

A Simple and Scalable Static Analysis for Bound Analysis and Amortized Complexity Analysis

Moritz Sinn, Florian Zuleger, and Helmut Veith *

TU Vienna

Abstract. We present the first scalable *bound analysis* that achieves *amortized complexity analysis*. In contrast to earlier work, our bound analysis is not based on general purpose reasoners such as abstract interpreters, software model checkers or computer algebra tools. Rather, we derive bounds directly from abstract program models, which we obtain from programs by comparatively simple invariant generation and symbolic execution techniques. As a result, we obtain an analysis that is more predictable and more scalable than earlier approaches. We demonstrate by a thorough experimental evaluation that our analysis is fast and at the same time able to compute bounds for challenging loops in a large real-world benchmark. Technically, our approach is based on lossy vector addition systems (VASS). Our bound analysis first computes a lexicographic ranking function that proves the termination of a VASS, and then derives a bound from this ranking function. Our methodology achieves amortized analysis based on a new insight how lexicographic ranking functions can be used for bound analysis.

1 Introduction

Automatic methods for computing bounds on the resource consumption of programs are an active area of research [22, 19, 7, 20, 28, 8, 5, 21, 6]. We present the first scalable *bound analysis* for imperative programs that achieves *amortized complexity analysis*. Our techniques can be applied for deriving upper bounds on how often loops can be iterated as well as on how often a single or several control locations can be visited in terms of the program input.

The majority of earlier work on bound analysis has focused on mathematically intriguing frameworks for bound analysis. These analyses commonly employ general purpose reasoners such as abstract interpreters, software model checkers or computer algebra tools and therefore rely on elaborate heuristics to work in practice. In this paper we take an orthogonal approach that complements previous research. We propose a bound analysis based on a simple abstract program model, namely *lossy vector addition systems with states*. We present a static analysis with four well-defined analysis phases that are executed one after each other: program abstraction, control-flow abstraction, generation of a lexicographic ranking function and bound computation.

* Supported by the Austrian National Research Network S11403-N23 (RiSE) of the Austrian Science Fund (FWF) and by the Vienna Science and Technology Fund (WWTF) through grants PROSEED and ICT12-059.

A main contribution of this paper is a thorough experimental evaluation. We compare our approach against recent bounds analysis tools [7, 5, 6, 12], and show that our approach is faster and at the same time achieves better results. Additionally, we demonstrate the scalability of our approach by a comparison against our earlier tool [28], which to the best of our knowledge represents the only tool evaluated on a large publicly available benchmark of C programs. We show that our new approach achieves better results while increasing the performance by an order of magnitude. Moreover, we discuss on this benchmark how our tool achieves amortized complexity analysis in real-world code.

Our technical key contribution is a new insight how lexicographic ranking functions can be used for bound analysis. Earlier approaches such as [7] simply count the number of elements in the image of the lexicographic ranking function in order to determine an upper bound on the possible program steps. The same idea implicitly underlies the bound analyses [15, 19, 16, 20, 28, 6, 12]. However, this reasoning misses arithmetic dependencies between the components of the lexicographic ranking function (see Section 2). In contrast, our analysis calculates how much a lexicographic ranking function component is increased when another component is decreased. This enables amortized analysis.

Related Work. An interesting line of research studies the *amortized analysis* of first-order functional programs (e.g. [22, 21]) formulated as type rules over a template potential function with unknown coefficients; these coefficients are then found by linear programming. It is not clear how to transfer this approach to an imperative setting. Promising first steps for the amortized analysis of imperative programs are reported in [8]. Quantifier elimination is applied for simplifying a constraint system over template cost functions. Since quantifier elimination is expensive, the technique does not yet scale to larger programs.

Lexicographic ranking functions in automated termination analysis have been pioneered by Bradley et al. (see [10] and follow-up papers) who employ an elaborate constraint solving technique. A recent paper experimentally compares termination analysis by lexicographic ranking and transition invariants [13] implemented on top of a software model checker. [7] iteratively constructs a lexicographic ranking function by solving linear constraint systems. [11] is a hybrid of the approaches [13] and [7]. [10], [13] and [11] compute a lexicographic ranking function for a *single* control location (i.e., one loop header) at a time, while the application of bound analysis requires to find a common lexicographic ranking function for *all* control locations. [7] computes such a ranking function, but is limited to fairly small programs. Our approach complements the cited approaches as it represents a simple and scalable construction of a lexicographic ranking function for all control locations.

Bound Analysis. The COSTA project (e.g. [5, 6]) studies the extraction of cost recurrence relations from Java bytecode programs and proposes new methods for solving them with the help of computer algebra systems. [15] proposes to extend the polyhedra abstract domain with max- and non-linear expressions. [19] introduces multiple counters and exploits their dependencies such that upper bounds have to be computed only for restricted program parts. [16] pro-




Fig. 1. Our running example, ‘?’ denotes non-determinism (arising from a condition not modeled in the analysis). On the right we state the lossy VASS obtained by abstraction, Id denotes $a' \leq a, b' \leq b, i' \leq i$.

poses an abstract interpretation-guided program transformation that separates the different loop phases such that bounds can be computed for each phase in isolation. [20] employs proof-rules for bound computation combined with disjunctive abstract domains for summarizing inner loops. [28] proposes a bound analysis based on the size-change abstract domain. [17, 12] discuss how to alternate between bound analysis and invariant analysis for the mutual benefit of the computed bounds and invariants.

2 Motivation and Overview

The example presented in Figure 1 (encountered during our experiments) is challenging for an automated bound analysis: (C1) There are loops whose loop counter is modified by an inner loop: the innermost loop modifies the counter variables a and b of the two outer loops. Thus, the inner loop *cannot be ignored* (i.e., cannot be sliced away) during the analysis of the two outer loops. (C2) The middle loop with loop counter b requires a *path-sensitive* analysis to establish the linear loop bound n : it is not enough to consider how often the innermost loop can be executed (at most n^2 times) but rather how often the if-branch of the innermost loop (on which b is actually incremented) can be executed (at most n times). (C3) Current bound analysis techniques cannot model *increments* and instead approximate increments by *resets*, e.g., approximate the increment of b by an assignment to a value between 0 and n (using the fact that n is an upper bound of b)! Because of this overapproximation no bound analysis from the literature is able to compute the linear loop bound n for the middle loop. We now illustrate the main steps of our analysis:

1. *Program Abstraction:* First, our analysis abstracts the program to the VASS depicted in Figure 1. We introduce VASSs in Section 3. In this paper we are using *parameterized* VASSs, where we allow increments that are symbolic but constant throughout the program (such as $n - 1$). We extract lossy VASSs from C programs using simple invariant generation and symbolic execution techniques (described in Section 7).

2. Control Flow Abstraction: We propose a new abstraction for bound analysis, which we call *control flow abstraction* (CA) (described in Section 4). CA abstracts the VASS from Figure 1 into a transition system with four transitions: $\rho_1 \equiv a' \leq a - 1 \wedge b' \leq b + 1 \wedge i' \leq i$, $\rho_2 \equiv a' \leq a \wedge b' \leq b - 1 \wedge i' \leq i + (n - 1)$, $\rho_3 \equiv a' \leq a \wedge b' \leq b \wedge i' \leq i - 1$, $\rho_4 \equiv a' \leq a - 1 \wedge b' \leq b + 1 \wedge i' \leq i - 1$. CA effectively merges loops at different control locations into a single loop creating one transition for every cyclic path of every loop (without unwinding inner loops). This significantly simplifies the design of the later analysis phases.

3. Ranking Function Generation: Our ranking function generation (Algorithm 2 stated in Section 5) finds an *order* on the transitions resulting from CA such that there is a variable for every transition, which decreases on that transition and does not increase on the transitions that are lower in the order. This results in the lexicographic ranking function $l = \langle a, a, b, i \rangle$ for the transitions $\rho_1, \rho_4, \rho_2, \rho_3$ in that order. Our soundness theorem (Theorem 1) guarantees that l proves the termination of Figure 1.

4. Bound Analysis: Our bound analysis (Algorithm 3 stated in Section 6) computes a bound for every transition ρ by adding for every other transition τ how often τ increases the variable of ρ and by how much. In this way, our bound analysis computes the bound n for ρ_2 , because ρ_2 can be incremented by ρ_1 and ρ_4 , but this can only happen n times, due to the initial value n of a . Further, our bound analysis computes the bound $n * (n - 1)$ for ρ_3 from the fact that only ρ_2 can increase the counter i by $n - 1$ and that ρ_2 has the already computed transition-bound n . Our soundness result (Theorem 2) guarantees that the bound n obtained for ρ_2 is indeed a bound on how often the middle loop of Figure 1 can be executed.

Our bound analysis solves the challenges (C1)-(C3): CA allows us to analyze all loops at once (C1) creating one transition for every loop path (C2). The abstract model of lossy VASS is precise enough to model counter increments, which is a key requirement for achieving amortized complexity analysis (C3).

2.1 Amortized Complexity Analysis

In his influential paper [27] Tarjan introduces amortized complexity analysis using the example of a stack, which supports two operations *push* (which puts an element on the stack) and *popMany* (which removes several elements from the stack). He assumes that the cost of *push* is 1 and the cost of *popMany* is the number of removed elements. We use his example (see Figure 2) to discuss how our bound analysis achieves amortized analysis: Our analysis first abstracts the program to a VASS and then applies CA. This results in the three transitions $\rho_1 \equiv i' = i - 1 \wedge n' = n + 1$, $\rho_2 \equiv i' = i - 1 \wedge n' = n$, $\rho_3 \equiv i' = i \wedge n' = n - 1$ (the first two transitions come from the outer loop, the last transition from the inner loop). Algorithm 2 then computes the lexicographic ranking function $\langle i, i, n \rangle$ for the transitions ρ_1, ρ_2, ρ_3 in that order. Our bound analysis (Algorithm 3) then computes the joint bound m for the transitions ρ_1 and ρ_2 . Our bound analysis further computes the bound m for transition ρ_3 from the fact that only ρ_1 can increase the counter n by 1 and that ρ_1 has the already computed bound m .

```

void main(int m) {
    int i=m, n = 0;    //stack = emptyStack();
l1:  while (i > 0) {
        i--;
        if (?) //push
            n++;      //stack.push(element);
        else //popMany
l2:    while (n > 0 && ?)
            n--;      //element = stack.pop();
    }
}

```

Fig. 2. Model of Tarjan’s stack example [27] for amortized complexity analysis

Adding these two bounds gives the amortized complexity bound $2m$ for Figure 2. We highlight that our analysis has actually used the variable n of transition ρ_3 as a *potential function* (see [27] for a definition)! A lexicographic ranking function $\langle x_1, \dots, x_n \rangle$ can be seen as a *multidimensional potential function*. Consider, for example, the ranking function $\langle a, a, b, i \rangle$ for the transitions $\rho_1, \rho_4, \rho_2, \rho_3$ of Figure 1. The potential of ρ_3 can be increased by ρ_2 whose potential in turn can be increased by ρ_1 and ρ_4 .

3 Lossy VASSs and Basic Definitions

In this section we define lossy VASSs (introduced in [9]) and state definitions that we need later on. We will often drop the ‘lossy’ in front of ‘VASS’ because we do not introduce non-lossy VASSs and there is no danger of confusion. In this paper, we will use VASSs as *minimal program model* for bound analysis of sequential programs without procedures. We leave the extension to concurrent and recursive programs for future work.

Definition 1 (Lossy Vector Addition System with States (VASS)). *We fix some finite set of variables $Var = \{x_1, \dots, x_n\}$. A lossy vector addition system with states (VASS) is a tuple $P = (L, E)$, where L is a finite set of locations, and $E \subseteq L \times \mathbb{Z}^n \times L$ is a finite set of transitions. We write $l_1 \xrightarrow{d} l_2$ to denote an edge (l_1, d, l_2) for some vector $d \in \mathbb{Z}^n$. We often specify the vector $d \in \mathbb{Z}^n$ by predicates $x'_i \leq x_i + d_i$ with $d_i \in \mathbb{Z}$.*

A path of P is a sequence $l_0 \xrightarrow{d_0} l_1 \xrightarrow{d_1} \dots$ with $l_i \xrightarrow{d_i} l_{i+1} \in E$ for all i . A path is cyclic, if it has the same start- and end-location. A path is simple, if it does not visit a location twice except for start- and end-location. We write $\pi = \pi_1 \cdot \pi_2$ for the concatenation of two paths π_1 and π_2 , where the end-location of π_1 is the start-location of π_2 . We say π' is a subpath of a path π , if there are paths π_1 and π_2 with $\pi = \pi_1 \cdot \pi' \cdot \pi_2$.

The set of valuations of Var is the set $V_{Var} = Var \rightarrow \mathbb{N}$ of mappings from Var to the natural numbers. A trace of P is a sequence $(l_0, \sigma_0) \xrightarrow{d_0} (l_1, \sigma_1) \xrightarrow{d_1} \dots$

such that $l_0 \xrightarrow{d_0} l_1 \xrightarrow{d_1} \dots$ is path of P , $\sigma_i \in V_{Var}$ and $\sigma_{i+1}(x_j) \leq \sigma_i(x_j) + d_i$ for all i and $1 \leq j \leq n$. P is terminating, if there is no infinite trace of P .

Values of VASS variables are always non-negative. We describe how to obtain VASSs from programs by abstraction in Section 7. The non-negativity of VASS values has two important consequences: (1) Transitions in VASSs contain implicit guards: for example a transition $x' \leq x + c$ can only be taken if $x + c \geq 0$. (2) VASS transitions can be used to model variable increments as well as variable resets: we replace the assignment $x := k$, where $k \in \mathbb{Z}$, by the VASS transition $x' \leq x + k$ during program abstraction (we point out that lossiness is essential for abstracting assignments). This only increases the set of possible program traces and thus provides a conservative abstraction.

Parameterized VASSs. In our implementation we use a slight generalization of lossy VASSs. We allow the increment n in a transition predicate $x' \leq x + n$ to be symbolic but constant; in particular, we require that n does not belong to the set of variables Var . Our bound algorithm works equally well with symbolic increments under the condition that we know the sign of n . We call these extended systems *parameterized VASSs*. See Figure 1 for an example.

In the following we introduce some standard terminology that allows us to precisely speak about loops and related notions.

Definition 2 (Reducible Graph, Loop Header, Natural Loop, Loop-nest Tree, e.g. [4]). Let $G = (V, E)$ be a directed graph with a unique entry point such that all nodes are reachable from the entry point. A node a dominates a node b , if every path from entry to b includes a . An edge $l_1 \rightarrow l_2$ is a back edge, if l_2 dominates l_1 . G is reducible, if G becomes acyclic after removing all back edges. A node is a loop header, if it is the target of a back edge. The (natural) loop of a loop header h in a reducible graph is the maximal set of nodes L such that for all $x \in L$ (1) h dominates x and (2) there is a back edge from some node n to h such that there is a path from x to node n that does not contain h .

In the rest of this paper we restrict ourselves to VASSs and programs whose control flow graph is reducible. This choice is justified by the fact that irreducible control flow is *very rare* in practice (e.g. see the study in [26]). For analyzing irreducible programs we propose to use program transformations that make the program reducible; we do not elaborate this idea further due to lack of space.

Next, we define a special case of path, which corresponds to the notion of bound used in this paper (defined below).

Definition 3 (Loop-path). A loop-path π is a simple cyclic path, which starts and ends at some loop header l , and visits only locations inside the natural loop of l .

Example: $l_2 \xrightarrow{\tau_2} l_3 \xrightarrow{Id} l_2$ is a loop-path for the VASS in Figure 1. However, $l_2 \xrightarrow{Id} l_1 \xrightarrow{\tau_1} l_2$ is not a loop-path because it does not stay inside the natural loop of l_2 . $l_2 \xrightarrow{\tau_2} l_3 \xrightarrow{Id} l_4 \xrightarrow{\tau_4} l_3 \xrightarrow{Id} l_2$ is not a loop-path, because it is not simple (l_3 is visited twice).

```

Procedure: CA( $P$ )
Input: a reducible VASS  $P$ 
Output: a transition system  $\mathcal{T}$ 
 $\mathcal{T} := \emptyset$ ;
foreach loop header  $l$  in  $P$  do
  foreach loop-path  $\pi = l \xrightarrow{d_1} l_1 \cdots l_{n-1} \xrightarrow{d_n} l$  do
     $\mathcal{T} := \mathcal{T} \cup \{d_1 + \cdots + d_n\}$ ;
return  $\mathcal{T}$ ;

```

Algorithm 1: CA creates a *transition system* from a given VASS

Definition 4 (Instance of a loop-path). Let $\pi = l_1 \xrightarrow{d_1} l_2 \xrightarrow{d_2} \cdots l_{n-1} \xrightarrow{d_{n-1}} l_1$ be a loop-path. A path ν is an instance of π iff ν is of the form $l_1 \xrightarrow{d_1} l_2 * l_2 \xrightarrow{d_2} l_3 * l_3 \cdots l_{n-1} * l_{n-1} \xrightarrow{d_{n-1}} l_n = l_1$, where $l_i * l_i$ denotes any (possibly empty) path starting and ending at location l_i which does not contain l_1 . A path p contains an instance ν of π iff ν is a subpath of p . Let ν be an instance of π contained in p ; a transition t on p belongs to ν , if t is on ν and $t = l_i \xrightarrow{d_i} l_{i+1}$ for some i .

We note the following facts about instances: Every transition in a path belongs to *at most one* instance of a loop-path. Every transition in a given cyclic path belongs to *exactly one* instance of a loop-path.

Example: There are four instances of loop-paths in the path $\pi = l_1 \xrightarrow{\tau_1} l_2 \xrightarrow{\tau_2} l_3 \xrightarrow{Id} l_4 \xrightarrow{\tau_3} l_3 \xrightarrow{Id} l_2 \xrightarrow{\tau_2} l_3 \xrightarrow{Id} l_2 \xrightarrow{Id} l_1$ of the VASS in Figure 1: $l_1 \xrightarrow{\tau_1} l_2 \xrightarrow{Id} l_1$, $l_2 \xrightarrow{\tau_2} l_3 \xrightarrow{Id} l_2$ (twice) and $l_3 \xrightarrow{Id} l_4 \xrightarrow{\tau_3} l_3$.

Definition 5 (Path-bound). A path-bound for a loop-path π is an expression b over Var such that for every trace $(l_0, \sigma_0) \xrightarrow{d_0} \cdots$ of P the path $l_0 \xrightarrow{d_0} \cdots$ contains at most $b(\sigma_0)$ instances of π .

Path-bounds have various applications in bound and complexity analysis: the *computational complexity* of a program can be obtained by adding the bounds of the loop-paths of all loops; a *loop bound* can be obtained by adding the bounds of all loop-paths of a given loop; the number of *visits* to a *single control location* l can be obtained by adding the bounds of the loop-paths that include l (our notion of a path-bound can be seen as a path-sensitive generalization of the notion of a “reachability-bound” [20]); similarly one can obtain a bound on the number of *visits* to a *set of control locations*. More generally, one can obtain resource bounds for a given *cost model* by multiplying the bound on the number of visits to a control location with the *cost* for visiting this location.

4 Control Flow Abstraction

Control flow abstraction (CA), stated in Algorithm 1, is based on two main ideas: (1) Given a program P , CA results into one transition for every loop-path

Procedure: $\text{Ranking}(\mathcal{T})$ Input: a transition system $\mathcal{T}$ Output: a lexicographic ranking function l, which has one component for every transition $\rho \in \mathcal{T}$ $\mathcal{S} := \mathcal{T};$ $l := \text{"lexicographic ranking function with no components"};$ while <i>there is a transition $\rho \in \mathcal{S}$ and a variable x such that $\rho \models x' < x$ and for all $\rho' \in \mathcal{S}$ we have $\rho' \models x' \leq x$</i> do $\mathcal{S} := \mathcal{S} \setminus \rho;$ $l := l.\text{append}(x);$ if $\mathcal{S} = \emptyset$ then return $l;$ else return "Transitions $\mathcal{S}$ maybe non-terminating";
--

Algorithm 2: Ranking computes a lexicographic ranking function

π for all loop headers l of P . This enables a path-sensitive analysis, which ensures high precision during ranking function generation and bound analysis. (2) The control structure is abstracted: effectively, all loops are merged into a single loop. This allows to compute a common lexicographic ranking function for all loops later on. CA maps VASSs to *transition systems*. Transition systems are not meant to be executed; their sole purpose is to be used for ranking function generation and bound analysis.

Definition 6 (Transition System). A transition system is a set of vectors $d \in \mathbb{Z}^n$. We often specify a transition $d \in \mathbb{Z}^n$ by predicates $x'_i \leq x_i + d_i$, where $d_i \in \mathbb{Z}$. We also write $d \models x'_i \leq x_i$ (resp. $d \models x'_i < x_i$) for $d_i \leq 0$ (resp. $d_i < 0$).

Loop-path Contraction. Algorithm 1 creates one transition for every loop-path $\pi = l \xrightarrow{d_1} l_1 \cdots l_{n-1} \xrightarrow{d_n} l$. The transition $d_1 + \cdots + d_n$ represents the accumulated effect of all variable increments along the path. The key idea of loop-path contraction is to *ignore any inner loop* on π . We will incorporate the effects of the inner loops only later on during the ranking function generation and bound analysis phase.

CA represents our choice of precision in the analysis: CA facilitates a high degree of disjunctiveness in the analysis, where we keep one disjunct for every loop-path. By encapsulating the level of precision in a single analysis phase, we achieve a modular analysis (only during CA we need to deal with the control structure of the VASS). This simplifies the design of the later termination and bound analysis and also allows us to easily adjust the analysis precision if the number of paths is prohibitively high (see the discussion on *path merging* in [25]).

5 Ranking Function Generation

In this section we introduce our algorithm for ranking function generation: Algorithm 2 reads in a transition system obtained from CA and returns a lexico-

graphic ranking function that provides a witness for termination. The key idea of the algorithm is to incrementally construct a lexicographic ranking function from local ranking functions. We call a variable x a *local ranking function* for a transition ρ , if $\rho \models x' < x$. A tuple $l = \langle y_1, y_2, \dots, y_k \rangle \in \text{Var}^k$ is a *lexicographic ranking function* for a transition system $\mathcal{T}$ iff for each $\rho \in \mathcal{T}$ there is a ranking function component y_i that is a local ranking function for ρ and $\rho \models y'_j \leq y_j$ for all $j < i$. Algorithm 2 maintains a candidate lexicographic ranking function l and a set of transitions $\mathcal{S}$ for which no ranking function component has been added to l . In each step the algorithm checks if there is a transition ρ in $\mathcal{S}$ and a variable x such that (1) x is a local ranking function for ρ and (2) no remaining transition increases the value of x , i.e., the condition $\forall \rho' \in \mathcal{S}. \rho' \models x' \leq x$ is satisfied. If (1) and (2) are satisfied, ρ is removed from the set of remaining transitions $\mathcal{S}$ and x is added as the component for ρ in the lexicographic ranking function l . Conditions (1) and (2) ensure that the transition ρ cannot be taken infinitely often if only transitions from $\mathcal{S}$ are taken. The algorithm stops, if no further transition can be removed. If $\mathcal{S}$ is empty, the lexicographic ranking function l is returned. Otherwise it is reported that the remaining transitions $\mathcal{S}$ might lead to non-terminating executions.

Next we state the correctness of the combined application of Algorithm 1 and Algorithm 2. The proof can be found in [25].

Theorem 1. *If Algorithm 2 returns a lexicographic ranking function l for the transition system $\mathcal{T}$ obtained from Algorithm 1 then VASS P is terminating.*

Reasons for Failure. There are two reasons why our ranking function generation algorithm may fail: (1) There is a transition ρ without a *local ranking function*, i.e., there is no variable x with $\rho \models x' < x$. Such a transition ρ will never be removed from $\mathcal{S}$. (2) There is a *cyclic dependency* between local ranking functions, i.e., for every transition $\rho \in \mathcal{S}$ there is a local ranking function x but the condition “ $\rho' \models x' \leq x$ for all $\rho' \in \mathcal{S}$ ” is never satisfied. We found cyclic dependencies to be very rare in practice (only 4 instances); we provide a discussion of the failures encountered in our experiments in [25].

Non-determinism. We note that in presence of transitions with more than one local ranking function, the result of Algorithm 2 may depend on the choice for x . However, it is straight-forward to extend Algorithm 2 to generate all possible lexicographic ranking functions.

6 Bound Computation

In this section we introduce our bound algorithm: Algorithm 3 computes a bound b for a transition ρ of the transition system $\mathcal{T}$. The main idea of Algorithm 3 is to rely only on the components of the lexicographic ranking function l for bound computation. Let x be the component of ρ in l . We recall that the termination algorithm has already established that x is a local ranking function for ρ and therefore we have $\rho \models x > x'$. Thus ρ can be executed at most $\text{InitialValue}(x)$

```

Procedure: Bound( $\rho$ )
Input: a transition  $\rho$ 
Output: a bound for transition  $\rho$ 
Global: transition system  $\mathcal{T}$ , lexicographic ranking function  $l$ 
 $x :=$  ranking function component of  $\rho$  in  $l$ ;
 $b := \text{InitialValue}(x)$ ;
foreach transition  $\rho' \in \mathcal{T}$  with  $\rho' \not\models x' \leq x$  do
  | Let  $k \in \mathbb{N}$  s.t.  $x' \leq x + k$  in  $\rho'$ ;
  |  $b := b + \text{Bound}(\rho') \cdot k$ ;
Let  $k \in \mathbb{N}$  s.t.  $x' \leq x - k$  in  $\rho$ ;
return  $b = b/k$ ;

```

Algorithm 3: Bound returns a bound for transition ρ

often unless x is increased by other transitions: Algorithm 3 initializes $b := \text{InitialValue}(x)$ and then checks for every other transition ρ' if it increases x , i.e., $\rho' \not\models x' \leq x$. For every such transition ρ' Algorithm 3 recursively computes a bound, multiplies this bound by the height of the increase k and adds the result to b . Finally, we divide b by the decrease k of x on transition ρ .

Termination. Algorithm 3 terminates because the recursive calls cannot create a cycle. This is because Algorithm 3 uses only the components of l for establishing bounds and the existence of the lexicographic ranking function l precludes cyclic dependencies.

Soundness. Our soundness result (Theorem 2, for a proof see [25]) rests on the assumption that the CFG of P is an SCC whose unique entry point is also its unique exit point. We can always ensure this condition by a program transformation that encloses P in a dummy while-loop $\text{while}(y > 0)\{P; y--;\}$, where y is a fresh variable with $\text{InitialValue}(y) = 1$. We point out that this program transformation enable us to compute path-bounds in terms of the program inputs for CFGs with multiple SCCs (e.g., a program with two successive loops).

Theorem 2. *Let b be a bound computed by Algorithm 3 for a transition ρ obtained from a loop-path π during CA. Then b is a path-bound for π .*

Greedy Bound Computation. The bound computed by Algorithm 3 depends on the lexicographic ranking function l . Clearly, it is possible to run the algorithm for multiple lexicographic ranking functions and choose the minimum over the generated bounds. However, we found the greedy approach to work well in practice and did not see a need for implementing the enumeration strategy.

Complexity of the Algorithm / Size of Bound Expressions. For ordinary VASS, the complexity of Algorithm 3 is polynomial in the size of the input with a small exponent (depending on the exact definition of the complexity parameters). Unfortunately, this statement does not hold for *parameterized* VASSs, for which

bound expressions can be exponentially big: We consider n transitions $\rho_1, \dots, \rho_n$ with the local ranking functions $x_1, \dots, x_n$ and the lexicographic ranking function $\langle x_1, \dots, x_n \rangle$. We assume that transition ρ_i increments x_j by some constant c_{ij} for $i < j$. Then, Algorithm 3 computes the bound stated in the following formula, which is exponentially big for symbolic coefficients c_{ij} :

$$b(\rho_n) = \sum_{k \in [0, n-1]} \prod_{i_1 < \dots < i_k \in [1, n-1]} \text{InitialValue}(x_{i_1})^{c_{i_1 i_2}} \dots c_{i_k n}$$

However, in practical examples the variable dependencies are *sparse*, i.e., most coefficients c_{ij} are zero (confirmed by our experiments). We highlight that Algorithm 3 exploits this sparsity as it does not compute the bound using the explicit formula stated above but rather computes the bound for the current transition ρ using only the bounds of the transitions that actually increase the counter of ρ (i.e., $c_{ij} > 0$). We note that in our experiments the computed bounds are small and the running time of Algorithm 3 is basically linear in the number of transitions. We conclude that in practice one should make use of the fine-grained precision offered by the possibly exponentially-sized bound expressions.

Preprocessing: Merging Transitions. Before the bound computation our analysis applies the following rule until no more transitions can be merged: Let ρ_1 and ρ_2 be two transitions with the same local ranking function x in l such that $x' \leq x + k \in \rho_1$ and $x' \leq x + k \in \rho_2$ for some k (i.e., both transitions decrement x by the same amount). We replace ρ_1 and ρ_2 by the transition $\rho = \{y' \leq y + \max\{k_1, k_2\} \mid y' \leq y + k_1 \in \rho_1 \wedge y' \leq y + k_2 \in \rho_2\}$. It is not difficult to see that merging transitions is sound and always improves the bound computed by Algorithm 3 (we do not give a formal proof here for lack of space).

Example: We have obtained the *loop bound* of the middle loop in Figure 1 from the path-bound n of its single transition ρ_2 (see Section 2). We have obtained $2m$ as the *amortized complexity* of Figure 2 by adding the path-bounds of its transitions ρ_1, ρ_2, ρ_3 applying merging to ρ_1 and ρ_2 (see Section 2.1).

7 Program Abstraction

In this section we describe how to abstract programs to VASSs.

Definition 7 (Program). Let Σ be a set of states. The set of transition relations $\Gamma = 2^{\Sigma \times \Sigma}$ is the set of relations over Σ . A program is a tuple $P = (L, E)$, where L is a finite set of locations, and $E \subseteq L \times \Gamma \times L$ is a finite set of transitions. We write $l_1 \xrightarrow{\rho} l_2$ to denote a transition (l_1, ρ, l_2) . We assume the set of reachable states $\text{Reach}(l)$ is defined for every location $l \in L$ in the standard way. Let $e_1, e_2 \in \Sigma \rightarrow \mathbb{Z}$ be integer-valued expressions over the states, and let $c \in \mathbb{Z}$ be some integer. We say $e_1 \geq 0$ is invariant for l , if $e_1(s) \geq 0$ holds for all $s \in \text{Reach}(l)$. We say $e_2' \leq e_1 + c$ is invariant for $l_1 \xrightarrow{\rho} l_2$, if $e_2(s_2) \leq e_1(s_1) + c$ holds for all $(s_1, s_2) \in \rho$ with $s_1 \in \text{Reach}(l_1)$. We say e_1 is a norm, if $e_1 \geq 0$ is invariant for every location l .

Definition 8 (Abstraction of a Program). A VASS $V = (L, E')$ with variables Var is an abstraction of a program $P = (L, E)$ iff (1) every $x \in Var$ is a norm and (2) for each transition $l_1 \xrightarrow{\rho} l_2 \in E$ there is a transition $l_1 \xrightarrow{d} l_2 \in E'$ such that every $x' \leq x + c \in d$ is invariant for $l_1 \xrightarrow{\rho} l_2$.

The above definition suggests a three-step methodology for abstracting programs: (1) Guess a set of norms $N \subseteq \Sigma \rightarrow \mathbb{Z}$. (2) For every $x \in N$ show that $x \geq 0$ is invariant at all locations l . If this is not the case, discard the norm x . (3) For every $x \in N$ and every transition $l_1 \xrightarrow{\rho} l_2$ find a constant expression c such that $x' \leq x + c$ is invariant for $l_1 \xrightarrow{\rho} l_2$. Next, we describe how we implement this methodology.

7.1 Abstracting Programs to VASSs: Our Implementation

Guessing Norms. The key idea of Algorithm 2 is to find a local ranking function for every transition. We recall that a transition is obtained from a loop-path during CA. For this reason, our main heuristic is to consider expressions as norms that are local ranking functions for at least one loop-path of the program under analysis. Our implementation iterates over all loop-paths $\pi = l \xrightarrow{\rho_1} l_1 \xrightarrow{\rho_2} \dots l_{n-1} \xrightarrow{\rho_n} l$: Let $\mathbf{rel}(\pi) = \rho_1 \circ \dots \circ \rho_n$ be the transition relation obtained by contracting all transition relations along π . We implement the computation of $\mathbf{rel}(\pi)$ by *symbolic backward execution*, which returns a set of *guards* $e \geq 0$ (we note that guards are normalized, e.g., $n \geq i$ is transformed into $n - i \geq 0$) and *updates* $x' = e$, where e is some expression over the program variables and x' denotes the value of x after executing π . A *local ranking function* is an expression r such that (a) $r \geq 0$ is a guard of $\mathbf{rel}(\pi)$ and (b) $\delta_r = r - r' > 0$, where r' denotes the expression r where every variable x is replaced by expression e according to the update $x' = e$ of $\mathbf{rel}(\pi)$. For every local ranking function r our implementation adds the expression $\max\{r + \delta_r, 0\}$ to the set of norms N . Clearly, all norms $x = \max\{r + \delta_r, 0\} \in N$ satisfy the invariant $x \geq 0$.

Abstracting Transitions. In our implementation we derive a transition predicate $x' \leq x + c$ for a given norm $x = \max\{e, 0\} \in N$ and transition $l_1 \xrightarrow{\rho} l_2$ as follows: We obtain the expression e' from e by replacing variables with their updates according to ρ . The expression e' either constitutes an *increment*, i.e., $e' = e + k_1$, or a *reset*, i.e., $e' = k_2$, for some expression k_i . For now, assume k_i is constant. We proceed by a case distinction: If $e' = e + k_1$ and $e + k_1 \geq 0$ is invariant for $l_1 \xrightarrow{\rho} l_2$, then our implementation derives the transition predicate $x' \leq x + k_1$. This derivation is sound, because of the invariant $e + k_1 \geq 0$. (We motivate this derivation rule as follows: assume $r \geq 0$ and $\delta_r = r - r'$ hold on ρ , we have $e' = e + (-\delta_r) \geq 0$ for $e = r + \delta_r$.) Otherwise, our implementation derives the transition predicate $x' \leq x + \max\{k_i, 0\}$. This derivation is sound because of properties of maxima. If k_i is not constant, we first search for an invariant $k_i \leq u$ with u constant, and then proceed as above (replacing k_i with u). We implement invariant analysis by symbolic backward execution (see [25]).

	Bounded	1	$\log n$	n	$n \log n$	n^2	n^3	$n^{>3}$	EXP	Time w/o Time-outs	# Time-outs
Loopus	383	131	0	151	0	81	16	4	0	437s	5
KoAT	321	121	0	142	0	54	0	3	0	682s	282
PUBS	279	116	5	129	5	15	4	0	6	1000s	58
Rank	84	56	0	19	0	8	1	0	0	173s	6

Fig. 3. Analysis results for the benchmark from [12].

Non-linear Local Ranking Functions. In our experiments we only found few loops that do not have a linear local ranking function. However, these loops almost always involve the iterated division or multiplication of a loop counter by a constant such as in the transition relation $\rho \equiv x > 1 \wedge x' = x/2$. For such loops we can introduce the logarithm of x as a norm, i.e., $y = \log x$, and then try to establish $y > 0$ from the condition $x > 1$ and derive the transition relation by $y' \leq y - 1$ from the update $x' = x/2$.

Data Structures. Previous approaches [18, 24] have described how to abstract programs with data structures to integer programs by making use of appropriate norms such as the length of a list or the number of elements in a tree. In our implementation we follow these approaches using a light-weight abstraction based on optimistic aliasing assumptions.

8 Experiments

We implemented the discussed approach as an intraprocedural analysis (we use function inlining) based on the LLVM [23] compiler framework. Our tool **Loopus** computes loop bounds (depending on a command-line parameter, see [1]) either in terms of (1) the inputs to the SCC to which the loop belongs, or (2) the function inputs (this is implemented by enclosing the function body in a dummy loop as described in Section 6). At the same time **Loopus** also computes the asymptotic complexity of the considered SCC. We use the Z3 SMT solver [14] for removing unsatisfiable paths during the analysis. Given a loop condition of form $a \neq 0$ **Loopus** heuristically decides to either assume $a > 0$ or $a < 0$ as loop-invariant; this assumption is reported to the user. Similarly, **Loopus** assumes $x > 0$ when an update of a loop counter of the form $x = x * 2$ or $x = x/2$ is detected. The task of validating these assumptions is orthogonal to our approach and can be performed by standard tools for invariant generation. **Loopus** and more details on our experimental evaluation are available at [1].

8.1 Comparison to Tools from the Literature

We compare **Loopus** against the tools KoAT [12], PUBS [5, 6] and Rank [7]. For the comparison we use the benchmark [2], which consists of small example programs from the bound analysis literature and the benchmark suite which was used to evaluate T2 [11]. Since **Loopus** expects C code but KoAT expects a transition system as input, we needed to obtain C programs for comparison, see [25] for details. Figure 3 states the results for the different tools (the results

	Analyzed	Outer Dep.	Inner Dep.	Paths > 1	Non-Trivial
Loops	4210	255	305	1276	1475
Bounded	3205 [3060]	120 [112]	148 [129]	744 [695]	831 [812]
/ Overall	76% [73%]	47% [44%]	49% [42%]	58% [54%]	56% [55%]
SCCs	2833	181	193	902	937
Bounded	2289	70	95	542	564
/ Overall	81%	39%	50%	60%	60%

Fig. 4. Loop and SCC Statistic of our current implementation for the cBench Benchmark, the results obtained with the implementation of [28] are given in square brackets

for KoAT, PUBS and Rank were taken from [12]). Columns 2 to 9 state the number of programs that were found to have the given complexity by the respective tool. The table shows that **Loopus** can compute bounds for more loops than the other tools. Moreover one can see significant differences in analysis time, which are due to time-outs (the table shows the analysis time without time-outs and the number of time-outs separately; the time-out is set to 60s for all tools). The detailed comparison available at [1] shows that there are 84 loops for which **Loopus** computes an asymptotically more precise bound than any of the 3 other tools, compared to 66 loops for which one of the 3 other tools computed an asymptotically more precise bound than **Loopus**.

8.2 Evaluation on Real-World Code

We evaluated **Loopus** on the program and compiler optimization benchmark *Collective Benchmark* [3](cBench), which contains a total of 1027 different C files (after removing code duplicates) with 211.892 lines of code.

Data Structures. For expressing bounds of loops iterating over arrays or recursive data structures **Loopus** introduces shadow variables representing appropriate norms such as the length of a list or the size of an array. **Loopus** makes the following optimistic assumptions which are reported to the user: Pointers do not alias, a recursive data structure is acyclic if a loop iterates over it, a loop iterating over an array of characters is assumed to be terminating if an inequality check on the string termination character ‘\0’ is found.¹ We made these assumptions in order to find interesting examples, a manual check on a sample of around 100 loops in the benchmark found the assumptions to be valid with respect to termination. The task of validating an assumption is orthogonal to our approach and can be performed by standard tools for shape analysis.

Results. In Figure 4, we give our results on different loop classes. We recall that our bound analysis is based on an *explicitly* computed termination proof. We do not list the results of our termination analysis separately, because a bound was computed for 98% of all loops for which termination was proven. Details on the reasons for failure of our termination analysis and bound analysis, which occurred during the experiments, are given in [25]. In column *Analyzed* we state

¹ This assumption is necessary since the type system of C does not distinguish between an array of characters and a string

the results over all loops in the benchmark. We summarize the results over all loop categories except *Analyzed* in the column *Non-Trivial*.

Challenging Loop Classes. The loop-class ‘*Outer Dependent*’ captures all outer loops whose termination behavior is affected by the executions of an inner loop. (E.g., in Figure 1 termination of loop l_1 depends on loop l_3 , while in Figure 2 termination of loop l_1 does *not* depend on loop l_2 .) We define an inner loop to be in the set of ‘*Inner Dependent*’ loops if it has a loop counter that is *not* reset before entering the loop. (E.g., in Figure 1 loop counter i of loop l_3 is always reset to $n - 1$ before entering the loop, while loop counter b of loop l_2 is never reset.) The loop-class ‘*Paths > 1*’ contains all loops which have more than 1 path left after program slicing (see [25]). The categories for the SCCs are the same as for the loops: we define an SCC to be in a certain category if it contains at least one loop which is in that category. Success ratios of around 50% in the difficult categories demonstrate that our method is able to handle non-trivial termination and complexity behavior of real world programs.

Amortization. For 107 loops out of the 305 loops in the class ‘*Inner Dependent*’, the bound that our tool computed was amortized in the sense that it is asymptotically smaller than one would expect from the loop-nesting depth of the loop. In 12 cases the amortization was caused by incrementing a counter of the inner loop in the outer loop as in Figures 1 and 2. The 12 loops are available at [1]. *For these loops a precise bound cannot be computed by any other tool (as discussed in the beginning of Section 2).*

Performance. The results were obtained on a Linux machine with a 3.2 Ghz dual core processor and 8 GB Ram. 92 loops of the 4302 loops in our benchmark are located in 44 SCCs with an irreducible control flow. We thus analyzed 4210 loops. The total runtime of our tool on the benchmark (more than 200.000 LoC) did not exceed 20 minutes. The time-out limit of maximal 420 seconds computation time per SCC was not reached. There were only 27 out of 2833 SCCs (174 out of 4210 loops) on which the analysis spent more than 10 seconds.

Experimental Comparison. For the purpose of a realistic comparison, we ran the tool of [28] on the same machine with an equal time out limit of 420 second. The results are given in square brackets in Figure 4. Note the significant increase in the number of loops bounded in each of the challenging categories. The execution of the tool [28] took an order of magnitude longer (nearly 13 hours) and we got 78 time-outs. The main reason for the drastic performance increase is our new reasoning on inner loops: The approach of [28] handles inner loops by inserting the transitive hull of an inner loop on a given path of the outer loop. This can blow up the number of paths exponentially. We avoid this exponential blow-up thanks to CA: CA allows us to analyze inner and outer loops at the same time and thus eliminates the need for transitive hull computation.

Acknowledgements. We thank Fabian Souczek and Thomas Pani for help with the experiments.

References

1. <http://forsyte.at/static/people/sinn/loopus/>.
2. <http://aprove.informatik.rwth-aachen.de/eval/IntegerComplexity>.
3. <http://ctuning.org/wiki/index.php/CTools:CBench>.
4. Alfred V. Aho, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, 1986.
5. Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. Cost analysis of object-oriented bytecode programs. *Theor. Comput. Sci.*, 413(1):142–159, 2012.
6. Elvira Albert, Samir Genaim, and Abu Naser Masud. On the inference of resource usage upper and lower bounds. *ACM Trans. Comput. Log.*, 14(3):22, 2013.
7. Christophe Alias, Alain Darte, Paul Feautrier, and Laure Gonnord. Multi-dimensional rankings, program termination, and complexity bounds of flowchart programs. In *SAS*, pages 117–133, 2010.
8. Diego Esteban Alonso-Blas and Samir Genaim. On the limits of the classical approach to cost analysis. In *SAS*, pages 405–421, 2012.
9. Ahmed Bouajjani and Richard Mayr. Model checking lossy vector addition systems. In *STACS*, pages 323–333, 1999.
10. Aaron R. Bradley, Zohar Manna, and Henny B. Sipma. Linear ranking with reachability. In *CAV*, pages 491–504, 2005.
11. Marc Brockschmidt, Byron Cook, and Carsten Fuhs. Better termination proving through cooperation. In *CAV*, pages 413–429, 2013.
12. Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Juergen Giesl. Alternating runtime and size complexity analysis of integer programs. In *TACAS*, page to appear, 2014.
13. Byron Cook, Abigail See, and Florian Zuleger. Ramsey vs. lexicographic termination proving. In *TACAS*, pages 47–61, 2013.
14. Leonardo Mendonça de Moura and Nikolaj Bjørner. Z3: An efficient smt solver. In *TACAS*, pages 337–340, 2008.
15. Bhargav S. Gulavani and Sumit Gulwani. A numerical abstract domain based on expression abstraction and max operator with application in timing analysis. In *CAV*, pages 370–384, 2008.
16. Sumit Gulwani, Sagar Jain, and Eric Koskinen. Control-flow refinement and progress invariants for bound analysis. In *PLDI*, pages 375–385, 2009.
17. Sumit Gulwani and Sudeep Juvekar. Bound analysis using backward symbolic execution. Technical Report MSR-TR-2004-95, Microsoft Research, 2009.
18. Sumit Gulwani, Tal Lev-Ami, and Mooly Sagiv. A combination framework for tracking partition sizes. In *POPL*, pages 239–251, 2009.
19. Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. Speed: precise and efficient static estimation of program computational complexity. In *POPL*, pages 127–139, 2009.
20. Sumit Gulwani and Florian Zuleger. The reachability-bound problem. In *PLDI*, pages 292–304, 2010.
21. Jan Hoffmann, Klaus Aehlig, and Martin Hofmann. Multivariate amortized resource analysis. *ACM Trans. Program. Lang. Syst.*, 34(3):14, 2012.
22. Martin Hofmann and Steffen Jost. Static prediction of heap space usage for first-order functional programs. In *POPL*, pages 185–197, 2003.
23. Chris Lattner and Vikram S. Adve. Llv: A compilation framework for lifelong program analysis & transformation. In *CGO*, pages 75–88, 2004.

24. Stephen Magill, Ming-Hsien Tsai, Peter Lee, and Yih-Kuen Tsay. Automatic numeric abstractions for heap-manipulating programs. In *POPL*, pages 211–222, 2010.
25. Moritz Sinn, Florian Zuleger, and Helmut Veith. A simple and scalable static analysis for bound analysis and amortized complexity analysis. *CoRR*, abs/1401.5842, 2014.
26. James Stanier and Des Watson. A study of irreducibility in c programs. *Softw., Pract. Exper.*, 42(1):117–130, 2012.
27. R. E. Tarjan. Amortized computational complexity. *SIAM Journal on Algebraic Discrete Methods*, 6(2):306–318, April 1985.
28. Florian Zuleger, Sumit Gulwani, Moritz Sinn, and Helmut Veith. Bound analysis of imperative programs with the size-change abstraction. In *SAS*, pages 280–297, 2011.