

Prolog-D-Linda : An Embedding of Linda in SICStus Prolog

Geoff Sutcliffe and James Pinakis

Dept of Computer Science, The University of Western Australia,
Nedlands, 6009, Western Australia
Email: geoff@cs.uwa.edu.au

1 Introduction

This paper presents an embedding of the Linda parallel programming paradigm in Prolog, resulting in a coarsely grained parallel Prolog. The work extends that reported in [Sutcliffe and Pinakis, 1990]¹. This embedding supports a distributed tuple space and a control hierarchy that provides remote I/O facilities for client processes. The embedding uses unification and Prolog style deduction in the tuple space. Two applications of Prolog-D-Linda are described.

2 The Linda paradigm

Linda is a programming framework of language-independent operators. These operators may be injected into the syntax of existing programming languages, such as Modula-II [Borrman, Herdieckerhoff, and Klein, 1988], C [Berndt, 1989], LISP [Yuen and Wong, 1990], Joyce [Pinakis and McDonald, 1991], and Russell [Butcher and Zedan, 1991], resulting in new parallel programming languages. Linda permits cooperation between parallel processes by controlling access to a shared data structure called the *tuple space*. The tuple space contains ordered collections of data called *tuples*. Manipulation of the tuple space is possible only by using the set of Linda operators.

2.1 Tuples and the Tuple Space

Tuples are collections of *fields*, of any arity. Every field has a data type drawn from the host language. The type of a tuple is the cross product of the types of its fields. A field can be a *formal* field or an *actual* field. A formal field has a type but no value, while an actual field has both a type and a value.

The tuple space contains any number of tuples, and identical tuples may exist in the tuple space. Processes communicate by inserting, removing and examining tuples in the tuple space. Thus the tuple space is a shared data object. All processes having access to the tuple space have access to all tuples in it.

¹ To make this paper self contained it has been necessary to duplicate a small portion of the information given in [Sutcliffe and Pinakis, 1990].

2.2 The Operators

The `out` operator inserts a tuple into the tuple space. The tuple is supplied as an argument to `out`.

The `in` operator removes a tuple from the tuple space. Its argument is a template against which tuples are matched. A template matches a tuple if all corresponding fields match. Two actual fields match if they have the same type and value. A formal field and an actual field match if they have the same type. Two formal fields cannot match. If a match for the template is found, the matched tuple is removed from the tuple space and formal fields in the template are given the values of the corresponding actual fields in the tuple. If more than one tuple matches a template, only one is chosen. If no matching tuple can be found in the tuple space, `in` will block and wait for a matching tuple to be inserted by an `out` operation.

The `rd` operator (pronounced read) is similar to `in`, but leaves the matched tuple in the tuple space. The `rd` operator is used for its binding and synchronization side-effects.

The `inp` and `rdp` operators perform tasks equivalent to `in` and `rd` but are non-blocking. Instead they return a boolean value which indicates the success of the operation. Leichter [1989] has argued against the use of these operators.

The final operator is `eval`. The `eval` operator is syntactically similar to `out`. When `eval` is called, a new process is created to evaluate each of the fields in the tuple argument. When the evaluation of all fields has terminated, the tuple is placed in the tuple space.

3 Prolog-D-Linda

Prolog-D-Linda represents tuples by Prolog clauses. Both Prolog rules and facts can exist in the tuple space. Facts correspond almost directly to standard Linda tuples. The necessity of a predicate symbol in a fact is analogous to requiring that the first field of a tuple be an actual field with a string literal value, as enforced by some Linda implementations [Leichter, 1989]. (This requirement does not reduce the generality of the system.) Formals in tuples are implemented by unbound variables. As data in Prolog is untyped (everything is a term) the data in Prolog-D-Linda's tuples is untyped. The effect of rules in the tuple space is discussed in section 5.

The Prolog-D-Linda tuple space is represented by a collection of Prolog databases. The tuples are distributed across the databases, as is described in section 4. Tuples are added to the tuple space (the `out` operation) using Prolog's `assertz` operation, and removed (the `in` operation) using `retract`. Tuples in the tuple space are examined (the `rd` operation) simply by using Prolog's query mechanism. The tuple matching method is thus generalised to Prolog's unification. As a consequence of this, formals can match and be extracted from the tuple space.

Prolog-D-Linda's `eval` operation differs from that of the original Linda paradigm. The `eval` operator is used to start a new Prolog environment, containing specified clauses and evaluating a specified Prolog query. The evaluation of the query may of course cause a tuple

to be inserted in the tuple space. This form of `eval` is more general than the original, and can implement the original.

4 The Implementation

Prolog-D-Linda has been implemented in SICStus Prolog [Carlsson, 1991]. SICStus Prolog provides an interface for accessing functions coded in C, and this is used for process creation and inter-process communication. Prolog-D-Linda runs on a network of Sun SPARC stations running SunOS 4.1.1, and connected via an Ethernet. This environment provides access to a shared file system via Sun's Network File System.

4.1 Overview

Prolog-D-Linda's tuple space and associated operations are implemented in *server* processes. Multiple servers can be used, each being responsible for part of the tuple space. Linda operations in *client* processes are translated into requests which are passed to an appropriate server. One of the servers is designated to be the *eval-server*. In addition to being responsible for part of the tuple space, the *eval-server* is responsible for processing all `eval` requests. Prolog-D-Linda is controlled by a single *controller* process, which must be associated with a terminal device. The controller is responsible for: (i) starting and stopping the server processes, (ii) for reading and displaying the terminal input and output of servers, and (iii) for reading and displaying the terminal input and output of clients that are started via an `eval` request.

Communication between the controller and servers, and between servers and clients, is via internet domain stream sockets. When a server is started, its terminal input and output streams are connected to a file descriptor in the controller. The descriptor is obtained from the `rexec()` system call used to start the server. Similarly, when a client is started by the *eval-server*, the client's terminal input and output streams are connected to a file descriptor in the *eval-server*. As well as the terminal I/O connection, every client establishes two further connections to each server. One connection is used for sending Linda operation requests to the server, and the other is used for receiving replies. The controller and servers monitor their input descriptors for incoming data, and process incoming data as described below.

The flow of requests, replies and I/O, via the sockets, is illustrated in figure 1.

4.2 The Configuration File

A configuration file must be supplied to Prolog-D-Linda. The configuration file is in the form of a Prolog program which specifies (i) the names of processors that will execute server processes, (ii) the name of the *eval-server*, and (iii) how the tuple space is to be partitioned amongst the servers. A sample configuration file is listed in the appendix. Prolog-D-Linda is started by executing the controller, which reads the configuration file to determine the names of the server processors. The names are stored in the `servers__/1` clause, whose argument is the list of processor names. When this clause has been found, the controller uses the `rexec()` system call to start each of the servers.

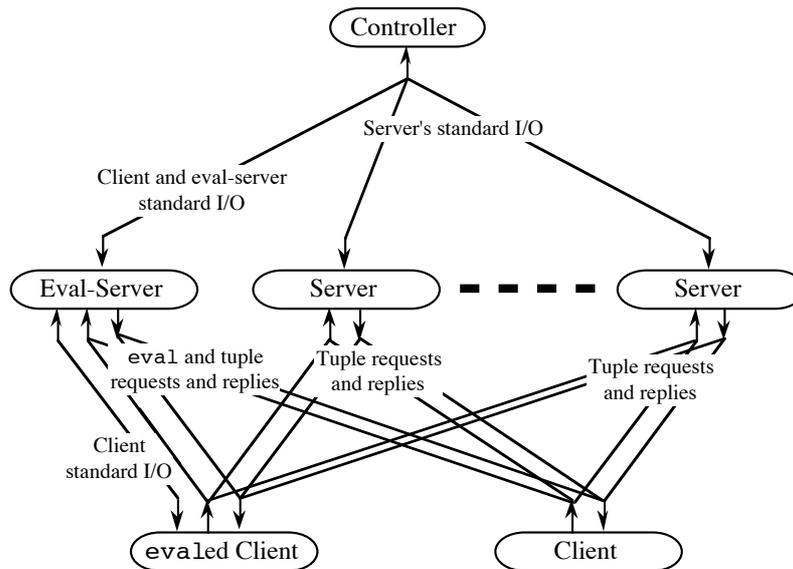


Figure 1 - The Prolog-D-Linda System

Client processes may be started independently at a terminal, or via the `eval` primitive. New clients consult the configuration file, and use the consulted clauses to determine how the tuple space is partitioned amongst the servers. The tuple space partition information is in the form of `select_server__/2` clauses. The first argument returns the name of the server processor that is responsible for the tuple that is supplied in the second argument. To determine where to send a tuple space operation request, a client simply evaluates an appropriate `select_server__/2` goal.

The eval-server starts new clients using an `rexec()` system call. The name of the eval-server is stored in the configuration file as the argument of the `eval_server__/1` clause.

4.3 The Operators

The servers read Linda operation requests from the request connections opened by the clients. The requests are serviced by evaluating them as Prolog queries. The requests are thus simply queries on Prolog procedures which implement the required operations. The use of Prolog evaluation to service requests is a general mechanism, and allows any query to be passed to a server for evaluation.

Tuple space `out` requests, received by servers, are implemented by `assertz`ing the supplied tuple into the server's database. The `in` requests attempt to `retract` the requested tuple template. The `rd` requests evaluate the requested tuple template as a Prolog goal. When an `in` or `rd` request is satisfied, the tuple template, with variables instantiated, is sent back to the requesting client on the reply connection. If a server is unable to satisfy an `in` or `rd` request, the request is placed on a global wait queue in the server, to wait for an appropriate tuple to be `outed`. After an `out` operation, all requests on the wait queue are re-evaluated.

The `inp` and `rdp` operations return the atom `fail`, to the requesting client, if the request cannot be immediately satisfied. This is used in the client to cause the operation to fail.

Prolog-D-Linda's `eval` operation takes three arguments : the name of a processor on which to execute the client, a Prolog goal, and a list of Prolog source file names. When the eval-server receives an `eval` request, it starts the new client by `reexecing` a SICStus Prolog saved state, on the specified processor. (The shared file system provides transparent access to files on remote processors.) The saved state has the client running so that it immediately consults the configuration file and opens the request and reply connections to the servers. In the interim, the eval-server places a tuple of the form `<client processor>(<the goal>,<the source files>)` into its tuple space. The client ins this tuple, consults the source files, and evaluates the goal. On completion of the goal the client closes the connections and terminates.

4.4 Terminal I/O

The terminal output of servers is read by the controller, off the descriptors obtained from its `reexec()` calls. The output is displayed on the controller's terminal, prefixed by the descriptor number from which it was read. This number uniquely identifies the server from which the output originated. Input to be sent to a server is entered at the controller's terminal, prefixed by the descriptor number which identifies the server (the descriptor number is obtained from previous output from that server). The controller strips the descriptor number from the input and forwards the remainder of the input to the server, on that descriptor.

The terminal output of clients that have been started via an `eval` request, is read by the eval-server off the descriptors obtained from its `reexec()` calls. The output is forwarded to the controller by writing the output to the eval-server's terminal output. The forwarded output is prefixed by the descriptor number from which it was read. The output is received by the controller and displayed on its terminal in the manner described above. Output from clients is therefore displayed with two prefixed descriptor numbers : firstly the descriptor number upon which it arrived at the controller (the descriptor number that identifies the eval-server), and secondly the descriptor number upon which it arrived at the eval-server (the descriptor number that identifies the client). These numbers uniquely identify the client from which the output originated. Input to be sent to a client is entered at the controller's terminal, prefixed by the two descriptor numbers which identify the client (the descriptor numbers are obtained from previous output from that client). The controller strips the first descriptor number from the input and forwards the remainder of the input to the eval-server, on that descriptor. The eval-server strips the second descriptor number from its input and forwards the remainder to the client, on that descriptor.

This I/O hierarchy permits all clients to be interactive, even though they may not be associated with a terminal device. If there is to be a lot of such client I/O, then the system can be configured so that the eval-server is not responsible for any part of the tuple space, i.e. the eval-server only deals with `eval` requests and client I/O. The I/O hierarchy is also used to halt the system. If the keyword `halt.` is entered at the controller's terminal, `halt.` is written to all servers' terminal inputs. When a server reads the `halt.` message, it

immediately closes all open descriptors and terminates. After writing `halt` to all the servers, the controller closes all its open descriptors and terminates.

5 Deductive tuple spaces

The Linda tuple space and associated operations are very similar to a standard concurrent access relational database system. The `in` and `out` operations effect database updates, and the `rd` operation effects database queries. The difference is that the Linda paradigm is viewed as providing communication between, and synchronization of, parallel processes, whereas a relational database is viewed only as storing data. Much research has been done on the generalisation of relational database to deductive database, in Prolog. Lloyd [1987] gives a good summary of this work. It is a logical step to extend the Prolog-Linda tuple space to a deductive tuple space. By allowing rules as well as facts to be added to and removed from the tuple space, the tuple space becomes deductive. Tuple space `rd` and `rdp` requests may be satisfied by facts, or using rules. Rules are evaluated using normal Prolog deduction, including backtracking. If a deductive tuple space is used it is necessary for all the required tuples (rules and facts) to be stored in the same partition of the tuple space, i.e. in one server.

A deductive tuple space greatly increases the capabilities of the tuple space, but not without some penalty. The first problem is the increased time required to evaluate deductive `rds`. The second problem, which is an extreme case of the first, is the danger of a server entering an infinite deduction. Client requests to that server will not be evaluated. In particular requests that could terminate the infinite deduction are not evaluated. A solution to this second problem is to restrict the nature of the deductive database to be hierarchical [Lloyd, 1987]. Despite the problems associated with a deductive tuple space, such a model provides facilities that are not available from a standard tuple space. Two examples are described here.

∑ In Linda it is awkward to simultaneously rd tuples of two different signatures. A method suggested in [Leichter, 1989] requires the outing process to know that the tuples will be requested in this way. A deductive tuple space provides a direct solution :

```
rd_either(Tuple1,Tuple2):-
    out((switch(Tuple1):-Tuple1)),
    out((switch(Tuple2):-Tuple2)),
    %----Wait for Tuple1 or Tuple2 to be outed
    rd(switch(Which)).
    %----Which contains the outed tuple
```

∑ A deductive tuple space has the potential for extreme space saving. There are indeed some groups of tuples that can only be finitely stored in a deductive manner. For example :

```
out_even:-
    out((even(Negative):-Negative < 0,!,fail)),
    out(even(0)),
    out((even(Number):-Number_less_2 is Number-2,
    even(Number_less_2))).
```

would effectively place all tuples even(X) into the tuple space, where X is an even natural number.

6 Applications

Two non-trivial applications have been developed using Prolog-D-Linda.

6.1 Automated deduction

Prolog-D-Linda has been used to implement two distributed automated deduction systems [Sutcliffe, 1991, 1992]. The deduction systems have multiple deduction components, which execute as separate client processes. Each component runs a different format of deduction system. Lemmas created in each of the deduction components are passed to the other components, via the tuple space. An extended version of one of the systems, in which the lemmas created are distributed via a separate 'lemma control' component, has also been developed. The speed-ups obtained by these distributed deduction systems are largely due to cross-fertilisation between the deduction components. The implementation of the systems is highly modular, and new deduction or other components can easily be added.

6.2 Genetic Algorithms

Prolog-D-Linda has been used to implement a genetic algorithm, in which multiple clients access and update the solution pool in parallel. Each candidate solution is stored as a tuple containing the solution and its objective value. Each client process repeatedly (i) rds two parent solution tuples from the tuple space, (ii) performs a crossover to produce two child solutions, (iii) for each child, ins a 'sucker' solution chosen at random, (iv) outs the child solution if $e^{(SuckerObjective - ChildObjective)/T} > \text{random}([0,1])$, otherwise outs the sucker. (I.e.

if the child has a better objective value than the sucker, then the child is always outed; if the child has a worse objective value, then the child may still be saved by virtue of the Boltzman distribution, with temperature T.) Some variants of this algorithm have also been implemented. It is the iterative nature of this genetic algorithm that permits it to be parallelised. Similar work has been done by Ackley [1987] and Robertson [1987].

7 Conclusion

Prolog-D-Linda is a truly distributed logic programming environment. This distribution allows applications to take advantage of the added computing power available, as well as to be structured in a parallel fashion. The parallelism obtained is acknowledged to be coarse. In the context of parallel Prolog architectures, it has been argued that "exploiting as much fine grain parallelism as possible may be a flawed strategy; any gains through increased parallelism are wasted due to communication overheads" [Wise, 1991, p 2].

The distribution of the tuple space in Prolog-D-Linda makes it superior to its predecessors. As the partitioning is user controlled, it is possible to tune the use of the tuple space so that bottlenecks are avoided. The introduction of a deductive tuple space is a significant enhancement to the capabilities of the Linda paradigm. A deductive tuple space provides direct solutions to problems that were previously difficult or impossible.

The Prolog-D-Linda embedding of Linda in Prolog is very natural : the pattern matching and database features of Prolog have been used directly in the embedding; garbage collection and hashing in the tuple space are provided free by the Prolog implementation; the implementation of formals in tuples is direct; the specification of how the tuple space is to be partitioned is done as a Prolog program. This naturalness contrasts with the FCP(≠) implementation described by Shapiro [1989]. The closest relations to Prolog-D-Linda appear to be Shared Prolog [Brogi and Ciancarini, 1991] and its successor, PoLiS Prolog [Ciancarini, 1992]. PoLiS Prolog extends upon Prolog-D-Linda, in that it provides multiple tuple spaces. In both Shared Prolog and PoLiS Prolog the tuple spaces are non-deductive, and the manner in which the tuple space operators can be used is restricted.

8 References

- Ackley D.H. (1987), *A Connectionist Machine for Genetic Hillclimbing*, Kluwer Academic Publishers, Dordrecht, The Netherlands.
- Berndt D.J. (1989), *C-Linda Reference Manual, Version 2.0*, Scientific Computing Associates Inc., New Haven, CT.
- Borrmann L., Herdieckerhoff M., and Klein A. (1988), Tuple Space Integrated into Modula-2, Implementation of the Linda Concept on a Hierarchical Multiprocessor, In Jesshope, Reinartz (Ed.), *Proceedings of CONPAR '88* (NEED, 1988), Cambridge University Press, Cambridge, England, 1-8.

- Brogi A., and Ciancarini P. (1991), The Concurrent Language, Shared Prolog, In *ACM Transactions on Programming Languages and Systems* 13(1), ACM Press, New York, NY, 99-123.
- Butcher P., and Zedan H. (1992), Lucinda - A Polymorphic Linda, In Banatre J.B., Le Metayer D. (Ed.), *Proceedings of the Workshop on Research Directions in High-Level Parallel Programming Languages* (Mont Saint-Michel, France, 1991), (Lecture Notes in Computer Science 574), Springer-Verlag, New York, NY, 110-125.
- Carlsson M., and Widen J. (1991), SICStus Prolog User's Manual, T91: 11B, Swedish Institute of Computer Science, Kista, Sweden.
- Ciancarini P. (1992), Parallel logic programming using the Linda model of computation, In Banatre J.B., Le Metayer D. (Ed.), *Proceedings of the Workshop on Research Directions in High-Level Parallel Programming Languages* (Mont Saint-Michel, France, 1991), (Lecture Notes in Computer Science 574), Springer-Verlag, New York, NY, 110-125.
- Leichter J.S. (1989), Shared Tuple Memories, Shared Memories, Buses and LAN's - Linda Implementations Across the Spectrum of Connectivity, Ph.D. Thesis, Yale University, Yale, CT.
- Lloyd J.W. (1987), *Foundations of Logic Programming, 2nd Edition*, Springer-Verlag, New York, NY.
- Pinakis J., and McDonald C. (1991), The Inclusion of the Linda Tuple Space Operations in a Pascal-based Concurrent Language, In Gupta G., Lions J. (Ed.), *Proceedings of the 14th Australian Computer Science Conference* (Kensington, Australia, 1991), Department of Computer Science, University of New South Wales, Kensington, Australia, 45.1-45.11.
- Robertson G. (1987), Parallel Implementation of Genetic Algorithms in a Classifier System, In Davis L. (Ed.), *Genetic Algorithms and Simulated Annealing*, (Research Notes in Artificial Intelligence Pitman Publishing, London, England, 129-140.
- Shapiro E. et al. (1989), Technical Correspondence, In *Communications of the ACM* 32(10), ACM Press, New York, NY, 1244-1258.
- Sutcliffe G. (1991), A Parallel Linear and UR-Derivation System, In Kanal L.N., Suttner C. B. (Ed.), *Informal Proceedings of PPAI-91, International Workshop on Parallel Processing for Artificial Intelligence* (Sydney, Australia, 1991), International Joint Conferences on Artificial Intelligence, Inc., Sydney, Australia, 211-215.
- Sutcliffe G. (1992), A Heterogeneous Parallel Deduction System, In Hasegawa R., Stickel M. E. (Ed.), *Proceedings of the Workshop on Automated Deduction: Logic Programming and Parallel Computing Approaches, FGCS'92* (Tokyo, Japan, 1992), Institute for New Generation Computer Technology, Tokyo, Japan,
- Sutcliffe G., and Pinakis J. (1990), Prolog-Linda - An Embedding of Linda in muProlog, In Tsang C.P. (Ed.), *Proceedings of AI'90 - the 4th Australian Conference on Artificial Intelligence* (Perth, Australia, 1990), World Scientific, Singapore, 331-340.

Wise M.J. (1991), MB-Prolog: A Distributed Prolog with Communication via Message-Brokers, Wise M.J., Distributed at Parallel Processing for AI Workshop, in IJCAI'91.

Yuen C.K., and Wong W.F. (1990), BaLinda Lisp: A Parallel Lisp Dialect for Biddle with the Concurrent Facilities of Linda, Technical Report TRA1/90, Department of Information Systems and Computer Science, National University of Singapore, Kent Ridge, Singapore.

Appendix

Listed below is a sample Prolog-D-linda configuration file. The configuration specifies two servers - `bison` and `budgie`, of which `bison` is nominated as the eval-server. The tuple space is partitioned so that `bison` maintains tuples of arity 0 and 1, and `budgie` maintains all other tuples.

```
/*----The tuple space is partitioned between two processors */
servers__([bison,budgie]).
```

```
/*----bison does the eval requests */
eval_server(bison).
```

```
/*----bison maintains tuple with 1 or 0 arguments */
select_server__(bison,Tuple):-
    functor(Tuple,_,Arity),
    Arity =< 1.
```

```
/*----budgie maintains all other tuples */
select_server__(budgie,Tuple):-
    functor(Tuple,_,Arity),
    Arity >= 2.
```