Theory and Practice of Logic Programming **1** (3): 359–380, May 2001.

Printed in the United Kingdom © 2001 Cambridge University Press

High-level networking with mobile code and first order AND-continuations

PAUL TARAU

Department of Computer Science, University of North Texas, P.O. Box 311366, Denton, TX 76203, USA
(*e-mail:* tarau@cs.unt.edu)

VERONICA DAHL

Logic and Functional Programming Group, Department of Computing Sciences, Simon Fraser University, Canada (*e-mail:* veronica@cs.sfu.ca)

Abstract

We describe a scheme for moving living code between a set of distributed processes coordinated with unification based Linda operations, and its application to building a comprehensive Logic programming based Internet programming framework. *Mobile threads* are implemented by capturing first order continuations in a compact data structure sent over the network. Code is fetched lazily from its original base turned into a server as the continuation executes at the remote site. Our code migration techniques, in combination with a dynamic recompilation scheme, ensure that heavily used code moves up smoothly on a speed hierarchy while volatile dynamic code is kept in a quickly updatable form. Among the examples, we describe how to build programmable client and server components (Web servers, in particular) and mobile agents.

KEYWORDS: mobile computations, remote execution, networking, Internet programming, first order continuations, Linda coordination, blackboard-based logic programming, mobile agents, dynamic recompilation, code migration

1 Introduction

Data mobility has been present since the beginning of networked computing, and is now used in numerous applications – from remote consultation of a database, to Web browsing.

Code mobility followed, often made transparent to users as with network files systems (i.e. Sun's NFS). Java's ability to execute applets directly in client browsers, can be seen as its most recent incarnation.

Migrating the state of the computation from one machine or process to another, still requires a separate set of tools. Java's remote method invocations (RMI) add *control mobility* and a (partially) automated form of *object mobility* i.e. integrated code (class) and data (state) mobility. The Oz 2.0 distributed programming proposal of Van Roy *et al.* (1997) makes *object mobility* more transparent, although the mobile entity is still the state of the objects, not 'live' code.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

Mobility of 'live code' is called *computation mobility* (Cardelli, 1997). It requires interrupting the execution, moving the state of a runtime system (stacks, for instance) from one site to another, and then resuming execution. Clearly, for some languages, this can be hard or completely impossible to achieve.

Telescript and General Magic's Odyssey (General Magic Inc., 1997) agent programming framework, IBM's Java based *aglets*, as well as Luca Cardelli's Oblique (Bharat and Cardelli, 1995), have pioneered implementation technologies with *computation mobility*.

This paper will show that we can achieve full *computation mobility* through our *mobile threads*, with no need for a specially designed new language. It is implemented by a surprisingly small, source level modification of the BinProlog system, which takes advantage of the availability of *first order continuations*[1] as well as of BinProlog's high level networking primitives. Mobile threads complete our Logic Programming based Internet programming infrastructure, built in view of creating Prolog components which can interoperate with mainstream languages and programming environments. *Mobile threads* can be seen as a refinement of *mobile computations*, as corresponding to *mobile partial computations* of any granularity. *Mobile agents* can be seen as a collection of synchronized *mobile threads* sharing common state (Tarau and Dahl, 1996). We achieve synchronization using a variant of the Linda coordination protocol.

The paper is organized as follows:

- Section 2 describes our networking infrastructure and Linda based client/server components
- Section 3 describes our code migration and code acceleration techniques (dynamic recompilation)
- Section 4 describes our mobile computation mechanism, as follows: section 4.2 introduces engines and threads, section 4.3) reviews the underlying binarization mechanism used to implement our first order continuations, subsection 4.4 explains how we implement thread mobility by capturing continuations (subsection 4.4.1) and moving them from their base to their target (section 4.4.2), how this can be emulated with remote predicate calls (section 4.4.3) and how mobile agents can be built within our framework (subsection 4.5)
- Section 5 discusses related work
- Section 6 presents our conclusions and future work.

The main 'paradigm independent' novelties of our contribution are:

- use of first order continuations for implementing *mobile computations*
- a flexible thread mobility algorithm expressed in terms of client-server role alternation and communication through Linda operations
- a technique, based on intuitionistic assumptions, for dealing with complex networking code component-wise

[1] That is, continuations (representations of future computations) accessible as an ordinary data structure – a Prolog term in this case.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

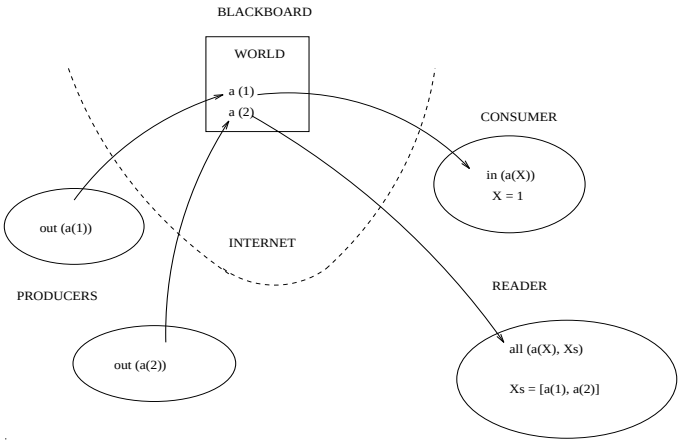


Fig. 1. Basic Linda operations.

Our contributions are synergetically integrated into a powerful agent building infrastructure, which brings together logic programming based knowledge processing, Linda-style coordination, and live code migration through mobile threads.

2 Basic Linda and remote execution operations

2.1 Coordination of Linda clients

Our networking constructs are built on top of the popular Linda (Carriero and Gelernter, 1989; Ciancarini, 1994; De Bosschere and Tarau, 1996; P. Ciancarini *et al*., 1996; Sun Microsystems, 1999) coordination framework, enhanced with unification based pattern matching, remote execution and a set of simple client-server components merged together into a scalable peer-to-peer layer, forming a 'web of interconnected worlds':

```
out(X): puts X on the server
in(X):  waits until it can take an object
        matching X from the server
all(X,Xs): reads the list Xs matching X
        currently on the server
```

The presence of the `all/2` collector avoids the need for backtracking over multiple remote answers. Note that the only blocking operation is `in/1`. Typically, distributed programming with Linda coordination follows consumer-producer patterns (see Figure 1) with added flexibility over message-passing communication through associative search. Blocking `rd/1`, which waits until a matching term becomes available, without removing it, is easily emulated in terms of `in/1` and `out/1`, while non-blocking `rd/1` is emulated with `all/2`.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

2.2 Remote execution mechanisms

The implementation of arbitrary remote execution is easy in a Linda + Prolog system, due to Prolog's *metaprogramming* abilities, which allow us to send arbitrary Prolog terms over the network in a uniform way, without the need for implementing complex serialization/remote object mechanisms. Our primitive remote call operation is:

```
host(Other_machine)=>>
   remote_run(Answer,RemoteGoal).
```

It implements deterministic *remote predicate calls* with (first)-answer or `'no'` returned to the calling site.

For instance, to iterate over the set of servers forming the receiving end of our 'Web of Worlds', after retrieving the list from a 'master server' which constantly monitors them making sure that the list reflects login/logout information, we simply override `host/1` and `port/1` with intuitionistic implication `=>>` (Tarau, 1998a; Dahl *et al.*, 1997):

```
ask_all_servers(Channel,Servers,Query):-
  member(server_id(Channel,H,P),Servers),
  host(H)=>>port(P)=>>
    ask_a_server(Query,_),
  fail;true.
```

Note that a `Channel` pattern is used to select a subset of relevant servers, and in particular, when `Channel` is a 'match all' free logical variable, all of them. By using term subsumption this allows building sophisticated 'publish/subscribe' communication pattern hierarchies.

2.3 Servants: basic Linda agents

A *servant* is one of the simplest possible *agents*, which pulls commands from a server and runs them locally:

```
servant:-
  in(todo(Task)),
  call(Task),
  servant.
```

Note that *servant* is started as a background thread. No 'busy wait' is involved, as the servant's thread blocks until `in(todo(Task))` succeeds. More generally, distributed event processing is implemented by creating a 'watching' agent attached to a thread for each pattern.

As *servants* pulling commands are operationally indistinguishable from *servers* acting upon clients' requests, they can be used as emulators for *servers*. A class of obvious applications of this ability is their use as pseudo-servers running on machines with dynamically allocated IP addresses (as offered by most ISPs today), laying behind firewalls. This mechanism also works when, because of security restrictions, server components cannot be reached from outside, as in the case of *Java applets* which cannot listen on ports of the client side machine.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

2.4 Server side code

Servants as well as other clients can connect to BinProlog *servers*. Higher order *call/N* (Mycroft and O'Keefe, 1984), combined with intuitionistic assumptions ''=>>'', are used to pass arbitrary *interactors* to generic server code:

```
run_server(Port):-
  new_server(Port,Server),
  register_server(Port),
  server(Server)=>>server_loop,
  close_socket(Server).

server_loop:-
  repeat,
    interact,
  assumed(server_done),
  !.

interact:-
  assumed(Interactor),
  assumed(Server),
  % higher-order call to interactor
  call(Interactor,Server).
```

Note the use of a specialized *server-side* interpreter `server_loop`, configurable through the use of higher-order 'question/answer' closures we have called *interactors*.

The components of a 'generic' default server can be overridden through the use of *intuitionistic implication* to obtain customized special purpose servers. The use of intuitionistic implications (pioneered by Miller's work (Miller, 1989)) helps to overcome (to some extent) Prolog's lack of object oriented programming facilities, by allowing us to 'inject' the right *interactor* into the generic (and therefore reusable) interpreter. BinProlog's ''=>>'' temporarily assumes a clause in `asserta` order, i.e. at the beginning of the predicate. The assumption is scoped to be only usable to prove its right side goal and vanishes on backtracking. We refer elsewhere (Tarau, 1998a; Tarau *et al*., 1996a; Dahl *et al*., 1997) for more information on assumptions and their applications.

2.5 Modular HTTP server component building: do or delegate

We will show in this section a typical application of our component based server building technology: how to enhance efficiently the HTTP protocol, to handle server side Prolog scripts directly, without using CGIs or vendor specific server side extensions.

The top goal of the HTTP server looks as follows:

```
run_http_server:-
  server_action(http_server_action)=>>run_server.

http_server_action(ServiceSocket):-
  socket(ServiceSocket)=>>http_server_step(ServiceSocket).
```

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

The http_server_action is passed to the inner server loop using intuitionistic implication. This allows reusing general server logic, independently of a particular protocol. The action itself is described in http_server_step: it consists of preparing a fall-back mechanism to a standard HTTP server, unless the request is for a file recognized as Prolog code, using the *redirection* facilities built in the HTTP protocol.

```
http_server_step(Socket):-
  ( assumed(fallback_server(FallBackServer))->true
  ; FallBackServer="http://localhost:80"
  ),
  server_try(Socket,sock_readln(Socket,Question)),
  http_get_client_header(Socket,Header),
  http_process_query(Socket,Question,Header,FallBackServer).
```

Our very simple query processor uses Assumption Grammars (Dahl *et al*., 1997). Their ability to handle multiple DCG streams (Tarau *et al*., 1996a) is instrumental, as we use more than one independent grammar processor in the process.

```
http_process_query(Socket,Qs,Css,FallBackServer):-
  #<Qs, % sets input string from grammar
  match_word("GET "),
  match_before(" ",PathFile,_),
  match_word("HTTP/"),
  #>_version, %
  !,
  split_path_file(PathFile,Ds,Fs),
  !,
  has_text_file_sufix(Fs,Suf),
  !,
  ( Suf=".pro"->http_process_local(Socket,Ds,Fs,Suf,Css)
  ; write('redirecting '),write_chars(Ds),write_chars(Fs),nl,
    http_send_line(Socket,"HTTP/1.0 302 Found"),
    make_cmd0(["Location: ",FallBackServer,Ds,Fs],Redirect),
    http_send_line(Socket,Redirect),
    http_send_line(Socket,"")
  ),
  close_socket(Socket).
```

Our HTTP server component fits in 76 lines of code, and can be used to set up Prolog based Web processing by simply starting it in a command window on any Unix machine or PC. This involves no execution overhead, as is the case with CGI scripts. It basically offers the advantages of Apache server side includes (SSIs) or Microsoft's Active Server Pages (ASPs), without requiring integration into a server, often subject to using the languages the server supports.

2.6 Master servers: Connecting a web of virtual places

A *virtual place* is implemented as a server listening on a port which can spawn clients in the same or separate threads interacting with other servers.

A master server on a 'well-known' host/port is used as a registration service to

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

exchange identification information among peers composed of clients and a server, usually running as threads of the same process.

As in the case of the HTTP server, we can derive a master server by specializing its interactor components through intuitionistic implications.

3 Code migration and code acceleration techniques

We have seen that setting up a self contained networking infrastructure (Web protocols included) is a fairly simple task. The next step is emphasizing mobile agent support, which is particularly promising in synergy with knowledge processing capabilities – a key strength of Logic Programming systems.

3.1 Lazy code fetching

In BinProlog, code is fetched lazily over the network, one predicate at a time, as needed by the execution flow.

Code is cached in a local database and then dynamically recompiled on the fly if usage statistics indicate that it is *not volatile* and it is *heavily used* locally.

The following operations

```
host(Other_machine)=>>rload(File).
host(Other_machine)=>>code(File)=>>TopGoal.
```

allow fetching remote files `rload/1` or on-demand fetching of a predicate at a time from a remote host during execution of `TopGoal`.

This is basically the same mechanism as the one implemented for Java applet code fetching, except that we have also implemented a caching mechanism, at predicate level (predicates are cached as dynamic code on the server to efficiently serve multiple clients).

3.2 Dynamic recompilation

Dynamic recompilation is used on the client side to speed-up heavily used, relatively non-volatile predicates. With dynamically recompiled consulted code, listing of sources and dynamic modification to any predicate is available, while the average performance stays close to statically compiled code (usually within a factor of 2-3).

Our implementation of dynamic recompilation for BinProlog is largely motivated by the difficulty/complexity of relying on the programmer to specify execution methods for remote code.

The intuition behind the dynamic recompilation algorithm of BinProlog is that *update* vs. *call* based *statistics* are associated to each predicate declared or detected as dynamic. Dynamic (re)compilation is triggered for relatively non-volatile predicates, which are promoted on the *'speed-hierarchy'* to a faster implementation method (interpreted -> bytecode -> native). The process is restarted from the 'easier to change' interpreted representation, kept in memory in a compact form, upon an update.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

We can describe BinProlog's dynamic *'recompilation triggering statistics'* through a simple 'thermostat' metaphor. *Updates* (assert/retract) to a predicate have the effect of increasing its associated 'temperature', while *Calls* will decrease it. Non-volatile ('cool') predicates are dynamically recompiled, while recompilation is avoided for volatile ('hot') predicates. A *ratio* based on cooling factors (number of calls, compiled/interpreted execution speed-up, etc.) and heating factors (recompilation time, number of updates etc.) smoothly adjusts for optimal overall performance, usually within a factor of two from static code.

4 Computation mobility with threads and continuations

4.1 Why do computations need to be mobile?

Advanced *mobile object* and *mobile agents* agent systems have been built on top of Java's dynamic class loading and its new reflection and remote method invocation classes. IBM Japan's Aglets or General Magic's Odyssey provide comprehensive mobility of code and data. Moreover, data is encapsulated as state of objects. This property allows us to protect its sensitive components more easily. Distributed Oz/Mozart provides fully transparent movement of objects over the network, giving the illusion that the same program runs on all the computers.

So *why do we need* the apparently more powerful concept of mobile 'live code', i.e. mobile execution state?

Our answer to this question is that live mobile code is needed because is still *semantically simpler* than mobile object schemes. Basically, all that a programmer needs to know is that his or her program has moved to a new site and it is executing there. A unique (in our case `move_thread`) primitive, with an intuitive semantics, is all that needs to be learned. When judging about how appropriate a language feature is, we think that the way it looks to the end user is among the most important ones. For this reason, mobile threads are competitive with sophisticated *object mobility* constructs on 'end-user ergonomy' grounds, while being fairly simple to implement, as we have shown, in languages in which continuations can be easily represented as data structures.

And *what if the host language does not offer first order continuations*? A simple way around this is to implement on top of it a script interpreter (e.g. a subset of Scheme or Prolog) which does support them! As it is a good idea to limit code migration to lightweight scripts anyway, this is a very practical solution for either C/C++ or Java based mobile code solutions, without requiring class-specific serialization mechanisms.

4.2 Engines and answer threads

Mobile computations really make sense only when multiple computation threads can coexist at a given place. We build this infrastructure in two steps: an engine encapsulating the state of an independent computation, and a thread actually running it. Note that engines can also be used without multi-threading, as a form of coroutining.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

4.2.1 Engines

BinProlog allows launching multiple Prolog engines having their own stack groups (heap, local stack and trail). An engine can be seen as an abstract data-type which produces a (possibly infinite) stream of solutions as needed. To create a new engine, we use:

```
create_engine(+HeapSize,+StackSize,+TrailSize,-Handle)
```

or, by using default parameters for the stacks:

```
create_engine(-Handle)
```

The `Handle` is a unique integer denoting the engine for further processing. To 'fuel' the engine with a goal and an expected answer variable we use:

```
load_engine(+Handle,+Goal,+AnswerVariable)
```

No processing, except the initialization of the engine takes place, and no answer is returned with this operation.

To get an answer from the engine we use:

```
ask_engine(+Handle,-Answer)
```

Each engine has its own heap garbage collection process and backtracks independently using its choice-point stack and trail during the computation of an answer. Once computed, an answer is copied from an engine to its 'master'.

When the stream of answers reaches its end, `ask_engine/2` will simply fail. The resolution process in an engine can be discarded at any time by simply loading another goal with `load_engine/3`. This allows avoiding the cost of backtracking, for instance in the case when a single answer is needed, as well as garbage collection costs.

If for some reason we are not interested in the engine any more, we can free the space allocated to the engine and completely discard it with:

```
destroy_engine(+Handle)
```

The following example[2] in the BinProlog distribution (Tarau, 1998a) shows a sequence of the previously described operations:

```
?-create_engine(E),
  load_engine(E,append(As,Bs,[1,2]),As+Bs),
  ask_engine(E,R1),write(R1),nl,
  ask_engine(E,R2),write(R2),nl,
  destroy_engine(E).
```

Multiple 'orthogonal engines' as shown in Figure 2 enhance the expressiveness of Prolog by allowing an AND-branch of an engine to collect answers from multiple OR-branches of another engine. They give to the programmer the means to see as an abstract sequence and control, the answers produced by an engine, in a way similar to Java's `Enumeration` interface.

[2] See more in files library/engines.pl, progs/engtest.pl

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

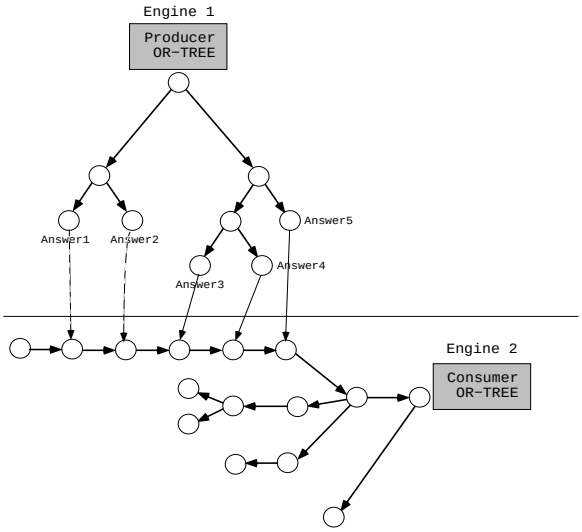


Fig. 2. Orthogonal engines.

4.2.2 Threads

Engines can be assigned to their own thread by using BinProlog's thread package (currently fully implemented on win32, Linux and Solaris platform). A unique primitive is needed,

```
ask_thread(E,R)
```

which launches a new thread `R` to perform the computation of an answer of engine `E`. On top of this facility each thread can implement a separate server or client, or become the base of a mobile agent.

Thread synchronization is provided through monitor objects, handled with:

```
synchronize_on(Monitor,Goal,Answer)
```

The thread waits until the monitor is free, executes Goal, frees the monitor and returns Answer.

The thread attached to an engine can be obtained with:

```
get_engine_thread(Engine,Thread)
```

and can be controlled directly with:

```
thread_suspend(Thread) % suspends the thread
thread_resume(Thread)  % resumes a suspended thread
thread_join(Thread)    % waits until a thread terminates
thread_cancel(Thread) % discards a thread
```

4.3 First order continuations through binarization

Having first order continuations largely simplifies the implementation of mobile code operations: a thread is suspended, its continuation is packed, sent over the network and resumed at a different place.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

We will shortly explain here BinProlog's continuation passing preprocessing technique, which results in the availability of continuations as data structures accessible to the programmer.

The binarization transformation Binary clauses have only one atom in the body (except for some in-line 'builtin' operations like arithmetics), and therefore they need no 'return' after a call. A transformation introduced by Tarau and Boyer (1990) allows us to faithfully represent logic programs with operationally equivalent binary programs.

To keep things simple, we will describe our transformations in the case of definite programs. We will follow here the notations of Tarau and De Bosschere (1993).

Let us define the *composition* operator $\oplus$ that combines clauses by unfolding the leftmost body-goal of the first argument.

Let $\mathtt{A_0:-A_1,A_2,\ldots,A_n}$ and $\mathtt{B_0:-B_1,\ldots,B_m}$ be two clauses (suppose $n > 0, m \geqslant 0$). We define

$$\mathtt{(A_0:-A_1,A_2,\ldots,A_n)} \oplus \mathtt{(B_0:-B_1,\ldots,B_m)} = \mathtt{(A_0:-B_1,\ldots,B_m,A_2,\ldots,A_n)}\theta$$

with $\theta = \mathrm{mgu}(\mathtt{A_1},\mathtt{B_0})$. If the atoms $\mathtt{A_1}$ and $\mathtt{B_0}$ do not unify, the result of the composition is denoted as $\perp$. Furthermore, as usual, we consider $\mathtt{A_0:-true,A_2,\ldots,A_n}$ to be equivalent to $\mathtt{A_0:-A_2,\ldots,A_n}$, and for any clause C, $\perp \oplus \mathtt{C} = \mathtt{C} \oplus \perp = \perp$. We assume that at least one operand has been renamed to a variant with variables standardized apart.

This Prolog-like inference rule is called LD-resolution. It has the advantage of giving a more accurate description of Prolog's operational semantics than SLD-resolution. Before introducing the binarization transformation, we describe two auxiliary transformations.

The first transformation converts facts into rules by giving them the atom `true` as body. For example, the fact `p` is transformed into the rule `p :- true`.

The second transformation, inspired by Warren (1981), eliminates the metavariables by wrapping them in a `call/1` goal. E.g., the rule `and(X,Y):-X, Y` is transformed into `and(X,Y) :- call(X), call(Y).`

The transformation of Tarau and Boyer (1990) (*binarization*) adds continuations as extra arguments of atoms in a way that preserves also first argument indexing.

Let P be a definite program and *Cont* a new variable. Let T and $E = p(T_1, ..., T_n)$ be two expressions.[3] We denote by $\psi(E, T)$ the expression $p(T_1, ..., T_n, T)$. Starting with the clause

(C) $\quad A :- B_1, B_2, ..., B_n.$

we construct the clause

(C') $\quad \psi(A, Cont) :- \psi(B_1, \psi(B_2, ..., \psi(B_n, Cont))).$

The set P' of all clauses C' obtained from the clauses of P is called the binarization of P.

The following example shows the result of this transformation on the well-known 'naive reverse' program:

[3] Atom or term.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

```
app([],Ys,Ys,Cont):-true(Cont).
app([A|Xs],Ys,[A|Zs],Cont):-
  app(Xs,Ys,Zs,Cont).

nrev([],[],Cont):-true(Cont).
nrev([X|Xs],Zs,Cont):-
  nrev(Xs,Ys,app(Ys,[X],Zs,Cont)).
```

The transformation preserves a strong operational equivalence with the original program with respect to the LD resolution rule, which is *reified* in the syntactical structure of the resulting program, i.e. each resolution step of an LD derivation on a definite program P can be mapped to an SLD-resolution step of the binarized program P'.

Clearly, continuations become explicit in the binary version of the program. We have devised a technique to access and manipulate them in an intuitive way, by modifying BinProlog's binarization preprocessor. Basically, the clauses constructed with `::-` instead of `:-` are considered as being already in binary form, and not subject therefore to further binarization. By explicitly accessing their arguments, a programmer is able to access and modify the current continuation as a 'first order object'. Note however that code *referring* to the continuation is also *part* of it, so that some care should be taken in manipulating the circular term representing the continuation from 'inside'.

4.4 Mobile threads: Take the **future** *and run*

As continuations (describing *future* computations to be performed at a given point) are first order objects in BinProlog, it is easy to extract from them a conjunction of goals representing *future* computations intended to be performed at another site, send it over the network and resume working on it at that site. The natural unit of mobility is a *thread* moving to a server executing multiple local and remotely originated threads. *Threads communicate with their local and remote counterparts, listening on ports through the Linda protocol, as described in De Bosschere and Tarau* (1996). This combination of Linda based coordination and thread mobility is intended to make building complex, pattern based agent scripts fairly easy.

4.4.1 Capturing continuations

Before moving to another site, the current continuation needs to be captured in a data structure (see Appendix I). For flexibility, a wrapper `capture_cont_for/1` is used first to restrict the scope of the continuation to a (deterministic) toplevel `Goal`. This avoids taking irrelevant parts of the continuation (like prompting the user for the next query) to the remote site inadvertently.

A unique logical variable is used through a backtrackable linear assumption `cont_marker(End)` in order to mark the end of the scope of the continuation with `end_cont(End)`.

From inside the continuation, `call_with_cont/1` is used to extract the relevant

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

segment of the continuation. Towards this end, `consume_cont(Closure,Marker)` extracts a conjunction of goals from the current continuation until Marker is reached, and then it applies `Closure` to this conjunction (calls it with the conjunction passed to `Closure` as an argument).

Extracting the continuation itself is easy, by using BinProlog's ability to accept user-defined binarized clauses (introduced with ::- instead of :-), that can access the continuation as a 'first order' object:

```
get_cont(Cont,Cont)::-true(Cont).
```

4.4.2 The continuation moving protocol

Our continuation moving protocol can be described easily in terms of synchronized *source side*[4], and *target side* operations.

Source side operations

- wrap a Goal with a unique terminator marking the end of the continuation to be captured, and call it with the current continuation available to it through a linearly assumed fact[5]
- reserve a free port P for the future code server
- schedule on the target server a sequence of actions which will lead to resuming the execution from right after the `move_thread` operation (see target side operations), return and become a code server allowing the mobile thread to fetch required predicates one a time.

Target side operations are scheduled as a sequence of goals extracted from the current continuation at the *source side*, and received over the network together with a small set of synchronization commands:

- schedule as delayed task a sequence of goals received from the source side and return,
- wait until the *source side* is in server mode,
- set up the back links to the source side as assumptions,
- execute the delayed operations representing the moved continuation,
- fetch code from the source side as needed for execution of the goals, of the moved continuations and their subcalls,
- shut down the code server on the source side.

Communication between the base and the target side proceeds through *remote predicate calls* protected with *dynamically generated passwords* shared between the two sides before the migratory component 'takes off'.

Initially the target side waits in server mode. Once the continuation is received on the target side, the source side switches to server mode (unless it already contains

[4] This will be also shortly called the *base* of the mobile thread.
[5] BinProlog's linear assumptions are backtrackable additions to the database, usable at most once.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

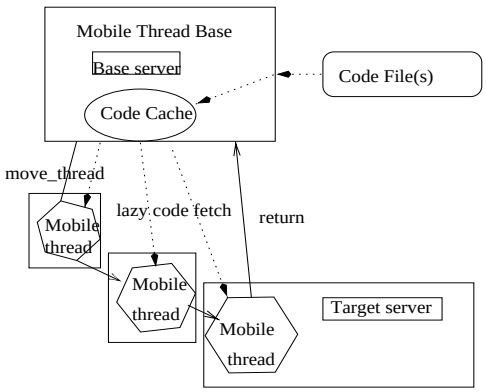


Fig. 3. Launching a mobile thread from its base.

a running server thread), ready to execute code fetching and persistent database update requests from its mobile counterpart on the target side.

Figure 3 shows the connections between a mobile thread and its base.

Note that when the base turns into a server, it offers its *own code* for remote use by the moved thread – a kind of virtual 'on demand' process cloning operation, one step at a time. Since the server actually acts as a code cache, multiple moving threads can benefit from this operation. Note also that only predicates needed for the migratory segment of the continuation are fetched. This ensures that migratory code is kept lightweight for most mobile applications. Synchronized communication, using Linda operations, can occur between the mobile thread and its base server, and through the server, between multiple mobile threads which have migrated to various places. Note that *servants* described in section 2.3 can also be used to *emulate servers* in case of the process is hosted in a component unable to listen on a server port.

As our networking infrastructure, our *mobile threads* are platform independent. Like Java, BinProlog (Tarau, 1998b) is a platform independent emulator based language. As a consequence, a thread can start on a Unix machine and move transparently to a Windows NT system and back. For faster, platform specific execution, BinProlog provides compilation to C of static code using an original partial translation technique described in Tarau *et al.* (1996b).

4.4.3 Emulating computation mobility through control mobility

As shown in Tarau *et al.* (1997b), part of the functionality of *mobile computations* can be emulated in terms of remote predicate calls combined with remote code fetching. An implicit *virtual place* (host+port) can be set as the target of the remote calls. Then, it is enough to send the top-level goal to the remote side and have it fetch the code as needed from a server at the site from where the code originates.

Note however that this is less efficient in terms of network transactions and less reliable than sending the full continuation at once as with our *mobile threads.*

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

4.5 Mobile agents

Mobile agents can be seen as a collection of synchronized *mobile threads* sharing common state (Tarau and Dahl, 1996). Agents execute a set of goals, possibly spread over a set of different processes. In addition to usual Prolog computations, agents perform remote and local transactions in coordination with other agents, and react to event driven changes occurring on blackboards.

Mobile agents are implemented by sending out *mobility threads* to a set of servers[6] registered to a given master server. A pyramidal deployment strategy can be used to efficiently implement, for instance, *push technology* through mobile agents. Communication with the base server is easily achieved through remote predicate calls with `remote_run`. Inter-agent communication can be achieved either by rendez-vous of two mobile threads at a given site, by communicating through a local blackboard, or through the base server known to all the deployed agents. Basic security of mobile agents is achieved with randomly generated passwords, required for `remote_run/1` operations, and by running them on a restricted BinProlog machine, without user-level file write and external process spawn operations. For more details on recent developments of mobile agent infrastructures and their applications, extending the framework presented in this paper, we refer elsewhere (Tarau, 1999c; Tarau, 1999b; Mikler *et al*., 1999; Tarau, 1999a; Rochefort *et al*., 1999).

In the latest version of BinProlog (Tarau, 1998a), two simple **move/0** and **return/0** operations can be used to transport a mobile agent thread to a server (listening on a default host/port) and back. The anaphoric *there* indicates that the mobile thread should be created on the server. Alternatively, *here* would indicate execution on a local thread. In this mobility scheme, the client simply waits until computation completes, when bindings for the first solution are propagated back.

```
Window 1: a mobile thread

?-there,move,println(on_server),member(X,[1,2,3]),
        return,println(back).
back
X=1

Window 2: a server

?-trust.
on_server
```

In case return is absent, computation proceeds to the end of the transported continuation.

Note that mobile computation is more expressive and more efficient than remote predicate calls as such. Basically, it *moves once*, and executes on the server *all future*

[6] Possibly filtered down to a relevant subset using a 'channel'-like pattern.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

computations of the current AND branch until a return instruction is hit, when it takes the remaining continuation and comes back. This can be seen by comparing real time execution speed for:

```
?-for(I,1,1000),remote_run(println(I)),fail.

?-there,move,for(I,1,1000),println(I),fail.
```

While the first query uses `remote_run/1` each time to send a remote task to the server, the second moves once the full computation to the server where it executes without further requiring network communications. Note that the `move/0`, `return/0` pair cut nondeterminism for the transported segment of the current continuation. This avoids having to transport state of the choice-point stack as well as implementation complexity of multiple answer returns and tedious distributed backtracking synchronization. Surprisingly, this is not a strong limitation, as the programmer can simply use something like:

```
?-there,move,findall(X,for(I,1,1000),Xs),return,member(X,Xs).
```

to emulate (finite!) nondeterministic remote execution, by collecting all solutions at the remote side and exploring them through (much more efficient) local backtracking after returning.

5 Related work

Remote execution and code migration techniques are pioneered by several authors (Almes *et al.*, 1985; Jul *et al.*, 1988; Stamos and Gifford, 1990). Support for Remote Procedure Calls (RPC) are part of major operating systems like Sun's Solaris and Microsoft's Windows NT.

A very large number of research projects have recently started on mobile computations/mobile agent programming. Among the pioneers, Kahn and Cerf's Knowbots (Kahn and Cerf, 1988). Among the most promising recent developments, Luca Cardelli's Oblique project at Digital and mobile agent applications (Bharat and Cardelli, 1995) and IBM Japan's aglets (IBM, 1999). Mobile code technologies are pioneered by General Magic's Telescript (see General Magic inc (1997) for their last Java based *mobile agent* product). General Magic's Portico software combines mobile code technologies and voice recognition based command language (MagicTalk).

Another mobility framework, sharing some of our objectives towards transparent high level distributed programming, is built on top of Distributed Oz (Van Roy *et al.*, 1997; Haridi *et al.*, 1997) a multi-paradigm language, also including a logic programming component. Although thread mobility is not implemented in Distributed Oz 2, some of this functionality can be emulated in terms of network transparent mobile objects. The illusion of a unique application which transparently runs on multiple sites makes implementing shared multi-user applications particularly easy. We can emulate this by implementing mobile agents (e.g. avatars) as mobile threads with parts of the shared world *visible* to an agent represented as dynamic facts, lazily replicated through our lazy code fetching scheme when the agent moves.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

Both Distributed Oz 2 and our BinProlog based infrastructure need a full language processor (Oz 2 or BinProlog) to be deployed at each node. However, assuming that a Java processor is already installed, our framework's Java client (Tarau *et al*., 1997b; Tarau *et al*., 1997a) allows this functionality to be available through applets attached to a server side BinProlog thread.

Implementation technologies for mobile code are studied in Adl-Tabatabai *et al*. (1996). Early work on the Linda coordination framework (Carriero and Gelernter, 1989; Castellani and Ciancarini, 1996; Brogi and Ciancarini, 1991) has shown its potential for coordination of multi-agent systems. The logical modeling and planning aspects of computational Multi-Agent systems have been pioneered (Cohen and Perrault, 1979; Cohen *et al*., 1989; Kowalski and Kim, 1991; Wooldridge, 1992; Cohen and Cheyer, 1994; Cohen and Levesque, 1995; Lésperance *et al*., 1996; Chaib-draa and Levesque, 1994). A survey of Logic Programming approaches to Web applications in terms of a classification into client-based systems, server-side systems, and peer-to-peer systems is provided in Loke (1998).

Let us point out a key difference between our framework and a typical Java based mobile code system. Conventional mobile code systems like IBM's Aglets (IBM, 1999) require serialization hints from the programmer and do not implement a fully generic reflective computation mobility infrastructure. Aglets do not provide code mobility as they assume that code is already available at the destination site. In practice this means that the mobile code/mobile computation layer is not really transparent to the programmer.

In contrast, our architecture is based on building an autonomous layer consisting of a reflective interpreter which provides the equivalent of implicit serialization and supports orthogonal transport mechanisms for data, code and computation state. The key idea is simply that by introducing interpreters spawned as threads by a server at each networked site, *computation mobility* at object-level is mapped to *data mobility* at meta-level in a very straightforward way. A nice consequence is transport independence coming from the unifrom representation of data, code and computation state allowing Corba, RMI, HLA, as well as plain or multicast sockets to be used interchangeably as our transport mechanism.

6 Conclusion

We have described how mobile threads are implemented by capturing first order continuations in a data structure sent over the network. Supported by *lazy code fetching* and *dynamic recompilation*, they have been shown to be an effective framework for implementing mobile agents.

The techniques presented here are not (Bin)Prolog specific. The most obvious porting target of our design is to functional languages featuring first order continuations and threads. Another porting target is Java and similar object oriented languages having threads, reflection classes and remote method invocation. We are working on a Java based component using an embedded continuation passing Prolog interpreter which is already able to interoperate with BinProlog (latest version at `http://www.binnetcorp.com/Jinni`). An interesting application is

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

using BinProlog as an accelerator for Java based threads through migration to BinProlog, execution of a computationally intensive task and return to the Java component.

Acknowledgements

We thank the anonymous referees for their valuable suggestions and comments. We thank NSERC (grants OGP0107411 and 611024) for support.

Appendix I: Capturing first order continuations in BinProlog

```
% calls Goal with current continuation available to its inner calls
capture_cont_for(Goal):-
  assumeal(cont_marker(End)),
    Goal,
  end_cont(End).

% passes Closure to be called on accumulated continuation
call_with_cont(Closure):-
  assumed(cont_marker(End)),
  consume_cont(Closure,End).

% gathers in conjunction goals from the current continuation
% until Marker is reached when it calls Closure on it
consume_cont(Closure,Marker):-
  get_cont(Cont),
  consume_cont1(Marker,(_,_,_,Cs),Cont,NewCont), % first _
  call(Closure,Cs),                              % second _
  % sets current continuation to leftover NewCont
  call_cont(NewCont).                            % third _

% gathers goals in Gs until Marker is hit in continuation Cont
% when leftover LastCont continuation (stripped of Gs) is returned
consume_cont1(Marker,Gs,Cont,LastCont):-
   strip_cont(Cont,Goal,NextCont),
   ( NextCont==true-> !,errmes(in_consume_cont,expected_marker(Marker))
   ; arg(1,NextCont,X),Marker==X->
     Gs=Goal,arg(2,NextCont,LastCont)
   ; Gs=(Goal,OtherGs),
     consume_cont1(Marker,OtherGs,NextCont,LastCont)
   ).

% this 'binarized clause' gets the current continuation
get_cont(Cont,Cont)::-true(Cont).

% sets calls NewCont as continuation to be called next
call_cont(NewCont,_) ::- true(NewCont).
```

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

Appendix II: Thread mobility in BinProlog

```
% wraps continuation of current thread to be taken
% by inner move_thread goal to be executed remotely
wrap_thread(Goal):-
  capture_cont_for(Goal).

% picks up wrapped continuation,
% jumps to default remote site and runs it there
move_thread:-
  call_with_cont(move_with_cont).

% moves to remote site goals Gs in current continuation
move_with_cont(Gs):-
  % gets info about this host
  detect_host(BackHost),
  get_free_port(BackPort),
  default_password(BackPasswd),
  default_code(BackCode),
  % runs delayed remote command (assumes is with +/1)
  remote_run(
     +todo(
       host(BackHost)=>>port(BackPort)=>>code(BackCode)=>>(
         sleep(5), % waits until server on BackPort is up
         % runs foals Gs picked up from current continuation
         (Gs->true;true), % ignores failure
         % stops server back on site of origin
         stop_server(BackPasswd)
       )
     )
  ),
  % becomes data and code server for mobile code until is
  % stopped by mobile code possessing password
  server_port(BackPort)=>>run_unrestricted_server.
```

References

Adl-Tabatabai, A.-R., Langdale, G., Lucco, S. and Wahbe, R. (1996) Efficient and language-independent mobile programs. *Proceedings ACM SIGPLAN'96 Conference on Programming Language Design and Implementation* (*PLDI*), 127–136.

Almes, G. T., Black, A. P., Lazowska, E. D. and Noe, J. D. (1985) The Eden system: A technical review. *IEEE Trans. Softw. Eng.*, **11**(1), 43–59.

Bharat, K. A. and Cardelli, L. (1995) Migratory applications. *Proceedings 8th Annual ACM Symposium on User Interface Software and Technology*. `http://gatekeeper.dec.com/pub/DEC/SRC/research-reports/ abstracts/src-rr-138.html`.

Brogi, A. and Ciancarini, P. (1991) The concurrent language, shared Prolog. *Toplas*, **13**(1), 99–123.

Cardelli, L. (1997) Mobile computation. In: Vitek, J. and Tschudin, C. (eds.), *Mobile Object Systems – Towards the Programmable Internet*, pp. 3–6. Springer-Verlag.

Carriero, N. and Gelernter, D. (1989) Linda in context. *Comm. ACM*, **32**(4), 444–458.

Castellani, S. and Ciancarini, P. (1996) Enhancing coordination and modularity mechanisms

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

for a language with objects-as-multisets. In: Ciancarini, P. and Hankin, C. (eds.), *Proc. 1st Int. Conf. on Coordination Models and Languages*, pp. 89–106. Springer-Verlag.

Chaib-draa, B. and Levesque, P. (1994) Hierarchical models and communication in multi-agent environments. *MAAMAW'94*, pp. 119–134.

Ciancarini, P. (1994) Distributed programming with logic tuple spaces. *New Generation Comput.*, **12**(3), 251–284.

Cohen, P. R. and Cheyer, A. (1994) An open agent architecture. In: Etzioni, O. (ed.), *Software Agents — Papers from the 1994 Spring Symposium* (*Technical Report SS–94–03*), pp. 1–8. AAAIP.

Cohen, P. R. and Levesque, H. J. (1995) Communicative actions for artificial agents. *Proceedings First International Conference on Multi-agent Systems* (*ICMAS'95*), pp. 65–72.

Cohen, P. R. and Perrault, C. R. (1979) Elements of a plan based theory of speech acts. *Cogsci*, **3**, 177–212.

Cohen, P. R., Greenberg, M. L., Hart, D. M. and Howe, A. E. (1989) Trial by fire: Understanding the design requirements for agents in complex environments. *AI Mag.*, **10**(3), 32–48.

Dahl, V., Tarau, P. and Li, R. (1997) Assumption grammars for processing natural language. In: Naish, L. (ed.), *Proceedings 14th International Conference on Logic Programming*, pp. 256–270.

De Bosschere, K. and Tarau, P. (1996) Blackboard-based extensions in Prolog. *Software — Practice and Experience*, **26**(1), 49–69.

General Magic Inc. (1997) *Odissey*. Technical report, available at `http://www.genmagic.com/`.

Haridi, S., Van Roy, P. and Smolka, G. (1997) An overview of the design of Distributed Oz. *Proceedings 2nd International Symposium on Parallel Symbolic Computation* (*PASCO'97*), pp. 176–187. ACM Press.

IBM (1999) *Aglets*. Technical report `http://www.trl.ibm.co.jp/aglets`.

Jul, E., Levy, H., Hutchinson, N. and Black, A. (1988) Fine-grained mobility in the Emerald system. *ACM Trans. Comput. Syst.*, **6**(1), 109–133.

Kahn, R. E. and Cerf, V. G. (1988) *The Digital Library Project, Volume I: The World of Knowbots*. Technical report. Unpublished manuscript, Corporation for National Research Initiatives, Reston, VA.

Kowalski, R. and Kim, J.-S. (1991) A metalogic programming approach to multi-agent knowledge and belief. In: V. Lifschitz (ed.), *AI and Mathematical Theory of Computation: Papers in Honour of John McCarthy*. Academic Press.

Lésperance, Y., Levesque, H. J., Lin, F., Marcu, D., Reiter, R. and Scherl, R. B. (1996) Foundations of a logical approach to agent programming. *Atal'95*, pp. 331–346 Springer-Verlag.

Loke, S. W. (1998) *Adding Logic Programming Behaviour to the World Wide Web*. PhD thesis, University of Melbourne, Australia.

Sun Microsystems (1999) *JavaSpaces*. Technical report, available from `http://www.javasoft.com/products/javaspaces`.

Mikler, A. R., Unger, H., Tarau, P., Hopper, A. S. and Chen, F. (1999) A mobile agent-based file system for distributed networks. *High Performance Computing'99, Special Session on Adaptive and Intelligent Computing Systems.*

Miller, D. A. (1989) Lexical scoping as universal quantification. In: Levi, G. and Martelli, M. (eds.), *Proceedings Sixth International Conference on Logic Programming*, pp. 268–283. MIT Press.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

Mycroft, A. and O'Keefe, R. A. (1984) A polymorphic type system for Prolog. *Artif. Intell.*, 295–307.

P. Ciancarini, P., Knoche, A., Tolksdorf, R. and Vitali, F. (1996) PageSpace: an architecture to coordinate distributed applications on the Web. *Proceedings 5th International World Wide Web Conference*. Available from `http://www5conf.inria.fr/fich_html/papers/P5/Overview.html`.

Rochefort, S., Dahl, V. and Tarau, P. (1999) A framework for virtual learning environments. *Proceedings EMEDIA'99 World Conference on Educational Multimedia, Hypermedia and Telecommunications.*

Stamos, J. W. and Gifford, D. K. (1990) Remote evaluation. *ACM Trans. Programming Languages and Systems*, **12**(4), 537–565.

Tarau, P. and De Bosschere, K. (1993) Memoing with abstract answers and Delphi lemmas. In: Deville, Y. (ed.), *Logic Program Synthesis and Transformation*, pp. 196–209. Springer-Verlag.

Tarau, P. (1998a) *BinProlog 7.0 Professional Edition: Advanced BinProlog Programming and Extensions Guide*. Technical report, BinNet Corp. (Available from `http://www.binnetcorp.com/BinProlog`.)

Tarau, P. (1998b) *BinProlog 7.0 Professional Edition: User Guide*. Technical report, BinNet Corp. (Available from `http://www.binnetcorp.com/BinProlog`.)

Tarau, P. (1999a) A logic programming based software architecture for reactive intelligent mobile agents. In: Van Roy, P. and Tarau, P. (eds.), *Proceedings DIPLCL'99*. `http://www.binnetcorp.com/wshops/ICLP99DistInetWshop.html`.

Tarau, P. (1999b) Inference and computation mobility with Jinni. In: Apt, K. R., Marek, V. W. and Truszczynski, M. (eds.), *The Logic Programming Paradigm: a 25 Year Perspective*, pp. 33–48. Springer-Verlag.

Tarau, P. (1999c) Intelligent mobile agent programming at the intersection of Java and Prolog. *Proceedings 4th International Conference on The Practical Application of Intelligent Agents and Multi-Agents*, pp. 109–123.

Tarau, P. and Boyer, M. (1990) Elementary logic programs. In: P. Deransart and J. Maluszyński (eds.), *Proceedings of Programming Language Implementation and Logic Programming: Lecture Notes in Computer Science 456*, pp. 159–173. Springer-Verlag.

Tarau, P. and Dahl, V. (1996) A coordination logic for agent programming in virtual worlds. In: W. Conen and G. Neumann (eds.), *Proceedings Asian'96 Post-Conference Workshop on Coordination Technology for Collaborative Applications.*

Tarau, P., Dahl, V. and Fall, A. (1996a) Backtrackable state with linear affine implication and assumption grammars. In: Jaffar, J. and Yap, R. H.C. (eds.), *Concurrency and Parallelism, Programming, Networking, and Security: Lecture Notes in Computer Science 1179*, pp. 53–64. Springer-Verlag.

Tarau, P., De Bosschere, K. and Demoen, B. (1996b) Partial translation: Towards a portable and efficient Prolog implementation technology. *J. Logic Programming*, **29**(1–3), 65–83.

Tarau, P., Dahl, V. and De Bosschere, K. (1997a) A logic programming infrastructure for remote execution, mobile code and agents. *Proceedings of WETICE'97*, pp. 106–112.

Tarau, P., Dahl, V. and De Bosschere, K. (1997b) Logic programming tools for remote execution, mobile code and agents. *Proceedings ICLP'97 Workshop on Logic Programming and Multi Agent Systems.*

Van Roy, P., Haridi, S., Brand, P., Smolka, G., Mehl, M. and Scheidhauer, R. (1997) Mobile objects in Distributed Oz. *ACM Trans. Programming Languages and Systems*, **19**(5), 804–851.

Van Roy, P., Haridi, S. and Brand, P. (1997) *Using Mobility to Make Transparent Distribution Practical*. Technical report.

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press

Warren, D. H. D. (1981) Higher-order extensions to Prolog – are they needed? In: D. Michie, J. Hayes and Y. H. Pao (eds.), *Machine intelligence 10*. Ellis Horwood.

Wooldridge, M. (1992) *The logical modelling of computational multi-agent systems*. PhD thesis, Department of Computation, UMIST, Manchester, UK. (Also available as Technical Report MMU–DOC–94–01, Department of Computing, Manchester Metropolitan University, UK.)

https://doi.org/10.1017/S1471068401001193 Published online by Cambridge University Press