Lower Bound Cost Estimation for Logic Programs *

Saumya Debray
Department of Computer Science
University of Arizona
Tucson, AZ 85721, U.S.A.
debray@cs.arizona.edu

P. López García   Manuel Hermenegildo
Dept. of Artificial Intelligence
Universidad Politecnica de Madrid
E-28600 Madrid, Spain
pedro@dia.fi.upm.es,herme@.fi.upm.es,

Nai-Wei Lin
Department of Computer Science and Information Engineering
National Chung Cheng University
Chiayi, 62107, Taiwan, R.O.C.
naiwei@cs.ccu.edu.tw

April 4, 1997

Keywords: Cost Analysis, Lower bound estimation, Granularity Control, Parallelism.

Abstract

It is generally recognized that information about the runtime cost of computations can be useful for a variety of applications, including program transformation, granularity control during parallel execution, and query optimization in deductive databases. Most of the work to date on compile-time cost estimation of logic programs has focused on the estimation of upper bounds on costs. However, in many applications, such as parallel implementations on distributed-memory machines, one would prefer to work with lower bounds instead. The problem with estimating lower bounds is that in general, it is necessary to account for the possibility of failure of head unification, leading to a trivial lower bound of 0. In this paper, we show how, given type and mode information about procedures in a logic program, it is possible to (semi-automatically) derive non-trivial lower bounds on their computational costs. We also discuss the cost analysis for the special and frequent case of divide-and-conquer programs and show how —as a pragmatic short-term solution —it may be possible to obtain useful results simply by identifying and treating divide-and-conquer programs specially.

1 Introduction

It is generally recognized that information about the runtime cost of computations can be useful for a variety of applications. For example, it is useful for granularity control, i.e., dynamic control of thread creation in parallel implementations of logic and functional languages [12, 3, 8, 13], and for query optimization in deductive databases [4]. In the context of logic programming, the work on cost estimation has generally focused on upper bound

*The work of S. Debray was supported in part by the National Science Foundation under grant CCR-9123520. The work of M. Hermenegildo and P. López-García was supported in part by ESPRIT project LTR 22532 "DiSCiPl" and CICYT proyect number TIC96-1012-C02-01.

cost analyses [5]. However, in many cases one would prefer to work with lower bounds instead. As an example, consider a distributed memory implementation of Prolog: suppose that the work involved in spawning a task on a remote processor takes 1000 instructions, and that we infer that a particular procedure call in a program will execute no more than 5000 instructions. This suggests that it *may be* worth executing this call on a remote processor, but provides no assurance that doing so will not actually produce a performance degradation relative to a sequential execution (the call might terminate after executing only a small number of instructions). On the other hand, if we know that a call will execute at least 5000 instructions, we can be assured that spawning a task on a remote processor to execute this call is worthwhile. Thus, while upper bound cost information is better than no information at all, lower bounds may be more useful than upper bounds.

The biggest problem with the inference of lower bounds on the computational cost of logic programs is the possibility of failure. Any attempt to infer lower bounds has to contend with the possibility that a goal may fail during head unification, yielding a trivial lower bound of 0. An obvious solution would be to try and rule out "bad" argument values by considering the types of predicates. However, most existing type analyses provide upper approximations, in the sense that the type of a predicate is a superset of the set of argument values that are actually encountered at runtime. Unfortunately, straightforward attempts to address this issue, for example by trying to infer lower approximations to the calling types of predicates, fail to yield nontrivial lower bounds for most cases.

In [2], we showed how, given mode and (upper approximation) type information, we can detect procedures and goals that can be guaranteed to not fail. Our technique is based on an intuitively very simple notion, that of a (set of) tests "covering" the type of a variable. We showed that the problem of determining a covering is undecidable in general, and is co-NP-hard even if we have only finite types and simple equality tests. We then gave an algorithm for checking whether a set of tests covers a type, that is efficient in practice. Based on this information, we show how to identify goals and procedures that can be guaranteed to not fail at runtime. Note that this information is interesting in its own right, in the context of program transformations (for example, we may want to execute possibly-failing goals ahead of non-failing goals where possible) and in systems that exploit speculative parallelism.

The main contributions of this paper are as follows: (i) we show how non-failure information can be used to infer lower bounds on the computational costs of goals; (ii) discuss how to bound the chromatic polynomial of a graph from below, and thereby show how to infer lower bounds on the number of solutions a predicate can generate (this information is useful, for example, for estimating communication costs in distributed-memory implementations); (iii) show how information about the number of solutions computed can be used to improve lower bound estimates when all solutions to a goal are required; and (iv) show how to obtain improved lower bound estimates for a simple but common class of divide-and-conquer programs. We discuss the application of our ideas to granularity control for parallel programs: in this case, the use of lower bound cost estimates guarantees that no slowdowns will occur, even in systems with significant overheads associated with parallel execution. Our ideas have been implemented within the CASLOG system, and the resulting lower bound cost estimates, given in Section 7, can be seen to be quite precise, especially for an automatic analysis tool. Experimental results with granularity control using lower bound cost estimates indicate that significant performance improvements can be attained using our approach.

2 Lower-Bound Cost Analysis: The One-Solution Case

If only one solution is required of any computation, it suffices to know whether a computation will generate at least one solution, i.e., will not fail. Assuming that this information is available, for example by using the technique mentioned in the previous section, cost analysis for a particular predicate can then proceed as follows:

1. We first determine the relative sizes of variable bindings at different program points in a clause by computing lower bounds on output argument sizes as functions of input argument sizes. This is done by solving (or estimating lower bound solutions to) the resulting difference equations: the approach is very similar to that discussed in [5], the only difference being that whereas [5] estimated upper bounds on argument sizes using the *max* function across the output sizes of different clauses in a cluster, we use the *min* function across clauses to estimate lower bounds on argument sizes.

2. The (lower bound) computational cost of a clause is then expressed as a function of the input argument size, in terms of the costs of the body literals in that clause.

   Consider a clause $C \equiv$ 'H :− $B_1, \ldots, B_m$'. Let the input argument size for the head of the clause be n, and let (lower bounds on) the input argument sizes for the body literals $B_1, \ldots, B_m$ be $\phi_1(n), \ldots, \phi_m(n)$ respectively. Assume that the cost of head unification and tests for this clause is at least $h(n)$, and let $Cost_{B_i}(x)$ denote a lower bound on the cost of the body literal B_i. Then, if B_k is the rightmost body literal that is guaranteed to not fail, the following gives a lower bound on the cost $Cost_C(n)$ of the clause C on an input of size n:

$$h(n) + \sum_{i=1}^{k} Cost_{B_i}(\phi_i(n)) \ \leq \ Cost_C(n).$$

3. A lower bound on the cost $Cost_p(n)$ of a predicate p on an input of size n is then given by

$$\min\{Cost_C(n) \mid C \text{ is a clause defining } p\} \ \leq \ Cost_p(n).$$

As discussed in [5], recursion is handled by expressing the cost of recursive goals symbolically as a function of the input size. From this, we can obtain a set of difference equations that can be solved (or approximated) to obtain a lower bound on the cost of a predicate in terms of the input size.

Given a predicate defined by m clauses $C_1, \ldots, C_m$, we can improve the precision of this analysis by noting that clause C_i will be tried only if clauses $C_1, \ldots, C_{i-1}$ fail to yield a solution. For an input of size n, let $\delta_i(n)$ denote the least amount of work necessary to determine that clauses $C_1, \ldots, C_{i-1}$ will not yield a solution and that C_i must be tried: the function δ_i obviously has to take into account the type and cost of the indexing scheme being used in the underlying implementation. In this case, the lower bound for p can be improved to:

$$\min\{Cost_{C_i}(n) + \delta_i(n) \mid 1 \leq i \leq m\} \ \leq \ Cost_p(n).$$

The pruning operator can also be taken into account, so that clauses which are after the first clause, say C_i, which has a non-failing sequence of literals just before the cut, are ignored, and the lower bound on the cost of the predicate is then the minimum of the costs of the clauses preceding the clause C_i and this clause itself.

3 Lower-Bound Cost Analysis: All Solutions

In many applications, it is reasonable to assume that all solutions are required. For example, in a distributed memory implementation of a logic programming system, the cost of sending or receiving a message is likely to be high enough that it makes sense for a remote computation to compute all the solutions to a query and return them in a single message instead of sending a large number of messages, each containing a single solution. For such cases, estimates of the computational cost of a goal can be improved greatly if we have lower bounds on the number of solutions—indeed, as the example of a distributed memory system suggests, in some cases the number of solutions may itself be a reasonable measure of cost.

If we obtain lower bounds on the number of solutions that can be generated by the literals in a clause (this problem is addressed in next section), we can use this information to improve lower bound cost estimates for the case where all solutions to a predicate are required. Consider a clause '$p(\bar{x})\ :-\ B_1, \ldots, B_n$' where B_k is the rightmost literal that is guaranteed to not fail. Let the input argument size for the head of the clause be n, and let (lower bounds on) the input argument sizes for the body literals $B_1, \ldots, B_m$ be $\phi_1(n), \ldots, \phi_m(n)$ respectively. Assume that the cost of head unification and tests for this clause is at least $h(n)$, and let $Cost_{B_i}(x)$ denote a lower bound on the cost of the body literal B_i. Now consider a body literal B_j, where $1 \leq j \leq k+1$, i.e., all the predecessors of B_j are guaranteed to not fail. The number of times B_j will be executed is given by the total number of solutions generated by its predecessors, i.e., the literals $B_1, \ldots, B_{j-1}$. Let this number be denoted by N_j : we can estimate N_j using Theorem 5.1 (or extensions thereof), e.g., by considering a clause whose body consists of the literals $B_1, \ldots, B_{j-1}$, and where the output variables in the head are given by $vars(B_1, \ldots, B_{j-1}) \cap vars(B_j, \ldots, B_n)$. Assume that the cost of head unification and tests for this clause is at least $h(n)$, and let $Cost_{B_i}(x)$ denote a lower bound on the cost of the body literal B_j. Then, a lower bound on the execution cost of the clause to obtain all solutions is given by

$$h(n) + \sum_{i=1}^{k}(N_i \times Cost_{B_i}(\phi_i(n)) \ \leq \ Cost_C(n).$$

4 Number of Solutions: The Single-Clause Case

In this section we address the problem of estimating lower bounds on the number of solutions which a predicate can generate.

4.1 Simple Conditions for Lower Bound Estimation

It is tempting to try and estimate a lower bound on the number of solutions generated by a clause '$H\ :-\ B_1, \ldots, B_n$' from lower bounds on the number of solutions generated by each of the body literals B_i, possibly using techniques analogous to those used in [5] for the estimation of upper bounds on the number of solutions. Unfortunately, this does not work.

For example, given a clause '$p(X)$:− $q(X), r(X)$', where X is an output variable, and assuming that q and r generate n_q and n_r bindings, respectively, for X, then $min(n_q, n_r)$ is not a lower bound on the number of solutions the clause can generate. To see this, consider the situation where q can bind X to either a or b, while r can bind X to either b or c: thus, $min(n_q, n_r) = min(2, 2) = 2$, but the number of solutions for the clause is 1.

The following gives a simple sufficient condition for estimating a lower bound on the number of solutions generated by a clause.

Theorem 4.1 *Let $x_1, \ldots, x_m$ be distinct unaliased output variables in the head of a clause such that each of the x_i occurs at most once in the body of the clause, and x_i and x_j do not occur in the same body literal for $i \neq j$. If n_i is a lower bound on the number of bindings that can be generated for x_i by the clause body, then $\prod_{i=1}^{m} n_i$ is a lower bound on the number of solutions that can be generated by the clause.*

This result can be generalized in various ways: we do not pursue them here due to space constraints. The utility of this theorem is shown in Example 5.1.

4.2 Handling Equality and Disequality Constraints

This section presents a simple algorithm for computing a lower bound on the number of solutions for predicates which can be "unfolded" into a conjunction of binary equality and disequality constraints on a set of variables. The constraints are in the form of $X = Y$ or $X \neq Y$ for any two variables X and Y. The types of the variables in a predicate are assumed to be the same and to be given as a finite set of atoms. The problem of computing the number of bindings that satisfy a set of binary equality and disequality constraints on a set of variables with the same type can be transformed into the problem of computing the *chromatic polynomial* of a graph G, denoted by $C(G, k)$, which is a polynomial in k and represents the number of different ways G can be colored by using no more than k colors (see [5]).

Unfortunately, the problem of computing the chromatic polynomial of a graph is NP-hard, because the problem of k-colorability of a graph G is equivalent to the problem of deciding whether $C(G, k) > 0$ and the problem of graph k-colorability is NP-complete [10]. Therefore, we will develop an approximation algorithm to compute a lower bound on the chromatic polynomial of a graph. The basic idea is to start with a subgraph that consists of only a single vertex of the graph, then repeatedly build larger and larger subgraphs by adding a vertex at a time into the previous subgraph. When a vertex is added, the edges connecting that vertex to vertices in the previous subgraph are also added. At each iteration, a lower bound on the number of ways of coloring the newly added vertex can be determined by the number of edges accompanied with the vertex. Accordingly, a lower bound on the chromatic polynomial for the corresponding subgraph can be determined using the bound on the polynomial for the previous subgraph and the bound on the number of ways of coloring the newly added vertex.

We now describe the algorithm more formally. The *order* of a graph $G = (V, E)$, denoted by $|G|$, is the number of vertices in V. Let G be a graph of order n. Suppose $\omega = v_1, \ldots, v_n$ is an ordering of V. We define two sequences of subgraphs of G according to ω. The first is a sequence of subgraphs $G_1, \ldots, G_n$, called *accumulating subgraphs*, where $G_i = (V_i, E_i)$, $V_i = \{v_1, \ldots, v_i\}$, and E_i is the set of edges of G that join the vertices of

Let $G = (V, E)$ be a graph of order n. The algorithm proceeds as follows:

begin
  compute the degree for each vertex in V;
  generate an ordering $\omega = v_1, \ldots, v_n$ of V by sorting the vertices in decreasing order
    on their degrees using the radix sort;
  $C(G, k) := k$;
  $G_1 := (\{v_1\}, \emptyset)$;
  for $i := 2$ **to** n **do**
    compute the order $|G'_i|$ of the interfacing subgraph G'_i;
    $C(G, k) := C(G, k) \times (k - |G'_i|)$;
    construct the accumulating subgraph G_i;
  od
end

Figure 1: An approximation algorithm for computing the chromatic polynomial of a graph

V_i, for $1 \leq i \leq n$, The second is a sequence of subgraphs $G'_2, \ldots, G'_n$, called *interfacing subgraphs*, where $G'_i = (V'_i, E'_i)$, V'_i is the set of vertices of G_{i-1} that are adjacent to vertex v_i, and E'_i is the set of edges of G_{i-1} that join the vertices of V'_i, for $2 \leq i \leq n$.

The algorithm for computing the chromatic polynomial of a graph, based on the construction of accumulating subgraphs and interfacing subgraphs, is shown in Figure 1. This algorithm constructs the accumulating subgraphs according to an ordering of the set of vertices. At each iteration, the number of ways of coloring the newly added vertex is computed based on the order of the corresponding interfacing subgraph.

Theorem 4.2 *Let $G = (V, E)$ be a graph of order n and ω be an ordering of V. Suppose the interfacing subgraphs of G corresponding to ω are $G'_2, \ldots, G'_n$. Then:*

$$k \prod_{i=2}^{n} (k - |G'_i|) \leq C(G, k).$$

The proof of this theorem is given in [11]; we omit it here due to space constraints. Since the bound obtained from this may depend on the ordering chosen for the vertices in the graph, we use a heuristic to find a "good" ordering. The intuition behind the heuristic is that if the maximum order of the interfacing subgraphs is smaller, then we can get a nontrivial lower bound ($\neq 0$) on $C(G, k)$ for more values of k. Therefore, we use the ordering that sorts the vertices in the decreasing order on the degrees of vertices.

Let the graph under consideration have n vertices and m edges. First, the computation for the degrees of vertices in the graph can be performed in $O(n + m)$. Second, since the degrees of vertices in the graph are at most $n - 1$, we can sort the vertices using radix sort in $O(n)$. Third, The total cost for the construction of accumulating subgraphs G_i, $i \leq i \leq n$, is $O(n + m)$ because each edge in the graph is examined only twice. Finally, since only the orders of the interfacing subgraphs are needed to compute the chromatic polynomial, it is not necessary to construct the interfacing graphs. The orders of the interfacing subgraphs can be obtained as a by-product of constructing the accumulating graphs. Therefore, the complexity of the whole algorithm is $O(n + m)$.

5 Number of Solutions: Multiple Clauses

The previous section discussed the estimation of lower bounds on the number of solutions computed by a single clause. In this section we discuss how we can estimate the number of solutions for a group of clauses.

Theorem 5.1 *Consider a set of clauses $S = \{C_1, \ldots, C_n\}$ that all have the same head unification and tests. If n_i is a lower bound on the number of solutions generated by C_i, $1 \leq i \leq n$, then $\sum_{i=1}^{n} n_i$ is a lower bound on the total number of (not necessarily distinct) solutions generated by the set of clauses S.*

The restrictions in this theorem can be relaxed in various ways: we do not pursue this here due to space constraints. We can use the result above to estimate a lower bound on the number of solutions generated by a predicate for an input of size n as follows: partition the clauses for the predicate into clusters such that the clauses in each cluster have the same head unification and tests, so that Theorem 5.1 is applicable, and compute lower bound estimates of the number of solutions for each cluster. Then, if a number of different clusters—say, clusters $C_1, \ldots, C_k$, with number of solutions at least $n_1, \ldots, n_k$ respectively, may be applicable to an input of size n, then the number of solutions overall for an input of size n is given by $min(n_1, \ldots, n_k)$. The utility of this approach is illustrated by the following example.

Example 5.1 Consider the following predicate to generate all subsets of a set represented as a list:

```
subset([], X) :- X = [].
subset([H|L], X) :- X = [H|X1], subset(L, X1).
subset([H|L], X) :- subset(L, X).
```

As discussed in Section 2, recursion is handled by initially using a symbolic representation to set up difference equations, and then solving, or estimating solutions to, these equations. In this case, let (a lower bound on) the number of solutions computed by `subset/2` on an input of size n be symbolically represented by $S(n)$. The first clause for the predicate yields the equation

$$S(0) = 1.$$

From Theorem 4.1, on an input of size n, $n > 0$, the second and third clauses each yield at least $S(n-1)$ solutions. Since they have the same head unification and tests, Theorem 5.1 is applicable, and the number of solutions given by these two clauses taken together is therefore at least $S(n-1) + S(n-1) = 2S(n-1)$. Thus, we have the equation

$$S(n) = 2S(n-1).$$

These difference equations can be solved to get the lower bound $S(n) = 2^n$ on the number of solutions computed by this predicate on an input of size n. $\square$

6 Cost Estimation for Divide-and-Conquer Programs

A significant shortcoming of the approach to cost estimation presented is its loss in precision in the presence of divide-and-conquer programs in which the sizes of the output arguments of the "divide" predicates are dependent. In the familiar quicksort program (see Section 6.1), for example, since either of the outputs of the partition predicate can be the empty list, the straightforward approach computes lower bounds under the assumption that both output can *simultaneously* be the empty list, and thereby significantly underestimates the cost of the program. In some sense, the reason for this loss of precision is that the approach outlined so far is essentially an independent attributes analysis [9]. However, even if we came up with a relational attributes analysis that kept track of relationships between the sizes of different output arguments of a predicate, it is not at all obvious how we might, systematically and from first principles, use this information to improve our lower bound cost estimates. For the quicksort program, for example, if the input list has length n, then the two output lists of the partition predicate have lengths m and $n - m - 1$ for some m, $0 \leq m < n$. The resulting cost equation for the recursive clause is of the form

$$C(n) = C(m) + C(n - m - 1) + \ldots \qquad (0 \leq m \leq n - 1)$$

In order to determine a worst-case lower bound solution to this equation we need to determine the value of m that maximizes the function $C(n)$, and doing this automatically, when we don't even know what $C(n)$ looks like, seems nontrivial. As a pragmatic solution, we argue that it may be possible to get quite useful results simply by identifying and treating common classes of divide-and-conquer programs specially.

In many of these programs, the sum of the sizes of the input for the "divide" predicates in the clause body is equal to the size of the input in the clause head minus some constant. This size relationship can be derived is some cases by the approach presented in [5]. However, this is not possible in other cases, since in this approach the size of each output argument is treated as a function only of the input sizes, independently of the sizes of other output arguments, and, as a result, relationships between the sizes of different output arguments are lost (consider for example the `partition`/4 predicate defined in example 6.1). A possible solution to gain in precision is to use one of the recently proposed approaches for inferring size relationships for this class of programs [1, 6].

Assuming that we have the mentioned size relationship for these programs, in the cost analysis phase we obtain an expression of the form:

$y(0) = C$,
$y(n) = y(n - 1 - k) + y(k) + g(n)$ for $n > 0$, where k is an arbitrary value such that $0 \leq k \leq n - 1$, C is a constant and $g(n)$ is any function.

in the cost analysis phase, where $y(n)$ denotes the cost of the divide-and-conquer predicate for an input of size n and $g(n)$ is the cost of the part of a clause body which does not contain any call to the divide-and-conquer predicate.

For each particular computation, we obtain a succession of values for k. Each succession of values for k yields a value for $y(n)$.

In the following we discuss how we can compute lower/upper bounds for expressions such as that for $\texttt{Cost}_{\texttt{qsort}}(n)$.

Consider the expression:

$y(0) = C$,
$y(n) = y(n-1-k) + y(k)$ for $n > 0$, where k is an arbitrary value such that $0 \le k \le n-1$ and C is a constant.

A *computation tree* for such an expression is a tree in which each non-terminal node is labeled with $y(n)$, $n > 0$, and has two children $y(n-1-k)$ and $y(k)$ (left- and right-hand-side respectively), where k is an arbitrary value such that $0 \le k \le n-1$. Terminal nodes are labeled with $y(0)$ and have no children. Assume that we construct a tree for $y(n)$ following a depth-first traversal. In each non-terminal node, we (arbitrarily) chose a value for k such that $0 \le k \le n-1$. We say that the *computation succession* of the tree is the succession of values that have been chosen for k in chronological order, as the tree construction proceeds.

Lemma 6.1 *Any computation tree corresponding to the expression:*

$y(0) = C$,
$y(n) = y(n-1-k) + y(k)$ *for* $n > 0$, *where k is an arbitrary value such that* $0 \le k \le n-1$ *and* C *is a constant,*

has $n+1$ *terminal nodes and* n *non-terminal nodes.*

Proof By induction on n. For $n = 0$ the theorem holds trivially. Let us assume that the theorem holds for all m such that $0 \le m \le n$, then, we can prove that for all m such that $0 \le m \le n+1$ the theorem also holds by reasoning as follows: we have that $y(n+1) = y(n-k) + y(k)$, where k is an arbitrary value such that $0 \le k \le n$. Since $0 \le k \le n$, we also have that $0 \le n-k \le n$, and, by induction hypothesis, the number of terminal nodes in any computation tree of $y(n-k)$ (respectively $y(k)$) is $n-k+1$ (respectively $k+1$). The number of terminal nodes in any computation tree of $y(n+1)$ is the sum of the number of terminal nodes in the children of the node labeled with $y(n+1)$, i.e. $(n-k+1)+(k+1) = n+2$. Also, the number of non-terminal nodes in any computation tree of $y(n-k)$ (respectively $y(k)$) is $n-k$ (respectively k). The number of non-terminal nodes of any computation tree of $y(n+1)$ is the sum of the number of non-terminal nodes of the children of the node labeled with $y(n+1)$ plus one (the node $y(n+1)$ itself, since it is non-terminal) i.e. $1 + (n-k) + k = n+1$. □

Theorem 6.2 *For any computation tree corresponding to the expression:*

$y(0) = C$,
$y(n) = y(n-1-k) + y(k)$ *for* $n > 0$, *where k is an arbitrary value such that* $0 \le k \le n-1$ *and* C *is a constant,*

it holds that $y(n) = (n+1) \times C$.

Proof By Lemma 6.1, any computation tree has $n+1$ terminal nodes labeled with $y(0)$ and the evaluation of each of these terminal nodes is C. □

Theorem 6.3 *Given the expression:*

$y(0) = C$,
$y(n) = y(n-1-k) + y(k) + g(k)$ *for* $n > 0$, *where k is an arbitrary value such that* $0 \le k \le n-1$, C *is a constant and* $g(k)$ *a function,*

for any computation tree corresponding to it, it holds that $y(n) = (n+1) \times C + \sum_{i=1}^{n} g(k_i)$*, where* $\{k_i\}_{i=1}^{n}$ *is the computation succession of the tree.*

Proof By Lemma 6.1, any computation tree has $n+1$ terminal nodes and n non-terminal nodes. The evaluation of each terminal node yields the value C and each time a non-terminal node i is evaluated, $g(k_i)$ is added. $\square$

In order to minimize (respectively maximize) $y(n)$ we can find a succession $\{k_i\}_{i=1}^{n}$ that minimizes (respectively maximizes) $\sum_{i=1}^{n} g(k_i)$. This is easy when $g(k)$ is a monotonic function, as the following corollary shows.

Corollary 6.1 *Given the expression:*

> $y(0) = C$,
> $y(n) = y(n-1-k) + y(k) + g(k)$ *for* $n > 0$*, where* k *is an arbitrary value such that* $0 \leq k \leq n-1$*,* C *is a constant and* $g(k)$ *an increasing monotonic function,*

Then, the succession $\{k_i\}_{i=1}^{n}$*, where* $k_i = 0$ *(respectively* $k_i = n-1$*) for all* $1 \leq i \leq n$ *gives the minimum (respectively maximum) value for* $y(n)$ *of all computation trees.*

Proof It follows from Theorem 6.3 and from the fact that $g(k)$ is an increasing monotonic function. $\square$

It follows from Corollary 6.1 that the solution of the difference equation (obtained by replacing k by 0):

> $y(0) = C$,
> $y(n) = y(n-1) + y(0) + g(0)$ for $n > 0$,

i.e. $(n+1) \times C + n \times g(0)$ is the minimum of $y(n)$, and the solution of the difference equation:

> $y(0) = C$,
> $y(n) = y(0) + y(n-1) + g(n-1)$ for $n > 0$,

i.e. $(n+1) \times C + n \times g(n-1)$ is the maximum of $y(n)$.

Note that we can replace $g(k)$ by any lower/upper bound on it to compute a lower/upper bound on $y(n)$. We can also take any lower/upper bound on each $g(k_i)$. For example, if $g(k)$ is an increasing monotonic function then $g(k_i) \leq g(n-1)$ and $g(k_i) \geq g(0)$ for $1 \leq i \leq n$, thus, $y(n) \leq (n+1) \times C + n \times g(n-1)$ and $y(n) \geq (n+1) \times C + n \times g(0)$.

Let's now assume that the function g depends on n and k:

Corollary 6.2 *Given the expression:*

> $y(0) = C$,
> $y(n) = y(n-1-k) + y(k) + g(n,k)$ *for* $n > 0$*, where* k *is an arbitrary value such that* $0 \leq k \leq n-1$*,* C *is a constant and* $g(n,k)$ *a function.*

Then, the solution of the difference equation:

> $f(0) = C$,
> $f(n) = f(n-1) + C + L$ *for* $n > 0$,

where L is a lower/upper bound on $g(n,k)$, is a lower/upper bound on $y(n)$ for all $n \geq 0$ and for any computation tree corresponding to $y(n)$. In particular, if $g(n,k)$ is an increasing monotonic function, then $L \equiv g(1,0)$ (respectively $L \equiv g(n, n-1)$) is a lower (respectively upper) bound on $g(n,k)$.

Example 6.1 Let us see how, using the described approach for divide-and-conquer programs, the lower-bound cost analysis can be improved. We first consider the analysis without the incorporation of the optimization, and then we compare with the result obtained when the optimization is used.

Consider the predicate `qsort/2` defined as follows:

```
qsort([], []).
qsort([First|L1], L2) :-
     partition(First, L1, Ls, Lg),
     qsort(Ls, Ls2), qsort(Lg, Lg2),
     append(Ls2, [First|Lg2], L2).

partition(F,[],[],[]).
partition(F,[X|Y],[X|Y1],Y2) :-
        X =< F,
        partition(F,Y,Y1,Y2).
partition(F,[X|Y],Y1,[X|Y2]) :-
        X > F,
        partition(F,Y,Y1,Y2).

append([],L,L).
append([H|L],L1,[H|R]) :- append(L,L1,R).
```

Let $\mathtt{Cost}_{\mathtt{p}}(n)$ denote the cost (number of resolution steps) of a call to predicate `p` with an input of size n.

The estimation of cost functions proceeds in a "bottom-up" way as follows:

The difference equation obtained for `append/3` is:

$\mathtt{Cost}_{\mathtt{append}}(0, m) = 1$ (the cost of head unification),
$\mathtt{Cost}_{\mathtt{append}}(n, m) = 1 + \mathtt{Cost}_{\mathtt{append}}(n-1, m)$.

The solution to this equation is: $\mathtt{Cost}_{\mathtt{append}}(n, m) = n + 1$.

The difference equation for `partition/4` is:

$\mathtt{Cost}_{\mathtt{partition}}(0) = 1$ (the cost of head unification),
$\mathtt{Cost}_{\mathtt{partition}}(n) = 1 + \mathtt{Cost}_{\mathtt{partition}}(n-1)$.

The solution to this equation is: $\mathtt{Cost}_{\mathtt{partition}}(n) = n + 1$.
For `qsort/2`, we have:

$\mathtt{Cost}_{\mathtt{qsort}}(0) = 1$ (the cost of head unification),
$\mathtt{Cost}_{\mathtt{qsort}}(n) = 1 + \mathtt{Cost}_{\mathtt{partition}}(n-1) + \mathtt{Cost}_{\mathtt{qsort}}(0) + \mathtt{Cost}_{\mathtt{qsort}}(0) + \mathtt{Cost}_{\mathtt{append}}(0)$

because the computed lower bound for the size of the input to the calls to `qsort` and `append` is 0. Thus, the cost function for `qsort/2` is given by:

$\texttt{Cost}_{\texttt{qsort}}(0) = 1$,
$\texttt{Cost}_{\texttt{qsort}}(n) = n + 4$, for $n > 0$.

Now, we use the described approach for divide-and-conquer programs. Assume that we use the expression:

$\texttt{Cost}_{\texttt{qsort}}(0) = 1$,
$\texttt{Cost}_{\texttt{qsort}}(n) =$
$1 + \texttt{Cost}_{\texttt{partition}}(n-1) + \texttt{Cost}_{\texttt{qsort}}(k) + \texttt{Cost}_{\texttt{qsort}}(n-1-k) + \texttt{Cost}_{\texttt{append}}(k)$,
for $0 \leq k \leq n-1$ and $n > 0$.

Replacing values, we obtain:

$$\texttt{Cost}_{\texttt{qsort}}(n) = 1 + n + \texttt{Cost}_{\texttt{qsort}}(k) + \texttt{Cost}_{\texttt{qsort}}(n-1-k) + k + 1, \text{ for } 0 \leq k \leq n-1.$$

$$\texttt{Cost}_{\texttt{qsort}}(n) = n + k + 2 + \texttt{Cost}_{\texttt{qsort}}(k) + \texttt{Cost}_{\texttt{qsort}}(n-1-k), \text{ for } 0 \leq k \leq n-1.$$

According to Corollary 6.2, by giving to n and k the minimum possible value, i.e. 1 and 0 respectively, we have that $n + k + 2 \geq 3$, and thus we replace $n + k + 2$ by 3 in order to obtain a lower bound on the former expression, which yields:

$$\texttt{Cost}_{\texttt{qsort}}(n) = 3 + \texttt{Cost}_{\texttt{qsort}}(k) + \texttt{Cost}_{\texttt{qsort}}(n-1-k), \text{ for } 0 \leq k \leq n-1.$$

which is equivalent to the difference equation:

$$\texttt{Cost}_{\texttt{qsort}}(n) = 3 + 1 + \texttt{Cost}_{\texttt{qsort}}(n-1), \text{ for } n > 0.$$

The solution of this equation is $\texttt{Cost}_{\texttt{qsort}}(n) = 4n + 1$, which is an improvement on the former lower bound. □

The previous results can be easily generalized to cover multiple recursive divide-and-conquer programs and programs where the sum of the sizes of the input for the "divide" predicates in the clause body is equal to the size of the input in the clause head minus some constant which is not necessarily 1.

Example 6.2 Consider the predicate `deriv/3` defined as:

```
deriv(U+V,X,DU+DV)            :- deriv(U,X,DU), deriv(V,X,DV).
deriv(U-V,X,DU-DV)            :- deriv(U,X,DU), deriv(V,X,DV).
deriv(U*V,X, DU*V+U*DV)       :- deriv(U,X,DU), deriv(V,X,DV).
deriv(U/V,X,(DU*V-U*DV)/V^2) :- deriv(U,X,DU), deriv(V,X,DV).
deriv(U^N,X, DU*N*U^N1)       :- integer(N), N1 is N-1, deriv(U,X,DU).
deriv(-U,X,-DU)               :- deriv(U,X,DU).
deriv(exp(U),X,exp(U)*DU)     :- deriv(U,X,DU).
deriv(log(U),X,DU/U)          :- deriv(U,X,DU).
deriv(X,X,1)                  :- atom(X).
deriv(C,X,0)                  :- atom(C), C \= X.
```

Let $\texttt{Cost}_{\texttt{deriv}}(n)$ denote the computational cost (number of resolution steps) of a call to deriv/3, where the size of the first argument (which is an input argument) is n. The size measure used is the number of nodes (function symbols) in a term.

Roughly, the analysis proceeds as follows:

From clauses 9 and 10 we derive the boundary condition $\texttt{Cost}_{\texttt{deriv}}(1) = 1$. For the rest of the clauses we derive difference equations, so that we choose the minimum of the solution functions to these equations:

- For clauses 1 to 4 we yield the equation:

  Equation: $\texttt{Cost}_{\texttt{deriv}}(n) = 1 + \texttt{Cost}_{\texttt{deriv}}(n-1-k) + \texttt{Cost}_{\texttt{deriv}}(k)$,
  whose solution is $\texttt{Cost}_{\texttt{deriv}}(n) = n$.

- For clause 5:

  Equation: $\texttt{Cost}_{\texttt{deriv}}(n) = 1 + 1 + 1 + \texttt{Cost}_{\texttt{deriv}}(n-2)$.
  Solution: $\texttt{Cost}_{\texttt{deriv}}(n) = 3n \div 2 + 1$.

- For clauses 6 and 7:

  Equation: $\texttt{Cost}_{\texttt{deriv}}(n) = 1 + \texttt{Cost}_{\texttt{deriv}}(n-1)$.
  Solution: $\texttt{Cost}_{\texttt{deriv}}(n) = n$.

Thus, the cost function inferred is $\texttt{Cost}_{\texttt{deriv}}(n) = n$, for all $n \geq 1$, which is the best lower bound.

Now suppose we do not use the described approach for divide and conquer programs:

For clauses 1 to 4 we have:

Equation: $\texttt{Cost}_{\texttt{deriv}}(n) = 1 + \texttt{Cost}_{\texttt{deriv}}(1) + \texttt{Cost}_{\texttt{deriv}}(1)$.
Solution: $\texttt{Cost}_{\texttt{deriv}}(n) = 3$.

This cost function is not very useful (the problem here is that the lower bounds computed for the size of the input of the two recursive calls is trivial).

For the rest of the clauses we obtain the same results as before, but when we choose the minimum of these functions we have that the final cost function is: $\texttt{Cost}_{\texttt{deriv}}(1) = 1$, and $\texttt{Cost}_{\texttt{deriv}}(n) = 3$, for $n > 1$ (which is not very useful). □

7 Implementation

We have implemented a prototype of a lower bound size/cost analyzer, by recoding the version of CASLOG [5] currently integrated in the CIAO system [7]. The analysis is fully automatic, and only requires type information for the program entry point. Types, modes and size measures are automatically inferred by the system. Table 1 shows some accuracy and efficiency results of the lower bound cost analyzer. The second column of the table shows the cost function (which depends on the size of the input arguments) inferred by the analysis. T_{tms} is the time required by the type, mode, and size measure analysis (SPARCstation 10, 55MHz, 64Mbytes of memory), T_{nf} the time required by the non-failure analysis, and T_{ca} the time required by the cost analysis (which includes a size analysis). **Total** is the total analysis time ($Total = T_{tms} + T_{nf} + T_{ca}$). All times are given in milliseconds.

Programs	**Cost function**	T_{tms}	T_{nf}	T_{ca}	**Total**
fib	$\lambda x.1.447 \times 1.618^x + 0.552 \times (-0.618)^x - 1$	90	20	70	180
hanoi	$\lambda x.x2^x + 2^{x-1} - 2$	860	60	210	1,130
qsort	$\lambda x.4x + 1$	440	80	200	720
nrev	$\lambda x.0.5x^2 + 1.5x + 1$	100	10	90	200
mmatrix	$\lambda\langle x, y\rangle.2xy + 2x + 1$	350	90	180	620
deriv	$\lambda x.x$	940	80	300	1,320
add_poly	$\lambda\langle x, y\rangle x + 1$	280	90	160	530
append	$\lambda x.x + 1$	130	60	70	260
subst	$\lambda\langle x, y, z\rangle.x$	120	60	140	320
flat	$\lambda x.0.5x^2 + 3.5x - 2$	310	30	270	410
intersect	$\lambda\langle x, y\rangle.x + 1$	140	70	130	340
diff	$\lambda\langle x, y\rangle.x + 1$	170	100	180	450
perm	$\lambda x.0.5x^2 + 1.5x + 1$	120	60	190	370

Table 1: Accuracy and efficiency of the lower bound cost analysis

8 Application to Automatic Parallelization

As briefly mentioned in the introduction, one of the most attractive applications of lower bound cost analysis is implementing granularity control in parallelizing compilers, an issue on which we expand in this section. The usefulness of granularity control is based on the fact that while logic programming languages offer a great deal of scope for parallelism, just because something *can* be done in parallel does not necessarily mean, in practice, that it *should* be done in parallel. This is because the parallel execution of a task incurs various overheads, e.g. overheads associated with process creation and scheduling, the possible migration of tasks to remote processors and the associated communication overheads, etc. In general, a goal should not be a candidate for parallel execution if its granularity, i.e., the "work available" underneath it, is less than the work necessary to create a separate task for that goal. While the overheads for spawning goals in parallel in some architectures are small (e.g. in small shared memory multiprocessors), in many other architectures (e.g. distributed memory multiprocessors, workstation "farms", etc.) they can be very significant. Automatic parallelization in general cannot be done realistically in the latter without granularity control.

Due to the considerations mentioned above it is desirable to devise a method whereby the granularity of a goal may be estimated. All of the previous work that we know of in this context involves estimating upper bounds on the cost of goals (see, for example, [3]). The use of upper bounds allows us to guarantee that, given a program that is already parallelized, we can make it run more efficiently by running some of the parallel goals sequentially. However, the problem faced by parallelizing compilers is in fact exactly the converse of the one tackled above: what needs to be guaranteed is that the parallel execution will be more efficient than the sequential one, rather than the other way around. This type of granularity control can be solved using essentially the same general approach, but we need a *lower bound* on the cost of each goal. The techniques presented in the paper directly address this problem, arguably more interesting in practice, and which could not be solved with the upper bound approximation of [3]. In fact, the usefulness of lower bounds was already clear when the work presented in [3] was developed, but the determination of useful lower bounds was deemed too difficult at the time. This approach allows us to guarantee that, given a sequential program, it will run more efficiently by running some of the goals in parallel. This in effect allows obtaining *guaranteed speedups* (or, at least, ensuring that no *slow-downs* will occur) from automatic parallelization, even in architectures (such as those mentioned above) for which parallel execution involves a significant overhead. We know of no other approach which can achieve this.

programs	**seq**	**ngc**	**gc**	**opt**	$\mathbf{e_1}$	$\mathbf{e_2}$
unb_matrix	0.579	16.247	0.511	0.057	+96.85 %	+97.20 %
fib	0.77	1.475	0.188	0.077	+87.24 %	+92.04 %
hanoi	1.440	1.467	0.592	0.144	+59.65 %	+66.15 %
qsort	0.476	0.472	0.293	0.047	+37.84 %	+42.09 %

Table 2: Granularity control results for benchmarks on ECL^{i}PS^e.

We have performed a series of experiments in this granularity control application by using the cost functions inferred by the proposed cost analysis. In this sense, we have used the granularity control system described in [12] (which is integrated in the CIAO system, and has an annotator which transforms programs to perform granularity control). Table 2 presents results of granularity control (showing execution times in seconds) for some benchmarks on the ECL^{i}PS^e system using 10 workers, and running on a SUN SPARC 2000 SERVER with 10 processors. Results are given for the versions which perform granularity control (`gc`), the sequential execution (`seq`) and the parallel execution without granularity control (`ngc`) for comparison. **opt** is a lower bound on the optimal time, i.e. $\mathbf{opt} = \frac{\mathbf{seq}}{10}$. $\mathbf{e_1}$ and $\mathbf{e_2}$ are computed according to the following expressions: $\mathbf{e_1} = \frac{\mathbf{ngc}-\mathbf{gc}}{\mathbf{ngc}} \times 100$, and $\mathbf{e_2} = \frac{\mathbf{ngc}-\mathbf{gc}}{\mathbf{ngc}-\mathbf{opt}} \times 100$. The program `unb_matrix` performs the multiplication of 4×2 and 2×1000 matrices.

The experiments performed show very promising results, in the sense that the granularity control does improve speedups in practice.

References

[1] F. Benoy and A. King. Inferring Argument Size Relationships with CLP(R). *Proc. 6th International Workshop on Logic Program Synthesis and Transformation*, Stockholm University/Royal Intitute of Technology, 1996, pp. 134–153.

[2] S. Debray, P. López García, and M. Hermenegildo. Non-Failure Analysis for Logic Programs. In *1997 International Conference on Logic Programming*, Leuven, Belgium, June 1997. MIT Press, Cambridge, MA.

[3] S. K. Debray, N.-W. Lin, and M. Hermenegildo. Task Granularity Analysis in Logic Programs. In *Proc. of the 1990 ACM Conf. on Programming Language Design and Implementation*, pages 174–188. ACM Press, June 1990.

[4] S.K. Debray and N.-W. Lin. Static estimation of query sizes in horn programs. In *Third International Conference on Database Theory*, Lecture Notes in Computer Science 470, pages 515–528, Paris, France, December 1990. Springer-Verlag.

[5] S.K. Debray and N.W. Lin. Cost analysis of logic programs. *ACM Transactions on Programming Languages and Systems*, 15(5):826–875, November 1993.

[6] R. Giacobazzi, S.K. Debray, and G. Levi. Generalized Semantics and Abstract Interpretation for Constraint Logic Programs. *Journal of Logic Programming*, 25(3):191–248, 1995.

[7] M. Hermenegildo, F. Bueno, M. García de la Banda, and G. Puebla. The CIAO Multi-Dialect Compiler and System: An Experimentation Workbench for Future (C)LP Systems. In *Proceedings of the ILPS'95 Workshop on Visions for the Future of Logic Programming*, Portland, Oregon, USA, December 1995. Available from `http://www.clip.dia.fi.upm.es/`.

[8] L. Huelsbergen, J. R. Larus, and A. Aiken. Using Run-Time List Sizes to Guide Parallel Thread Creation. In *Proc. ACM Conf. on Lisp and Functional Programming*, June 1994.

[9] Neil D. Jones and Steven S. Muchnick. Complexity of flow analysis, inductive assertion synthesis, and a language due to Dijkstra. In Steven S Muchnick and Neil D Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 12, pages 380–393. Prentice-Hall, 1981.

[10] R. M. Karp. Reducibility among Combinatorial Problems. *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher (eds), Plenum Press, New York, 1972, pp. 85–103.

[11] N.-W. Lin. Approximating the Chromatic Polynomial of a Graph. *Proc. Nineteenth International Workshop on Graph-Theoretic Concepts in Computer Science*, Amsterdam, June 1993.

[12] P. López García, M. Hermenegildo, and S.K. Debray. A Methodology for Granularity Based Control of Parallelism in Logic Programs. *Journal of Symbolic Computation, Special Issue on Parallel Symbolic Computation*, 11(3–4):217–242, 1996. In press.

[13] F. A. Rabhi and G. A. Manson. Using Complexity Functions to Control Parallelism in Functional Programs. Res. Rep. CS-90-1, Dept. of Computer Science, Univ. of Sheffield, England, Jan 1990.