

The Intelligent Channel: A Scheme For Result Sharing in Logic Programs

Simon Kasf and Jack Minker
University of Maryland

1. Introduction

The separation of logic and control in logic programs allows the programmer to write programs whose execution is determined by the interpreter. This characteristic of logic programs spurred research towards diversifying the means for controlling the execution of logic programs, and towards understanding and exploring the value of parallelism in logic programming. Much of these efforts belong to the study of AND/OR-parallelism, i.e., parallel execution of conjunctive/disjunctive goals respectively.

A brute force approach to AND/OR-parallelism, i.e., executing every possible subgoal in a conjunctive goal has two major drawbacks: combinatorial explosion of processes and the need for many active binding environments. The latter arises due to the interaction of shared variables in conjunctive goals.

We propose a scheme to alleviate the above difficulties. For simplicity, we demonstrate our approach with examples where atoms share a single variable. The approach is applicable to Horn-clause logic programs [1].

2. A Model for AND/OR Parallelism

We first introduce a model that serves as a unifying basis for constructing various schemes for parallel execution of logic programs, (see [1] for details).

OR-nodes are created when a conjunctive goal of n atoms ($n > 1$) is split into n subgoals, that may be executed in parallel. When such OR-nodes are created as descendants of an AND-node, a new environment is created at each of the newly created sons. The environment of an OR-node is a set of variable templates v_i , where each variable template is a pair:

$$v_i : \langle in_i, out_i \rangle,$$

v_i is a new private variable id, in_i is the location of the current instantiation of the variable v_i , and out_i is the location where the partial bindings for v_i are stored.

The environment of AND-node which is the direct ancestor of the new OR-nodes is modified with an equation:

$$x_i = UNIF(v_1, v_2, \dots, v_n)$$

where v_i -s are all the newly created variables (Fig. 1.a). The functional semantics of UNIF is that of determining if the bindings for the v_i -s unify. The communication

among conjunctive nodes is determined using the procedural specification of UNIF.

AND-nodes are created as a result of a successful unification of a goal represented as an OR-node and some procedure head in the program.

Analogously to OR-nodes, the environment of AND-nodes is a set of variable templates. Let v be some variable that occurred in an OR-node α , and let t_1, \dots, t_n be the partial bindings for v and α obtained by creating all possible sons of α . We add the equation

$$v = MERGE(t_1, \dots, t_n)$$

to the environment of α (Fig. 1.b). MERGE performs a set union of the bindings to the variable v . As before in the case of UNIF, the operational specification of procedure MERGE is left unspecified. MERGE-UNIF equations allow to separate the process of resolution/reduction from the process of obtaining the bindings for shared variables. This separation of process control (literal sequencing) and communication control (management of bindings) allows the interpreter to define various communication strategies among conjunctive goals [1]. In particular, the communication structure may be determined by the setting of the variable templates in the environments of an AND/OR tree, and enforcing lazy/eager evaluation of the equations.

3. Intelligent Channel

Let $P(x), Q(x)$ be an AND-node, and let $P(x)$ and $Q(x)$ be its two OR-sons. If t is a binding for $P(x), Q(x)$ then there exists at least two paths $I(P, P')$ and $I(Q, Q')$ such that $I(P, P')$ and $I(Q, Q')$ are V -compatible, i.e., the composed substitution along each path is unifiable with t for x .

Thus, it is possible to pursue compatible computations of P and Q in parallel until they both reach the first OR-node where a decision must be made as to which path to pursue next. At the conflict point a binding is performed at each of the parallel processes to its own private copy of the shared AND-variable x . This binding is reported to a channel process assigned to a shared variable. All newly created nodes are suspended until a decision is made as to which partial OR-binding to pursue. The selection criterion is based on the compatibility of all parallel OR-computations that originated from the conjunctive goals with the shared vari-

able. We refer to partial bindings so obtained as *candidate bindings*. Once a candidate binding is chosen, only those suspended nodes whose execution is consistent with this binding are fired. When some partial binding is established to be a failure binding, this information is propagated up, and all the computations contingent on this partial binding are terminated.

For example, consider the logic program below for the goal $\leftarrow P(x), Q(x)$.

$$\begin{aligned} P(x) &\leftarrow P1(x). \quad Q(a) \leftarrow Q1(a). \quad P1(b). \\ P(a) &\leftarrow P2(a). \quad Q(c) \leftarrow Q2(c). \quad Q3(b). \\ P(b) &\leftarrow P3(b). \quad Q(c) \leftarrow Q3(c). \end{aligned}$$

A parallel execution of the program is given in Fig. 1. In Fig. 1.b $I(P,P1)$ is compatible with $I(Q,Q1)$, $I(Q,Q2)$ and $I(Q,Q3)$. The path $I(P,P2)$ is compatible with $I(Q,Q1)$, and $I(P,P3)$ is compatible with $I(Q,Q3)$. The Interpreter selects 'a' as a candidate binding for x. In Fig. 1.c 'a' was recognized to be a failure for P and subsequently for $\langle P, Q \rangle$. Similarly, 'c' is chosen for execution and is marked as a failure binding for $\langle P, Q \rangle$. As a result all the computations depending on a/x and c/x are removed yielding the snapshot in Fig. 1.d. Finally, the binding 'b/x' is selected yielding a successful parallel search, Fig. 1.e.

The evident benefits that stem from maintaining compatibility of parallel executions of OR-paths are not immediately implementable within the framework of existing control constructs. Below we present an algorithm, termed Intelligent Channel 1 (IC-1), that combines AND-parallelism with backtracking on partial bindings. IC-1 is based on a new control construct, *SELECT BINDING* that restricts the parallel execution of multiple branches in an AND/OR tree to the branches compatible with the selected binding.

Algorithm Intelligent Channel 1 (IC-1)

To clarify the presentation of the algorithm we present a semi-synchronous algorithm in which expansions of an OR-node are obtained simultaneously.

IC-1 consists of three parts: OR/AND-node execution and Channel Process execution. The creation of AND/OR nodes is as described in Section 2. OR/AND-node executions extend the creation of AND/OR-nodes with communication among conjunctive goals, and with pruning branches based on failure bindings. The Channel Process is delegated to the solution of MERGE-UNEF equations accumulated during a partial execution of a program and to the selection of candidate bindings. IC-1 follows:

OR-node execution

1. If the node is null or solved report success to the parent and terminate. Otherwise:

2. If the node is enabled, expand it by generating ALL possible successors. The enabling criteria for AND/OR nodes is the selection of a candidate binding which is consistent with the composed substitution used to derive the node. The successors are generated as

usual, by unifying the node with the set of axioms in the program. However, the instantiation of free variables in the node is determined by the candidate binding selected by the CHANNEL assigned to the free variable, and is obtained by accessing the "In" field of the variable template. This is how communication takes place among conjunctive goals

3. Suspend all newly created sons until the Channel process selects a candidate binding which is consistent with the composed substitution used to generate the new sons.

4. Report ALL newly computed partial bindings to the Channel Process; i.e., all partial bindings obtained by a substitution of the form: t/x are reported to the CHANNEL x. For simplicity, in the first version we insisted on mutual exclusion in Step 4 among all active AND-node processes, i.e., only one AND-node at a time communicates its partial bindings to the CHANNEL.

AND-node execution

1. If a node is *enabled*, create all its OR-sons, otherwise, suspend. Sons of an AND-node are created by splitting the conjunctive goal associated with the node into independent subgoals that interface through the UNIF equation.

2. For each new shared variable x create a shared channel data-structure denoted as CHANNEL x, and a Channel process that monitors the nondeterministic execution of its sons. This is achieved by setting the variable template of v_i to be $v_i : \langle in, out_i \rangle$ for all such variables v_i that emerged from the variable x. Thus, all v_i s that emerge from x share the same current instantiation.

3. Propagate up final/failure bindings for the node.

Channel Process Execution

1. For each CHANNEL check newly computed partial bindings against the set of candidate-bindings, thus, performing a partial solution of the MERGE-UNIF equations.

2. Whenever a failure binding is discovered, terminate the failure paths.

3. Whenever the set of candidate bindings is empty, *SELECT* a new candidate binding and activate all AND-nodes suspended on this binding.

ALGORITHM IC-1

Assign goal to AND-process and enable it.

Repeat until the goal is solved:

Execute all active AND-node processes.

Execute all active OR-node processes.

Execute all Channel processes.

4. Discussion

The major advantages of our approach are that only compatible computational paths are executed simultaneously, i.e., several branches of an AND/OR tree are explored in parallel while maintaining only *one active*

environment; The degree of parallelism may be monitored by modifying Step 1 in the Channel process to select multiple bindings. The combinatorial explosion that arises as a result of expanding the complete AND/OR tree is reduced, and many useless computations are eliminated.

Step 1 in the Channel Process is the key to completeness and efficiency of 1C-1. Analogous to procedure Select node in a conventional logic programming interpreter [2], select binding is the control component of 1C-1. It is not difficult to verify that if we select bindings in a breadth-first fashion, 1C-1 is a complete search

strategy. However, if a "depth-first-like selection strategy is used completeness is not guaranteed.

References

[1] Kasif, S., Analysis of Parallelism in Logic Programs, Ph.D. Thesis, Computer Science Dept., Univ. of Md., 1984.
 [2] Kowalski, R. A., *Logic for Problem Solving*, North-Holland, 1979.

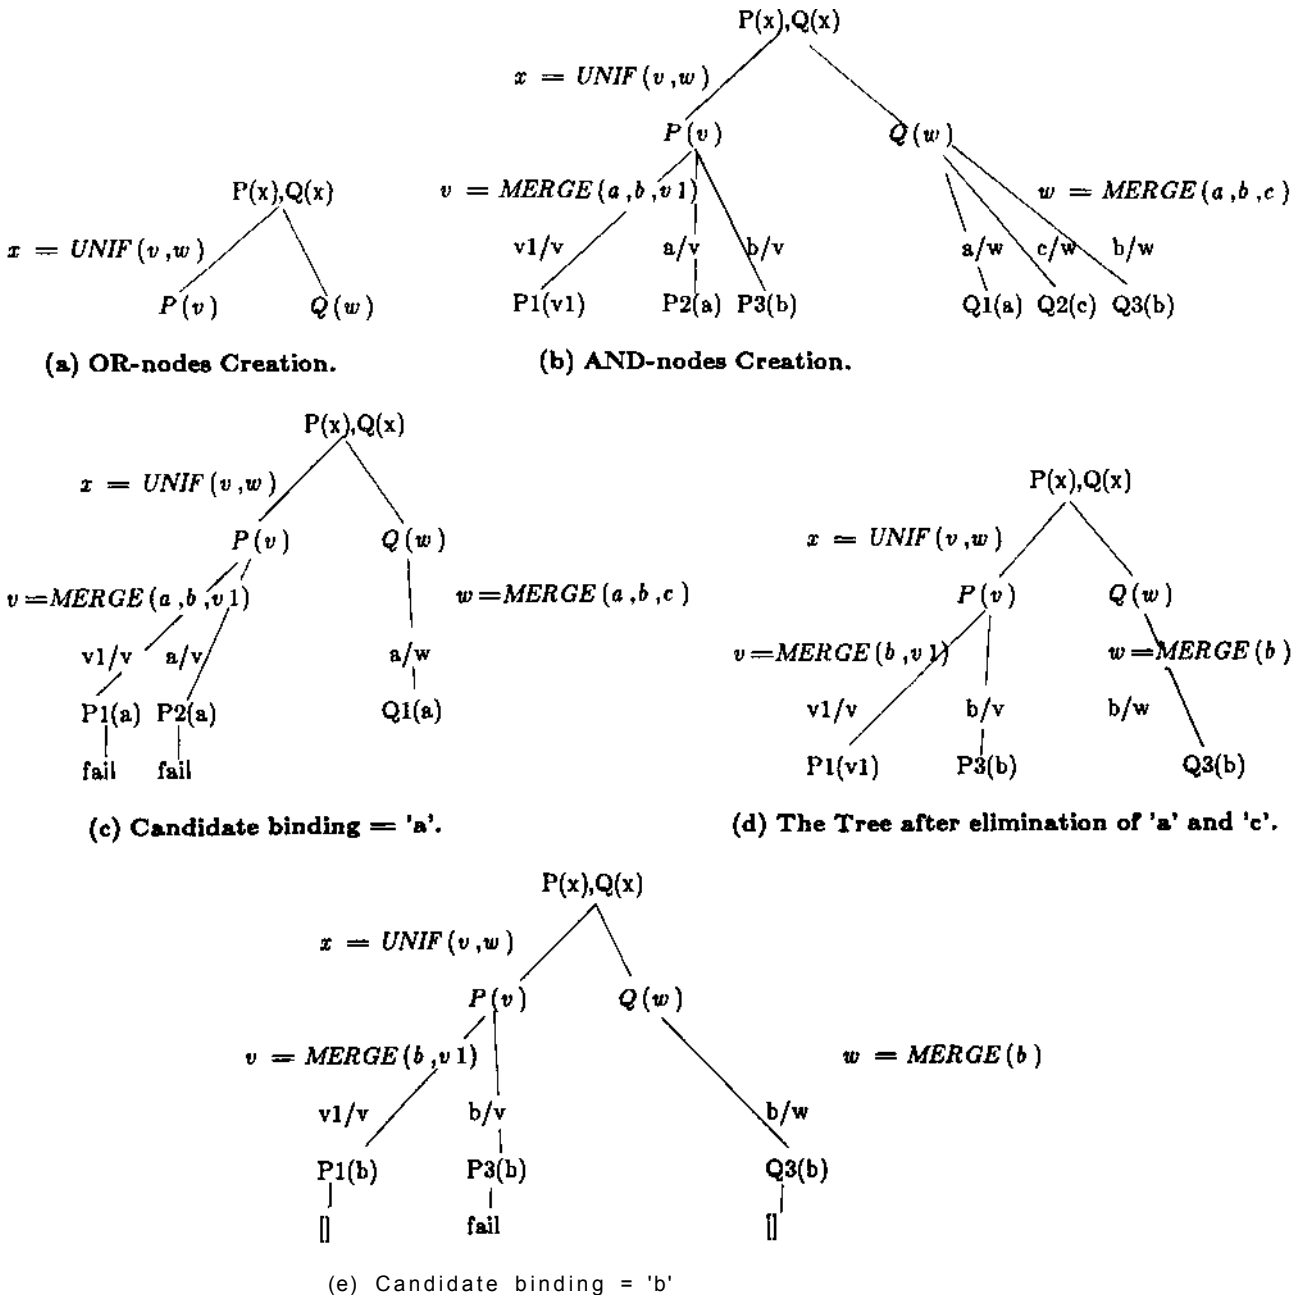


Figure 1: Intelligent Channel Execution.