TPLP **19** (5–6): 722–739, 2019. © Cambridge University Press 2019
doi:10.1017/S1471068419000152

Resource Analysis driven by (Conditional) Termination Proofs[*]

ELVIRA ALBERT
DSIC, Complutense University of Madrid (UCM), E-28040 Madrid, Spain

MIQUEL BOFILL
IMAE, University of Girona (UdG), E-17003 Girona, Spain

CRISTINA BORRALLERAS
University of Vic - Central University of Catalonia (UVic-UCC), 08500 Vic (Barcelona), Spain

ENRIQUE MARTIN-MARTIN and ALBERT RUBIO
DSIC, Complutense University of Madrid (UCM), E-28040 Madrid, Spain

submitted 24 July 2019; accepted 31 July 2019

Abstract

When programs feature a complex control flow, existing techniques for resource analysis produce *cost relation systems* (CRS) whose cost functions retain the complex flow of the program and, consequently, might not be solvable into closed-form *upper bounds*. This paper presents a novel approach to resource analysis that is driven by the result of a termination analysis. The fundamental idea is that the termination proof encapsulates the flows of the program which are relevant for the cost computation so that, by driving the generation of the CRS using the termination proof, we produce a *linearly-bounded* CRS (LB-CRS). A LB-CRS is composed of cost functions that are guaranteed to be *locally* bounded by linear ranking functions and thus greatly simplify the process of CRS solving. We have built a new resource analysis tool, named MaxCore, that is guided by the VeryMax termination analyzer and uses CoFloCo and PUBS as CRS solvers. Our experimental results on the set of benchmarks from the Complexity and Termination Competition 2019 for C Integer programs show that MaxCore outperforms all other resource analysis tools.

KEYWORDS: resource analysis, termination analysis, cost relation systems, upper bounds

1 Motivation and Related Work

The classical approach to resource analysis by Wegbreit consists of two steps: (1) the generation of a cost relation system (CRS) from the program that defines by means of recursive cost functions its resource consumption, (2) solving the CRS into a closed-form expression that bounds its cost. This approach is generic w.r.t. the *cost model* that defines the type of resource that is being measured, e.g., it has been applied to estimate number of execution steps, memory, energy (Liqat et al. 2015; Grech et al. 2015),

[*] This work was funded partially by the Spanish MICINN/FEDER, UE projects RTI2018-094403-B-C31, RTI2018-094403-B-C33 and RTI2018-095609-B-I00, the MINECO project TIN2015-69175-C4-2-R, the MINECO/FEDER, UE projects TIN2015-69175-C4-3-R and TIN2015-66293-R, and by the CM project S2018/TCS-4314.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

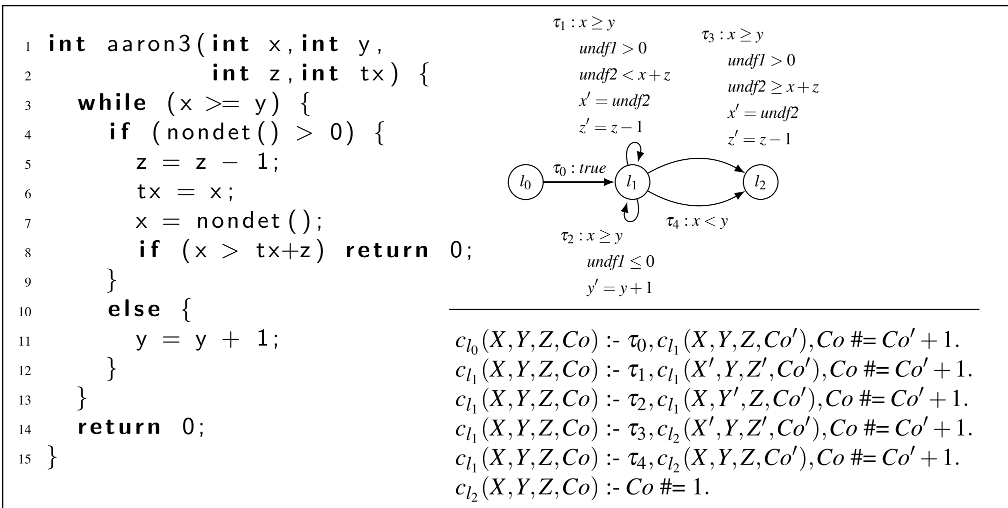

Fig. 1. Motivating example (left). Direct TS (upper-right). Non-solvable CRS (bottom-right)

user-defined cost models (Navas et al. 2007). W.l.o.g., we use the cost model adopted in the Complexity and Termination competition `http://termination-portal.org/wiki/Termination_Competition_2019` (abbreviated as TermComp) which simply estimates the asymptotic complexity order (e.g., by accumulating constant values in cost functions). This classical resource analysis approach has been applied to a wide variety of declarative and imperative programming languages: earlier work applied it to functional (Wegbreit 1975) and logic languages (Debray and Lin 1993; Debray et al. 1994), later work to imperative languages such as Java and Java bytecode (Albert et al. 2007), concurrent programs (Garcia et al. 2015; Albert et al. 2018), LLVM (Grech et al. 2015; Liqat et al. 2015), among others. In most cases, the program written in any imperative/declarative, source/bytecode language is first transformed into a simpler intermediate representation (IR) that works only on Integer data, which is the starting point of our work. For this, a size abstraction is applied on the program to transform all data into their sizes (e.g., by using the well-known term-size/term-depth abstractions, or the path-length norm by Spoto et al. (2010) for heap-allocated data structures, etc). This step is followed by a size analysis (Cousot and Halbwachs 1978) that infers size relations among the program variables. Therefore, step (1) above can be conceptually split into two parts: (1a) the transformation of the program into an Integer IR using a language-specific size abstraction, and (1b) the generation of a CRS from the IR using the gathered size relations. The IR we adopt in the paper are Integer Transition Systems (abbreviated as TS) which are an official input language for TermComp. For the sake of generality, our work assumes that the input program (written in any language) has been already transformed into a TS and a language-specific size analysis has been applied, and focuses on (1b).

An important limitation of this classical approach to resource analysis is that CRS inherit the structure of the input program from which they are generated or, equivalently, of its IR. When the program features a complex control flow, this might lead to CRS that cannot be solved in step (2). Our motivating example is aaron3, borrowed from the set of benchmarks used in TermComp'19. Fig. 1 shows the C implementation for this program (left) and the TS directly obtained from it (up-right). The TS will be explained

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

in further detail later, by now, we only want to emphasize that it comprises the different paths in the execution flow and that its arrows are labeled with the *constraints* that are gathered along each path (*undef* variables are fresh variables used to represent the unknown result of function nondet). For instance, the upper arrow from l_1 represents the iteration of the loop that executes the then branch of the first if statement and it accumulates the constraints gathered from those instructions in τ_1 (the guard contains $undef < x+z$ instead of $undef \leq x+z$ because the condition in line 8 is evaluated after the 3 assignments, so the variable z refers to the original value minus one). Note that, regardless of the programming language used to implement aaron3, a similar TS would be produced. The CRS, written as a $CLP(\mathbb{Z})$ program, that has been obtained by a standard cost analysis from this TS is shown in the figure (down-right). We can observe that the structure of the cost functions (i.e., the predicates) corresponds directly to the flow in the original TS, with one cost function per location in the TS and the constraints guarding the cost equations (i.e., the clauses). The cost accumulated by each function is calculated in the last parameter of the predicates. While one could execute this $CLP(\mathbb{Z})$ program for concrete input values, our purpose is to obtain an upper bound for Co that is sound for any possible execution, i.e., solve the CRS into a closed-form upper bound. However, this CRS is not solvable by existing systems (e.g., CoFloCo, PUBS) due to two reasons: (1) they rely on linear ranking functions to bound the number of iterations that loops (i.e., the recursive predicates) perform, while c_{l_1} requires the lexicographic ranking function $\langle z, x-y\rangle$, and (2) they cannot find the phases in the execution flow for the different increase/decrease of variables. Concretely, the loop presents two phases. In the first phase (when $z > 0$), at each iteration either z decreases and x takes an arbitrary value smaller than or equal to $x+z$, or y increases by one. In the worst case x increases, and after every increment of x there may be $x-y$ increments of y followed by a new update of x. However, these potential increments in x can only happen z times, and then the loop enters the second phase where $z \leq 0$. In this other phase, x decreases or y increases, therefore reducing the difference $x-y$ at each iteration.

The problem of the non-solvability of the CRS obtained from complex flow programs was observed in (Flores-Montoya and Hähnle 2014), which proposes to partition all possible executions of the program into a finite set of execution patterns, named *chains*, so that more precise constraints can be inferred for each of the chains, that results in simpler ranking functions and more upper bounds being found. However, the computation of the chains is not guided by semantic criteria, rather a full partitioning is carried out, that might lead to inaccuracy as our example shows. Indeed, CoFloCo (Flores-Montoya 2017) —implementing the chains— is not able to infer an upper bound: it detects 5 different chains for the loop in aaron3 but can only infer a bound for 2. One of those detected chains is the loop formed by the transitions with constraints τ_1 and then τ_2. This chain is detected for the precondition $x \geq y$, which is not strong enough to obtain a linear ranking function. Since the chain detection was not able to extract the finer phases above depending on the value of z, CoFloCo cannot find an upper bound. Further related work based on finding phases includes (Gulwani et al. 2009; Sharma et al. 2011). The former is based on size-change constraints that are less expressive than the general linear constraints used by (Flores-Montoya and Hähnle 2014) and us. The latter computes rather sophisticated phases but its main target is on proving safety properties, and it is unclear how effectively it would perform for cost.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

The main idea of our approach that differs from such previous work is to use, as semantic criterion to guide the CRS generation, the termination proofs inferred by a powerful termination analyzer as they comprise the actual phases needed to compute resource bounds. This idea is materialized in our analysis by transforming the TS into a *hierarchically loop-nested* TS that witnesses all components in the termination proof (e.g., the one for aaron3 appears later in Fig. 7). The benefit of hierarchically loop-nested TS is that they allow us to produce CRS that are *Linearly-Bounded* (LB), as shown later in Fig. 8. Cost functions in the LB-CRS are guaranteed to have linear ranking functions. Thus, the solving process is greatly simplified, e.g., we indeed find an upper bound of $O(n^3)$ for aaron3, where n is the maximum of the parameters x, y, z, and tx. Interestingly, we rely on a *conditional* termination analysis (Borralleras et al. 2017) that, when it cannot prove termination unconditionally, tries to infer preconditions under which termination is guaranteed. Conditional termination proofs allow us to generalize our results to conditional upper bounds. Finally, another work related to ours is (Sinn et al. 2014). The similarity with our approach is that both can use lexicographic ranking functions to bound the cost but our technique is more general as it allows more powerful termination arguments, besides not being limited to difference constraints as (Sinn et al. 2014). According to our experimental results, the precision of our system significantly outperforms their system Loopus.

Summary of contributions. Briefly, the main contributions of our work are: (i) We define the concept of *lexicographic phase-level termination proof*, *Proof*, to store information on the phases which have been considered during the conditional termination proof and unfold the TS accordingly. (ii) We present a transformation which takes the unfolded TS together with the *Proof*s of its phases and produces a hierarchically loop-nested TS^h which explicitly represents the different components of the termination proof. The CRS generated from TS^h is *locally* LB, although still needs to be globally bounded in the solving step. (iii) We propose extensions of the basic framework: to embed the ranking functions into the CRS; and to embed the preconditions inferred by the termination analysis so that conditional upper bounds can be generated. (iv) We implement MaxCore (standing for Max-SMT based termination analyzer + COst Recurrence Equation solver), that makes use of VeryMax (Borralleras et al. 2017) to generate the conditional termination proofs from which our implementation produces CRS, and uses both CoFloCo and PUBS as backend solvers. (v) We prove experimentally on the benchmarks from TermComp'19 for C Integer programs that MaxCore outperforms all existing resource analyzers in number of: problems solved, unique problems solved, more accurate solutions, and overall score.

2 Lexicographic Phase-Level Termination Proofs and Unfolded TS

In this section we present an overview of (Borralleras et al. 2017) and propose how to adapt the results of this analysis to guide the generation of the CRS. Essentially, Borralleras et al. (2017) describe a template-based method for proving conditional termination, and then show how to use conditional proofs to advance towards an (unconditional) termination proof. The key idea is that conditional termination proofs show termination for a subset of states which can be excluded in the rest of the termination analysis, i.e., the rest of the proof can concentrate on the complementary states. This way, the method allows generating not only a termination proof, but also a characterization of the

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

execution phases in a program. An execution *phase* characterizes a subset of states in which termination follows from a different conditional invariant.

We assume programs are given as *(Linear) Integer Transition Systems* (TSs). A TS is a control-flow graph with transitions τ of the form (l_s, ρ, l_t), where l_s and l_t are locations and ρ is a conjunction of *linear* inequalities describing the transition relation (by abuse of notation we sometimes use τ to express only its associated ρ). When the input program contains non-linear instructions that are not handled by our analysis, they are translated into *undefined* variables within the inequalities to express the loss of information. For instance, if the condition in the `if` statement in line 4 was `x*y`, this is transformed into a call to function `nondet` that has led to the introduction of the undefined variable *undef1* (representing the unknown value `x*y`) in the constraints τ_1, τ_2 and τ_3. The formula ρ can contain primed variables v', which represent the value of a variable v after the transition (equalities $v' = v$ are omitted). A *program component* C of a program P is the set of transitions of a *strongly connected component* (SCC) of the CFG of P. For example, in the TS of Fig. 1 there are two trivial (i.e., single node) SCCs; the transitions τ_1 and τ_2 form a non-trivial program component.

Termination of a program is proven component-by-component, and termination of a program component is proven iteratively by removing transitions that can only be finitely executed. A *ranking function* for a component C and a transition $\tau = (l_s, \rho, l_t) \in C$ is a function $R : \mathbb{Z}^n \to \mathbb{Z}$ such that it is bounded from below $\rho \models R \geq 0$, it strictly decreases $\rho \models R > R'$ and, for every $(\hat{l}_s, \hat{\rho}, \hat{l}_t) \in C$, it is non-increasing $\hat{\rho} \models R \geq R'$, where R' is the version of R using primed variables. The key property of ranking functions is that if one transition admits one, then it cannot be executed infinitely. In our setting, proving termination of a component C is based on finding a linear ranking function together with some supporting invariants that ensure the conditions for being a ranking function. Invariants are described by a function $Q : \mathscr{L}(C) \to \mathscr{F}(V)$, where $\mathscr{L}(C)$ is the set of locations of C and $\mathscr{F}(V)$ are conjunctions of linear inequalities over the variables V of the program. Then, strictly decreasing transitions w.r.t. this ranking function can be removed and the process is iterated over the remaining SCCs. However, although all supporting invariants are inductive in (Borralleras et al. 2017), they are not necessarily initiated in all computations. In this case, those invariants are called *conditional invariants* as they yield a precondition for termination, i.e., they prove termination for a subset of initial states. Therefore, the rest of the proof can be restricted to the remaining states. This makes the proof method more powerful and, as a by-product, loops with different execution phases can be handled naturally.

In this paper, we propose to store information on the phases which have been considered during the termination proof, together with the lexicographic termination proof of each phase. This information will capture all the possible execution flows in the execution of the program and will be used for guiding the generation of the CRS.

Definition 1 (*lexicographic phase-level termination proof, Proof*)
Let C be a component and R a ranking function for C with a supporting conditional invariant Q. Then C can be split into $C^{>} \uplus C^{subSCC} \uplus C^{noSCC}$ where:

- $C^{>}$ contains the strictly-decreasing transitions in C w.r.t. R assuming Q,
- C^{subSCC} contains the transitions that belong to an *SCC* in $C \setminus C^{>}$, and
- $C^{noSCC} = C \setminus (C^{>} \uplus C^{subSCC})$ contains the transitions that after removing the strictly decreasing transitions do not belong to any *SCC*.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

We denote by C^R the set of transitions $C^> \uplus C^{noSCC}$. A lexicographic phase-level termination proof for (a phase of) C can be represented by a tree-like structure $\mathit{Proof}(C) = \langle R, Q, C^R, \langle \mathit{Proof}\ (C_1), \ldots, \mathit{Proof}(C_k)\rangle\rangle$ where $C_1, \ldots, C_k$ are the new SCCs in C^{subSCC}.

The information kept for the termination proof of (a phase of) a component in the definition above is (i) the ranking function used; (ii) its supporting conditional invariants; (iii) the set of transitions removed, either because they strictly decrease wrt. the ranking function or they do not belong to any SCC after removing the strictly decreasing ones; and, recursively, (iv) the information corresponding to the termination proof of the remaining SCCs after transition removal.

Example 1

Let us consider the non-trivially terminating component $C = \{\tau_1, \tau_2\}$ of Fig. 1, where $\tau_1 = (l_1, \rho_1, l_1)$, $\tau_2 = (l_1, \rho_2, l_1)$, $\rho_1 = x \geq y \wedge \mathit{undf1} > 0 \wedge \mathit{undf2} < x + z \wedge x' = \mathit{undf2} \wedge z' = z - 1$ and $\rho_2 = x \geq y \wedge \mathit{undf1} \leq 0 \wedge y' = y + 1$. In this case a possible ranking function is $x - y$, with supporting conditional invariant $z \leq 0$. In particular, we have $x - y \geq 0$ both in τ_1 and τ_2, and $x - y$ strictly decreases in τ_2, as well as in τ_1 assuming $z \leq 0$. Therefore we have $C^> = \{\tau_1, \tau_2\}$, $C^{subSCC} = \emptyset$ and $C^{noSCC} = \emptyset$, giving us $C^R = \{\tau_1, \tau_2\}$ and $\mathit{Proof}(C) = \langle x - y, Q, \{\tau_1, \tau_2\}, \langle\rangle\rangle$ where $Q(l_1) = z \leq 0$. This is a conditional termination proof for C with supporting conditional invariant $z \leq 0$. To complete the termination proof, we have to analyze the rest of the states where $z \geq 1$. For this, we will assume an entry transition of the form $(l_0, z \geq 1, l_1)$ instead of the original $(l_0, \mathit{true}, l_1)$, and a strengthened version of C defined by $C' = \{\tau_1', \tau_2'\}$, with $\tau_1' = (l_1, \rho_1', l_1)$, $\tau_2' = (l_1, \rho_2', l_1)$, $\rho_1' = \rho_1 \wedge z \geq 1$ and $\rho_2' = \rho_2 \wedge z \geq 1$. In this new phase, $z - 1$ is a ranking function for C' and τ_1' without the need of any additional supporting invariant, since $z - 1 \geq 0$ in τ_1', $z - 1$ strictly decreases in τ_1' and it is non-increasing in τ_2'. Therefore, we have $C'^> = \{\tau_1'\}$, $C'^{subSCC} = \{\tau_2'\}$ and $C'^{noSCC} = \emptyset$ and $\mathit{Proof}(C') = \langle z - 1, Q', \{\tau_1'\}, \mathit{Proof}(\{\tau_2'\})\rangle$, with $Q'(l_1) = \mathit{true}$. Finally, $x - y$ is a ranking function for τ_2', giving $\mathit{Proof}(\{\tau_2'\}) = \langle x - y, Q', \{\tau_2'\}, \langle\rangle\rangle$.

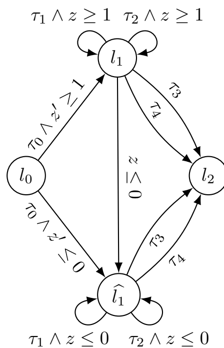

Fig. 2. Unfolded TS

Lexicographic phase-level termination proofs can be considered to use a semantically equivalent unfolded version of the TS. The unfolding goes as follows. For each transition (l_s, ρ, l_t) of a component C, on the one hand (l_s, ρ, l_t) is strengthened with the negation of the conditional invariant Q for C; more precisely, the transition is replaced by $(l_s, \rho \wedge \neg Q(l_s), l_t)$, or a set of transitions if $\neg Q(l_s)$ has disjunctions. On the other hand, a transition $(\widehat{l_s}, \rho \wedge Q(l_s), \widehat{l_t})$ is added between two fresh locations $\widehat{l_s}$ and $\widehat{l_t}$. Transitions strengthened with the negation of the conditional invariant correspond to a phase for which termination has not yet been proven, whereas transitions strengthened with the conditional invariant correspond to a phase for which termination has already been proven. Under this assumption, the remaining proof can be restricted to transitions strengthened with the negated invariant. A single transition $(l_s, Q(l_s), \widehat{l_s})$ is added to connect the two phases, i.e., to allow switching to a phase for which termination has already been proven. Finally, to preserve semantic equivalence of the unfolded transition system, the entry transitions (l_s, ρ, l_t) of C are

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

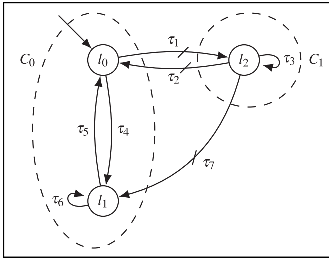

Fig. 3. First *split* without h

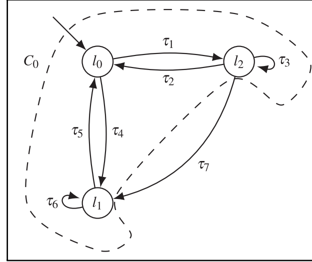

Fig. 4. Final *split*

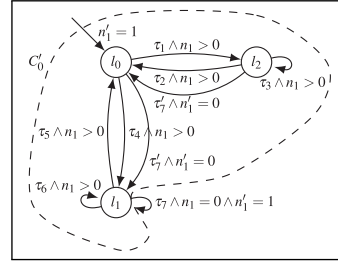

Fig. 5. Move source of τ_7

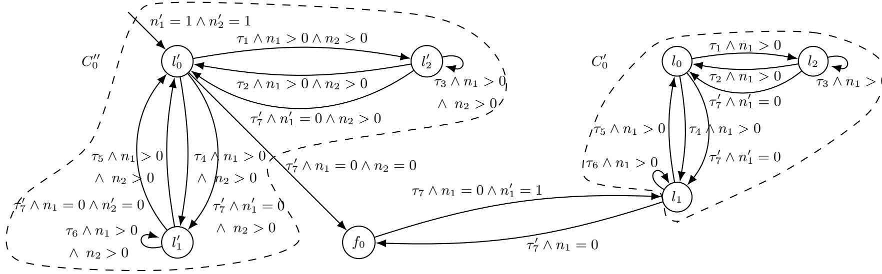

Fig. 6. Example of complex transformation

unfolded into $(l_s, \rho \wedge \neg Q(l_t)', l_t)$ and $(l_s, \rho \wedge Q(l_t)', \widehat{l_t})$, while exit transitions are unfolded into (l_s, ρ, l_t) and $(\widehat{l_s}, \rho, l_t)$. It is worth noticing that this unfolding is equivalent to the one described in (Borralleras et al. 2017) but in general leads to a simpler TS. In what follows, we assume the original TS has been unfolded as described above, and denote it TS_u. Fig. 2 shows the unfolded TS corresponding to the termination proof of the program in Fig. 1 with: τ_0 : *true*; $\tau_1 : x \geq y$, *undf1* > 0, *undf2* $< x + z$, $x' =$ *undf2*, $z' = z - 1$; $\tau_2 : x \geq y$, *undf1* ≤ 0, $y' = y + 1$; $\tau_3 : x \geq y$, *undf1* > 0, *undf2* $\geq x + z$, $x' =$ *undf2*, $z' = z - 1$; $\tau_4 : x < y$. Note that we have also strengthened the transitions looping in $\widehat{l_1}$ with its conditional invariant $z \leq 0$. This graph visualizes the loop phases in the program, which have been described in Sec. 1.

3 Linearly-Bounded Hierarchically-Loop-Nested Integer Transition Systems

The goal of this section is to soundly transform each phase of an unfolded transition system TS_u, which is given as an SCC C with its corresponding *Proof*(C) using linear ranking functions (see Sec. 2) into a TS composed of linearly-bounded hierarchically loop-nested SCCs as defined below. Let us introduce some notation. By *entryT*(C), we denote all entry transitions to C, i.e. with target location in C and source location out of C, and by *exitT*(C) we denote all exit transitions from C, i.e. with source location in C and target location out of C. A location l in C is said to be an *entry location* if there is a transition in *entryT*(C) with l as target. A location l in C is said to be an *exit location* if there is a transition in *exitT*(C) with l as source. In what follows we assume that when we are given a component C we also have *entryT*(C) and *exitT*(C).

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

Definition 2 (*linearly-bounded hierarchically-loop-nested SCC/TS*)
An SCC C is said to be *hierarchically loop-nested* if (i) it has a single entry and exit location e; (ii) there is a set of locations $l_0, \ldots, l_n$ with $e = l_0$ s.t. if l_i is connected (with one or more transitions) to another l_j then $j > i$ or $j = 0$ and (iii) for all l_i with $i \geq 0$, either l_i has no more connections than these or it is the entry location of a sub-SCC that is also a hierarchically loop-nested TS. A TS is hierarchically loop-nested if all its subSCCs also are. In addition, it is said to be *linearly bounded* if the loop (with all transitions between locations in) $l_0, \ldots, l_n$ is bounded by a linear ranking function and all sub-SCCs are linearly bounded.

Therefore, from C and *Proof*(C), we aim at generating a transformed TS that has a representation of nested loops, where every loop has a single location that is both the entry and the exit location. W.l.o.g., we assume that the component C has a single entry location (if there are several we simply clone C for every entry location, by renaming locations). Every cloned component C_i for the entry location i will have as entries those of C that have i as target. Regarding exit transitions, it is easy to transform any component C with exit locations that are different from the entry location into one TS that has only exits from the entry location. Furthermore, this transformation can be done introducing only transitions from l_i to l_j if a transition from l_i to l_j already exists. This transformation, that we call in what follows $exitToentry$ can be, in general, done to change the source location of a set of transitions from one location to another (and in particular from one exit to an entry). This more general construction, that we call $moveSourceLocation$, takes a component C (including entries), a set of transitions T with the same source location l and a location e, and introduces a fresh variable to encode the move from l to the new location e when the transitions in T can be applied, and then changes T to have e as source. The transition system in Fig. 5 is the result of applying $moveSourceLocation$ to C_0, $T = \{\tau_7\}$ and $e = l_1$ in Fig. 4.

Now, we describe how to transform any SCC C and *Proof*(C) into a LB hierarchically-loop-nested one. As it is a general transformation procedure for any possible component C, its formal description is quite involved. However, in practice, in most cases the transformation is not that complex, as we show later in Ex. 2 for our running example. We will also provide some examples of the application of the more involved steps. We first define the following auxiliary function *split* on C which roughly uses *Proof*(C) to extract a set of sub-SCCs (maybe including a single location without transitions) which represent the inner loops and a subset of C^R (the removed transitions in the first step of *Proof*(C)) that are the transitions performed to go from one inner loop to another and that form the outer loop. It is important to note that if we remove any of these selected transitions, the only SCCs of the remaining graph are the ones we have extracted. The splitting has a DAG-like shape of components whose leaves return to the unique initial component C_0, and C_0 has the same target location for all returning transitions.

Definition 3
Let C be a terminating SCC with *Proof*(C). Procedure *split*(C) extracts subcomponents $C_0, \ldots, C_n$ and disjoint non-empty sets of transitions $T_0, \ldots, T_n$ with $n \geq 0$, such that

a. the transitions in $C_0, \ldots, C_n$ union $T_0, \ldots, T_n$ coincide with C; when C_i has no transition, we say it includes the single source location of all transitions in T_i,

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

b. every C_i is included in $C'_{j_1} \cup \ldots \cup C'_{j_m} \cup C^R$ for some $m \geq 0$, where $Proof(C'_{j_k})$ is a subproof of $Proof(C)$ for $k \in \{1 \ldots m\}$,
c. every T_i is included in C^R,
d. C_0 includes the entry location,
e. every C_i is an SCC and no location is shared between C_i components,
f. the source location of all transitions in T_i for $i \in \{0 \ldots n\}$ belongs to C_i,
g. the target location of every transition in T_i belongs to some C_j with $j > i$ or to C_0,
h. all transitions in $T_0 \cup \ldots \cup T_n$ having target location in C_0 have the same target location.

As a simple example, the *split* of the phase where $z \geq 1$ in our running example (see Fig. 1) has one sub-SCC $C_0 = \{\tau_2 \wedge z \geq 1\}$ and $T_0 = \{\tau_1 \wedge z \geq 1\}$. On the other hand, the *split* of the phase where $z \leq 0$ has C_0 as the trivial SCC containing $\widehat{l_1}$ and $T_0 = \{\tau_1 \wedge z \leq 0, \tau_2 \wedge z \leq 0\}$.

Given a terminating SCC C with $Proof(C)$, the result of $split(C)$ can always be built. As a possible way to obtain $split(C)$, we can make a first selection of $C_0, \ldots, C_n$ and $T_0, \ldots, T_n$ as follows: (i) we take a set of transitions in C^R having the same source and target location and remove them from the SCC; (ii) then we recompute the (maximal) SCCs of the remaining graph, obtaining $C_0, \ldots, C_n$; (iii) transitions that are not in any of the obtained subSCCs (included those initially removed) must be in T as they could be removed in the termination proof, and are the selected set of transitions that belong to the corresponding T_i depending on where is the source location. Fig. 3 shows a first split if we start removing τ_1, as we obtain C_0 and C_1 and $T_0 = \{\tau_1\}$ and $T_1 = \{\tau_2, \tau_7\}$. After this, it is easy to see that we have conditions a–g. However it may happen that condition h does not hold, as it is the case in the example since τ_2 and τ_7 have different target locations in C_0. Then, as shown in Fig. 4 we can join some components until the condition holds again. In this case we join C_0 and C_1 into a single component C_0 and T_0 contains only τ_7.

In what follows, if $C_0, \ldots, C_n$ and $T_0, \ldots, T_n$ is $split(C)$ then we define $\mathsf{split\text{-}exits}(C_i) = T_i$ and $\mathsf{split\text{-}entries}(C_i)$ to all transitions in $T_0, \ldots, T_n$ with target location in C_i. We call split-entry locations of C_i to the set of target locations of $\mathsf{split\text{-}entries}(C_i)$ and split-exit locations of C_i to the set of source locations of $\mathsf{split\text{-}exits}(C_i)$. The following recursive procedure soundly transforms a given component with a termination proof only containing linear ranking functions into a non-cycling set (forming a tree-like structure) of hierarchically connected loop-nested SCCs (with the same single entry and exit location) all of them being bounded by a linear ranking function.

Definition 4 (*transformation to linearly-bounded hierarchically loop-nested SCCs*) Let C be a terminating SCC with $Proof(C)$ and single entry and exit location e. Procedure $nestedLoopTrans(C)$ transforms C by first computing $C_0, \ldots, C_n$ and $T_0, \ldots, T_n$ with $split(C)$. If $n = 0$ and C_0 is a single location, then return C. Otherwise we perform the following steps:

1. Build all $Proof(C_i)$ from $Proof(C)$ following Def. 1 for all non-trivial SCC C_i. Note that some C_i can include more than one component in $Proof(C)$ and some transitions in C^R.
2. Clone all C_i (and T_i and $Proof(C_i)$) with $i > 0$ such that C_i has more than one split-entry location. After this, all components have a single split-entry location. Then

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

apply *moveSourceLocation* to the resulting components (including the cloned ones) and to C_0 to move all transitions in T_i (maybe cloned) with source different from the single split-entry location of C_i. After this, we have components $C'_0, \dots, C'_m$ and $T'_0, \dots, T'_m$, with all C'_i with $i \geq 0$ having a single location as both split-entry and split-exit. Let s be such location of C'_0, which may be different from e.

3. Let f_i be a fresh location if C'_i is a non-trivial SCC or the single location in C'_i otherwise.
4. If C'_0 is a trivial SCC only including e, add all entries of C to the transformation. Otherwise clone C'_0 obtaining C''_0 with a mapping μ (from old locations to fresh locations) and $Proof(C''_0)$. For every entry transition $\langle o, \rho, e\rangle$ in C add an entry transition $\langle o, \rho, \mu(e)\rangle$ to C''_0 and if $e = s$ a transition $\langle o, \rho, f_0\rangle$ to the transformation. Then, for every transition $\langle s, \rho, t\rangle$ in T'_0, add a transition $\langle \mu(s), \rho', f_0\rangle$ as exit in C''_0 and for every exit $\langle e, \rho, o\rangle$ of C add as exit in C''_0 a transition $\langle \mu(e), \rho', f_0\rangle$ if $e = s$ and $\langle \mu(e), \rho, o\rangle$ otherwise (where, in all cases, ρ' does not include any of the conjuncts with primed variables of ρ). If $e \neq s$ then apply *exitToentry* to the resulting C''_0 considering that $\mu(e)$ is the entry location. Finally, compute *nestedLoopTrans*(C''_0).
5. Replace every transition in T'_i of the form $\langle s, \rho, t\rangle$ by $\langle f_i, \rho, t\rangle$ and $\langle s, \rho', f_i\rangle$, where ρ' does not include any of the conjuncts with primed variables. Note that these transitions are new entries and exits of $C'_0, \dots, C'_m$.
6. Add all exits $\langle e, \rho, o\rangle$ of C to C'_0 as exit transitions. If $e \neq s$ then apply *exitToentry* to the resulting C'_0. Then, replace every new exit transition $\langle s, \rho, o\rangle$ of C'_0 by $\langle s, \rho', f_0\rangle$ and add $\langle f_0, \rho, o\rangle$ as exit of the transformation of C, where, again, ρ' is ρ without primed variables.
7. Compute *nestedLoopTrans*(C'_i) on the resulting C'_i for all $i \geq 0$, and add the result to the transformation of C.

Intuitively, the steps of the transformation can be understood as follows. After computing the *split*(C), step 1 builds the termination proofs associated to the chosen subcomponents and transitions. Then, in step 2 we turn the components into components with a single split-entry and split-exit location. For instance, Fig. 5 shows the result of applying this step to the split given in Fig. 4. In this case, we do not need to clone any component but, as can be seen, we apply *moveSourceLocation* to C_0, $\{\tau_7\}$ and the split-entry location l_1 (in Fig. 4) since the split-entry of C_0 is l_1 and the split-exit of C_0 is l_2. After the step the split-entry and the split-exit s of C'_0 is l_1, which is different from the entry location e which is l_0. Note also that *moveSourceLocation* has changed the entry transition to C'_0 adding $n'_1 = 1$ (which is now the new version of the entry to C). Fig. 6, shows the result of applying the transformation steps to the C'_0 and T'_0 in Fig. 5 but without applying *nestedLoopTrans* recursively. In step 3, we define the locations that are used to express the outer loop (i.e. the loop of all sub-SCCs C'_i) of the transformation. This is f_0 in Fig. 6. In step 4 we connect the loop with the entries of the component. This step is crucial as it includes an initial use of the first sub-SCC C'_0, before entering the outer loop. The reason for that is that there must be paths in the original C that run some transitions in C'_0 before running any of the transitions in T'_0, which are used as soon as we enter the main loop. In Fig. 6, we can see the resulting C''_0, which is the result of first cloning C'_0, and then, since $l'_0 = e \neq s = l'_1$, we add a transition from l'_1 to f_0 and apply *exitToentry* to move this transition to l'_0. In step 5, we connect the sub-SCCs using the locations f_i to create the outer loop (which is represented by

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

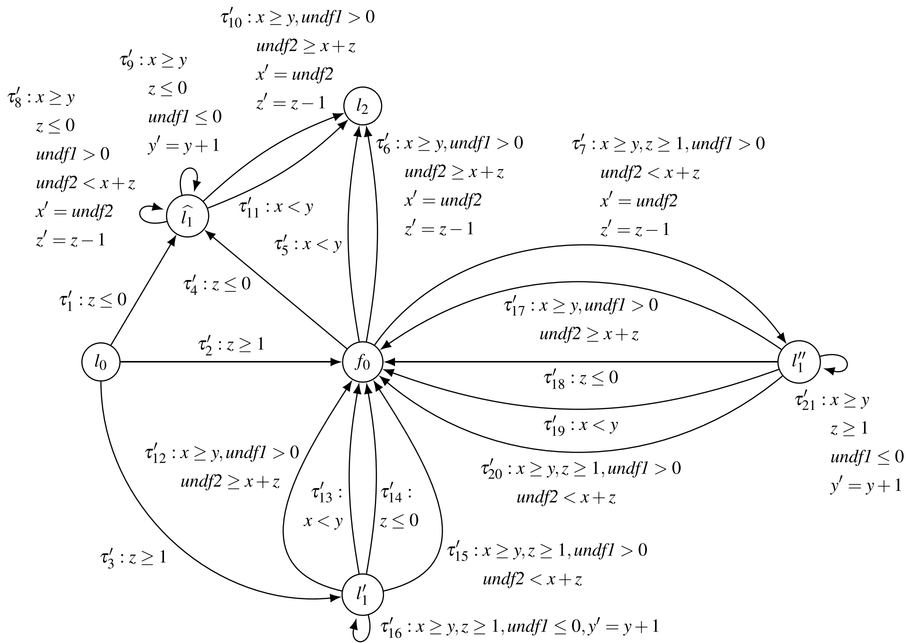

Fig. 7. Transformed CRS

the connections of C'_0 to f_0 in the right-hand-side of Fig. 6), and in 6 we introduce the needed exit transitions. For simplicity, Fig. 6 does not include the exits, but they would be leaving from $l_0 = e \neq s = l_1$, and hence *exitToentry* is applied to move them to l_1 and then connected to f_0. Finally, in the last step we apply the transformation recursively. To further illustrate how the transformation works, the following example shows the complete application of the transformation to our running example.

Example 2

Let us show how it works starting from the transformed graph in Fig. 2 and with the termination proofs for each phase given in Ex. 1. The resulting transition system is shown in Fig. 7. Its key feature is that it is ready to generate a linearly-bounded CRS in next section. There are two SCCs in Fig. 2. The SCC that cycles in the location $\widehat{l_1}$ is proved with a single ranking function where all transitions are removed, and hence our transformation does not change anything, since C_0 is the location $\widehat{l_1}$ and $T'_0 = T_0$ contains both transitions. The SCC that cycles in location l_1, needs a lexicographic combination of two ranking functions, with each component removing one transition, firstly removing $\tau_1 \wedge z \geq 1$ and secondly $\tau_2 \wedge z \geq 1$. Then *split* gives $T_0 = \{\tau_1 \wedge z \geq 1\}$ and C_0 is the SCC including location l_1 and transition $\tau_2 \wedge z \geq 1$. Therefore, first of all we compute *Proof*(C_0) according to step 1 of Def. 4. Step 2 does not change anything, since we have a single component $C'_0 = C_0$ with a single location. Step 3 delivers a fresh location f_0. Next, since $C'_0 = C_0$ is a non-trivial SCC, we clone it to C''_0 in step 4. This new SCC corresponds to location l'_1 and transition $\tau'_{16} = \tau_2 \wedge z \geq 1$ in Fig. 7. Transition τ'_3 is the entry added to C''_0. A transition $\tau'_2 = z \geq 1$ entering f_0 is also added since in this case $e = s$. Transitions τ'_{12}, τ'_{13}, τ'_{14} and τ'_{15} are the transitions added as exits. Note that τ'_3 expresses the same transition relation as the entry $\tau_0 \wedge z' \geq 1$ of l_1, since $\tau_0 = true$.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

Transition τ'_{15} is the same as the transition $\tau_1 \wedge z \geq 1$ in T_0 without conjuncts with primed variables, while τ'_{12}, τ'_{13} and τ'_{14} express the same transition relation as the original exits of l_1 with τ_3, τ_4 and $z \leq 0$, respectively, except for the conjuncts with primed variables again. In step 5, the only transition in T_0 is unfolded into the transitions τ'_7 and τ'_{20} of Fig. 7. Step 6 adds transitions τ'_{17}, τ'_{18} and τ'_{19}, corresponding to the exit transitions of l_1 but without conjuncts with primed variables, as well as transitions τ'_4, τ'_5 and τ'_6 as exits of the transformation of C. Finally, notice that l_1 has been renamed to l''_1 to avoid confusion with the original location, and τ'_{21} corresponds to $\tau_2 \wedge z \geq 1$. Note that all recursive calls to *nestedLoopTrans* trivially terminate in this example.

The transformation provided in this section is sound for resource analysis.

Theorem 1 (*soundness and linearly-bounded*)
Given a component C with *Proof*(C), then every SCC in *nestedLoopTrans*(C) is a *linearly-bounded hierarchically loop-nested transition system*, and for every path π from an initial location s to a final location t in C, there is a path π' from s to t in *nestedLoopTrans*(C) with $\#(\pi') \geq \#(\pi)$, where $\#(\pi)$ is the number of operations involved in π.

4 Linearly-Bounded Cost Relation Systems

A CRS is a set of cost equations of the form $c(\overline{x}) = 1 + c_1(\overline{x_1}) + \ldots + c_n(\overline{x_n})\{Ct\}$, where the *constraints* Ct define the applicability conditions for the equation and state size relations among $\overline{x}, \overline{x_1}, \ldots, \overline{x_n}$. As stated in Sec. 1, w.l.o.g., we always accumulate a constant unitary cost. The set of cost equations for $c(\overline{x})$ defines the (possibly non-deterministic) cost function c. Even if the input language from which the CRS are produced is deterministic, due to the loss of information implicit to static analysis (e.g., when *undef* variables appear), the associated CRS will typically be non-deterministic. CRS can be considered as constraint logic programs over integers that accumulate costs, e.g., the above equation can be written as the clause $c(\overline{X}, Co)$:- $Ct,\ c_1(\overline{X_1}, Co_1),\ \ldots,\ c_n(\overline{X_n}, Co_n),\ Co \#= 1 + Co_1 + \ldots + Co_n$ (see also the CRS in Fig. 1).

The following definition presents the generation of a CRS from a TS with possible multiple nested loops. As explained in Sec. 1, we assume that a language-specific size analysis has been already applied such that entry locations of SCC/sub-SCC in the TS are annotated with *size*$(l, \langle\overline{x}\rangle, \langle\overline{x'}\rangle)$: *the size relations between the values of the variables when reaching* $(\overline{x})$ *and leaving* $(\overline{x'})$ *a location* l. For example, for the TS in Fig. 7, the size analysis of location l''_1 will infer the relations *size*$(l''_1, \langle x, y, z\rangle, \langle x', y', z'\rangle) = \{x = x', y' \geq y, z = z'\}$ as any path $l''_1 \to^* l''_1$ using τ'_{21} will not modify x and z, and y can only increase. Size analyses for various languages can be found e.g. at (Albert et al. 2007; Serrano et al. 2013; Brockschmidt et al. 2016).

Definition 5 (*linearly-bounded CRS*)
Given a linearly-bounded hierarchically-loop-nested TS *ts*, let G be the set of locations in *ts*, $\overline{x}$ be the variables involved in *ts* and *scc*(L) the list of the SCCs in *ts* considering only the locations in the set L. Let *entry_loc*(L) denote the entry location of a set of locations L (i.e., the only location receiving transitions from outside L). The LB-CRS for *ts* is made up of the cost equations generated by *eqs*(G) that, for every SCC $L \in$ *scc*(G), proceeds as follows:

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

$$
\begin{array}{ll}
① \; c_{l_0}(x,y,z) = 1 + c_{f_0}(x,y,z) \; \{z \geq 1\} & ⑧ \; c_{f_0}(x,y,z) = 1 + c_{l_2}(x',y',z') \; \{x \geq y, u_1 > 0, \\
② \; c_{l_0}(x,y,z) = 1 + c_{l'_1}(x,y,z) \; \{z \geq 1\} & \quad u_2 \geq x+z, x' = u_2, y' = y, z' = z-1\} \\
③ \; c_{l_0}(x,y,z) = 1 + c_{\widehat{l_1}}(x,y,z) \; \{z \leq 0\} & ⑨ \; c_{l''_1}(x,y,z) = 1 + c_{l''_1}(x',y',z') \; \{x \geq y, z \geq 1, \\
④ \; c_{l_2}(x,y,z) = 1 \; \{\} & \quad u_1 \leq 0, x' = x, y' = y+1, z' = z\} \\
⑤ \; c_{f_0}(x,y,z) = 1 + c_{l''_1}(x_1,y_1,z_1) + c_{f_0}(x_2,y_2,z_2) & ⑩ \; c_{l''_1}(x,y,z) = 1 \; \{x \geq y, u_1 > 0, u_2 \geq x+z\} \\
\quad \{x_1 = x_2, y_2 \geq y_1, z_1 = z_2, x \geq y, z \geq 1, u_1 > 0, & ⑪ \; c_{l''_1}(x,y,z) = 1 \; \{z \leq 0\} \\
\quad u_2 < x+z, x_1 = u_2, y_1 = y, z_1 = z-1\} & ⑫ \; c_{l''_1}(x,y,z) = 1 \; \{x < y\} \\
⑥ \; c_{f_0}(x,y,z) = 1 + c_{\widehat{l_1}}(x,y,z) \; \{z \leq 0\} & ⑬ \; c_{l''_1}(x,y,z) = 1 \; \{x \geq y, z \geq 1, u_1 > 0, u_2 < x+z\} \\
⑦ \; c_{f_0}(x,y,z) = 1 + c_{l_2}(x,y,z) \; \{x < y\} & \quad \ldots.
\end{array}
$$

Fig. 8. Fragment of Linearly-bounded CRS obtained from transformed TS

- If $L = \{l_o\}$, i.e., the SCC contains only one location l_o, then every transition $l_o \xrightarrow{\tau} l_d$ generates one cost equation $c_{l_o}(\overline{x}) = 1 + c_{l_d}(\overline{x}') \; \{\tau\}$. If there are no transitions from l_o, a dummy equation $c_{l_o}(\overline{x}) = 1\{\}$ is generated for uniformity.
- If $|L| > 1$ and $entry_loc(L) = l_o$, as the entry location is part of the principal loop in L, we remove it to detect and, transitively, translate the remaining components. Every $D_i \in scc(L \setminus \{l_o\})$ is translated by $eqsC(D_i)$—defined below—and every cycle in the *component graph* (each SCC D_i is condensed into a single vertex d_i) starting from l_o, i.e., cycles paths of the form $l_o \xrightarrow{\tau} d_1 \rightarrow d_2 \ldots \rightarrow d_j \rightarrow l_o$, generates an equation (where $n_i = entry_loc(D_i)$):

$$
\begin{array}{rcl}
c_{l_o}(\overline{x}) & = & 1 + c_{n_1}(\overline{x'}) + c_{n_2}(\overline{x^{(2)}}) + \ldots + c_{n_j}(\overline{x^{(j)}}) + c_{l_o}(\overline{x^{(j+1)}}) \\
 & & \{\tau\} \cup size(l_{n_1}, \langle\overline{x'}\rangle, \langle\overline{x^{(2)}}\rangle) \cup \ldots \cup size(l_{n_j}, \langle\overline{x^{(j)}}\rangle, \langle\overline{x^{(j+1)}}\rangle)
\end{array}
$$

  Every outgoing transition $l_o \xrightarrow{\tau} l_k$ ($l_k \notin L$) generates an equation $c_{l_o}(\overline{x}) = 1 + c_{l_k}(\overline{x'})\{\tau\}$.

eqsC proceeds as *eqs* with one difference: in both cases ($L = \{l_o\}$ and $|L| > 1$) the outgoing transitions $l_o \xrightarrow{\tau} l_k$ generate cost equations $c_{l_0}(\overline{x}) = 1\{\tau\}$, i.e., without any call.

Let us give the intuition behind the above definition. For each location l in the TS, we produce a corresponding cost function c_l that captures its cost, and every transition produces an equation. The labels of the transitions become the constraints of the CRS. The equation for the cycle in the component graph collects the costs of all the sequential components and finishes with a recursive call to express the loop. The size analysis allows us to track the changes in the variables after every function call in order to express the cost in terms of the initial parameter values. Note that the constraints of the transitions $d_i \rightarrow d_{i+1}$ and $d_j \rightarrow l_o$ are not needed in the equation because they have been already used when generating the equations for every SCC D_i recursively. Finally, as *eqsC* is applied to components in an inner loop, the flow represented by these outgoing transitions is incorporated in the cost equation of the outer loop and no function call is needed.

Example 3
In the transformed TS from Fig. 7 there are 5 SCCs: $\{l_0\}$, $\{l_2\}$, $\{l'_1\}$, $\{\widehat{l_1}\}$ and $\{f_0, l''_1\}$. The first four are unitary, thus they generate equations directly. For example in l_0, the transitions $l_0 \rightarrow f_0$, $l_0 \rightarrow l'_1$, and $l_0 \rightarrow \widehat{l_1}$ generate equations with calls to the corresponding cost function (see equations #1–3 in Fig. 8). On the other hand, l_2 has no outgoing transition hence it creates a dummy equation for c_{l_2} (#4). Considering the non-unitary $\{f_0, l''_1\}$, the only SCC after removing the entry location f_0 is $\{l''_1\}$, thus $eqsC(\{l''_1\})$ generates 5 equations for $c_{l''_1}$: τ'_{21} creates the recursive equation (#9) and

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

τ'_{17}, τ'_{18}, τ'_{19}, and τ'_{20} generate 4 equations without any call (#10–13). As f_0 is the entry location, its outgoing transitions generate the non-recursive equations of c_{f_0} that invoke $c_{\widehat{l_1}}$ and c_{l_2} (#6–8). In the component graph l''_1 is condensed into $d_{l''_1}$ by removing the transition τ'_{22}, hence there is only one cycle $f_0 \rightarrow d_{l''_1} \rightarrow f_0$ that generates the recursive equation of c_{f_0} (#5). The size analysis relates the input and output values after invoking function $c_{l''_1}$ in this equation ($x_1 = x_2, y_2 \geq y_1, z_1 = z_2$), which allows tracking the changes from the initial values x, y, z to the ones used in the recursive call x_2, y_2, z_2 when solving the CRS. This CRS is solvable because, thanks to the transformation of the TS, all cost functions have now a linear ranking function that bounds the number of calls. The solver can hence use that information to generate the overall cost. Concretely, $c_{\widehat{l_1}}$, $c_{l'_1}$, and $c_{l''_1}$ are invoked $x - y$ times; c_{f_0} is invoked z times; and c_{l_0} and c_{l_2} are invoked only once. This contrasts with the original CRS in Fig. 1, where the lexicographic ranking function $\langle z, x - y\rangle$ cannot be used by the backend solvers to compute a loop bound.

The next corollary easily follows from the soundness of the TS transformation in Th. 1.

Corollary 1 (*soundness of linearly-bounded CRS*)
Let $N = nestedLoopTrans(C)$ be the *hierarchically loop-nested transition system* obtained from a component C with $Proof(C)$. The CRS obtained from N applying Def. 5 soundly overapproximates the cost of C (for the considered cost model), and all its functions are linearly bounded.

4.1 Embedding the Ranking Functions from Termination Proofs within CRS

CRS solving —step (2) of resource analysis— requires finding ranking functions for all recursive cost functions (i.e. cycles) to bound the number of iterations they might perform. As the termination analyzer must have already found ranking functions for all cycles, it is desirable to pass this information to the CRS solver (e.g., the resource analyzer might implement less powerful algorithms to find ranking functions). However, existing solvers are not prepared to receive this information. Our proposal does not require implementing any extension to the existing solvers. We can embed the constraints that define the ranking functions within the CRS as follows.

Definition 6 (*CRS with ranking functions*)
We assume that every location l in the TS is annotated with the linear ranking function contained in the termination proof. The main idea is to add a new parameter to the cost function c_l representing the ranking function, which is bound in the initial call and decreases in every recursive call. Therefore, the generation of the cost equations is as in Def. 5 with the following differences:

- Cost functions c_l (except those for the initial location l_0) are extended with one additional parameter r representing the ranking function: $c_l(\overline{x}, r)$.
- Cost equations invoking a cost function with ranking function rf bound the extra parameter: $\{r = rf\}$.
- Cost equations with recursive calls are extended to express that the extra parameter containing the ranking function is positive and strictly decreasing: $\{r \geq 0, r' < r\}$.

Example 4
Let us explain the above definition using our running example. In the transformed TS of Fig. 7, the termination analyzer detects that f_0, l''_1, and $l_{\widehat{l_1}}$ have ranking functions z,

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

Table 1. *Experimental results on C programs from TermComp'19 complexity competition*

	PUBS$_C$	MaxCore(P)	CoFloCo$_C$	MaxCore(C)	AProVE	CoFloCo$_C$	Loopus	MaxCore(C)
Solved	158 (32.6%)	280 (57.9%)	288 (59.5%)	311 (64.3%)	278 (57.4%)	288 (59.5%)	239 (49.4%)	311 (64.3%)
Only	6	128	21	44	2	5	3	33
Best	20	128	32	47	2	7	9	38
Score	316	546	573	611	1075	1147	946	1228
Time(s)	836.8	gen: 837.5 sol: 1574.4	1175.6	gen: 838.6 sol: 1272.3	2350.9	1175.6	9.38	gen: 838.6 sol: 1272.3

$x-y$, and $x-y$ resp. Then, the cost equations of l_0, f_0, and l_1'' will be modified as follows (we show only a fragment):

$$
\begin{array}{llll}
\text{①'} & c_{l_0}(x,y,z) & = & 1 + c_{f_0}(x,y,z,r)\ \{z \geq 1, \mathbf{r} = \mathbf{z}\} \\
& & \cdots & \\
\text{⑤'} & c_{f_0}(x,y,z,r) & = & 1 + c_{l_1''}(x_1,y_1,z_1,r_1) + c_{f_0}(x_2,y_2,z_2,r_2)\ \{x_1 = x_2, y_2 \geq y_1, z_1 = z_2, x \geq y, z \geq 1, \\
& & & u_1 > 0, u_2 < x+z, x_1 = u_2, y_1 = y, z_1 = z-1, \mathbf{r} \geq \mathbf{0}, \mathbf{r_1} = \mathbf{x_1} - \mathbf{y_1}, \mathbf{r_2} < \mathbf{r}\} \\
\text{⑥'} & c_{f_0}(x,y,z,r) & = & 1 + c_{\widehat{l_1}}(x,y,z,r')\ \{z \leq 1, \mathbf{r'} = \mathbf{x} - \mathbf{y}\} \\
& & \cdots & \\
\text{⑨'} & c_{l_1''}(x,y,z,r) & = & 1 + c_{l_1''}(x_1,y_1,z_1,r_1)\{x \geq y, z \geq 1, u_1 \leq 0, x' = x, y' = y+1, z' = z, \mathbf{r} \geq \mathbf{0}, \mathbf{r_1} < \mathbf{r}\} \\
\text{⑬'} & c_{l_1''}(x,y,z,r) & = & 1\ \{x \geq y, z \geq 1, u_1 > 0, u_2 < x+z\} \\
& & \cdots &
\end{array}
$$

The remarked constraints in the equations represent the changes. In equations $1'$ and $6'$ the ranking functions of f_0 (z) and $l_{\widehat{l_1}}$ $(x-y)$ are bound. In the recursive equations $5'$ and $9'$, the extra parameter is set to positive and decreasing. Additionally, in $5'$ the extra parameter r_1 of l_1'' is bound to $x_1 - y_1$. Finally, cost equations without invocations (as equation $13'$) are not modified. Note that c_{l_0} (equation $1'$) is not extended with any parameter because it is the entry location of the program.

4.2 Extension to Conditional Upper Bounds

The termination analysis we use (Borralleras et al. 2017) is able to infer preconditions *Pre* under which the program terminates when it cannot prove termination for all inputs. Such preconditions *Pre* may also be valid for the upper bounds. As in Sec. 4.1, the idea is to embed *Pre* into the CRS (and enable a flag cond=on) so that an unconditional solver can be used and, as preconditions are assumed, a *conditional upper bound* U can now be found. Then, when reporting the results, if cond=on, we output that U is an upper bound if preconditions *Pre* hold.

Definition 7 (*conditional CRS*)
Let G be the set of locations in the TS and l_0 be its entry location. Then the cost equations of the conditional CRS are $eqs(G)$ plus the additional equation $c^e_{l_0}(\overline{x}) = 1 + c_{l_0}(\overline{x})\{Pre\}$. In this case, the upper bound obtained for $c^e_{l_0}$ is the valid upper bound for c_{l_0} under the conditions in *Pre*.

5 Implementation and Experimental Evaluation

Our implementation, MaxCore(X) where X instantiates the CRS solver, achieves the cooperation of three advanced tools for complexity and termination analysis: VeryMax (winner of TermComp'19 for C programs) produces the termination proofs, our implementation generates from them LB-CRS that are fed: (X=C) to CoFloCo (Flores-Montoya 2017) or

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

(X=P) to PUBS (Albert et al. 2008) to produce the upper bounds. MaxCore can be used online from a web interface https://costa.fdi.ucm.es/maxcore, where the benchmarks used for our experiments can also be found. This section evaluates the effectiveness and efficiency of MaxCore by analyzing all C programs from the TermComp'19 complexity competition, that in total are 484 benchmarks containing also non-terminating programs. Experiments have been performed on an Intel Core i7-4790 at 3.6GHz x 8 and 16GB of memory, running Ubuntu 18.04. The row Solved in Table 1 shows the number of benchmarks that each system is able to bound. The row Only shows the number of benchmarks that only the corresponding system can solve and no other system can. The row Best counts the times that a system has obtained the best upper bound, and the remaining systems have larger bounds. The Score represents the points obtained by the systems following the competition rules http://cbr.uibk.ac.at/competition/rules.php. Finally, we show the overall time in seconds in the row Time(s). As in TermComp'19, systems only have 300 seconds to analyze every program. For MaxCore's instantiations we show 2 values: the time needed to generate the CRS (*gen*) and the time required to obtain a closed upper bound (*sol*). Detailed results for every system and benchmark can be found at https://costa.fdi.ucm.es/maxcore/benchmarks/.

The left part of Table 1 compares PUBS_C and MaxCore(P). PUBS and CoFloCo are CRS solvers. To avoid confusion, we use PUBS_C and $\mathsf{CoFloCo}_C$ for the systems that translate C programs to CRS using clang (http://clang.llvm.org/) and llvm2KITTeL (https://github.com/s-falke/llvm2kittel), and then use the respective CRS solver to obtain an upper bound. Unlike CoFloCo, PUBS only works with linear size relations, but it is able to obtain logarithmic upper bounds. As TermComp'19 only supports polynomial bounds, in the comparisons we have considered $O(n^k \times log^p(n))$ equal to $O(n^{k+1})$ in Best and Score. Since PUBS is a solver that does not perform any additional analysis on the CRS (unlike CoFloCo, which tries to detect *chains*), this comparison plainly shows the large improvement that can be only attributed to the proposed generation of our LB-CRS. Concretely, MaxCore(P) almost doubles the number of programs solved (280 vs. 158), and there are 128 programs that MaxCore(P) solves that PUBS_C cannot, while only 6 programs are uniquely solved by PUBS_C. Regarding time, MaxCore(P) is about three times slower than PUBS_C but the gains clearly justify the additional analysis time.

The central part of Table 1 shows the comparison between $\mathsf{CoFloCo}_C$ and MaxCore(C). Here the difference is not as large as with PUBS, but it is still important: MaxCore(C) solves 23 programs more than $\mathsf{CoFloCo}_C$ (311 vs. 288), 44 of them that $\mathsf{CoFloCo}_C$ cannot handle. However, $\mathsf{CoFloCo}_C$ solves 21 programs that MaxCore(C) cannot and obtains better upper bounds in 32 programs. Since MaxCore(C) uses VeryMax to build the termination proof that guides the generation of the LB-CRS, the system returns ∞ if VeryMax cannot find that proof. This happens in 6 of these 21 unsolved programs. Moreover, other 5 programs are not solved because their termination proof presents some features not yet integrated in the system, but are planned to be integrated soon. For the remaining 10 unsolved programs, the termination proofs found by VeryMax are too involved, thus leading to (unnecessarily) more complex CRS that CoFloCo cannot handle. Note that VeryMax can find different proofs for a program, and currently we simply use the first one. In the future, we plan to investigate on finding the best suited proofs for the solving step. MaxCore(C) has a running time 1.8 larger than $\mathsf{CoFloCo}_C$, so the difference is smaller than with PUBS, and clearly it pays off as well. Comparing MaxCore(P) and Max-

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

Core(C), the latter is about 300 seconds faster and obtains better results in all metrics. Indeed, all programs that can be bound with MaxCore(P) can be bound with equal or smaller bounds by MaxCore(C) except for one program. Therefore, we have selected MaxCore(C) to compare to the two systems participating in TermComp'19: CoFloCo$_C$, winner of TermComp'19 for complexity of ITS and C programs, and AProVE (Giesl et al. 2004), a system that implements an alternative approach which alternates between finding runtime bounds and finding size bounds (Brockschmidt et al. 2016). Additionally, we have also considered the Loopus system (Sinn et al. 2014) described in Sec. 1. The results of this comparison appear in the right part of Table 1, where it can be seen that MaxCore(C) outperforms the other systems in all metrics: besides number of problems solved, more importantly MaxCore(C) solves 33 programs that no other system can bound, generates better bounds in 38 programs, and obtains 81 more points than CoFloCo$_C$, 153 more than AProVE, and 282 more than Loopus. Moreover, it is slightly faster than AProVE, requiring 0.9 times its running time. Finally, note that Loopus is extremely fast compared to the rest of systems (it requires 0.008 times the running time of CoFloCo$_C$, the second fastest system). The reason is that Loopus relies on difference logic, a more limited domain for obtaining bounds than the linear integer arithmetic used in the rest of systems, for which very efficient algorithms exist.

6 Conclusions

This paper brings the important advances achieved in the field of termination analysis, where programs featuring complex control flow can be automatically proven to terminate, to the field of resource analysis, where there is more limited support for such kind of complex-flow programs. The success of our approach is the use of termination proofs as semantic guidance to generate linearly-bounded CRS that can be fed to an off-the-shelf CRS solver. Our experimental results on the TermComp'19 benchmarks show that our tool, MaxCore, outperforms the standalone resource analyzers CoFloCo, AProVE, and Loopus significantly both in accuracy, number of problems solved, and uniquely solved. As future work, we plan to apply precondition inference techniques (Kafle et al. 2018) to improve the precision of the termination proof when assertions are provided.

References

Albert, E., Arenas, P., Genaim, S., and Puebla, G. 2008. Automatic inference of upper bounds for recurrence relations in cost analysis. In *Proc. of SAS 2008*. LNCS, vol. 5079. Springer, 221–237.

Albert, E., Arenas, P., Genaim, S., Puebla, G., and Zanardini, D. 2007. Cost Analysis of Java Bytecode. In *Proc. of ESOP'07*. LNCS, vol. 4421. Springer, 157–172.

Albert, E., Correas, J., Ka I Pun, E. B. J., and Román-Díez, G. 2018. Parallel Cost Analysis. *ACM Trans. Comput. Log. 19,* 4, 1–37.

Borralleras, C., Brockschmidt, M., Larraz, D., Oliveras, A., Rodríguez-Carbonell, E., and Rubio, A. 2017. Proving termination through conditional termination. In *Proc. TACAS 2017*. LNCS, vol. 10205. Springer, 99–117.

Brockschmidt, M., Emmes, F., Falke, S., Fuhs, C., and Giesl, J. 2016. Analyzing runtime and size complexity of integer programs. *ACM Trans. Program. Lang. Syst. 38,* 4, 13:1–13:50.

Cousot, P. and Halbwachs, N. 1978. Automatic discovery of linear restraints among variables of a program. In *Proc. POPL 1978*. ACM, 84–96.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press

Debray, S. K. and Lin, N. 1993. Cost analysis of logic programs. *ACM Trans. Program. Lang. Syst. 15,* 5, 826–875.

Debray, S. K., López-García, P., Hermenegildo, M. V., and Lin, N. 1994. Estimating the computational cost of logic programs. In *Proc. SAS 1994*. LNCS, vol. 864. Springer, 255–265.

Flores-Montoya, A. 2017. Cost analysis of programs based on the refinement of cost relations. Ph.D. thesis, Darmstadt University of Technology, Germany.

Flores-Montoya, A. and Hähnle, R. 2014. Resource analysis of complex programs with cost equations. In *Proc. APLAS 2014*. LNCS, vol. 8858. Springer, 275–295.

Garcia, A., Laneve, C., and Lienhardt, M. 2015. Static analysis of cloud elasticity. In *Proc. PPDP 2015*. ACM, 125–136.

Giesl, J., Thiemann, R., Schneider-Kamp, P., and Falke, S. 2004. Automated termination proofs with aprove. In *Proc. RTA 2004*. LNCS, vol. 3091. Springer, Aachen, Germany, 210–220.

Grech, N., Georgiou, K., Pallister, J., Kerrison, S., Morse, J., and Eder, K. 2015. Static analysis of energy consumption for LLVM IR programs. In *Proc. SCOPES 2015*. ACM, 12–21.

Gulwani, S., Jain, S., and Koskinen, E. 2009. Control-flow refinement and progress invariants for bound analysis. In *Proc. of PLDI 2009*. ACM, 375–385.

Kafle, B., Gallagher, J. P., Gange, G., Schachte, P., Søndergaard, H., and Stuckey, P. J. 2018. An iterative approach to precondition inference using constrained horn clauses. *Theory Pract. Log. Program. 18,* 3-4, 553–570.

Liqat, U., Georgiou, K., Kerrison, S., López-García, P., Gallagher, J. P., Hermenegildo, M. V., and Eder, K. 2015. Inferring parametric energy consumption functions at different software levels: ISA vs. LLVM IR. In *Proc. FOPARA 2015, Selected Papers*. LNCS, vol. 9964. Springer, 81–100.

Navas, J. A., Mera, E., López-García, P., and Hermenegildo, M. V. 2007. User-definable resource bounds analysis for logic programs. In *Proc. ICLP 2007*. LNCS, vol. 4670. Springer, 348–363.

Serrano, A., López-García, P., Bueno, F., and Hermenegildo, M. V. 2013. Sized type analysis for logic programs. *Theory Pract. Log. Program. 13,* 4-5-Online-Supplement.

Sharma, R., Dillig, I., Dillig, T., and Aiken, A. 2011. Simplifying loop invariant generation using splitter predicates. In *Proc. of CAV 2011*. Springer, 703–719.

Sinn, M., Zuleger, F., and Veith, H. 2014. A simple and scalable static analysis for bound analysis and amortized complexity analysis. In *Proc. CAV 2014*. LNCS, vol. 8559. Springer, 745–761.

Spoto, F., Mesnard, F., and Payet, É. 2010. A termination analyzer for java bytecode based on path-length. *ACM Trans. Program. Lang. Syst. 32,* 3, 8:1–8:70.

Wegbreit, B. 1975. Mechanical Program Analysis. *Communications ACM 18,* 9, 528–539.

https://doi.org/10.1017/S1471068419000152 Published online by Cambridge University Press