Location Directory scales better than Forwarding Pointers and Central Servers. Also, some kinds of fault can be handled more easily (see Awerbuch and Peleg [1995], and there is also a lightweight crash recovery in the Globe system [Ballintijn et al. 1999]).

4.7 Arrow Directory

Some algorithms can be devised for a particular communication pattern. For example, if agents do not require instant messaging, a simple mail-box infrastructure can be used, where senders send messages to static mailboxes and all agents periodically check mailboxes for incoming messages.

Demmer and Herlihy [1998] describe the Arrow Distributed Directory protocol for distributed shared object systems. The algorithm is devised for a particular object migration pattern; it assumes that the whole object is always sent to the object requester. The arrow directory imposes an optimal distributed queue of object requests, with no point of bottleneck.

The protocol was motivated by emerging *active network* technology, in which programmable network switches are used to implement customized protocols, such as application-specific packet routing.

Algorithm The arrow directory is given by a minimum spanning tree for a network, where the network is modelled as a connected graph. Each vertex models a node (site), and each edge a reliable communication link. A node can send messages directly to its neighbours, and indirectly to non-neighbours along a path. The directory tree is initialised so that following arrows (pointers) from any node leads to the node where the object resides.

When a node wants to acquire exclusive access to the object, it sends a message *find* which is forwarded via arrows and sets its own arrow to itself. When the other node receives the message, it immediately “flips” the arrow to point back to the immediate neighbour who forwarded the message. If the node does not hold the object, it forwards the message. Otherwise, it buffers the message *find* until it is ready to release the object to the object requester. The node releases the object by sending it directly to the requester, without further interaction with the directory.

If two *find* messages are issued at about the same time, one will eventually cross the other’s path and be “diverted” away from the object, following arrows towards the node (say v) where the other *find* message was issued. Then, the message will be blocked at v until the object reaches v , is accessed and eventually released.

5. EXAMPLE INFRASTRUCTURE: FORWARDING-POINTERS ALGORITHM

In this section we give a forwarding-pointers algorithm, in which daemons on each site maintain chains of forwarding pointers for agents that have migrated from their site. It removes the single bottleneck of the centralised-server solution in Section 3; it is thus a step closer to algorithms that may be of wide practical use. The algorithm is more delicate, so expressing it as a translation provides a more rigorous test of the framework.

The daemons are implemented as static agents; the translation $\mathcal{FP}_\Phi \llbracket LP \rrbracket$ of a located process $LP = \mathbf{new} \Delta \mathbf{in} @_{a_1} P_1 \mid \dots \mid @_{a_n} P_n$, (well-typed with respect to Φ) then consists roughly of the daemon agent (one on each site s_j , named DS_j) in parallel with a compositional translation $\llbracket P_i \rrbracket_{a_i}$ of each source agent:

$$\mathcal{FP}_\Phi \llbracket LP \rrbracket \stackrel{\text{def}}{=} \mathbf{new} \Delta, \Phi_{aux} \mathbf{in} \\ @_{DS_1} (Daemon_{s_1} \mid \mathbf{lock!}m) \mid \dots \mid @_{DS_m} (Daemon_{s_m} \mid \mathbf{lock!}m) \\ \mid @_{a_1} \llbracket P_1 \rrbracket_{a_1} \mid \dots \mid @_{a_n} \llbracket P_n \rrbracket_{a_n}$$

```

Daemons
def let [S DS] = s in
  *register?[B rack] → lock?m →
    lookup[Agents ~rw[Site Agents]] B in m with
      found(Bstate) →
        Bstate?[_-] →
          Bstate![S DS] | lock!m | ⟨B⟩rack![]
      notfound →
        new Bstate : ~rw[Site Agents] in
          Bstate![S DS] | ⟨B⟩rack![]
          | let[Agents [Site Agents]] m' = (m with B ↦ Bstate) in
            lock!m'
  | *migrating?[B mack] → lock?m →
    lookup[Agents ~rw[Site Agents]] B in m with
      found(Bstate) →
        Bstate?[_-] → (lock!m | ⟨B⟩mack![])
      notfound → 0
  | *migrated?[B [U DU] ack] → lock?m →
    lookup[Agents ~rw[Site Agents]] B in m with
      found(Bstate) →
        lock!m | ⟨B@U⟩ack![] | Bstate![U DU]
      notfound → 0
  | *message? {X} [[B U DU] c v] → lock?m →
    lookup[Agents ~rw[Site Agents]] B in m with
      found(Bstate) →
        lock!m
        | Bstate?[R DR] →
          iflocal ⟨B⟩c!v then Bstate![R DR]
          else ⟨DR@R⟩message! {X} [[B U DU] c v]
          | Bstate![R DR]
      notfound → lock!m
        | ⟨DU@U⟩message! {X} [[B U DU] c v]

Φaux def DS1 : Agents@s1, ..., DSm : Agents@sm,
  lock : ~rwMap[Agents ~rw[Site Agents]]
  register : ~rw[Agents ~w[]],
  migrating : ~rw[Agents ~w[]],
  migrated : ~rw[Agents [Site Agents] ~w[]],
  message : ~rw {X} [[Agents Site Agents] ~wX X],
  currentloc : ~rw[Site Agents]

```

Fig. 4. A Forwarding-Pointers Translation: the Daemon

where m is a map such that $m(a) = [s_j DS_j]$ if $\Phi, \Delta \vdash a@s_j$. For each term P_i of the source language $\text{n}\pi_{\text{LD,LI}}$, considered as the body of an agent named a_i , the result $\llbracket P_i \rrbracket_{a_i}$ of the translation is a term of the target language $\text{n}\pi_{\text{LD}}$. As before, the translation consists of a compositional encoding of the bodies of agents, given in Figure 5, and daemons, defined in Figure 4. Note that in terms of the target language, each site name s_i is re-bound to the pair $[s_i DS_i]$ of the site name together with the respective daemon name; the agent name a_i is re-bound to the triple $[A_i s_i DS_i]$ of the low-level agent name A_i together with the initial site and daemon names. The low-level agent A_i is defined by the agent encoding; it contains

```

[[⟨b@?⟩c!v]]A
 $\stackrel{\text{def}}{=} \text{currentloc?}[S DS] \rightarrow$ 
  iflocal ⟨DS⟩message! {T}[b c v]
    then currentloc![S DS]
    else currentloc![S DS]

[[createZ b = P in Q]]A
 $\stackrel{\text{def}}{=} \text{currentloc?}[S DS] \rightarrow$ 
  new pack : rw[], rack : rw[] in
  createZ B =
    let b = [B S DS] in
    ⟨DS⟩register![B rack]
    | rack?[]  $\rightarrow$  iflocal ⟨A⟩pack![] then
      currentloc![S DS] | [[P]]B
  in
  let b = [B S DS] in
  pack?[]  $\rightarrow$  (currentloc![S DS] | [[Q]]A)

[[migrate to s  $\rightarrow$  P]]A
 $\stackrel{\text{def}}{=} \text{currentloc?}[S DS] \rightarrow$ 
  let [U DU] = s in
  if [S DS] = [U DU] then
    currentloc![U DU] | [[P]]A
  else
    new mack : rw[] in
    ⟨DS⟩migrating![A mack]
    | mack?[]  $\rightarrow$  migrate to U  $\rightarrow$ 
    new rack : rw[] in
    ⟨DU⟩register![A rack]
    | rack?[]  $\rightarrow$  new ack : rw[] in
    ⟨DS@S⟩migrated![A [U DU] ack]
    | ack?[]  $\rightarrow$  (currentloc!s | [[P]]A)

[[iflocal (b)c!v then P else Q]]A
 $\stackrel{\text{def}}{=} \text{let } [B \_ \_] = b \text{ in}$ 
  iflocal ⟨B⟩c!v then [[P]]A else [[Q]]A

```

Fig. 5. A Forwarding-Pointers Translation: the Compositional Encoding (selected clauses)

the body P_i of agent a_i . Agents and daemons interact using channels of an *interface context* Φ_{aux} , also defined in Figure 4, which in addition declares lock channels and the daemon names $DS_1 \dots DS_m$. It uses a map type constructor, which (together with the map operations) can be translated into the core language.

Daemons are created, one on each site. These will each maintain a collection of forwarding pointers for all agents that have migrated away from their site. To keep the pointers current, agents synchronize with their local daemons on creation and migration. Location independent communications are implemented via the daemons, using the forwarding pointers where possible. If a daemon has no pointer for the destination agent of a message then it will forward the message to the daemon on the site where the destination agent was created; to make this possible an agent name is encoded by a triple of an agent name and the site and daemon

of its creation. Similarly, a site name is encoded by a pair of a site name and the daemon name for that site. There is a translation of types with clauses

$$\begin{aligned} \llbracket \text{Agent}^Z \rrbracket &\stackrel{\text{def}}{=} [\text{Agent}^Z \text{ Site Agent}^Z] \\ \llbracket \text{Site} \rrbracket &\stackrel{\text{def}}{=} [\text{Site Agent}^Z] \end{aligned}$$

We generally use lower case letters for site and agent names occurring in the source program and upper case letters for sites and agents introduced by its encoding.

Looking first at the compositional encoding, in Figure 5, each agent uses a `currentloc` channel as a lock, as before. It is now also used to store both the site where the agent is and the name of the daemon on that site. The three interesting clauses of the encoding, for location-independent output, creation, and migration, each begin with an input on `currentloc`. They are broadly similar to those of the simple Central-Forwarding-Server translation in Section 3.


Turning to the body of a daemon, defined in Figure 4, it is parametric in a pair s of the name of the site S where it is and the daemon's own name DS . It has four replicated inputs, on its `register`, `migrating`, `migrated`, and `message` channels. Some partial mutual exclusion between the bodies of these inputs is enforced by using the `lock` channel. The data stored on the `lock` channel now maps the name of each agent that has ever been on this site to a lock channel (e.g. `Bstate`) for that agent. These agent locks prevent the daemon from attempting to forward messages to agents that may be migrating. Each stores the site and daemon (of that site) where the agent was last seen by this daemon — i.e. either this site/daemon, or the site/daemon to which it migrated from here. The use of agent locks makes this algorithm rather more concurrent than the previous one — rather than simply sequentialising the entire daemon, it allows daemons to process inputs while agents are migrating, so many agents can be migrating away from the same site, concurrently with each other and with delivery of messages to other agents at the site.

Location-independent output A location-independent output $\langle b@? \rangle c!v$ in agent A is implemented by requesting the local daemon to deliver it. (Note that A cannot migrate away before the request is sent to the daemon and a lock on `currentloc` is released.)

The `message` replicated input of the daemon gets the map m , from agent names to agent lock channels. If the destination agent B is not found, the message is forwarded to the daemon DU on the site U where B was created. Otherwise, if B is found, the agent lock `Bstate` is grabbed, obtaining the forwarding pointer $[R DR]$ for B . Using `iflocal`, the message is then either delivered to B , if it is here, or to the daemon DR , otherwise. Note that the `lock` is released before the agent lock is requested, so the daemon can process other inputs even if B is currently migrating; it also prevents deadlock. In particular, in order to complete any migration of B the daemon should be able to process message `migrated` that requires to acquire `lock`.

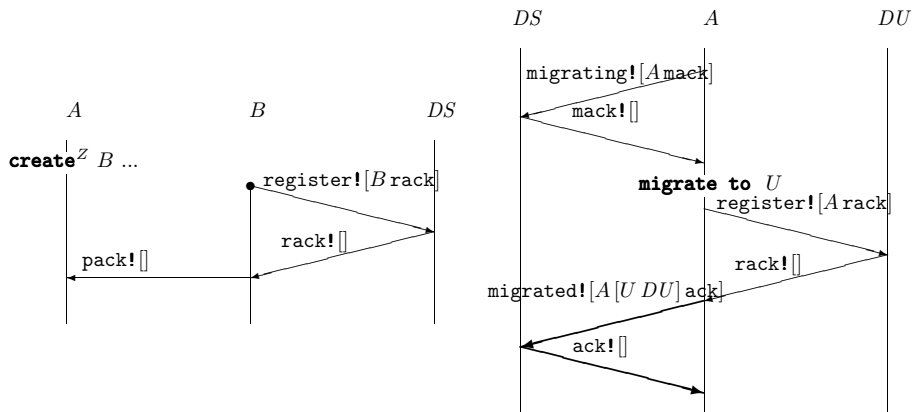
A single location-independent output, forwarded once between daemons (if the target agent is not at the local site), involves inter-agent messages as below. (Communications that are guaranteed to be between agents on the same site are drawn

with thin arrows.)



Creation The compositional encoding for **create**^Z is similar to that of the encoding in Section 3. It differs in two main ways. Firstly the source language name *b* of the new agent must be replaced by the actual agent name *B* tupled with the names *S* of this site and *DS* of the daemon on this site. Secondly, the internal forwarder, receiving on **deliver**, is no longer required — the final delivery of messages from daemons to agents is now always local to a site, and so can be done using **iflocal**. An explicit acknowledgement (on **dack** in the simple translation) is likewise unnecessary.

A single creation involves inter-agent messages as on the left below.



Migration Degenerate migrations, of an agent to the site it is currently on, must now be identified and treated specially; otherwise, the Daemon can deadlock. An agent *A* executing a non-degenerate migration now synchronises with the daemon *DS* on its starting site *S*, then migrates, registers with the daemon *DU* on its destination site *U*, then synchronises again with *DS*. In between the first and last synchronisations the agent lock for *A* in daemon *DS* is held, preventing *DS* from attempting to deliver messages to *A*.

A single migration involves inter-agent messages as on the right above.

Local communication The translation of **iflocal** must now extract the real agent name *B* from the triple *b*, but is otherwise trivial.

6. EXAMPLE INFRASTRUCTURE: QUERY SERVER WITH CACHING ALGORITHM

In this final example we take a further step towards a realistic algorithm, demonstrating that nontrivial optimisations can be cleanly expressed within the Nomadic Pict framework.

The central forwarding server described in Section 3 is a bottleneck for all agent communication; further, all application messages must make two hops (and these messages are presumably the main source of network load). The forwarding pointers algorithm described in Section 5 removes the bottleneck, but there application messages may have to make many hops, even in the common case. Adapting the central forwarding server so as to reduce the number of application-message hops required, we have the central query server algorithm, first described in Section 4. It has a server that records the current site of every agent; agents synchronise with it on migrations. In addition, each site has a daemon. An application message is sent to the local daemon, which then queries the server to discover the site of the target agent; the message is then sent to the daemon on the target site. If the agent has migrated away, a notification is returned to the original daemon to try again. In the common case application messages will here take only one hop. The obvious defect is the large number of control messages between daemons and the server; to reduce these each site's daemon can maintain a cache of location data. The *Query Server with Caching (QSC)* [Wojciechowski and Sewell 2000] does this. When a daemon receives a mis-delivered message, for an agent that has left its site, the message is forwarded to the server. The server both forwards the message on to the agent's current site and sends a cache-update message to the originating daemon. In the common case application messages are therefore delivered in only one hop.

The QSC encoding in Appendix C makes the algorithm precise, reusing the main design patterns from the encodings of Sections 4 and 3. Each class of agents maintains some explicit state as an output on a lock channel. The query server maintains a map from each agent name to the site (and daemon) where the agent is currently located. This is kept accurate when agents are created or migrate. Each daemon maintains a map from some agent names to the site (and daemon) that they guess the agent is located at. This is updated only when a message delivery fails. The encoding of each high-level agent records its current site (and daemon).

The algorithm is very asynchronous and should have good performance, with most application-level messages delivered in a single hop and none taking more than three hops (though 5 messages). The query server is involved only between a migration and the time at which all relevant daemons receive a cache update; this should be a short interval. Some additional optimisations are feasible, e.g. updating the daemon's cache more frequently.

The algorithm does, however, depend on reliable machines. The query server has critical state; the daemons do not, and so in principle could be re-installed after a site crash, but it is only possible to reboot a machine when no other daemons have pointers (that they will use) to it. In a refined version of the protocol the daemons and the query server would use a store-and-forward protocol to deliver all messages reliably in spite of failures, and the query server would be replicated. In order to extend collaboration between clusters of domains (e.g. over a wide-area network),

a federated architecture of interconnected servers must be adopted. In order to avoid long hops, the agents should register and unregister with the local query server on changing domains (see Wojciechowski [2006] for an example algorithm: the Federated Query Server with Caching).

7. NOMADIC PICT: THE PROGRAMMING LANGUAGE AND ITS IMPLEMENTATION

Nomadic Pict is a prototype distributed programming language, based on the Nomadic π calculus of Section 2 and on the Pict language of Pierce and Turner [2000]. Pict is a concurrent, though not distributed, language based on the asynchronous π calculus [Milner et al. 1992]. It supports fine-grain concurrency and the communication of asynchronous messages, extending the π calculus with a rich type system, a range of convenient forms for programming (such as function declarations) that can be compiled down to π calculus, and various libraries.

Low-level Nomadic Pict adds the Nomadic π calculus primitives for programming mobile computations from Section 2: agent creation, migration of agents between sites, and communication of location-dependent asynchronous messages between agents. In addition to these, Nomadic Pict adds timeouts, a facility for initiating communication between separate programs with a trader for type dynamic values, and labelled variant types. High-level Nomadic Pict adds location-independent communication; we can express an arbitrary infrastructure for implementing this as a user-defined translation into the low-level language. The rest of the language is taken directly from Pict, with the front-end of the Nomadic Pict compiler based on the Pict compiler.

The language inherits a rich type system from Pict, including simple record types, higher-order polymorphism, simple recursive types and subtyping. It has a partial type inference algorithm, and many type annotations can in practice be inferred by the compiler.

Names play a key rôle in the Nomadic Pict language, as in Nomadic π . New names of agents and channels can be created dynamically. These names are *pure*, in the sense of Needham [1989]; no information about their creation is visible within the language (in our current implementation they do contain site IDs, but could equally well be implemented by choosing large random numbers). Site names contain an IP address and TCP port number of the runtime system which they represent. Channel, agent, and site names are first-class values and they can be freely sent to processes which are located at other agents. As in the π calculus, names can be scope-extruded.

Programs in high-level Nomadic Pict are compiled in the same way as they are formally specified, by translating the high-level program into the low-level language. That in turn is compiled to a core language executed by the runtime. The core language is architecture-independent; its constructs correspond approximately to those of the low-level Nomadic π calculus, extended with value types and system function calls. The runtime system executes in steps, in each of which the closure of the agent at the front of the *agent queue* is executed for a fixed number of interactions. An agent closure consists of a *run queue*, of Nomadic π process/environment pairs waiting to be scheduled (round-robin), *channel queues* of terms that are blocked on

internal or inter-agent communication, and an environment that records bindings of variables to channels and basic values. The process at the front of the run queue is evaluated according to the abstract machine designed for Pict [Turner 1996]. It ensures fair execution of the fine-grain parallelism in the language. The compiler and runtime are written in OCaml [Leroy 1995].

In Appendix D we give a more detailed overview of the language. To make the paper self-contained, we include both the Nomadic-Pict-specific features and some aspects of Pict. We also describe some useful syntactic sugar and distributed programming idioms, such as remote procedure calls (RPC) and distributed objects. The language implementation is described in Appendix E. For concreteness, the full syntax of the language is included as Appendix F. The implementation is available on-line, together with a tutorial, library documentation, and examples [Wojciechowski 2000a].

8. CORRECTNESS: NOMADIC π CALCULUS SEMANTIC DEFINITION

We now return to the calculus of Section 2. This section defines its semantics—the type system and operational semantics—and gives the basic metatheoretic results. The following Section 9 develops proof techniques over the semantics, which are then used in Section 10 to prove correctness of the Central Forwarding Server algorithm we gave in Section 3. Throughout we give outline proofs, highlighting the main points, and refer the reader to the PhD thesis of Unyapoth [2001] for full details.

8.1 Type System

The type system is based on a standard simply typed π calculus, with channels carrying (possibly tuple-structured) first-order values. This is extended with input/output subtyping, as in Pierce and Sangiorgi [1996]: channel types have capabilities \mathbf{r} (only input is allowed), \mathbf{w} (output only), or \mathbf{rw} (both), with \mathbf{r} covariant and \mathbf{w} contravariant. Additionally, the type of agent names has a capability \mathbf{m} or \mathbf{s} , with $\mathbf{Agent}^{\mathbf{s}} \leq \mathbf{Agent}^{\mathbf{m}}$; only the latter supports migration. There is a standard subsumption rule

$$\frac{\Gamma \vdash e \in S \quad \Gamma \vdash S \leq T}{\Gamma \vdash e \in T}$$

The main judgements are $\Gamma \vdash_a P$, for well-formedness of a basic process as part of agent a , and $\Gamma \vdash LP$, for well-formedness of located processes. There is also a judgement $\Gamma \vdash x@z$, taking the location z of x from a located type context Γ . Sometimes we use unlocated type contexts, also written Γ , and there are standard rules for pattern and expression formation. The typing rules are given in full in Appendix A; a few of the most interesting rules are below.

$$\frac{\Gamma \vdash a \in \mathbf{Agent}^{\mathbf{m}} \quad \Gamma \vdash s \in \mathbf{Site} \quad \Gamma \vdash_a P}{\Gamma \vdash_a \mathbf{migrate\ to\ } s \rightarrow P} \quad \frac{a \neq b \quad \Gamma, b : \mathbf{Agent}^Z \vdash_b P \quad \Gamma, b : \mathbf{Agent}^Z \vdash_a Q}{\Gamma \vdash_a \mathbf{create}^Z b = P \mathbf{in\ } Q}$$

$\frac{}{\Gamma \Vdash_a c!v \xrightarrow{c!v} @_a \mathbf{0}}$	$\frac{\Gamma \vdash c \in \sim^r T \quad \Gamma, \Delta \vdash v \in T \quad \text{dom}(\Delta) \subseteq \text{fv}(v) \quad \Delta \text{ extensible}}{\Gamma \Vdash_a c!p \rightarrow P \xrightarrow{c!v} @_a \text{match}(p, v) P}$
$\frac{\Gamma \Vdash_a P \xrightarrow{c!v} LP \quad \Gamma \Vdash_a Q \xrightarrow{c!v} LQ}{\Gamma \Vdash_a P \mid Q \xrightarrow{\tau} \mathbf{new} \Delta \mathbf{in} LP \mid LQ}$	$\frac{(\Gamma, x : T) \Vdash_a P \xrightarrow{c!v} LP \quad x \in \text{fv}(v) \quad x \neq c}{\Gamma \Vdash_a \mathbf{new} x : T \mathbf{in} P \xrightarrow{\Delta, x : T} LP}$
$\Gamma \Vdash_a \mathbf{migrate to } s \rightarrow P \xrightarrow{\text{migrate to } s} @_a P$	<p style="text-align: center;">.....</p>
$\frac{\Gamma, a : \mathbf{Agent}^m @_s \Vdash LP \xrightarrow{@_a \text{migrate to } s'} LQ}{\Gamma \Vdash \mathbf{new} a : \mathbf{Agent}^m @_s \mathbf{in} LP \xrightarrow{\tau} \mathbf{new} a : \mathbf{Agent}^m @_{s'} \mathbf{in} LQ}$	

Fig. 6. Selected LTS Rules

$$\frac{\Gamma \vdash a, b \in \mathbf{Agent}^s \quad \Gamma \vdash s \in \mathbf{Site} \quad \Gamma \vdash c \in \sim^w T \quad \Gamma \vdash v \in T}{\Gamma \vdash_a \langle b @_s \rangle c!v} \quad \frac{\Gamma \vdash_a P}{\Gamma \vdash @_a P}$$

The system also includes type variables and existential packages, deconstructed by pattern matching.

A type context is *extensible* if all term variables are of agent or channel types, and therefore may be new-bound.

8.2 Reduction Semantics

The reduction semantics was introduced informally in Section 2.5. Its formal definition involves structural congruence relations $P \equiv Q$ and $LP \equiv LQ$, defined in Appendix B.1, and a reduction relation $\Gamma \Vdash LP \rightarrow \Gamma' \Vdash LP'$ over pairs of located type contexts and located processes, defined in Appendix B.2.

8.3 Labelled Transition Semantics

The reduction semantics describes only the internal behaviour of complete systems of located processes — for compositional reasoning we need also a typed labelled transition semantics, expressing how processes can interact with their environment. This lifts the development of corresponding reduction and labelled transition semantics in the π calculus [Milner 1992] to Nomadic π . Transitions are defined inductively on process structure, without the structural congruence. The transition relations have the following forms, for basic and located process:

$$\Gamma \Vdash_a P \xrightarrow{\alpha} LP \quad \Gamma \Vdash LP \xrightarrow{\beta} LQ$$

Here the *unlocated labels* α are of the following forms:

τ	internal computation
migrate to s	migrate to the site s
$c!v$	send value v along channel c
$c?v$	receive value v from channel c

The *located labels* β are of the form τ or $@_a\alpha$ for $\alpha \neq \tau$. Private names (together with their types, which may be annotated with an agent's current site) may be exchanged in communication and are made explicit in the transition relation by the extruded context Δ . Selected rules are given in Figure 6, and the full definition in Appendix B.3.

Adding **migrate to** s to the standard input/output and τ labels is an important design choice, made for the following reasons.

- Consider a located process LP in some program context. If an agent a in LP migrates, the location context is consequently updated with a associated to its new site. This change of location context has an effect on both LP and its environment, since it can alter their execution paths (especially those involving location-dependent communication with a). Migration of an agent must therefore be thought of as a form of interaction with the environment.
- We observe, in the reduction rules, that the location context in the configuration *after* the transition can only be modified by migration of an agent. Including this migrating action allows the location context on the right hand side to be omitted.

Execution of other agent primitives (i.e. **create** and **iflocal**) is regarded as internal computation, since it does not have an immediate effect on program contexts. In the case of **create**, the newly created agent remains unknown to the environment unless its name is extruded by an output action.

8.4 Basic Metatheory

In a typed semantics, the type system should prevent a mismatch between the value received and the shape expected in communication. However, matching a value and a pattern of the same type does not always yield a substitution. For example, taking Γ to be $x : [\square \square]$, a pattern $[y z]$ may have type $[\square \square]$ w.r.t. Γ , but $\text{match}([y z], x)$ is undefined. A similar situation occurs when matching a name x of an existential type to an existential pattern $\{X\}p$. To prevent this, we define *ground* and *closed* type contexts as follows.

DEFINITION 8.1 (GROUND TYPE CONTEXT)

A type context Γ is ground if, for all $x \in \text{dom}(\Gamma)$, $\Gamma \vdash x \in T$ implies $T \neq [T_1 \dots T_n]$ and $T \neq \{X\}S$, for any T_1, \dots, T_n, X, S .

DEFINITION 8.2 (CLOSED TYPE CONTEXT)

A type context Γ is closed if it is ground and $\text{fv}(\Gamma) \cap \mathcal{TV} = \emptyset$ and, for all $x \in \text{dom}(\Gamma)$, $\Gamma \vdash x \in T$ implies $T \notin \mathcal{T}$.

It is easy to show that each name declared in a closed type context is either a site, an agent, or a channel.

We may now state the type preservation result.

THEOREM 8.1 (TYPE PRESERVATION)

For any well-formed closed located type context Γ , if $\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LQ$ then $\Gamma, \Delta \vdash LQ$.

PROOF SKETCH An induction on the derivations of $\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LP$ and $\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LQ$. Γ needs to be closed so that, matching a pattern with a value of the same type always yields a type-preserving substitution, whenever the transition involving matching occurs. \square

THEOREM 8.2 (REDUCTION/LTS CORRESPONDENCE)

For any well-formed located type context Γ and located process LP such that $\Gamma \vdash LP$, we have: $\Gamma \Vdash LP \rightarrow \Gamma' \Vdash LQ$ if and only if either

- $\Gamma \Vdash LP \xrightarrow{\tau} LQ$ with $\Gamma' = \Gamma$, or
- $\Gamma \Vdash LP \xrightarrow{@_a \text{migrate to } s} LQ$ with $\Gamma' = \Gamma \oplus a \mapsto s$.

PROOF SKETCH We need to show this in two parts: that a reduction implies a silent transition or a migrate action, and vice versa. Each of the two parts is shown by an induction on reduction/transition derivations. The case where the silent transition of LP is derived by the communication rule needs the following lemma, which can easily be proved by an induction on transition derivations.

LEMMA 8.3

- If $\Gamma \Vdash LP \xrightarrow[\Xi]{@_a c!v} LQ$ then $LP \equiv \mathbf{new} \Delta, \Xi \mathbf{in} (@_a c!v \mid LP')$ for some Δ and LP' . Moreover, $LQ \equiv \mathbf{new} \Delta \mathbf{in} LP'$.
- If $\Gamma \Vdash LP \xrightarrow[\Xi]{@_a c?v} LQ$ then, for some Δ, p and LP', Q , with $\text{dom}(\Delta) \cap \text{dom}(\Xi) = \emptyset$, either:
 - $LP \equiv \mathbf{new} \Delta \mathbf{in} (@_a c?p \rightarrow Q \mid LP')$ and $LQ \equiv \mathbf{new} \Delta \mathbf{in} (@_a (\text{match}(p, v)Q) \mid LP')$; or
 - $LP \equiv \mathbf{new} \Delta \mathbf{in} (@_a *c?p \rightarrow Q \mid LP')$ and $LQ \equiv \mathbf{new} \Delta \mathbf{in} (@_a (\text{match}(p, v)Q) \mid @_a *c?p \rightarrow Q \mid LP')$.
- If $\Gamma \Vdash LP \xrightarrow{@_a \text{migrate to } s} LQ$ then

$$LP \equiv \mathbf{new} \Delta \mathbf{in} (@_a \mathbf{migrate to } s \rightarrow P \mid LP')$$

for some Δ and LP', P . Moreover, $LQ \equiv \mathbf{new} \Delta \mathbf{in} (@_a P \mid LP')$.

As in Theorem 8.1, Γ needs to be closed so that, matching a pattern with a value of the same type always yields a type-preserving substitution, whenever the transition involving matching occurs. \square

The next two lemmas ensure the absence of two kinds of runtime errors: mismatching of values exchanged in channel communication, and non-evaluable expressions.

LEMMA 8.4 (RUNTIME SAFETY: CHANNELS)

Given that Γ is a closed type context, and $(\Gamma, \Delta)(c) = \sim^I T$, we have:

- (1) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a c!v \mid LP)$ then $I \leq \mathbf{w}$;

- (2) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a c? p \rightarrow P \mid LP)$ then $I \leq r$;
- (3) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a *c? p \rightarrow P \mid LP)$ then $I \leq r$;
- (4) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a(c!v|c? p \rightarrow P) \mid LP)$ then $\mathbf{match}(p, v)$ is defined; and
- (5) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a(c!v|*c? p \rightarrow P) \mid LP)$ then $\mathbf{match}(p, v)$ is defined.

LEMMA 8.5 (RUNTIME SAFETY: EXPRESSIONS)

Given that Γ is a closed type context, we have:

- (1) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a(\mathbf{if} v \mathbf{then} P \mathbf{else} Q) \mid LP)$ then $v \in \{\mathbf{true}, \mathbf{false}\}$;
- (2) if $\Gamma \vdash \mathbf{new} \Delta \mathbf{in} (@_a(\mathbf{let} p = ev \mathbf{then} P) \mid LP)$ then $\mathbf{eval}(ev)$ and $\mathbf{match}(p, \mathbf{eval}(ev))$ are defined.

9. CORRECTNESS: NOMADIC π CALCULUS SEMANTIC TECHNIQUES

In this section we describe the Nomadic π techniques used for stating and proving correctness. This is not specific to the particular CFS algorithm, although examples are taken from it. The next section describes the large-scale structure of the correctness proof, using these techniques.

Correctness Statement We are expressing distributed infrastructure algorithms as encodings from a high-level language to its low-level fragment, so the behaviour of a source program and its encoding can be compared directly with some notion of *operational equivalence* — our main theorem will be roughly of the form

$$\forall P. P \simeq \mathcal{C} \llbracket P \rrbracket \quad (\dagger)$$

where P ranges over well-typed programs of the high-level language (P may use location-independent communication whereas $\mathcal{C} \llbracket P \rrbracket$ will not). Now, what equivalence \simeq should we take? The stronger it is, the more confidence we gain that the encoding is correct. At first glance, one might take some form of weak bisimulation since (modulo divergence) it is finer than most notions of testing [de Nicola and Hennessy 1984] and is easier to work with; see also the discussion of Sewell [1997] on the choice of an appropriate equivalence for a Pict-like language. However, as in Nestmann’s work on choice encodings [Nestmann and Pierce 1996], (\dagger) would not hold, as the encodings $\mathcal{C} \llbracket P \rrbracket$ tend to involve *partial commitment* of some nondeterministic choices. In particular, migration steps and acquisitions of the daemon or agent locks involve nondeterministic internal choices, and lead to partially committed states — target level terms which are not bisimilar to any source level term. We therefore take \simeq to be an adaptation of *coupled simulation* [Parrow and Sjödin 1992] to our language. This is a slightly coarser relation, but it is expected to be finer than any reasonable notion of observational equivalence for Nomadic π (again modulo questions of divergence and fairness). This is discussed further below in Section 9.1.

Dealing with House-keeping steps Our example infrastructure introduces many τ steps, each of which induces an intermediate state — a target level term which is not a literal translation of any source level term. Some of these steps are the partial commitment steps mentioned above. Many, however, are deterministic *house-keeping* steps; they can be reduced to certain normal forms, and related to them by

expansions (defined in Section 9.3). For example, consider the following fragment of code from the \mathcal{C} -encoding (after some reduction steps).

```

new  $\Phi_{aux, m} : \text{Map}[\text{Agent}^s \text{ Site}], \Delta$  in
  @D(Daemon
    | lookup[\text{Agent}^s \text{ Site}] a in m with
      found(s)  $\rightarrow$  new dack :  $\hat{\text{rw}}[]$  in
         $\langle a@s \rangle$  deliver! {X} [c v dack] | dack?[]  $\rightarrow$  lock!m
        notfound  $\rightarrow 0$ )
    | @a([[P]]a | Deliverer | ...)
    | @b1([[Q1]]b1 | ...) | ... | @bn([[Qn]]bn | ...)
  where Deliverer  $\stackrel{\text{def}}{=} \text{*deliver? } \{X\} [c \ v \ \text{dack}] \rightarrow ((D@SD)\text{dack!}[] | c!v)$ 

```

This is a state of the encoded whole system in which an agent has sent a message forwarding request (to agent a) to the daemon, and the daemon's request code has acquired the daemon lock, which contains the site map m . The subsequent steps performed by the daemon D , and by the **Deliverer** process in the agent a , are house-keeping steps. They include the map lookup operation, sending the message to the **Deliverer** process in a (with a location-dependent message to channel **deliver** there), and communication along the **dack** channel.

To prove these give rise to expansions requires a variety of techniques, some novel and some straightforward adaptations of earlier work.

- Maps. We use a π calculus encoding of finite maps, similar to the encoding of lists with persistent values [Milner 1993]. We prove that the encoding is correct, and that map lookup and update operations yield expansions.
- The location-dependent **deliver** message, sent to agent a , is guaranteed to arrive because a cannot migrate until the daemon lock is released by **lock!** m , which does not occur until agent a returns a **dack** to the daemon. To capture this, we introduce a notion of *temporarily immobile* located process — one in which no migration can take place until an input on a lock channel. This is discussed in Section 9.5.

Certain reductions, such as the location-dependent message delivery step, are *deterministic*, as defined in Section 9.4. The key property of a temporarily immobile process is that such deterministic reductions still give rise to expansions when in parallel with temporarily immobile processes.

Proving that processes are temporarily immobile involves a coinductive characterisation and preservation results (under parallel and new-binders).
- The reaction of the **deliver** message and the **Deliverer** process, in agent a , is essentially functional. We adapt the notion of uniform receptiveness [Sangiorgi 1999], showing that the reaction induces an expansion by showing that the **deliver** channel is uniformly receptive — it always has a single replicated input in each agent, and no other input. The details are omitted here.
- The location-dependent **dack** message, from agent a to the daemon, is guaranteed to arrive for the simple reason that the daemon cannot migrate — it has the static

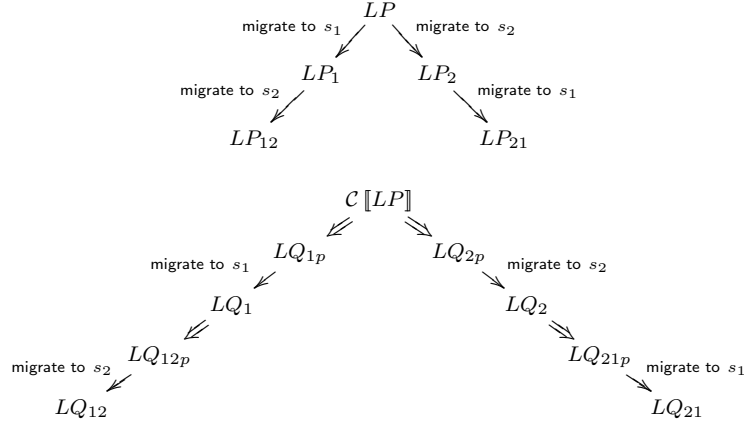


Fig. 7. An example of partial committed state

type \mathbf{Agent}^s . The reduction step is therefore deterministic, and hence induces an expansion.

- The **dack** acknowledgement channel is fresh for each request, so the daemon contains exactly one input and at most one output. It is straightforward to show that the communication is deterministic and hence gives an expansion.
- In all of the above, we make essential use of congruence results for expansion to pull out the interesting part of the process, allowing the b_i agents and parts of the daemon to be neglected. The presence of agent mobility and location-dependent communication means these results must take account of the possible migrations of agents; in Section 9.2 we define *translocating* bisimulations that do so; translocating expansions are similar.

9.1 Partial Commitment and Coupled Simulation

As an example, consider the encoding $\mathcal{C}[[LP]]$ of an agent a which sends message $c!v$ to agent b at the current site of a , and in parallel visits the sites s_1 and s_2 (in any order).

$$LP \stackrel{\text{def}}{=} @_a(\langle b \rangle c!v \mid \mathbf{migrate\ to\ } s_1 \mid \mathbf{migrate\ to\ } s_2)$$

Assuming a and b are initially at the same site, parts of the reduction graphs of LP and $\mathcal{C}[[LP]]$ can be represented as in Figure 7. If the **migrate to** s_1 process in $\mathcal{C}[[LP]]$ successfully acquires the local lock (a partial commitment step) the resulting process (LQ_{1p} in Figure 7) does not correspond exactly to any state of LP . LQ_{1p} cannot correspond to LP_1 since executing $\langle b \rangle c!v$ at this point means that $c!v$ will reach b (which is not the case for node LP_1); it cannot correspond to LP either, since we know that a will eventually end up in s_2 .

To address this phenomenon, coupled simulation [Parrow and Sjödin 1992] relaxes the bisimulation clauses somewhat. A pair $(\mathcal{S}_1, \mathcal{S}_2)$ of type-context-indexed relations is a *coupled simulation* if:

- \mathcal{S}_1 and $(\mathcal{S}_2)^{-1}$ are weak simulations (the standard coinductive notion of simula-

tion, indexed by located type contexts).

- if $(LP, LQ) \in (\mathcal{S}_1)_\Gamma$ then there exists LQ' such that $\Gamma \Vdash LQ \xrightarrow{\tau} LQ'$ and $(LP, LQ') \in (\mathcal{S}_2)_\Gamma$.
- if $(LP, LQ) \in (\mathcal{S}_2)_\Gamma$ then there exists LP' such that $\Gamma \Vdash LP \xrightarrow{\tau} LP'$ and $(LP', LQ) \in (\mathcal{S}_1)_\Gamma$.

Two processes LP, LQ are *coupled similar* wrt Γ , written $LP \Leftarrow_\Gamma LQ$, if they are related by both components of some coupled simulation.

Intuitively “ LQ coupled simulates LP ” means that “ LQ is at most as committed as LP ” with respect to internal choices and that LQ may internally evolve to a state LQ' where it is at least as committed as LP , i.e., where LP coupled simulates LQ' .

In this paper, coupled simulation will be used for relating whole systems, which cannot be placed in any program context. For this reason, we do not need to incorporate translocation into the definition above.

9.2 Translocating Equivalences and Congruence Result

To prove our main result (†) we need compositional techniques, allowing separate parts of the protocols to be treated separately. In particular, we need operational congruences (both equivalences and preorders) that are preserved by program contexts involving parallel composition and new-binding. In Nomadic π the behaviour of location-dependent communications depends on the relative location of agents: if a and b are at the same site then the location-dependent message $@_b(a@s)c!v$ reduces to (and in fact is weakly equivalent to) the local output $@_ac!v$, whereas if they are at different sites then the location-dependent message is weakly equivalent to $\mathbf{0}$. A parallel context, for example $[\cdot]@_a\mathbf{migrate\ to\ }s$, can migrate the agent a , so to obtain a congruence we need refined equivalences, taking into account the possibility of such changes of agent location caused by the environment.

Relocators, ranged over by δ , can be applied to located type contexts in order to relocate agents in such contexts. A valid relocator for (Γ, M) is a type-respecting partial function from M to site names of Γ , formally defined below.

DEFINITION 9.1 (VALID RELOCATORS)

A relocator δ is said to be *valid for* (Γ, M) if $\text{dom}(\delta) \subseteq M$ and, for all $x \in M$, $\Gamma \vdash x \in \mathbf{Agent}^m$ and $\Gamma \vdash \delta(x) \in \mathbf{Site}$.

We write $\Gamma\delta$ for the result of applying δ to Γ and $\Gamma\delta\beta$ for $(\Gamma\delta)\beta$.

Allowing arbitrary relocations would give too strong a notion, though. We introduce *translocating* relations that are parameterised by a set of agents that the environment may move.

DEFINITION 9.2 (TRANSLOCATING INDEXED RELATION)

A *translocating indexed relation* is a binary relation between $n\pi_{LD,LI}$ processes, indexed by closed well-formed located type contexts Γ and sets $M \subseteq \text{mov}(\Gamma)$, where

$\text{mov}(\Gamma)$ is the set of names of type Agent^m in Γ :

$$\begin{aligned} \text{mov}(\bullet) &\stackrel{\text{def}}{=} \emptyset \\ \text{mov}(\Gamma, X) &\stackrel{\text{def}}{=} \text{mov}(\Gamma) \\ \text{mov}(\Gamma, x : T@z) &\stackrel{\text{def}}{=} \begin{cases} \text{mov}(\Gamma) \cup \{x\} & T = \text{Agent}^m \\ \text{mov}(\Gamma) & \text{otherwise} \end{cases} \end{aligned}$$

Channel communication introduces further problems since it allows extrusion of new agent names to and from the environment. Consider an output of a new-bound agent name a to the environment. Other components in the environment may then send messages to a , but cannot migrate it, so when checking a translocating equivalence we do not need to consider relocation of a . On the other hand, a new agent name received from the environment by an input process is the name of an agent created in the environment, so (if created with the mobile capability) it may be migrated at any time.

Therefore the translocating index of the bisimulation only needs to be updated when an input action occurs. For this we define the set $M_1 \uplus_{\beta} M_2$ to be $M_1 \cup M_2$ whenever β is an input, and to be M_1 otherwise.

$$M_1 \uplus_{\beta} M_2 \stackrel{\text{def}}{=} \begin{cases} M_1 \cup M_2 & \exists a, c, v, \beta = @_a c?v \\ M_1 & \text{otherwise} \end{cases}$$

The notion of translocating bisimulation can therefore be formalised as follows.

DEFINITION 9.3 (TRANSLOCATING SIMULATIONS)

(1) A translocating indexed relation \mathcal{S} on $n\pi_{\text{LD}, \text{LI}}$ is a *translocating strong simulation* if $(LP, LQ) \in \mathcal{S}_{\Gamma}^M$ implies the following:

- $\Gamma \vdash LP$ and $\Gamma \vdash LQ$;
- $M \subseteq \text{mov}(\Gamma)$; and

— For any relocater δ valid for (Γ, M) , if $\Gamma \delta \Vdash LP \xrightarrow[\Delta]{\beta} LP'$ then there exists LQ' such that $\Gamma \delta \Vdash LQ \xrightarrow[\Delta]{\beta} LQ'$ and $(LP', LQ') \in \mathcal{S}_{\Gamma \delta \beta, \Delta}^{M \uplus_{\beta} \text{mov}(\Delta)}$.

\mathcal{S} is called a *translocating strong bisimulation* if all of its indexed relations are symmetric. Two located processes LP and LQ are translocating strongly bisimilar w.r.t. Γ, M , written $LP \sim_{\Gamma}^M LQ$, if there exists a translocating strong bisimulation which when indexed by Γ and M , contains the pair (LP, LQ) .

(2) Replacing $\Gamma \delta \Vdash LQ \xrightarrow[\Delta]{\beta} LQ'$ in the final item of this definition with

$\Gamma \delta \Vdash LQ \xrightarrow[\Delta]{\hat{\beta}} LQ'$ yields the *weak* version of translocating simulation. A located process LQ weak translocating bisimulates LP w.r.t. Γ, M , denoted $LP \approx_{\Gamma}^M LQ$, if there exists a weak translocating bisimulation which when indexed by Γ, M , contains the pair (LP, LQ) .

Some simple examples of translocating bisimulations are the following.

$$\begin{aligned} @_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q &\sim_{\Gamma}^{M_1} @_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q' \\ @_a \langle b@s \rangle c!v &\approx_{\Gamma}^{M_2} @_b c!v \end{aligned}$$

where $M_1 \subseteq \text{mov}(\Gamma)/\{a, b\}$ and $M_2 \subseteq \text{mov}(\Gamma)/\{b\}$; we assume that the processes above are well-typed w.r.t. Γ , and that $\Gamma \vdash a@s$ and $\Gamma \vdash b@s$.

We prove congruence results for both strong and weak translocating bisimulation, stating the result here only for the strong version. It uses a further auxiliary definition: the set $\text{mayMove}(LP)$ is the set of agents in LP syntactically containing **migrate to**.

THEOREM 9.1 (TRANSLOCATING CONGRUENCE)

Given a closed located type context Γ, Θ with Θ extensible, if

- $LP \dot{\sim}_{\Gamma, \Theta}^{M_P} LP'$ and $LQ \dot{\sim}_{\Gamma, \Theta}^{M_Q} LQ'$,
- $\text{mayMove}(LQ, LQ') \subseteq M_P$,
- $\text{mayMove}(LP, LP') \subseteq M_Q$, and
- $M \stackrel{\text{def}}{=} M_P \cap M_Q \cap \text{agents}(\Gamma)$

then

$$\mathbf{new} \Theta \mathbf{in} (LP \mid LQ) \dot{\sim}_{\Gamma}^M \mathbf{new} \Theta \mathbf{in} (LP' \mid LQ').$$

PROOF SKETCH The proof deals with derivatives of $\mathbf{new} \Theta \mathbf{in} LP \mid LQ$ w.r.t. Γ , which have the general form of

$$LR_k = \mathbf{new} \Theta, \Theta_{comm} \mathbf{in} (LP_k \mid LQ_k)$$

well-typed w.r.t. $\Gamma, \Theta_{in}, \Theta_{out}$. Here we classify new names bound in the derivative, and those extruded to or from the environment as follows.

- Θ_{comm} consists of names exchanged by communication between LP and LQ . This can be classified further as Θ_{comm}^{LP} , the private names of LP extruded by output actions to LQ , and vice versa for Θ_{comm}^{LQ} .
- Θ_{out} consists of names extruded by output actions to the environment. Again, this can be classified further as Θ_{out}^{LP} , for the names extruded by LP , and vice versa for Θ_{out}^{LQ} .
- Θ_{in} consists of names received from the environment.

Using this classification of names, the set $\text{mov}(\Theta_{in})$ anticipates the movements of agents received from the environment (i.e. the context of LR_k), and the set $M_P \cup \text{mov}(\Theta_{comm}^{LQ}, \Theta_{out}^{LQ})$ anticipates the movements of free agents in LQ_k . Since the environment of LP_k comprises LQ_k and the context of LR_k as a whole, the translocating index of the bisimulation relations between LP_k and LP'_k must include the set below.

$$M_{P_k} = M_P \cup \text{mov}(\Theta_{comm}^{LQ}, \Theta_{out}^{LQ}, \Theta_{in})$$

The premises of Theorem 9.1 can therefore be generalised in the coinduction as follows.

- $LP_k \dot{\sim}_{\Gamma, \Theta_{in}, \Theta_{out}, \Theta_{comm}, \Theta}^{M_{P_k}} LP'_k$, and $LQ_k \dot{\sim}_{\Gamma, \Theta_{in}, \Theta_{out}, \Theta_{comm}, \Theta}^{M_{Q_k}} LQ'_k$, where M_{Q_k} is defined in the similar way as M_{P_k} ;
- $\text{mayMove}(LP_k, LP'_k) \subseteq M_Q \cup \text{mov}(\Theta_{comm}^{LP}, \Theta_{out}^{LP})$; and
- $\text{mayMove}(LQ_k, LQ'_k) \subseteq M_P \cup \text{mov}(\Theta_{comm}^{LQ}, \Theta_{out}^{LQ})$.

The proof of this theorem relies on the invariance under labelled transitions of the above premises. \square

By using the techniques outlined in the beginning of Section 9, we may prove that

$$\begin{aligned} & \mathbf{new} \Phi_{aux, m} : \text{Map}[\text{Agent}^s \text{ Site}] \mathbf{in} \\ & \quad @_D(\text{Daemon} \\ & \quad \quad | \mathbf{lookup}[\text{Agent}^s \text{ Site}] a \mathbf{in} m \mathbf{with} \\ & \quad \quad \quad \mathbf{found}(s) \rightarrow \mathbf{new} \text{dack} : \hat{\sim}^{rw} [] \mathbf{in} \\ & \quad \quad \quad \langle a@s \rangle \text{deliver!} \{X\} [c v \text{dack}] | \text{dack?}[] \rightarrow \text{lock!} m \\ & \quad \quad \quad \mathbf{notfound} \rightarrow \mathbf{0}) \\ & \quad | @_a([\![P]\!]_a | \text{Deliverer} | \dots) \\ & \approx_{\Gamma, \Delta}^{\{b_1, \dots, b_n\}} \end{aligned}$$

$$\begin{aligned} & \mathbf{new} \Phi_{aux, m} : \text{Map}[\text{Agent}^s \text{ Site}] \mathbf{in} \\ & \quad @_D(\text{Daemon} | \text{lock!} m) \\ & \quad | @_a([\![P]\!]_a | \text{Deliverer} | \dots) \end{aligned}$$

where the processes above are well-typed with respect to Γ, Δ . Applying the congruence result, the fragment of code from the \mathcal{C} -encoding given in the beginning of this section can be proved to translocating weak bisimulate the following process.

$$\begin{aligned} & \mathbf{new} \Phi_{aux, m} : \text{Map}[\text{Agent}^s \text{ Site}], \Delta \mathbf{in} \\ & \quad @_D(\text{Daemon} | \text{lock!} m) \\ & \quad | @_a([\![P]\!]_a | \text{Deliverer} | \dots) \\ & \quad | @_{b_1}([\![Q_1]\!]_{b_1} | \dots) | \dots | @_{b_n}([\![Q_n]\!]_{b_n} | \dots) \end{aligned}$$

9.3 Expansion

To construct the coupled simulation, we use an *expansion* relation $\hat{\sim}$ [Nestmann and Pierce 1996] and the “up to” technique of Sangiorgi and Milner [1992], adapted with translocation, to allow elimination of target processes that are in intermediate/house-keeping stages.


A definition of expansion uses two refinements of weak simulation: progressing and strict simulation. We adapt the definitions from Nestmann [1996], adding type contexts and translocation.

DEFINITION 9.4 (PROGRESSING AND STRICT SIMULATION)

A weak translocating simulation \mathcal{S} is called

- strict* if, for all $(LP, LQ) \in \mathcal{S}_\Gamma^M$ and valid δ for (Γ, M) , $\Gamma\delta \Vdash LP \xrightarrow[\Delta]{\beta} LP'$ implies there exists LQ' such that $\Gamma\delta \Vdash LQ \xrightarrow[\Delta]{\hat{\beta}} LQ'$ with $(LP', LQ') \in \mathcal{S}_{\Gamma\delta\beta, \Delta}^{M \uplus \beta \text{mov}(\Delta)}$.
 - progressing* if, for all $(LP, LQ) \in \mathcal{S}_\Gamma^M$ and valid δ for (Γ, M) , $\Gamma\delta \Vdash LP \xrightarrow[\Delta]{\beta} LP'$ implies there exists LQ' such that $\Gamma\delta \Vdash LQ \xrightarrow[\Delta]{\hat{\beta}} LQ'$ with $(LP', LQ') \in \mathcal{S}_{\Gamma\delta\beta, \Delta}^{M \uplus \beta \text{mov}(\Delta)}$;
- and

LQ is said to *progressing simulate* (or *strictly simulate*) LP w.r.t. Γ, M if there exists a progressing simulation \mathcal{S} (or a strict simulation) such that $(LP, LQ) \in \mathcal{S}_\Gamma^M$.



The above diagrams show progressing and strict simulations. Informally, LQ strictly simulates LP means that LQ weakly simulates LP , but LQ never introduces more internal steps and may ignore the silent transitions of LP . On the other hand, LQ progressing simulates LP means that LQ weakly simulates LP , but LQ introduces more internal steps and never ignores a silent action, hence the absence of $\hat{\cdot}$ in the weak transition of LQ in the definition. The definition of expansion simply makes use of these two refinements.

DEFINITION 9.5 (EXPANSION)

An indexed binary relation \mathcal{S} is a *translocating expansion* if \mathcal{S} is a strict simulation and \mathcal{S}^{-1} is a progressing simulation.

LP translocating expands LQ w.r.t. Γ under M , written $LP \succeq_\Gamma^M LQ$, if there exists an expansion \mathcal{S} with $(LP, LQ) \in \mathcal{S}_\Gamma^M$. Moreover, if $LP \succeq_\Gamma^{\text{mov}(\Gamma)} LQ$ then LP and LQ are said to be related by *expansion congruence*, written $LP \succeq_\Gamma LQ$.

We depend on a congruence result, analogous to that above, for expansion. The proof of this result is similar to that for translocating bisimulations.

9.4 Deterministic Reduction

A component in a system of concurrent processes may be deterministic, in the sense that its next computational step can be determined. An example of this is a location-dependent message $\langle b@s \rangle c!v$, executed in an agent a ; if the agent b is static, and is located at s , then all the subsequent transitions are determined, eventually moving the output $c!v$ to b . We define deterministic reduction as follows.

DEFINITION 9.6 (DETERMINISTIC REDUCTION)

Given a closed located type context Γ and $M \subseteq \text{mov}(\Gamma)$, a located processes LP is said to *deterministically reduce* to LQ w.r.t. (Γ, M) , written $\Gamma \Vdash LP \xrightarrow[M]{\text{det}} LQ$, if, for any valid δ for (Γ, M) , the following hold:

- $\Gamma\delta \Vdash LP \xrightarrow{\tau} LQ$; and
- $\Gamma\delta \Vdash LP \xrightarrow[\Delta]{\beta} LQ'$ implies $\beta = \tau$, $\Delta = \bullet$ and $LQ' \sim_{\Gamma\delta}^M LQ$.

We also define the relation $\xrightarrow[M]{\text{det}}$ to be the transitive closure of $\xrightarrow[M]{\text{det}}$; that is $\Gamma \Vdash LP \xrightarrow[M]{\text{det}} LQ$ implies there exists LP_1, \dots, LP_n such that, letting $LP = LP_0$

and $LQ = LP_{n+1}$, we have

$$\Gamma \Vdash LP_i \xrightarrow[M]{\text{det}} LP_{i+1} \quad 0 \leq i \leq n$$

A process LP is said to be τ -deterministic w.r.t. Γ, M if there exists LQ such that $\Gamma \Vdash LP \xrightarrow[M]{\text{det}} LQ$.

The lemma below states the key property of τ -determinacy: that a deterministic reduction induces an expansion.

LEMMA 9.2 (DETERMINISTIC REDUCTION INDUCES EXPANSION)

If $\Gamma \Vdash LP \xrightarrow[M]{\text{det}} LQ$ then $LP \dot{\prec}_{\Gamma}^M LQ$.

From the example fragment of code from \mathcal{C} -encoding, when the agent a received the message forwarded from the daemon, it sends an acknowledgement back to the daemon using $\langle D@SD \rangle \text{dack!} []$. Since the location-dependent sugar output is τ -deterministic, we have:

$$@_a \langle D@SD \rangle \text{dack!} [] \dot{\prec}_{\Gamma}^M @_D \text{dack!} []$$

for any $M \subseteq \text{mov}(\Gamma)$ such that $D \notin M$. However, since the daemon D is static, $@_a \langle D@SD \rangle \text{dack!} []$ is related by expansion congruence to $@_D \text{dack!} []$; and hence placing the location-dependent output in any program context yields an expansion.

9.5 Temporary Immobility

At many points in the execution of an encoded program, it is intuitively clear that an agent cannot migrate — while waiting for an acknowledgement from the daemon, or for either `currentloc` or `lock` to be released in the agent or daemon. To capture such an intuition, we consider derivatives of a process LP — if an input action on a lock channel l always precedes any (observable) migration action then LP can be said to be temporarily immobile, blocked by l . Care must be taken, however, to ensure that the lock l is not released by the environment. This can be made precise by the following definitions.

As in the case of translocating equivalences, we need to consider the possibility of agents being moved by the environment.

DEFINITION 9.7 (TRANSLOCATING PATH)

A *translocating path* of LP_0 wrt (Γ, M) is a sequence

$$\frac{\beta_1}{\Delta_1} \rightarrow \dots \rightarrow \frac{\beta_n}{\Delta_n}$$

for which there exist LP_1, \dots, LP_n and $\delta_0, \dots, \delta_{n-1}$ such that for each $i \in 0 \dots n-1$:

— δ_i is a valid relocater for $(\hat{\Gamma}, \hat{M})$, where

$$\begin{aligned} \hat{\Gamma} &\stackrel{\text{def}}{=} \Gamma, \Delta_1, \dots, \Delta_i \\ \hat{M} &\stackrel{\text{def}}{=} M \uplus_{\beta_1} \text{mov}(\Delta_1) \dots \uplus_{\beta_i} \text{mov}(\Delta_i), \text{ and} \end{aligned}$$

— $((\Gamma \delta_0, \Delta_1) \delta_1 \beta_1, \Delta_2 \dots \beta_i, \Delta_i) \delta_i \Vdash LP_i \xrightarrow[\Delta_{i+1}]{\beta_{i+1}} LP_{i+1}$.

DEFINITION 9.8 (TEMPORARY IMMOBILITY)

Given a closed located type context Γ , a located process LP with $\Gamma \vdash LP$, and a translocating index $M \subseteq \text{agents}(\Gamma)$, LP is *temporarily immobile* under lock l wrt (Γ, M) if, for all translocating paths

$$\frac{\beta_1}{\Delta_1} \dots \frac{\beta_n}{\Delta_n}$$

of LP wrt (Γ, M) which do not contain an input action $\beta_i = @_a c!v$ with $l \in \text{fv}(c, v)$, the following hold for all $i \leq n$, b, c, v and s :

- $\beta_i = @_b c!v$ implies $l \notin \text{fv}(\beta_i)$; and
- $\beta_i \neq @_b \text{migrate to } s$.

Consider for example the process below.

$$LQ \stackrel{\text{def}}{=} \mathbf{new} \Omega_{aux} \mathbf{in} \\ @_D \text{Daemon} \\ | @_a ([P]_a | \text{currentloc!s} | \text{Deliverer})$$

Here agent a cannot migrate until the daemon lock `lock` is successfully acquired, so LQ is temporarily immobile under `lock` with respect to any type-correct (Γ, M) that does not admit environmental relocation of a , i.e. with $a \notin M$. Assume further that a is at s and that the daemon is forwarding an LI message to a , i.e. the above is in parallel with

$$LP \stackrel{\text{def}}{=} @_D \langle a@s \rangle \text{deliver!}[c v \text{ack}]$$

This parallel composition, with a surrounding new-binder for `lock`, expands to

$$\mathbf{new} \text{lock} : \text{ } \wedge^{\text{rw}} \text{Map}[\text{Agent}^s \text{ Site}] \mathbf{in} \\ LQ | @_a \text{deliver!}[c v \text{ack}]$$

The proof of this expansion relies on the fact that the reductions of LP cannot release `lock`, so a cannot migrate, and hence the reductions of LP are deterministic, successfully delivering the message to a at s . It uses the following lemma.

LEMMA 9.3

Given that LQ is temporarily immobile under l with respect to Γ, Δ and M , with Δ extensible and $l \in \text{dom}(\Delta)$, if $\Gamma, \Delta \Vdash LP_1 \xrightarrow[M]{\text{det}} LP_2$ then

$$\mathbf{new} \Delta \mathbf{in} LP_1 | LQ \underset{\Gamma}{\succeq} \overset{M \cap \text{dom}(\Gamma)}{\cdot} \mathbf{new} \Delta \mathbf{in} LP_2 | LQ$$

Proofs of temporary immobility can be hard, since they involve quantification over derivatives. We formulate a coinductive definition of temporary immobility (which is equivalent to the one given here). A process is temporarily immobile if it belongs to a blocking set: a set which is closed under transitions that are not inputs on the lock channel, and in which no migration can occur. This alternative definition allows temporary immobility to be proved by analysing single step transitions. Moreover, since temporary immobility is preserved by weak bisimulation, we may apply “up to” techniques [Sangiorgi and Milner 1992], so that we may work with sets which are a blocking set when closed up under weak bisimulation. Proving

that the process LQ given above is temporarily immobile, for example, involves analysing its transitions, which can be classified into two groups:

- local computation: execution of the process P in a , which does not involve the daemon. The result of this type of transition is in the same form as LQ .
- daemon computation: execution of the process P in a , which involves the daemon. The result of this type of transition expands a process which is of the same form as LQ . (Sending location-dependent message to the static daemon, for example, induces expansion.)

Temporary immobility is preserved by parallel composition and **new** binding. This can be used for proving that the process below is temporarily immobile.

$$\begin{aligned}
 LR = \mathbf{new} \Omega_{aux} \mathbf{in} \\
 & @_D \mathit{Daemon} \\
 & | @_{b_1} (\llbracket P_1 \rrbracket_{b_1} | \mathit{currentloc!s_1} | \mathit{Deliverer}) | \dots \\
 & | @_{b_n} (\llbracket P_n \rrbracket_{b_n} | \mathit{currentloc!s_n} | \mathit{Deliverer})
 \end{aligned}$$

Since LR is strongly bisimilar to $LQ_1 | \dots | LQ_n$, where LQ_i is obtained from LQ by replacing the name of the agent a by b_i and the process P by P_i . The proof of the strong bisimulation uses a result similar to a proposal of Milner [1993, p.29].

10. CORRECTNESS: PROOF FOR THE CENTRAL FORWARDING SERVER

This section outlines the strategies taken in order to prove the correctness of the example CFS encoding $\mathcal{C}[\cdot]$, defined in Section 3, using the techniques from Section 9.

10.1 Factoring the Proof

We simplify the construction of the main coupled simulation (between an arbitrary source program, in $n\pi_{LD,LI}$, and its encoding, in $n\pi_{LD}$) by factoring the encoding through an *intermediate language* IL , with states ranged over by Sys , that is specific to this encoding. The infrastructure encoding $\mathcal{C}[\cdot]$ is factored into the composition of a *loading* encoding \mathcal{L} , mapping source terms to corresponding systems in the intermediate language, and an *unloading* encoding \mathcal{F} , mapping systems in the intermediate language to their corresponding target terms.

$$\begin{array}{ccc}
 n\pi_{LD,LI} & \xrightarrow{\mathcal{L}[\cdot]} & IL \\
 & \searrow \mathcal{C}[\cdot] & \downarrow \mathcal{F}[\cdot] \\
 & & n\pi_{LD}
 \end{array}$$

In proving correctness of the loading encoding, we essentially deal with all the house-keeping steps, relating terms introduced by such steps to some normal forms. Such normal forms allow house-keeping steps to be abstracted away, so that in proving correctness of the unloading encoding, we can concentrate on relating partially committed terms to target-level terms. This helps us manage the complexity of the state-space of the encoding, by:

- (1) reducing the size of the coupled simulation relations, omitting states which reduce by house-keeping steps to certain normal forms (which have no house-keeping steps);
- (2) dealing with states in which many agents may be partially committed simultaneously; and
- (3) capturing some invariants, e.g. that the daemon's site-map is correct, in a type system for IL.

The cost is that the typing and labelled transition rules for IL must be defined. For lack of space we only outline the essential points here, referring the reader again to Unyapoth [2001] for the full development.

We use two functions mapping intermediate language states back into the source language. The *undo* and *commit* decoding functions, \mathcal{D}^b and \mathcal{D}^\sharp respectively, undo and complete partially committed migrations.

$$\mathfrak{n}\pi_{\text{LD,LI}} \xrightleftharpoons[\mathcal{D}^\sharp[\cdot]]{\mathcal{D}^b[\cdot]} \text{IL}$$

It suffices to have both functions commit creations and LI messages, as these are somewhat confluent.

We shall not define the loading, unloading and decoding functions here. Instead we illustrate the correspondence between steps in the source, intermediate and the target languages in the creation, migration and location-independent messaging cases in Figure 8. In the figure, some τ communication steps are annotated with the command or the name of the channel involved. The figure also shows how partially committed states are mapped to terms in the source language by the decoding functions.

10.2 Intermediate Language

Each term of the intermediate language represents a normal form of target-level derivatives, possibly in a partially committed state. It describes the state of the daemon as well as that of the encoded agent. The syntax is:

$$\text{Sys} ::= \text{eProg}(\Delta; \mathbf{D}; \mathbf{A})$$

Each term $\text{eProg}(\Delta; \mathbf{D}; \mathbf{A})$ is parameterised by Δ , a located type context corresponding to all names dynamically created during the execution of the program, and \mathbf{D} and \mathbf{A} , the state of the daemon and of the agents. Δ is binding in $\text{eProg}(\Delta; \mathbf{D}; \mathbf{A})$ and is therefore subject to alpha-conversion. The latter two parameters are described in more detail below:

—The state \mathbf{D} of the daemon is described by the following syntax:

$$\begin{aligned} \mathbf{D} & ::= [\text{map } \text{mesgQ}] \\ \text{mesgQ} & ::= \prod_{i \in I} \text{mesgReq}(\{T_i\} [a_i \ c_i \ v_i]) \end{aligned}$$

Each daemon state $[\text{map } \text{mesgQ}]$ consists of a site map map , expressed as a list of pairs, and an unordered queue of message forwarding requests mesgQ . A message forwarding request $\text{mesgReq}(\{T\} [a \ c \ v])$ requires the daemon to forward $c!v$ to the agent a , where T is the type of v .




Fig. 8. Relationships Between Source, Intermediate, and Target

—The state \mathbf{A} of the agents is a partial function mapping agent names to agent states. Each agent state, represented as $[P \mathbf{E}]$, consists of a *main body* P and a *pending state* \mathbf{E} . The syntax of \mathbf{E} is given below:

$$\begin{aligned} \mathbf{E} ::= & \text{FreeA}(s) \mid \text{RegA}(b \ Z \ s \ P \ Q) \\ & \mid \text{MtingA}(s \ P) \mid \text{MrdyA}(s \ P) \end{aligned}$$

If an agent a has pending state $\text{FreeA}(s)$, the local lock of a is free and is ready to initiate a **create** or **migrate to** process from its main body. Otherwise, a is in a partially committed state, with a pending execution of **create**^Z $b = P$ **in** Q (when its state is $\text{RegA}(b \ Z \ s \ P \ Q)$) or **migrate to** $s \rightarrow P$ (when its state is $\text{MtingA}(s \ P)$ or $\text{MrdyA}(s \ P)$). In $\text{FreeA}(s)$ and $\text{RegA}(b \ Z \ s \ P \ Q)$, s denotes the current site of a , internally recorded and maintained by the agent itself.

In $\text{RegA}(b \ Z \ s \ P \ Q)$, the name b is bound in P and Q and is subject to alpha-conversion.

Informally, each transition of a system originates either from an agent or the daemon. A process from the main body of an agent may be executed immediately if it is either an **iflocal**, **if**, **let** or a pair of an output and a (replicated) input on the same channel. The result of such an execution (governed by $n\pi_{\text{LD,LI}}$ LTS rules) is placed in parallel with other processes in the main body, except for execution of an LI output $\langle b \rangle c!v$, which results in the message forwarding request $\text{msgReq}(\{T\} [b \ c \ v])$ being added to the message queue of the daemon (T is the type of v). These steps correspond exactly to those taken by source- and target-level terms. A process **create**^Z $b = P$ **in** Q or **migrate to** $s \rightarrow P$ from the main body of a may proceed (in fact initiate) if the local lock is free, i.e. the pending state is $\text{FreeA}(s')$. The result of such initiation turns the pending state to $\text{RegA}(b \ Z \ s' \ P \ Q)$ or $\text{MtingA}(s \ P)$ respectively. Translating into target-level terms, an agent in such a state has successfully acquired its local lock and sent a registration or migrating request to the daemon.

A system with registration request $\text{RegA}(b \ Z \ s \ P \ Q)$ is executed in a single reduction step, corresponding in the target-level to acquiring the daemon lock, updating the site map and sending the acknowledgement to b . After completion, the declaration $b : \text{Agent}^Z @s$ is placed at the top level and, at the same time, the site map is extended with the new entry (b, s) . The new agent b with state $[P \text{FreeA}(s)]$ now commences its execution, and so does its parent. The top third of Figure 8 gives the correspondences between steps in the source, intermediate and the target languages in the creation case. In the figure, some τ communication steps are annotated with the command or the name of the channel involved.

Likewise, a system with a message forwarding request $\text{msgReq}(\{T\} [b \ c \ v])$ is executed in a single reduction step, corresponding in the target-level to acquiring the daemon lock, looking up the site of b , delivering the message, and receiving an acknowledgement from b . After completion, the message $c!v$ is added to the main body of b .

Serving a migrating request $\text{MtingA}(s \ P)$ from an agent a , however, involves two steps. The first step acquires the daemon lock, initialising the request and turning the pending state of a to $\text{MrdyA}(s \ P)$. In the second step, the agent a migrates to s (hence changes the top-level declaration) and the site map updates a with the

entry (a, s) . The first step corresponds in the target-level to acquiring the daemon lock, looking up the site of a in the site map, and sending an acknowledgement, permitting a to migrate. The second step corresponds to a migrating to s and sending an acknowledgement back to the daemon, which updates its site map and then sends the final acknowledgement to a , allowing it to proceed.

10.3 Proof Outline

Note that our encoding is not *uniform* [Palamidessi 1997], as it introduces a centralised daemon at top level. This means that our reasoning must largely be about the whole system, dealing with interactions between encoded agents and the daemon. We cannot use simple induction on source program syntax.

We prove the coupled simulation over programs which are well-typed with respect to a *valid system context*: a type context in which all agents are declared as static (in order to use the standard definition of coupled simulation) and channels are not used for sending or receiving agent names (in order to make sure the daemon has a record of all agents in the system). Dynamically created new-bound agents may be mobile, of course.

The main lemmas can now be stated.

LEMMA 10.1 (SYNTACTIC FACTORISATION)

For any LP well-typed with respect to a valid system context Φ

$$\begin{aligned} -\mathcal{C}_\Phi \llbracket LP \rrbracket &\equiv \mathcal{F} \llbracket \mathcal{L}_\Phi \llbracket LP \rrbracket \rrbracket, \text{ and} \\ -LP &\equiv \mathcal{D}^\flat \llbracket \mathcal{L}_\Phi \llbracket LP \rrbracket \rrbracket \equiv \mathcal{D}^\sharp \llbracket \mathcal{L}_\Phi \llbracket LP \rrbracket \rrbracket. \end{aligned}$$

This follows from the definitions of the encoding and decoding functions.

LEMMA 10.2 (SEMANTIC CORRECTNESS OF IL)

For any Sys well-formed with respect to Φ , $\mathcal{F} \llbracket Sys \rrbracket \dot{\succeq}_\Phi Sys$.

This lemma is the heart of the correctness argument. The proof uses expansion up to expansion to relate each well-formed term in the intermediate language with its corresponding target term. We use the techniques of Section 9 — part of the reasoning for the LI message-delivery case was outlined there. In broad, we heavily employ the congruence properties of translocating expansion for factoring out program contexts which are not involved in house-keeping reductions of the target terms. Temporary immobility is used whenever we need to guarantee that location-dependent messages to partially committed agents are safely delivered.

The following two lemmas relate intermediate language states to source terms, by weak simulation relations using either the undo or commit decodings. Their proofs are relatively straightforward.

LEMMA 10.3 (\mathcal{D}^\flat IS A STRICT SIMULATION)

For any Sys well-formed for Φ , if $\Phi \Vdash Sys \xrightarrow[\Xi]{\beta} Sys'$ then $\Phi \Vdash \mathcal{D}^\flat \llbracket Sys \rrbracket \xrightarrow[\Xi]{\hat{\beta}} \mathcal{D}^\flat \llbracket Sys' \rrbracket$.

LEMMA 10.4 ($\mathcal{D}^{\sharp^{-1}}$ IS A PROGRESSING SIMULATION)

For any Sys well-formed with respect to Φ , if $\Phi \Vdash \mathcal{D}^\sharp \llbracket Sys \rrbracket \xrightarrow[\Xi]{\beta} LP$ then there exists a well-formed state Sys' such that $LP \equiv \mathcal{D}^\sharp \llbracket Sys' \rrbracket$ and $\Phi \Vdash Sys \xrightarrow[\Xi]{\beta} Sys'$.

These two lemmas are proved by direct constructions of simulation relations. The analysis of possible transitions is made feasible by the factoring out of housekeeping steps.

These results are combined to give a direct relation between the source and the target terms, proving that a source term LP and its translation $\mathcal{C} \llbracket LP \rrbracket$ are related by a coupled simulation.

THEOREM 10.5 (ENCODING CORRECTNESS)

For any LP well-formed with respect to a valid system context Φ , $LP \Leftarrow_{\Phi} \mathcal{C} \llbracket LP \rrbracket$.

PROOF. The proof puts together the operational correspondence results developed, as can be summarised in the diagram below.

$$\begin{array}{ccc}
 n\pi_{LD,LI} & LP \xrightarrow[(10.1)]{\equiv} \mathcal{D} \llbracket \mathcal{L}_{\Phi} \llbracket LP \rrbracket \rrbracket & \\
 & \vdots & \Leftarrow_{\Phi} \left. \begin{array}{l} (10.3, 10.4) \\ \mathcal{L}_{\Phi} \llbracket LP \rrbracket \end{array} \right\} \\
 IL & \Leftarrow_{\Phi} & \\
 & \vdots & \Leftarrow_{\Phi}^{\emptyset} \left. \begin{array}{l} (10.2) \\ \mathcal{F} \llbracket \mathcal{L}_{\Phi} \llbracket LP \rrbracket \rrbracket \end{array} \right\} \\
 n\pi_{LD} & \mathcal{C}_{\Phi} \llbracket LP \rrbracket \xrightarrow[(10.1)]{\equiv} \mathcal{F} \llbracket \mathcal{L}_{\Phi} \llbracket LP \rrbracket \rrbracket &
 \end{array}$$

□

11. RELATED WORK

A range of work on mobility from different perspectives (process migration within a cluster, mobile computing, and wide-area migration in mobile agent languages) is surveyed in the collection edited by Milojević et al. [1999].

The direct precursors of our work on Nomadic Pict were programming languages closely based on process calculi. The collection of Nielson [1997] describes CML, FACILE, LCS, and the Poly/ML concurrency primitives, all of which draw on channel based communication as in Milner's CCS [1989]. With the exception of FACILE, these are focussed on local concurrency, without support for distributed programming. Milner, Parrow, and Walker generalised CCS to the π calculus [1992; 1992], allowing channel names to be themselves sent over channels, and with an elegant operational semantics for fresh generation of new channel names. The π calculus is small but very expressive, allowing data structures, functions, objects, locks, and other constructs of sequential and concurrent programming to be encoded with asynchronous message-passing. This was demonstrated in the Pict language of Pierce and Turner [Pierce and Turner 2000; Turner 1996], which was an experiment in building a concurrent (but again not distributed) programming language based closely on the π calculus, by analogy to the development of functional programming languages such as ML and Haskell above the λ -calculus.

The distributed join-calculus, of Fournet et al. [1996], aimed to redesign the π calculus to make a better foundation for distributed programming, as developed in the subsequent JoCaml programming language [Conchon and Le Fessant 1999]. The distributed join-calculus ensures syntactically that there is a unique receiver for each channel, and then regains expressivity by allowing receivers to synchronise

on multiple messages. It also distributes processes over a hierarchical structure of abstract locations, which may be freshly generated and which may migrate to different points in the hierarchy. In implementations one can think of the first level of this hierarchy as physical machines, with lower levels as migratable running computations. Implementations had an elaborate overlay network built-in, with forwarding pointer chains (as in our algorithm of §5) and mechanisms to collapse those chains. The hidden complexity of this algorithm, and the fact that its behaviour under failure had either to be exposed to the programmer or concealed by a high-level semantics in which reconnection was prohibited, was the immediate spur for our development of the lower level of abstraction of low-level Nomadic Pict, in which the semantics under failure is clear and in which one can see and analyse the design of such higher-level algorithms.

The π calculus is an attractive starting point for calculi for distributed computation, from its clear treatment of concurrency, the elegant treatment of names, and the similarity between π asynchronous message passing and asynchronous network communication¹. This led to a wide variety of distributed process calculi, adding notions of distribution, locality, mobility and security. Some parts of the rather large design space are surveyed in Sewell [2000] and Cardelli [1999], and we mention a few prominent examples below.

- The early π_l calculus of Amadio and Prasad [1994], used for modelling the notions of locality and failure presented in the programming language FACILE [Thomsen et al. 1996].
- The dpi of Sewell [1998], used for studying a type system in which the input and output capabilities of channels may be either global or local.
- The Seal calculus of Vitek, Castagna, and Zappa Nardelli [Vitek and Castagna 1998; Castagna et al. 2005], intended as a framework for writing secure distributed applications over large scale open networks such as the Internet, and the Box- π calculus of Sewell and Vitek [2003], used for studying *wrappers*: secure environments that provide fine-grain control of the allowable interaction between them, and between components and other system resources.
- The various $D\pi$ calculi of Hennessy, Riely, and others [Riely and Hennessy 1998; Hennessy 2007], used for studying *partially typed* semantics, designed for mobile agents in open distributed systems in which some sites may harbor malicious intentions. These typically address code mobility but not computation mobility, with a focus on type-based enforcement of desirable properties.
- The Ambient calculus of Cardelli and Gordon [1998], a calculus for describing the movement of processes and devices, including movement through administrative domains. This prompted further work on semantics, e.g. [Merro and Nardelli 2005], and several variant calculi.
- The extension of TyCO with distribution and code mobility [Vasconcelos et al.

¹In practice, one would typically implement π -style asynchronous messaging above TCP connections, not UDP, as UDP does not provide retransmission and has a fixed upper-bound datagram size. In the absence of migration, such an implementation would provide a FIFO property that is not reflected in the π calculus semantics.

1998], a name-passing process calculus, which allows asynchronous communication between concurrent objects via labelled messages carrying names.

The above systems address a variety of distributed-systems problems with semantically well-founded approaches, generally focussing on the dynamics of interaction (as one would expect from their process-calculus origins) and in some cases on type systems. One can also take a more logical view, as in the P2 system [Loo et al. 2005] and SD3 [Jim 2001], both of which describe distributed algorithms declaratively. Murphy [2008] develops a language based on a Curry-Howard correspondence for a modal logic, focussing on type-theoretic guarantees that mobile code will never access resources that are not present at the current site.

There are many related programming languages, not based on a π calculus semantics but supporting some form of mobility, including Kali Scheme [Cejtin et al. 1995], Obliq [Cardelli 1995], and Mozart [Van Roy and Haridi 2004].

As for mobility at the virtual machine level, Xen live migration [Clark et al. 2005] deals with the special case of migration over a single switched LAN. In that setting, one can arrange for the migrating VM to carry its IP address with it, with an unsolicited ARP reply. This results in the loss of some in-flight IP packets, but (as higher-level protocols such as TCP must be resilient to such loss in any case) the migration is essentially transparent. Migration in the wide-area setting, without additional support from the IP layer, would presumably need overlay networks of the kind we describe, though perhaps a TCP-connection-based approach would be a better fit to applications than the asynchronous messages that we consider here.

Turning now to verification, Moreau [Moreau 2002; 2001] develops a fault-tolerant directory service and message routing algorithm, based on forwarding pointers, and verifies the correctness of the abstract algorithm (mechanised in Coq). Verification of mobile communication infrastructures has also been considered in the Mobile UNITY setting, by McCann and Roman [1997]. There is, of course, a great deal of other work on verification of distributed algorithms in general, and on proof techniques for π calculi. Roughly speaking, the verification and proof techniques of process calculi can be classified as those based on types and those based on the dynamic behaviour of processes. A type system for the π calculus was first proposed by Milner [Milner 1993], giving the notions of *sort* and *sortings*, which ensure uniformity of the kind of names that can be sent or received by channels. Many refinements on the type system have subsequently been proposed, including polymorphism [Turner 1996; Pierce and Sangiorgi 1997], subtyping [Pierce and Sangiorgi 1996], linear types [Kobayashi et al. 1996], objects [Walker 1995], and a generic type system [Igarashi and Kobayashi 2001]. Adding the notions of locality and distribution to the π calculus admits further refinements to be made. Sewell [1998] formulated *dpi* for studying a type system where each channel is located at an agent and can be given global/local usage capability as well as that for input/output. An approximation to the join-style of interaction, for example, can be obtained by giving them global-output and local-input capabilities. This type system retains the expressiveness of channel communication, yet admits optimisation at compile time. Yoshida and Hennessy [1999] formulated a type system for $D\pi\lambda$ which emulates the join-style of interaction using input/output subtyping. The presence of higher-order processes makes this formulation challenging. The

type system of $D\pi_1^r$ extends the concept of uniform receptiveness [Sangiorgi 1999] to ensure that each output (perhaps an inter-agent message) is guaranteed to react with a (unique) input process at its destination. The techniques of refining channel types are also used in ensuring security-related properties. For example, the partial typing of Riely and Hennessy [1999] ensures that resources of trusted sites are not abused by untrusted sites; Sewell and Vitek introduced causality types [Sewell and Vitek 2003] for reasoning about information flow between security domains; and Cardelli et al. [2000] introduced a notion of *groups* which can be used for ensuring that the boundary of an ambient may only be dissolved by trusted groups of ambients.

The behavioural theories of these distributed variants of process calculi are generally adapted from those of the π calculus, which are based around operational semantics and operational equivalences. A reduction semantics is given for all of the cited calculi. This, together with some notions of barbs, allows a definition of barbed bisimulation to be given, as is the case for the Distributed Join-calculus [Fournet and Gonthier 1996], the Seal calculus [Castagna and Vitek 1999], and the Ambient calculus [Gordon and Cardelli 1999]. A labelled transition semantics is also given for the π_{ll} , $D\pi$, $D\pi_1^r$, Ambient, Seal and Box π calculi, allowing some notions of bisimilarity to be given. These definitions of labelled transition semantics often involve refining that of the standard π with location annotation ($@_l$ for $D\pi_1^r$ and “relative location” tags for Seal and Box- π). The labelled transition semantics of $D\pi$ [Riely and Hennessy 1998] extends the standard π input and output actions with labels that indicate movements and failures of locations. The style of transition systems for the Ambient calculus [Gordon and Cardelli 1999; Merro and Nardelli 2005] is quite different from those for π calculus, as they involve relative locations of ambients. The definitions of labelled transition semantics and operational equivalences of distributed CCS [Riely and Hennessy 1997] and π_{ll} [Amadio 1997] also take location failures into account.

Several authors have used process-calculus proof techniques to verify the correctness of implementations or abstract machines for various ambient calculi. Fournet et al. [2000] give a translation of Ambients into the Join calculus. As in our central forwarding server proof, they build an intermediate language to capture intermediate states of the translation and use coupled simulations, though they work in a barbed reduction-semantics setting rather than the labelled transition setting we adopt. They also describe an implementation in JoCaml based on this translation, though with some significant differences. Giannini et al. [2006] give an abstract machine (PAN) for Safe Ambients [Levi and Sangiorgi 2000], a restricted calculus in which ambient movement depends on agreement between both parties, and ambients are either single-threaded or immobile. This is rather different from Nomadic π , in which an agent can migrate to another site at any time. They prove the abstract machine has the same barbs as the source. Hirschhoff et al. [2007] refine this abstract machine, optimising the treatment of forwarders, and prove it weakly bisimilar to PAN. They also describe an OCaml implementation loosely based on their abstract machine.

The work on Nomadic Pict described in the current paper led to two substantial subsequent lines of research. Firstly, in a production language, one would like

to express high-level abstractions such as that of high-level Nomadic Pict using a general-purpose module system rather than the special-purpose encodings of the Nomadic Pict implementation. An ML-style module system [Milner et al. 1997] is a good fit for this: one can express the high- and low-level abstractions as signatures, with abstract types of site name, agent name, etc., and operations to send messages, migrate, etc., and express a particular overlay network implementation as a functor from one to the other. However, when one imagines this in a wide-area setting, it quickly becomes obvious that one will need multiple different overlay network implementations, and that they will inevitably exist in multiple simultaneous versions. This observation prompted work on type equality for abstract types in the distributed setting [Sewell 2001; Leifer et al. 2003; Sewell et al. 2005; 2007; Billings et al. 2006; Deniélou and Leifer 2006], and the Acute and HashCaml prototype languages. The former provides a slightly lower level of abstraction than low-level Nomadic Pict: instead of migration, it has a primitive for freezing a group of threads into a thunk (together with support for modules, versions, etc.). This makes it possible to implement low-level Nomadic Pict itself as an Acute module [Sewell et al. 2007, §11], and high-level Nomadic Pict overlays could be implemented as further modules above that.

Secondly, recall that the low-level Nomadic Pict abstraction was designed to be implementable with a clear semantics in the presence of failure (site failure, message loss, or disconnection): each low-level reduction step is implementable with at most one asynchronous inter-site message. Later work took this further, characterising the exact semantics (including failure cases) not for simple asynchronous messages, but instead for the communication primitives provided by the Sockets API to the UDP and TCP protocols [Serjantov et al. 2001; Wansbrough et al. 2002; Bishop et al. 2005; 2006; Ridge et al. 2008]. Work by Compton [2005] (above that UDP model) and Ridge [2009] (above a simplified TCP model) demonstrates that it is feasible to verify, fully formally, executable distributed code above such models.

12. CONCLUSION

We have studied the overlay networks required for communication between mobile computations. By expressing such distributed algorithms as Nomadic Pict encodings, between carefully chosen (and well-defined) levels of abstraction, we have descriptions of them that are:

- executable: one can rapidly prototype the algorithms, and applications written above them in the high-level language;
- concise: with the details of concurrency, locking, name-generation etc. made clear; and
- precise: with a semantics that we can use for formal reasoning and that gives a solid understanding of the primitives for informal reasoning.

We discussed the design space of possible algorithms, and implemented a programming language that lets the algorithms (and applications above them) be executed. We developed semantics and proof techniques for proving correctness of such algorithms. The techniques were illustrated by a proof that an example algorithm is

correct with respect to coupled simulation. This algorithm, though non-trivial, is relatively simple, but we believe that more sophisticated algorithms could be dealt with using the same techniques (albeit with new intermediate languages, tailored to particular algorithms).

More generally, the work is a step towards semantically-founded engineering of wide-area distributed systems. Here we dealt with the combination of migration and communication, and for a complete treatment one must also simultaneously address failure and malicious attack.

ELECTRONIC APPENDIX

The electronic appendix for this article can be accessed in the ACM Digital Library by visiting the following URL: <http://www.acm.org/pubs/citations/journals/toplas/20YY-V-N/p1-URLend>.

ACKNOWLEDGMENTS

The initial work on the Nomadic Pict project was in collaboration with Benjamin C. Pierce [Sewell et al. 1998; 1999]. We also owe special thanks to Benjamin and to David N. Turner for allowing us to use the source code of Pict, and thank Robin Milner and Ken Moody for their support. This work was funded in part by EPSRC grants GR/L62290, GR/N24872, GR/T11715, EP/C510712, EP/F036345, and EP/H005633, a Wolfson Foundation Scholarship (Wojciechowski), a Royal Thai Government Scholarship (Unyapoth) and a Royal Society University Research Fellowship (Sewell).

REFERENCES

- AMADIO, R. M. 1997. An asynchronous model of locality, failure, and process mobility. In *COORDINATION '97: Proceedings of the Second International Conference on Coordination Languages and Models, LNCS 1282*. Springer-Verlag, London, UK, 374–391.
- AMADIO, R. M. AND PRASAD, S. 1994. Localities and failures (extended abstract). In *FSTTCS '94: Proceedings of the 14th Conference on Foundations of Software Technology and Theoretical Computer Science, LNCS 880*. Springer-Verlag, London, UK, 205–216.
- APPEL, A. W. 1992. *Compiling with Continuations*. Cambridge University Press.
- ARNOLD, K., WOLLRATH, A., O'SULLIVAN, B., SCHEIFLER, R., AND WALDO, J. 1999. *The Jini specification*. Addison-Wesley, Reading, MA, USA.
- AWERBUCH, B. AND PELEG, D. 1995. Online tracking of mobile users. *Journal of the ACM* 42, 5 (Sept.), 1021–1058. A shorter version appeared as Technical Memo MIT/LCS/TM-410, Massachusetts Institute of Technology, Laboratory for Computer Science, July 1989.
- BALLINTJN, G., VAN STEEN, M., AND TANENBAUM, A. 1999. Simple crash recovery in a wide-area location service. In *PDCS '99: Proceedings of the 11th IASTED International Conference on Parallel and Distributed Computing Systems, Fort Lauderdale, FL, USA*. IASTED, 87–93.
- BILLINGS, J., SEWELL, P., SHINWELL, M., AND STRNIŠA, R. 2006. Type-safe distributed programming for OCaml. In *ML '06: Proceedings of the 2006 ACM SIGPLAN Workshop on ML*. ACM, New York, NY, USA, 20–31.
- BISHOP, S., FAIRBAIRN, M., NORRISH, M., SEWELL, P., SMITH, M., AND WANSBROUGH, K. 2005. Rigorous specification and conformance testing techniques for network protocols, as applied to TCP, UDP, and sockets. In *SIGCOMM '05: Proceedings of the ACM SIGCOMM 2005 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications*. ACM, New York, NY, USA, 265–276.

- BISHOP, S., FAIRBAIRN, M., NORRISH, M., SEWELL, P., SMITH, M., AND WANSBROUGH, K. 2006. Engineering with logic: HOL specification and symbolic-evaluation testing for TCP implementations. In *POPL '06: Proceedings of the 33rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 55–66.
- CARDELLI, L. 1995. A language with distributed scope. In *POPL '95: Proceedings of the 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 286–297.
- CARDELLI, L. 1999. Abstractions for mobile computation. In *Secure Internet Programming: Security Issues for Mobile and Distributed Objects, LNCS 1603*, J. Vitek and C. D. Jensen, Eds. State-of-the-Art Survey. Springer-Verlag, London, UK, 51–94.
- CARDELLI, L., GHELLI, G., AND GORDON, A. D. 2000. Ambient groups and mobility types. In *TCS '00: Proceedings of the 16th IFIP International Conference on Theoretical Computer Science, Exploring New Frontiers of Theoretical Informatics, LNCS 1872*. Springer-Verlag, London, UK, 333–347.
- CARDELLI, L. AND GORDON, A. D. 1998. Mobile ambients. In *FoSSaCS '98: Proceedings of the First International Conference on Foundations of Software Science and Computation Structure, LNCS 1378*. Springer-Verlag, London, UK, 140–155.
- CARDELLI, L., GORDON, A. D., AND GHELLI, G. 1999. Mobility types for mobile ambients. In *ICALP '99: Proceedings of the 26th International Colloquium on Automata, Languages and Programming, LNCS 1644*. Springer-Verlag, London, UK, 230–239.
- CASTAGNA, G. AND VITEK, J. 1999. Commitment and confinement for the Seal calculus. Trusted objects, Centre Universitaire d'Informatique, University of Geneva. July.
- CASTAGNA, G., VITEK, J., AND ZAPPA NARDELLI, F. 2005. The Seal calculus. *Information and Computation* 201, 1, 1–54.
- CEJTIN, H., JAGANNATHAN, S., AND KELSEY, R. 1995. Higher-order distributed objects. *ACM Transactions on Programming Languages and Systems (TOPLAS)* 17, 5 (September), 704–739.
- CHANDRA, T. D. AND TOUEG, S. 1996. Unreliable failure detectors for reliable distributed systems. *Journal of the ACM* 43, 2, 225–267.
- CHERITON, D. 1988. The V distributed system. *Communications of the ACM* 31, 3 (Mar.), 314–333.
- CLARK, C., FRASER, K., HAND, S., HANSEN, J. G., JUL, E., LIMPACH, C., PRATT, I., AND WARFIELD, A. 2005. Live migration of virtual machines. In *NSDI '05: Proceedings of the 2nd USENIX/ACM Symposium on Networked Systems Design and Implementation*. Boston, MA, 273–286.
- COMPTON, M. 2005. Stenning's protocol implemented in UDP and verified in Isabelle. In *CATS '05: Proceedings of the 2005 Australasian Symposium on Theory of Computing*. Australian Computer Society, Inc., Darlinghurst, Australia, 21–30.
- CONCHON, S. AND LE FESSANT, F. 1999. Jocaml: mobile agents for Objective-Caml. In *ASA/MA '99: Proceedings of the First International Symposium on Agent Systems and Applications/Third International Symposium on Mobile Agents*. IEEE Computer Society, Palm Springs, CA, USA, 22–29.
- DE NICOLA, R. AND HENNESSY, M. C. B. 1984. Testing equivalences for processes. *Theoretical Computer Science* 34, 1-2 (Nov.), 83–133.
- DEMMER, M. J. AND HERLIHY, M. P. 1998. The arrow distributed directory protocol. In *DISC '98: Proceedings of the 12th International Symposium on Distributed Computing, Andros, Greece, LNCS 1499*. Springer-Verlag, 119–133.
- DENIÉLOU, P.-M. AND LEIFER, J. J. 2006. Abstraction preservation and subtyping in distributed languages. In *ICFP '06: Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming*. ACM, New York, NY, USA, 286–297.
- DOUGLIS, F. AND OUSTERHOUT, J. 1991. Transparent process migration: Design alternatives and the Sprite implementation. *Software – Practice and Experience* 21, 8 (Aug.), 757–785.
- FISCHER, M. J., LYNCH, N. A., AND PATERSON, M. S. 1985. Impossibility of distributed consensus with one faulty processor. *Journal of the ACM* 32, 2, 374–382.

- FOURNET, C. AND GONTHIER, G. 1996. The reflexive CHAM and the join-calculus. In *POPL '96: Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM Press, New York, USA, 372–385.
- FOURNET, C., GONTHIER, G., LÉVY, J.-J., MARANGET, L., AND RÉMY, D. 1996. A calculus of mobile agents. In *CONCUR '96: Proceedings of the 7th International Conference on Concurrency Theory, LNCS 1119*. Springer-Verlag, London, UK, 406–421.
- FOURNET, C., LÉVY, J.-J., AND SCHMITT, A. 2000. An asynchronous, distributed implementation of mobile ambients. In *TCS '00: Proceedings of the 16th IFIP International Conference on Theoretical Computer Science, Exploring New Frontiers of Theoretical Informatics, LNCS 1872*. Springer-Verlag, London, UK, 348–364.
- GIANNINI, P., SANGIORGI, D., AND VALENTE, A. 2006. Safe Ambients: Abstract machine and distributed implementation. *Science of Computer Programming* 59, 3, 209–249.
- GORDON, A. D. AND CARDELLI, L. 1999. Equational properties of mobile ambients. In *FoSSaCS '99: Proceedings of the Second International Conference on Foundations of Software Science and Computation Structure, LNCS 1578*. Springer-Verlag, London, UK, 212–226.
- GUERRAOU, R. AND SCHIPER, A. 1996. Fault-tolerance by replication in distributed systems. In *Ada-Europe '96: Proceedings of the 1996 Ada-Europe International Conference on Reliable Software Technologies, LNCS 1088*. Springer-Verlag, London, UK, 38–57.
- HENNESSY, M. 2007. *A Distributed Pi-calculus*. Cambridge University Press.
- HIRSCHKOFF, D., POUS, D., AND SANGIORGI, D. 2007. An efficient abstract machine for Safe Ambients. *Journal of Logic and Algebraic Programming* 71, 2, 114–149.
- IGARASHI, A. AND KOBAYASHI, N. 2001. A generic type system for the pi-calculus. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 128–141.
- JIM, T. 2001. SD3: A trust management system with certified evaluation. In *SP '01: Proceedings of the 2001 IEEE Symposium on Security and Privacy*. IEEE Computer Society, Washington, DC, USA, 106–115.
- JUL, E., LEVY, H., HUTCHINSON, N., AND BLACK, A. 1988. Fine-grained mobility in the Emerald system. *ACM Transactions on Computer Systems* 6, 1 (Feb.), 109–133.
- KOBAYASHI, N., PIERCE, B. C., AND TURNER, D. N. 1996. Linearity and the pi-calculus. In *POPL '96: Proceedings of the 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 358–371.
- LANGE, D. B. AND ARIDOR, Y. 1997. *Agent Transfer Protocol – ATP/0.1*. IBM Tokyo Research Laboratory.
- LEIFER, J. J., PESKINE, G., SEWELL, P., AND WANSBROUGH, K. 2003. Global abstraction-safe marshalling with hash types. In *ICFP '03: Proceedings of the 8th ACM SIGPLAN International Conference on Functional Programming*. ACM, New York, NY, USA, 87–98.
- LEROY, X. 1995. Le système Caml Special Light: modules et compilation efficace en Caml. Technical Report RR-2721, INRIA, Institut National de Recherche en Informatique et en Automatique.
- LEVI, F. AND SANGIORGI, D. 2000. Controlling interference in Ambients. In *POPL '00: Proceedings of the 27th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 352–364.
- LOO, B. T., CONDIE, T., HELLERSTEIN, J. M., MANIATIS, P., ROSCOE, T., AND STOICA, I. 2005. Implementing declarative overlays. *ACM SIGOPS Operating Systems Review (Proceedings of the 20th ACM Symposium on Operating Systems Principles, Brighton, UK)* 39, 5, 75–90.
- MCCANN, P. J. AND ROMAN, G.-C. 1997. Mobile UNITY coordination constructs applied to packet forwarding for mobile hosts. In *COORDINATION '97: Proceedings of the Second International Conference on Coordination Languages and Models, LNCS 1282*. Springer-Verlag, London, UK, 338–354.
- MERRO, M. AND NARDELLI, F. Z. 2005. Behavioral theory for mobile ambients. *Journal of the ACM* 52, 6, 961–1023.
- MILNER, R. 1989. *Communication and Concurrency*. International Series in Computer Science. Prentice Hall.

- MILNER, R. 1992. Functions as processes. *Journal of Mathematical Structures in Computer Science* 2, 2, 119–141. Previous version as Rapport de Recherche 1154, INRIA Sophia-Antipolis, 1990, and in *Proceedings of ICALP '91*, LNCS 443.
- MILNER, R. 1993. The polyadic π -calculus: A tutorial. In *Logic and Algebra of Specification*, F. L. Bauer, W. Brauer, and H. Schwichtenberg, Eds. Series F: Computer and System Sciences, vol. 94. NATO Advanced Study Institute, Springer-Verlag. Available as Technical Report ECS-LFCS-91-180, University of Edinburgh, October 1991.
- MILNER, R., PARROW, J., AND WALKER, D. 1992. A calculus of mobile processes, Parts I and II. *Information and Computation* 100, 1, 1–77.
- MILNER, R., TOFTE, M., HARPER, R., AND MACQUEEN, D. 1997. *The Definition of Standard ML (Revised)*. The MIT Press.
- MILOJČIĆ, D., DOUGLIS, F., AND WHEELER, R., Eds. 1999. *Mobility: Processes, Computers, and Agents*. Addison-Wesley, Reading, MA, USA.
- MOREAU, L. 2001. Distributed Directory Service and Message Router for Mobile Agents. *Science of Computer Programming* 39, 2–3, 249–272.
- MOREAU, L. 2002. A fault-tolerant directory service for mobile agents based on forwarding pointers. In *SAC '02: Proceedings of the 17th ACM Symposium on Applied Computing*. ACM, New York, NY, USA, 93–100.
- MULLENDER, S. J. AND VITÁNYI, P. M. B. 1988. Distributed match-making. *Algorithmica* 3, 367–391.
- MURPHY, VII, T. 2008. Modal types for mobile code. Ph.D. thesis, Carnegie Mellon University. Available as Tech. Rep. CMU-CS-08-126.
- NEEDHAM, R. M. 1989. Names. In *Distributed Systems*, S. Mullender, Ed. Addison-Wesley, 89–101.
- NESTMANN, U. 1996. On determinacy and nondeterminacy in concurrent programming. Ph.D. thesis, Technische Fakultät, Universität Erlangen.
- NESTMANN, U. AND PIERCE, B. C. 1996. Decoding choice encodings. In *CONCUR '96: Proceedings of the 7th International Conference on Concurrency Theory*, LNCS 1119. Springer-Verlag, London, UK, 179–194.
- NIELSON, F., Ed. 1997. *ML with Concurrency: Design, Analysis, Implementation, and Application*. Monographs in Computer Science. Springer-Verlag, New York, NY, USA.
- PALAMIDESI, C. 1997. Comparing the expressive power of the synchronous and the asynchronous π -calculus. In *POPL '97: Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 256–265.
- PARROW, J. AND SJÖDIN, P. 1992. Multiway synchronization verified with coupled simulation. In *CONCUR '92: Proceedings of the Third International Conference on Concurrency Theory*, LNCS 630. Springer-Verlag, London, UK, 518–533.
- PIERCE, B. C. AND SANGIORGI, D. 1996. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science* 6, 5, 409–454. An extract appeared in the *Proceedings of LICS '93*: 376–385.
- PIERCE, B. C. AND SANGIORGI, D. 1997. Behavioral equivalence in the polymorphic pi-calculus. In *POPL '97: Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 242–255.
- PIERCE, B. C. AND TURNER, D. N. 1995. Concurrent objects in a process calculus. In *Theory and Practice of Parallel Programming: Proceedings of the International Workshop TPPP '94, Sendai, Japan, 1994*, LNCS 907, T. Ito and A. Yonezawa, Eds. Springer-Verlag, 187–215.
- PIERCE, B. C. AND TURNER, D. N. 1997. *Pict Language Definition*. Available electronically as part of the Pict distribution.
- PIERCE, B. C. AND TURNER, D. N. 2000. Pict: A programming language based on the pi-calculus. In *Proof, Language and Interaction: Essays in Honour of Robin Milner*, G. Plotkin, C. Stirling, and M. Tofte, Eds. Foundations of Computing. MIT Press, 455–494.
- POPEK, G. J. AND WALKER, B. J. 1986. *The LOCUS Distributed System Architecture*. Computer Systems Series. MIT Press, Cambridge, MA, USA.

- RIDGE, T. 2009. Verifying distributed systems: the operational approach. In *POPL '09: Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 429–440.
- RIDGE, T., NORRISH, M., AND SEWELL, P. 2008. A rigorous approach to networking: TCP, from implementation to protocol to service. In *FM '08: Proceedings of the 15th International Symposium on Formal Methods, LNCS 5014*. Springer-Verlag, Berlin, Heidelberg, 294–309.
- RIELY, J. AND HENNESSY, M. 1997. Distributed processes and location failures (extended abstract). In *ICALP '97: Proceedings of the 24th International Colloquium on Automata, Languages and Programming, LNCS 1256*. Springer-Verlag, London, UK, 471–481.
- RIELY, J. AND HENNESSY, M. 1998. A typed language for distributed mobile processes (extended abstract). In *POPL '98: Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 378–390.
- RIELY, J. AND HENNESSY, M. 1999. Trust and partial typing in open systems of mobile agents. In *POPL '99: Proceedings of the 26th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 93–104.
- SANGIORGI, D. 1999. The name discipline of uniform receptiveness. *Theoretical Computer Science* 221, 1–2, 457–493. An abstract appeared in the *Proceedings of ICALP '97*, LNCS 1256, pages 303–313.
- SANGIORGI, D. AND MILNER, R. 1992. The problem of “weak bisimulation up to”. In *CONCUR '92: Proceedings of the Third International Conference on Concurrency Theory, LNCS 630*. Springer-Verlag, London, UK, 32–46.
- SERJANTOV, A., SEWELL, P., AND WANSBROUGH, K. 2001. The UDP calculus: Rigorous semantics for real networking. In *TACS '01: Proceedings of the 4th International Symposium on Theoretical Aspects of Computer Software, Sendai, Japan, LNCS 2215*. Springer-Verlag, London, UK, 535–559.
- SEWELL, P. 1997. On implementations and semantics of a concurrent programming language. In *CONCUR '97: Proceedings of the 8th International Conference on Concurrency Theory, LNCS 1243*. Springer-Verlag, London, UK, 391–405.
- SEWELL, P. 1998. Global/local subtyping and capability inference for a distributed pi-calculus. In *ICALP '98: Proceedings of the 25th International Colloquium on Automata, Languages and Programming, LNCS 1443*. Springer-Verlag, London, UK, 695–706.
- SEWELL, P. 2000. A brief introduction to applied π . Tech. Rep. 498, Computer Laboratory, University of Cambridge, Cambridge, UK.
- SEWELL, P. 2001. Modules, abstract types, and distributed versioning. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 236–247.
- SEWELL, P., LEIFER, J. J., WANSBROUGH, K., ZAPPA NARDELLI, F., ALLEN-WILLIAMS, M., HABOUZIT, P., AND VAFEIADIS, V. 2005. Acute: high-level programming language design for distributed computation. In *ICFP '05: Proceedings of the 10th ACM SIGPLAN International Conference on Functional Programming*. ACM, New York, NY, USA, 15–26.
- SEWELL, P., LEIFER, J. J., WANSBROUGH, K., ZAPPA NARDELLI, F., ALLEN-WILLIAMS, M., HABOUZIT, P., AND VAFEIADIS, V. 2007. Acute: High-level programming language design for distributed computation. *Journal of Functional Programming* 17, 4–5 (July), 547–612. Invited submission for an ICFP 2005 special issue.
- SEWELL, P. AND VITEK, J. 2003. Secure composition of untrusted code: Box- π , wrappers and causality types. *Journal of Computer Security* 11, 2, 135–188. Invited submission for a CSFW '00 special issue.
- SEWELL, P. AND WOJCIECHOWSKI, P. T. 2008. Verifying overlay networks for relocatable computations (or: Nomadic Pict, relocated). In *Proceedings of the Joint HP-MSR Research Workshop on The Rise and Rise of the Declarative Datacentre, Cambridge, UK*. <http://research.microsoft.com/riseandrise> (accessed 13 February 2010).
- SEWELL, P., WOJCIECHOWSKI, P. T., AND PIERCE, B. C. 1998. Location independence for mobile agents. In *Proceedings of IFL 98: the Workshop on Internet Programming Languages, in conjunction with IEEE ICCL '98, Chicago, IL, USA*. 1–6.

- SEWELL, P., WOJCIECHOWSKI, P. T., AND PIERCE, B. C. 1999. Location-independent communication for mobile agents: a two-level architecture. In *Internet Programming Languages, LNCS 1686*. Springer-Verlag, 1–31.
- THOMSEN, B., LETH, L., AND KUO, T.-M. 1996. A Facile tutorial. In *CONCUR '96: Proceedings of the 7th International Conference on Concurrency Theory, LNCS 1119*. Springer-Verlag, London, UK, 278–298.
- TURNER, D. N. 1996. The polymorphic pi-calculus: Theory and implementation. Ph.D. thesis, University of Edinburgh.
- UNYAPOTH, A. 2001. Nomadic π -calculi: Expressing and verifying communication infrastructure for mobile computation. Ph.D. thesis, University of Cambridge. Also appeared as Technical Report UCAM-CL-TR-514, Computer Laboratory, University of Cambridge.
- UNYAPOTH, A. AND SEWELL, P. 2001. Nomadic Pict: correct communication infrastructure for mobile computation. In *POPL '01: Proceedings of the 28th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. ACM, New York, NY, USA, 116–127.
- VAN ROY, P. AND HARIDI, S. 2004. *Concepts, Techniques, and Models of Computer Programming*. MIT Press.
- VAN STEEN, M., HAUCK, F. J., BALLINTIJN, G., AND TANENBAUM, A. S. 1998. Algorithmic design of the Globe wide-area location service. *The Computer Journal* 41, 5, 297–310.
- VASCONCELOS, V. T., LOPES, L., AND SILVA, F. 1998. Distribution and mobility with lexical scoping in process calculi. In *HLCL '98: Proceedings of the 3rd International Workshop on High-Level Concurrent Languages*. Electronic Notes in Theoretical Computer Science (ENTCS), vol. 16.3. Elsevier Science Publishers.
- VITEK, J. AND CASTAGNA, G. 1998. Towards a calculus of secure mobile computations. In *Proceedings of the Workshop on Internet Programming Languages, in conjunction with IEEE ICCL '98. Chicago, IL, USA*.
- WALKER, D. 1995. Objects in the π -calculus. *Information and Computation* 116, 2, 253–271.
- WANSBROUGH, K., NORRISH, M., SEWELL, P., AND SERJANTOV, A. 2002. Timing UDP: Mechanized semantics for sockets, threads, and failures. In *ESOP '02: Proceedings of the 11th European Symposium on Programming Languages and Systems, LNCS 2305*. Springer-Verlag, London, UK, 278–294.
- WOJCIECHOWSKI, P. T. 2000a. *Nomadic Pict. Documentation and User's Manual*. Available as part of the Nomadic Pict distribution from <http://www.cs.put.poznan.pl/pawelw/npict> (accessed 13 February 2010).
- WOJCIECHOWSKI, P. T. 2000b. Nomadic Pict: Language and infrastructure design for mobile computation. Ph.D. thesis, University of Cambridge. Also appeared as Technical Report UCAM-CL-TR-492, Computer Laboratory, University of Cambridge.
- WOJCIECHOWSKI, P. T. 2000-2010. *The Nomadic Pict System*. <http://www.cs.put.poznan.pl/pawelw/npict> (accessed 13 February 2010).
- WOJCIECHOWSKI, P. T. 2001. Algorithms for location-independent communication between mobile agents. In *Proceedings of the 2001 AISB Symposium on Software Mobility and Adaptive Behaviour, York, UK*. Also published as Technical Report IC-2001-13, School of Computer and Communication Sciences, Ecole Polytechnique Fédérale de Lausanne (EPFL).
- WOJCIECHOWSKI, P. T. 2006. Scalable message routing for mobile software assistants. In *EUC '06: Proceedings of the 4th IFIP International Conference on Embedded And Ubiquitous Computing, LNCS 4096*. Springer-Verlag, 355–364.
- WOJCIECHOWSKI, P. T. AND SEWELL, P. 1999. Nomadic Pict: Language and infrastructure design for mobile agents. In *ASA/MA '99: Proceedings of the First International Symposium on Agent Systems and Applications/Third International Symposium on Mobile Agents, Palm Springs, CA, USA*. IEEE Computer Society.
- WOJCIECHOWSKI, P. T. AND SEWELL, P. 2000. Nomadic Pict: Language and infrastructure design for mobile agents. *IEEE Concurrency* 8, 2 (April–June), 42–52. Invited submission for ASA/MA '99.

YOSHIDA, N. AND HENNESSY, M. 1999. Subtyping and locality in distributed higher order processes (extended abstract). In *CONCUR '99: Proceedings of the 10th International Conference on Concurrency Theory, LNCS 1664*. Springer-Verlag, London, UK, 557–572.

Received December 2008; revised June 2009; accepted September 2009.

THIS DOCUMENT IS THE ONLINE-ONLY APPENDIX TO:

Nomadic Pict: Programming Languages, Communication Infrastructure Overlays, and Semantics for Mobile Computation

PETER SEWELL

Computer Laboratory, University of Cambridge

and

PAWEŁ T. WOJCIECHOWSKI

Institute of Computing Science, Poznań University of Technology

and

ASIS UNYAPOTH

Government Information Technology Services, Thailand

ACM Transactions on Programming Languages and Systems, Vol. V, No. N, Month 20YY, Pages 1–62.

A. NOMADIC π CALCULUS TYPE SYSTEM

A.1 Subtyping $\Gamma \vdash S \leq T$

$$\frac{\Gamma \vdash B}{\Gamma \vdash B \leq B} \quad \frac{\Gamma \vdash X}{\Gamma \vdash X \leq X} \quad \frac{\vdash \Gamma \text{ ok}}{\Gamma \vdash \text{Site} \leq \text{Site}}$$

$$\frac{\vdash \Gamma \quad Z \leq Z'}{\Gamma \vdash \text{Agent}^Z \leq \text{Agent}^{Z'}} \quad \frac{I = I' = \text{rw} \Rightarrow (\Gamma \vdash S \leq T \wedge \Gamma \vdash T \leq S) \quad I \leq I' = \text{r} \Rightarrow \Gamma \vdash S \leq T \quad I \leq I' = \text{w} \Rightarrow \Gamma \vdash T \leq S}{\Gamma \vdash \sim^I S \leq \sim^{I'} T}$$

$$\frac{\Gamma, X \vdash S \leq T}{\Gamma \vdash \{X\} S \leq \{X\} T} \quad \frac{\Gamma \vdash S_1 \leq T_1 \quad \dots \quad \Gamma \vdash S_n \leq T_n}{\Gamma \vdash [S_1 \dots S_n] \leq [T_1 \dots T_n]}$$

A.2 Unlocated Type Context Formation $\vdash \Gamma \text{ ok}$

$$\vdash \bullet \text{ ok} \quad \frac{\vdash \Gamma \text{ ok} \quad X \notin \text{dom}(\Gamma)}{\vdash \Gamma, X \text{ ok}} \quad \frac{\vdash \Gamma \text{ ok} \quad \Gamma \vdash T \quad x \notin \text{dom}(\Gamma)}{\vdash \Gamma, x : T \text{ ok}}$$

A.3 Located Type Context Formation $\vdash_L \Gamma \text{ ok}$

$$\vdash_L \bullet \text{ ok} \quad \frac{\vdash_L \Gamma \text{ ok} \quad X \notin \text{dom}(\Gamma)}{\vdash_L \Gamma, X \text{ ok}}$$

$$\frac{\vdash_L \Gamma \quad \Gamma \vdash s \in \text{Site} \quad \Gamma \vdash \text{Agent}^Z \quad x \notin \text{dom}(\Gamma)}{\vdash_L \Gamma, x : \text{Agent}^Z @ s \text{ ok}}$$

$$\frac{\vdash_L \Gamma \quad T \neq \text{Agent}^Z \quad \Gamma \vdash T \quad x \notin \text{dom}(\Gamma)}{\vdash_L \Gamma, x : T \text{ ok}}$$

A.4 Type Formation $\Gamma \vdash T$

$$\frac{\vdash \Gamma \text{ ok } B \in \mathcal{T}}{\Gamma \vdash B} \quad \frac{\vdash \Gamma \text{ ok } Z \in \{\mathbf{m}, \mathbf{s}\}}{\Gamma \vdash \mathbf{Agent}^Z} \quad \frac{\vdash \Gamma \text{ ok}}{\Gamma \vdash \mathbf{Site}}$$

$$\frac{\vdash \Gamma \text{ ok } X \in \text{dom}(\Gamma)}{\Gamma \vdash X} \quad \frac{\Gamma \vdash T \quad I \in \{\mathbf{r}, \mathbf{w}, \mathbf{rw}\}}{\Gamma \vdash \sim^I T}$$

$$\frac{\Gamma \vdash T_1 \quad \dots \quad \Gamma \vdash T_n}{\Gamma \vdash [T_1 \dots T_n]} \quad \frac{\Gamma, X \vdash T \quad X \notin \text{dom}(\Gamma)}{\Gamma \vdash \{X\} T}$$

 A.5 Value and Expression Formation $\Gamma \vdash e \in T$

$$\frac{\vdash \Gamma, x : T, \Gamma' \text{ ok}}{\Gamma, x : T, \Gamma' \vdash x \in T} \quad \frac{\vdash \Gamma \text{ ok } t \in B}{\Gamma \vdash t \in B}$$

$$\frac{\Gamma \vdash e \in S \quad \Gamma \vdash S \leq T}{\Gamma \vdash e \in T} \quad \frac{\Gamma \vdash e \in \{S/X\}T \quad \Gamma \vdash S \quad \Gamma, X \vdash T}{\Gamma \vdash \{S\} e \in \{X\} T}$$

$$\frac{\Gamma \vdash e \in T \quad \Gamma \vdash e' \in T}{\Gamma \vdash e = e' \in \mathbf{Bool}} \quad \frac{\Gamma \vdash e_1 \in T_1 \quad \dots \quad \Gamma \vdash e_n \in T_n}{\Gamma \vdash [e_1 \dots e_n] \in [T_1 \dots T_n]}$$

$$\frac{f \in \mathcal{F} \quad f \text{ is not } = \quad f : B_1 \times \dots \times B_n \rightarrow B \quad \Gamma \vdash e_1 \in B_1 \quad \dots \quad \Gamma \vdash e_n \in B_n}{\Gamma \vdash f(e_1, \dots, e_n) \in B}$$

 A.6 Pattern Formation $\Gamma \vdash p \in T \triangleright \Delta$

$$\frac{\vdash \Gamma \text{ ok } \quad \Gamma \vdash T}{\Gamma \vdash x \in T \triangleright x : T} \quad \frac{\Gamma \vdash p_1 \in T_1 \triangleright \Delta_1 \quad \dots \quad \Gamma \vdash p_n \in T_n \triangleright \Delta_n}{\Gamma \vdash [p_1 \dots p_n] \in [T_1 \dots T_n] \triangleright \Delta_1, \dots, \Delta_n}$$

$$\frac{\vdash \Gamma \text{ ok } \quad \Gamma \vdash T}{\Gamma \vdash _ \in T \triangleright \bullet} \quad \frac{\Gamma, X \vdash p \in T \triangleright \Delta \quad X \notin \text{dom}(\Gamma)}{\Gamma \vdash \{X\} p \in \{X\} T \triangleright X, \Delta}$$

 A.7 Basic Process Formation $\Gamma \vdash_a P$

$$\frac{\Gamma \vdash a \in \mathbf{Agent}^s \quad \Gamma \vdash c \in \sim^w T \quad \Gamma \vdash v \in T}{\Gamma \vdash_a c!v}$$

$$\frac{\Gamma \vdash x \in \sim^r T \quad \Gamma \vdash p \in T \triangleright \Delta \quad \Gamma, \Delta \vdash_a P}{\Gamma \vdash_a x?p \rightarrow P \quad \text{and} \quad \Gamma \vdash_a *x?p \rightarrow P}$$

$$\frac{\Gamma \vdash a, b \in \mathbf{Agent}^s \quad \Gamma \vdash c \in \sim^w T \quad \Gamma \vdash v \in T}{\Gamma \vdash_a \langle b@? \rangle c!v}$$

$$\frac{\Gamma \vdash ev \in T \quad \Gamma \vdash p \in T \triangleright \Delta \quad \Gamma, \Delta \vdash_a P}{\Gamma \vdash_a \mathbf{let } p = ev \mathbf{ in } P}$$

$$\frac{\Gamma \vdash v \in \mathbf{Bool} \quad \Gamma \vdash_a P \quad \Gamma \vdash_a Q}{\Gamma \vdash_a \mathbf{if } v \mathbf{ then } P \mathbf{ else } Q}$$

$$\frac{a \neq b \quad \Gamma, b : \text{Agent}^Z \vdash_b P \quad \Gamma, b : \text{Agent}^Z \vdash_a Q}{\Gamma \vdash_a \mathbf{create}^Z b = P \mathbf{in} Q}$$

$$\frac{\Gamma \vdash a \in \text{Agent}^m \quad \Gamma \vdash s \in \text{Site} \quad \Gamma \vdash_a P}{\Gamma \vdash_a \mathbf{migrate to} s \rightarrow P}$$

$$\frac{\Gamma \vdash b \in \text{Agent}^s \quad \Gamma \vdash c \in \hat{\sim}^w T \quad \Gamma \vdash v \in T \quad \Gamma \vdash_a P \quad \Gamma \vdash_a Q}{\Gamma \vdash_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q}$$

$$\frac{\vdash \Gamma \text{ ok} \quad \Gamma \vdash a \in \text{Agent}^s}{\Gamma \vdash_a \mathbf{0}} \quad \frac{\Gamma \vdash_a P \quad \Gamma \vdash_a Q}{\Gamma \vdash_a P \mid Q} \quad \frac{\Gamma, x : \hat{\sim}^I T \vdash_a P}{\Gamma \vdash_a \mathbf{new} x : \hat{\sim}^I T \mathbf{in} P}$$

A.8 Located Process Formation $\Gamma \vdash LP$

$$\frac{\Gamma \vdash_a P \quad \Gamma \vdash a \in \text{Agent}^s \quad \Gamma, c : \hat{\sim}^I T \vdash LP}{\Gamma \vdash @_a P} \quad \frac{\Gamma, c : \hat{\sim}^I T \vdash LP}{\Gamma \vdash \mathbf{new} c : \hat{\sim}^I T \mathbf{in} LP}$$

$$\frac{\Gamma \vdash LP \quad \Gamma \vdash LQ}{\Gamma \vdash LP \mid LQ} \quad \frac{\Gamma, a : \text{Agent}^Z \vdash LP \quad \Gamma \vdash s \in \text{Site}}{\Gamma \vdash \mathbf{new} a : \text{Agent}^Z @s \mathbf{in} LP}$$

B. NOMADIC π CALCULUS REDUCTION AND LABELLED TRANSITION SEMANTICSB.1 Structural Congruence $P \equiv Q$ and $LP \equiv LQ$

Axioms

$$P \equiv P \mid \mathbf{0} \quad P \mid Q \equiv Q \mid P \quad P \mid (Q \mid R) \equiv (P \mid Q) \mid R$$

$$LP \mid LQ \equiv LQ \mid LP \quad LP \mid (LQ \mid LR) \equiv (LP \mid LQ) \mid LR$$

$$\frac{x \notin \text{fv}(P)}{P \mid \mathbf{new} x : T \mathbf{in} Q \equiv \mathbf{new} x : T \mathbf{in} P \mid Q}$$

$$\frac{x \notin \text{fv}(LP)}{LP \mid \mathbf{new} x : T @z \mathbf{in} LQ \equiv \mathbf{new} x : T @z \mathbf{in} LP \mid LQ}$$

$$@_a (P \mid Q) \equiv @_a P \mid @_a Q$$

$$\frac{x \neq a}{@_a \mathbf{new} x : \hat{\sim}^I T \mathbf{in} P \equiv \mathbf{new} x : \hat{\sim}^I T \mathbf{in} @_a P}$$

$$\frac{x, x' \text{ are all distinct}}{\mathbf{new} x : T \mathbf{in} \mathbf{new} x' : T' \mathbf{in} P \equiv \mathbf{new} x' : T' \mathbf{in} \mathbf{new} x : T \mathbf{in} P}$$

$$\frac{x, x' \text{ are all distinct}}{\mathbf{new} x : T @z \mathbf{in} \mathbf{new} x' : T' @z' \mathbf{in} LP \equiv \mathbf{new} x' : T' @z' \mathbf{in} \mathbf{new} x : T @z \mathbf{in} LP}$$

Congruence Rules

$$\begin{array}{c}
 P \equiv P \quad \frac{P \equiv Q \quad Q \equiv R}{P \equiv R} \quad \frac{P \equiv R}{P \mid Q \equiv R \mid Q} \\
 \\
 \frac{P \equiv Q}{\mathbf{new} \Delta \mathbf{in} P \equiv \mathbf{new} \Delta \mathbf{in} Q} \quad \frac{P \equiv Q}{c?p \rightarrow P \equiv c?p \rightarrow Q} \quad \frac{P \equiv Q}{*c?p \rightarrow P \equiv *c?p \rightarrow Q} \\
 \\
 \frac{P \equiv Q}{\mathbf{if} v \mathbf{then} P \mathbf{else} R \equiv \mathbf{if} v \mathbf{then} Q \mathbf{else} R} \quad \frac{P \equiv Q}{\mathbf{if} v \mathbf{then} R \mathbf{else} P \equiv \mathbf{if} v \mathbf{then} R \mathbf{else} Q} \\
 \\
 \frac{P \equiv Q}{\mathbf{iflocal} \langle a \rangle c!v \mathbf{then} P \mathbf{else} R \equiv \mathbf{iflocal} \langle a \rangle c!v \mathbf{then} Q \mathbf{else} R} \\
 \\
 \frac{P \equiv Q}{\mathbf{iflocal} \langle a \rangle c!v \mathbf{then} R \mathbf{else} P \equiv \mathbf{iflocal} \langle a \rangle c!v \mathbf{then} R \mathbf{else} Q} \\
 \\
 \frac{P \equiv Q}{\mathbf{let} p : T = ev \mathbf{in} P \equiv \mathbf{let} p : T = ev \mathbf{in} Q} \\
 \\
 \frac{P \equiv Q}{\mathbf{create}^m b = R \mathbf{in} P \equiv \mathbf{create}^m b = R \mathbf{in} Q} \quad \frac{P \equiv Q}{\mathbf{create}^m b = P \mathbf{in} R \equiv \mathbf{create}^m b = Q \mathbf{in} R} \\
 \\
 \frac{P \equiv Q}{\mathbf{migrate} \mathbf{to} s \rightarrow P \equiv \mathbf{migrate} \mathbf{to} s \rightarrow Q} \\
 \\
 \frac{P \equiv Q}{@_a P \equiv @_a Q} \\
 \\
 LP \equiv LP \quad \frac{LP \equiv LQ \quad LQ \equiv LR}{LP \equiv LR} \quad \frac{LP \equiv LR}{LP \mid LQ \equiv LR \mid LQ} \\
 \\
 \frac{LP \equiv LQ}{\mathbf{new} \Delta \mathbf{in} LP \equiv \mathbf{new} \Delta \mathbf{in} LQ}
 \end{array}$$

 B.2 Reduction $\Gamma \Vdash LP \rightarrow \Gamma' \Vdash LP'$

$$\begin{array}{l}
 \Gamma \Vdash @_a \mathbf{migrate} \mathbf{to} s \rightarrow P \rightarrow (\Gamma \oplus a \mapsto s) \Vdash @_a P \\
 \Gamma \Vdash @_a \mathbf{if} \mathbf{true} \mathbf{then} P \mathbf{else} Q \rightarrow \Gamma \Vdash @_a P \\
 \Gamma \Vdash @_a \mathbf{if} \mathbf{false} \mathbf{then} P \mathbf{else} Q \rightarrow \Gamma \Vdash @_a Q \\
 \Gamma \Vdash @_a \langle b \ @? \rangle c!v \rightarrow \Gamma \Vdash @_b c!v \\
 \Gamma \Vdash @_a (c!v \mid c?p \rightarrow P) \rightarrow \Gamma \Vdash @_a \mathbf{match}(p, v)P \\
 \Gamma \Vdash @_a (c!v \mid *c?p \rightarrow P) \rightarrow \Gamma \Vdash @_a ((\mathbf{match}(p, v)P) \mid *c?p \rightarrow P) \\
 \\
 \frac{\Gamma \vdash a @ s}{\Gamma \Vdash @_a \mathbf{create}^Z b = P \mathbf{in} Q \rightarrow \Gamma \Vdash \mathbf{new} b : \mathbf{Agent}^Z @s \mathbf{in} (@_b P \mid @_a Q)}
 \end{array}$$

$$\frac{\text{eval}(ev) \text{ defined}}{\Gamma \Vdash @_a \mathbf{let} p = ev \mathbf{in} P \rightarrow \Gamma \Vdash @_a \text{match}(p, \text{eval}(ev)) P}$$

$$\frac{\Gamma \vdash a@s \quad \Gamma \vdash b@s}{\Gamma \Vdash @_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q \rightarrow \Gamma \Vdash @_a P \mid @_b c!v}$$

$$\frac{\Gamma \vdash a@s \quad \Gamma \vdash b@s' \quad s \neq s'}{\Gamma \Vdash @_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q \rightarrow \Gamma \Vdash @_a Q}$$

$$\frac{LP \equiv LP' \quad \Gamma \Vdash LP \rightarrow \Gamma' \Vdash LQ \quad LQ \equiv LQ'}{\Gamma \Vdash LP' \rightarrow \Gamma' \Vdash LQ'}$$

$$\frac{\Gamma \Vdash LP \rightarrow \Gamma' \Vdash LR}{\Gamma \Vdash LP \mid LQ \rightarrow \Gamma' \Vdash LR \mid LQ}$$

$$\frac{\Gamma, x : \mathbf{Agent}^Z@s \Vdash LP \rightarrow \Gamma', x : \mathbf{Agent}^Z@s' \Vdash LQ}{\Gamma \Vdash \mathbf{new} x : \mathbf{Agent}^Z@s \mathbf{in} LP \rightarrow \Gamma' \Vdash \mathbf{new} x : \mathbf{Agent}^Z@s' \mathbf{in} LQ}$$

$$\frac{\Gamma, x : \sim^I T \Vdash LP \rightarrow \Gamma', x : \sim^I T \Vdash LQ}{\Gamma \Vdash \mathbf{new} x : \sim^I T \mathbf{in} LP \rightarrow \Gamma' \Vdash \mathbf{new} x : \sim^I T \mathbf{in} LQ}$$

B.3 Labelled Transitions $\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LQ$ and $\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LQ$

$$\frac{\text{eval}(ev) \text{ defined}}{\Gamma \Vdash_a \mathbf{let} p = ev \mathbf{in} P \xrightarrow{\tau} @_a \text{match}(p, \text{eval}(ev)) P}$$

$$\frac{\Gamma \vdash a@s}{\Gamma \Vdash_a \mathbf{create}^Z b = P \mathbf{in} Q \xrightarrow{\tau} \mathbf{new} b : \mathbf{Agent}^Z@s \mathbf{in} (@_b P \mid @_a Q)}$$

$$\Gamma \Vdash_a \mathbf{migrate to} s \rightarrow P \xrightarrow{\text{migrate to } s} @_a P$$

$$\frac{\Gamma \vdash a@s \quad \Gamma \vdash b@s}{\Gamma \Vdash_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q \xrightarrow{\tau} @_a P \mid @_b c!v}$$

$$\frac{\Gamma \vdash a@s \quad \Gamma \vdash b@s' \quad s \neq s'}{\Gamma \Vdash_a \mathbf{iflocal} \langle b \rangle c!v \mathbf{then} P \mathbf{else} Q \xrightarrow{\tau} @_a Q}$$

$$\Gamma \Vdash_a c!v \xrightarrow{c!v} @_a \mathbf{0}$$

$$\frac{\Gamma \vdash c \in \sim^r T \quad \Gamma, \Delta \vdash v \in T \quad \text{dom}(\Delta) \subseteq \text{fv}(v) \quad \Delta \text{ extensible.}}{\Gamma \Vdash_a c?p \rightarrow P \xrightarrow[\Delta]{c?v} @_a \text{match}(p, v) P}$$

$$\text{and } \Gamma \Vdash_a *c?p \rightarrow P \xrightarrow[\Delta]{c?v} @_a (\text{match}(p, v) P \mid *c?p \rightarrow P)$$

$$\begin{array}{c}
 \frac{\Gamma \Vdash_a P \xrightarrow[\Delta]{c!v} LP \quad \Gamma \Vdash_a Q \xrightarrow[\Delta]{c?v} LQ}{\Gamma \Vdash_a P \mid Q \xrightarrow{\tau} \mathbf{new} \Delta \mathbf{in} (LP \mid LQ)} \\
 \\
 \Gamma \Vdash_a \mathbf{if true then} P \mathbf{else} Q \xrightarrow{\tau} @_a P \\
 \\
 \Gamma \Vdash_a \mathbf{if false then} P \mathbf{else} Q \xrightarrow{\tau} @_a Q \\
 \\
 \Gamma \Vdash_a \langle b@? \rangle c!v \xrightarrow{\tau} @_b c!v \\
 \\
 \frac{\Gamma, x:T \Vdash_a P \xrightarrow[\Delta]{c!v} LP \quad x \in \text{fv}(v) \quad x \neq c}{\Gamma \Vdash_a \mathbf{new} x:T \mathbf{in} P \xrightarrow[\Delta, x:T]{c!v} LP} \\
 \\
 \frac{\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LP \quad LP \equiv LQ}{\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LQ} \\
 \\
 \frac{\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LP}{\Gamma \Vdash_a P \mid Q \xrightarrow[\Delta]{\alpha} LP \mid @_a Q} \\
 \\
 \frac{\Gamma, x:T \Vdash_a P \xrightarrow[\Delta]{\alpha} LP \quad x \notin \text{fv}(\alpha)}{\Gamma \Vdash_a \mathbf{new} x:T \mathbf{in} P \xrightarrow[\Delta]{\alpha} \mathbf{new} x:T \mathbf{in} LP} \\
 \\
 \frac{\Gamma, x:T@z \Vdash LP \xrightarrow[\Delta]{@_a c!v} LQ \quad x \in \text{fv}(v) \quad x \neq c \quad x \neq a}{\Gamma \Vdash \mathbf{new} x:T@z \mathbf{in} LP \xrightarrow[\Delta, x:T@z]{@_a c!v} LQ} \\
 \\
 \frac{\Gamma \Vdash_a P \xrightarrow[\Delta]{\alpha} LP \quad \beta = \begin{cases} \tau & \text{if } \alpha = \tau \\ @_a \alpha & \text{otherwise} \end{cases}}{\Gamma \Vdash @_a P \xrightarrow[\Delta]{\beta} LP} \\
 \\
 \frac{\Gamma \Vdash LP \xrightarrow[\Delta]{@_a c!v} LP' \quad \Gamma \Vdash LQ \xrightarrow[\Delta]{@_a c?v} LQ'}{\Gamma \Vdash LP \mid LQ \xrightarrow{\tau} \mathbf{new} \Delta \mathbf{in} (LP' \mid LQ')} \\
 \\
 \frac{\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LQ \quad LQ \equiv LR}{\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LR}
 \end{array}$$

$$\frac{\Gamma \Vdash LP \xrightarrow[\Delta]{\beta} LR}{\Gamma \Vdash LP \mid LQ \xrightarrow[\Delta]{\beta} LR \mid LQ}$$

$$\frac{\Gamma, x : T@z \Vdash LP \xrightarrow[\Delta]{\beta} LQ \quad x \notin \text{fv}(\beta)}{\Gamma \Vdash \mathbf{new} \ x : T@z \ \mathbf{in} \ LP \xrightarrow[\Delta]{\beta} \mathbf{new} \ x : T@z \ \mathbf{in} \ LQ}$$

$$\frac{\Gamma, a : \mathbf{Agent}^m@s \Vdash LP \xrightarrow{\text{@}_a \text{migrate to } s'} LQ}{\Gamma \Vdash \mathbf{new} \ a : \mathbf{Agent}^m@s \ \mathbf{in} \ LP \xrightarrow{\tau} \mathbf{new} \ a : \mathbf{Agent}^m@s' \ \mathbf{in} \ LQ}$$

C. EXAMPLE INFRASTRUCTURE: QUERY SERVER WITH CACHING ALGORITHM

In this appendix we give the full definition of the *Query Server with Caching* (QSC) algorithm introduced in Section 6.

The QSC encoding consists of three parts, a top-level translation (applied to whole programs), an auxiliary compositional translation $\llbracket P \rrbracket$ of subprograms P , defined phrase-by-phrase, and an encoding of types. The encoding involves three main classes of agents: the query server Q itself (on a single site), the daemons (one on each site), and the translations of high-level application agents (which may migrate). The top-level translation of a program P launches the query server and all the daemons before executing $\llbracket P \rrbracket$. The query server and the code for a single daemon are given in Figure 9; the interesting clauses of the compositional translation are in the text below. The compositional translation is parametric on a triple $[\mathbf{Agent}^Z \ \mathbf{Agent}^s \ \text{Site}]$.

Interactions between the query server, the per-site daemons, and the encodings of high-level agents, are via the channels below.

```

register   :  $\hat{\text{rw}}[\mathbf{Agent}^s \ [\text{Site} \ \mathbf{Agent}^Z]]$ 
migrating :  $\hat{\text{rw}}\mathbf{Agent}^s$ 
migrated  :  $\hat{\text{rw}}[\text{Site} \ \mathbf{Agent}^Z]$ 
message   :  $\hat{\text{rw}}\{X\} [\mathbf{Agent}^s \ \text{Site} \ \mathbf{Agent}^Z \ \hat{\text{w}}X \ X]$ 
trymessage :  $\hat{\text{rw}}\{X\} [\mathbf{Agent}^s \ \hat{\text{w}}X \ X]$ 
trydeliver :  $\hat{\text{rw}}\{X\} [\mathbf{Agent}^s \ \text{Site} \ \mathbf{Agent}^Z \ \hat{\text{w}}X \ X \ \text{Bool}]$ 
update    :  $\hat{\text{rw}}[\mathbf{Agent}^s \ [\text{Site} \ \mathbf{Agent}^Z]]$ 
dack      :  $\hat{\text{w}}[]$ 
ack       :  $\hat{\text{w}}[]$ 

```

In addition, the translations of high-level agents use a channel name `currentloc` internally.

The messages sent between agents fall into three groups, implementing the high-level agent creation, agent migration, and location-independent messages. Typical executions are illustrated in Figure 10 and below. Correspondingly, only these cases of the compositional translation are non-trivial.

Each class of agents maintains some explicit state as an output on a lock channel. The query server maintains a map from each agent name to the site (and daemon) where the agent is currently located. This is kept accurate when agents are created

or migrate. Each daemon maintains a map from some agent names to the site (and daemon) that they guess the agent is located at. This is updated only when a message delivery fails. The encoding of each high-level agent records its current site (and daemon).

To send a location-independent message the translation of a high-level agent simply asks the local daemon to send it. The compositional translation of $\langle b@? \rangle c!v$, ‘send v to channel c in agent b ’, is below.

$$\begin{aligned} \llbracket \langle b@? \rangle c!v \rrbracket_{[a \ Q \ SQ]} &\stackrel{\text{def}}{=} \\ \text{currentloc?}[S \ DS] \rightarrow & \\ \mathbf{iflocal} \langle DS \rangle \text{trymessage!} \{T\} [b \ c \ v] \mathbf{then} & \\ \text{currentloc!}[S \ DS] & \\ \mathbf{else} \ \mathbf{0} & \end{aligned}$$

This first reads the name S of the current site and the name DS of the local daemon from the agent’s lock channel `currentloc`, then sends $[b \ c \ v]$ on the channel `trymessage` to DS , replacing the lock after the message is sent. The translation is parametric on the triple $[a \ Q \ SQ]$ of the name a of the agent containing this phrase, the name Q of the query server, and the site SQ of the query server — for this phrase, none are used. We return later to the process of delivery of the message.

To migrate while keeping the query server’s map accurate, the translation of a **migrate to** in a high-level agent synchronises with the query server before and after actually migrating, with `migrating`, `migrated`, and `ack` messages.

$$\begin{aligned} \llbracket \mathbf{migrate\ to} \ u \rightarrow P \rrbracket_{[a \ Q \ SQ]} &\stackrel{\text{def}}{=} \text{currentloc?}[S \ DS] \rightarrow \\ \mathbf{let} \ [U \ DU] = u \ \mathbf{in} & \\ \langle Q@SQ \rangle \text{migrating!} a & \\ | \ \text{ack?}_- \rightarrow \mathbf{migrate\ to} \ U \rightarrow & \\ \langle Q@SQ \rangle \text{migrated!}[U \ DU] & \\ | \ \text{ack?}_- \rightarrow & \\ \text{currentloc!}[U \ DU] & \\ | \ \llbracket P \rrbracket_{[a \ Q \ SQ]} & \end{aligned}$$

```


QueryServersq  $\stackrel{\text{def}}{=}$ 
let [SQ Q] = sq in
new lock :  $\sim^{\text{rw}}$ Map[Agents [Site Agents]]
lock!(map.make ==)
| *register?[a [S DS]] → lock?m →
  let[Agents [Site Agents]] m' = (m with a ↦ [S DS]) in
  lock!m'
  | <a@S>ack![]
| *migrating?a → lock?m →
  lookup[Agents [Site Agents]] a in m with
  found([S DS]) →
  <a@S>ack![]
  | migrated?[S' DS'] →
    let[Agents [Site Agents]] m' = (m with a ↦ [S' DS']) in
    lock!m'
    | <a@S'>ack![]
  notfound → 0
| *message? {X} [DU U a c v] → lock?m →
  lookup[Agents [Site Agents]] a in m with
  found([R DR]) →
  <DU@U>update![a [R DR]]
  | <DR@R>trydeliver! {X} [Q SQ a c v true]
  | dack?_ → lock!m
  notfound → 0

Daemons  $\stackrel{\text{def}}{=}$ 
let [S D] = s in
new lock :  $\sim^{\text{rw}}$ Map[Agents [Site Agents]]
<toplevel@firstSite>ndack! [S D]
| lock!(map.make ==)
| *trymessage? {X} [a c v] → lock?m →
  lookup[Agents [Site Agents]] a in m with
  found([R DR]) →
  <DR@R>trydeliver! [D S a c v false]
  | lock!m
  notfound →
  <Q@SQ>message! [D S a c v]
  | lock!m
| *trydeliver? {X} [DU U a c v ackme] →
  iflocal <a>c!v then
  if ackme then <DU@U>dack![] else 0
  else <Q@SQ>message! {X} [DU U a c v]
| *update?[a s] → lock?m →
  let[Agents [Site Agents]] m' = (m with a ↦ s) in
  lock!m'


```

Fig. 9. Parts of the Top Level — the Query Server and a site Daemon


The best scenario: good guess in the D cache. This should be the common case.



No guess in the D cache.




The worst scenario: wrong guess in the D cache.



Horizontal arrows are synchronised communications within a single machine (using `iflocal`); slanted arrows are asynchronous messages.

Fig. 10. The Delivery of Location-Independent Message $\langle b@? \rangle c!v$ from a to b .

A sample execution is below.



The query server's lock is kept during the migration. The agent's own record of its current site and daemon must also be updated with the new data $[U \ DU]$ when the agent's lock is released. Note that in the body of the encoding the name DU of the daemon on the target site must be available. This is achieved by encoding site names in the high-level program by pairs of a site name and the associated daemon name; there is a translation of types


$$\begin{aligned} \llbracket \text{Agent}^Z \rrbracket &\stackrel{\text{def}}{=} \text{Agent}^Z \\ \llbracket \text{Site} \rrbracket &\stackrel{\text{def}}{=} [\text{Site Agent}^Z] \end{aligned}$$

Similarly, a high-level agent a must synchronise with the query server while creating a new agent b , with messages on **register** and **ack**.

$$\begin{aligned} &\llbracket \text{create}^Z b = P \text{ in } P' \rrbracket_{[a \ Q \ SQ]} \stackrel{\text{def}}{=} \\ &\quad \text{currentloc?}[S \ DS] \rightarrow \\ &\quad \quad \text{create}^Z b = \\ &\quad \quad \langle Q@SQ \rangle \text{register!}[b \ [S \ DS]] \\ &\quad \quad | \text{ack?}_- \rightarrow \text{iflocal} \langle a \rangle \text{ack!}[] \text{ then} \\ &\quad \quad \quad \text{currentloc!}[S \ DS] \\ &\quad \quad \quad | \llbracket P \rrbracket_{[b \ Q \ SQ]} \\ &\quad \quad \text{else } 0 \\ &\quad \text{in} \\ &\quad \quad \text{ack?}_- \rightarrow \\ &\quad \quad \quad \text{currentloc!}[S \ DS] \\ &\quad \quad \quad | \llbracket P' \rrbracket_{[a \ Q \ SQ]} \end{aligned}$$

The current site/daemon data for the new agent must be initialised to $[S \ DS]$; the creating agent is prevented from migrating away until the registration has taken place by keeping its **currentloc** lock until an **ack** is received from b . A sample

execution is below.



Returning to the process of message delivery, there are three cases (see Figure 10). Consider the implementation of $\langle b@? \rangle c!v$ in agent a on site S , where the daemon is D . Suppose b is on site R , where the daemon is DR . Either D has the correct site/daemon of b cached, or D has no cache data for b , or it has incorrect cache data. In the first case D sends a **trydeliver** message to DR which delivers the message to b using **iflocal**. For the PA application this should be the common case; it requires only one network message.

In the cache-miss case D sends a **message** message to the query server, which both sends a **trydeliver** message to DR (which then delivers successfully) and an **update** message back to D (which updates its cache). The query server's lock is kept until the message is delivered, thus preventing b from migrating until then.

Finally, the incorrect-cache-hit case. Suppose D has a mistaken pointer to $DU@U$. It will send a **trydeliver** message to DU which will be unable to deliver the message. DU will then send a **message** to the query server, much as before (except that the cache update message still goes to D , not to DU).

D. THE NOMADIC PICT PROGRAMMING LANGUAGE

This appendix gives a more detailed overview of the Nomadic Pict prototype distributed programming language, as outlined in Section 7, beginning with the low-level language. The language implementation is described in the following Appendix E.

D.1 Primitives

We will introduce the low-level primitives in groups. They fall into two main syntactic categories of *processes* and *declarations*. A program is simply a series of declarations, which may contain processes. The other main syntactic categories are abstractions, patterns, values, paths, types, and constants.

Declarations Declarations D include definitions of types, channels, process abstrac-

tions, agents, and also a migration primitive.

type T = T'	type abbreviation
new c:T	new channel name creation
def f[...] = P and ... and f' [...] = P'	process abstraction
agent a = P and ... and a' = P'	agent creation
migrate to s P	agent migration

The declaration **type** T = T' introduces a new name T for complex type T'. Execution of **new** c:T creates a new unique channel name c for carrying values of type T. The execution of the construct **agent** a=P spawns a new agent on the current site, with body P. A group of agent definitions introduced by **agent** and separated by **and** can be mutually recursive, i.e. each of the bodies P can refer to any of the defined agent names. Agents can migrate to named sites — the execution of **migrate to** s P as part of an agent results in the whole agent migrating to site s. The **def** declarations are used to program in the functional style (examples will be given in Section D.6).

Processes Processes P, Q, ... form a separate syntactic category.

(P Q)	parallel composition
(D P)	local declaration
()	null process

The body of an agent may consist of many process terms in parallel, i.e. essentially of many lightweight threads. They will interact only by message passing. We can write a composition of more than two processes as (P1 | ... | Pn). Large programs often contain processes with long sequences of declarations like (**new** x1:T1 ... (**new** x2:T2 P)). To avoid many nested parentheses this can be written simply as (**new** x1:T1 ... **new** x2:T2 P). In sequences of declarations, it is often convenient to start some process running in parallel with the evaluation of the remainder of the declarations. The Pict declaration keyword **run** can be used for this, e.g. a program

```
(new x:T
  run print!"Hello"
  new y:T
  P)
```

will be transformed into (**new** x:T (print!"Hello" | (**new** y:T P))). The next process forms are the π calculus-style interaction primitives of Pict.

c!v	output v on channel c in the current agent
c?p = P	input from channel c
c?*p = P	replicated input from channel c
if v then P else Q	conditional

An output c!v (of value v on channel c) and an input c?p=P in the same agent may match, resulting in P with the appropriate parts of the value v bound to the formal parameters in the pattern p. The communication is asynchronous, i.e. the output is never blocked. The implementation uses local environments to store bindings of parameters to actual values. A replicated input c?*p=P behaves similarly except

that it persists after the matching, and so may receive another value. In both $c?p=P$ and $c?*p=P$ the names in p are binding in P . Finally, the low-level language includes primitives for interaction between agents.

```

iflocal <a>c!v then P else Q  test-and-send to agent a on this site
<a>c!v                               send to agent a on this site
<a@s>c!v                             send to agent a on site s
wait c?p=P timeout t -> Q  input with timeout

```

The semantics of the first three constructs has been described in Section 2. For implementing infrastructures that are robust under some level of failure, or support disconnected operation, some timed primitive is required. The low-level language includes a single timed input as above, with timeout value n . If a message on channel c is received within t seconds then P will be started as in a normal input, otherwise Q will be. The timing is approximate, as the runtime system may introduce some delays.

We also include a construct for explicit agent-level garbage collection.

```

terminate                               terminate execution of the current agent

```

The execution of `terminate` terminates the current agent and removes its closure from the heap, releasing memory occupied by the agent. Any I/O operations (e.g. input from a keyboard) will be abruptly killed.

The high-level language is obtained by extending the low-level language with a single location-independent communication primitive.

```

c@a!v                               LI output to channel c at agent a

```

The intended semantics of an output $c@a!v$ is that its execution will reliably deliver the message $c!v$ to agent a , irrespective of the current site of a and of any migrations. The low-level communication primitives are also available, for interacting with application agents whose locations are predictable. The actual semantics of $c@a!v$ depends on the compile-time encoding (or translation) of this primitive into the low-level language, from language libraries.

D.2 Names and Scope Extrusion

Names play a key rôle in the Nomadic Pict language. New names of agents and channels can be created dynamically. These names are *pure*, in the sense of Needham [1989]; no information about their creation is visible within the language (in our current implementation they do contain site IDs, but could equally well be implemented by choosing large random numbers). Site names contain an IP address and TCP port number of the runtime system which they represent.

Channel, agent, and site names are first-class values and they can be freely sent to processes which are located at other agents. As in the π calculus, names can be *scope-extruded* — here channel and agent names can be sent outside the agent in which they were created. For example, if the body of agent a is

```

agent b =
(
  new d : T
  iflocal <a>c!d then () else ()

```

```

)
in
  c?x=x! []

```

then channel name `d` is created in agent `b`. After the output message `c!d` has been sent from `b` to `a` (**iflocal**) and has interacted with the input `c?x=x! []` there will be an output `d! []` in agent `a`.

D.3 Types

All bound variables are explicitly typed. In practice, however, many of these type annotations can be inferred by the compiler. Therefore we do not include them in the syntax above. Types are required in definitions, e.g. execution of `new c:~T P` creates a new unique channel name for carrying values of type `T`. The language inherits a rich type system from Pict, including simple record types, higher-order polymorphism, simple recursive types and subtyping. It has a partial type inference algorithm. Below we summarise the key types, referring the reader to Pierce and Turner [1997] for details.

Base Types The base types include `String` of strings, `Char` of characters, `Int` of integers, and `Bool` of Booleans. They are two predefined Boolean constants `false` and `true` of type `Bool`. Nomadic Pict adds new base types `Site` and `Agent` of site and agent names.

Channel Types and Subtyping Pict's four channel types are as follows: `~T` is the type of input/output channels carrying values of type `T`, `!T` is the type of output channels accepting `T`, `?T` is the type of input channels yielding `T`, and `/T` is the type of *responsive* output channels carrying `T` in process abstractions and functions. The first three correspond to channels with capabilities `rw`, `w`, and `r` of Section 2. The type `~T` is a subtype of both `!T` and `?T`. That is, a channel that can be used for both input and output may be used in a context where just one capability is needed. The type `/T` is a subtype of `!T` and it was introduced to define process abstractions and channels carrying results in a functional style (see examples in Section D.6). Type `/T` guarantees that there is exactly one (local) receiver. We define a type `Sig` as `/[]` to mark responsive channels which are used to signal completion rather than for exchanging data.

Records, Polymorphic and Recursive types We can use tuples `[T1...Tn]` of types `T1...Tn` and existential polymorphic types such as `[#X T1...Tn]` in which the type variable `X` may occur in the field types `T1...Tn`. We can add labels to tuples obtaining records such as `[label1=T1...labeln=Tn]`. Recursive types are constructed as (**rec** `X=T`), in which the type variable `X` occurs in type `T`.

Variant and Dynamic Types In Nomadic Pict we have added a *variant* type `[label1>T1...labeln>Tn]` and a type `Dyn` of *dynamic* values. The variant type `[label1>T1...labeln>Tn]` has values `[label>v:T]`. The dynamic type is useful for implementing traders, i.e. maps from string keywords (or textual descriptions) to dynamic values. Dynamic values are implemented as pairs (v, T) of a value and its type.

Defining Types and Type Operators We can use a declaration keyword **type** to

define new types and type operators, e.g. `type (Double X) = [X X]` denotes a new type operator `Double` which can be used as in `new c:^(Double Int)`. In Nomadic Pict programs, we often use a type operator `Map` from the libraries, taking two types and giving the type of maps, or lookup tables, from one to the other (we have already used maps in our translations).

D.4 Values and Patterns

Values Channels allow the communication of first-order values, consisting of channel, agent, and site names, constants, integers, strings, characters, tuples `[v1 . . . vn]` of the n values `v1 . . . vn`, packages of existential types `[#T v1 . . . vn]`, elements of variant types `[label>v]`, and dynamic values. A dynamic value can be constructed using a constructor `dynamic`, as in `(dynamic v)`. Values are deconstructed by pattern matching on input or, in the case of variants and dynamic values, using syntactic sugar `switch` and `typecase`, which we define in Section D.6.

Patterns `p` are of the same shapes as values (but names cannot be repeated), with the addition of a wildcard pattern `_`.

D.5 Expressing Encodings

A Nomadic Pict program is organised as a file containing a sequence of declarations, preceded by a number of `import` clauses:

```
import "name" {- Imports and any global declarations -}
...
program arg : T =
(
  {- A user-defined program in the high-level language -}
)
```

After imports, we can have any global declarations, such as constants and global functions. Then, in the body of `program`, we include the actual program, which can be expressed using the high-level language constructs. The program accepts an argument `arg` of type `T`, which is usually a list of sites in the system.

The (optional) compositional translation of the high-level location-independent language into the low-level language, is structured as follows:

```
{- Any global declarations of the compositional translation -}

{toplevel P arg} T' =
(
  {- A top-level definition -}
)
{- A translation of types -}
{- A compositional translation of primitives -}
```

Firstly, we declare any global constants, functions, and channel names used by the compositional translation. Then, we use `toplevel` to declare a top-level program (in the low-level language), which declares actions that are executed before the `program` is executed, e.g. spawning daemons and servers on remote sites. The names `P` and `arg`, denote correspondingly the main program declared with `program`,

and its argument (see the **program** declaration above). The type T' is the type of the translation parameter (e.g. **Agent** in our example translation in Figure 3).

Encodings of high-level types and language primitives can be expressed using a rudimentary module language, allowing the translation of each interesting phrase (all those involving agents or communication) to be specified and type checked; the translation of a whole program (including the translation of types) can be expressed using this compositional translation. If the definition of some high-level language construct is missing, the compiler will use the low-level language construct. A concrete syntax of the language is in Wojciechowski [2000a]; the example infrastructures in the previous sections should give the idea.

This special-purpose scheme for expressing encodings is sufficient for the purpose, but can be replaced in a general-purpose language by use of a module system, as the ML-style module system of Acute [Sewell et al. 2007] was used to express Nomadic Pict-style encodings.

D.6 Syntactic Sugar

The core language described in Section D.1 lacks some constructs which are useful in programming. In order to avoid complicating the semantics of the core language, additional programming features are provided as *syntactic sugar*, i.e. there is an unambiguous translation of the code with the additions into code without them. Below we describe some syntactic sugar. Most are standard Pict forms; some are new. Interested readers are directed to a formal description of the source-to-source translations in Pict in Pierce and Turner [1997], where all Pict forms are described in detail.

Process Abstractions and Functions In Pict, we can define *process abstractions*, i.e. process expressions prefixed by patterns, via the declaration keyword **def**, as in

```
def f [x:T1 y:T2] = (x!y | x!y)
```

and instances are created using the same syntax as output expressions, as in $f![a\ b]$. The name f has type $/[T1\ T2]$. Recursive and mutually recursive definitions

```
def f [...] = ... g![...] ...
and g [...] = ... f![...] ...
```

are also allowed.

A *functional style* is supported by a small extension to the syntactic class of abstractions. For example, we can replace a process abstraction $\mathbf{def}\ f\ [a1:T1\ a2:T2\ r:/T] = r!v$, where v is some complex value, by a ‘function definition’

```
def f (a1:T1 a2:T2) : T = v
```

and avoid explicitly giving a name to the result channel r . For simplicity, we often confuse process abstractions as above and process abstractions which do not return any values, using a single term “functions”.

We can define anonymous abstractions as in Pict

```
\[...] = ...
```

For example, below is a function f which accepts process abstractions of type $\mathbf{String} \rightarrow \mathbf{Sig}$

```
def f g:[String Sig] = ((g "foo"); ())
```

We can create an instance of `f` passing an anonymous function which prints an argument `s` and sends an acknowledgement signal on channel `r` as follows

```
f!\[s:String r:Sig] = ((pr s); r![])
```

Functions can be effectively used in Nomadic Pict by *all* agents which have been defined in the lexical scope of function definitions. So formally it looks as though each agent has a private copy of each function it might ever use. Similarly, any public library function imported by a program can be invoked in all agents defined by the program. All functions defined inside an agent are private to this agent.

Declaration Values and Applications The syntactic category of values is extended with *declaration values* of the form $(D \ v)$, as in

```
c!(new d:T d)
```

The complex value is *always* evaluated to yield a simple value, which is substituted for the complex expression; the process above creates a fresh channel `d` and sends it off along `c`, as in `(new d:T c!d)`.

In value expressions, we allow the *application* syntax $(v \ v1 \ \dots \ v2)$. For example, we can define a process abstraction

```
def double [i:Int r:/Int] = +![i i r]
```

and then, in the scope of the declaration, write `(double i)` as a value, dropping the explicit result channel `r`, e.g. `printi!(double 2)` would compute 4 and print it out on the console, using the library channel `printi`.

Value Declarations A declaration

```
val p=v
```

evaluates a complex value `v` and names its result. Formally, a **val** declaration `(val p=v e)` is translated using the continuation-passing translation, so that the body `e` appears inside an input prefix on the continuation channel which is used to communicate a simple value evaluated from the complex value `v`. This means that **val** declarations are *blocking*: the body `e` cannot proceed until the bindings introduced by the **val** have actually been established.

Other Syntactic Sugar The idiom “invoke an operation, wait for a signal (i.e. a null value `[]`) as a result, and continue” appears frequently, so it is convenient to introduce `;` (semi-colon), as in

```
(v1 ...); (v2 ...)
```

for the sequence of operations `v1` and `v2`.

Matching Variants and Dynamic Values In Nomadic Pict programs we use a variant type `[label1> T1 ... labeln> Tn]` so often (e.g. in our infrastructure translations), that it is convenient to introduce a new construct **switch**, as in

```
c?v= switch v of
  (
    label1> p1 -> P1
```

```

...
  labeln> pn -> Pn
)

```

that matches a value v of variant type with all the variants, chooses the one which has the same label as v , and proceeds with a process P of the matched variant.

We can compare dynamic values at runtime using the construct **typecase**, as in

```

c?v= typecase v of
  s:String -> print!s
  [s:String d:~String] -> d!s
  else print!"Type not recognised."

```

where c has type $\sim\text{Dyn}$. Instances of dynamic values are created using (**dynamic** v). For example, $c!(\text{dynamic } ["\text{ala"} \ x])$ in parallel with the process term above may synchronise, resulting in "ala" being sent along the channel x , $c!(\text{dynamic } "\text{ala}")$ would print "ala", but any other (dynamic) value sent on c would print an error message "Type not recognised". The constructs **switch** and **typecase** are desugared using the value equality testing primitive. In the examples above, **switch** and **typecase** are process terms but we can also use these constructs in expressions yielding a value.

D.7 Procedures

Within a single agent one can express 'procedures' in Nomadic Pict as simple replicated inputs. Replicated inputs are useful to express server agents. Below is a first attempt at a pair-server, that receives values x on channel **pair** and returns two copies of x on channel **result**, together with a single invocation of the server.

```

new pair   : ~T
new result : ~[T T]
( pair?*x = result![x x]
| pair!v
| result?z = ... z ... )

```

This pair-server can only be invoked sequentially—there is no association between multiple requests and their corresponding results. A better idiom is below, in which new result channels are used for each invocation. The pair-server has a polymorphic type (X is a type variable), instantiated to Int by a client process.

```

type (Res X) = ~[X X]
new pair : ~[#X X (Res X)]
( pair?*[#X x r] = r![x x]
| (new result:(Res Int) (pair![1 result] | result?z =... z ...))
| (new result:(Res Int) (pair![2 result] | result?z =... z ...)))

```

The example can easily be lifted to remote procedure calls between agents. We show two versions, firstly for location-dependent RPC between static agents and secondly for location-independent RPC between agents that may be migrating. In the first case, the server becomes

```

new pair : ~[#X X (Res X) Agent Site]
pair?*[#X x r b s] = <b @ s>r![x x]

```

which returns the result using location-dependent communication to the agent **b** on site **s** received in the request. If the server is part of agent **a1** on site **s1** it would be invoked from agent **a2** on site **s2** by

```
new result : (Res Int)
  ( <a1 @ s1>pair![7 result a2 s2]
    | result?z = ...z... )
```

If agents **a1** or **a2** can migrate this can fail. A more robust idiom is easily expressible in the high-level language—the server becomes

```
new pair : ^[#X X (Res X) Agent]
  pair?*[#X x r b] = r@b![x x]
```

which returns the result using location-independent communication to the agent **b**. If the server is part of agent **a1** it would be invoked from agent **a2** by

```
new result : (Res Int)
  ( pair@a1![3 result a2]
    | result?z= ...z... )
```

D.8 Mobile Agents

Nomadic Pict agents are located at sites and they can freely migrate to other named sites. Agents carry their computation state with themselves and their execution is resumed on a new site from the point where they stopped on previous site. Mobile agents can exchange messages on channels. A channel name can be created dynamically and sent to other agents which can use it for communication.

Below is a program in the high-level language showing how a mobile agent can be expressed.

```
new answer : ^String
def spawn [s:Site] =
  (agent b =
    (migrate to s
     answer@a!(sys.read "How are you?"))
  in
    ())
  ( spawn![s1]
    | spawn![s2]
    | answer ?* s = print!s)
```

In the main part of the program, assumed to be part of some agent **a**, a function **spawn** is invoked twice. The function spawns a new agent **b**, which migrates to a site passed as the function argument. After migration, the agent outputs a location-independent message to agent **a**, on channel **answer**, containing a string read from a standard input on the target site; the answer is printed out on the current site of **a**.

D.9 Locks, Methods and Distributed Objects

The language inherits a common idiom for expressing concurrent objects from Pict [Pierce and Turner 1995]. The process

```

new lock: ^StateType
  ( lock!initialState
    | method1?*arg = (lock?state = ... lock!state' ...)
    ...
    | methodn?*arg = (lock?state = ... lock!state'' ...) )

```

is analogous to an object with methods `method1..methodn` and a state of type `StateType`. Mutual exclusion between the bodies of the methods is enforced by keeping the state as an output on a lock channel; the lock is free if there is an output and taken otherwise. For a more detailed discussion of object representations in process calculi, the reader is referred to Pierce and Turner [1995]. It contains an example program illustrating how a simple *reference cell* abstraction can be defined in Pict.

Below we rewrite the example of a reference cell abstraction, showing how *distributed objects* can be expressed in Nomadic Pict. The program uses mobile agents and many of the derived forms described in previous sections.

A reference cell can be modeled by an agent with two procedures connecting it to the outside world — one for receiving `set` requests and one for receiving `get` requests. Below is a cell, which holds an integer value (in channel `contents`) that initially contains 0.

```

type RefInt =
[
  set=/[Agent Int Sig]
  get=/[Agent /Int]
]

def refInt [s:Site r:/RefInt] =
(
  new set:^[Agent Int Sig]
  new get:^[Agent !Int]

  agent refIntAg =
  (
    new contents: ^Int
    run contents!0
    migrate to s
    ( set?*[a:Agent v:Int c:Sig]= contents?_ = (contents!v | c![])
      | get?*[a:Agent res:!Int]= contents?v = (contents!v | res@a!v)
    )
  )
  r![
    set = \[a:Agent v:Int c:Sig] = set@refIntAg![a v c]
    get = \[a:Agent res:!Int] = get@refIntAg![a res]
  ]
)

```

A function `refInt` defines two method channels `set` and `get` and creates a cell agent `refIntAg` which immediately migrates to site `s`. The cell agent maintains the invariant that, at any given moment, there is at most one process ready to send on `contents` and when methods `set` and `get` are not active, there is exactly one value in `contents`. The function `refInt` returns a record which defines an interface

to procedures of the cell agent. The record contains two labelled fields, `set` and `get`, with anonymous functions implementing the location-independent access to the corresponding procedures.

Now, we can create two instances (objects) `cell1` and `cell2` of our cell, one on site `s1` and second on site `s2` and use them in some agent `a`, as below.

```

val cell1 = (refInt s1)
val cell2 = (refInt s2)

agent a =
  (
    (cell2.set ag 5);
    (prNL (int.toString (cell1.get a)));
    (prNL (int.toString (cell2.get a)));
    ()
  )

```

The agent `a` first stores 5 at object `cell2`, then gets stored values from both objects and prints them out (with a newline). Distributed objects are used in some Nomadic Pict libraries.

D.10 Trading Names and Values

Nomadic Pict was designed as a language for *prototyping* distributed applications and we almost never needed to split programs into many files, compiled and executed separately on different machines. We simply spawned different parts of distributed programs dynamically on “empty” Nomadic Pict runtime systems, using agents and migration. However, occasionally it is convenient to compile and execute server and client programs (likely to be on different machines) separately and at different times.

The `Nstd/Sys` library offers two functions `publish` and `subscribe` that can be used in order to exchange channel and agent names, basic values, and any complex values which can be sent along channels at runtime, thus making possible to set up connection between different programs. Below is an example program which is split into files `server.pi` and `client.pi`.

```

{- server.pi -}
new c : ^String
val s = (this_site)
agent b = ((publish "foo" (dynamic [b s c]));
           c?p= print!p)

```

In file `server.pi`, the program creates a new channel name `c`, assigns the current site name to `s`, creates agent `b`, and publishes a record containing `c`, `s`, and `b` at the system trader. After the names are published, the program waits for a message on `c` and prints the message out. The function `publish` takes as arguments a value to be published (which must be converted to a type `Dyn`) and a string keyword (“foo” in our example) to identify the value.

```

{- client.pi -}
agent a =
  typecase (subscribe "foo" a) of

```

```
[ag:Agent si:Site ch:~String] ->
  <ag@si>ch!"Hello world!"
else print!"Type mismatch for foo"
```

In file `client.pi`, the program creates agent `a` and subscribes for the value published by the server in file `server.pi`. The function `subscribe` takes two parameters: the string keyword `"foo"` which was used to publish the value at the trader, and the name of the current agent. The function blocks until the value is available. The value returned by `subscribe` is a dynamic value which can be matched against expected types using `typecase`. If the dynamic typechecking succeeds, then basic values extracted from the dynamic value are used in our example for communication (a string “Hello world!” is passed to the server).

When the runtime system starts up, we have to specify — using options `-trader` and `-tport`, an address and port number for the runtime system selected to be a trader. By default the current runtime system is chosen.

E. NOMADIC PICT IMPLEMENTATION

Programs in high-level Nomadic Pict are compiled in the same way as they are formally specified, by translating the high-level program into the low-level language. That in turn is compiled to a core language executed by the runtime (see Figure 11). The core language is architecture-independent; its constructs correspond approximately to those of the Low Level Nomadic π calculus, extended with value types and system function calls.

The compiler and runtime are written in OCaml [Leroy 1995]. The former is around 15 000 lines (roughly double that of the Pict compiler). The runtime is only around 1700 lines, not including distributed infrastructure algorithms and standard libraries, which are, of course, written in Nomadic Pict itself.

In this appendix we describe the compiler and runtime system in more detail.

E.1 Architecture of the Compiler

The compilation of a Nomadic Pict program has the following phases: parsing the high-level program and infrastructure encoding; importing separately compiled units (e.g. standard libs); scope resolution and typechecking the high-level program and meta-definitions of the encoding; applying the encoding to generate low-level code; scope resolution and typechecking the low-level code; continuation-passing translation of the low-level code to the core language; joining imported code (if there are any bindings exported from a unit); and incremental optimisation of the core language.

Below, we describe briefly the more interesting phases. The generation of the core language from the low-level language is based on Pierce and Turner’s Pict compiler, extended with rules for the Nomadic Pict constructs; see the Pict definition [Pierce and Turner 1997] for a formal description of this translation for Pict constructs.

Importing A program consists of a collection of named *compilation units*, each comprising a sequence of `import` statements, followed by a sequence of declarations. Individual units can be compiled separately. Compilation begins with the unit that has been designated as the *main unit*. A program defined in the main unit can use the high-level constructs. If this is the case, there will be also included: a top-level

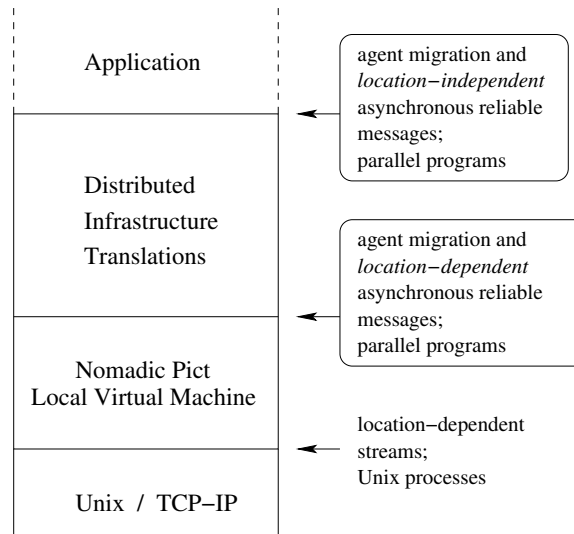


Fig. 11. The Nomadic Pict Two-Levels of Abstraction

and a compositional translation of high-level constructs. The program begins the execution from the top-level clause, which creates all the necessary daemons, and initializes any parameters of the language translation.

Scope Resolution The process of resolving variable scopes yields an alpha-renamed copy of the original term. The alpha-renamed term has the property that every bound variable is unique, so that a simplified implementation of substitution and inlining can be used.

Typechecking The typechecker performs partial type inference. Typechecking is performed twice, before and after an encoding is applied, for more precise error reporting. In the last phases, any separately compiled modules are joined and the compiler incrementally optimises the resulting core language code.

Some languages, such as ML and Haskell, which are based on the Hindley-Milner type system, can automatically infer all necessary type annotations. Pict's type system, however, is significantly more powerful than the Hindley-Milner type system (e.g. it allows higher-order polymorphism and subtyping). Thus, a simple *partial* type inference algorithm is used (the algorithm is partial, in the sense that it may sometimes have to ask the user to add more explicit type information rather than determine the types itself). The algorithm is formalised in Pierce and Turner [1997]. It exploits the situations where the type assigned to a bound variable can be completely determined by the surrounding program context. The inference is local, in the sense that it only uses the immediately surrounding context to try to fill in a missing type annotation. For example, the variable x in the input expression $c?x=e$ has type Int if the channel c is known to have type Int .

Types are erased before execution and so there is no way that type annotations in the program could affect its behaviour, except for uses of type `Dyn`, which allows data that are created dynamically to be used safely.

Applying Encodings Each high-level construct in a program is replaced by its meta definition, in such a way that free occurrences of variables in the meta definition are substituted by current variables from the program. Also certain types, such as `Agent` and `Site` defined in the program are replaced by their encodings.

Continuation Passing Style (CPS) The compiler uses a continuation passing style translation to reduce the overheads of interpreting the source program. In particular, the CPS translation is used to simplify complex expressions of the low-level language so that they fall within the core language. The complex expressions are complex values, value declarations (`val x = v P`), application (`v v1 ... vn`), and abstractions such as a “function definition” `def f (a1 a2) = v`. The CPS conversion in Pict is similar to those used in some compilers for functional languages (e.g. [Appel 1992]). In essence, it transforms a complex value expression into a process that performs whatever computation is necessary and sends the final value along a designated *continuation channel*.

A complex value is always evaluated “strictly” to yield a simple value, which is substituted for the complex expression. For example, when we write `c! [13 (v v1 v2)]`, we do not mean to send the *expression* `[13 (v v1 v2)]` along `c` but to send a simple value evaluated from this complex value. Thus, the expression must be interpreted as a core language expression that evaluates first the ‘function’ value `v`, followed by the argument values `v1` and `v2`, then calls the function instructed to return its result along the application expression’s continuation channel, and finally packages the result received along the continuation channel into a simple tuple along with the integer 13 and sends the tuple along `c`.

Optimisations In the last phase, all separately compiled units are joined with the main unit, and the compiler incrementally optimises the resulting core language program. It does a static analysis and partial evaluation of a program, reducing π computations whenever possible and removing inaccessible fragments of code. The remaining computations make up the generated or “residual” program executed by the runtime system. The Pict optimiser also checks the program’s consistency — the following conditions must hold: no unbound variables (every variable mentioned in the program must be in scope), all bound variables must be unique, static variables (i.e. ones whose value is known to be a compile-time constant) are represented as global variables in the generated code. In the current implementation of Nomadic Pict, global variables are dynamically copied to a local agent environment upon agent creation; other solutions are plausible in a more optimised version of the compiler and runtime system.

Architecture-Independent Core Language The compiler generates the core language which is executed by the Nomadic Pict runtime system. The core language is architecture-independent; its constructs correspond approximately to those of the Low Level Nomadic π calculus (extended with value types and system function calls). *Process* terms are output atoms, input and migrate prefixes, parallel compositions, processes prefixed by declarations, terminate, test-and-send, and conditional processes. There is no separate primitive for cross-network communication — these are all encoded by terms of agent migration and test-and-send. *Declarations* introduce new channels and agents. Finally, *Values* (i.e. entities that can be




Fig. 12. Architecture of the Nomadic Pict Runtime System. Abbreviations: a_i , agent IDs; c_i , channel IDs; n_i , names; v , values; P or Q , processes; e_i , local environments; s_i , strings

communicated on channels) include variables, agent and channel names, records of values, and constants (such as `String`, `Char`, `Int`, and `Bool`). Record values generalise tuple values (since the labels in a record are optional).

If a program uses only the Pict language, then it is compiled to a subclass of the core language, and an original Pict backend can be chosen to translate it to a C program, which is then compiled and executed on a single machine; see Turner [1996] for a detailed description of generating C code from Pict core language.

E.2 Architecture of the Runtime System

Because much of the system functionality, including all distributed infrastructure, is written in Nomadic Pict, the runtime has a very simple architecture (illustrated in Figure 12). It consists of two layers: the Virtual Machine and I/O server, above TCP. The implementation of the virtual machine builds on the abstract machine designed for Pict [Turner 1996].

Virtual Machine and Execution Fairness The virtual machine maintains a state consisting of an *agent store* of agent closures; the agent names are partitioned into an *agent queue*, of agents waiting to be scheduled, and a *waiting room*, of agents whose process terms are all blocked. An agent closure consists of a *run queue*, of Nomadic π process/environment pairs waiting to be scheduled (round-robin), *channel queues* of terms that are blocked on internal or inter-agent communication, and

an environment. Environments record bindings of variables to channels and basic values.

The virtual machine executes in steps, in each of which the closure of the agent at the front of the agent queue is executed for a fixed number of interactions. This ensures fair execution of the fine-grain parallelism in the language. Agents with an empty run queue wait in the waiting room. They stay suspended until some other agent sends an output term to them. The only operations that remove agent closures from the agent store are **terminate** and **migrate**. An operation **migrate** moves an agent to a remote site. On the remote site, the agent is placed at the end of the agent queue.

The agent scheduler provides *fair* execution, guaranteeing that runnable concurrent processes of all non-terminating agents will eventually be executed, and that processes waiting to communicate on a channel will eventually succeed (of course, if sufficient communication partners become available on a local or remote site). The implementation is deterministic and the language parallel operations are interleaved fairly. Non-deterministic behaviour will naturally arise because of time-dependent interactions between the abstract machine, the I/O server, and the system function calls to the operating system.

Interaction with an Operating System and User For many library functions execution consists of one or more calls to corresponding Unix I/O routines. For example, processing `print!"foo"` involves an invocation of the OCaml library call `output_string`. All interaction between the abstract behaviour of a Nomadic Pict library function and its environment (the operating system and user) occurs via invocations of *system function* calls. When a system function call reaches the front of the run queue some special processing takes place. The interpreter invokes the system function, passing all the function parameters and a result channel. The functions which can block for some time or can potentially never return (such as input from a user) will be executed within a separate execution thread, so that they do not block parallel computation. The agent operations **migrate** and **terminate** are special cases — they have to wait until all threads that execute system functions invoked inside the agent have terminated. If the system function returns any value, the Nomadic Pict program will receive it along the result channel.

I/O Server The multithreaded I/O server receives incoming agents, consisting of an agent name and an agent closure; they are unmarshalled and placed in the agent store. Note that an agent closure contains the entire state of an agent, allowing agent execution to be resumed from the point where it was suspended. Agent communication uses standard network protocols (TCP in our first implementation). The runtime system does not support any reliable protocols that are tailored for agents, such as the Agent Transfer Protocol of Lange and Aridor [1997]. Such protocols must be encoded explicitly in an infrastructure encoding — the key point in our experiments is to understand the dependencies between machines (both in the infrastructure and in application programs); we want to understand exactly how the system behaves under failure, not simply to make things that behave well under very partial failure. This is assisted by the purely local nature of the runtime system implementation.

When the runtime system starts up, the user has to specify an address for the runtime system selected to maintain the trader's map from strings to published names and values. The library functions `publish` and `subscribe`, written in Nomadic Pict, implement the whole distributed protocol which is necessary to contact the trading runtime system (so, the implementation of the I/O Server remains purely local).

F. NOMADIC PICT SYNTAX

This chapter describes the syntax of Nomadic Pict programs (for description of lexical rules and Pict syntax we use extracts from Pierce and Turner [1997], by courtesy of Benjamin C. Pierce).

F.1 Lexical Rules

Whitespace characters are space, newline, tab, and formfeed (control-L). Comments are bracketed by `{-` and `-}` and may be nested. A comment is equivalent to whitespace.

Integers are sequences of digits (negative integers start with a `-` character). Strings can be any sequence of characters and escape sequences enclosed in double-quotes. Sites can be any sequence of characters and escape sequences enclosed in double single-quote characters (`' '`), used to denote the IP address, followed by a colon and integer, to denote a port number. The escape sequences `\"`, `\n`, and `\\` stand for the characters double-quote, newline, and backslash. The escape sequence `\ddd` (where `d` denotes a decimal digit) denotes the character with code `ddd` (codes outside the range `0..255` are illegal). Character constants consist of a single quote character (`'`), a character or escape sequence, and another single quote.

Alphanumeric identifiers begin with a symbol from the following set:

```
a b c d e f g h i j k l m n o p q r s t u v w x y z
A B C D E F G H I J K L M N O P Q R S T U V W X Y Z
```

Subsequent symbols may contain the following characters in addition to those mentioned above:

```
0 1 2 3 4 5 6 7 8 9 ' ,
```

Symbolic identifiers are non-empty sequences of symbols drawn from the following set:

```
~ * % \ + - < > = & | @ $ , ' ,
```

F.2 Reserved Words

The following symbols are reserved words:

Agent	agent	and	Bool	ccode	Char	def
dynamic	else	false	if	iflocal	import	inline
Int	in	migrate	new	now	of	program
rec	run	Site	String	terminate	then	timeout
to	Top	toplevel	true	Type	type	typecase
val	switch	wait	where	with	@	^
\	/	.	;	:	=	
!	#	?	?*	-	<	>
->	{	([})]

F.3 Concrete Grammar

For each syntactic form, we note whether it is part of the core language (C), the language for expressing encodings (T), a derived form (D), an optional type annotation that is filled in during type reconstruction if omitted by the programmer (R), or an extra-linguistic feature (E). Syntactic forms characteristic for the Nomadic Pict language are marked by N .

Compilation units

TopLevel ::=

<i>Import ... Import Dec ... Dec</i>	E	Compilation unit
<i>Import ... Import TopDec ... TopDec</i>	EN	Compilation unit

Import ::=

<i>import String</i>	E	Import statement
----------------------	-----	------------------

Top-level declarations

TopDec ::=

<i>Dec</i>	Declaration
<i>{ Agent } = Type</i>	TN Agent type
<i>{ Site } = Type</i>	TN Site type
<i>program Id : Type = Proc</i>	TN Program declaration
<i>{ topLevel Id Id } Type = Proc</i>	TN Toplevel declaration
<i>{ def Id } Id Abs</i>	TN Process abstraction
<i>{ agent Id = Id in Id } Id = Proc</i>	TN Agent creation
<i>{ migrate to Id Id } Id = Proc</i>	TN Agent migration
<i>{ Id ?* Id = Id } Id = Proc</i>	TN Replicated input
<i>{ < Id @ Id > Id ! Id } Id = Proc</i>	TN Output to agent on site
<i>{ < Id > Id ! Id } Id = Proc</i>	TN Output to adjacent agent
<i>{ iflocal < Id > Id ! Id then Proc else Proc } Id = Proc</i>	TN Test-and-send to agent
<i>{ Id @ Id ! Id } Id = Proc</i>	TN Location-independent output
<i>{ do String Id in Id } Id = Proc</i>	TN Macro definition

Declarations

Dec ::=

<i>new Id</i> : <i>Type</i>	<i>C</i>	Channel creation
<i>val Pat</i> = <i>Val</i>	<i>D</i>	Value binding
<i>run Proc</i>	<i>D</i>	Parallel process
<i>Val</i> ;	<i>D</i>	Sequential execution
<i>inline def Id Abs</i>	<i>D</i>	Inlinable definition
<i>def Id₁ Abs₁ and ... and Id_n Abs_n</i>	<i>C</i>	Recursive definition ($n \geq 1$)
<i>type Id</i> = <i>Type</i>	<i>D</i>	Type abbreviation
<i>type (Id KindId₁ ... KindId_n) = Type</i>	<i>D</i>	Type operator abbrev ($n \geq 1$)
<i>now (Id Flag ... Flag)</i>	<i>E</i>	Compiler directive
<i>agent Id₁ = Proc₁ and ... and Id_n = Proc_n</i>	<i>CN</i>	Agent creation ($n \geq 1$)
<i>agent Id₁ = Proc₁ and ... and Id_n = Proc_n in</i>	<i>CN</i>	Agent creation ($n \geq 1$)
<i>migrate to Val</i>	<i>CN</i>	Migrate to site
<i>do String Val</i>	<i>TN</i>	Macro inlining
<i>do String Val in</i>	<i>TN</i>	Macro inlining
<i>{ Id } Val</i>	<i>TN</i>	Declaration inlining
Flag ::=		
<i>Id</i>	<i>E</i>	Ordinary flag
<i>Int</i>	<i>E</i>	Numeric flag
<i>String</i>	<i>E</i>	String flag

Abstractions

Abs ::=		
<i>Pat</i> = <i>Proc</i>	<i>C</i>	Process abstraction
<i>(Label FieldPat ... Label FieldPat) RType</i> = <i>Val</i>	<i>D</i>	Value abstraction

Patterns

Pat ::=		
<i>Id RType</i>	<i>C</i>	Variable pattern
<i>[Label FieldPat ... Label FieldPat]</i>	<i>C</i>	Record pattern
<i>(rec RType Pat)</i>	<i>C</i>	Rectype pattern
<i>_ RType</i>	<i>C</i>	Wildcard pattern
<i>Id RType @ Pat</i>	<i>C</i>	Layered pattern
<i>! Id</i>	<i>T</i>	Reference pattern
FieldPat ::=		
<i>Pat</i>	<i>C</i>	Value field
<i># Id Constr</i>	<i>C</i>	Type field

Type constraints

Constr ::=		
<i>(empty)</i>	<i>D</i>	No constraint
<i>< Type</i>	<i>C</i>	Subtype constraint
<i>= Type</i>	<i>C</i>	Equality constraint

Processes

Proc ::=		
<i>Val ! Val</i>	<i>C</i>	Output atom
<i>Val ? Abs</i>	<i>C</i>	Input prefix
<i>Val ?* Abs</i>	<i>CN</i>	Replicated input
<i>wait Val ? Abs timeout Val -> Proc</i>	<i>CN</i>	Timed input
<i>< Val @ Val > Val ! Val</i>	<i>DN</i>	Output to agent on site
<i>< Val > Val ! Val</i>	<i>DN</i>	Output to adjacent agent
<i>iflocal < Val > Val ! Val then Proc else Proc</i>	<i>CN</i>	Test-and-send to agent
<i>Val @ Val ! Val</i>	<i>DN</i>	Location-independent output
<i>()</i>	<i>C</i>	Null process
<i>(Proc₁ ... Proc_n)</i>	<i>C</i>	Parallel composition ($n \geq 2$)
<i>(Dec₁ ... Dec_n Proc)</i>	<i>C</i>	Local declarations ($n \geq 1$)
<i>if Val then Proc else Proc</i>	<i>C</i>	Conditional
<i>terminate</i>	<i>C</i>	Agent termination
<i>typecase Val of Pat₁ -> Proc₁ ... Pat_n -> Proc_n else Proc_{n+1}</i>	<i>DN</i>	Type matching ($n \geq 1$)
<i>switch RType Val of (Id₁ > Pat₁ -> Proc₁ ... Id_n > Pat_n -> Proc_n)</i>	<i>DN</i>	Variant matching ($n \geq 1$)
<i>{ Id } Val</i>	<i>TN</i>	Process inlining

Values

<i>Val ::=</i>		
<i>Const</i>	<i>C</i>	Constant
<i>Path</i>	<i>C</i>	Path
<i>\ Abs</i>	<i>D</i>	Process abstraction
[<i>Label FieldVal ... Label FieldVal</i>]	<i>C</i>	Record
if <i>RType Val then Val else Val</i>	<i>D</i>	Conditional
(<i>Val RType with Label FieldVal ... Label FieldVal</i>)	<i>D</i>	Field extension
(<i>Val RType where Label FieldVal ... Label FieldVal</i>)	<i>D</i>	Field override
(<i>RType Val Label FieldVal ... Label FieldVal</i>)	<i>D</i>	Application
(<i>Val > Val₁ ... Val_n</i>)	<i>D</i>	Right-assoc application ($n \geq 2$)
(<i>Val < Val₁ ... Val_n</i>)	<i>D</i>	Left-assoc application ($n \geq 2$)
(<i>rec RType Val</i>)	<i>C</i>	Rectype value
(<i>Dec₁ ... Dec_n Val</i>)	<i>D</i>	Local declarations ($n \geq 1$)
(<i>ccode Int Id String FieldVal ... FieldVal</i>)	<i>E</i>	Inline C code (Pict only)
(<i>ccode Int Id String FieldVal ... FieldVal</i>)	<i>EN</i>	System function call
(<i>dynamic Val RType</i>)	<i>DN</i>	Typed value
[<i>Id > Val</i>]	<i>DN</i>	Variant
<i>typecase RType Val of Pat₁ -> Val₁ ... Pat_n -> Val_n else Val_{n+1}</i>	<i>DN</i>	Type matching ($n \geq 1$)
<i>switch RType Val of (Id₁ > Pat₁ -> Val₁ ... Id_n > Pat_n -> Val_n)</i>	<i>DN</i>	Variant matching ($n \geq 1$)
{ { <i>Id</i> } }	<i>TN</i>	Value inlining
<i>Path ::=</i>		
<i>Id</i>	<i>C</i>	Variable
<i>Path . Id</i>	<i>C</i>	Record field projection
<i>FieldVal ::=</i>		
<i>Val</i>	<i>C</i>	Value field
# <i>Type</i>	<i>C</i>	Type field
<i>Const ::=</i>		
<i>String</i>	<i>C</i>	String constant
<i>Char</i>	<i>C</i>	Character constant
<i>Int</i>	<i>C</i>	Integer constant
<i>true</i>	<i>C</i>	Boolean constant
<i>false</i>	<i>C</i>	Boolean constant

Types

<i>Type ::=</i>		
<i>Top</i>	<i>C</i>	Top type
<i>Id</i>	<i>C</i>	Type identifier
<i>~ Type</i>	<i>C</i>	Input/output channel
<i>! Type</i>	<i>C</i>	Output channel
<i>/ Type</i>	<i>C</i>	Responsive output channel
<i>? Type</i>	<i>C</i>	Input channel
<i>Int</i>	<i>C</i>	Integer type
<i>Char</i>	<i>C</i>	Character type
<i>Bool</i>	<i>C</i>	Boolean type
<i>String</i>	<i>C</i>	String type
[<i>Label FieldType ... Label FieldType</i>]	<i>C</i>	Record type
(<i>Type with Label FieldType ... Label FieldType</i>)	<i>D</i>	Record extension
(<i>Type where Label FieldType ... Label FieldType</i>)	<i>D</i>	Record field override
<i>\ KindedId₁ ... KindedId_n = Type</i>	<i>C</i>	Type operator ($n \geq 1$)
(<i>Type Type₁ ... Type_n</i>)	<i>C</i>	Type application ($n \geq 1$)
(<i>rec KindedId = Type</i>)	<i>C</i>	Recursive type
<i>Agent</i>	<i>CN</i>	Agent type
<i>Site</i>	<i>DN</i>	Site type
<i>Dyn</i>	<i>DN</i>	Dynamic type
[<i>Id₁ > Type₁ ... Id_n > Type_n</i>]	<i>DN</i>	Variant type
{ <i>Id</i> }	<i>TN</i>	Type inlining
<i>FieldType ::=</i>		
<i>Type</i>	<i>C</i>	Value field
# <i>Id Constr</i>	<i>C</i>	Type field

RType ::=

$\langle \text{empty} \rangle$	R	Omitted type annotation
$: \text{Type}$	C	Explicit type annotation

Kinds

Kind ::=

$(\text{Kind}_1 \dots \text{Kind}_n \rightarrow \text{Kind})$	C	Operator kind ($n \geq 1$)
Type	C	Type kind

KindedId ::=

$\text{Id} : \text{Kind}$	C	Explicitly-kinded identifier
Id	D	Implicitly-kinded identifier

Labels

Label ::=

$\langle \text{empty} \rangle$	C	Anonymous label
$\text{Id} =$	C	Explicit label

G. THE INTERMEDIATE LANGUAGE (EXCERPT)

G.1 Typing Rules for the Intermediate Language

$$\frac{\Theta \vdash \text{map} \in \text{List} [\text{Agent}^s \text{Site}] \wedge \text{consolidate}(\text{map}) = [a_1 s_1] :: \dots :: [a_n s_n] :: \mathbf{nil} \wedge \{a_1, \dots, a_n\} = \text{agents}(\Theta) \wedge \forall i \in \{1, \dots, n\}. \Theta \vdash a_i @ s_i}{\Theta \vdash \text{map} \text{ ok}}$$

$$\frac{\text{msgQ} = \prod_{i \in I} \text{msgReq}(\{T_i\} [a_i c_i v_i]) \quad \forall i \in I. \Theta \vdash [a_i c_i v_i] \in [\text{Agent}^s \sim^w T_i T_i]}{\Theta \vdash \text{msgQ} \text{ ok}}$$

$$\frac{\Theta \vdash a @ s}{\Theta \vdash_a \text{FreeA}(s) \text{ ok}}$$

$$\frac{\Theta, b : \text{Agent}^Z @ s \vdash_b P \quad \Theta, b : \text{Agent}^Z @ s \vdash_a Q}{\Theta \vdash a @ s \vdash \Theta, b : \text{Agent}^Z @ s} \quad \frac{}{\Theta \vdash_a \text{RegA}(b Z s P Q) \text{ ok}}$$

$$\frac{\Theta \vdash_a P \quad \Theta \vdash s \in \text{Site}}{\Theta \vdash_a \text{MtingA}(s P) \text{ ok}}$$

$$\frac{\Theta \vdash_a P \quad \Theta \vdash s \in \text{Site}}{\Theta \vdash_a \text{MrdyA}(s P) \text{ ok}}$$

$$\frac{\forall a \in \text{dom}(\mathbf{A}) \exists P, \mathbf{E}. \mathbf{A}(a) = [P \mathbf{E}] \wedge \Theta \vdash_a P \wedge \Theta \vdash_a \mathbf{E} \text{ ok} \quad \exists_{\leq 1} a \in \text{dom}(\mathbf{A}). \exists Q, s, R. \mathbf{A}(a) = [Q \text{MrdyA}(s R)]}{\Theta \vdash \mathbf{A} \text{ ok}}$$

$$\frac{\Phi, \Delta \vdash \text{map} \text{ ok} \quad \Phi, \Delta \vdash \text{msgQ} \text{ ok} \quad \Phi, \Delta \vdash \mathbf{A} \text{ ok} \quad \vdash \Phi \text{ ok} \quad \text{dom}(\mathbf{A}) = \text{dom}(\text{map})}{\Phi \vdash \text{eProg}(\Delta; [\text{map} \text{msgQ}]; \mathbf{A}) \text{ ok}}$$

G.2 LTS for the Intermediate Language

$$\begin{array}{c}
\frac{\mathbf{A}(a) = [P|Q \mathbf{E}] \quad \Phi, \Delta \Vdash_a P \xrightarrow{\tau} @_a P'}{P \equiv \mathbf{if} \ v \ \mathbf{then} \ P_1 \ \mathbf{else} \ P_2 \ \vee \ P \equiv \mathbf{let} \ p = ev \ \mathbf{in} \ P_1 \ \vee} \\
\frac{P \equiv (c!v \mid c?p \rightarrow R) \quad \vee \ P \equiv (c!v \mid *c?p \rightarrow R)}{\Phi \Vdash \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A} \oplus a \mapsto [P'|Q \mathbf{E}])} \\
\\
\frac{\mathbf{A}(a) = [((\mathbf{create}^Z \ b = P \ \mathbf{in} \ Q) \mid R) \ \mathbf{FreeA}(s)] \quad b \notin \text{dom}(\Phi, \Delta)}{\Phi \Vdash \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A} \oplus a \mapsto [R \ \mathbf{RegA}(b \ Z \ s \ P \ Q)])} \\
\\
\frac{\mathbf{A}(a) = [((\mathbf{migrate} \ \mathbf{to} \ s \rightarrow P) \mid Q) \ \mathbf{FreeA}(s')]}{\Phi \Vdash \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A} \oplus a \mapsto [Q \ \mathbf{MtingA}(s \ P)])} \\
\\
\frac{\mathbf{A}(a) = [((b@?)c!v \mid P) \ \mathbf{E}] \quad (\Phi, \Delta)(c) = \sim^I T}{\Phi \Vdash \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ} \mid \mathit{msgReq}(\{T\} [b \ c \ v]); \mathbf{A} \oplus a \mapsto [P \ \mathbf{E}])} \\
\\
\frac{\mathbf{A}(a) = [R \ \mathbf{RegA}(b \ Z \ s \ P \ Q)] \quad b \notin \text{dom}(\Phi, \Delta)}{\mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A}) \ \text{idle}} \\
\frac{\Phi \Vdash \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta, b : \mathbf{Agent}^Z @s; [[b \ s] :: \mathit{map} \ \mathit{msgQ}]; \mathbf{A} \oplus a \mapsto [Q|R \ \mathbf{FreeA}(s)] \oplus b \mapsto [P \ \mathbf{FreeA}(s)])}{\Phi \Vdash \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A}) \ \text{idle}} \\
\\
\frac{\mathbf{A}(a) = [R \ \mathbf{MtingA}(s \ P)] \quad \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \ \text{idle}}{\Phi \Vdash \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A} \oplus a \mapsto [R \ \mathbf{MrdyA}(s \ P)])} \\
\\
\frac{\mathbf{A}(a) = [R \ \mathbf{MrdyA}(s \ P)]}{\Phi \Vdash \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta \oplus a \mapsto s; [[a \ s] :: \mathit{map} \ \mathit{msgQ}]; \mathbf{A} \oplus a \mapsto [P|R \ \mathbf{FreeA}(s)])} \\
\\
\frac{\mathbf{eProg}(\Delta; \mathbf{D}; \mathbf{A}) \ \text{idle} \quad \mathbf{A}(a) = [P \ \mathbf{E}]}{\Phi \Vdash \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ} \mid \mathit{msgReq}(\{T\} [a \ c \ v]); \mathbf{A}) \xrightarrow{\tau} \mathbf{eProg}(\Delta; [\mathit{map} \ \mathit{msgQ}]; \mathbf{A} \oplus a \mapsto [c!v|P \ \mathbf{E}])}
\end{array}$$