

A General Framework for Static Cost Analysis of Parallel Logic Programs^{*}

M. Klemen^{1,2}, P. López-García^{1,3}, J.P. Gallagher^{1,4}, J.F. Morales¹, and M.V. Hermenegildo^{1,2}

¹ IMDEA Software Institute, Spain

² ETSI Informáticos, Universidad Politécnica de Madrid (UPM), Spain

³ Spanish Council for Scientific Research (CSIC), Spain

⁴ Roskilde University, Denmark

{maximiliano.klemen, pedro.lopez, john.gallagher, josef.morales, manuel.hermenegildo}@imdea.org

Abstract. The estimation and control of *resource usage* is now an important challenge in an increasing number of computing systems. In particular, requirements on timing and energy arise in a wide variety of applications such as internet of things, cloud computing, health, transportation, and robots. At the same time, parallel computing, with (heterogeneous) multi-core platforms in particular, has become the dominant paradigm in computer architecture. Predicting resource usage on such platforms poses a difficult challenge. Most work on static resource analysis has focused on sequential programs, and relatively little progress has been made on the analysis of parallel programs, or more specifically on parallel logic programs. We propose a novel, general, and flexible framework for setting up cost equations/relations which can be instantiated for performing resource usage analysis of parallel logic programs for a wide range of resources, platforms and execution models. The analysis estimates both lower and upper bounds on the resource usage of a parallel program (without executing it) as functions on input data sizes. In addition, it also infers other meaningful information to better exploit and assess the potential and actual parallelism of a system. We develop a method for solving cost relations involving the *max* function that arise in the analysis of parallel programs. Finally, we instantiate our general framework for the analysis of logic programs with Independent And-Parallelism, report on an implementation within the CiaoPP system, and provide some experimental results. To our knowledge, this is the first approach to the cost analysis of *parallel logic programs*.

Keywords: Resource Usage Analysis, Parallelism, Static Analysis, Complexity Analysis, (Constraint) Logic Programming, Prolog.

1 Introduction

Estimating in advance the resource usage of computations is useful for a number of applications; examples include granularity control in parallel/distributed systems, automatic program optimization, verification of resource-related specifications and

^{*} Research partially funded by Spanish MINECO TIN2015-67522-C3-1-R *TRACES* project, and the Madrid P2018/TCS-4339 *BLOQUES-CM* program. We are also grateful to the anonymous reviewers for their useful comments.

detection of performance bugs, as well as helping developers make resource-related design decisions. Besides *time* and *energy*, we assume a broad concept of resources as numerical properties of the execution of a program, including the number of *execution steps*, the number of *calls* to a procedure, the number of *network accesses*, number of *transactions* in a database, and other user-definable resources. The goal of automatic static analysis is to estimate such properties without running the program with concrete data, as a function of input data sizes and possibly other (environmental) parameters.

Due to the heat generation barrier in traditional sequential architectures, parallel computing, with (heterogeneous) multi-core processors in particular, has become the dominant paradigm in current computer architecture. Predicting resource usage on such platforms poses important challenges. Most work on static resource analysis has focused on sequential programs, and relatively little progress has been made on the analysis of parallel programs, or on parallel logic programs in particular. The significant body of work on static analysis of sequential logic programs has already been applied to the analysis of other programming paradigms, including imperative programs. This is achieved via a transformation into *Horn clauses* [23]. In this paper we concentrate on the analysis of parallel Horn clause programs, which could be the result of such a translation from a parallel imperative program or be themselves the source program. Our starting point is the well-developed technique of setting up recurrence relations representing resource usage functions parameterized by input data sizes [28, 25, 10, 8, 11, 24, 2, 26], which are then solved to obtain (exact or safely approximated) closed forms of such functions (i.e., functions that provide upper or lower bounds on resource usage). We build on this and propose a novel, general, and flexible framework for setting up cost equations/relations which can be instantiated for performing static resource usage analysis of parallel logic programs for a wide range of resources, platforms and execution models. Such an analysis estimates both lower and upper bounds on the resource usage of a parallel program as functions on input data sizes. We have instantiated the framework for dealing with Independent And-Parallelism (IAP) [16, 12], which refers to the parallel execution of conjuncts in a goal. However, the results can be applied to other languages and types of parallelism, by performing suitable transformations into Horn clauses.

The main contributions of this paper can be summarized as follows:

- We have extended a general static analysis framework for the analysis of sequential Horn clause programs [24, 26], to deal with parallel programs.
- Our extensions and generalizations support a wide range of resources, platforms and parallel/distributed execution models, and allow the inference of both lower and upper bounds on resource usage. This is the first approach, to our knowledge, to the cost analysis of *parallel logic programs* that can deal with features such as backtracking, multiple solutions (i.e., non-determinism), and failure.
- We have instantiated the developed framework to infer useful information for assessing and exploiting the potential and actual parallelism of a system.
- We have developed a method for finding closed-form functions of cost relations involving the *max* function that arise in the analysis of parallel programs.
- We have developed a prototype implementation that instantiates the framework for the analysis of logic programs with Independent And-Parallelism within the CiaoPP system [15, 24, 26], and provided some experimental results.

2 Overview of the Approach

Prior to explaining our approach, we provide some preliminary concepts. Independent And-Parallelism arises between two goals when their corresponding executions do not affect each other. For pure goals (i.e., without side effects) a sufficient condition for the correctness of IAP is the absence of variable sharing at run-time among such goals. IAP has traditionally been expressed using the $\&/2$ meta-predicate as the constructor to represent the parallel execution of goals. In this way, the conjunction of goals (i.e., literals) \mathbf{p} & \mathbf{q} in the body of a clause will trigger the execution of goals \mathbf{p} and \mathbf{q} in parallel, finishing when both executions finish.

Given a program \mathcal{P} and a predicate $\mathbf{p} \in \mathcal{P}$ of arity k and a set Π of k -tuples of calling data to \mathbf{p} , we refer to the (*standard*) *cost* of a call $\mathbf{p}(\bar{\mathbf{e}})$ (i.e., a call to \mathbf{p} with actual data $\bar{\mathbf{e}} \in \Pi$), as the resource usage (under a given cost metric) of the complete execution of $\mathbf{p}(\bar{\mathbf{e}})$. The *standard cost* is formalized as a function $\mathcal{C}_{\mathbf{p}} : \Pi \rightarrow \mathcal{R}_{\infty}$, where \mathcal{R}_{∞} is the set of real numbers augmented with the special symbol ∞ (which is used to represent non-termination). We extend the function $\mathcal{C}_{\mathbf{p}}$ to the powerset of Π , i.e., $\hat{\mathcal{C}}_{\mathbf{p}} : 2^{\Pi} \rightarrow 2^{\mathcal{R}_{\infty}}$, where $\hat{\mathcal{C}}_{\mathbf{p}}(E) = \{\mathcal{C}_{\mathbf{p}}(\bar{\mathbf{e}}) \mid \bar{\mathbf{e}} \in E\}$. Our goal is to abstract (safely approximate, as accurately as possible) $\hat{\mathcal{C}}_{\mathbf{p}}$ (note that $\mathcal{C}_{\mathbf{p}}(\bar{\mathbf{e}}) = \hat{\mathcal{C}}_{\mathbf{p}}(\{\bar{\mathbf{e}}\})$). Intuitively, this abstraction is the composition of two abstractions: a size abstraction and a cost abstraction. The goal of the analysis is to infer two functions $\hat{\mathcal{C}}_{\mathbf{p}}^{\downarrow}$ and $\hat{\mathcal{C}}_{\mathbf{p}}^{\uparrow} : \mathcal{N}_{\top}^m \rightarrow \mathcal{R}_{\infty}$ that give lower and upper bounds respectively on the cost function $\hat{\mathcal{C}}_{\mathbf{p}}$, where \mathcal{N}_{\top}^m is the set of m -tuples whose elements are natural numbers or the special symbol \top , meaning that the size of a given term under a given size metric is *undefined*. Such bounds are given as a function of tuples of data sizes (representing the concrete tuples of data of the concrete function $\hat{\mathcal{C}}_{\mathbf{p}}$). Typical size metrics are the actual value of a number, the length of a list, the size (number of constant and function symbols) of a term, etc. [24, 26].

We now enumerate different metrics used to evaluate the performance of parallel versions of a logic program, compared against its corresponding sequential version. When possible, we define these metrics parameterized with respect to the resource (e.g., number of *resolution steps*, *execution time*, or *energy consumption*) in which the cost is expressed.

- **Sequential cost (*Work*)**: It is the standard cost of executing a program, assuming no parallelism.
- **Parallel cost (*Depth*)**: It is the cost of executing a program in parallel, considering an unbounded number of processors.
- **Maximum number of processes running in parallel ($Proc_{max}(P)$)**: The maximum number of processes that can run simultaneously in a program. This is useful, for example, to determine what is the minimum number of processors that are required to actually run all the processes in parallel.

The following example illustrates our approach.

Example 1. Consider the predicate `scalar/3` below, and a calling mode to it with the first argument bound to an integer n and the second one bound to a list of integers $[x_1, x_2, \dots, x_k]$. Upon success, the third argument is bound to the list of products $[n \cdot x_1, n \cdot x_2, \dots, n \cdot x_k]$. Each product is recursively computed by predicate

`mult/3`. The calling modes are automatically inferred by CiaoPP (see [15] and its references): the first two arguments of both predicates are input, and their last arguments are output.

```

1 scalar(_, [], []).
2 scalar(N, [X|Xs], [Y|Ys]) :-
3     mult(N, X, Y) & scalar(N, Xs, Ys).
4
5 mult(0, _, 0).
6 mult(N, X, Y) :-
7     N > 1,
8     N1 is N - 1,
9     mult(N1, X, Y0),
10    Y is Y0 + X.

```

The call to the parallel `&/2` operator in the body of the second clause of `scalar/3` causes the calls to `mult/3` and `scalar/3` to be executed in parallel.

We want to infer the cost of such a call to `scalar/3`, in terms of the number of resolution steps, as a function of its input data sizes. We use the CiaoPP system to infer size relations for the different arguments in the clauses, as well as dealing with a rich set of size metrics (see [24, 26] for details). Assume that the size metrics used in this example are the *actual value* of `N` (denoted `int(N)`), for the first argument, and the *list-length* for the second and third arguments (denoted `length(X)` and `length(Y)`). Since size relations are obvious in this example, we focus only on the setting up of cost relations for the sake of brevity. Regarding the number of solutions, in this example all the predicates generate at most one solution. For simplicity we assume that all builtin predicates, such as `is/2` and the comparison operators have zero cost (in practice they have a “trust” assertion that specifies their cost as if it had been inferred by the system). As the program contains parallel calls, we are interested in inferring both total resolution steps, i.e., considering a sequential execution (represented by the `seq` identifier), and the number of parallel steps, considering a parallel execution, with infinite number of processors (represented by `par`). In the latter case, the definition of this resource establishes that the aggregator of the costs of the parallel calls that are arguments of the `&/2` meta-predicate is the `max/2` function. Thus, the number of parallel resolution steps for `p & q` is the maximum between the parallel steps performed by `p` and the ones performed by `q`. However, for computing the *total resolution steps*, the aggregation operator we use is the addition, both for parallel and sequential calls. For brevity, in this example we only infer upper bounds on resource usages.

We now set up the cost relations for `scalar/3` and `mult/3`. Note that the cost functions have two arguments, corresponding to the sizes of the input arguments⁵. In the equations, we underline the operation applied as cost aggregator for `&/2`.

For the sequential execution (`seq`), we obtain the following cost relations:

$$\begin{aligned}
 C_{\text{scalar}}(n, l) &= 1 && \text{if } l = 0 \\
 C_{\text{scalar}}(n, l) &= C_{\text{mult}}(n) \underline{+} C_{\text{scalar}}(n, l - 1) + 1 && \text{if } l > 0 \\
 C_{\text{mult}}(n) &= 1 && \text{if } n = 0 \\
 C_{\text{mult}}(n) &= C_{\text{mult}}(n - 1) + 1 && \text{if } n > 0
 \end{aligned}$$

⁵ For the sake of clarity, we abuse of notation in the examples when representing the cost functions that depend on data sizes.

After solving these equations and composing the closed-form solutions, we obtain the following closed-form functions:

$$\begin{aligned} C_{\text{scalar}}(n, l) &= (n + 2) \times l + 1 && \text{if } n \geq 0 \wedge l \geq 0 \\ C_{\text{mult}}(n) &= n + 1 && \text{if } n \geq 0 \end{aligned}$$

For the parallel execution (**par**), we obtain the following cost relations:

$$\begin{aligned} C_{\text{scalar}}(n, l) &= 1 && \text{if } l = 0 \\ C_{\text{scalar}}(n, l) &= \underline{\text{max}}(C_{\text{mult}}(n), C_{\text{scalar}}(n, l - 1)) + 1 && \text{if } l > 0 \\ C_{\text{mult}}(n) &= 1 && \text{if } n = 0 \\ C_{\text{mult}}(n) &= C_{\text{mult}}(n - 1) + 1 && \text{if } n > 0 \end{aligned}$$

After solving these equations and composing the closed forms, we obtain the following closed-form functions:

$$\begin{aligned} C_{\text{scalar}}(n, l) &= n + l + 1 && \text{if } n \geq 0 \wedge l \geq 0 \\ C_{\text{mult}}(n) &= n + 1 && \text{if } n \geq 0 \end{aligned}$$

By comparing the complexity order (in terms of resolution steps) of the sequential execution of **scalar/3**, $O(n \cdot l)$, with the complexity order of its parallel execution (assuming an ideal parallel model with an unbounded number of processors) $O(n+l)$, we can get a hint about the maximum achievable parallelization of the program.

Another useful piece of information about **scalar/3** that we want to infer is the maximum number of processes that may run in parallel, considering all possible executions. For this purpose, we define a resource in our framework named **sthreads**. The operation that aggregates the cost of both arguments of the meta-predicate **&/2**, **count_process/3** for the **sthreads** resource, adds the maximum number of processes for each argument plus one additional process, corresponding to the one created by the call to **&/2**. The sequential cost aggregator is now the *maximum* operator, in order to keep track of the maximum number of processes created along the different instructions of the program executed sequentially. Note that if the instruction **p** executes at most Pr_p processes in parallel, and the instruction **q** executes at most Pr_q processes, then the program **p**, **q** will execute at most $\text{max}(Pr_p, Pr_q)$ processes in parallel, because all the parallel processes created by **p** will finish before the execution of **q**. Note also that for the sequential execution of both **p** and **q**, the cost in terms of the **sthreads** resource is always zero, because no additional process is created.

The analysis sets up the following recurrences for the **sthreads** resource and the predicates **scalar/3** and **mult/3** of our example:

$$\begin{aligned} C_{\text{scalar}}(n, l) &= 0 && \text{if } l = 0 \\ C_{\text{scalar}}(n, l) &= C_{\text{mult}}(n) + C_{\text{scalar}}(n, l - 1) + 1 && \text{if } l > 0 \\ C_{\text{mult}}(n) &= 0 && \text{if } n \geq 0 \end{aligned}$$

After solving these equations and composing the closed forms, we obtain the following closed-form functions:

$$\begin{aligned} C_{\text{scalar}}(n, l) &= l && \text{if } n \geq 0 \wedge l \geq 0 \\ C_{\text{mult}}(n) &= 0 && \text{if } n \geq 0 \end{aligned}$$

As we can see, this predicate will execute, in the worst case, as many processes as there are elements in the input list.

3 The Parametric Cost Relations Framework for Sequential Programs

The starting point of our work is the standard general framework described in [24] for setting up parametric relations representing the resource usage (and size relations) of programs and predicates.⁶

The framework is doubly parametric: first, the costs inferred are functions of input data sizes, and second, the framework itself is parametric with respect to the type of approximation made (upper or lower bounds), and to the resource analyzed. Each concrete resource r to be tracked is defined by two sets of (user-provided) functions, which can be constants, or general expressions of input data sizes:

1. *Head cost* $\varphi_{[ap,r]}(H)$: a function that returns the amount of resource r used by the unification of the calling literal (subgoal) \mathbf{p} and the head H of a clause matching \mathbf{p} , plus any preparation for entering a clause (i.e., call and parameter passing cost).
2. *Predicate cost* $\Psi_{[ap,r]}(\mathbf{p}, \bar{\mathbf{x}})$: it is also possible to define the *full cost* for a particular predicate \mathbf{p} for resource r and approximation ap , i.e., the function $\Psi_{[ap,r]}(\mathbf{p}) : \mathcal{N}_\top^m \rightarrow \mathcal{R}_\infty$ (with the sizes of \mathbf{p} 's input data as parameters, $\bar{\mathbf{x}}$) that returns the usage of resource r made by a call to this predicate. This is specially useful for built-in or external predicates, i.e., predicates for which the source code is not available and thus cannot be analyzed, or for providing a more accurate function than analysis can infer. In the implementation, this information is provided to the analyzer through *trust assertions*.

For simplicity we only show the equations related to our standard definition of cost. However, our framework has also been extended to allow the inference of a more general definition of cost, called accumulated cost, which is useful for performing static profiling, obtaining a more detailed information regarding how the cost is distributed among a set of user-defined *cost centers*. See [13, 22] for more details.

Consider a predicate \mathbf{p} defined by clauses C_1, \dots, C_m . Assume $\bar{\mathbf{x}}$ are the sizes of \mathbf{p} 's input parameters. Then, the resource usage (expressed in units of resource r with approximation ap) of a call to \mathbf{p} , for an input of size $\bar{\mathbf{x}}$, denoted as $C_{pred[ap,r]}(\mathbf{p}, \bar{\mathbf{x}})$, can be expressed as:

$$C_{pred[ap,r]}(\mathbf{p}, \bar{\mathbf{x}}) = \bigodot_{1 \leq i \leq m} (C_{cl[ap,r]}(C_i, \bar{\mathbf{x}})) \quad (1)$$

where $\bigodot = ClauseAggregator(ap, r)$ is a function that takes an approximation identifier ap and returns a function which applies over the cost of all the clauses, $C_{cl[ap,r]}(C_i, \bar{\mathbf{x}})$, for $1 \leq i \leq m$, in order to obtain the cost of a call to the predicate \mathbf{p} . For example, if ap is the identifier for approximation "upper bound" (ub), then a possible conservative definition for $ClauseAggregator(ub, r)$ is the \sum function. In this case, and since the number of solutions generated by a predicate that will

⁶ We give equivalent but simpler descriptions than in [24], which are allowed by assuming that programs are the result of a normalization process that makes all unifications explicit in the clause body, so that the arguments of the clause head and the body literals are all unique variables. We also change some notation for readability and illustrative purposes.

be demanded is generally not known in advance, a conservative upper bound on the computational cost of a predicate is obtained by assuming that all solutions are needed, and that all clauses are executed (thus the cost of the predicate is assumed to be the sum of the costs of all of its clauses). However, it is straightforward to take mutual exclusion into account to obtain a more precise estimate of the cost of a predicate, using the maximum of the costs of mutually exclusive groups of clauses, as done in [26].

Let us see now how to compute the resource usage of a clause. Consider a clause C of predicate \mathbf{p} of the form $H :- L_1, \dots, L_k$ where $L_j, 1 \leq j \leq k$, is a literal (either a predicate call, or an external or builtin predicate), and H is the clause head. Assume that $\psi_j(\bar{\mathbf{x}})$ is a tuple with the sizes of all the input arguments to literal L_j , given as functions of the sizes of the input arguments to the clause head. Note that these $\psi_j(\bar{\mathbf{x}})$ size relations have previously been computed during size analysis for all input arguments to literals in the bodies of all clauses. Then, the cost relation for clause C and a single call to \mathbf{p} (obtaining all solutions), is:

$$C_{cl[ap,r]}(C, \bar{\mathbf{x}}) = \varphi_{[ap,r]}(\text{head}(C)) + \sum_{j=1}^{\text{lim}(ap,C)} \text{sols}_j(\bar{\mathbf{x}}) \times C_{lit[ap,r]}(L_j, \psi_j(\bar{\mathbf{x}})) \quad (2)$$

where $\text{lim}(ap,C)$ gives the index of the last body literal that is called in the execution of clause C , $\psi_j(\bar{\mathbf{x}})$ are the sizes of the input parameters of literal L_j , and sols_j represents the product of the number of solutions produced by the predecessor literals of L_j in the clause body:

$$\text{sols}_j(\bar{\mathbf{x}}) = \prod_{i=1}^{j-1} s_{pred}(L_i, \psi_i(\bar{\mathbf{x}})) \quad (3)$$

where $s_{pred}(L_i, \psi_i(\bar{\mathbf{x}}))$ gives the number of solutions produced by L_i , with arguments of size $\psi_i(\bar{\mathbf{x}})$. The number of solutions and size relations are both inferred automatically by the framework (we refer the reader to [10, 8, 11, 26] for descriptions of this process).

Finally, $C_{lit[ap,r]}(L_j, \psi_j(\bar{\mathbf{x}}))$ is replaced by one of the following expressions, depending on L_j :

- If L_j is a call to a predicate q which is in the same strongly connected component as \mathbf{p} (the predicate under analysis), then $C_{lit[ap,r]}(L_j, \psi_j(\bar{\mathbf{x}}))$ is replaced by the symbolic call $C_{pred[ap,r]}(q, \psi_j(\bar{\mathbf{x}}))$, giving rise to a recurrence relation that needs to be bounded with a closed-form expression by the solver afterwards.
- If L_j is a call to a predicate q which is in a different strongly connected component than \mathbf{p} , then $C_{lit[ap,r]}(L_j, \psi_j(\bar{\mathbf{x}}))$ is replaced by the closed-form expression that bounds $C_{pred[ap,r]}(q, \psi_j(\bar{\mathbf{x}}))$. The analysis guarantees that this expression has been inferred beforehand, due to the fact that the analysis is performed for each strongly connected component, in a reverse topological order.
- If L_j is a call to a predicate q , whose cost is specified (with a trust assertion) as $\Psi_{[ap,r]}(q, \bar{\mathbf{x}})$, then $C_{lit[ap,r]}(L_j, \psi_j(\bar{\mathbf{x}}))$ is replaced by the expression $\Psi_{[ap,r]}(q, \psi_j(\bar{\mathbf{x}}))$.

4 Our Extended Resource Analysis Framework for Parallel Programs

In this section, we describe how we extend the resource analysis framework detailed above, in order to handle logic programs with Independent And-Parallelism, using the binary parallel $\&/2$ operator. First, we introduce a new general parameter that indicates the execution model the analysis has to consider. For our current prototype, we have defined two different execution models: standard *sequential* execution, represented by seq , and an abstract parallel execution model, represented by $par(n)$, where $n \in \mathbb{N} \cup \{\infty\}$. The abstract execution model $par(\infty)$ is similar to the *work* and *depth* model, presented in [6] and used extensively in previous work such as [18]. Basically, this model is based on considering an unbounded number of available processors to infer bounds on the depth of the computation tree. The *work* measure is the amount of work to be performed considering a sequential execution. These two measures together give an idea on the impact of the parallelization of a particular program. The abstract execution model $par(n)$, where $n \in \mathbb{N}$, assumes a finite number n of processors.

In order to obtain the cost of a predicate, equation (1) remains almost identical, with the only difference of adding the new parameter to indicate the execution model.

Now we address how to set the cost for clauses. In this case, equation (2) is extended with the execution model ex , and also the Σ default sequential cost aggregation is replaced by a parametric associative operator \oplus , that depends on the resource being defined, the approximation, and the execution model. For $ex \equiv par(\infty)$ or $ex \equiv seq$, the following equation is set up:

$$C_{cl[ap,r,ex]}(C, \bar{x}) = \varphi_{[ap,r]}(\text{head}(C)) + \overset{\text{lim}(ap,ex,C)}{\bigoplus_{j=1}} (sols_j(\bar{x}) \times C_{lit[ap,r,ex]}(L_j, \psi_j(\bar{x}))) \quad (4)$$

Note that the cost aggregator operators must depend on the resource r (besides the other parameters). For example, if r is *execution time*, then the cost of executing two tasks in parallel must be aggregated by taking the maximum of the execution times of the two tasks. In contrast, if r is *energy consumption*, then the aggregation is the addition of the energy of the two tasks.

Finally, we extend how the cost of a literal L_i , expressed as $C_{lit[ap,r,ex]}(L_i, \psi_i(\bar{x}))$, is set up. The previous definition is extended considering the new case where the literal is a call to the *meta-predicate* $\&/2$. In this case, we add a new parallel aggregation associative operator, denoted by \otimes . Concretely, if $L_i = B_1 \& B_2$, where B_1 and B_2 are two sequences of goals, then:

$$C_{lit[ap,r,ex]}(B_1 \& B_2, \bar{x}) = C_{body[ap,r,ex]}(B_1, \bar{x}) \otimes C_{body[ap,r,ex]}(B_2, \bar{x}) \quad (5)$$

$$C_{body[ap,r,ex]}(B, \bar{x}) = \overset{\text{lim}(ap,ex,B)}{\bigoplus_{j=1}} (sols_j(\bar{x}) \times C_{lit[ap,r,ex]}(L_j^B, \psi_j(\bar{x}))) \quad (6)$$

where $B = L_1^B, \dots, L_m^B$.

Consider now the execution model $ex \equiv par(n)$, where $n \in \mathbb{N}$ (i.e., assuming a finite number n of processors), and a recursive parallel predicate p that creates a parallel task q_i in each recursion i . Assume that we are interested in obtaining an upper bound on the cost of a call to p , for an input of size \bar{x} . We first infer the number k of parallel tasks created by p as a function of \bar{x} . This can be easily done by using our cost analysis framework and providing the suitable assertions for inferring a resource named “*ptasks*.” Intuitively, the “counter” associated to such resource must be incremented by the (symbolic) execution of the $\&/2$ parallel operator. More formally, $k = C_{pred[ub,ptasks]}(p, \bar{x})$. To this point, an upper bound m on the number of tasks executed by any of the n processors is given by $m = \lceil \frac{k}{n} \rceil$. Then, an upper bound on the cost (in terms of resolution steps, i.e., $r = steps$) of a call to p , for an input of size \bar{x} can be given by:

$$C_{pred[ub,r,par(n)]}(p, \bar{x}) = C^u + Spaw^u \quad (7)$$

where C^u can be computed in two possible ways: $C^u = \sum_{i=1}^m C_i^u$; or $C^u = m C_1^u$, where C_i^u denotes an upper bound on the cost of parallel task q_i , and C_1^u, \dots, C_k^u are ordered in descending order of cost. Each C_i^u can be considered as the sum of two components: $C_i^u = Sched_i^u + T_i^u$, where $Sched_i^u$ denotes the cost from the point in which the parallel subtask q_i is created until its execution is started by a processor (possibly the same processor that created the subtask), i.e. the cost of task preparation, scheduling, communication overheads, etc. T_i^u denotes the cost of the execution of q_i disregarding all the overheads mentioned before, i.e., $T_i^u = C_{pred[ub,r,seq]}(q, \psi_q(\bar{x}))$, where $\psi_q(\bar{x})$ is a tuple with the sizes of all the input arguments to predicate q in the body of p . $Spaw^u$ denotes an upper bound on the cost of creating the k parallel tasks q_i . It will be dependent on the particular system in which p is going to be executed. It can be a constant, or a function of several parameters, (such as input data size, number of input arguments, or number of tasks) and can be experimentally determined.

In addition, we propose a method for finding closed-form functions for cost relations that arise in the analysis of parallel programs, where the *max* function usually plays a role both as parallel and sequential cost aggregation operation, i.e., as \otimes and \oplus respectively. In the following subsection, we detail these methods.

4.1 Solving Cost Recurrence Relations Involving *max* Operation

Automatically finding closed-form upper and lower bounds for recurrence relations is an uncomputable problem. For some special classes of recurrences, exact solutions are known, for example for linear recurrences with one variable. For some other classes, it is possible to apply transformations to fit a class of recurrences with known solutions, even if this transformation obtains an appropriate approximation rather than an equivalent expression.

Particularly for the case of analyzing independent and-parallel logic programs, nonlinear recurrences involving the *max* operator are quite common. For example, if we are analyzing elapsed time of a parallel logic program, a proper parallel aggregation operator is the maximum between the times elapsed for each literal running in parallel. To the best of our knowledge, no general solution exists for recurrences of this particular type. However, in this paper we identify some common cases of this type of recurrences, for which we obtain closed forms that are proven

to be correct. In this section, we present these different classes, together with the corresponding method to obtain a correct bound.

Consider the following function $f : \mathbb{N}^m \rightarrow \mathbb{N}$, defined as a general form of a first-order recurrence equation with a *max* operator:

$$f(\bar{x}) = \begin{cases} \max(C, f(\bar{x}_{|i} - 1)) + D & x_i > \Theta \\ B & x_i \leq \Theta \end{cases} \quad (8)$$

where $\Theta \in \mathbb{N}$. C, D and B are arbitrary expressions possibly depending on \bar{x} . Note that $\bar{x} = x_1, x_2, \dots, x_m$. We define $\bar{x}_{|i} - 1 = x_1, \dots, x_i - 1, \dots, x_m$, for a given i , $1 \leq i \leq m$. If C and D do not depend on x_i , then C and D do not change through the different recursive instances of f . In this case, a closed-form upper bound is defined by the following theorem (whose proof is included in A):

Theorem 1. *Given $f : \mathbb{N}^m \rightarrow \mathbb{N}$ as defined in (8), where C and D are non-decreasing functions of $\bar{x} \setminus x_i$. Then, $\forall \bar{x}$:*

$$f(\bar{x}) = f'(\bar{x}) = \begin{cases} \max(C, B) + (x_i - \Theta) \cdot D & x_i > \Theta \\ B & x_i \leq \Theta \end{cases}$$

For the case where $C = g(\bar{x})$ and $D = h(\bar{x})$ are functions non-decreasing on x_i , then the upper bound is given by the following closed form:

Theorem 2. *Given $f : \mathbb{N}^m \rightarrow \mathbb{N}$ as defined in (8), where g and h are functions of \bar{x} , non-decreasing on x_i . Then, $\forall \bar{x}$:*

$$f(\bar{x}) \leq f'(\bar{x}) = \begin{cases} \max(g(\bar{x}), B) + (x_i - \Theta - 1) \times \max(g(\bar{x}), h(\bar{x}_{|i} - 1)) + h(\bar{x}_{|i}) & x_i > \Theta \\ B & x_i \leq \Theta \end{cases}$$

The proof of this Theorem is included in B.

For the remaining cases, where a $\max(e_1, e_2)$ appears, we try to eliminate the *max* operator by proving either $e_1 \leq e_2$ or $e_2 \leq e_1$, for any input. In order to do that, we use the function comparison capabilities of CiaoPP, presented in [20, 21]. In cases where e_1 and/or e_2 contains non-closed recurrence functions, we use the Z3 SMT solver [7] to find, if possible, a proof either for $e_1 \leq e_2$ or $e_2 \leq e_1$, treating the non-closed functions as uninterpreted functions, assuming that they are positive and non-decreasing. As the algorithm used by SMT solvers in this case is not guaranteed to terminate, we set a timeout. In the worst case, when no proof is found, then we replace the *max* operator with an addition, losing precision but still finding safe upper bounds.

5 Implementation and Experimental Results

We have implemented a prototype of our approach, leveraging the existing resource usage analysis framework of CiaoPP. The implementation basically consists of the parameterization of the operators used for sequential and parallel cost aggregation, i.e., for the aggregation of the costs corresponding to the arguments of $\&/2$ and $\&/2$, respectively. This allows the user to define resources in a general way, taking into account the underlying execution model.

Table 1. Description of the benchmarks.

map_add1/2	Parallel increment by one of each element of a list.
fib/2	Parallel computation of the nth Fibonacci number.
mmatrix/3	Parallel matrix multiplication.
blur/2	Generic parallel image filter.
add_mat/3	Matrix addition.
intersect/3	Set intersection.
union/3	Set union.
diff/3	Set difference.
dyade/3	Dyadic product of two vectors.
dyade_map/3	Dyadic product applied on a set of vectors.
append_all/3	Appends a prefix to each list of a list of lists.

Table 2. Resource usage inferred for Independent And-Parallel Programs.

Bench	Res	Bound Inferred	BigO	TA (ms)
map_add1(x)	SCost	$2 \cdot l_x + 1$	$\mathcal{O}(l_x)$	35.57
	PCost	$2 \cdot l_x + 1$	$\mathcal{O}(l_x)$	
	SThreads	l_x	$\mathcal{O}(l_x)$	
fib(x)	SCost	$F(i_x) + L(i_x) - 1$	$\mathcal{O}(2^{i_x})$	52.66
	PCost	$i_x + 1$	$\mathcal{O}(i_x)$	
	SThreads	$F(i_x) + L(i_x) - 1$	$\mathcal{O}(2^{i_x})$	
mmatrix(m_1, n_1, m_2, n_2)	SCost	$i_{n_2} \cdot i_{m_2} \cdot i_{m_1} + 2 \cdot i_{m_2} \cdot i_{m_1} + 2 \cdot i_{m_1} + 1$	$\mathcal{O}(i_{n_2} \cdot i_{m_2} \cdot i_{m_1})$	220.9
	PCost	$2 \cdot i_{m_1} + i_{n_1} + 1$	$\mathcal{O}(i_{n_1} + i_{m_1})$	
	SThreads	$i_{m_2} \cdot i_{m_1} + i_{m_1}$	$\mathcal{O}(i_{m_2} \cdot i_{m_1})$	
blur(m,n)	SCost	$2 \cdot i_m \cdot i_n + 2 \cdot i_n + 1$	$\mathcal{O}(i_m \cdot i_n)$	123.321
	PCost	$2 \cdot i_m + 2 \cdot i_n + 1$	$\mathcal{O}(i_m + i_n)$	
	SThreads	i_n	$\mathcal{O}(i_n)$	
add_mat(m,n)	SCost	$i_m \cdot i_n + 2 \cdot i_n + 1$	$\mathcal{O}(i_m \cdot i_n)$	62.72
	PCost	$i_m + 2 \cdot i_n + 1$	$\mathcal{O}(i_m + i_n)$	
	SThreads	i_n	$\mathcal{O}(i_n)$	
intersect(a,b)	SCost	$l_a \cdot l_b + 3 \cdot l_a + 3$	$\mathcal{O}(l_x)$	191.16
	PCost	$l_b + 2 \cdot l_a + 3$	$\mathcal{O}(i_n)$	
	SThreads	l_a	$\mathcal{O}(l_x)$	
union(a,b)	SCost	$l_a \cdot l_b + 3 \cdot l_a + 3$	$\mathcal{O}(l_a \cdot l_b)$	193.37
	PCost	$2 \cdot l_b + 2 \cdot l_a + 3$	$\mathcal{O}(l_a + l_b)$	
	SThreads	l_a	$\mathcal{O}(l_a)$	
diff(a,b)	SCost	$l_a \cdot l_b + 3 \cdot l_a + 3$	$\mathcal{O}(l_a \cdot l_b)$	191.16
	PCost	$l_b + 2 \cdot l_a + 3$	$\mathcal{O}(l_a + l_b)$	
	SThreads	l_a	$\mathcal{O}(l_a)$	
dyade(a,b)	SCost	$l_a \cdot l_b + 2 \cdot l_a + 1$	$\mathcal{O}(l_a \cdot l_b)$	71.08
	PCost	$l_b + l_a + 1$	$\mathcal{O}(l_a + l_b)$	
	SThreads	l_a	$\mathcal{O}(l_a)$	
dyade_map(l,m)	SCost	$i_{max(m)} \cdot l_m \cdot l_l + 2 \cdot l_m \cdot l_l + 2 \cdot l_m + 1$	$\mathcal{O}(i_{max(m)} \cdot l_m \cdot l_l)$	248.39
	PCost	$i_{max(m)} + l_m + l_l + 1$	$\mathcal{O}(i_{max(m)} + l_m + l_l)$	
	SThreads	$l_l \cdot l_m + l_l$	$\mathcal{O}(l_m \cdot l_l)$	
append_all(l,m)	SCost	$l_l \cdot l_m + 2 \cdot l_m + 1$	$\mathcal{O}(l_l \cdot l_m)$	108.4
	PCost	$l_l + l_m + 1$	$\mathcal{O}(l_l + l_m)$	
	SThreads	l_m	$\mathcal{O}(l_m)$	

$F(n)$ represents the nth element of the Fibonacci sequence.

$L(n)$ represents the nth Lucas number.

l_n, i_n represent the size of n in terms of the metrics *length* and *int*, respectively.

We selected a set of benchmarks that exhibit different common parallel patterns, briefly described in Table 1, together with the definition of a set of resources that help understand the overall behavior of the parallelization. Table 2 shows some results of the experiments that we have performed with our prototype implementation. Column **Bench** shows the main predicates analyzed for each benchmark. Set operations (**intersect**, **union** and **diff**), as well as the programs **append_all**,

Table 3. Resource usage inferred for a bounded number of processors.

Bench	Bound Inferred	BigO	T _A (ms)
map_add1(x)	$2 \cdot \lceil \frac{l_x}{p} \rceil + 1$	$\mathcal{O}(\lceil \frac{l_x}{p} \rceil)$	54.36
blur(m,n)	$2 \cdot \lceil \frac{l_m}{p} \rceil \cdot i_m + 2 \cdot \lceil \frac{l_n}{p} \rceil + 1$	$\mathcal{O}(\lceil \frac{l_n}{p} \rceil \cdot i_m)$	205.97
add_mat(m,n)	$\lceil \frac{l_n}{p} \rceil \cdot i_m + 2 \cdot \lceil \frac{l_m}{p} \rceil + 1$	$\mathcal{O}(\lceil \frac{l_n}{p} \rceil \cdot i_m)$	185.89
intersect(a,b)	$\lceil \frac{l_a}{p} \rceil \cdot l_b + 2 \cdot \lceil \frac{l_a}{p} \rceil + l_a + 2$	$\mathcal{O}(\lceil \frac{l_a}{p} \rceil \cdot l_b)$	330.47
union(a,b)	$\lceil \frac{l_a}{p} \rceil \cdot l_b + 2 \cdot \lceil \frac{l_b}{p} \rceil + l_a + l_b + 2$	$\mathcal{O}(\lceil \frac{l_a}{p} \rceil \cdot l_b)$	311.3
diff(a,b)	$\lceil \frac{l_a}{p} \rceil \cdot l_b + 2 \cdot \lceil \frac{l_a}{p} \rceil + l_a + 2$	$\mathcal{O}(\lceil \frac{l_a}{p} \rceil \cdot l_b)$	339.01
dyade(a,b)	$\lceil \frac{l_a}{p} \rceil \cdot l_b + 2 \cdot \lceil \frac{l_a}{p} \rceil + 1$	$\mathcal{O}(\lceil \frac{l_a}{p} \rceil \cdot l_b)$	120.93
append_all(l,m)	$\lceil \frac{l_m}{p} \rceil \cdot l_l + 2 \cdot \lceil \frac{l_m}{p} \rceil + 1$	$\mathcal{O}(\lceil \frac{l_m}{p} \rceil \cdot l_l)$	117.8

p is defined as the minimum between the number of processors and **SThreads**.

`dyade` and `add_mat`, are Prolog versions of the benchmarks analyzed in [18], which is the closest related work we are aware of.

Column **Res** indicates the name of each of the resources inferred for each benchmark: *sequential resolution steps* (**SCost**), *parallel resolution steps* assuming an unbounded number of processors (**PCost**), and *maximum number of processes executing in parallel* (**SThreads**). The latter gives an indication of the maximum parallelism that can potentially be exploited. Column **Bound Inferred** shows the upper bounds obtained for each of the resources indicated in Column **Res**. While in the experiments both upper and lower bounds were inferred, for the sake of brevity, we only show upper-bound functions. Column **BigO** shows the complexity order, in big O notation, corresponding to each resource. For all the benchmarks in Table 2 we obtain the exact complexity orders. We also obtain the same complexity order as in [18] for the Prolog versions of the benchmarks taken from that work. Finally, Column **T_A(ms)** shows the analysis times in milliseconds. The results show that most of the benchmarks have different asymptotic behavior in the sequential and parallel execution models. In particular, for `fib(x)`, the analysis infers an exponential upper bound for sequential execution steps, and a linear upper bound for parallel execution steps. As mentioned before, this is an upper bound for an ideal case, assuming an unbounded number of processors. Nevertheless, such upper-bound information is useful for understanding how the cost behavior evolves in architectures with different levels of parallelism. In addition, this *dual* cost measure can be combined together with a bound on the number of processors in order to obtain a general asymptotic upper bound (see for example Brent’s Theorem [14], which is also mentioned in [18]). The program `map_add1(1)` exhibits a different behavior: both sequential and parallel upper bounds are linear. This happens because we are considering *resolution steps*, i.e., we are counting each head unification produced from an initial call `map_add1(1)`. Even under the parallel execution model, we have a chain of head unifications whose length depends linearly on the length of the input list. It follows from the results of this particular case that this simple, non-associative parallelization will not be useful for improving the number of resolution steps performed in parallel.

Another useful information inferred in our experiments is the maximum number of processes that can be executed in parallel, represented by the resource named **SThreads**. We can see that for most of our examples the analysis obtains a linear upper bound for this resource, in terms of the size of some of the inputs. For example, the execution of `intersect(a,b)` (parallel set intersection) will create *at most* l_a processes, where l_a represents the length of the list a . For other examples,

the analysis shows a quadratic upper bound (as in `mmatrix`), or even exponential bounds (as in `fib`). The information about upper bounds on the maximum level of parallelism required by a program is useful for understanding its scalability in different parallel architectures, or for optimizing the number of processors that a particular call will use, depending on the size of the input data.

Finally, the results of our experiments considering a bounded number of processors are shown in Table 3.

6 Related Work

Our approach is an extension of an existing cost analysis framework for sequential logic programs [11, 13, 21], which extends the classical cost analysis techniques based on setting up and solving recurrence relations, pioneered by [28], with solutions for relations involving `max` and `min` functions. The framework handles characteristics such as backtracking, multiple solutions (i.e., non-determinism), failure, and inference of both upper and lower bounds including non-polynomial bounds. These features are inherited by our approach, and are absent from other approaches to parallel cost analysis in the literature.

The most closely-related work to our approach is [18], which describes an automatic analysis for deriving bounds on the worst-case evaluation cost of first order functional programs. The analysis derives bounds under an abstract *dual* cost model based on two measures: *work* and *depth*, which over-approximate the sequential and parallel evaluation cost of programs, respectively, considering an unlimited number of processors. Such an abstract cost model was introduced by [6] to formally analyze parallel programs. The work is based on type judgments annotated with a cost metric, which generate a set of inequalities which are then solved by linear programming techniques. Their analysis is only able to infer multivariate resource polynomial bounds, while non-polynomial bounds are left as future work. In [17] the authors propose an automatic analysis based on the *work* and *depth* model, for a simple imperative language with explicit parallel loops.

There are other approaches to cost analysis of parallel and distributed systems, based on different models of computation than the independent and-parallel model in our work. In [3] the authors present a static analysis which is able to infer upper bounds on the maximum number of *active* (i.e., not finished nor suspended) processes running in parallel, and the total number of processes created for imperative *async-finish* parallel programs. The approach described in [1] uses recurrence (cost) relations to derive upper bounds on the cost of concurrent object-oriented programs, with shared-memory communication and future variables. They address concurrent execution for loops with semi-controlled scheduling, i.e., with no arbitrary interleavings. In [4] the authors address the cost of parallel execution of object-oriented distributed programs. The approach is to identify the synchronization points in the program, use serial cost analysis of the blocks between these points, and then, exploiting the techniques mentioned, construct a graph structure to capture the possible parallel execution of the program. The path of maximal cost is then computed. The allocation of tasks to processors (called “locations”) is part of the program in these works, and thus, although independent and-parallel programs could be modeled in this computation style, it is not directly comparable to our more abstract model of parallelism.

Solving, or safely bounding recurrence relations with `max` and `min` functions has been addressed mainly for recurrences derived from divide-and-conquer algorithms [5, 27, 19]. In [2] the authors present solutions for Cost Relation Systems by obtaining upper bounds for both the number of nodes and the cost added in each node, in the derived evaluation tree. These bounds are then combined in order to obtain a closed-form, upper-bound expression. This closed form possibly contains maximization operations to express upper bounds for a set of subexpressions. However, each cost relation is defined as a summatory of costs, while in our work we consider other operations for aggregating the costs, including `max` operators. The presence of these operators often generates recurrence relations where the recursive calls are under the scope of such a `max` operator, for which we present a method to obtain a closed-form bound. This class of recurrences are not handled by most of the current computer algebra systems, as the authors in [2] mention.

7 Conclusions

We have presented a novel, general, and flexible analysis framework that can be instantiated for estimating the resource usage of parallel logic programs, for a wide range of resources, platforms, and execution models. To the best of our knowledge, this is the first approach to the cost analysis of *parallel logic programs*. Such estimations include both lower and upper bounds, given as functions on input data sizes. In addition, our analysis also infers other information which is useful for improving the exploitation and assessing the potential and actual parallelism of a program. We have also developed a method for solving the cost relations that arise in this particular type of analysis, which involve the *max* function. Finally, we have developed a prototype implementation of our general framework, instantiated it for the analysis of logic programs with Independent And-Parallelism, and performed an experimental evaluation, obtaining very encouraging results w.r.t. accuracy and efficiency.

References

1. E. Albert, P. Arenas, S. Genaim, M. Gómez-Zamalloa, and G. Puebla. Cost Analysis of Concurrent OO programs. In *Proc. of APLAS'11*, volume 7078 of *LNCS*, pages 238–254. Springer, 2011.
2. E. Albert, P. Arenas, S. Genaim, and G. Puebla. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, 46(2):161–203, 2011.
3. E. Albert, P. Arenas, S. Genaim, and D. Zanardini. Task-Level Analysis for a Language with Async-Finish parallelism. In *Proc. of LCTES'11*, pages 21–30. ACM Press, 2011.
4. E. Albert, J. Correas, E. Johnsen, K.I. Pu, and G. Román-Díez. Parallel cost analysis. *ACM Trans. Comput. Logic*, 19(4), November 2018.
5. L. Alonso, E.M. Reingold, and R. Schott. Multidimensional divide-and-conquer maximin recurrences. *SIAM J. Discret. Math.*, 8(3):428–447, August 1995.
6. Guy E. Blelloch and John Greiner. A provable time and space efficient implementation of NESL. In *ACM Int'l. Conf. on Functional Programming*, pages 213–225, May 1996.
7. L. Mendonça de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *TACAS*, volume 4963 of *LNCS*, pages 337–340. Springer, 2008.
8. S. K. Debray and N. W. Lin. Cost analysis of logic programs. *ACM TOPLAS*, 15(5):826–875, November 1993.

9. 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.
10. S. K. Debray, N.-W. Lin, and M. V. Hermenegildo. Task Granularity Analysis in Logic Programs. In *Proc. PLDI'90*, pages 174–188. ACM, June 1990.
11. S. K. Debray, P. Lopez-Garcia, M. V. Hermenegildo, and N.-W. Lin. Lower Bound Cost Estimation for Logic Programs. In *ILPS'97*, pages 291–305. MIT Press, 1997.
12. G. Gupta, E. Pontelli, K. Ali, M. Carlsson, and M. V. Hermenegildo. Parallel Execution of Prolog Programs: a Survey. *ACM TOPLAS*, 23(4):472–602, July 2001.
13. R. Haemmerlé, P. Lopez-Garcia, U. Liqat, M. Klemen, J. P. Gallagher, and M. V. Hermenegildo. A Transformational Approach to Parametric Accumulated-cost Static Profiling. In *FLOPS'16*, volume 9613 of *LNCS*, pages 163–180. Springer, 2016.
14. Robert Harper. *Practical Foundations for Programming Languages*. Cambridge University Press, 2 edition, 2016.
15. M. Hermenegildo, G. Puebla, F. Bueno, and P. Lopez Garcia. Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor). *Science of Computer Programming*, 58(1–2):115–140, 2005.
16. M. Hermenegildo and F. Rossi. Strict and Non-Strict Independent And-Parallelism in Logic Programs: Correctness, Efficiency, and Compile-Time Conditions. *Journal of Logic Programming*, 22(1):1–45, 1995.
17. T. Hoefler and G. Kwasniewski. Automatic complexity analysis of explicitly parallel programs. In *26th ACM Symp. on Parallelism in Algorithms and Architectures*, SPAA '14, pages 226–235, 2014.
18. J. Hoffmann and Z. Shao. Automatic static cost analysis for parallel programs. In Jan Vitek, editor, *Programming Languages and Systems*, pages 132–157, Berlin, Heidelberg, 2015. Springer Berlin Heidelberg.
19. H. Hwang and T.-H. Tsai. An asymptotic theory for recurrence relations based on minimization and maximization. *Theoretical Computer Science*, 290(3):1475 – 1501, 2003.
20. P. Lopez-Garcia, L. Darmawan, and F. Bueno. A Framework for Verification and Debugging of Resource Usage Properties. In *Technical Communications of ICLP*, volume 7 of *LIPICs*, pages 104–113. Schloss Dagstuhl, July 2010.
21. P. Lopez-Garcia, L. Darmawan, M. Klemen, U. Liqat, F. Bueno, and M. V. Hermenegildo. Interval-based Resource Usage Verification by Translation into Horn Clauses and an Application to Energy Consumption. *TPLP*, 18:167–223, March 2018.
22. P. Lopez-Garcia, M. Klemen, U. Liqat, and M. V. Hermenegildo. A General Framework for Static Profiling of Parametric Resource Usage. *TPLP*, 16(5-6):849–865, 2016.
23. M. Méndez-Lojo, J. Navas, and M. Hermenegildo. A Flexible (C)LP-Based Approach to the Analysis of Object-Oriented Programs. In *LOPSTR*, volume 4915 of *LNCS*, pages 154–168. Springer-Verlag, August 2007.
24. J. Navas, E. Mera, P. Lopez-Garcia, and M. Hermenegildo. User-Definable Resource Bounds Analysis for Logic Programs. In *Proc. of ICLP'07*, volume 4670 of *LNCS*, pages 348–363. Springer, 2007.
25. M. Rosendahl. Automatic Complexity Analysis. In *Proc. of FPCA '89*, pages 144–156. ACM Press, 1989.
26. A. Serrano, P. Lopez-Garcia, and M. V. Hermenegildo. Resource Usage Analysis of Logic Programs via Abstract Interpretation Using Sized Types. *TPLP, ICLP'14 Special Issue*, 14(4-5):739–754, 2014.
27. B.-F. Wang. Tight bounds on the solutions of multidimensional divide-and-conquer maximin recurrences. *Theoretical Computer Science*, 242(1):377 – 401, 2000.
28. B. Wegbreit. Mechanical Program Analysis. *Comm. of the ACM*, 18(9):528–539, 1975.

Appendices

A Proof for Theorem 1

Theorem Given $f : \mathbb{N}^m \rightarrow \mathbb{N}$ as defined in (8), where C and D are non-decreasing functions of $\bar{x} \setminus x_i$. Then, $\forall \bar{x}$:

$$f(\bar{x}) = f'(\bar{x}) = \begin{cases} \max(C, B) + (x_i - \Theta + 1) \cdot D & x_i > \Theta \\ B & x_i \leq \Theta \end{cases}$$

Proof. The proof for the case $x_i \leq \Theta$ is trivial.

In the following, we prove the theorem for $x_i > \Theta$, or equivalently, for $x_i \geq \Theta + 1$. The proof is by induction on this subset.

Base Case. We have to prove that $f(x_1, \dots, x_{i-1}, \Theta + 1, \dots, x_m) = f'(x_1, \dots, x_{i-1}, \Theta + 1, \dots, x_m)$. Using the definition of f and f' we have that

$$\begin{aligned} f(x_1, \dots, x_{i-1}, \Theta + 1, \dots, x_m) &= \max(C, f(x_1, \dots, x_{i-1}, \Theta, \dots, x_m)) + D \\ &= \max(C, B) + D \\ f'(x_1, \dots, x_{i-1}, \Theta + 1, \dots, x_m) &= \max(C, B) + (\Theta + 1 - \Theta) \cdot D \\ &= \max(C, B) + D \end{aligned}$$

General Case. Assuming

$f(x_1, \dots, x_{i-1}, x_i, \dots, x_m) = f'(x_1, \dots, x_{i-1}, x_i, \dots, x_m)$, we need to prove that $f(x_1, \dots, x_{i-1}, x_i + 1, \dots, x_m) = f'(x_1, \dots, x_{i-1}, x_i + 1, \dots, x_m)$. By induction hypothesis we have that:

$$\begin{aligned} f(x_1, \dots, x_{i-1}, x_i + 1, \dots, x_m) &= \max(C, f(x_1, \dots, x_{i-1}, x_i, \dots, x_m)) + D \\ &= \max(C, \max(C, B) + (x_i - \Theta) \cdot D) + D \\ &= \max(C, B) + (x_i - \Theta) \cdot D + D \\ &= \max(C, B) + (x_i - \Theta + 1) \cdot D \\ &= f'(x_1, \dots, x_{i-1}, x_i + 1, \dots, x_m) \end{aligned}$$

B Proof of Theorem 2

For all $a, b, c \in \mathbb{N} \cup \{0\}$, the following properties hold:

- Commutative: $\max(a, b) = \max(b, a)$
- Associative: $\max(a, \max(b, c)) = \max(\max(a, b), c)$
- Idempotent: $\max(a, a) = a$

Lemma 1. $\forall a, b, c \in \mathbb{N} : \max(a, b + c) \leq \max(a, b) + \max(a, c)$

Lemma 2. $\forall a, b, c, d \in \mathbb{N} : a \leq c \wedge b \leq d \implies \max(a, b) \leq \max(c, d)$

Theorem Given $f : \mathbb{N}^m \rightarrow \mathbb{N}$ as defined in (8), where g and h are functions of \bar{x} , non-decreasing on x_i . Then, $\forall \bar{x}$:

$$f(\bar{x}) \leq f'(\bar{x}) = \begin{cases} \max(g(\bar{x}), B) + (x_i - \Theta - 1) \times \max(g(\bar{x}), h(\bar{x}_i - 1)) + h(\bar{x}_i) & x_i > \Theta \\ B & x_i \leq \Theta \end{cases}$$

Proof. The proof for the case $x_i \leq \Theta$ is trivial.

In the following, we prove the theorem for $x_i > \Theta$, or equivalently, for $x_i \geq \Theta + 1$. The proof is by induction on this subset. For brevity, we only show the argument corresponding to the position of x_i in \bar{x} . However, the proof is still valid considering all of the arguments.

Base Case. We have to prove that $f(\Theta + 1) \leq f'(\Theta + 1)$. Using the definition of f and f' we have that

$$\begin{aligned} f(\Theta + 1) &= \max(g(\Theta + 1), f(\Theta)) + h(\Theta + 1) \\ &= \max(g(\Theta + 1), B) + h(\Theta + 1) \\ f'(\Theta + 1) &= \max(g(\Theta + 1), B) + ((\Theta + 1) - \Theta - 1) \times \max(g(\Theta + 1), h(\Theta)) + h(\Theta + 1) \\ &= \max(g(\Theta + 1), B) + h(\Theta + 1) \end{aligned}$$

General Case. Assuming $f(x) \leq f'(x)$, we need to prove that $f(x + 1) \leq f'(x + 1)$. By induction hypothesis and Lemma 2 we have that:

$$\begin{aligned} f(x + 1) &= \max(g(x + 1), f(x)) + h(x + 1) \\ &\leq \max(g(x + 1), \max(g(x), B) + (x - \Theta - 1) \times \max(g(x), h(x - 1)) + h(x)) + h(x + 1) \end{aligned}$$

By Lemma 1 we have that:

$$\begin{aligned} f(x + 1) &\leq \max(g(x + 1), \max(g(x), B) \\ &\quad + \max(g(x + 1), (x - \Theta - 1) \times \max(g(x), h(x - 1)))) \\ &\quad + \max(g(x + 1), h(x)) \\ &\quad + h(x + 1) \end{aligned} \tag{9}$$

Consider now the first term appearing in the sum of the right hand side of the inequality (9). Since \max is associative, and it holds that $\forall x : g(x+1) \geq g(x)$ (which follows from the hypothesis of the theorem), we obtain:

$$\begin{aligned} \max(g(x+1), \max(g(x), B)) &= \max(\max(g(x+1), g(x)), B) \\ &= \max(g(x+1), B) \end{aligned} \quad (10)$$

We consider now the second term in (9). By Lemma 1 we obtain:

$$\begin{aligned} &\max(g(x+1), (x - \Theta - 1) \times \max(g(x), h(x-1))) \\ &\leq (x - \Theta - 1) \times \max(g(x+1), \max(g(x), h(x-1))) \end{aligned}$$

As before, by associativity of \max , this is equivalent to:

$$(x - \Theta - 1) \times \max(g(x+1), h(x-1))$$

By Lemma 2, and $h(x-1) \leq h(x)$ (by hypothesis), we have that:

$$(x - \Theta - 1) \times \max(g(x+1), h(x)) \quad (11)$$

Replacing the results of (10) and (11) in (9):

$$\begin{aligned} f(x+1) &\leq \max(g(x+1), B) \\ &\quad + (x - \Theta - 1) \times \max(g(x+1), h(x)) \\ &\quad + \max(g(x+1), h(x)) + h(x+1) \\ &= \max(g(x+1), B) \\ &\quad + (x - \Theta) \times \max(g(x+1), h(x)) + h(x+1) \\ &= f'(x+1) \end{aligned}$$

$$\therefore f(x+1) \leq f'(x+1)$$