Efficient Logic Variables for Distributed Computing

Seif Haridi[*], Peter Van Roy[†], Per Brand[‡], Michael Mehl[§], Ralf Scheidhauer[¶] and Gert Smolka[||]

September 15, 1998

Abstract

We define a practical algorithm for distributed rational tree unification and prove its correctness in both the off-line and on-line cases. We derive the distributed algorithm from a centralized one, showing clearly the trade-offs between local and distributed execution. The algorithm is used to implement logic variables in the Distributed Oz system. Oz has a very general execution model that consists of concurrent constraints extended with explicit state and higher-orderness. We show that logic variables can easily be added to the more restricted models of Java and ML. We present common distributed programming idioms in a network-transparent way using logic variables. We show that in common cases the algorithm maintains the same message latency as explicit message passing. In addition, it is able to handle uncommon cases that arise from the properties of latency tolerance and third-party independence. This is evidence that using logic variables in distributed computing is beneficial at both the system and language levels. At the system level, they improve latency tolerance and third-party independence. At the language level, they help make network-transparent distribution practical.

1 Introduction

Logic variables were first studied in the context of logic programming. They remain an essential part of logic programming and constraint programming systems [36, 18]. In the context of the Distributed Oz project [14, 31], we have come to realize their usefulness to distribution. Logic variables express dependencies between computations without imposing an execution order. This property can be exploited in distributed computing:

- Two basic problems in distributed computing are latency tolerance and third-party independence. We show that using logic variables instead of explicit message passing can improve these two aspects of an application with little programming effort.
- Using logic variables we can express common distributed programming idioms in a network-transparent manner that results in optimal or near-optimal message latency.

The main contribution of this article is a practical distributed algorithm for rational tree unification which realizes these benefits. The algorithm efficiently implements logic variables in the

[*]`seif@sics.se`, Swedish Institute of Computer Science, S-164 28 Kista, Sweden
[†]`pvr@info.ucl.ac.be`, Dép. INGI, Université catholique de Louvain, B-1348 Louvain-la-Neuve, Belgium
[‡]`perbrand@sics.se`, Swedish Institute of Computer Science, S-164 28 Kista, Sweden
[§]`mehl@dfki.de`, German Research Center for Artificial Intelligence (DFKI), D-66123 Saarbrücken
[¶]`scheidhr@dfki.de`, German Research Center for Artificial Intelligence (DFKI), D-66123 Saarbrücken
[||]`smolka@ps.uni-sb.de`, DFKI and Universität des Saarlandes, D-66123 Saarbrücken, Germany

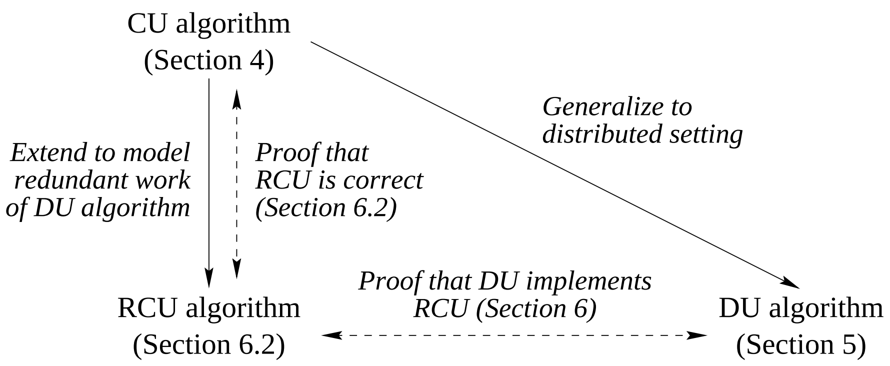

Figure 1: Defining the algorithm and proving it correct

Distributed Oz system. We formally define the algorithm and prove that it satisfies safety and liveness properties in both the off-line and on-line cases.

From the programmer's point of view, the use of logic variables adds a dataflow component to program execution. In a first approximation, this component can be completely ignored. That is, it is invisible to the programmer whether or not a thread temporarily blocks while waiting for a variable's value to arrive. Programs can be developed using well-known techniques of concurrent object-oriented programming [20]. In a second approximation, the dataflow component greatly simplifies many concurrent programming tasks [12, 5].

This article consists of two main parts which may be read independently of each other. The first part, Section 2, motivates and discusses in depth the use of logic variables in concurrent and distributed programming. Section 2.1 introduces a general execution model, its distributed extension, and the concept of logic variable. Section 2.2 gives the key ideas of the distributed unification algorithm. Section 2.3 shows how to express basic concepts in concurrent programming using logic variables. Section 2.4 expresses common distributed programming idioms in a network-transparent manner with logic variables. We show that the algorithm provides good network behavior for these examples. Finally, Section 2.5 shows how to add logic variables in an orthogonal way to Java and ML, taken as representative examples of object-oriented and functional languages.

The second part, Section 3 and following, defines the distributed unification algorithm, proves its correctness, and discusses its implementation (see Figure 1). Section 3 defines the formal representation of logic variables and data structures. This section also defines configurations and executions and introduces the reduction rule notation used to define algorithms. Section 4 defines the CU algorithm, which implements off-line centralized unification, and summarizes well-known results about its correctness. By *off-line* we mean that the set of equations is finite and initially known. Section 5 defines the DU algorithm, which implements off-line distributed unification. Section 6 defines the RCU algorithm, which modifies the centralized algorithm to reflect the redundant work done by the DU algorithm. The section then proves that the DU algorithm is a correct implementation of the CU and RCU algorithms. Section 7 defines on-line versions of the CU and DU algorithms. By *on-line* we mean that new equations can nondeterministically be introduced at any moment. We define the *finite size* property, and prove that given weak fairness, every introduced equation that satisfies this property is eventually entailed by the store for both algorithms. Section 8 defines the algorithm used in the Distributed Oz implementation, which is an optimized refinement of the on-line DU algorithm.

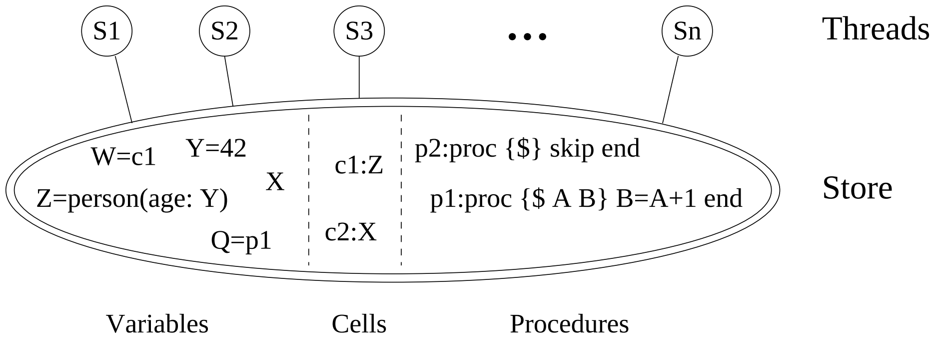

Figure 2: The Oz execution model

2 Logic variables in concurrent and distributed settings

This section motivates our unification algorithm by showing its usefulness to distributed programming. First Section 2.1 introduces our execution model and its notation. Then Section 2.2 gives the key ideas of the algorithm. This is followed by Section 2.3 which gives programming examples showing the usefulness of logic variables for basic tasks in concurrent programming. Section 2.4 continues with tasks in distributed programming. We explain in detail the network behavior of our algorithm for these tasks. Finally, Section 2.5 shows how to add logic variables to other languages including the Java and ML families.

2.1 Basic concepts and notation

As a framework for logic variables, we introduce a general execution model that can accommodate most programming languages. Underlying it is a formal model called *concurrent constraints* [29, 27] that contains logic variables as a basic ingredient. Some uses of logic variables, e.g., synchronization and communication, already appear in this model. The general execution model, called Oz execution model, extends the concurrent constraint model with explicit state and higher-orderness. Other, novel uses of logic variables, e.g., locks, are possible only when state is introduced.

This section gives the essential characteristics of the Oz execution model and how it is distributed. Later on, we show how to add logic variables to the more restricted models of Java and ML. The advantage of using a general formal model is that it allows us to define precisely what the examples do. It is straightforward to compile Java or ML to Oz; the converse is not as easy.

2.1.1 The Oz execution model

The Oz execution model consists of dataflow threads observing a shared store (see Figure 2). Threads contain statement sequences S_i and communicate through shared references into the store. A thread is *dataflow* if it only executes its next statement when all the values the statement needs are available. If the statement needs a value that is not yet available, then the thread automatically blocks until it can access that value. We add the fairness condition that every thread that is not blocked eventually executes its next statement. As we shall see, data availability in the Oz model is implemented using logic variables.

The shared store is not physical memory, rather it is an abstract store that only allows legal operations for the entities involved, i.e., there is no direct way to inspect their internal representations. The store consists of three compartments, namely logic variables (with binding,

$$
\begin{array}{rlr}
S ::= & S\ S & \text{Sequence} \\
\mid & X\texttt{=f(}l_1\texttt{:}Y_1\ \texttt{...}\ l_n\texttt{:}Y_n\texttt{)}\ \mid & \text{Value} \\
 & X\texttt{=<number>}\ \mid\ X\texttt{=<atom>}\ \mid\ \texttt{\{NewName}\ X\texttt{\}} & \\
\mid & \textbf{local}\ X_1\ \texttt{...}\ X_n\ \textbf{in}\ S\ \textbf{end}\ \mid\ X\texttt{=}Y & \text{Variable} \\
\mid & \textbf{proc}\ \texttt{\{}X\ Y_1\ \texttt{...}\ Y_n\texttt{\}}\ S\ \textbf{end}\ \mid\ \texttt{\{}X\ Y_1\ \texttt{...}\ Y_n\texttt{\}} & \text{Procedure} \\
\mid & \texttt{\{NewCell}\ Y\ X\texttt{\}}\ \mid\ \texttt{\{Exchange}\ X\ Y\ Z\texttt{\}}\ \mid\ \texttt{\{Access}\ X\ Y\texttt{\}} & \text{State} \\
\mid & \textbf{if}\ X\ \textbf{then}\ S\ \textbf{else}\ S\ \textbf{end} & \text{Conditional} \\
\mid & \textbf{thread}\ S\ \textbf{end}\ \mid\ \texttt{\{GetThreadId}\ X\texttt{\}} & \text{Thread} \\
\mid & \textbf{try}\ S\ \textbf{catch}\ X\ \textbf{then}\ S\ \textbf{end}\ \mid\ \textbf{raise}\ X\ \textbf{end} & \text{Exception}
\end{array}
$$

Figure 3: The Oz kernel language

if bound), cells (named mutable pointers, i.e., explicit state), and procedures (named lexically-scoped closures, i.e., higher-orderness). Variables can reference the names of procedures and cells. Cells point to variables. The external references of threads and procedures are variables. When a variable is bound, it disappears, that is, all threads that reference it will automatically reference the binding instead. Variables can be bound to any entity, including other variables. The variable and procedure stores are monotonic, i.e., information can only be added to them, not changed or removed. Because of monotonicity, a thread that is not blocked is guaranteed to stay not blocked until it executes its next statement.

2.1.2 The Oz language

All Oz execution can be defined in terms of a kernel language whose semantics are outlined in [12, 38]. The current Oz language is called Oz 2 to distinguish it from an earlier language, Oz 1, whose kernel language is called Oz Programming Model (OPM) [29]. We do not discuss Oz 1 further in this paper. Figure 3 defines the abstract syntax of a statement S in the Oz kernel language. Statement sequences are reduced sequentially inside a thread. All variables are logic variables, declared in an explicit scope defined by the **`local`** statement. Values (records, numbers, etc.) are introduced explicitly and can be equated to variables. Procedures are defined at run-time with the **`proc`** statement and referred to by a variable. Procedure applications block until the first argument references a procedure name. State is created explicitly by `NewCell`, which creates a *cell*, an updatable pointer into the variable store. Cells are updated by `Exchange` and read by `Access`. The **`if`** statement defines a conditional that blocks until its condition is true or false in the variable store. Threads are created explicitly with the **`thread`** statement and have their own identifier. Exception handling is dynamically scoped. The **`try`** statement defines a scope and the **`raise`** statement raises an exception that is caught by the innermost enclosing scope.

The full Oz language is defined by transforming all its statements into this kernel language. Oz supports idioms such as objects, classes, reentrant locks, and a variety of channel called "ports" [29, 38]. The system implements them efficiently while respecting their definitions. We give a brief summary of each idiom's definition. For clarity, we have made small conceptual simplifications. Full definitions are given in [12].

- **Object**. An object is essentially a one-argument procedure `{Obj M}` that references a cell, which is hidden by lexical scoping. The cell holds the object's state. The argument `M` indexes into the method table. A method is a procedure that is given the message and the object state, and calculates the new state.

- **Class**. A class is essentially a record that contains the method table and attribute names. A class is defined through multiple inheritance, and any conflicts are resolved at definition time when building its method table.

- **Reentrant lock**. A reentrant lock is essentially a one-argument procedure `{Lck P}` used for explicit mutual exclusion, e.g., of method bodies in objects used concurrently. `P` is a zero-argument procedure defining a critical section. Reentrant means that the same thread is allowed to reenter the lock. Calls to the lock may therefore be nested. The lock is released automatically if the thread in the body terminates or raises an exception that escapes the lock body.

- **Port**. A port is an asynchronous channel that supports many-to-one communication. A port `P` encapsulates a stream `S`. A *stream* is a list with unbound tail. The operation `{Send P M}` adds `M` to the end of `S`. Successive sends from the same thread appear in the order they were sent.

2.1.3 The distribution model

Distributed Oz is a conservative extension to the centralized Oz language [7] that completely separates functionality from distribution structure. That is, Oz language semantics are unchanged[1] while adding predictable and programmable control over network communication patterns. Porting existing Oz programs to Distributed Oz requires essentially no effort.

Allowing a successful separation of functionality from distribution structure puts severe restrictions on a language. It would be almost impossible in C++ because the semantics are informal and unnecessarily complex and because the programmer has full access to all underlying representations [32]. It is possible in Oz because of four properties:

- Oz has a simple formal foundation that does not sacrifice expressiveness or efficient implementation. Oz appears to the programmer as a concurrent object-oriented language whose basic functionality is comparable to modern languages such as Java. The current emulator-based implementation is competitive with Java emulators [16, 15]. Standard techniques for concurrent object-oriented design apply to Oz [20]. Furthermore, Oz introduces powerful new techniques that are not supported by Java [12]. Some of these techniques are presented here.

- Oz is both a state-aware and dataflow language. That is, language entities can naturally be classified into stateless, single assignment, and stateful. This helps give the programmer control over network communication patterns in a natural manner. Stateless data includes procedures and values, which can safely be copied to many sites [3]. Stateful data includes objects, which at any instant must reside on just one site [38]. Dataflow synchronization allows to decouple calculating a value from sending it across the network.

- Oz is a fully dynamic and compositional language. That is, Oz is dynamically typed and all entities are first-class. This makes it easy to implement fully open distribution, in which independent computations can connect and disconnect at will. When connected they can communicate as if they were in the same centralized process. For example, it is possible to define a class `C` in one computation, pass `C` to an *independent* computation that has never before heard of `C`, let the independent computation define a class `D` inheriting from `C`, and pass `D` back to the original computation.

- Oz provides language security. That is, all references to language entities are created and passed explicitly. An application cannot forge references nor access references that have not been explicitly given to it. The underlying representation of language entities is inaccessible to the programmer. This is a consequence of the abstract store and a kernel language with lexical scoping and first-class procedures. These are essential properties to implement a capability-based security policy, which is important in open distribution.

[1] Only ports are changed slightly to better model asynchronous FIFO communication between sites [38].

The Distributed Oz execution model extends the Oz execution model by giving a distributed semantics to each language entity. The distributed semantics defines the network behavior when language entities are shared between sites. The semantics is chosen carefully to give predictable control over network communication patterns. The centralized semantics is unchanged: we say the model is *network-transparent*. In the current system, language entities are put in four categories. Each category is implemented by a distributed protocol:

- **Stateless**: records, numbers, procedures, and classes. Since they do not change, these entities can be copied at will. There is a trade-off between when to copy, how many times to copy per site, and access speed [3].
- **Single assignment**: logic variables. Assignment is done by a distributed unification algorithm, which is the subject of this paper. To be precise, logic variables provide *consistent multiple assignment*, that is, assignments must be consistent with each other. We keep the phrase "single assignment" to avoid multiplying terminology.
- **Stateful**: cells, objects, ports, reentrant locks, and threads. For efficiency reasons, these entities' state pointers are localized to a site. If the state pointer's site can change, we say that the entity is mobile. Currently there are two mobility behaviors: a mobile state protocol (cells, objects, ports, locks) and a stationary access protocol (threads). The mobile state protocol ensures coherent state updates with controlled mobility of the state pointer [38]. The stationary access protocol is used for entities that cannot move.
- **Resource**: entities external to the language. These may be executed only on the site at which they were created. External entities are defined by the operating system in which the language executes. This includes computational and memory resources, which can be made visible in the language, e.g., as *virtual sites* [13].

The single-assignment category can be seen as an optimization of the stateful category in which a variable is bound to only one value, instead of repeatedly to different values. That is, the distributed unification algorithm is more efficient than the mobile state protocol. However, it turns out that logic variables have many more uses than simply as an optimization of stateful entities. These uses are explained below.

2.1.4 Logic variables

A logic variable conceptually has a fixed value from the moment of its creation. The value is unknown at first, and it remains unknown until the variable is bound. At all times, the variable can be used as if it were the value. If the value is not needed then execution continues. If the value is needed, then the thread requiring the value will suspend until the variable is bound.

A logic variable can be passed among sites arbitrarily. At all times, it "remembers its origins", that is, when the value becomes known then the variable will receive it. The communication needed to bind the variable is part of the variable, and not part of the program manipulating the variable. This means that the variable can be passed around at will, and the value will always arrive at the variable. This is one key reason why logic variables are useful in distributed computing.

Logic variables can replace standard (assignable) variables in all cases where they are assigned only one value, i.e., where they are used as placeholders for values. The algorithm used to bind logic variables must ensure that the result is independent of binding order. In a centralized system, the algorithm is called *unification* and is usually implemented as an extension of a union-find algorithm. Union-find handles only the binding of variables with variables [23]. Unification generalizes this to handle nonvariable terms as well. In a good implementation, binding a new variable to a nonvariable (the common case) compiles to a single register move or store operation [36].

A logic variable may be bound to another logic variable. A legitimate question is whether variable-variable binding is useful in a practical system. As we shall see, one reason that

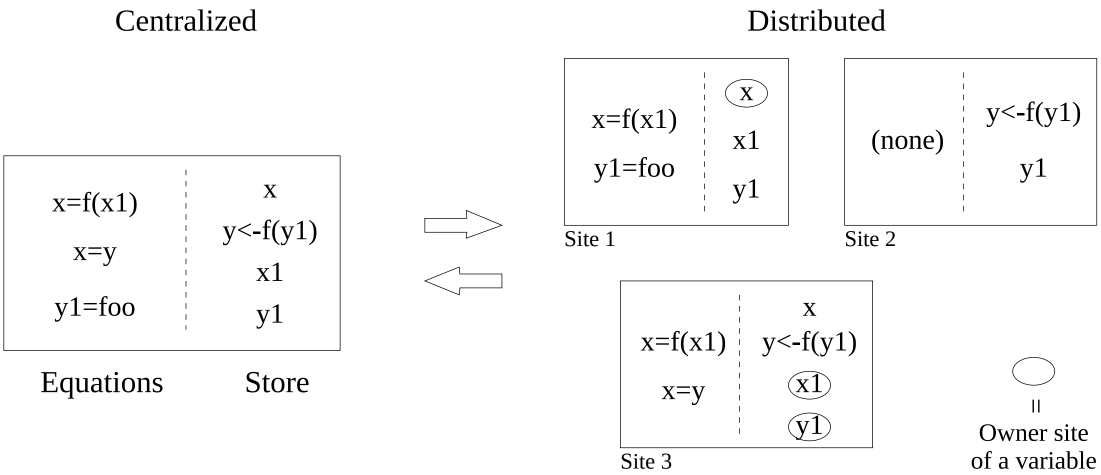

Figure 4: Initial configuration of example

variable-variable binding is important is that it allows to maintain maximum latency tolerance and third-party independence when communicating among more than two sites, independent of fluctuating message delays. A second reason is that it has a very simple logical semantics.

It is possible to disallow variable-variable binding to obtain a simpler implementation. The simpler implementation blocks any attempt to do variable-variable binding until at least one of the variables is bound to a value. The price of the simpler implementation is that third-party dependencies are not removed in all cases. Futures [11] and I-structures [17] resemble this weaker version of logic variables. There remains a crucial difference with logic variables, namely that futures and I-structures can be assigned only once, whereas logic variables can be assigned more than once, as long as the assignments are consistent with each other. The efficiency difference between full and weak logic variables is small when variable-variable binding is not done. The distributed binding algorithm is almost identical for the full and weak versions. Furthermore, the full version has a simple logical semantics. For these three reasons we have implemented the full version in the Distributed Oz implementation.

2.2 Distributed unification

For logic variables to be practical in a distributed setting they must be efficiently implementable. This section gives the key ideas of the distributed algorithm that realizes this goal. We explain the algorithm in just enough detail so that its network behavior becomes clear. This will allow us to infer the algorithm's behavior in the programming examples that follow. A complete definition of the algorithm and proofs of its correct behavior are given starting from Section 3.

The two basic operations on logic variables are binding and waiting until bound. Waiting until bound is easy: the variable has a list containing suspended threads that need its value. When the value arrives, the threads are awoken. Binding is harder: it requires cooperation between sites. If a variable exists on several sites, then it must be bound to the same value on all these sites, despite concurrent binding attempts. Unification implements the binding operation. At any instant there can be any number of bindings in various stages of completion. Both the centralized and distributed algorithms cause each binding request to be eventually incorporated into the store, if it is consistent with the store.

The basic distributed operation is binding a variable to a value. This is implemented by making one site the "owner" of the variable. The site that declares the variable is its owner. A binding request is sent to the owner, and the owner forwards the binding to each site that

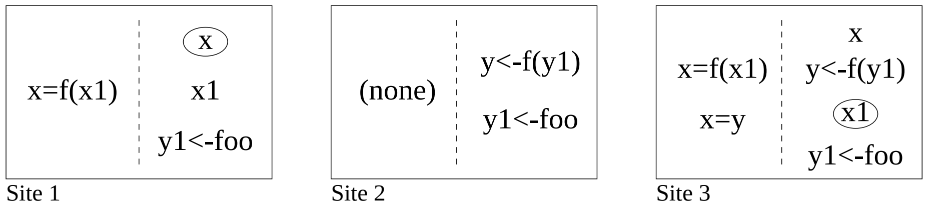

Figure 5: Configuration after executing the equation Y1=foo

knows the variable. In terms of network behavior, one message is sent to the owner, and one message is sent by the owner to each site that knows the variable. The owner's sends can be done by a reliable multicast, if the network supports it efficiently. The owner accepts the first binding request and ignores all subsequent binding requests. An ignored request will be retried by its initiating site after it receives the binding. As we will see in the programming examples, in the majority of cases a variable is declared either on a site that will need its value or on the site that will bind the variable. In both of these cases, the network behavior of the algorithm is very good.

A logic variable X can be bound to a data structure or to another variable. The algorithm is the same in both cases. By default the binding is *eager*, i.e., the new value is immediately sent to all sites that know about X. This means that a bound variable is guaranteed to eventually disappear from the system. The same binding eventually appears on each site that has the variable. For example, executing the equation X=f(X1) causes the binding X $\leftarrow$ f(X1) to appear in the store on all sites containing X. Furthermore, X1 is added to the known variables of all of these sites that did not know X1.

2.2.1 An example

We illustrate the algorithm with an example. Figure 4 shows a centralized configuration on the left and on the right one possible way it can be distributed. Each configuration has a set of equations and a store. Equations are requests to perform bindings. The store contains the variables and their bindings, if the latter exist.

In the distributed case, each site has a set of equations and a store. A variable can exist on more than one site. One of the sites is the variable's owner. In Figure 4, site 1 is the owner of X and site 3 is the owner of both X1 and Y1. Variable Y is already bound to f(Y1) on sites 2 and 3. If there is only one site, the distributed algorithm is identical to the centralized algorithm. Each occurrence on a site is called a "proxy". When the variable is bound then each proxy will eventually see the same binding.

Site 1 requests the binding Y1=foo. This sends a message to site 3, the owner of Y1. The owner sends a message to all proxies of Y1, on sites 1, 2, and 3. When the messages arrive, the binding Y1 $\leftarrow$ 3 appears on all these sites (see Figure 5). Since the owner is on site 3, its message to site 3 does not need any network operations.

2.2.2 Lazy and eager variables

Logic variables can have different distributed behaviors, as long as network transparency is satisfied in each case. As explained above, by default a logic variable is *eager* on all sites, i.e., its binding is sent to all sites that reference the variable. This gives maximal latency tolerance and third-party independence. However, this may cause the binding to be sent to sites that do not need it. We say that a logic variable is *lazy* on a site if its value is only sent to that site when the site needs it, e.g., when a thread is waiting for the variable. Binding a lazy variable typically needs fewer messages, since not all sites that know the variable needs its value. In some cases, such as implementing barrier synchronization using a short-circuit technique (see Section 2.4.4), lazy variables are preferable. Both eager and lazy variables obey the on-line DU algorithm of Section 5. They differ only in the scheduling of one reduction rule. The Distributed Oz implementation currently only provides eager variables; with a minor change it can provide both. A programmer annotation can then decide whether a variable is eager or lazy. The implementation issues are further explored in Section 8.5.6.

2.3 Examples of concurrent programming

It turns out that logic variables suffice to express most concurrent programming idioms in an intuitive and concise manner. Additional concepts such as semaphores, critical sections, or monitors, are not needed. Bal *et al* [5] conclude that logic variables are "spectacularly expressive" in concurrent programming even without explicit state. We give examples of four important idioms, namely synchronization, communication, mutual exclusion, and first-class channels. Many other idioms can be found in the concurrent logic programming literature [28, 12].

Synchronization and communication. The following fragment creates two threads and synchronizes their execution:

```
local X in
   thread {Print a} X=unit end
   thread {Wait X} {Print b} end
end
```

The statement `{Wait X}` blocks until `X`'s value is known. Therefore "a" is always printed before "`b`". The value of `X` is communicated from the first to the second thread. `{Wait X}` is not a new notion; it can be defined as **`if`** `X=1` **`then skip else skip end`**.

Mutual exclusion. A critical section can be defined by means of logic variables and one cell. The cell is used to manage access to a token, which is passed from one thread to the next. Assume that a cell `C` exists with initial content **`unit`**, e.g., defined by `{NewCell` **`unit`** `C}`. Then the following fragment defines a critical section:

```
local X Y in
   {Exchange C X Y} {Wait X} % Enter
   {InnerSanctum}            % Body
   Y=unit                    % Exit
end
```

We show that only one thread at a time can be executing the body. A thread that tries to enter is given `C`'s previous state X_n and current state Y_n. The thread then waits on X_n. When the previous thread leaves, it binds Y_{n-1}=**`unit`**. Since $Y_{n-1} = X_n$, this allows the next thread to enter. This works even if many threads try to enter concurrently, since the exchanges are serialized. Section 2.4.5 uses this idea to define a procedure `NewSimpleLock` that can create any number of locks.

First-class channels. A simple kind of FIFO channel is the *stream*, a list with unbound tail. Reading from the stream is receiving from the channel. Appending to the stream is sending through the channel. Each element of the stream can be a record containing both the query and an answer channel. For example, here is a fragment that handles queries appearing on the stream `X0`:

```
case X0 of query(Q1 A1)|X1 then  % Wait for query Q1 and channel A1
   A1={CalcAnswerStream Q1}      % Calculate a stream of answers on A1
   case X1 of query(Q2 A2)|X2 then  % Wait for Q1 and A2
      A2={CalcAnswerStream Q2}      % Calculate a stream of answers
```

```
      ...
   end
end
```

The **case** statement is a useful idiom: it waits until `X0` is sufficiently bound to read the pattern `query(Q1 A1)|X1`. The pattern variables `Q1`, `A1`, `X1`, are declared implicitly. Typically, the above fragment would be written as part of a loop:

```
local P in
   proc{P M}
      case M of query(Q A) then A={CalcAnswerStream Q}
   end
   {Forall X0 P}
end
```

`Forall` takes a list `X0` and a one-argument procedure, and applies the procedure to all the list's elements. The above example can be written more compactly as a nested statement with exactly the same meaning:

```
{Forall X0 proc {$ M} ... end}    % Same procedure body as above
```

This is just a short-cut that avoids explicitly declaring `P`. Using `Forall` is efficient; there are no memory leaks and the stream is not consumed faster than it is produced. It may be produced faster than it is consumed, however. Usually, a stream is associated to an Oz *port* and one writes to the port (see Section 2.1.2).

2.4 Examples of distributed programming

The purpose of this section is to show the usefulness of logic variables when extending concurrent object-oriented programming to a distributed setting. Sections 2.4.1-2.4.6 present a series of common programming idioms in distributed programming. We show how to express them in a concurrent object-oriented language that contains logic variables. The resulting solutions have two properties:

- The solutions perform correctly independently of how they are partitioned among sites. That is, the programming idioms underlying the communication patterns can be expressed in a network-transparent manner.
- No matter how the solutions are partitioned among sites, the resulting message traffic is optimal or near-optimal (in the common cases) or at least reasonable (in the uncommon cases). That is, given logic variables, the same programming idioms perform well in both centralized and distributed settings.

This shows that, at least in the cases presented here, using logic variables allows to keep useful programming idioms of centralized object-oriented programming, while allowing the implementation to extend efficiently to a distributed setting. This is evidence that controlling execution through data availability, which is what logic variables provide, is a natural way to keep good performance while mapping a program to arbitrary distribution structures.

The examples show a variety of distributed programming techniques using logic variables. Some of them, e.g., barrier synchronization and distributed locking, will normally be provided as primitives by the system. Others, e.g., stream communication, will normally be programmed by the user. We do not distinguish between the two cases since our goal is to show the expressiveness of logic variables.

2.4.1 Latency tolerance and third-party independence

From the viewpoint of execution order of basic language operations, a distributed execution cannot be distinguished from a concurrent execution. Distinguishing them requires looking at the effects of partitioning an execution over several sites. This affects system properties such

```
proc {Generate N Max ?L} % Return list of integers from N to Max-1
   if N < Max then L1 in
      L=N|L1
      {Generate N+1 Max L1}
   else L=nil end
end

fun {Sum L A}            % Return (A + sum of elements of list L)
   case L
   of nil then A
   [] X|Ls then {Sum Ls A+X}
   end
end

local CS L S in          % Generate a list and sum its elements
   CS={NewComputeServer ´sinuhe.sics.se´}  % Create remote compute server
   thread L = {Generate 0 150000} end  % Producer thread (exec. locally)
   {CS proc {$} S={Sum L 0} end}       % Consumer thread (exec. remotely)
   {Wait S} {Print S}                  % Print result (exec. locally)
end
```

Figure 6: Stream communication

as network properties (e.g., delays and limited bandwidth) and site resources (e.g., disks and main memory). At the language level, the latter shows up as the restriction of some operations to be local or remote only (such as local memory operations and remote message sends).

Logic variables decouple the declaration of a variable from its binding. Once a variable is declared, it can be passed to other sites, even before it is bound. When it is bound, the binding will be transferred automatically and efficiently to the sites needing it. This decoupling allows programs to provide a degree of *latency tolerance*, i.e., their execution is less affected by changes in network latency. For example, the following fragment:

```
local Ans in
   {DataBase query("how far is up?" Ans)}  % Remote query
   {RemoteClient inform(Ans)}              % Send answer to client
end
```

The database query and the client transfer are initiated concurrently. Assume that they are on different sites. The initiator site owns `Ans`. As soon as `Ans` is bound, the binding will be sent from the database site to the initiator site, which forwards it to the client site. This is all done independently of the initiator.

A logic variable can be bound to another logic variable. This allows programs to provide a degree of *third-party independence*, i.e., their execution is not affected by what happens at sites that are no longer involved in their execution. For example, assume variable `X` exists on sites 1 and 2 and variable `Y` exists on sites 2 and 3. Assume that `X` and `Y` are bound together on site 2. Then binding `X` to 23 on site 1 should be visible on site 3 independent of what happens to site 2.

2.4.2 Stream communication

The first example program consists of a producer that creates a data stream and a consumer that reads this stream. We first examine the program in a centralized setting. Then we explain what happens when the producer and consumer are on different sites.

```
proc {Generate N L} % Create list L of counting numbers starting with N
   case L of X|Ls then
      X=N
      {Generate N+1 Ls}
   else skip end
end

fun {Sum N L A}     % Return (A + sum of first N elements of L)
   if N>0 then X L1 in
      L=X|L1
      {Sum N-1 L1 A+X}
   else
      A
   end
end

local CS L S in
   CS={NewComputeServer ´sinuhe.sics.se´}  % Create remote compute server
   {CS proc {$} {Generate 0 L} end}    % Producer thread (exec. remotely)
   thread S={Sum 150000 L 0} end       % Consumer thread (exec. locally)
   {Wait S} {Print S}                  % Print result (exec. locally)
end
```

Figure 7: Stream communication with flow control

As we saw before, a stream is a list whose tail is a logic variable. The producer thread incrementally binds the tail to a pair of an element and a new tail. The consumer thread can start reading the stream while the producer is still writing to it. In the program of Figure 6, the producer generates a list of numbers and the consumer sums them. The consumer will print the number 11250075000. This example has no flow control, i.e., the producer will create elements eagerly. Unless the list's maximum size is small, flow control is needed to avoid problems with memory utilisation. This is true in both centralized and distributed settings.

Flow control can be added by letting the consumer inform the producer when it is able to receive the next element. We do this by letting the consumer bind the tail to a pair of a logic variable and a new tail. The producer waits until this pair exists and then binds the logic variable to the next element. The producer and consumer then execute in lock step (see Figure 7). An n-element buffer can be programmed with minor changes. The consumer can terminate the producer if the **`skip`** is replaced with `L=nil`. The producer detects the end of `L` and terminates.

To distribute these examples, let the producer thread and consumer thread run on different sites. For example, Figure 6 executes the producer locally and the consumer remotely, and Figure 7 does it the other way around. In both cases, remote execution is initiated by designating the site on which a calculation starts.

The distributed behavior is as follows. In the first example, the consumer is started by sending a request, i.e., a zero-argument procedure, to a remote compute server `CS`. This procedure and the function `Sum` are defined locally. Their compiled code is sent across the network. Because of network transparency, the procedure body can be any expression at all. The logic variable `S` is shared between the local and remote sites, and therefore transparently becomes a distributed variable. When `Sum` finishes its calculation then the result is bound to `S`. This sends the result back to the local site.

In terms of network operations, both examples are efficient. Figure 6 will send one message from the producer to the consumer for each element of the stream, exactly as desired. Figure 7

```
local CS1 CS2 CS3 L S1 S2 S3 in
   CS1={NewComputeServer ´sinuhe.sics.se´}
   CS2={NewComputeServer ´norge.info.ucl.ac.be´}
   CS3={NewComputeServer ´tinman.ps.uni-sb.de´}
   thread {Generate 0 L} end              % Producer (exec. locally)
   {CS1 proc {$} S1={Sum L 150000 0} end} % Consumer 1 (on Site 1)
   {CS2 proc {$} S2={Sum L 150000 0} end} % Consumer 2 (on Site 2)
   {CS3 proc {$} S3={Sum L 150000 0} end} % Consumer 3 (on Site 3)
end
```

Figure 8: Stream communication with multiple readers

will send one message in each direction for each element of the stream, exactly as desired.

2.4.3 Stream communication with multiple readers

Now let the stream be read by multiple consumers. Figure 8 shows how to do it with consumers on three sites. We assume three compute servers referenced by `CS1`, `CS2`, and `CS3`. Both examples of the previous section (with and without flow control) will work with multiple consumers. This is an excellent illustration of the difference between logic variables and I-structures. It is allowed for multiple readers to bind the list's tail, since they bind it in a consistent way. This would not work with ordinary single assignment, e.g., as provided by I-structures.

The example without flow control is straightforward: one message is sent to each consumer per element. The example with flow control is more interesting; it is shown in Figure 8. In this case, each consumer sends a message to request the next element when it is needed. The network behavior is as follows. To make things interesting, we assume a fast, a medium, and a slow consumer. The fast consumer sends a message to the producer, which is the owner of the first stream variable. The message contains two variables: one for the element and one for the next stream variable. Both of these variables are owned by the fast consumer. It follows that from this point on, the fast consumer will be the owner of all stream variables. Therefore all further stream elements will be sent by the producer to the fast consumer, who will multicast them to the other consumers. After the first message, the medium consumer will send requests to the fast consumer, since it is the owner. These requests will be ignored, since the fast consumer will already have bound the stream variable. The slow consumer will send no requests at all; it receives the elements before asking for them.

2.4.4 Barrier synchronization

We would like to create a set of concurrent tasks and be informed as soon as all tasks have finished. This should work efficiently independently of how the tasks are partitioned over a set of sites. Figure 9 gives a simple solution that works well in both centralized and distributed settings. To explain how it works, we need first of all to understand how to synchronize on the termination of a single thread. This is done as follows, where statement `S` represents a task:

```
local X in thread S X=unit end ... {Wait X} end
```

The main thread creates a new thread whose body is `S X=unit`. The new thread will bind `X` after `S` is finished, and the main thread detects this with a `{Wait X}`. A statement `S` *finishes* when it reduces to **`skip`** in its thread. Other threads may be created during the execution of `S`; these are independent of `S`. If the task is executed remotely, then binding `X` sends a single message to the main site, which owns `X`. This informs the thread of the task's completion. The message sent back to the task's site is a simple acknowledgement that does not affect the barrier's latency, which is one message.

```
proc {BarrierSync Ps}
   proc {Conc Ps L}
      case Ps of P|Pr then X Ls in
         L=X|Ls
         thread {P} X=unit end
         {Conc Pr Ls}
      else
         L=nil
      end
   end
   L
in
   {Conc Ps L}
   {Forall L proc {$ X} {Wait X} end}
end

{BarrierSync [proc {$} E1 end   % Task 1
              proc {$} E2 end   % Task 2
              proc {$} E3 end]} % Task 3
```

Figure 9: Barrier synchronization

We generalize this idea to multiple tasks. The general scheme is as follows:

```
local X1 ... Xn in
   thread S1 X1=unit end
   thread S2 X2=unit end
   ...
   thread Sn Xn=unit end
   {Wait X1} ... {Wait Xn} S
end
```

The main thread waits until all `Xi` are bound. When `Si` terminates then its thread binds `Xi=`**`unit`**. When all tasks terminate then all `Xi` are bound, so the main thread runs `S`.

Assume now that the tasks are distributed over a set of sites. Each `Xi` is owned by the main thread's site. Therefore binding `Xi=`**`unit`** sends a message from the task site to the main site. When all variables are bound, the main thread resumes execution. Concurrently, the main site sends a message back for each message it received. These messages do not affect the barrier's latency.

2.4.5 Distributed locking

If a program fragment may be executed by many threads, then it is important to be able to guarantee mutual exclusion. A thread that attempts to execute the fragment should block and be queued. Multiple requests should be correctly queued and blocked independent of whether the threads are on the same site or on another site. We show that it is possible to implement this concisely and efficiently in the language. As explained in Section 2.3, Figure 10 shows one way to implement a lock that handles exceptions correctly.[2] If multiple threads attempt to access the lock body, then only one is given access and the others are queued. The queue is a sequence of logic variables. Each thread blocks on one variable in the sequence, and will bind the next variable after it has executed the lock body. Each thread desiring the lock therefore references two logic variables: one to wait for the lock, and one to pass the lock to the next thread. Each logic variable is referenced by two threads.

[2] A thread-reentrant lock is defined in [38].

```
proc {NewSimpleLock ?Lock}
   Cell = {NewCell unit}
in
   proc {Lock Code}
      Old New in
      try
         {Exchange Cell Old New} {Wait Old}  % Enter
         {Code}                              % Body
      finally New=unit end                   % Exit
   end
end
```

Figure 10: Distributed locking

If the threads are on different sites, then the queue is distributed. A single message will be sent to transfer the lock from one site to another. This implements distributed token passing, which is a well-known distributed algorithm for mutual exclusion [6]. We explain how it works. When a thread tries to enter the lock body, the `Exchange` gives it access to the previous thread's `New` variable. The previous thread's site is `New`'s owner. When the previous thread binds `New`, the owner sends the binding to the next thread's site. This requires a single message.

2.4.6 Remote method invocation (RMI)

Let us invoke an object from within a thread on a given site. Where will the object execute? On a network-transparent system there are several possible answers to this question. In Distributed Oz, objects migrate to the invoking site by default [37, 38]. Therefore the object executes locally with respect to the invoking thread. This makes it easy for the object to synchronize with respect to the thread. If the object raises an exception, then it is passed to the thread. Object migration is implemented by a lightweight mobility protocol that serializes the path of the object's concurrent state pointer among the invoking sites.

It is possible in Oz to define a generic procedure that takes any object and returns a stationary object, i.e., such that all its methods will execute on the same site. This works because Oz has first-class messages and dynamic typing. This is not possible in Java. Figure 11 defines `NewStationary`, which given any object class, creates a stationary object of that class. It works by wrapping the object inside a port, which is a stationary entity to which messages can be sent asynchronously. Therefore the object always executes on the same site, namely the site on which it was created. As before, the object synchronizes with respect to the invoking thread, and exceptions are passed to the invoking thread. The logic variable `R` is used both to synchronize and to pass back exceptions.

Figure 12 defines `Obj` remotely and invokes it from the local site. For example, `{Obj get(X)} {Print X}` queries the object and prints the result on the local site. The object responds by binding the variable `X` with the answer. Since the local site owns `X`, the binding request sends one message from the remote site to the local site. With the initial invocation, this gives a total message latency of two for the remote call, just like an RPC. There is a third message back to the remote site that does not affect the message latency.

2.5 Adding logic variables to other languages

This section shows how to add logic variables in an orthogonal way to Java and ML, representative examples of object-oriented and functional languages.

```
proc {NewStationary Class Init ?StatObj}
   Obj={New Class Init}
   S P={NewPort S}
   N={NewName}
in
   proc {StatObj M}
      R in
      {Send P M#R}
      if R=N then skip
      else raise R end
      end
   end
   thread
     {ForAll S
      proc {$ M#R}
         thread
            try {Obj M} R=N
            catch X then R=X end
         end
      end}
   end
end
```

Figure 11: RMI part 1: Create a stationary object from any class

```
% Create class Counter on local site
class Counter
   attr i
   meth init i <- 0 end
   meth inc  i <- @i+1 end
   meth get(X) X=@i end
   meth error raise e(some_error) end end
end

% Create object Obj on remote site
{CS proc {$} Obj={NewStationary Counter init} end}

% Invoke object from local site
{Obj inc}
{Obj inc}
local X in {Obj get(X)} {Print X} end
try {Obj error} catch X then {Print X} end
```

Figure 12: RMI part 2: A stationary counter object

```
public class List {
   final unknown int car;
   final unknown List cdr;

   List(unknown int car, unknown List cdr) {
      this.car=:=car;
      this.cdr=:=cdr;
   }
   public void cons(unknown int car, unknown List cdr) {
      this.car=:=car;
      this.cdr=:=cdr;
   }
}
```

Figure 13: List implementation in CC-Java

2.5.1 Java

Sundström [33] has recently defined and implemented a Java variant, CC-Java (Concurrent Constraint Java), which replaces monitors by logic variables and adds statement-level thread creation. Except for these differences, CC-Java has the same syntax and semantics as Java.

CC-Java provides logic variables through a single new modifier, `unknown`, which can be used in declarations of local variables, fields, formal parameters, and functions. For example, a variable `i` declared as `unknown int i;` is initially assigned an unknown value. Standard Java variables can be freely replaced by unknown variables. The result is always a legal CC-Java program. Variables with Java types will never be assigned unknown values–any attempt will suspend the thread until the value is known.

An unknown variable is bound by the new operator "`=:=`", which does unification. Each of the two operands can be known (i.e., be a standard Java variable) or unknown. Doing `i=:=23` binds `i` to 23. For correctness, the assignment operator "`=`" must overwrite (not unify) an unknown variable on the left-hand side. Declaring an unknown variable as `final` means that it is only assigned once, i.e., when it is declared. An `final unknown` variable is therefore equivalent to an Oz logic variable. An `unknown` variable is equivalent to an Oz cell that points to a logic variable.

Figure 13 shows how to implement lists in CC-Java. Each list pair contains two logic variables, and therefore lists can be partially instantiated just like in Oz. Using logic variables does not imply any memory penalty for lists: when compiled to Oz, a CC-Java list pair uses just two memory words. Threads can synchronize on the instantiation state of lists.

Figure 14 uses these lists to write the stream communication example of Figure 6 in CC-Java (see Section 2.4.2). The `thread` statement of CC-Java is used to generate the list in another thread. The example has been written in a natural style in Oz and CC-Java, where Oz uses recursion and CC-Java uses iteration to define the `generate` and `sum` functions. Comparing the two examples, we see that there is very little difference in clarity between these two styles. Their run-time efficiencies are comparable.

It is instructive to examine more closely the CC-Java program. First, there are two synchronization points: the statements `ptr.cons(x,ls)` and `sum+=x` inside the `sum` function. The former waits until `ptr` contains a list pair and the latter waits until `x` is an integer. Second, both `final unknown` and `unknown` variables are useful. The former are used as part of data structures that should not change. The latter are used in loops that need a new logic variable in each iteration. The new variable is created with an assignment statement, e.g., `ptr=ls` creates

```
public class StreamExample {
   // Return list of integers from n to max-1
   static List generate(int n, int max) {
      final unknown List l;
      unknown List ptr=l;
      for (int i=n; i<max; i+=1) {
         final unknown List tail;
         ptr=:=new List(i,tail);
         ptr=tail;
      }
      ptr=:=null;
      return l;
   }

   // Return (a + sum of elements of list l)
   static int sum(unknown List l, int a) {
      int sum=a;
      unknown List ptr=l;
      while (ptr!=null) {
         final unknown int x;
         final unknown List ls;
         ptr.cons(x,ls);
         sum+=x;
         ptr=ls;
      }
      return sum;
   }

   // Generate a list and sum its elements
   public static void main(String[] args) {
      unknown List l;
      int sum;
      thread l=:=generate(0,1500);
      sum=sum(l,0);
      System.out.println(sum);
   }
}
```

Figure 14: Stream communication in CC-Java

a new variable referred to by `ptr`.

It is straightforward to compile CC-Java to either Oz or Java. A prototype CC-Java to Oz compiler has been implemented that understands the full Java syntax and compiles most of the Java language. Benchmarks show that CC-Java and Oz have comparable performance on the Mozart implementation of Distributed Oz. Both CC-Java and Oz on Mozart have performance comparable to Java on JDK 1.1.4, except that threads are much faster in Mozart [16, 15].

We outline how to implement a CC-Java to Java compiler. All Java code that does not use logic variables is unchanged. For each class `C` of which `unknown` instances are declared, the compiler adds a second class definition `UnknownC` to the Java code. The class `UnknownC` includes all methods of `C` and additional methods to unify the variable and to obtain its value. At each point where the value of an object of class `UnknownC` is needed, the compiler inserts a call to obtain the value. If the value is not yet available, then the calling thread is suspended until the value becomes available through unification.

2.5.2 ML

Smolka [30] has recently shown how logic variables can be added as a conservative extension to a functional language very similar to Standard ML. We outline how the extension is done. Several new operations are added, including the following:

- `lvar: unit -> 'a`. The operation `lvar()` creates a fresh logic variable and returns its address.
- `<-: 'a * 'a -> 'a`. The operation `x <- y` binds `x`, which must be a logic variable, to `y`.
- `==: 'a * 'a -> 'a`. The operation `x == y` unifies `x` and `y`. This raises an exception if `x` and `y` are not unifiable.
- `wait: 'a -> 'a`. The operation `wait x` is an identity function that blocks until its argument is nonvariable.
- `spawn` e. This operation spawns a new thread evaluating expression e and returns `()`.

Execution states map addresses to any ML entity including primitive values, records, and reference cells. Execution states are extended so that a state may also map an address to a logic variable or to another address. The entity corresponding to an address is obtained by iterating the state function until the result is no longer an address. This iteration is the dereferencing operation. If a thread needs an entity and encounters a logic variable, then it blocks until the entity is available.

With this extension, existing ML programs continue to work and logic variables may be freely intermixed with ML entities. ML provides explicit stateful entities which are called *references* and behave like typed Oz cells. As in Oz and CC-Java, the combination of logic variables and state allows to easily express powerful concurrent programming techniques. Smolka outlines the semantics of the extension and illustrates some of these programming techniques.

3 Basic concepts and notation

This section introduces the basic concepts and notation used for the centralized algorithms. Most of this notation remains valid for the distributed algorithms. The extra notation they need will be given later on.

3.1 Terms and constraints

In the usual case, a variable will be bound to a data structure. However, because of unpredictable network behavior, it may also be necessary to bind variables to variables or data structures to data structures. The result should not depend on the order in which the bindings occur. This justifies using a constraint system (D,C) to model the data structures and their bindings [18]. The domain D is the set of data structures of interest; for generality we assume these are rational trees, i.e., rooted directed graphs. The constraints C model bindings; we assume they are equalities between terms that describe sets of trees.

Constraints represent bindings that give partial information. For example, the constraint $x = f(y)$ means that the trees described by the variable x all have a root labelled f and a single subtree, which is a tree described by the variable y. In this way, we express clearly what it means to bind terms that may contain unbound variables. If y is not bound, then nothing is known about the subtree of x.

We introduce a uniform notation for terms, which can be either variables or trees that may contain variables. Terms are denoted by u, v, w. Variables are denoted by x, y, z. Nonvariable terms are denoted by t, t_1, t_2. A term can either be a variable or a nonvariable. A nonvariable term is a record of the form $f(x_1, ..., x_n)$ with arity $n \geq 0$, where x_1, ..., x_n are variables, and where the label f is an atomic constant taken from a given set of constants. A constraint has the general form $\bigwedge_i u_i = v_i$ where u_i and v_i are terms. A *basic* constraint has the form $x = u$.

To bind u and v means to add the constraint $u = v$ to the system. This is sometimes called *telling* the constraint. The operation of binding u and v is called *unification*. This is implementable in a time and space essentially equivalent to that needed for manipulating data structures in imperative languages [36]. For more on the constraint-solving aspects of unification see [18].

For the purpose of variable-variable unification, we assume a partial order between terms such that all variables are in a total order and all nonvariable terms are less than all variables. That is, we assume a transitive antisymmetric relation $\text{less}(u, v)$ such that for any distinct variables x and y, exactly one of $\text{less}(x, y)$ or $\text{less}(y, x)$ holds. In addition, for any nonvariable term t and any variable x, $\text{less}(t, x)$ holds. The algorithm uses the order to avoid creating binding cycles (e.g., x bound to y and vice versa). This is especially important in a distributed setting.

3.2 Configurations

A *configuration* $c = (\alpha; \sigma; \mu)$ of a centralized execution is a triple containing an action α, a store σ, and a memo table μ:

$$\alpha = \bigwedge_i u_i = v_i \wedge \bigwedge_i \text{false} \wedge \bigwedge_i \text{true} \qquad \sigma = \bigcup_i x_i \leftarrow u_i \qquad \mu = \bigcup_i x_i = y_i$$

The action α is a multiset of three kinds of primitive actions, of the form $u = v$, false, and true. The notation $x \leftarrow u$ represents the binding of x to u. The store is a set of bindings. All variables x_i in the store are distinct and there are no cycles $x_{a_1} \leftarrow x_{a_2}, ..., x_{a_{n-1}} \leftarrow x_{a_n}, x_{a_n} \leftarrow x_{a_1}$. The notation $x \leftarrow u, \sigma$ will be used as shorthand for $\{x \leftarrow u\} \cup \sigma$. The function $\text{lhs}(\sigma) = \bigcup_i x_i$ gives the set of bound variables in σ, which are exactly the variables on the left-hand sides of the binding arrows.

The memo table μ is used to store previously-encountered variable-variable equalities so that the algorithm does not go into an infinite loop when unifying terms representing cyclic rational trees. For example, consider the equation $x = y$ with store $x \leftarrow f(x) \wedge y \leftarrow f(y)$. Dereferencing x and y and decomposing the resulting equation $f(x) = f(y)$ gives $x = y$ again (see Section 4). This loop is broken by putting $x = y$ in the memo table and testing for its presence. The memo table has two operations, *ADD* and *MEM*, defined as follows:

$$
\begin{array}{lcl}
ADD(x = y, \mu) & = & \mu \cup \{x = y\} \\
ADD(x = t, \mu) & = & \mu \\
MEM(x = u, \mu) & = & \text{true if } x = u \in \mu \\
MEM(x = u, \mu) & = & \text{false otherwise}
\end{array}
$$

If the number of variables is finite, then the number of possible variable-variable equations in the memo table is finite also. Therefore all possible loops are broken. In the off-line case this is always true. In the on-line case this is true if the finite size property holds (see Section 7).

3.3 Adapting unification to reactive systems

A store represents a conjunction of constraints. Adding a constraint that is inconsistent with the store therefore results in the store collapsing to a "false" or "failed" state. This behavior is incompatible with a long-lived reactive system. Furthermore, it is expensive in a distributed system since it requires a global synchronization. Rather, we want an inconsistent constraint to be flagged as such (e.g., by raising an exception), without actually being put in the store. This requires two modifications to the unification algorithm:

- **Incremental tell**, i.e., information that is inconsistent with the store is not put in the store [29]. The CU and DU algorithms both implement incremental tell by decomposing a rational tree constraint into the basic constraints $x = y$ and $x = t$ and by not collapsing the store when an inconsistency is detected with a basic constraint. Inconsistency is represented as a new action "false" instead of being incorporated into the store. This new action can be used to inform the program, e.g., by raising an exception.
- **Full failure reporting**, i.e., every case of inconsistency is flagged to the program. Neither the CU nor the DU algorithms do this. If the inconsistent equation $x = y$ is present more than once, then the CU algorithm flags this only once. The DU algorithm flags the inconsistency of $(x = y)_s$ only once per site s. That is, both algorithms flag a given inconsistency only once per memo table. Inconsistencies can be flagged more often at the price of sometimes doing redundant work. For example, if equations reduce within a thread and each thread is given its own memo table, then multiple occurrences of an inconsistent equation will be flagged once per thread. The Distributed Oz implementation will flag an inconsistency at least once per program-invoked unification (see Section 8).

3.4 Rules and algorithms

We define an algorithm by a set of reduction rules, where a rule defines a transition relation between configurations. The algorithms in this article all have a straightforward translation into an efficient imperative pseudocode. We do not define the algorithms in this way since it complicates reasoning about them. Rule reduction is an atomic operation. If more than one rule is applicable in a given configuration, then one is chosen nondeterministically. A rule is defined according to the following diagram:

$$
\begin{array}{c|c}
\alpha & \alpha' \\
\hline
\sigma;\, \mu & \sigma';\, \mu'
\end{array}
\quad C
$$

A rule becomes applicable for a given action α when the actual store matches the store σ given in the rule and the optional condition C is satisfied. The rule's reduction atomically replaces the current configuration (α;σ;μ) by the result configuration (α';σ';μ').

Both the centralized and the distributed algorithms are defined in the context of a structure rule and a congruence.

Structure $\quad \dfrac{\alpha_1 \wedge \alpha_2}{\sigma;\ \mu} \Bigg\| \dfrac{\alpha_1' \wedge \alpha_2}{\sigma';\ \mu'} \quad \textbf{if} \quad \dfrac{\alpha_1}{\sigma;\ \mu} \Bigg\| \dfrac{\alpha_1'}{\sigma';\ \mu'}$

Congruence $\quad \begin{cases} \alpha_1 \wedge \alpha_2 \equiv \alpha_2 \wedge \alpha_1 \\ \text{true} \wedge \alpha \equiv \alpha \end{cases}$

Because of the congruence, the primitive action true may occur an arbitrary number of times. This is not true of the other primitive actions, which form a multiset.

3.5 Executions

An *execution* e of a given algorithm is a (possibly infinite) sequence of configurations such that the first configuration is an initial configuration and each transition corresponds to the reduction of one rule:

$$c_1 \xrightarrow{R_1} c_2 \xrightarrow{R_2} \cdots \xrightarrow{R_{n-1}} c_n$$

In any execution, distributed or centralized, we assume that all rule reductions are done in some global order. This is done without loss of generality since the results we prove in this article will hold for *all* possible executions. Therefore, reductions that are not causally related may be assumed to execute in any order.

An *initial* configuration of the CU algorithm is of the form $(\alpha_1;\emptyset;\emptyset)$, where α_1 is a finite conjunction of equations and the store and memo table are both empty. A *final* configuration (if it exists) is a configuration where no rules are applicable. The last configuration of a (finite) execution is not necessarily a final configuration; it is simply the configuration that has no successors. A *valid* configuration is one that is reachable by an execution of a given algorithm. We speak of centralized executions (using the CU or RCU algorithms, see Sections 4 and 6.2) and distributed executions (using the DU algorithm, see Section 5).

4 Centralized unification (CU algorithm)

This section defines the CU algorithm, an algorithm for rational tree unification. The definition is given as a set of transition rules with an interleaving semantics. That is, rule reductions do not overlap in time. Only their order is important. Nothing is lost by using an interleaving semantics, since we prove properties that are true for all interleavings that are valid executions [2].

The CU algorithm is the base case for the DU algorithm of Section 5. There are many possible variant definitions of unification. The definition presented here is deterministic and constrained by the requirement that it must be extensible to a practical distributed algorithm. That is, the use of global conditions is minimized (see Section 5.1).

4.1 Definition

We define the CU algorithm by the following seven rules.

Interchange $\quad \dfrac{u = x}{\sigma;\ \mu} \Bigg\| \dfrac{x = u}{\sigma;\ \mu} \quad \text{less}(u, x)$

Bind $\quad \dfrac{x = u}{\sigma;\ \mu} \Bigg\| \dfrac{\text{true}}{x \leftarrow u, \sigma;\ \mu} \quad \text{less}(u, x), x \notin \text{lhs}(\sigma)$

Memo Hit $\quad \dfrac{x = u}{x \leftarrow v, \sigma;\ \mu} \Bigg\| \dfrac{\text{true}}{x \leftarrow v, \sigma;\ \mu} \quad \text{less}(u, x),\ \mathit{MEM}(x = u, \mu)$

Dereference $\quad \dfrac{x = u}{x \leftarrow v, \sigma;\ \mu} \Bigg\| \dfrac{v = u}{x \leftarrow v, \sigma;\ \mathit{ADD}(x = u, \mu)} \quad \text{less}(u, x),\ \neg\mathit{MEM}(x = u, \mu)$

IDENTIFY	$\dfrac{x = x \;\Vert\; \text{true}}{\sigma;\,\mu \;\Vert\; \sigma;\,\mu}$	
CONFLICT	$\dfrac{t_1 = t_2 \;\Vert\; \text{false}}{\sigma;\,\mu \;\Vert\; \sigma;\,\mu}$	$t_1 = f_1(x_1, ..., x_m),\, t_2 = f_2(y_1, ..., y_n)$ $(f_1 \neq f_2 \vee m \neq n)$
DECOMPOSE	$\dfrac{t_1 = t_2 \;\Vert\; \bigwedge_{1\leq i\leq n} x_i = y_i}{\sigma;\,\mu \;\Vert\; \sigma;\,\mu}$	$t_1 = f(x_1, ..., x_n),\, t_2 = f(y_1, ..., y_n)$

4.2 Properties

For the correctness proof of the on-line algorithm, we need the following entailment property of the CU algorithm. We will prove that the DU algorithm implements the CU algorithm. It follows that the entailment property also holds for the DU algorithm.

Logical formula of a configuration. A configuration $c = (\alpha;\sigma;\_)$ has an associated logical formula $\varepsilon(c) = \varepsilon_a(\alpha) \wedge \varepsilon_s(\sigma)$, where:

$$
\begin{array}{lcl}
\varepsilon_a(u = v \wedge \alpha) & = & u = v \wedge \varepsilon_a(\alpha) \\
\varepsilon_a(\text{true}) & = & \text{true} \\
\varepsilon_a(\text{false}) & = & \text{false} \\
\varepsilon_s(\{x \leftarrow u\} \cup \sigma) & = & x = u \wedge \varepsilon_s(\sigma) \\
\varepsilon_s(\emptyset) & = & \text{true}
\end{array}
$$

Theorem 1 (Logical equivalence property) *In every transition $c_i \rightarrow c_{i+1}$ of every execution of the CU algorithm, the logical equivalence $\varepsilon(c_i) \leftrightarrow \varepsilon(c_{i+1})$ holds under the standard equality theory.*

PROOF. By standard equality theory we mean the theory $\mathcal{E}$ given by Lloyd [21] (page 79), minus the acyclicity axioms (i.e., rule 4). This theory has the usual axioms implying nonequality of distinct symbols, term equality implies argument equality and vice versa, substitution by equals is allowed, and identity. What we want to prove is $\mathcal{E} \models \bar{\forall}\varepsilon(c_i) \leftrightarrow \varepsilon(c_{i+1})$, where the quantification is over all free variables. This is a standard result in unification theory. □

Corollary 2 (Entailment property or CU correctness) *Given any initial configuration $c_1 = (\alpha_1;\emptyset;\emptyset)$ of the CU algorithm. Then the algorithm terminates in a final configuration $c_n = (\alpha_n;\sigma_n;\_)$ with $\varepsilon(c_n) \leftrightarrow \varepsilon(c_1)$. Furthermore, α_n consists of zero or more false actions, and if there are zero, then $\varepsilon_s(\sigma_n) \leftrightarrow \varepsilon_a(\alpha_1)$.*

PROOF. This follows by induction from the previous theorem. □

5 Distributed unification (DU algorithm)

This section defines a distributed algorithm for rational tree unification. The section is organized as follows. Section 5.1 explains how to generalize the CU algorithm to a distributed setting. Section 5.2 sets the stage by giving the basic concepts and notation needed in the DU algorithm. Section 5.3 defines the DU algorithm in two parts: the non-bind rules, which are the local part of the algorithm, and the bind rules, which are the distributed part. Finally, Section 5.4 compares the CU and DU algorithms from the viewpoint of the dereference operation.

5.1 Generalizing CU to a distributed setting

A distributed algorithm must be defined by reduction rules that do local operations only, since these are the only rules we can implement directly. To be precise, two conditions must be satisfied. First, testing whether a rule is applicable should require looking only at one site. Second, reducing the rule should modify only that site, except that the rule is allowed to create actions annotated with other sites. In the distributed system these actions correspond to messages. Rules that satisfy these two conditions are called *local* rules. A distributed algorithm defined in terms of local rules is a *transition system* [35] in an asynchronous network that is not FIFO.

We would like to extend each CU rule to become a local rule in the distributed setting. In this way, we maintain a close correspondence between the centralized and distributed algorithms, which simplifies analysis of the distributed case. Furthermore, this minimizes costly communication between sites.

The first step is to annotate the rule's actions and bindings with sites. Each CU rule reduces an input action and may inspect a binding in the store. We annotate the input action by its site and the binding by the same site. This is correct if we assume that a binding will eventually appear on each site that references the variable. We annotate the output action by the same site as the input action. A "true" output action does not need a site. Actions may remain unannotated, in which case the DU algorithm does not specify where they are reduced. These annotations suffice for the rules INTERCHANGE, IDENTIFY, CONFLICT, and DECOMPOSE to become DU rules. An important property of CONFLICT is that an inconsistency is always flagged on the site that causes it.

The three remaining CU rules cannot be so easily extended since they have global conditions. To be precise, BIND has the unboundness condition $x \notin \text{lhs}(\sigma)$,[3] and MEMO HIT and DEREFERENCE both have the memoization condition $MEM(x = u, \mu)$. It turns out that the memoization condition can be relaxed in the distributed algorithm, so that it becomes a local condition there. In this way, the MEMO HIT and DEREFERENCE rules become local rules. We give each site its own memo table, which is independent of the other memo tables. Section 6.2 proves that this relaxation is correct, but that redundant local work may be done.

The unboundness condition of BIND cannot be eliminated in this way. Implementing it requires communication between sites. The single BIND rule therefore becomes a *set* of local rules in the distributed setting. The BIND rule is replaced by four rules that exchange messages to implement a coherent variable elimination algorithm.

The resulting DU algorithm consists of ten local rules, namely six non-bind rules (Section 5.3.1) and four bind rules (Section 5.3.2). The six non-bind rules do not send any messages. Of the four bind rules, only INITIATE and WIN send messages. All rules test applicability by looking at one site only, except for WIN and LOSE, which use information tied to a variable but not tied to any particular site, namely a flag unbound(x)/bound(x) and a binding request $(x \sim u)$.

5.2 Basic concepts and notation

We introduce a set $S = \{1, ..., k\}$ of k sites, where $k \geq 1$. We model distributed execution by *placing* each action and each binding on a site. A primitive action or binding ξ is placed on site s by adding parentheses and a subscript $(\xi)_s$. The same ξ may be placed on several sites, in which case the resulting actions or bindings are considered separately.

A configuration of a distributed execution is a triple $(A;\Sigma;M)$ consisting of an action A, a store Σ, and a memo table M. We denote the action, store, and memo table on site s by A_s, Σ_s, and M_s, respectively.

The store Σ contains sited bindings $(x \leftarrow u)_s$, sited binding initiations $(x \leftarrow \perp)_s$, and flags to denote whether a variable is bound or not (bound(x) or unbound(x)). This is organized as

[3] The opposite condition, confirming the existence of a binding, is local.

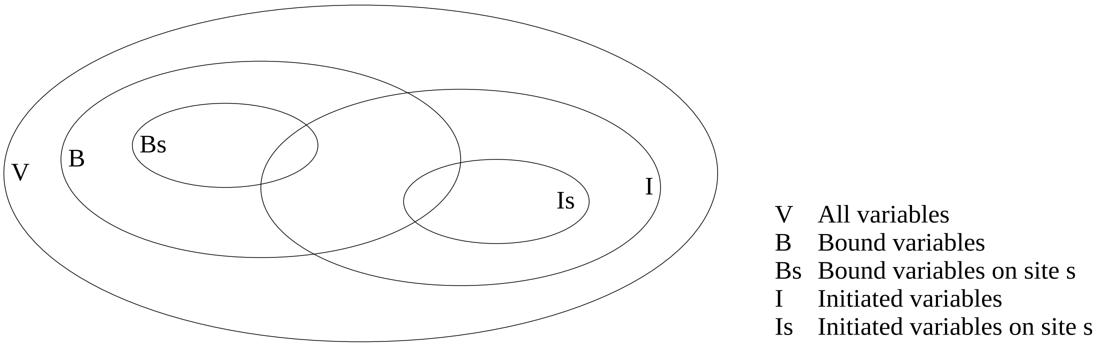

Figure 15: Bound variables and initiated variables

follows:

$$\begin{aligned}
\Sigma &= \Gamma \cup \bigcup_{s \in S} \Sigma_s \\
\Sigma_s &= \bigcup_{x_i \in B_s} (x_i \leftarrow u_i)_s \cup \bigcup_{x_i \in I_s} (x_i \leftarrow \perp)_s \\
\Gamma &= \bigcup_{x_i \in B} \mathrm{bound}(x_i) \cup \bigcup_{x_i \in V - B} \mathrm{unbound}(x_i)
\end{aligned}$$

The set V consists of all variables in A and Σ. The set $B \subseteq V$ contains all the *bound* variables. The set $B_s \subseteq B$ contains all the bound variables whose binding is known on site s. The set $I \subseteq V$ contains all the variables whose binding has been *initiated* on some site but whose binding (if it exists) is not yet known on that site. The set $I_s \subseteq I$ contains all the variables whose binding has been initiated on site s but is not yet known on that site. In terms of the bind rules of Section 5.3.2, B corresponds to those variables for which the WIN rule has reduced. I corresponds to those variables for which the INITIATE rule has reduced but the corresponding ARRIVE rule has not yet reduced. Figure 15 illustrates the relationship between these five sets.

The initial configuration is $(A^{\mathrm{init}};\Sigma^{\mathrm{init}};\emptyset)$, with initial actions A^{init} and $B = \emptyset$, $I = \emptyset$. Therefore $\Sigma^{\mathrm{init}} = \{\mathrm{unbound}(x_i) \mid x_i \in V\}$.

The action A is a multiset containing two new primitive actions in addition to the three primitive actions of the centralized setting. The new actions are needed to implement the distributed operations of the algorithm. This gives the following five primitive actions:

true	Null action
false_s	Failure notification on site s
$(u = v)_s$	Equation on site s
$x \sim u$	Binding request
$(x \Leftarrow u)_s$	Binding in transit to site s

The exact meaning of these actions is defined in the next section by the reduction rules that manipulate them. Intuitively, the action $x \sim u$ represents a message requesting the binding of x to u. For a given x, exactly one such action will cause a binding to be made; all others are discarded. The algorithm does not specify where the binding decision is made. An actual implementation can make the decision in a centralized or distributed way. In the Distributed Oz implementation, the decision is centralized; it is made on the site at which the variable was initially declared (see Section 8). The action $(x \Leftarrow u)_s$ represents a message containing a binding to x that may eventually become visible in Σ_s. As long as x is not in $\mathrm{var}(\Sigma_s)$ then the binding will stay "in the network".

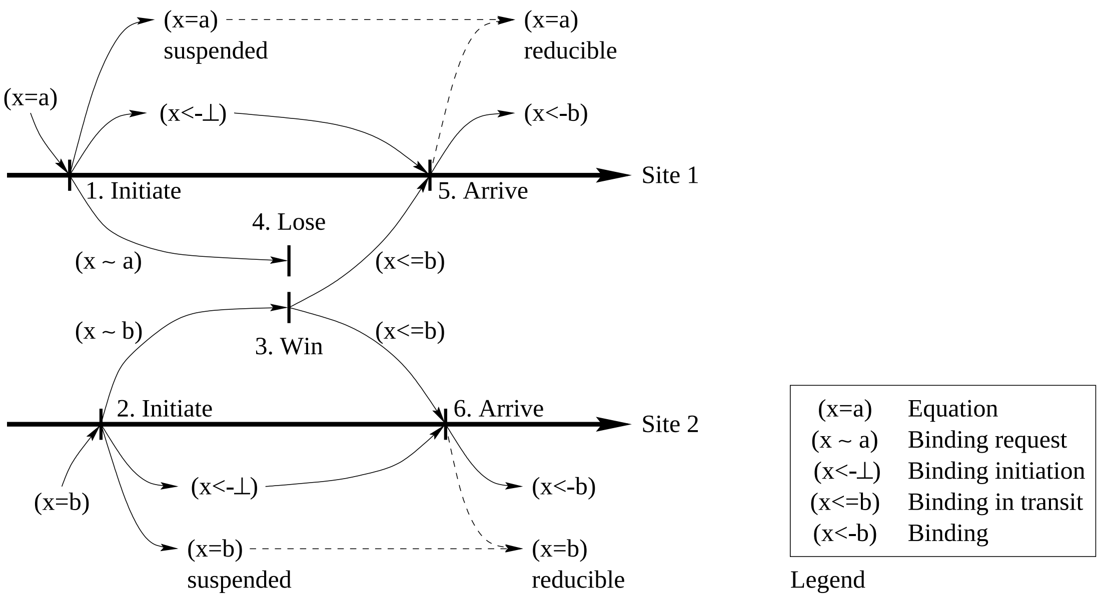

Figure 16: Distributed unification with $(x = a)_1$ and $(x = b)_2$

The global memo table M is the juxtaposition of all the local memo tables. That is, $M = \{(x = y)_s | x = y \in M_s\}$. Each local memo table M_s is a set of variable-variable equalities that has identical structure to the centralized memo table μ.

The following two functions on a local store Σ_s will be used in defining the algorithm:

$$\begin{aligned}
\mathrm{lhs}(\Sigma_s) &= \{x | \exists u.(x \leftarrow u)_s \in \Sigma_s \vee (x \leftarrow \perp)_s \in \Sigma_s\} \\
\mathrm{var}(\Sigma_s) &= \mathrm{lhs}(\Sigma_s) \cup \{x | \exists y, u.((y \leftarrow u)_s \in \Sigma_s \wedge u \equiv f(..., x, ...))\}
\end{aligned}$$

The function $\mathrm{lhs}(\Sigma_s)$ returns all bound and initiated variables of Σ_s. It generalizes the function $\mathrm{lhs}(\sigma)$ defined in Section 3.2. The function $\mathrm{var}(\Sigma_s)$ returns all variables mentioned in Σ_s, even if not bound.

5.3 Definition

The DU algorithm has a close relationship to the CU algorithm. The structure and congruence rules continue to hold in the distributed setting. The only centralized rule that is changed is the BIND rule; it is replaced by the four bind rules below. It is clear from inspection that all six non-bind DU rules have identical behavior to the corresponding CU rules, if the CU rules are considered as acting on a single site.

Figure 16 gives an example execution that does distributed variable elimination. In this figure, thin solid arrows represent actions or bindings. Vertical bars "**|**" denote rule reductions, which happen in the numbered order shown. Thin dotted arrows represent causal links.

Initially, site 1 has equation $(x = a)_1$ and site 2 has equation $(x = b)_2$. Both sites do the INITIATE rule, which puts binding initiations $(x \leftarrow \perp)_s$ in both local stores. This ensures that the two equations cannot reduce until the binding arrives. We say that the equations are *suspended*. Equation $(x = b)_2$ is the first to do the WIN rule, and b therefore becomes the global binding of x. The other equation is discarded by the LOSE rule. The binding $(x \Leftarrow b)$ is sent to all sites. It arrives at each site through the ARRIVE rule. At that point, the suspended equations $(x = a)_1$ and $(x = b)_2$ become reducible again. The equation $(x = a)_1$ will cause an inconsistency to be flagged on site 1.

5.3.1 Non-bind rules

These rules correspond exactly to the non-bind rules of the centralized algorithm. An inconsistency is flagged on the site that causes it by the action false_s.

$$\textsc{Interchange} \qquad \frac{(u = x)_s}{\Sigma;\ M} \;\Big\|\; \frac{(x = u)_s}{\Sigma;\ M} \qquad \text{less}(u, x)$$

$$\textsc{Memo Hit} \qquad \frac{(x = u)_s}{\begin{array}{c}(x \leftarrow v)_s, \Sigma;\\ M_s \cup M\end{array}} \;\Big\|\; \frac{\text{true}}{\begin{array}{c}(x \leftarrow v)_s, \Sigma;\\ M_s \cup M\end{array}} \qquad \begin{array}{l}\text{less}(u, x),\\ \mathit{MEM}(x = u, M_s)\end{array}$$

$$\textsc{Dereference} \qquad \frac{(x = u)_s}{\begin{array}{c}(x \leftarrow v)_s, \Sigma;\\ M_s \cup M\end{array}} \;\Big\|\; \frac{(v = u)_s}{\begin{array}{c}(x \leftarrow v)_s, \Sigma;\\ \mathit{ADD}(x = u, M_s) \cup M\end{array}} \qquad \begin{array}{l}\text{less}(u, x),\\ \neg \mathit{MEM}(x = u, M_s)\end{array}$$

$$\textsc{Identify} \qquad \frac{(x = x)_s}{\Sigma;\ M} \;\Big\|\; \frac{\text{true}}{\Sigma;\ M}$$

$$\textsc{Conflict} \qquad \frac{(t_1 = t_2)_s}{\Sigma;\ M} \;\Big\|\; \frac{\text{false}_s}{\Sigma;\ M} \qquad \begin{array}{l}t_1 = f_1(x_1, ..., x_m),\ t_2 = f_2(y_1, ..., y_n)\\ (f_1 \neq f_2 \vee m \neq n)\end{array}$$

$$\textsc{Decompose} \qquad \frac{(t_1 = t_2)_s}{\Sigma;\ M} \;\Big\|\; \frac{\bigwedge_{1 \le i \le n} (x_i = y_i)_s}{\Sigma;\ M} \qquad \begin{array}{l}t_1 = f(x_1, ..., x_n),\\ t_2 = f(y_1, ..., y_n)\end{array}$$

5.3.2 Bind rules

These rules replace the $\textsc{Bind}$ rule of the centralized algorithm. The binding initiation $(x \leftarrow \perp)_s$ and the condition $x \notin \text{lhs}(\Sigma_s)$ in the $\textsc{Initiate}$ rule together ensure that only one binding attempt can be made per site.

$$\textsc{Initiate} \qquad \frac{(x = u)_s}{\Sigma;\ M} \;\Big\|\; \frac{x \sim u \wedge (x = u)_s}{(x \leftarrow \perp)_s, \Sigma;\ M} \qquad \text{less}(u, x),\ x \notin \text{lhs}(\Sigma_s)$$

$$\textsc{Win} \qquad \frac{x \sim u}{\text{unbound}(x), \Sigma;\ M} \;\Big\|\; \frac{\bigwedge_{s \in S} (x \Leftarrow u)_s}{\text{bound}(x), \Sigma;\ M}$$

$$\textsc{Lose} \qquad \frac{x \sim u}{\text{bound}(x), \Sigma;\ M} \;\Big\|\; \frac{\text{true}}{\text{bound}(x), \Sigma;\ M}$$

$$\textsc{Arrive} \qquad \frac{(x \Leftarrow u)_s}{\Sigma;\ M} \;\Big\|\; \frac{\text{true}}{(x \leftarrow u)_s, \Sigma - \{(x \leftarrow \perp)_s\};\ M} \qquad x \in \text{var}(\Sigma_s)$$

5.4 Dereference chains

A *dereference chain* in a store σ is a sequence of bindings $x_1 \leftarrow x_2$, ..., $x_{n-1} \leftarrow u_n$ with $n \geq 1$ and u_n unbound. We say the value of x_1 in store σ is u_n. To find u_n given x_1, it is necessary to follow the chain. A major difference between CU and DU is that CU always constructs dereference chains, whereas DU with eager variables forbids dereference chains to cross sites. Instead, DU copies remote terms to make them local. In a centralized setting, pointer dereferencing is fast. The penalty of using dereference chains is small. This allows to maximize sharing between terms. CU guarantees that any term exists at most once in the store. In a distributed setting, pointer dereferencing is slow if the target is on another site. It also makes the current site dependent on the other site. Therefore copying is preferable.

Copying terms instead of creating dereference chains introduces redundant work. It is possible to reduce this work at the cost of more network operations. For example one can

	Distributed	Centralized
Action	true	true
	false_s	false
	$(u = v)_s$	$u = v$
	$x \sim u$	true
	$(x \Leftarrow u)_s$	$x \leftarrow u$
Store	$\text{bound}(x)$	true
	$\text{unbound}(x)$	true
	$(x \leftarrow \perp)_s$	true
	$(x \leftarrow u)_s$	$x \leftarrow u$

Table 1: Mapping from distributed to centralized configurations

eagerly bind to variables and lazily bind to (big) nonvariable terms. This guarantees that a cross-site dereference chain has a maximum length of one.

DU with lazy variables allows dereference chains to cross sites. When the value is needed, the binding is requested from the owner. If the binding is another variable, then the process is repeated. Each iteration of this process corresponds to a dereference operation.

Taylor [34] presents a centralized binding algorithm that avoids all dereference chains. Variables that are bound together are put into a circular linked list. When one of them is bound to a value, the list is traversed and all members are bound. This makes accessing a variable's value a constant-time operation, at the price of making binding more expensive. Taylor finds no significant performance difference between this algorithm and the standard algorithm when both are embedded in a high-performance Prolog system.

6 Off-line correctness

This section proves that the DU algorithm behaves as expected. We first define a mapping from any distributed to a centralized execution. Then we define a modification of the CU algorithm, the RCU algorithm, that models the redundant work done by the distributed algorithm. Then we prove safety and liveness properties by reducing the DU algorithm to the RCU algorithm. From this we show that the DU algorithm is correct.

We distinguish between the off-line correctness and the on-line correctness. In the off-line case, we have to show that the distributed algorithm terminates and gives correct results for any placement of a fixed set of initial equations. This can be done without any fairness assumptions. This is not true for the on-line case, which is handled in Section 7.

6.1 Mapping from distributed to centralized executions

The proofs in this section are based on a mapping m from any distributed configuration $(A;\Sigma;M)$ to a corresponding centralized configuration $(\alpha;\sigma;\mu)$. Distributed executions are mapped to centralized executions by mapping each of their configurations. The mapping m was designed according to the reasoning of Section 5.1. We show that m has very strong properties that lead directly to a proof that the distributed algorithm implements the centralized algorithm.

A primitive action is mapped to either a primitive action or a binding. A binding is always mapped to a binding. Other store contents map to true. In this way we can map any distributed configuration to a centralized one:

$$(A; \Sigma; M) \xrightarrow{m} (\alpha; \sigma; \mu) = (m_a(A); m_s(A, \Sigma); m_m(M))$$

Table 1 defines the mappings m_a and m_s for primitive actions and store contents. The mapping for all of A and Σ is the union of these mappings for all primitive actions in A and all store contents in Σ. The centralized memo table is the union of the tables on each site, namely $m_m(M) = \bigcup_{s\in S} M_s$.

The following diagram relates a distributed execution e and its corresponding centralized execution $m(e)$:

$$\begin{array}{ccc} (A;\Sigma;M) & \xrightarrow{e} & (A';\Sigma';M') \\ m\downarrow & & \downarrow m \\ (\alpha;\sigma;\mu) & \xrightarrow{m(e)} & (\alpha';\sigma';\mu') \end{array}$$

To show correctness, i.e., that the distributed algorithm is an implementation of unification, we need to show both safety and liveness properties. A sufficient safety property is proved in Section 6.3: given any distributed execution e, the corresponding $m(e)$ is a correct centralized execution. A sufficient liveness property is proved in Section 6.4: given any e, its execution eventually makes progress. That is, if the last configuration of $m(e)$ not final, then the last configuration of e is not final and continuing e will always eventually advance the centralized execution. In the distributed execution, the non-bind rules and the WIN rule are called *progressing* rules since they advance the centralized execution (see Table 2). The other rules are called *non-progressing*.

6.2 Redundancy in distributed unification (RCU algorithm)

This section defines and justifies a revised version of the CU algorithm, the RCU algorithm, that models the redundant work introduced by distributing rational tree unification. There are two sources of redundant work in the DU algorithm. The first source is due to the decoupling of binding initiation from binding arrival. A binding initiation for $(x = u)_s$ inhibits reduction of all equations of the form $(x = v)_s$. When a binding arrives on a site, these reductions become possible again, including the reduction of the original equation $(x = u)_s$. To make the original equation disappear, several rule reductions are needed including a DEREFERENCE, one or more IDENTIFY, and possibly a DECOMPOSE. This redundant work can be avoided in the implementation (see Section 8.5.4).

The second source of redundant work cannot be avoided in the implementation. It is due to each site having its own local memo table. Memo tables are needed because of rational trees with cycles. However, they are in fact a general caching technique that avoids doing any unification more than once. In the distributed algorithm, the information stored in each site's memo table is not seen by the other sites. Therefore each site has to reconstruct the part of the centralized memo table that it needs.

To model the local memo tables it suffices to weaken the memo table membership check. This affects the two rules MEMO HIT and DEREFERENCE. Assume there are k sites. We introduce a weaker membership check MEM_k that is true only if the equation has been entered at least k times. This is implemented by extending the memo table to store pairs of an equation and the number of times the equation has been entered:

$$\begin{array}{lcl} ADD(x=y,\mu) & = & \mu \cup \{(x=y,1)\} \text{ if } (x=y,\_) \notin \mu \\ ADD(x=y,\mu) & = & \mu - \{(x=y,i)\} \cup \{(x=y,i+1)\} \text{ if } (x=y,i) \in \mu \\ ADD(x=t,\mu) & = & \mu \\ MEM(x=u,\mu) & = & \text{true if } (x=u,\_) \in \mu \\ MEM(x=u,\mu) & = & \text{false otherwise} \\ MEM_k(x=u,\mu) & = & \text{true if } (x=u,i) \in \mu \wedge i \geq k \\ MEM_k(x=u,\mu) & = & \text{false otherwise} \end{array}$$

The R-MEMO HIT rule uses the new definition of MEM. The R-DEREFERENCE rule uses MEM_k and the new definition of ADD. If an equation has been entered from 1 to $k-1$ times then both rules are applicable.

Now we can update the CU algorithm to model the two sources of redundant work. We model memo table redundancy by replacing MEMO HIT and DEREFERENCE by R-MEMO HIT

Distributed rule	Centralized rule
Memo Hit	R-Memo Hit
Dereference	R-Dereference
Interchange	Interchange
Identify	Identify
Conflict	Conflict
Decompose	Decompose
Initiate	Skip
Win	R-Bind
Lose	Skip
Arrive	Skip

Table 2: Correspondence between distributed and centralized rules

and R-Dereference. We model bind redundancy by replacing Bind by R-Bind, as defined below. The three new rules are as follows:

R-Bind
$$\frac{x = u}{\sigma;\ \mu} \;\Bigg\|\; \frac{x = u}{x \leftarrow u, \sigma;\ \mu} \qquad \mathrm{less}(u, x), x \notin \mathrm{lhs}(\sigma)$$

R-Memo Hit
$$\frac{x = u}{x \leftarrow v, \sigma;\ \mu} \;\Bigg\|\; \frac{\mathrm{true}}{x \leftarrow v, \sigma;\ \mu} \qquad \begin{array}{l}\mathrm{less}(u, x),\\ \mathit{MEM}(x = u, \mu)\end{array}$$

R-Dereference
$$\frac{x = u}{x \leftarrow v, \sigma;\ \mu} \;\Bigg\|\; \frac{v = u}{x \leftarrow v, \sigma;\ \mathit{ADD}(x = u, \mu)} \qquad \begin{array}{l}\mathrm{less}(u, x),\\ \neg \mathit{MEM}_k(x = u, \mu)\end{array}$$

Theorem 3 (RCU correctness) *Given any initial configuration. Then the following two statements hold:*

1. *The RCU algorithm terminates.*
2. *All final configurations of the RCU and CU algorithms are logically equivalent to each other according to the definition of Section 4.2.*

Proof.

1. We know that CU terminates. The redundant work introduced by RCU has the following effect:
   - **Bind redundancy**. The R-Bind rule introduces extra rule reductions. For each binding $x = u$, the number of extra rules is bounded from above. The bound is determined by the term u; it is always at least 2 and it is increased by the arity of u if u is nonvariable.
   - **Memo table redundancy.** The memo table size for RCU is at most k times that of CU, which is finite. Hence only a finite number of extra rule reductions can be done.
2. For both bind and memo table redundancy, the additional equations are always duplicates of existing equations or equations of some previous configuration. Therefore they add no additional information and the Entailment property still holds.

This completes the proof. □

6.3 Safety

Theorem 4 (DU safety) *If e is any execution of the DU algorithm, then $m(e)$ is an execution of the RCU algorithm, and the sequence of rules reduced in $m(e)$ can be constructed from e.*

PROOF. We will prove that Table 2 correctly gives the centralized rule of $m(e)$ corresponding to a distributed rule in e. A "SKIP" rule means that no rule is executed. The proof is by induction on the length of execution e. In the base case, the initial configuration c_1 of e has an empty store and memo table, and a set of equations placed on different sites. Therefore $m(c_1)$ is a valid initial configuration for the centralized algorithm.

In the induction case, we assume that the theorem holds for an execution e. We need to show that for each distributed rule applicable in the last configuration of e, that doing this rule maps to doing a corresponding centralized rule. We do a case analysis over the distributed rules. We first consider the DECOMPOSE rule and then handle the other non-bind rules similarly (Section 6.3.1). We then consider the bind rules (Section 6.3.2).

6.3.1 Non-bind rules

Decompose. Assume that the distributed execution reduces a DECOMPOSE rule. Mapping the initial and final configurations of the decomposition gives the following diagram:

$$\begin{array}{ccc}
\dfrac{(t_1 = t_2)_s \wedge A}{\Sigma; M} & \overset{DEC}{\longrightarrow} & \dfrac{\bigwedge_i (x_i = y_i)_s \wedge A}{\Sigma; M} \\
m \downarrow & & \downarrow m \\
\dfrac{t_1 = t_2 \wedge m_a(A)}{m_s(A, \Sigma); m_m(M)} & \overset{X}{\longrightarrow} & \dfrac{\bigwedge_i x_i = y_i \wedge m_a(A)}{m_s(A, \Sigma); m_m(M)}
\end{array}$$

It is clear from inspection that rule X is a centralized decomposition.

Interchange, Identify, Conflict. These three rules are handled in the same way as the DECOMPOSE rule.

Memo Hit. It is clear that the MEMO HIT rule maps to an R-MEMO HIT rule, since from $M_s \subseteq \mu$ it follows that that $MEM(x = u, M_s) \Rightarrow MEM(x = u, \mu)$.

Dereference. We now show that the DEREFERENCE rule maps to an R-DEREFERENCE rule. We have the following diagram (where $\Sigma_x = (x \leftarrow v)_s, \Sigma)$:

$$\begin{array}{ccc}
\dfrac{(x = u)_s \wedge A}{\Sigma_x; M_s \cup M} & \overset{DRF}{\longrightarrow} & \dfrac{(v = u)_s \wedge A}{\Sigma_x; ADD(x = u, M_s) \cup M} \\
m \downarrow & & \downarrow m \\
\dfrac{x = u \wedge m_a(A)}{m_s(A, \Sigma_x); m_m(M_s \cup M)} & \overset{X}{\longrightarrow} & \dfrac{v = u \wedge m_a(A)}{m_s(A, \Sigma_x); m_m(ADD(x = u, M_s) \cup M)}
\end{array}$$

We know $\text{less}(u, x) \wedge \neg MEM(x = u, M_s)$. Since each site has its own local memo table, there can be only one redundant equation per site. Therefore $\neg MEM(x = u, M_s)$ implies that $\neg MEM_k(x = u, \mu)$. That is, if at least one local memo table does not contain $x = u$, then the centralized memo table contains $x = u$ less than k times. It follows that rule X is exactly an R-DEREFERENCE rule. This shows that relaxing the memoization condition to do only a check on the local part of the memo table is correct but may introduce one redundant equation per site and per original equation.

6.3.2 Bind rules

Initiate.

$$\begin{array}{ccc}
\dfrac{(x = u)_s \wedge A}{\Sigma; M} & \stackrel{INI}{\longrightarrow} & \dfrac{x \sim u \wedge (x = u)_s \wedge A}{(x \leftarrow \bot)_s \Sigma; M} \\
m \downarrow & & \downarrow m \\
\dfrac{x = u \wedge m_a(A)}{m_s(A, \Sigma); m_m(M)} & \stackrel{X}{\longrightarrow} & \dfrac{x = u \wedge m_a(A)}{m_s(A, \Sigma); m_m(M)}
\end{array}$$

It is clear from inspection that X is a SKIP rule. After the SKIP, we know that less(u, x).

Win.

$$\begin{array}{ccc}
\dfrac{x \sim u \wedge A}{\mathrm{unbound}(x), \Sigma; M} & \stackrel{WIN}{\longrightarrow} & \dfrac{\bigwedge_{s \in S} (x \Leftarrow u)_s \wedge A}{\mathrm{bound}(x), \Sigma; M} \\
m \downarrow & & \downarrow m \\
\dfrac{m_a(A)}{m_s(A, \Sigma); m_m(M)} & \stackrel{X}{\longrightarrow} & \dfrac{m_a(A)}{(x \leftarrow u), m_s(A, \Sigma); m_m(M)}
\end{array}$$

It is not immediately clear that transition X maps to a rule. We will show that X maps to the R-BIND rule. First, we show that $x = u$ is in $m_a(A)$. From $x \sim u$ we know that an INITIATE has been done and unbound(x) means no WIN has yet been done; together this means that A contains $(x = u)_s$, which maps to $x = u$. Second, less(u, x) because the INITIATE requires this and its truth is never altered. Third, $x \notin \mathrm{lhs}(m_s(A, \Sigma))$ since no WIN has been done and the only way to get a centralized binding for x is through a WIN. Taken together, these three statements imply that the centralized transition maps to an R-BIND rule.

Lose, Arrive. These rules trivially map to a SKIP rule.

This proves the theorem. $\square$

6.4 Liveness

What is left is to show that the DU algorithm always terminates in a configuration that maps to a terminal configuration of the RCU algorithm. The main step is to show that the DU algorithm will always "make progress", in the sense of this section. As a corollary, it follows that the DU algorithm is deadlock-free. We first prove a small lemma about the non-progressing rules.

Lemma 1 (Finiteness of non-progressing DU execution) *Given any valid configuration d of the DU algorithm, then the number of consecutive non-progressing rules that can be reduced starting from d is finite. The resulting configuration d' satisfies $m(d') = m(d)$.*

PROOF. The proof is by induction on the length of the execution e of which d is the last configuration. We assume that the lemma holds for all configurations of e before d. We show that it holds for d by enumerating all possible executions of non-progressing rules. Consider all rules that manipulate actions based on the same variable-term pair (say, x and u). Denote a configuration in which a set of non-progressing rules is applicable by the set containing those rules. Denote by X a configuration in which no rules or only progressing rules are applicable. By inspecting the rules, we deduce the following graph of causal relationships:

$$\begin{tikzcd}
\{\text{Arrival}\} \arrow[d] & \{\text{Initiate}\} \arrow[r] \arrow[d] & \{\text{Lose}\} \arrow[r] & \text{X} \\
\text{X} & \text{X} & &
\end{tikzcd}$$

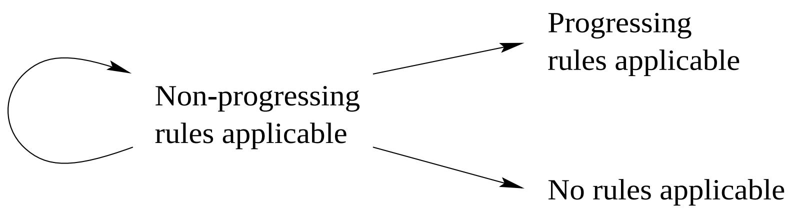

Figure 17: Non-progressing transitions in the DU algorithm

That is, applying INITIATE possibly leads to a configuration in which LOSE is applicable, and so forth. A graph identical to this one exists for all variable-term pairs. In all resulting sequences there are no cycles. Therefore a configuration in which some non-progressing rules are applicable will eventually lead to one in which no non-progressing rules are applicable.

Theorem 5 (DU liveness) *Assume that e is any execution of the DU algorithm such that the last configuration of $m(e)$ is not final in the RCU algorithm. Then the last configuration of e is not final and any execution that has e as the initial segment will eventually reduce a progressing rule in its continuation beyond e.*

PROOF. Assume a distributed execution e with last configuration d_i such that $c = m(d_i)$ is not final. We must show that $\exists j > i : d_i \to \cdots \to d_{j-1} \to d_j$ where $d_{j-1} \to d_j$ is an application of a progressing rule. Any execution starting from d_i and doing non-progressing rules as long as possible must initially follow the state diagram of Figure 17. Applying the lemma, we can assume that no non-progressing rules are applicable in d_{j-1}. It remains to show that a progressing rule is always applicable there. We do a case analysis over the RCU rules. Let the RCU configuration be c, so that $m(d_{j-1}) = c$. For each rule, we apply the inverse of mapping m, and we attempt to infer whether a progressing rule is applicable.

Interchange. Assume that the INTERCHANGE rule is applicable in c. Therefore $\mathrm{less}(u,x)$ holds and c contains $u = x$. For some site s, d_{j-1} contains $(u = x)_s$. Therefore the INTERCHANGE rule is applicable in d_{j-1}.

R-Memo Hit, R-Dereference. Except for the memo table, the conditions for these two rules are identical. For both R-MEMO HIT and R-DEREFERENCE, c contains $x = u$ and $x \leftarrow v$ and we know $\mathrm{less}(u,x)$, Therefore for some site s, d_{j-1} contains $(x = u)_s$. For this site, d_{j-1} contains one of $(x \leftarrow v)_s$ or $(x \Leftarrow v)_s$. The case $(x \Leftarrow v)_s$ is impossible by the lemma, since in that case one of ARRIVE or INITIATE is applicable depending on whether or not $(x \leftarrow \perp)_s$ is in the store. If $\mathit{MEM}(x = u, M_s)$ then a MEMO HIT is applicable. Otherwise, a DEREFERENCE is applicable.

R-Bind. Both $\mathrm{less}(u,x)$ and $x \notin \mathrm{lhs}(\sigma)$ hold. For some site s, d_{j-1} contains $(x = u)_s$. This site must also contain $(x \leftarrow \perp)_s$, since otherwise an INITIATE is applicable. Since $x \notin \mathrm{lhs}(\sigma)$, we know $\mathrm{unbound}(x)$, so the $x \sim u$ of the INITIATE still exists and a WIN rule is applicable.

Identify, Conflict, Decompose. These are straightforward.

This proves the theorem. □

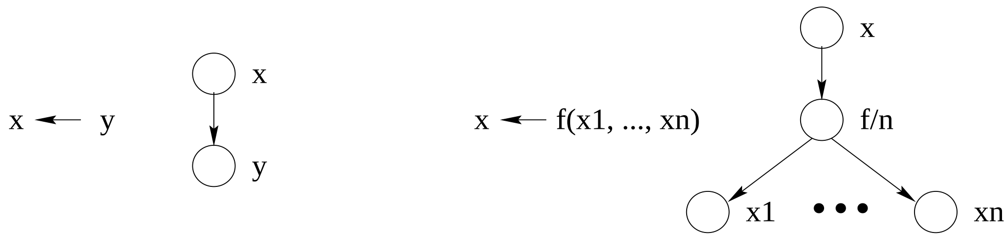

Figure 18: Mapping the store to its graph

6.5 Correctness

Theorem 6 (DU correctness) *Given any finite set of equations, then placing them on arbitrary sites and executing the DU algorithm will terminate and result in a configuration that maps to a configuration equivalent to that of a terminating CU execution.*

PROOF. From DU safety, any results obtained are correct results for the RCU algorithm. From the finiteness lemma and DU liveness, the DU algorithm will terminate and reach this correct result. From RCU correctness, the result is equivalent to the result of a terminating CU execution. □

7 On-line correctness

In a real system, it will never be the case that the unification operation is initiated with a fixed set of equations and runs to termination without any interaction with the external environment. Rather, the algorithm will be running indefinitely, and from time to time an equation will be added to the current action. This is the *on-line* case. The interesting property is not termination, but whether the equation will be entailed by the current store in a finite number of reductions (finite entailment). In this section, we extend the CU and DU algorithms to the on-line case, and we show that the extended algorithms satisfy the finite entailment property. We use the standard weak fairness assumption that any rule instance applicable infinitely often will eventually be reduced. We show that this is not enough to guarantee that the equation will be entailed, but that we need an additional property, the finite size property, to bound the amount of work needed to incorporate the equation in the store.

7.1 On-line CU and DU algorithms

We extend the CU algorithm (from now on called the *off-line* CU algorithm) with a new rule:

$$\textsc{Introduction} \qquad \frac{\text{true}}{\sigma;\ \mu} \Bigg\| \frac{u = v}{\sigma;\ \mu}$$

This rule is always applicable and adds a new equation to the action when it reduces. The extended algorithm, also called the *on-line* CU algorithm, therefore does not terminate. We extend the DU algorithm in a similar way by an INTRODUCTION rule that introduces $(u = v)_s$ for an arbitrary site s.

7.2 Finite size property

Any store σ can be mapped to a graph with two kinds of nodes, variables and records. The graph is defined in a straightforward way from the store's bindings (see Figure 18):

- A binding $x \leftarrow y$ maps to variable nodes x and y, with a directed edge from x to y.
- A binding $x \leftarrow f(x_1, ..., x_n)$ maps to a set of variable nodes x, x_1, ..., x_n and a record node f/n, with directed edges from x to f/n and from f/n to every x_1, ..., x_n.

Given a variable node x, we define graph(x, σ) as the subgraph of σ's graph whose nodes and edges are reachable from x. We also define size(x, σ) as the number of edges in this subgraph. This quantifies the size of the data structure attached to x.

Given a valid configuration with store σ, it is clear that size(x, σ) is finite. However, the size may be unbounded when considering all configurations of a given execution. This leads us to define the following property. We say a variable x has the *finite size property* for the execution e if:

$$\exists n \geq 0 : \forall(\_; \sigma_k; \_) \in e : \text{size}(x, \sigma_k) \leq n$$

That is, there is a bound on the size of x that holds for the whole execution. The finite size property is used to avoid infinite executions caused by race conditions in two cases:

1. Dereference chains that indefinitely increase in length. For example, consider the equation $x_0 = y_0$, which is accompanied by the infinite sequence of pairs of equations $x_i = x_{i+1}$ and $y_i = y_{i+1}$, starting with $i = 0$. These equations will be introduced by the Introduction rule at the appropriate times. We assume that the ordering condition enforces that lower-indexed variables are bound to higher-indexed variables. For each i starting with 0, if $x_i = x_{i+1}$ and $y_i = y_{i+1}$ are both introduced and bound before $x_i = y_i$ is dereferenced, then the store will never entail $x_0 = y_0$.

2. Nested terms that indefinitely increase in depth. For example, consider the equation $x_0 = y_0$, which is accompanied by the equations $x_i = f(x_{i+1})$ and $y_i = f(y_{i+1})$, starting with $i = 0$. For each i starting with 0, if $x_i = f(x_{i+1})$ and $y_i = f(y_{i+1})$ are both introduced and bound before $x_i = y_i$ is decomposed, then the store will never entail $x_0 = y_0$.

It is remarkable that these two infinite executions are possible even with the weak fairness assumption. One way to avoid infinite executions would be to give the Introduction rule lower priority than the others, i.e., as long as another rule is applicable, do not reduce an Introduction rule. But this does not model the real world, in which equations can arrive at any time. The finite size property does not have this deficiency. It does not restrict in any way *when* new equations are introduced. Rather, it forbids introducing any sequence of equations that would cause a problem.

The finite size property can be enforced easily for dereference chains by requiring that all new variables have higher order than all existing variables. The length of all dereference chains traversed to add an equation to the store is then bounded by the number of variables in the system when the equation is introduced.

In the case of nested structures, the finite size property can be enforced by avoiding unifying terms whose nesting depth is potentially unbounded. This seems to be a reasonable condition because when a potentially infinite unification is necessary in practice, then it is sufficient that it always makes progress, not that it completes. The weak fairness assumption is enough by itself to guarantee progress of infinite unifications and eventual termination of finite unifications. The finite size property guarantees that a unification that is intended by the programmer to be finite will actually be finite during the execution. These two conditions suffice for all practical programs we know of.

7.3 Correctness

To set the stage we first define the descendant sets of an equation and we generalize the graph and size functions to descendant sets.

Descendant set. A descendant of an equation is an action obtained from the equation after performing a sequence of rule reductions. Consider an execution e and $u = v \in \alpha_1$, where α_1 is the initial action of e. Then the multiset D_i is a *descendant set* of equation $u = v$ in execution e if it is either $D_1 = \{u = v\}$, or it is derived from D_{i-1} in the following way. Consider the transition $c_{i-1} \rightarrow c_i$, which reduces the primitive action a by means of a rule R. If $a \in D_{i-1}$ then D_i is derived according to the following rules:

$$
\begin{array}{lcl}
D_i & = & D_{i-1} - \{u = x\} + \{x = u\} \text{ if R=}\textsc{Interchange} \text{ and } a \equiv u = x \\
D_i & = & D_{i-1} - \{a\} \text{ if R=}\textsc{Bind}, \textsc{Memo Hit}, \text{ or } \textsc{Identify} \\
D_i & = & D_{i-1} - \{x = u\} + \{v = u\} \text{ if R=}\textsc{Dereference} \text{ and } a \equiv x = u \\
D_i & = & D_{i-1} - \{a\} + \sum_{1 \le j \le n} \{x_j = y_j\} \text{ if R=}\textsc{Decompose} \text{ and } a \equiv t_1 = t_2 \\
D_i & = & D_{i-1} - \{a\} + \{\text{false}\} \text{ if R=}\textsc{Conflict}
\end{array}
$$

Otherwise, if $a \notin D_{i-1}$ or R=INTRODUCTION, then $D_i = D_{i-1}$. For the proofs below we will reason about descendant sets derived from equations $u = v$ introduced in arbitrary configurations, not just in the initial configuration. Without loss of generality, we assume that $u = v \in \alpha_1$, so that $D_i \subseteq \alpha_i$. It suffices to remove the INTRODUCTION rule that introduces $u = v$ in configuration c and to wait until c before considering $u = v$.

Descendant sequence. We say D_1, D_2, ..., D_i, ... as defined above is a *descendant sequence* of equation $u = v$ in execution e.

Lemma 2 (Reachable nodes) *Given a store σ and a descendant sequence D_1, D_2, ..., D_i, ... for the execution e. Then in the graph of σ, every variable node of D_i is reachable from at least one variable node of D_1.*

PROOF. By induction on i using case analysis over all rules. □

Graph and size of a descendant set. We define graph(D, σ) to be the subgraph of σ's graph whose nodes and edges are reachable by at least one variable occurrence in D. We define size(D, σ) to be the number of edges in graph(D, σ). With this definition, we generalize the finite size property to descendant sets.

Theorem 7 (Finite entailment of on-line CU) *Given any execution e of the on-line CU algorithm satisfying weak fairness and with empty initial configuration (true;$\emptyset$;$\emptyset$). Assume that in some configuration an equation $u = v$ is introduced that has the finite size property for e. Then eventually the descendant set of $u = v$ contains only false actions. If the set is empty, then $u = v$ is entailed by the store.*

PROOF. We prove the theorem in two steps. Given e, we consider the descendant sequence D_1, D_2, ..., D_i, ... where $D_1 = \{u = v\} \subseteq \alpha_1$. Let $S_i = \text{size}(D_i, \sigma_i)$. First we use the finite size property to show that there exists a configuration from which point onwards there is no more interference from other reductions. Second, we show that there exists a mapping from the sequence $D_j, D_{j+1}, ...$ to a CU execution e'. The result then follows from properties of the CU algorithm.

Assume an on-line CU execution e such that $u = v \in \alpha_1$. Consider the transition $c_{i-1} \rightarrow c_i$ in e, which reduces rule R with action a. Then we have the following possibilities:

- If R is not BIND then $S_i \le S_{i-1}$.
- If R is BIND of $x = w$ and $x \notin \text{graph}(D_i, \sigma_{i-1})$ then $S_i \le S_{i-1}$. From $\sigma_i = \sigma_{i-1} \cup \{x \leftarrow w\}$ and $x \notin \text{graph}(D_i, \sigma_{i-1})$ it follows that $x \notin \text{graph}(D_i, \sigma_i)$.
- If R is BIND of $x = w$ and $x \in \text{graph}(D_i, \sigma_{i-1})$ then $S_i > S_{i-1}$. We now show that given the finite size property, this case can only occur a finite number of times. Assume that $\text{size}(D_i, \sigma_i) > \text{size}(D_{i-1}, \sigma_{i-1})$. We know that $\sigma_i = \sigma_{i-1} \cup \{x \leftarrow w\}$. Therefore x must be

a leaf node of graph(D_{i-1}, σ_{i-1}). From the lemma, we know that x is also a leaf node of graph(D_1, σ_{i-1}). Therefore, size$(D_1, \sigma_i) >$ size(D_1, σ_{i-1}). Given the finite size property for e, this can only happen a finite number of times.

So there is a configuration $c_j \in e$ such that $\forall k > j$: if a binding is done to an equation $x = w$ of D_{k-1} then $x \notin$ graph(D_k, σ_k). By case analysis over all rules, we deduce that graph$(D_k, \sigma_k) \subseteq$ graph(D_{k-1}, σ_{k-1}). This means that there is no more interference from other reductions. In the case analysis only BIND needs the condition $x \notin$ graph(D_k, σ_k).

Now from the execution e we construct the sequence of configurations c'_i for $i \geq 1$, where $c'_i = (\alpha'_i; \sigma'_i; \mu'_i)$ with $\alpha'_i = D_{j+i-1}$, $\sigma'_1 = \sigma_j$, and $\mu'_1 = \mu_j$. This sequence of configurations corresponds to a valid execution of the CU algorithm augmented with a SKIP rule. Removing the SKIP reductions gives a CU execution, which terminates. Because of weak fairness, it follows that in execution e either the store eventually entails $u = v$, i.e., $\exists(\_; \sigma_k; \_) \in e : \varepsilon(\sigma_k) \rightarrow u = v$, or there is eventually a false action, i.e., $\exists(\alpha_k; \_; \_) \in e : \varepsilon(\alpha_k) \rightarrow \text{false}$. □

We now extend this result to the distributed case. First we extend the DU algorithm to an on-line DU algorithm by an INTRODUCTION rule that introduces an equation on any site. It is easy to see that safety continues to hold. We now show liveness and finite entailment for the on-line DU algorithm.

Theorem 8 (Liveness of on-line DU) *Given weak fairness and any distributed execution e of the on-line DU algorithm such that $m(e)$ is nonterminal, then continuing e will always eventually reduce a progressing rule.*

PROOF. Minor modification of the proof of DU liveness, using weak fairness to compensate for the INTRODUCTION rule. □

Theorem 9 (Finite entailment of on-line DU) *Given any execution e of the on-line DU algorithm satisfying weak fairness and with empty initial configuration (true;$\emptyset$;$\emptyset$). Assume that in some configuration an equation $(u = v)_s$ is introduced that has the finite size property for e. Then eventually the descendant set of $(u = v)_s$ contains only false$_{s'}$ actions. If the set is empty, then $u = v$ is entailed by the store.*

PROOF. We reason on the corresponding centralized execution. By Theorem 7, the result holds for the corresponding centralized execution. Since advancing the distributed execution also advances the centralized execution (by liveness), therefore the result holds also for the distributed execution. □

8 The Distributed Oz implementation

Distributed Oz implements a refined version of the on-line DU algorithm, called "Oz algorithm" in what follows. Section 8.1 summarizes how the implementation differs with respect to the on-line DU algorithm. Section 8.2 introduces the distribution graph, which is the framework in which the Oz algorithm is defined. Then Section 8.3 defines the properties of the network and the notation used to define the distributed algorithm. After these preliminaries, the algorithm itself is defined. Section 8.4 defines the local algorithm and Section 8.5 defines the distributed algorithm.

8.1 Differences with on-line DU

The Oz algorithm makes concrete decisions regarding several aspects left open by the on-line DU algorithm. Furthermore, the Oz algorithm does several optimizations and adds a model for failure detection and handling. This section summarizes these refinements, optimizations, and extensions.

8.1.1 Refinements

Separation into local and distributed algorithms. The Oz algorithm consists of two parts: a purely local algorithm (corresponding to the DU non-bind rules, see Section 8.4) and a distributed algorithm (corresponding to the DU bind rules, see Section 8.5). A thread wishing to tell an equation invokes the local algorithm. To bind a distributed variable to another one or to a record, the local algorithm invokes the distributed algorithm. The thread blocks, waiting for a reply. When the variable binding is known locally, then the thread continues.

The manager site. Each distributed variable is managed from a special site, namely the variable's creation site. This site contains the variable's unbound/bound flag and other information, e.g., the register list (see below). The distributed variable is dependent on this site.

Variable ordering. The Oz algorithm implements the order relation less(u, v) as follows. Records are less than distributed variables are less than local variables. Distributed variables are totally ordered, local variables are totally ordered per site, and records are unordered. Local variables are ordered according to a per-site index i that is incremented for each new variable.[4] Distributed variables are ordered according to a pair (s, i) where s is the site number on which the variable was initially created and i is the index of the variable on that site. From this ordering relation it follows that if the number of sites is finite and data structures with unbounded depth are not created, then the Oz algorithm satisfies the finite size property (see Section 7.2).

8.1.2 Optimizations

Globalization. For efficiency reasons, the Oz algorithm distinguishes between local and distributed variables (see Section 8.5.1).

Variable registration. For efficiency reasons, a variable binding is not sent to all sites, but only to registered sites (see Section 8.5.2).

Grouping nested data structures. For efficiency reasons, binding a nested data structure to a distributed variable is done by the Oz algorithm as a single operation (see Section 8.5.3).

Winner optimization. When a variable is bound to a term, then the term does not have to be sent back to the site that initiated the binding (see Section 8.5.4).

Asynchronous streams. To allow streams to be created asynchronously, variables are given a set of registered sites as soon as they are globalized (see Section 8.5.5).

8.1.3 Extensions

The Oz algorithm extends on-line DU in the following ways.

Lazy and eager variables. This property affects when the variable is registered. Eager proxies are registered immediately. Lazy proxies delay registration until a binding attempt is made (see Section 8.5.6).

Read-only logic variables. Standard logic variables have two operations, reading the value and binding. For security reasons, it is often useful to remove the right to bind the variable, for example, when passing the variable to a less-trusted site.

[4] For local variables, the index is simply the variable's address.

Garbage collection. Distributed garbage collection is based on a credit mechanism that collects all garbage except cross-site cycles between stateful entities (see Section 8.2.1).

The failure model. The Oz algorithm is conservatively extended with a basic failure model that reflects network and site failures to the language level (see Section 8.2.2).

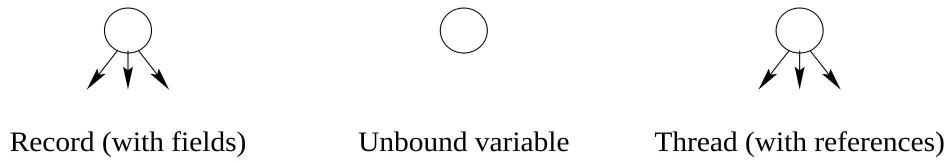

Figure 19: The three node types of the language graph

8.2 The distribution graph

We model distributed executions in a simple but precise manner using the concept of *distribution graph*. We obtain the distribution graph in two steps from an arbitrary execution state of the system. The first step is independent of distribution. We model the execution state by a directed graph, called *language graph*, in which a record, an unbound variable, and a thread each correspond to one node (see Figure 19). The edges in this graph denote the node's references: a record and a thread refer to other nodes; an unbound variable has no references.

In the second step, we distribute the execution over a set of sites. Assume a finite set of sites and annotate each node by its site (see Figure 20). If a variable node, e.g., `N2`, is referenced by at least one node on another site, then map it to a *set* of nodes, e.g., {`P1`,`P2`,`P3`,`M`}. This set is called the *access structure* of the original node. An access structure consists of one *proxy node* `Pi` for each site that referenced the original node and one *manager node* `M` for the whole structure. The resulting graph, containing both local nodes and access structures where necessary, is called the *distribution graph*. The execution of the distributed algorithm is defined in terms of this graph.

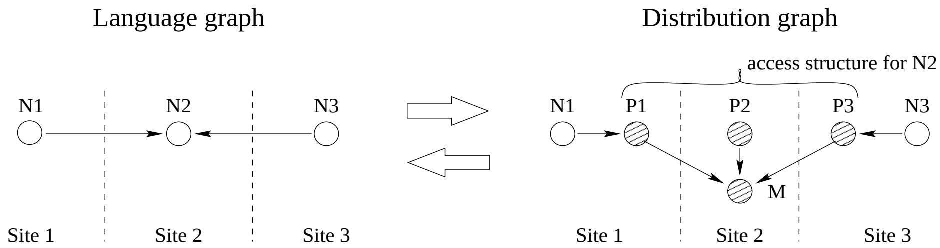

Figure 20: An access structure in the distribution graph

Each access structure is given a global name n that is unique system-wide. In Distributed Oz, n is the pair (s, i) that is also used to order the distributed variables. The global name n encodes (among other things) the manager site s. Furthermore, a proxy node is uniquely identified by the pair (n, s'), which contains the proxy's site s'. On each site, n indexes into a table that refers to the proxy. This allows to enforce the invariant that each site has at most one proxy.

Messages are sent between nodes in access structures. In terms of sites, a message is sent from the source node's site to the destination node's site. The message may contain a subgraph of the distribution graph. Just before the message leaves the source site, a new access structure is created for each local variable in the subgraph. The message refers only to proxy nodes, not to local variables. When the message arrives, the subgraph becomes known at the destination site. Each proxy node is looked up in the site table. If it is not present, then a new proxy node is created and entered in the table. This extends an existing access structure with one new proxy. The process of creating or extending an access structure is called *globalization*.

The behavior of a distributed variable is defined as a protocol between the nodes of its access structure. In general, nodes other than variable nodes can also have access structures, and therefore obey a protocol. The Distributed Oz implementation uses four non-trivial protocols. Three are designed for specific language entities, namely variables, object records, and state pointers. Variables use a *variable binding* protocol, which is part of the distributed unification algorithm and is presented in this article. Object records use a *lazy replication* protocol. State pointers use a *mobile state* protocol. The fourth protocol is a distributed garbage collection algorithm using a credit mechanism. Garbage collection is part of the management of access structures, and it therefore underlies the other three protocols. See [3, 38, 37, 13] for more information.

8.2.1 Distributed garbage collection

Distributed garbage collection is based on a credit mechanism, a variant of reference counting [25]. The credit mechanism interfaces with the local garbage collectors on each site. All distributed garbage is removed except for cross-site cycles between stateful entities.

The global name of an access structure is associated with a pool of *credits*. The manager site lends credit to sites and messages that refer to the global name. A manager site has an integer corresponding to the total number of credits lent. A proxy site must hold at least one credit for the access structure. The proxy site keeps a count of how many credits it has borrowed from the manager site. If local garbage collection removes the proxy node, then its credit is returned to the manager site. The global name can be reclaimed at the manager site when no more borrowed credits exist. The manager node itself can be reclaimed if it has both no global name and no local references.

8.2.2 The failure model

The failure model is designed according to the principle that a program should be able to make all decisions regarding failure behavior [13]. That is, the implementation does not make any irrevocable decisions by default. Full treatment of the model is beyond the scope of this paper. We briefly summarize the main ideas. The failure model detects permanent site failures and temporary and permanent network failures. The failure model has two effects at the language level:

- It extends each operation on an entity to have a transactional semantics. That is, when an operation is invoked, either it succeeds or it aborts. If it aborts then an exception is raised. It aborts if the manager site fails or if there is a network time-out. The system ensures that if the network failure is temporary, then the program can retry the operation.

- It allows to post "watchers" that are invoked by the system as soon as particular error conditions are detected. Watchers are explicitly installed to watch for given error conditions. These conditions do not necessarily keep the entity from working, for example, a remote proxy site failure will often have no effect at all, but it may nonetheless be important to be notified of this failure.

This failure model is part of the Distributed Oz implementation. We are currently developing more powerful abstractions in Oz on top of this model.

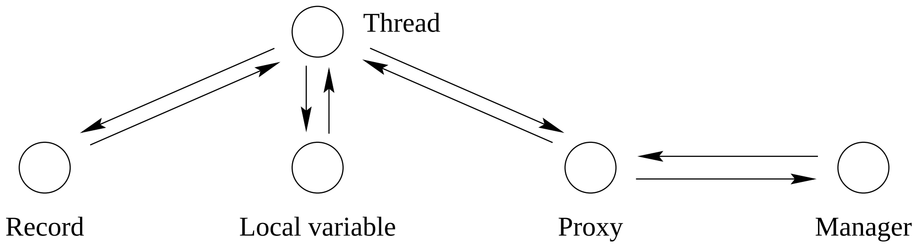

Figure 21: The five node types of the distribution graph and their message interfaces

8.3 Basic concepts and notation

8.3.1 Network

Consider a single manager node M, a set of k proxy nodes P_i with $1 \leq i \leq k$, and a set of m thread nodes T_i with $1 \leq i \leq m$. All nodes have state and interact according to Figure 21. Thread, proxy, and manager nodes send messages to each other and perform internal operations. Record and local variable nodes interact only with thread nodes. Let these nodes be linked together by a network N that is a multiset containing messages of the form $d : m$ where d identifies a destination (proxy, manager, or thread node) and where m is a message.

The Oz algorithm is defined using reduction rules of the form

$$\frac{\text{Condition}}{\text{Action}} \quad .$$

Each rule is defined in the context of a single node. Execution follows an interleaving model. The local algorithm imposes an order on how its rules are fired; see the pseudocode definition of Section 8.4. The distributed algorithm imposes no such order.

At each reduction step, a rule with valid condition is selected. Its associated actions are reduced atomically. A rule condition consists of boolean conditions on the node state and one optional receive condition **Receive**(d,m). The condition **Receive**(d,m) means that $d : m$ has arrived at d. Executing a rule with a receive condition removes $d : m$ from the network and performs the action part of the rule. A rule action consists of a sequence of operations on the node state with optional sends. The action **Send**(d,m) asynchronously sends message m to node d, i.e., it adds the message $d : m$ to the network.

We assume that the network and the nodes are *fair* in the following sense. The network is asynchronous, and messages to a given node take arbitrary finite time and may arrive in arbitrary order. All rules that are applicable infinitely often will eventually reduce.

8.3.2 Node state

Table 3 defines the state of the five node types by listing the attributes of each node. All nodes have attributes "id" and "type", which have constant values. The types NodeId, Atom, and Integer are atomic types implemented in a straightforward way. In a real system, threads, proxies, and managers would have more attributes, for example, threads have an execution state and proxies and managers maintain information for distributed garbage collection.

8.3.3 Utility operations

The memo table uses the function newmemo(), the procedure add(N_1,N_2,M) and the boolean function mem(N_1,N_2,M). The latter two are exactly the *ADD* and *MEM* operations defined in Section 3.2. The other operations are defined as follows:

Node	Attribute	Type
Any node	id	NodeId
	type	{RECORD,LOCVAR,PROXY,MANAGER,THREAD}
Record	label	Atom
(N.type=RECORD)	arity	Integer
	args	**array**[1..arity] **of** Node
Local variable	state	{UNBOUND, BOUND(Node)}
(N.type=LOCVAR)	eager	{FALSE, TRUE}
Proxy	state	{UNBOUND,INITIATED,BOUND(Node)}
(N.type=PROXY)	eager	{FALSE, TRUE}
	reg	{FALSE, TRUE}
	manager	NodeId
Manager	state	{UNBOUND,BOUND(Node)}
(N.type=MANAGER)	reglist	**set of** NodeId
Thread		
(N.type=THREAD)		

Table 3: Node state

$$
\begin{array}{lcl}
\mathrm{eq}(\mathrm{N}_1,\mathrm{N}_2) & = & \mathrm{N}_1.\mathrm{id}=\mathrm{N}_2.\mathrm{id} \\
\mathrm{locvar}(\mathrm{N}) & = & \mathrm{N.type{=}LOCVAR} \\
\mathrm{disvar}(\mathrm{N}) & = & \mathrm{N.type{=}PROXY} \\
\mathrm{nonvar}(\mathrm{N}) & = & \mathrm{N.type{=}RECORD} \\
\mathrm{var}(\mathrm{N}) & = & \mathrm{locvar}(\mathrm{N}) \vee \mathrm{disvar}(\mathrm{N}) \\
\mathrm{bound}(\mathrm{N}) & = & \textbf{if}\ \mathrm{var}(\mathrm{N}) \rightarrow \mathrm{N.state{=}BOUND}(\_)\ \textbf{fi} \\
\mathrm{initiated}(\mathrm{N}) & = & \textbf{if}\ \mathrm{disvar}(\mathrm{N}) \rightarrow \mathrm{N.state{=}INITIATED}\ \textbf{fi} \\
\mathrm{deref1}(\mathrm{N}) & = & \textbf{if}\ \mathrm{var}(\mathrm{N}) \wedge \mathrm{bound}(\mathrm{N}) \rightarrow \\
 & & \quad \mathrm{N}_1\ \textbf{where}\ \mathrm{N.state{=}BOUND}(\mathrm{N}_1)\ \textbf{fi} \\
\mathrm{compatible}(\mathrm{N}_1,\mathrm{N}_2) & = & \textbf{if}\ \mathrm{nonvar}(\mathrm{N}_1) \wedge \mathrm{nonvar}(\mathrm{N}_2) \rightarrow \\
 & & \quad \mathrm{N}_1.\mathrm{arity}=\mathrm{N}_2.\mathrm{arity} \wedge \mathrm{N}_1.\mathrm{label}=\mathrm{N}_2.\mathrm{label}\ \textbf{fi} \\
\mathrm{proxyids}(\mathrm{N}) & = & \textbf{if}\ \mathrm{disvar}(\mathrm{N}) \rightarrow \{\mathrm{N.id}\} \\
 & & [\,]\ \mathrm{locvar}(\mathrm{N}) \rightarrow \{\} \\
 & & [\,]\ \mathrm{nonvar}(\mathrm{N}) \rightarrow \bigcup_{1\le i\le \mathrm{N.arity}} \mathrm{proxyids}(\mathrm{N.args}[i])\ \textbf{fi}
\end{array}
$$

8.4 The local algorithm

Figure 22 defines the local algorithm, which executes in each thread that does a unification. The definition follows closely the non-bind rules of Section 5.3.1, where a rule corresponds to a guard and its body. The two main differences are that the local algorithm maintains a memo table that is shared among the rules and that a sequential order is imposed on rule reductions. The **if** is a guarded command that suspends until at least one of its guards is true.

In the implementation, executions of the local and distributed algorithms are not interleaved. Rather, the local algorithm is executed atomically until it exits (either normally or through an exception) or until it blocks when waiting for a binding request to complete.

Each invocation of unify is done within a thread. The unification completes in two ways: when it terminates normally or when an inconsistency is detected, in which case a failure exception is raised. During the unification, the thread will block when waiting for a binding request to complete. This works as follows. The first thread that tries to bind a variable will send a binding request (INITIATE rule). All other threads on the same site that try to bind the same variable do not send any message. Both the first and the other threads then block at the **if**: no guard is true because the variable satisfies disvar(N_1) $\wedge$ **not** bound(N_1) $\wedge$ initiated(N_1). As soon as the binding arrives, bound(N_1) is true and the **if** becomes reducible. All threads can then continue.

The local algorithm is optimized to bind local variables locally. With a suitable data representation, the algorithm is implementable very efficiently [10, 36]. The implementation further optimizes the definition by doing binding in place and dereferencing inline. Common cases of unification such as parameter passing and local variable initialization do not need binding nor dereferencing, but reduce to single register moves or stores.

In the implementation, the local memo table is implemented by means of *forwarding pointers* between variables. That is, when the equation $x = y$ is encountered for the first time, a forwarding pointer is installed from x to y. This allows a very fast check of memo table membership. Namely, if $x = y$ or $y = x$ is encountered later on, then dereferencing will reduce the equation to $y = y$, which does no further work. The forwarding pointers are installed in the context of a single atomic unification operation. To maintain consistency of the distribution graph, the forwarding pointers must be removed whenever execution leaves the local algorithm. Therefore the memo table lives only within the context of a contiguous sequence of rules executing within the local algorithm and not sending any messages. This is modeled in Figure 22 by emptying the memo table at each invocation of unify and after doing an INITIATE rule. This means that the Oz algorithm potentially does more redundant work than the on-line DU algorithm, because in the latter the local memo tables are never emptied.

8.5 The distributed algorithm

Figure 23 defines the distributed algorithm, which extends the DU bind rules of Section 5.3.2 with globalization and variable registration. The implementation does three other important optimizations, namely grouping nested data structures, the winner optimization, and asynchronous streams. For clarity, we do not define the latter formally, but rather show how to extend the protocol to include them. We also explain how to extend the protocol for lazy and eager variables.

8.5.1 Globalization

Newly-created variables are always local. When a message is sent referencing a local variable, then a new distributed variable is created and the local variable is bound to it. This is called *globalizing* the local variable. An access structure is created when a local variable is globalized. When the message arrives then a new proxy will be created for the distributed variable if none exists on the arrival site. Therefore globalization is part of the **Send** and **Receive** operations [3].

All variables have a boolean attribute "eager" that determines whether the variable is eager or lazy. This attribute affects only the network operations of the distributed algorithm. Assume we have a local variable L with L.state=UNBOUND and L.eager=b. After globalizing, the original site contains three nodes, L, P, and M, with the following states:

L.state=BOUND(P)
P.state=UNBOUND, P.eager=b, P.reg=b, P.manager=M.id
M.state=UNBOUND, M.reglist=**if** b **then** {P.id} **else** {} **endif**

8.5.2 Variable registration

In the DU algorithm, a binding arrives on a site if the variable exists in the site's store. In the Oz algorithm, the variable's manager keeps track of the sites that reference the variable. A site that receives a distributed variable for the first time (i.e., when a term containing the variable first arrives on the site) has to register with the manager in order to receive the variable's binding. In Figure 23, first the ARRIVE rule reduces, which sends reg messages to all proxies in the binding. The reg message causes all unregistered proxies to register with their manager. When a variable is bound, then a binding_in_transit message is sent to all registered sites. If the variable is already bound when the register message arrives then the binding is sent back immediately.

```
procedure unify(N1,N2,TId)
   define memotable M
   define procedure inner_unify(N1,N2)
      if /**** INTERCHANGE ****/
         var(N2), less(N1,N2) → inner_unify(N2,N1)
      [] /**** IDENTIFY ****/
         var(N1), eq(N1,N2) → skip
      [] /**** MEMO HIT ****/
         var(N1), less(N2,N1), bound(N1), mem(N1,N2,M) → skip
      [] /**** DEREFERENCE ****/
         var(N1), less(N2,N1), bound(N1), not mem(N1,N2,M) →
         add(N1,N2,M)
         inner_unify(deref1(N1),N2)
      [] /**** BIND ****/
         locvar(N1), less(N2,N1), not bound(N1) →
         N1.state ← BOUND(N2)
      [] /**** INITIATE ****/
         disvar(N1), less(N2,N1), not bound(N1), not initiated(N1) →
         N1.state ← INITIATED
         Send(N1.manager,binding_request(N2))
         newmemo(M)
         inner_unify(N1,N2)
      [] /**** DECOMPOSE ****/
         nonvar(N1), nonvar(N2), compatible(N1,N2) →
         for i:=1 to N1.arity do inner_unify(N1.args[i],N2.args[i])
      [] /**** CONFLICT ****/
         nonvar(N1), nonvar(N2), not compatible(N1,N2) →
         raise failure_exception
      endif
   end
in
   newmemo(M)
   inner_unify(N1,N2)
end
```

Figure 22: Distributed unification part 1: Local algorithm

/**** WIN ****/

$$\frac{\textbf{Receive}(\text{M.id,binding\_request(N)}) \;\wedge\; \text{M.state=UNBOUND}}{\begin{array}{l}\forall i \in \text{M.reglist: } \textbf{Send}(i, \text{ binding\_in\_transit(N)}) \\ \text{M.state} \leftarrow \text{BOUND(N)}\end{array}}$$

/**** LOSE ****/

$$\frac{\textbf{Receive}(\text{M.id,binding\_request(\_)}) \;\wedge\; \text{M.state=BOUND(\_)}}{\textbf{skip}}$$

/**** ARRIVE ****/

$$\frac{\textbf{Receive}(\text{P.id,binding\_in\_transit(N)}) \;\wedge\; (\text{P.state=UNBOUND} \vee \text{P.state=INITIATED})}{\begin{array}{l}\forall i \in \text{proxyids(N): } \textbf{Send}(i, \text{ reg}) \\ \text{P.state} \leftarrow \text{BOUND(N)}\end{array}}$$

/**** Variable registration ****/

$$\frac{\textbf{Receive}(\text{P.id,reg}) \;\wedge\; \text{P.reg=FALSE}}{\begin{array}{l}\text{P.reg} \leftarrow \text{TRUE} \\ \textbf{Send}(\text{P.manager,register(P.id)})\end{array}}$$

$$\frac{\textbf{Receive}(\text{P.id,reg}) \;\wedge\; \text{P.reg=TRUE}}{\textbf{skip}}$$

$$\frac{\textbf{Receive}(\text{M.id,register(PId)}) \;\wedge\; \text{M.state=UNBOUND}}{\text{M.reglist} \leftarrow \text{M.reglist} \cup \{\text{PId}\}}$$

$$\frac{\textbf{Receive}(\text{M.id,register(PId)}) \;\wedge\; \text{M.state=BOUND(N)}}{\textbf{Send}(\text{PId,binding\_in\_transit(N)})}$$

Figure 23: Distributed unification part 2: Distributed algorithm

8.5.3 Grouping nested data structures

The DU algorithm binds only a single record, namely the top level of the structure, and the rest is transferred when the top level binding arrives at a site. The Oz algorithm binds a complete tree in a single operation. In this way, it avoids the creation of distributed variables for the intermediate nodes. For example, the unification $x_1 = f(g(a))$, is represented in the DU algorithm as three actions $x_1 = f(x_2) \wedge x_2 = g(x_3) \wedge x_3 = a$. In the DU algorithm, the arrival of $x_1 \Leftarrow f(x_2)$, enables the arrival of $x_2 \Leftarrow g(x_3)$, and similarly for x_3. In the Oz algorithm, the binding $x_1 \Leftarrow f(g(a))$ arrives in one step, so the variables x_2 and x_3 are never created.

8.5.4 Winner optimization

The winner is the proxy that sent a successful binding_request(N). This proxy does not need to be sent N since it already exists on the proxy's site. The proxy can be sent a simple acknowledgement that its binding request was successful.

This optimization avoids the redundant work done by the R-BIND rule (see Section 6.2). It requires the following protocol extensions: the proxy state INITIATED becomes INITIATED(N), the modified message binding_request(N,PId) where PId identifies the winning proxy, and the new message binding_ack from the manager to the winning proxy. When the proxy receives binding_ack, then it retrieves N from the INITIATED(N) state.

8.5.5 Asynchronous streams

A variable that is exported from its manager site can be *preregistered*. That is, the destination site is added to the manager's reglist without waiting for a registration message. This is correct if there is a FIFO connection to the destination site. Preregistering variables allows elements to be added to streams asynchronously. The example of Section 2.4.2 relies on this behavior.

Let us look closely to see what happens. Assume variable `X0` exists on sites 1 and 2. Binding `X0=m1|X1` on site 1 causes `m1|X1` to be sent to site 2. `X1` will be preregistered, i.e., M_{X1}.reglist $\leftarrow M_{X1}$.reglist $\cup\, M_{X0}$.reglist when the binding leaves site 1. If `X1` is bound on site 1, then its binding will be sent immediately to site 2 without waiting for a reply.

If preregistration is not done, then adding elements to a stream requires a round trip message delay for each element. This is because remote proxies have to be registered before they can receive a binding. In our example, binding `X0=m1|X1` on site 1 causes `m1|X1` to be sent to site 2. When it arrives, an `X1` proxy is created on site 2 and promptly registers with site 1. Binding `X1=m2|X2` on site 1 will not send the binding to site 2 until the registration arrives on site 1. Therefore each new element appears on site 2 only after a round trip.

8.5.6 Lazy and eager variables

Lazy and eager logic variables are defined informally in Section 2.2.2. In terms of the on-line DU algorithm, they differ only in the scheduling of the ARRIVE rule. To be precise, laziness is a property of a variable proxy. A variable proxy is *lazy* if the reduction of ARRIVE is delayed until after INITIATE reduces on that site. If no such delay is enforced then the variable proxy is *eager*.

In terms of the Oz algorithm, this is implemented by registering lazy and eager proxies at different times. Eager proxies are registered as soon as they appear on a site (see ARRIVE rule in Figure 23). Lazy proxies are only registered after the INITIATE rule is reduced (see Figure 22), that is, when a binding request is made.

When two proxies are bound together, the result must be eager if at least one of the two was eager. When a local variable is bound to a proxy, the proxy must become eager if the local variable was eager. Implementing this requires replacing the reg message by three messages: (1) in the ARRIVE rule, reg becomes reg_if_eager, (2) in the INITIATE rule, a new message reg_always is sent, and (3) in the BIND rule, a new message reg_and_make_eager is sent if the local variable is eager.

9 Related work

As far as we know, this article presents the first complete and practical algorithm for distributed rational tree unification.[5] However, many partial implementations exist that provide some of the benefits for concurrent and distributed programming.

A basic operation in distributed unification is sending a bound term across the network. Lamma et al [19] investigate the costs of this operation and sketch an algorithm to send only that part of a term required by a consumer. Sending too little increases the message latency, since multiple requests will be done. Sending too much increases network load and memory consumption at the consumer. The proposed algorithm sends exactly that part of a term required by a consumer. For example, a list appending procedure requires only the spine of the list, and not the terms in the list. The algorithm uses "consumption specifications", simple tree grammars extended with an additional terminal `Remote`. These specifications can be given by static analysis or by programmer annotation.

9.1 Concurrent logic languages

Several concurrent logic languages have been extended to distributed settings. These systems do not use logic variables primarily to improve latency tolerance and network transparency. Rather, logic variables are integral parts of their execution models, and the distributed extensions must therefore take them into account. We summarize the distributed extensions to unification in DRL and KLIC.

DRL [8] (Distributed Real-time Logic language) is a concurrent logic language extended with features for distribution and soft real-time control. Distribution is introduced by allowing computations on different sites to communicate through shared logic variables. In DRL, the representative of a logic variable on a site is called a *logic channel*. A logic channel is always statically marked with a direction, which is either output or input. For a given logic variable, only one channel is marked output. Binding the output channel to a term causes the term to appear at all corresponding input channels. The binding blocks until the term contains only ground subterms and logic channels. It follows that variables can be transferred between sites only if they are statically declared as such.

Logic channels can be connected together. This operation is called "unification" in DRL, but the shared logic variables are not actually unified together. To be precise, no variable elimination is done, but communication links are set up between variables. Connecting two output channels causes a future binding of one of them to be sent also to the other. Connecting an input channel to another channel suspends until the input channel receives a value. It follows that dependencies on intermediate sites are not removed.

KLIC [9] is an efficient portable implementation of the concurrent logic language KL1 (Kernel Language 1) for distributed and shared-memory machines. KLIC achieves these goals by compiling into C. On one processor running a series of representative benchmarks, the performance of KLIC approaches that of C and C++ implementations. The distributed implementation of KLIC does distributed unification [26], including binding variables to variables. However, the algorithm has several curious properties: binding cycles can be created when binding variables to variables, inconsistencies are ignored, and a variable may be bound to different values on different sites. Apparently, the algorithm is only intended to be used in settings where there is no possibility of inconsistency.

9.2 Languages not based on logic

Systems with support for distributed computing commonly provide two-level addressing, namely the ability to transparently pass references between sites. Typical examples include Java

[5] To the referees: we would be happy to be corrected on this point!

RMI [22], CORBA [24], and the Ericsson OTP (Open Telecom Platform) [4, 39]. This basic mechanism can be extended to provide weak logic variables (see also Section 2.1.4): (1) if the value is available, it is should be used immediately without programmer intervention, (2) threads block on availability of the value, (3) when the value is available it is sent to all remote references, thus deblocking any blocked threads, and (4) no forwarding chains are created if a reference travels among many sites. There should be no overhead if the variable is on one site only. To provide full logic variables this must further be extended with dynamic typing and variable-variable unification.

10 Conclusions

This article has studied the use of logic variables to simplify distributed programming. We have shown that if the logic variables are well-implemented, then common distributed programming idioms can be written in a network-transparent manner, and they behave efficiently when distributed. We have defined the CU algorithm, a centralized algorithm for rational tree unification, and the DU algorithm, its conservative extension to a distributed setting. The DU algorithm replaces the centralized BIND rule by four rules that do coherent variable elimination. We show that the DU algorithm has good network behavior for common distributed programming idioms. We prove that the DU algorithm is a correct implementation of unification, and we bound the amount of extra work it can do compared to the CU algorithm. We show that both lazy and eager logic variables obey the DU algorithm. They differ only in the scheduling of a single reduction rule.

We extend both the CU and DU algorithms to the on-line case, in which new equations can be introduced indefinitely during execution. We show that if a weak fairness condition holds and if all variables in the equation satisfy the *finite size property*, then any introduced equation will eventually be entailed by the store.

An optimized version of the on-line DU algorithm, called the Oz algorithm, is part of the Mozart system, which implements the Distributed Oz language [1]. An official release of Mozart is planned for late 1998. Distributed Oz, also known as Oz 3, is an efficient distributed network-transparent extension to Oz 2 [13, 38, 14]. Oz 2 is the successor to Oz 1 [29], which was designed for fine-grained concurrency and implicit exploitation of parallelism. Oz 2 abandons this model in favor of explicit control over concurrency by means of a thread creation construct. Oz 2 has a robust centralized implementation that was officially released in February 1998 [7]. Oz 3 keeps the same language semantics as Oz 2 and extends it with support for mobile computations and orthogonal failure detection and handling within the language. Oz 2 programs are portable to Oz 3 almost immediately.

Acknowledgements

This research is partially financed by the Walloon Region. Michael Mehl implemented the distributed unification algorithm in the Mozart system. Per Brand and Erik Klintskog have extended it to incorporate orthogonal failure detection and handling. Ilies Alouini and Mustapha Hadim gave many valuable comments on this article. We thank the other members of the Distributed Oz projects at SICS, DFKI, and UCL. Finally, we thank the referees for comments that let us vastly improve the presentation.

References

[1] Mozart system (Oz 3), 1998. To be released.

[2] Mack W. Alford, Leslie Lamport, and Geoff P. Mullery. *Basic Concepts*, chapter 2. Lecture Notes in Computer Science, vol. 190. Springer Verlag, 1985.

[3] Ilies Alouini and Peter Van Roy. Le protocole réparti du language Distributed Oz (in French). September 1998. Submitted for publication.

[4] J. Armstrong, M. Williams, C. Wikström, and R. Virding. *Concurrent Programming in Erlang*. Prentice-Hall, Englewood Cliffs, N.J., 1996.

[5] Henri E. Bal, Jennifer G. Steiner, and Andrew S. Tanenbaum. Programming languages for distributed computing systems. *ACM Computing Surveys*, 21(3):261–322, September 1989.

[6] Randy Chow and Theodore Johnson. *Distributed Operating Systems and Algorithms*. Addison-Wesley, San Francisco, Calif., 1997.

[7] DFKI Oz version 2.0, February 1998. Available at `http://www.ps.uni-sb.de`.

[8] M. Diaz, B. Rubio, and J. M. Troya. DRL: A distributed real-time logic language. *Comput. Lang.*, 23(2–4):87–120, 1997.

[9] Tetsuro Fujise, Takashi Chikayama, Kazuaki Rokusawa, and Akihiko Nakase. KLIC: A portable implementation of KL1. In *Fifth Generation Computing Systems (FGCS '94)*, pages 66–79, December 1994.

[10] David Gudeman. Representing type information in dynamically typed languages. Technical Report TR93-27, University of Arizona, Department of Computer Science, September 1993.

[11] Robert H. Halstead. MultiLisp: A language for concurrent symbolic computation. *ACM Transactions on Programming Languages and Systems*, 7(4):501–538, October 1985.

[12] Seif Haridi. An Oz 2.0 tutorial. Available at `http://www.sics.se/~seif/oz.html`, 1996.

[13] Seif Haridi, Peter Van Roy, Per Brand, and Christian Schulte. Programming languages for distributed applications. *New Generation Computing*, 16(3), May 1998.

[14] Seif Haridi, Peter Van Roy, and Gert Smolka. An overview of the design of Distributed Oz. In the *2nd International Symposium on Parallel Symbolic Computation (PASCO 97)*. ACM, July 1997.

[15] Martin Henz. *Objects for Concurrent Constraint Programming*, volume 426 of *The Kluwer International Series in Engineering and Computer Science*. Kluwer Academic Publishers, Boston, November 1997.

[16] Martin Henz. *Objects in Oz*. PhD thesis, Universität des Saarlandes, Fachbereich Informatik, Saarbrücken, Germany, June 1997.

[17] Robert A. Iannucci. *Parallel Machines: Parallel Machine Languages. The Emergence of Hybrid Dataflow Computer Architectures*. Kluwer, Dordrecht, the Netherlands, 1990.

[18] Joxan Jaffar and Michael Maher. Constraint logic programming: A survey. *J. Log. Prog.*, 19/20:503–581, May/July 1994.

[19] Evelina Lamma, Paola Mello, Cesare Stefanelli, and Pascal Van Hentenryck. Improving distributed unification through type analysis. In *Euro-Par '97 Parallel Processing*, Lecture Notes in Computer Science, vol. 1300, pages 1181–1190. Springer-Verlag, 1997.

[20] Doug Lea. *Concurrent Programming in Java*. Addison-Wesley, 1997.

[21] John Lloyd. *Foundations of Logic Programming, Second Edition*. Springer-Verlag, 1987.

[22] Sun Microsystems. *The Remote Method Invocation Specification*, 1997. Available at `http://www.javasoft.com`.

[23] K. Mehlhorn and A. Tsakalidis. Data structures. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science – Volume A: Algorithms and Complexity*, pages 301–341. Elsevier MIT Press, 1990.

[24] Randy Otte, Paul Patrick, and Mark Roy. *Understanding CORBA: The Common Object Request Broker Architecture*. Prentice-Hall PTR, Upper Saddle River, N.J., 1996.

[25] David Plainfossé and Marc Shapiro. A survey of distributed garbage collection techniques. In the *International Workshop on Memory Management*, Lecture Notes in Computer Science, vol. 986, pages 211–249, Berlin, September 1995. Springer-Verlag.

[26] Kazuaki Rokusawa, Akihiko Nakase, and Takashi Chikayama. Distributed memory implementation of KLIC. *New Generation Computing*, 14(3):261–280, 1996.

[27] Vijay A. Saraswat. *Concurrent Constraint Programming*. MIT Press, 1993.

[28] Ehud Shapiro. The family of concurrent logic programming languages. *ACM Computing Surveys*, 21(3):413–510, September 1989.

[29] Gert Smolka. The Oz programming model. In *Computer Science Today*, Lecture Notes in Computer Science, vol. 1000, pages 324–343. Springer-Verlag, Berlin, 1995.

[30] Gert Smolka. Concurrent constraint programming based on functional programming. In Chris Hankin, editor, *Programming Languages and Systems*, Lecture Notes in Computer Science, vol. 1381, pages 1–11, Lisbon, Portugal, 1998. Springer-Verlag.

[31] Gert Smolka, Christian Schulte, and Peter Van Roy. PERDIO—Persistent and distributed programming in Oz. BMBF project proposal. Available at `http://www.ps.uni-sb.de`, February 1995.

[32] Bjarne Stroustrup. *The C++ Programming Language, Third Edition*. Addison-Wesley, 1997.

[33] Andreas Sundström. Comparative study between Oz 3 and Java. Technical report, Uppsala University and Swedish Institute of Computer Science, 1998.

[34] Andrew Taylor. *High-Performance Prolog Implementation*. PhD thesis, Basser Department of Computer Science, University of Sydney, June 1991.

[35] Gerard Tel. *An Introduction to Distributed Algorithms*. Cambridge University Press, Cambridge, United Kingdom, 1994.

[36] Peter Van Roy. 1983–1993: The wonder years of sequential Prolog implementation. *J. Log. Prog.*, 19/20:385–441, May/July 1994.

[37] Peter Van Roy, Seif Haridi, Per Brand, and Gert Smolka. Three moves are not as bad as a fire. In the *Workshop on Internet Programming Languages, International Conference on Computer Languages (ICCL 98)*, May 1998.

[38] Peter Van Roy, Seif Haridi, Per Brand, Gert Smolka, Michael Mehl, and Ralf Scheidhauer. Mobile objects in Distributed Oz. *ACM Transactions on Programming Languages and Systems*, 19(5):804–851, September 1997.

[39] Claes Wikström. Distributed programming in Erlang. In the *1st International Symposium on Parallel Symbolic Computation (PASCO 94)*, pages 412–421, Singapore, September 1994. World Scientific.

Contents