

Ranking Abstractions

Aziem Chawdhary¹, Byron Cook², Sumit Gulwani², Mooly Sagiv³, and
Hongseok Yang¹

¹ Queen Mary, University of London

² Microsoft Research

³ Tel Aviv University

Abstract. We propose an abstract interpretation algorithm for proving that a program terminates on all inputs. The algorithm uses a novel abstract domain which uses ranking relations to conservatively represent relations between intermediate program states. One of the attractive aspects of the algorithm is that it abstracts information that is usually not important for proving termination such as program invariants and yet it distinguishes between different reasons for termination which are not usually maintained in existing abstract domains. We have implemented a prototype of the algorithm and shown that in practice it is fast and precise.

1 Introduction

This paper develops sound algorithms for inferring that C programs terminate on all possible inputs. The oldest trick in the book of termination proofs for programs (e.g., [18]) is the ranking function proof. In this method, we find a function p that maps program states into a well-founded ordered set, such that $p(\sigma) > p(\sigma')$ whenever σ' is a state reachable from state σ .

Despite the enormous progress in synthesizing ranking functions (e.g., [3]), modern programming language features such as nested loops lead to non-linear behaviours which make it hard to apply existing techniques to synthesize ranking functions in a sound and precise way directly to the C code.

Recently, [17] introduced the *disjunctive well-foundedness* principle in order to split the termination argument into multiple ranking relations, corresponding to different situations in the program. The main idea is to use a finite set of ranking functions r_1, r_2, \dots, r_n each of which is well-founded, and to require in addition that the relation between any two intermediate states in the program is included in one of the relations, i.e.,

$$\tau^+ \subseteq \bigcup_{i=1}^n r_i \quad (1)$$

where τ is the transition system describing the meaning of the program and τ^+ is the non-reflexive transitive closure of τ . This principle localizes termination proofs by allowing the use of simpler ranking function synthesizers to handle more complicated termination proofs.

However, [17] leaves two open problems: (a) what is the best way to find the set of ranking functions r_1, r_2, \dots, r_n and (b) how to effectively check the condition in

Eq. 1. Notice that this is a safety question which can be attacked by any abstract interpreter [10]. However it may be expensive to check the condition by the abstract interpreter or the interpreter may fail due to imprecision.

In this paper we solve these two problems together in a novel way. The first problem is solved by developing abstract domains which are parameterized by sets of ranking functions. The meaning of each of the relations (ranking functions) overapproximates the relations between intermediate states in the program. We employ standard iterative fixpoint computations to compute a set of ranking functions or determine that the program may diverge. The ranking synthesizer is invoked with larger and larger relations obtained by composing the current approximation with every possible command. Notice that calling the ranking synthesizer allows us to abstract away information that is not necessary for termination, but maintains enough distinctions between different ranking functions. When a fixpoint is reached the condition in Eq. 1 is guaranteed to hold and thus there is no need to perform the inclusion check above. The efficiency provided by our domain is underlined by result which we prove, that, for a particular base abstract domain, fixpoint calculations are guaranteed to converge, at most, in two steps. For more refined domains we lose the guarantee of two, but in our experimental results we find that fixpoints converge in few iterations.

Related Work Program termination has been studied extensively with many impressive algorithms for automatically inferring termination for functional (e.g., [13]), logic (e.g., [6, 4]) and imperative programs (e.g., [3, 7, 19, 1]). The result in [17] encourages the use of existing safety analyzers in order to prove termination (e.g., SLAM [7] or Octagon [2]). The point of departure of this work is to define a new abstract domain, designed with termination in mind, rather than to re-use existing domains for safety. Termination analysis requires a precise treatment of disjunction, and information about well-foundedness, and we suggest that domains which target these properties will be more appropriate for termination analysis than domains designed for wholly other purposes. Our work follows [7, 2] by employing the disjunctive well foundedness principle [17] in order to split the termination argument into multiple ranking relations corresponding to different situations in the program.

By tailoring our abstract domain to termination we obtain a very efficient termination prover for imperative programs. In particular it is faster than TERMINATOR, which relies on SLAM [7], and variance analyses based on Octagon or Polyhedra [2]. The variance analysis we describe in this paper uses rank functions natively, in contrast to the non-native variance analyses proposed in [2] which were constructed from existing domains for invariance. In contrast to [2] we directly abstract ranking relations which allow us to be more precise in the cases where the underlying abstract domain used for invariance analysis is too coarse (e.g., non-disjunctive) and our analysis can be more efficient when the underlying domain records complicated invariants that are not needed for proving termination. In contrast to [7], we iteratively compute ranking functions without the use of counterexample guided refinement.

Our abstract domain is related to the abstraction used in size-change termination [13]. In both cases, program fragments are abstracted in terms of measures decreased or preserved by the fragments. The major difference is that our domain contains only those abstract elements that mean terminating program fragments (unless the elements

are \top), whilst size-change termination analyses can have an (non- \top) abstract element that denotes a diverging program fragment. As a result, size-change termination analyses have to check whether (the concretization of) an abstracted program terminates, whereas our analysis can skip this rather expensive checking.

2 Informal description of the analysis

In this section we informally describe the new analysis using an example. Later, in Section 3, we provide a more formal description.

Consider the program:

```

1   while (x>0 ∧ y>0) {
2       if (*) then { x=x-1; y=*; } else { y=y-1; }
3   }
```

This program illustrates the limitation of known termination analyses. The Octagon-based and Polyhedra-based termination analyses from [2] can quickly (*i.e.* in 0.02s) infer that the relation $'x \geq 0 \wedge 'x \geq x$ holds between any state at $\ell=2$ and any previous state at $\ell=2$, where $'x$ and x denote previous and current values of x respectively. (Note that $'x$ is denoting *some* previous value of x , and not necessarily the *last* value). Unfortunately, this relation is insufficient to prove termination of the loop, as it is not (disjunctively) well-founded—the condition sufficient for proving termination as described in [2].

In contrast TERMINATOR can prove the example terminating, but at a great cost (16s). TERMINATOR finds the following disjunctively well-founded relation at $\ell=2$:

$$('x \geq 0 \wedge 'x-1 \geq x) \vee ('y \geq 0 \wedge 'y-1 \geq y)$$

To find this relation TERMINATOR performs three rounds of refinement on the relation itself and 9 rounds of abstraction/refinement for the checking of the 3 candidate assertions, resulting in the discovery of 21 transition predicates.

The termination analysis in this paper gives us TERMINATOR's accuracy at the speed of the Octagon-based termination analysis. The new analysis finds the relation

$$('x \geq 0 \wedge 'x-1 \geq x) \vee ('y \geq 0 \wedge 'y-1 \geq y \wedge 'x=x)$$

in 0.02s.

Concretely, the new analysis uses a disjunctive domain of ranking relations conjoined with the information about unchanged variables. That is: disjunctions of relations of the form $T_e \wedge V_X$, where

$$V_X \stackrel{\text{def}}{=} \bigwedge_{x \in X} 'x=x, \quad T_e \stackrel{\text{def}}{=} 'e \geq 0 \wedge 'e-1 \geq e,$$

and $'e$ is the expression e with all variables x replaced by their corresponding pre-primed versions $'x$. Let R represent the transition relation of the loop body of our program in DNF:

$$R \stackrel{\text{def}}{=} C_1 \vee C_2, \\ C_1 \stackrel{\text{def}}{=} 'x > 0 \wedge 'y > 0 \wedge x='x-1, \quad C_2 \stackrel{\text{def}}{=} 'x > 0 \wedge 'y > 0 \wedge x='x \wedge y='y-1.$$

Our analysis begins by taking each disjunct in R and performing rank-function synthesis on it. In this case we get

$$\text{RFS}(C_1) = x \quad \text{and} \quad \text{RFS}(C_2) = y.$$

For each disjunct, the analysis also computes a set of variables whose values do not change. In this example, it determines that C_1 can change both x and y , but C_2 does not change variable x . Thus, we begin our analysis with the initial abstract state $A_0 \stackrel{\text{def}}{=} T_x \vee (T_y \wedge V_{\{x\}})$, that is,

$$A_0 = ('x \geq 0 \wedge 'x-1 \geq x) \vee ('y \geq 0 \wedge 'y-1 \geq y \wedge 'x=x).$$

Note that A_0 overapproximates the loop body R .

The meaning of this initial abstract state (*i.e.* $\gamma(A_0)$) is set of all finite sequences of program states $s_i s_{i+1} \dots s_{i+n}$ such that

$$(s_i(x) \geq 0 \wedge s_i(x) - 1 \geq s_{i+n}(x)) \vee (s_i(y) \geq 0 \wedge s_i(y) - 1 \geq s_{i+n}(y) \wedge s_i(x) = s_{i+n}(x)).$$

The analysis then computes the next abstract state A_1 that overapproximates the relational composition of A_0 and R . It takes each disjunction from A_0 and each disjunction from R , composes them, performs rank function synthesis, infers variables that do not change, and constructs the union of the new ranking relations together with A_0 . In this case we find:

$$\begin{array}{ll} \text{RFS}(T_x; C_1) = x & \text{RFS}(T_x; C_2) = x \\ \text{RFS}((T_y \wedge V_{\{x\}}); C_1) = x & \text{RFS}((T_y \wedge V_{\{x\}}); C_2) = y \end{array}$$

We also find that the last composition $(T_y \wedge V_{\{x\}}; C_2)$ does not change x . Thus,

$$A_1 = (A_0 \vee T_x \vee T_x \vee T_x \vee (T_y \wedge V_{\{x\}})) = A_0.$$

Since A_0 is a fixpoint and A_0 overapproximates R , we know that $\forall i > 0. R^i \subseteq A_0$, that is, $R^+ \subseteq A_0$. Thus, because A_0 is disjunctively well-founded, [17] tells us that R is well-founded—meaning that the loop of our program guarantees termination.

Note that rank function synthesis is extremely efficient, meaning that our implementation of the analysis can compute the relation A_0 for $\ell = 2$ as fast as the Octagon-based termination analyzer (*i.e.* in 0.02s) [2]. In contrast to the Octagon-based analyzer, however, we compute a relation that is sufficiently strong to establish termination.

To sum up, the essence of our method is that we symbolically execute the body of the loop, and then perform abstraction by calling a rank synthesis engine. This in effect abstracts all information except those that are relevant to termination.

3 Formal description

In this section we provide a rigorous description of the proposed termination analysis.

3.1 Programming language

We consider a simple while language in the paper. Let Vars be a finite set of program variables x, y, z, \dots and let r represent real numbers.

$$\begin{aligned} e &::= x \mid r \mid e + e \mid r \times e \\ b &::= e = e \mid e \neq e \mid b \wedge b \mid b \vee b \mid \neg b \\ a &::= x := e \mid x := * \mid \text{assume}(b) \\ c &::= a \mid c; c \mid \text{while } b \text{ } c \mid c \square c \end{aligned}$$

Note that the language has two forms of assignments, normal assignment $x := e$ and nondeterministic random assignment $x := *$. The nondeterministic assignment is used to model some features of a common programming language, for example C, that are not covered by our language above. Also notice that the language does not include the conditional statement. It can be encoded with `assume` and the nondeterministic choice operator \square : $(\text{if } b \text{ } c_0 \text{ } c_1) \stackrel{\text{def}}{=} ((\text{assume}(b); c_0) \square (\text{assume}(\neg b); c_1))$.

The semantics of our language is standard. We remind the reader of only the storage model used in the semantics:

$$\text{St} \stackrel{\text{def}}{=} \text{Vars} \rightarrow \text{Real}.$$

This model shows that we assume real variables in this paper. However, changing the type of variables from reals to integers or rationals will not affect the results of the paper, except the ones for the fast termination in Lemma 1 and Theorem 2.

3.2 Abstract domain

Our analysis is parameterized by a domain for representing relations on states. The domain is specified by the following data:

1. A set D and a monotone function $\gamma_r : D \rightarrow \mathcal{P}(\text{St} \times \text{St})$ (where the target $\mathcal{P}(\text{St} \times \text{St})$ is ordered by the subset relation).
2. An abstract identity element d_{id} in D , that satisfies

$$\Delta_{\text{St}} \subseteq \gamma_r(d_{\text{id}})$$

where Δ_{St} is the identity relation on St .

3. An operator $\text{RFS} : D \rightarrow \mathcal{P}_{\text{fin}}(D) \uplus \{\top\}$, which synthesizes ranking functions. We assume the following two conditions for this operator:
 - (a) RFS computes an overapproximation:

$$\text{RFS}(d) \neq \top \implies \gamma_r(d) \subseteq \bigcup \{ \gamma_r(d') \mid d' \in \text{RFS}(d) \}.$$

- (b) $\text{RFS}(d)$ denotes a well-founded relation:

$$\text{RFS}(d) \neq \top \implies \bigcup \{ \gamma_r(d') \mid d' \in \text{RFS}(d) \} \text{ is well-founded.}$$

4. An abstract transfer function $\text{trans}(a)$ for each atomic commands a (i.e., assignments or assume statements). The function $\text{trans}(a)$ has type $D \rightarrow \mathcal{P}_{\text{fin}}(D)$, and satisfies

$$\forall d \in D. (\gamma_r(d); \llbracket a \rrbracket) \subseteq \bigcup \{ \gamma_r(d') \mid d' \in \text{trans}(a)(d) \}$$

where the semicolon means the usual composition of relations and $\llbracket a \rrbracket$ is the standard relational meaning of the atomic command a .

5. An abstract composition operator $\text{comp}: D \times D \rightarrow D$ such that

$$\gamma_r(d); \gamma_r(d') \subseteq \gamma_r(\text{comp}(d, d')).$$

Intuitively, the data above means that we have a set D of relations, some of which are well-founded. This set comes with an algorithm RFS, which overapproximates a relation by a ranking relation. It also has operators, trans and comp , that soundly model all the atomic commands and concrete relation composition. One example of D is the set of conjunction of linear constraints. In this case, we can use a linear rank synthesis engine, which we denote `LINEARRANKSYN`, and define RFS as will be shown in Section 3.4.

The abstract domain \mathcal{A} of our analyzer is:

$$\mathcal{A} \stackrel{\text{def}}{=} (\mathcal{P}_{\text{fin}}(D))^\top \quad (\text{i.e., } \mathcal{P}(D) \uplus \{\top\}).$$

It is ordered by the the subset order \sqsubseteq extended with \top . That is, $A \sqsubseteq A'$ iff

$$A' = \top, \quad \text{or} \quad (A, A' \in \mathcal{P}_{\text{fin}}(D) \text{ and } A \subseteq A').$$

Each abstract element A in \mathcal{A} denotes a set of finite or infinite sequences of states, which we call *traces*. The element \top denotes the set of all traces, including infinite ones, and non- \top elements A denote a set of *finite* nonempty traces whose initial and final states are related by some d in A . Let $\gamma_r(A)$ be $\bigcup \{ \gamma_r(d) \mid d \in A \}$, the disjunction of d 's in A , and define \mathcal{T} to be the set of all nonempty traces:

$$\mathcal{T} \stackrel{\text{def}}{=} \text{St}^+ \cup \text{St}^\infty.$$

The formal meaning of A is given by a concretization function γ :

$$\begin{aligned} \gamma &: \mathcal{A} \rightarrow \mathcal{P}(\mathcal{T}) \\ \gamma(A) &\stackrel{\text{def}}{=} \text{if } (A = \top) \text{ then } \mathcal{T} \text{ else } \{ \tau \mid \tau \text{ is nonempty, finite, and } \tau_0[\gamma_r(A)]\tau_{|\tau|-1} \} \end{aligned}$$

where $|\tau|$ is the length of the trace τ , and τ_n is the n -th state of the trace τ , and notation $s[r]s'$ means that s, s' are related by r . For instance, when $[x : n, y : m]$ is a state mapping x and y to n and m , a finite trace

$$[x : 1, y : 1][x : 2, y : 2][x : 5, y : 3][x : -2, y : 2]$$

belongs to $\gamma(\{ 'x-1 \geq x, 'y-1 \geq y \})$, because x has a smaller value in the final state than in the initial state.

Our domain \mathcal{A} is a complete semi-lattice. The join of a family $\{A_i\}_{i \in I}$ of elements in \mathcal{A} is given by the union of all A_i 's, if none of A_i 's is \top and the union is finite. Otherwise, the join is \top .

3.3 Generic analysis

Our generic analyzer is an abstract interpretation, defined in a denotational style.

For functions $f: D \rightarrow \mathcal{A}$ and $g: D \times D \rightarrow D$, let f^\dagger, g^\dagger be their liftings on \mathcal{A} :

$$\begin{aligned} f^\dagger &: \mathcal{A} \rightarrow \mathcal{A} & g^\dagger &: \mathcal{A} \times \mathcal{A} \rightarrow \mathcal{A} \\ f^\dagger(A) &\stackrel{\text{def}}{=} \text{if } (A = \top) \text{ then } \top \text{ else } \bigsqcup_{d \in A} f(d) \\ g^\dagger(A, B) &\stackrel{\text{def}}{=} \text{if } (A = \top \vee B = \top) \text{ then } \top \text{ else } \bigsqcup_{d \in A, d' \in B} \{g(d, d')\}. \end{aligned}$$

Using these liftings, we define the generic analyzer as follows:⁴

$$\begin{aligned} \llbracket c \rrbracket^\# &: \mathcal{A} \rightarrow \mathcal{A} \\ \llbracket a \rrbracket^\# A &\stackrel{\text{def}}{=} (\text{trans}(a))^\dagger A \\ \llbracket c_0; c_1 \rrbracket^\# A &\stackrel{\text{def}}{=} (\llbracket c_1 \rrbracket^\# \circ \llbracket c_0 \rrbracket^\#) A \\ \llbracket c_0 \sqcup c_1 \rrbracket^\# A &\stackrel{\text{def}}{=} \llbracket c_0 \rrbracket^\# A \sqcup \llbracket c_1 \rrbracket^\# A \\ \llbracket \text{while } b \text{ c} \rrbracket^\# A &\stackrel{\text{def}}{=} \text{let } F \stackrel{\text{def}}{=} \lambda A'. \llbracket \text{assume}(b); c \rrbracket^\# (\{d_{\text{id}}\} \sqcup A') \\ &\quad \text{in } \llbracket \text{assume}(\neg b) \rrbracket^\# (\text{comp}^\dagger(A, \text{fix}(\text{RFS}^\dagger \circ F))) \end{aligned}$$

Intuitively, the argument A represents a set of finite or infinite traces that happen before the command c . The analyzer computes an overapproximation of all traces that are obtained by continuing the execution of c from the end of traces in A .

Our definition assumes an operator fix . The fix operator takes a function of the form $\text{RFS}^\dagger \circ F: \mathcal{A} \rightarrow \mathcal{A}$, and returns an abstract element A in the image of RFS^\dagger such that

$$A = \top \vee (A \neq \top \wedge (\text{RFS}^\dagger \circ F)(A) \neq \top \wedge \gamma_r((\text{RFS}^\dagger \circ F)(A)) \subseteq \gamma_r(A)).$$

One can use the standard fixpoint iteration to define fix ,⁵ because the above condition holds for all post fixpoints A of $(\text{RFS}^\dagger \circ F)$ (that are in the image of RFS^\dagger). However, this is not mandatory. In fact, a more optimized fix operator is used in the analysis of Section 3.4, which in some cases does not even compute a post fixpoint.

The most interesting case of the analysis is the loop. The best way to understand this case is to assume that fix is the standard fixpoint operator and to see a sequence generated during the iterative fixpoint computation:

$$\begin{aligned} A_0 &= \{\}, \\ A_1 &= A_0 \sqcup (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\} \\ &= (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\} \\ A_2 &= A_1 \sqcup (\text{RFS}^\dagger \circ F)(\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\}) \\ &= (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)^2\{d_{\text{id}}\}, \\ A_3 &= A_2 \sqcup (\text{RFS}^\dagger \circ F)(\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)^2\{d_{\text{id}}\}) \\ &= (\text{RFS}^\dagger \circ F)\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)^2\{d_{\text{id}}\} \sqcup (\text{RFS}^\dagger \circ F)^3\{d_{\text{id}}\}, \\ &\quad \dots \end{aligned}$$

⁴ In the definition, we view RFS , $\text{trans}(a)$ as functions of type $D \rightarrow (\mathcal{P}_{\text{fin}}(D))^\top$.

⁵ In this case, $\text{fix}(\text{RFS}^\dagger \circ F)$ is defined by the limit of the sequence $\{A_n\}$ where $A_0 = \{\}$ and $A_{n+1} = A_n \sqcup (\text{RFS}^\dagger \circ F)(A_n)$.

Here we used the fact that $\text{RFS}^\dagger \circ F$ preserves \sqcup . Note that in each step, we apply the lifted rank-synthesis algorithm RFS^\dagger to the analysis result of the loop body $F(A_n)$. This application of RFS throws away all the information from $F(A_n)$, except the one necessary for proving termination. Another thing to note is that the input A is not used in this fixpoint computation at all. As the expansion of A_3 shows, the fixpoint computation effectively starts with $(\text{RFS}^\dagger \circ F)\{d_{\text{id}}\}$, which means the results of running the loop body once on all states. The input A is pre-composed later to the computed fixpoint. This change of the starting point is crucial for the soundness of our analysis, because it ensures that the analyzer overapproximates the relation between any states (not just initial states) at a loop and the following states at the same loop (so that we can apply a known termination proof rule based on disjunctively well-founded relations [17]).

Given a program c , the analyzer works as follows:

$\text{ANALYSIS}(c) \stackrel{\text{def}}{=} \text{let } A = \llbracket c \rrbracket^\#(\{d_{\text{id}}\})$
 $\quad \text{in if } (A \neq \top) \text{ then (return "Terminates")} \text{ else (return "Unknown").}$

Theorem 1. *If $\text{ANALYSIS}(c)$ returns “Terminates”, then c terminates on all states.*

The proof of this theorem is given in the full version of the paper [5]. There we also clarify what we mean by “terminates on all states”, by defining a concrete trace semantics of commands based on Cousot’s work [9].

3.4 Linear Rank Abstraction

The linear rank abstraction is an instance of our generic analysis, by the domain of linear constraints and a linear ranking synthesis algorithm `LINEARRANKSYN`

Let r represent real numbers. Consider constraints C defined by the grammar below:

$$\begin{aligned} E &::= x \mid 'x \mid x' \mid r \mid E + E \mid r \times E \\ P &::= E = E \mid E \neq E \mid E < E \mid E > E \mid E \leq E \mid E \geq E \\ C &::= P \mid \text{true} \mid C \wedge C \end{aligned}$$

This grammar ensures that all the constraints are the conjunction of linear constraints. Note that a constraint can have three kinds of variables; a normal variable x denoting the current value of program variable x ; a pre-primed variable $'x$ storing the initial value of x ; post-primed variables x' that usually denotes values which were once stored in program variables during computation. We assume that there are finitely many normal variables (`Vars`) and finitely many pre-primed variables (`'Vars`), and that there is a one-to-one correspondence between these two kinds of variables. For post-primed variables, however, we assume an infinite set.

Each constraint means a relation on `St`. For each state s , let $'s$ be a function from `'Vars` to `Real` such that for every pre-primed variable $'x$, $'s('x)$ is $s(x)$ for the corresponding normal variable x . The meaning function γ_r of constraints C is defined as follows:

$$\gamma_r(C) \stackrel{\text{def}}{=} \{(s_0, s_1) \mid ('s_0, s_1 \models \exists X'. C)\}$$

where X' is the set of post-primed variables in C and \models is the usual satisfaction relation in first-order logic. Note that all post-primed variables in the constraint C are implicitly existentially-quantified.

The linear rank abstraction uses the set of constraints C as the parameter set D of the generic analysis. The identity element d_{id} is the identity relation

$$d_{id} \stackrel{\text{def}}{=} \bigwedge_{x \in \text{Vars}} 'x=x.$$

Assume that we are given an enumeration x_0, \dots, x_n of all program variables in Vars. Call an expression E *normalized*, when (1) E does not contain any pre or post primed variables and (2) it is of the form $a_{i_0} \times x_{i_0} + \dots + a_{i_k} \times x_{i_k} + a$ with $a_{i_0} = 1$ or -1 and $i_0 < i_1 \dots < i_k$. Note that in a normalized expression E , the coefficient of the first variable in E according to the given enumeration is 1 or -1 . Conceptually, LINEARRANKSYN implements a function of the type:⁶

$$D \rightarrow (\{(E, r) \mid E \text{ is normalized and } r \text{ is a positive real}\}) \uplus \{\top\}.$$

The output \top indicates that the algorithm fails to discover a ranking function, because (the implementation of) the algorithm is incomplete or the input constraint defines a non-well-founded relation between pre-primed variables and normal variables. The other output (E, r) means that the algorithm succeeds to find a ranking function which overapproximates the given constraint. Concretely, for a normalized expression E and a positive real r , let

$$T_{E,r} \stackrel{\text{def}}{=} ('E \geq 0 \wedge 'E-r \geq E),$$

where expression $'E$ is E with all normal variables x replaced by corresponding pre-primed variables $'x$. The output (E, r) of LINEARRANKSYN(C) means that

$$(\exists X'. C) \implies T_{E,r}$$

where X' is the set of all post-primed variables in C .

Assume that we have chosen a fixed positive real dec for the analysis, which is very small (in particular smaller than 1). Using LINEARRANKSYN and dec , we define the operator RFS as follows:

$$\text{RFS}(C) \stackrel{\text{def}}{=} \begin{cases} \{\} & \text{if } C \vdash \text{false} \\ \{T_{E,\text{dec}}\} & \text{else if LINEARRANKSYN}(C) = (E, r) \text{ and } r \geq \text{dec} \\ \top & \text{otherwise} \end{cases}$$

where \vdash is a sound (but not necessarily complete) theorem prover. Note that the result of RFS is always of the form $T_{E,\text{dec}}$, so the second subscript of T is not necessary. From now on, we write T_E for $T_{E,\text{dec}}$.

⁶ Usually the implementation of linear rank synthesis returns a tuple (E, r, b) where E is an expression without any pre or post primed variable whose value is decreasing, r is a decrement, and b is a lower bound of E . Our analysis picks the absolute value a of the coefficient of the first variable x_i in E , transforms E/a to a normal form E' , and regards $(E' - b/a, r/a)$ as an output from LINEARRANKSYN.

The abstract transfer functions for atomic commands are defined following Floyd's strongest postcondition semantics:

$$\begin{aligned} \llbracket x:=* \rrbracket^\# C &\stackrel{\text{def}}{=} \{C[x'/x]\} \quad (x' \text{ is fresh}) \\ \llbracket x:=e \rrbracket^\# C &\stackrel{\text{def}}{=} \{C[x'/x] \wedge x=(e[x'/x])\} \quad (x' \text{ is fresh}) \\ \llbracket \text{assume}(b) \rrbracket^\# C &\stackrel{\text{def}}{=} \text{if } (C \wedge b \vdash \text{false}) \text{ then } \{\} \\ &\quad \text{else } \{C_0, \dots, C_n \mid C_0 \vee \dots \vee C_n = \text{norm}(C \wedge b)\}. \end{aligned}$$

Here norm is the standard transformation that takes a formula in the propositional logic and transforms the formula to disjunctive normal form.

Next, we define the abstract composition comp . Let fresh be an operator on constraints C that renames all post-primed variables fresh. Let 'Vars be the set of pre-primed variables. The abstract composition is defined as follows

$$\text{comp}(C_0, C_1) \stackrel{\text{def}}{=} \text{let } (C_2 = \text{fresh}(C_1)) \text{ in } (C_0[Y'/\text{Vars}] \wedge C_2[Y'/\text{'Vars}]).$$

The variable set Y' in the definition denotes a set of fresh post-primed variables, that has as many elements as Vars . The two substitutions there replace a normal variable x and the corresponding pre-primed variable $\text{'}x$ by the same post-primed variable x' .

Finally, we specify a fix operator. For each function $(\text{RFS}^\dagger \circ F)$ on sets of constraints C , let $\{G_n\}_n$ be the standard fixpoint iteration sequence: $G_0 = \{\}$ and $G_{n+1} = G_n \sqcup (\text{RFS}^\dagger \circ F)(G_n)$. Given G , our fix operator returns the first G_n such that

$$G_n = \top \quad \vee \quad (G_n \neq \top \wedge G_{n+1} \neq \top \wedge \forall C \in G_{n+1}. \exists C' \in G_n. C \vdash C').$$

This definition assumes that some G_n satisfies the above property. If such a G_n does not exist, the fix operator is not defined, so the analysis can diverge during the fixpoint computation. In Theorem 2, we will discharge this assumption and prove the termination of the linear rank abstraction.

Example 1. Consider the program c below:

$$\text{while } (x > 0 \wedge y > 0) \ (x:=x-1 \ \square \ y:=y-1).$$

Given c , the analysis starts the fixpoint computation from the empty set $A_0 = \{\}$. The first iteration of the fixpoint computation is done in two steps. First, it applies the abstract transfer function of the loop body to $\{d_{\text{id}}\} \cup A_0 = \{d_{\text{id}}\}$:

$$\begin{aligned} &\llbracket \text{assume}(x>0 \wedge y>0); (x:=x-1 \ \square \ y:=y-1) \rrbracket^\# (\{d_{\text{id}}\}) \\ &= \llbracket x:=x-1 \ \square \ y:=y-1 \rrbracket^\# \{d_{\text{id}} \wedge x>0 \wedge y>0\} \\ &= \llbracket x:=x-1 \rrbracket^\# \{d_{\text{id}} \wedge x>0 \wedge y>0\} \cup \llbracket y:=y-1 \rrbracket^\# \{d_{\text{id}} \wedge x>0 \wedge y>0\} \\ &= \llbracket x:=x-1 \rrbracket^\# \{x=x \wedge y=y \wedge x>0 \wedge y>0\} \cup \llbracket y:=y-1 \rrbracket^\# \{x=x \wedge y=y \wedge x>0 \wedge y>0\} \\ &= \{x=x' \wedge y=y \wedge x'>0 \wedge y>0 \wedge x=x'-1, \quad x=x \wedge y=y' \wedge x>0 \wedge y'>0 \wedge y=y'-1\}. \end{aligned}$$

Next, the analysis calls LINEARRANKSYN twice with each of the two elements in the result set above. These function calls return x and y , from which the analysis constructs

two ranking relations below:

$$T_x \stackrel{\text{def}}{=} ('x \geq 0 \wedge 'x\text{-dec} \geq x) \quad \text{and} \quad T_y \stackrel{\text{def}}{=} ('y \geq 0 \wedge 'y\text{-dec} \geq y).$$

The result A_1 of the first iteration is $\{T_x, T_y\}$.

The second fixpoint iteration computes:

$$A_1 \sqcup (\text{RFS}^\dagger \circ \llbracket \text{assume}(x > 0 \wedge y > 0); (x := x - 1 \parallel y := y - 1) \rrbracket^\#) A_1.$$

We show that the abstract element on the right hand side of the join, denoted A'_2 , is again A_1 , so that the fixpoint computation converges here. To compute A'_2 , the analyzer first transforms A_1 according to the abstract meaning of the loop body. This results in a set with four elements:

$$\{ T_x[x'/x] \wedge x' > 0 \wedge y > 0 \wedge x = x' - 1, \quad T_x[y'/y] \wedge x > 0 \wedge y' > 0 \wedge y = y' - 1, \\ T_y[x'/x] \wedge x' > 0 \wedge y > 0 \wedge x = x' - 1, \quad T_y[y'/y] \wedge x > 0 \wedge y' > 0 \wedge y = y' - 1 \}.$$

The first two elements come from transforming T_x according to the left and right branches of the loop body. The other two elements are obtained similarly from T_y . Next, the analysis calls LINEARRANKSYN with all the four elements above. These four calls return x, x, y and y , which represent well-founded relations T_x, T_x, T_y, T_y . Thus, A'_2 is the same as T_x and T_y , and the fixpoint computation stops here.

After the fixpoint computation, the analysis composes the identity relation $\{d_{\text{id}}\}$ with the result of the fixpoint computation:

$$\text{comp}^\dagger(\{d_{\text{id}}\}, \{T_x, T_y\}) = \{ 'x = x'_0 \wedge 'y = y'_0 \wedge T_x[x'_0/'x], 'x = x'_0 \wedge 'y = y'_0 \wedge T_y[y'_0/'y] \} \\ = \{ T_x, T_y \}.$$

Finally, we apply $\llbracket \text{assume}(\neg(x > 0 \wedge y > 0)) \rrbracket^\#$ to the set above, which gives a set with four constraints:

$$\{ T_x \wedge x \leq 0, \quad T_x \wedge y \leq 0, \quad T_y \wedge x \leq 0, \quad T_y \wedge y \leq 0 \}.$$

Since the result is not \top , the analysis concludes that the given program c terminates. \square

In the example above, the fixpoint computation converges after two iterations. In the first iteration, which computes A_1 , it finds ranking functions, and in the next iteration, it confirms that the ranking functions are preserved by the loop. In fact, we can prove that the fixpoint computation of the analysis always follows the same pattern, and finishes in two iterations. Suppose that LINEARRANKSYN is well-behaved, such that

1. RFS always computes an optimal ranking function, in the sense that

$$(\text{RFS}(C) = \{T_E\} \wedge \gamma_r(C) \subseteq \gamma_r(T_{E+b})) \implies b \geq 0,$$

2. RFS depends only on the (relational) meaning of its argument.

Lemma 1. *For all commands c and normalized expressions E , if there is a constraint $C \in \llbracket c \rrbracket^\# \{T_E\}$ such that $\text{RFS}(C) = \{T_F\}$ and $\gamma_r(C) \neq \emptyset$, then F is of the form $E - b$ for some nonnegative b .*

Proof. The proof appears in the full version of this paper [5]. □

Theorem 2 (Fast Convergence). *Suppose that the theorem prover \vdash is complete. Then, for all commands c , the fixpoint iteration of*

$$G = \lambda A. (\text{RFS}^\dagger \circ \llbracket c \rrbracket^\#)(\{d_{\text{id}}\} \sqcup A)$$

terminates at most in two steps. Specifically, $G^2(\{\})$ is \top , or the result of $\text{fix } G$ is $\{\}$ or $G(\{\})$.

Proof. Suppose that $G^2(\{\})$ is not \top . This implies that both $G(\{\})$ and $G^2(\{\})$ are finite sets of T_E 's for normalized expressions E , because $G(= \text{RFS}^\dagger \circ \llbracket c \rrbracket^\#)$ preserves \top . If $G(\{\})$ is empty, $\{\}$ is the fixpoint of G , thus becoming the result of $\text{fix } G$, as claimed in the theorem. To prove the other nonempty case, suppose that $G(\{\})$ is a nonempty finite collection $A = \{T_{E_1}, \dots, T_{E_n}\}$. We need to show that for each T_F in $G(A)$, there exists $T_{E_i} \in A$ such that $T_F \vdash T_{E_i}$, which is equivalent to $\gamma_r(T_F) \subseteq \gamma_r(T_{E_i})$ due to the completeness assumption about the prover. Pick T_F in $G(A)$. Since $G(= \text{RFS}^\dagger \circ \llbracket c \rrbracket^\#)$ preserves the join operator, there exists T_{E_i} in A such that $T_F \in G(\{T_{E_i}\})$. This means that $\text{RFS}(C) = \{T_F\}$ for some constraint C in $\llbracket c \rrbracket^\#(T_{E_i})$. Note that since RFS filters out all the provably inconsistent constraints and the prover is assumed complete, $\gamma_r(C)$ is not empty. Thus, by Lemma 1, there is a nonnegative b such that $F = E - b$. This gives the required $\gamma_r(T_F) \subseteq \gamma_r(T_{E_i})$. □

Note that the theorem suggests that we could have used a different fix operator that does not call the prover at all and just returns $G^2(\{\})$. We do not take this alternative in the paper, since it is too specific for the RFS operator in this section; if RFS also keeps track of equality information, this two-step convergence result no longer holds.

Refinement with simple equalities The linear rank abstraction cannot prove the termination of the program in Section 2. When the linear rank abstraction is run for the program, it finds the ranking functions x and y for the true and false branches of the program, but loses the information that the else branch does not change the value of x , which is crucial for the termination proof. As a result, the linear rank abstraction returns \top , and reports, incorrectly, the possibility of nontermination.

One way to solve this problem and improve the precision of the linear rank abstraction is to use a more precise RFS operator that additionally keeps simple forms of equalities. Concretely, this refinement keeps all the definitions of the linear rank abstraction, except that it replaces the rank synthesizer RFS of the linear rank abstraction by RFS' below:

$$\text{RFS}'(C) \stackrel{\text{def}}{=} \text{if } (\text{RFS}(C)=\top) \text{ then } \top \text{ else } \{T_E \wedge (\wedge_{(C \vdash \cdot x=x)} \cdot x=x) \mid T_E \in \text{RFS}(C)\}.$$

When this refined analysis is given the program in Section 2, it follows the informal description in that section and proves the termination of the program.

4 Experimental evaluation

In order to evaluate the utility of our approach we have implemented the analysis in this paper, and then compared it to several known termination tools. The tools used in the experiments are as follows:

- LR)** LINEARRANKTERM is the new variance analysis that implements the linear rank abstraction with simple equalities in Section 3.4. This tool is implemented using CIL [15] allowing the analysis of programs written in C. However, no notion of shape is used in these implementations, restricting the input to only arithmetic programs. The tool uses RANKFINDER [16] as its linear rank synthesis engine and uses the Simplify prover [11] to filter out inconsistent states and check the implication between abstract states.
- O)** OCTATERM is the variance analysis [2] induced by the octagon analysis OCTANAL [14].
- P)** POLYTERM is the variance analysis [2] similarly induced from the polyhedra analysis POLY based on the New Polka Polyhedra library [12].
- T)** TERMINATOR [8].

These tools, except for TERMINATOR, were all run on a 2GHz AMD64 processor using Linux 2.6.16. TERMINATOR was executed on a 3GHz Pentium 4 using Windows XP SP2. Using different machines is unfortunate but somewhat unavoidable due to constraints on software library dependencies, etc. Note, however, that TERMINATOR running on the faster machine was still slower overall, so the qualitative results are meaningful. In any case, the running times are somewhat incomparable since on failed proofs TERMINATOR produces a counterexample path, but LINEARRANKTERM, OCTATERM and POLYTERM give a suspect pair of states

Fig. 1 contains the results from the experiments performed with these analyses.⁷ For example, Fig. 1(a) shows the outcome of the provers on example programs included in the OCTANAL distribution. Example 3 is an abstracted version of heapsort, and Example 4 of bubblesort.

Fig. 1(b) contains the results of experiments on fragments of Windows device drivers. These examples are small because we currently must hand-translate them before applying all of the tools but TERMINATOR.

Fig. 1(c) contains the results from experiments with the 4 tools on examples from the POLYRANK distribution.⁸ The examples can be characterized as small but famously difficult (*e.g.* McCarthy's 91 function). Note that LINEARRANKTERM performs poorly on these examples because of the limitations of RANKFINDER. Many of these examples involve phase changes or tricky arithmetic in the algorithm.

From these experiments we can see that LINEARRANKTERM is very fast and precise. The prototype we have developed indicates that a termination analyzer using abstractions based on ranking functions shows a lot of promise.

⁷ The programs used in our experiments except the ones for drivers are available in <http://www.dcs.qmul.ac.uk/~aziem/esop>. Unfortunately, we could not put the driver examples in the web page, because that might cause a problem related to intellectual property.

⁸ Note also that there is no benchmark number 5 in the original distribution. We have used the same numbering scheme as in the distribution so as to avoid confusion.

	1	2	3	4	5	6
LR	0.01 ✓	0.01 ✓	0.08 ✓	0.09 ✓	0.02 ✓	0.06 ✓
O	0.11 ✓	0.08 ✓	6.03 ✓	1.02 ✓	0.16 ✓	0.76 ✓
P	1.40 ✓	1.30 ✓	10.90 ✓	2.12 ✓	1.80 ✓	1.89 ✓
T	6.31 ✓	4.93 ✓	T/O -	T/O -	33.24 ✓	3.98 ✓

(a) Results from experiments with termination tools on arithmetic examples from the Octagon Library distribution.

	1	2	3	4	5	6	7	8	9	10
LR	0.23 ✓	0.20 ⊗	0.00 ⊗	0.04 ✓	0.00 ✓	0.03 ✓	0.07 ✓	0.03 ✓	0.01 ⊗	0.03 ✓
O	1.42 ✓	1.67 ⊗	0.47 ⊗	0.18 ✓	0.06 ✓	0.53 ✓	0.50 ✓	0.32 ✓	0.14 ⊗	0.17 ✓
P	4.66 ✓	6.35 ⊗	1.48 ⊗	1.10 ✓	1.30 ✓	1.60 ✓	2.65 ✓	1.89 ✓	2.42 ⊗	1.27 ✓
T	10.22 ✓	31.51 ⊗	20.65 ⊗	4.05 ✓	12.63 ✓	67.11 ✓	298.45 ✓	444.78 ✓	T/O -	55.28 ✓

(b) Results from experiments with termination tools on small arithmetic examples taken from Windows device drivers. Note that the examples are small as they must currently be hand-translated for the three tools.

	1	2	3	4	6	7	8	9	10	11	12
LR	0.19 ✓	0.02 ✓	0.01 †	0.02 †	0.02 †	0.01 †	0.04 †	0.01 †	0.03 †	0.02 †	0.01 †
O	0.30 †	0.05 †	0.11 †	0.50 †	0.10 †	0.17 †	0.16 †	0.12 †	0.35 †	0.86 †	0.12 †
P	1.42 ✓	0.82 ✓	1.06 †	2.29 †	2.61 †	1.28 †	0.24 †	1.36 ✓	1.69 †	1.56 †	1.05 †
T	435.23 ✓	61.15 ✓	T/O -	T/O -	75.33 ✓	T/O -	10.31 †				

(c) Results from experiments with termination tools on arithmetic examples from the POLYRANK distribution.

Fig. 1. Experiments with 4 termination provers/analyses. **LR** is used to represent LINEARRANK-TERM, **O** is used to represent OCTATERM, an Octagon-based variance analysis. **P** is POLYTERM, a Polyhedra-based variance analysis. The **T** represents TERMINATOR [8]. Times are measured in seconds. The timeout threshold was set to 500s. ✓ = “a proof was found”. † = “false counterexample returned”. T/O = “timeout”. ⊗ = “termination bug found”. Note that pointers and aliasing from the device driver examples were removed by a careful hand translation when passed to the tools **O**, **P** and **LR**. Note that a time of 0.00 means that the analysis was too fast to be measured by the timing utilities used.

Acknowledgements. We would like to thank Peter O’Hearn for encouragements and insightful comments on our work, Andrey Rybalchenko for explaining the subtleties of linear ranking functions and RANKFINDER, and Amir Ben-Amram and Neil Jones for helping us to understand the size-change termination. We also acknowledge detailed comments on the paper from Amir and anonymous referees, which help us to improve the presentation of the paper. Chawdhary was supported by a Microsoft PhD studentship, and Yang was supported by EPSRC.

References

1. I. Balaban, A. Pnueli, and L. Zuck. Ranking abstraction as companion to predicate abstraction. In *FORTE’05*, 2005.
2. J. Berdine, A. Chawdhary, B. Cook, D. Distefano, and P. O’Hearn. Variance analyses from invariance analyses. In *POPL’07*, 2007.
3. A. Bradley, Z. Manna, and H. Sipma. Termination of polynomial programs. In *VMCAI’05*, 2005.
4. M. Bruynooghe, M. Codish, J. Gallagher, S. Genaim, and W. Vanhoof. Termination analysis through combination of type based norms. *ACM Trans. Program. Lang. Syst.*, 29(2), 2007.
5. A. Chawdhary, B. Cook, S. Gulwani, M. Sagiv, and H. Yang. Ranking abstractions. Manuscript, 2008. Available at <http://www.dcs.qmul.ac.uk/~aziem/paper/esop08-full.pdf>.
6. M. Codish and C. Taboch. A semantic basis for the termination analysis of logic programs. *The Journal of Logic Programming*, 41(1), 1999.
7. B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *PLDI’06*, 2006.
8. B. Cook, A. Podelski, and A. Rybalchenko. Terminator: Beyond safety. In *CAV’06*, 2006.
9. P. Cousot. Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theoretical Comput. Sci.*, 277(1–2):47–103, 2002.
10. P. Cousot and R. Cousot. Systematic design of program analysis frameworks. In *POPL’79*, 1979.
11. D. Detlefs, G. Nelson, and J. Saxe. Simplify: A theorem prover for program checking, 2003.
12. B. Jeannot. NewPolka polyhedra library. <http://pop-art.inrialpes.fr/people/bjeannot/newpolka/index.html>.
13. C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. In *POPL’01*, 2001.
14. A. Miné. The Octagon abstract domain. *Higher-Order and Symbolic Comput.*, 19:31–100, 2006.
15. G. Necula, S. McPeak, S. Rahul, and W. Weimer. CIL:intermediate language and tools for analysis and transformation of C programs. In *CC’02*, 2002.
16. A. Podelski and A. Rybalchenko. A complete method for the synthesis of linear ranking functions. In *VMCAI’04*, 2004.
17. A. Podelski and A. Rybalchenko. Transition invariants. In *LICS’04*, 2004.
18. A. M. Turing. Checking a large routine. In *Report of a Conference on High Speed Automatic Calculating Machines*, pages 67–69, 1948. Reprinted in: *The early British computer conferences*, vol. 14 of Charles Babbage Institute Reprint Series for the History of Computing, MIT Press, 1989.
19. E. Yahav, T. Reps, M. Sagiv, and R. Wilhelm. Verifying temporal heap properties specified via evolution logic. *Logic Journal of IGPL*, Sept. 2006.

Appendix

A Soundness

In this appendix, we prove the soundness of our analysis, stated in Theorem 1.

To express the soundness formally, we consider a trace semantics of commands. Recall that \mathcal{T} is $\text{St}^+ \cup \text{St}^\infty$, the set of all traces. Let T_s be the set of all singleton traces $\{s \mid s \in \text{St}\}$. Define a composition operator seq for trace sets:

$$\begin{aligned} \text{seq} &: \mathcal{P}(\mathcal{T}) \times \mathcal{P}(\mathcal{T}) \rightarrow \mathcal{P}(\mathcal{T}) \\ \text{seq}(T, T') &\stackrel{\text{def}}{=} \{\tau s \tau' \mid \tau s \in (T \cap \text{St}^+) \wedge s \tau' \in T'\} \cup (T \cap \text{St}^\infty) \end{aligned}$$

Following Cousot's work [9], we define the concrete collecting trace semantics below:

$$\begin{aligned} \llbracket c \rrbracket &\in \mathcal{P}(\mathcal{T}) \\ \llbracket a \rrbracket &\stackrel{\text{def}}{=} \{ss' \mid (s, s') \in \llbracket a \rrbracket\} \\ \llbracket c_1; c_2 \rrbracket &\stackrel{\text{def}}{=} \text{seq}(\llbracket c_1 \rrbracket, \llbracket c_2 \rrbracket) \\ \llbracket c_1 \sqcap c_2 \rrbracket &\stackrel{\text{def}}{=} \llbracket c_1 \rrbracket \cup \llbracket c_2 \rrbracket \\ \llbracket \text{while } b \text{ } c \rrbracket &\stackrel{\text{def}}{=} \mathbf{let } F = \lambda T. T_s \cup \text{seq}(\llbracket \text{assume}(b); c \rrbracket, T) \\ &\quad \mathbf{in } \text{seq}(\text{gfix } F, \llbracket \text{assume}(\neg b) \rrbracket) \end{aligned}$$

Note that we use greatest fixpoints in the semantics of loops. This is because the greatest fixpoints correctly include all infinite execution traces of the loops, if exist. If we computed the least fixpoints there, we would exclude all infinite traces, so that all programs would be considered terminating. Also note that in the semantics of atomic commands a , $\llbracket a \rrbracket$ on the right is the relational meaning of a , and $\llbracket a \rrbracket$ on the left is the trace semantics of a .

We re-state the soundness of the analysis (i.e., Theorem 1), formally this time using the trace semantics that we have just defined.

Theorem 1 (Soundness). *If $\text{ANALYSIS}(c)$ returns “Terminate”, then all traces in $\llbracket c \rrbracket$ are finite.*

One challenge for proving the above soundness theorem is to deal with greatest fixpoints in the semantics of loops, because our analyzer uses an overapproximation of least fixpoints. Our first step of the proof is, therefore, to rewrite the concrete trace semantics such that the new semantics does not change the meaning of any commands, but interprets loops in terms of least fixpoints.

Define an operator repeat which maps a trace set T to the set of *infinite* traces τ satisfying the following condition: there are infinitely many finite traces $s_0 \tau_0 s_1, s_1 \tau_1 s_2, \dots$ in T such that

$$(\tau = s_0 \tau_0 s_1 \tau_1 \dots) \wedge (\forall n. |s_n \tau_n s_{n+1}| \geq 2) \wedge (\forall n. s_n \tau_n s_{n+1} \in T).$$

Call a trace set T *progressing* if and only if all traces in T are of length at least 2.

Proposition 1 (Fixpoints Correspondence). *For all functions F on $\mathcal{P}(T)$, if F is of the form $\lambda T. T_s \cup \text{seq}(T_0, T)$ for some trace set T_0 , and T_0 is progressing, then*

$$(\text{gfix } F) = (\text{lfix } F) \cup \text{repeat}(T_0).$$

Proof. We first prove that

$$(\text{gfix } F) \supseteq (\text{lfix } F) \cup \text{repeat}(T_0).$$

We will show that $(\text{lfix } F) \cup \text{repeat}(T_0)$ is a pre fixpoint of F . Then, the required relationship will follow, because the greatest fixpoint of the monotone function F is also the greatest pre fixpoint.

We notice two facts about seq , which follow from the definitions of seq and repeat :

1. $\text{seq}(T_0, \bigcup_{i \in I} T'_i) = \bigcup_{i \in I} \text{seq}(T_0, T'_i)$ for all nonempty I .
2. $\text{seq}(T_0, \text{repeat}(T_0)) \supseteq \text{repeat}(T_0)$.

Using these facts, we show that $(\text{lfix } F) \cup \text{repeat}(T_0)$ is a pre fixpoint of F :

$$\begin{aligned} & F((\text{lfix } F) \cup \text{repeat}(T_0)) \\ &= T_s \cup \text{seq}(T_0, (\text{lfix } F) \cup \text{repeat}(T_0)) \quad (\because \text{Def. of } F) \\ &= T_s \cup \text{seq}(T_0, \text{lfix } F) \cup \text{seq}(T_0, \text{repeat}(T_0)) \quad (\because \text{Fact 1 of seq}) \\ &= F(\text{lfix } F) \cup \text{seq}(T_0, \text{repeat}(T_0)) \\ &= (\text{lfix } F) \cup \text{seq}(T_0, \text{repeat}(T_0)) \\ &\supseteq (\text{lfix } F) \cup \text{repeat}(T_0) \quad (\because \text{Fact 2 of seq}). \end{aligned}$$

Next, we prove

$$(\text{gfix } F) \subseteq (\text{lfix } F) \cup \text{repeat}(T_0).$$

Pick a trace τ in $\text{gfix } F$. Since $\text{gfix } F$ is a fixpoint of F ,

$$\tau \in T_s \cup \text{seq}(T_0, (\text{gfix } F)).$$

If τ is in T_s , it should belong to $\text{lfix } F$, because T_s is a subset of $\text{lfix } F$. Otherwise, there are two possibilities. The first possibility is that τ is an infinite trace in T_0 . In this case, τ should again belong to $\text{lfix } F$ as well, because T_0 is a subset of $\text{seq}(T_0, T_s)$, which is included in $\text{lfix } F$. The second possibility is that τ can be decomposed into a finite prefix $\tau's$ and a suffix τ'' such that

$$\tau = \tau's\tau'' \quad \wedge \quad \tau's \in T_0 \quad \wedge \quad |\tau's| \geq 2 \quad \wedge \quad s\tau'' \in \text{gfix } F.$$

The third conjunct $|\tau's| \geq 2$ above is obtained from the second, using the fact that T_0 is progressing. If we apply the same reasoning to $s\tau''$ recursively, then either we generate a finitely many traces $s_0\tau_0s_1, s_1\tau_1s_2, \dots, s_m\tau_m$ such that

$$\begin{aligned} \tau &= s_0\tau_0s_1\tau_1\dots s_m\tau_m \quad \wedge \\ &(\forall n < m. |s_n\tau_n s_{n+1}| \geq 2 \quad \wedge \quad s_n\tau_n s_{n+1} \in T_0) \quad \wedge \quad s_m\tau_m \in \text{lfix } F. \end{aligned}$$

This implies that τ is in the result of applying $\text{seq}(T_0, -)$ to $\text{lfix } F$ m -times; since this result is a subset of $\text{lfix } F$, trace τ has to be in $\text{lfix } F$. Or, there are infinitely many finite traces $s_0\tau_0s_1, s_1\tau_1s_2, \dots$, such that

$$\tau = s_0\tau_0s_1\tau_1s_2\dots \quad \wedge \quad (\forall n. |s_n\tau_n s_{n+1}| \geq 2 \quad \wedge \quad s_n\tau_n s_{n+1} \in T_0).$$

This implies that τ belongs to $\text{repeat}(T_0)$. Thus, in both cases, we have shown that trace τ is in $(\text{lfix } F) \cup \text{repeat}(T_0)$, as required. \square

Using the proposition, we simplify the semantics of the loops, which we will use in the remainder of this appendix:

Corollary 1. *For all loops $\text{while } b \ c$, we have that*

$$\begin{aligned} \llbracket \text{while } b \ c \rrbracket &= \text{let } T_0 = \llbracket \text{assume}(b); c \rrbracket \\ &\quad F = \lambda T. T_s \cup \text{seq}(T_0, T) \\ &\text{in } \text{seq}((\text{lfix } F) \cup \text{repeat}(T_0), \llbracket \text{assume}(\neg b) \rrbracket). \end{aligned}$$

Proof. The trace set $\llbracket \text{assume}(b); c \rrbracket$ is progressing. Thus, this corollary follows from Proposition 1. \square

The remainder of the soundness proof consists of three steps. First, it builds a relational semantics, and shows how the relational semantics is related to the original trace semantics. Next, we prove that our analysis overapproximates the relational semantics. Finally, we combine the results of the previous two steps and derive the soundness of the analysis. The following three subsections explain these three steps separately.

A.1 Relational semantics

The relational semantics is an abstraction of the concrete trace semantics, based on the domain of relations on St . We factor the soundness proof of the generic analyzer through the soundness of this relational semantics, because firstly this factoring simplifies the proof and secondly the relational semantics describes the limit of our generic analyzer; it beats all instances of our analyzer for accuracy.

The relational semantics interprets commands using the domain of relations extended with \top :

$$\text{Rels}^\top \stackrel{\text{def}}{=} \mathcal{P}^\top(\text{St} \times \text{St}) \quad (= \mathcal{P}(\text{St} \times \text{St}) \uplus \top).$$

This domain is ordered by the subset order extended with \top , and forms a complete lattice. The meaning of each element in the domain is given by the concretization map γ :

$$\begin{aligned} \gamma &: \text{Rels}^\top \rightarrow \mathcal{P}(\mathcal{T}) \\ \gamma(r) &\stackrel{\text{def}}{=} \begin{cases} \{\tau \mid \tau \text{ is finite and } (\tau_0, \tau_{|\tau|-1}) \in r\} & \text{if } r \in \text{Rels} \\ \mathcal{T} & \text{otherwise} \end{cases} \end{aligned}$$

Define a binary operator rseq on Rels^\top , which overapproximates the sequential composition operator seq for trace sets:

$$\begin{aligned} \text{rseq} &: \text{Rels}^\top \times \text{Rels}^\top \rightarrow \text{Rels}^\top \\ \text{rseq}(r, r') &\stackrel{\text{def}}{=} \begin{cases} r; r' & \text{if } r \neq \top \text{ and } r' \neq \top \\ \top & \text{otherwise.} \end{cases} \end{aligned}$$

Lemma 2. For all r and r' in Rels^\top ,

$$\text{seq}(\gamma(r), \gamma(r')) \subseteq \gamma(\text{rseq}(r, r')).$$

Proof. If r or r' is \top , so is $\text{rseq}(r, r')$. The lemma, then, easily follows. Suppose that both r, r' are relations on St . Pick τ from $\text{seq}(\gamma(r), \gamma(r'))$. Since $\gamma(r)$ and $\gamma(r')$ both contain only finite traces, τ should be of the form $\tau' s \tau''$, where $\tau' s$ and $s \tau''$ are finite traces belonging, respectively, to $\gamma(r)$ and $\gamma(r')$. By the definition of γ , the first and last states of $\tau' s$ are related by r , and those states of $s \tau''$ are related by r . Thus, the first and last states of τ has to be related by $r; r'$. This means that $\tau \in \gamma(\text{rseq}(r, r'))$, as required. \square

Call a relation $r \in \text{Rels}$ *disjunctively well-founded* if and only if there are finitely many well-founded relations r_1, \dots, r_n satisfying

$$r \subseteq r_1 \cup \dots \cup r_n.$$

Let $\text{DISJWELLFOUNDED}(r)$ be a predicate on Rels that holds precisely when its argument r is disjunctively well-founded. Recall that Δ_{St} is the identity relation on states. The interpretation of commands is given below:

$$\begin{aligned} \llbracket c \rrbracket &\in \text{Rels}^\top \\ \llbracket a \rrbracket &\stackrel{\text{def}}{=} \llbracket a \rrbracket \\ \llbracket c_0; c_1 \rrbracket &\stackrel{\text{def}}{=} \text{rseq}(\llbracket c_0 \rrbracket, \llbracket c_1 \rrbracket) \\ \llbracket c_0 \parallel c_1 \rrbracket &\stackrel{\text{def}}{=} \llbracket c_0 \rrbracket \sqcup \llbracket c_1 \rrbracket \\ \llbracket \text{while } b \text{ c} \rrbracket &\stackrel{\text{def}}{=} \text{let } r = \text{lfix } \lambda r. \text{rseq}(\llbracket \text{assume}(b); c \rrbracket, \Delta_{\text{St}} \sqcup r) \\ &\quad \text{in if } (r \neq \top \wedge \text{DISJWELLFOUNDED}(r)) \\ &\quad \quad \text{then rseq}(\Delta_{\text{St}} \sqcup r, \llbracket \text{assume}(\neg b) \rrbracket) \\ &\quad \quad \text{else } \top \end{aligned}$$

Here $\llbracket a \rrbracket$ is the standard relational meaning of atomic command a .

The soundness of the relational semantics relies on the result by Podelski and Rybalchenko, applied to traces directly. We recall the result below:

Lemma 3 (Podelski and Rybalchenko). *Let r be a disjunctively well-founded relation. For every trace τ , if $\tau_n[r]\tau_m$ for all $n < m < |\tau|$ (where $|\tau|$ is ∞ in case that τ is infinite), the trace τ is finite.*

Proof. The proof of the lemma follows from the known proof of Podelski and Rybalchenko. We put the proof here, in order to make this appendix self-contained. Consider finitely many well-founded relations r_0, \dots, r_k . Pick r such that $r \subseteq r_0 \cup \dots \cup r_k$. For the sake of contradiction, suppose that there is an *infinite* trace τ satisfying

$$\forall n, m. n < m \implies \tau_n[r]\tau_m.$$

Define a function f that maps each pair (n, m) of indices with $n < m$ to an integer i such that $\tau_n[r_i]\tau_m$. Such a function should be well-defined, because $\tau_n[r]\tau_m$ and $r \subseteq$

$r_0 \cup \dots \cup r_k$. Since there are only finitely many r_i 's, by infinite Ramsey's theorem, there exist infinite subtrace τ' of τ and r_l such that

$$\forall n, m. (n < m) \implies \tau'_n[r_l]\tau'_m.$$

Note that for all indices n , states τ'_n and τ'_{n+1} are related by r_l . Thus, r_l cannot be well-founded, which contradicts our assumption on r_l . \square

Proposition 2 (Soundness of Relational Semantics). *For all commands c ,*

$$\llbracket c \rrbracket \subseteq \gamma(\langle c \rangle).$$

Proof. Define a relation $\mathcal{L} \subseteq \mathcal{P}(\mathcal{T}) \times \text{Rels}^\top$ by

$$T[\mathcal{L}]r \iff T \subseteq \gamma(r).$$

With this relation, we can rewrite the main claim of this proposition by

$$\forall c. \llbracket c \rrbracket [\mathcal{L}] \langle c \rangle.$$

We will show this in the proof.

Our proof relies on the four important properties of \mathcal{L} . The first is that \mathcal{L} is closed, downward for the left parameter and upward for the right parameter:

$$(T' \subseteq T \wedge T[\mathcal{L}]r \wedge r \sqsubseteq r') \implies T'[\mathcal{L}]r'.$$

The other three are the preservation of \mathcal{L} by various operators. Relation \mathcal{L} is preserved by the interpretations of atomic commands in the trace and relational semantics, the sequential composition operators for $\mathcal{P}(\mathcal{T})$ and Rels^\top , and the join operators for $\mathcal{P}(\mathcal{T})$ and Rels^\top . That is,

1. $\llbracket a \rrbracket [\mathcal{L}] \langle a \rangle$ for all atomic commands a ;
2. for all $T, T' \in \mathcal{P}(\mathcal{T})$ and $r, r' \in \text{Rels}^\top$,

$$T[\mathcal{L}]r \wedge T'[\mathcal{L}]r' \implies \text{seq}(T, T')[\mathcal{L}]r \text{seq}(r, r');$$

3. for all families $\{T_i\}_{i \in I}$ and $\{r_i\}_{i \in I}$ with the same index set I ,

$$(\forall i \in I. T_i[\mathcal{L}]r_i) \implies \left(\bigcup_{i \in I} T_i \right) [\mathcal{L}] \left(\bigcap_{i \in I} r_i \right).$$

Now, we prove the main claim of the proposition, by induction on the structure of c . All cases except the `while` loops follow from the four properties of \mathcal{L} and the induction hypothesis. For instance, consider the case that c is a nondeterministic choice $c_0 \sqcup c_1$. By the induction hypothesis, $\llbracket c_0 \rrbracket [\mathcal{L}] \langle c_0 \rangle$ and $\llbracket c_1 \rrbracket [\mathcal{L}] \langle c_1 \rangle$. Since the join operators preserve \mathcal{L} , these two relationships imply that

$$(\llbracket c_0 \rrbracket \cup \llbracket c_1 \rrbracket) [\mathcal{L}] (\langle c_0 \rangle \sqcup \langle c_1 \rangle).$$

This gives the proposition, because the left and right hand sides of \mathcal{L} are, respectively, $\llbracket c_0 \rrbracket \sqcup \llbracket c_1 \rrbracket$ and $\langle c_0 \rangle \sqcup \langle c_1 \rangle$.

Finally, consider the remaining case that c is `while` b c . Let

$$\begin{aligned} T_0 &= \llbracket \text{assume}(b); c \rrbracket, & F &= \lambda T. T_s \cup \text{seq}(T_0, T), \\ r_0 &= \langle \text{assume}(b); c \rangle, & G &= \lambda r. \Delta_{\text{St}} \sqcup \text{rseq}(r_0, r). \end{aligned}$$

where $T_s = \{s \mid s \in \text{St}\}$ and Δ_{St} is the identity relation on St . Define G' to be $\lambda r. \text{rseq}(r_0, \Delta_{\text{St}} \sqcup r)$ and r_i to be $\text{lfix } G'$. We will prove that

1. $\text{lfix } F \llbracket \mathcal{L} \rrbracket \text{lfix } G$;
2. $(\text{lfix } G) \sqsubseteq (\Delta_{\text{St}} \sqcup r_i)$;
3. $\text{repeat}(T_0) = \emptyset$ if r_i is disjunctively well-founded.

The proposition follows from these three. To see this, first consider the case that r_i is \top or it is not disjunctively well-founded. In this case, $\langle c \rangle$ is \top , so $T_0 \llbracket \mathcal{L} \rrbracket \langle c \rangle$ for all trace sets T_0 . The proposition follows from this. Next consider the other case that r_i is a disjunctively well-founded relation on states. By the third property above, $\text{repeat}(T_0) \llbracket \mathcal{L} \rrbracket \emptyset$. Since \mathcal{L} is upward closed on the right and it preserves the join operators,

$$(\text{lfix } F \cup \text{repeat}(T_0)) \llbracket \mathcal{L} \rrbracket (\Delta_{\text{St}} \sqcup r_i \sqcup \emptyset).$$

This implies the required

$$\text{seq}\left(\text{lfix } F \cup \text{repeat}(T_0), \llbracket \text{assume}(-b) \rrbracket\right) \llbracket \mathcal{L} \rrbracket \text{rseq}\left(\Delta_{\text{St}} \sqcup \text{lfix } G', \langle \text{assume}(-b) \rangle\right),$$

because the two interpretations of $\text{assume}(-b)$ are related by \mathcal{L} and the sequential composition operators preserve \mathcal{L} .

The first about the least fixpoints of F and G is a standard result. It holds because (1) F and G map \mathcal{L} -related values to \mathcal{L} -related values; (2) the empty trace set \emptyset and the empty relation \emptyset are related by \mathcal{L} ; (3) both finitary and infinitary join operators preserve \mathcal{L} ; and (4) F and G are continuous. The second holds because $\Delta_{\text{St}} \sqcup r_i$ is a fixpoint of G :

$$\begin{aligned} G(\Delta_{\text{St}} \sqcup r_i) &= \Delta_{\text{St}} \sqcup \text{seq}(r_0, \Delta_{\text{St}} \sqcup r_i) && (\because \text{Def. of } G) \\ &= \Delta_{\text{St}} \sqcup \text{seq}(r_0, \Delta_{\text{St}} \sqcup \text{lfix } G') && (\because \text{Def. of } r_i) \\ &= \Delta_{\text{St}} \sqcup \text{lfix } G' && (\because G' = \lambda r. \text{seq}(r_0, \Delta_{\text{St}} \sqcup r)) \\ &= \Delta_{\text{St}} \sqcup r_i && (\because \text{Def. of } r_i). \end{aligned}$$

For the third, assume that r_i is disjunctively well-founded. For the sake of contradiction, suppose that $\text{repeat}(T_0)$ is not empty. Pick a trace τ in $\text{repeat}(T_0)$. By the definition of repeat , there are infinitely many finite traces $s_0\tau_0s_1, s_1\tau_1s_2, \dots$ such that

$$\tau = s_0\tau_1s_1\tau_2s_2\dots \wedge (\forall n. |s_n\tau_n s_{n+1}| \geq 2 \wedge s_n\tau_n s_{n+1} \in T_0).$$

We will prove that the subtrace $s_0s_1\dots$ of τ satisfy

$$\forall n, m. n < m \implies s_n[r_i]s_m.$$

This gives the desired contradiction; it implies that τ is finite because r_i is disjunctively well-founded (Lemma 3), but τ belongs to $\text{repeat}(T_0)$ that contains only infinite traces. Let H be $\lambda T. \text{seq}(\llbracket \text{assume}(b); c \rrbracket, T)$. Recall that $T_s \llbracket \mathcal{L} \rrbracket \Delta_{\text{St}}$. Notice that functions H

and G' maps \mathcal{L} -related values to \mathcal{L} -related ones, because of the induction hypothesis and the preservation and closedness properties of \mathcal{L} . Thus, we have that

$$\forall n. H^n(T_s) [\mathcal{L}] G'^n(\Delta_{St}).$$

Also note that for all n, m with $n < m$, finite trace

$$s_n \tau_{n+1} \dots \tau_m s_m$$

is in $H^{m-n}(T_s)$. From these two observations and the definition of \mathcal{L} , it follows that

$$s_n \tau_{n+1} s_{n+1} \tau_{n+1} \dots \tau_m s_m \in \gamma(G'^{m-n}(\Delta_{St})).$$

The right hand side of this inequality is the same as $\gamma(G'^{m-n}(\emptyset))$, so that it is a subset of $\gamma(\text{fix } G')$. Therefore,

$$s_n \tau_{n+1} s_{n+1} \tau_{n+2} \dots \tau_m s_m \in \gamma(r_i),$$

which means $s_n[r_i]s_m$, as desired. \square

A.2 Overapproximation result

Consider a generic analysis. Let \mathcal{A} be the abstract domain of the analysis. For each $A \in \mathcal{A}$, define $\gamma_r(A) \in \text{Rels}^\top$ by

$$\gamma_r(A) \stackrel{\text{def}}{=} \begin{cases} \top & \text{if } A = \top \\ \bigcup \{ \gamma_r(d) \mid d \in A \} & \text{otherwise.} \end{cases}$$

Proposition 3. *The generic analysis overapproximates the relational semantics. That is, for all commands c and abstract values $A \in \mathcal{A}$,*

$$\text{rseq}(\gamma_r(A), \llbracket c \rrbracket) \sqsubseteq \gamma_r(\llbracket c \rrbracket^\#(A)).$$

Proof. We prove the proposition by induction on the structure of c . Since $\llbracket c \rrbracket^\#$ preserves \top , when A is \top , the inequality of the proposition holds. Thus, in this proof, we focus on non- \top abstract elements. Pick an abstract element A in $\mathcal{A} - \{\top\}$. Consider the case that c is an atomic command a . Then,

$$\begin{aligned} \text{rseq}(\gamma_r(A), \llbracket a \rrbracket) &= \gamma_r(A); \llbracket a \rrbracket && (\because \text{Def. of rseq}) \\ &= (\bigcup \{ \gamma_r(d) \mid d \in A \}); \llbracket a \rrbracket && (\because \text{Def. of } \gamma_r(A)) \\ &= \bigcup \{ \gamma_r(d); \llbracket a \rrbracket \mid d \in A \} && (\because -; r \text{ distributes over } \cup) \\ &\sqsubseteq \bigsqcup \{ \gamma_r(d') \mid d \in A \wedge d' \in \text{trans}(a)(d) \} && (\because \text{Condition on trans}) \\ &= \gamma_r(\text{trans}(a)^\dagger A) \\ &= \gamma_r(\llbracket a \rrbracket^\# A) && (\because \text{Def. of the Analyzer}). \end{aligned}$$

The cases that c is a sequential composition or a nondeterministic choice follow easily from the induction hypothesis, the associativity of rseq , and the preservation of \sqsubseteq by rseq . In the below, we prove both cases:

$$\begin{aligned} \text{rseq}(\gamma_r(A), \llbracket c_0; c_1 \rrbracket) &\sqsubseteq \text{rseq}(\gamma_r(A), \text{rseq}(\llbracket c_0 \rrbracket, \llbracket c_1 \rrbracket)) \\ &\sqsubseteq \text{rseq}(\text{rseq}(\gamma_r(A), \llbracket c_0 \rrbracket), \llbracket c_1 \rrbracket) && (\because \text{Associativity of rseq}) \\ &\sqsubseteq \text{rseq}(\gamma_r(\llbracket c_0 \rrbracket^\#(A)), \llbracket c_1 \rrbracket) && (\because \text{Ind. Hypo.}) \\ &\sqsubseteq \gamma_r(\llbracket c_1 \rrbracket^\#(\llbracket c_0 \rrbracket^\#(A))) && (\because \text{Ind. Hypo.}). \end{aligned}$$

$$\begin{aligned}
\text{rseq}(\gamma_r(A), \llbracket c_0 \sqcup c_1 \rrbracket) &\sqsubseteq \text{rseq}(\gamma_r(A), \llbracket c_0 \rrbracket \sqcup \llbracket c_1 \rrbracket) \\
&= \text{rseq}(\gamma_r(A), \llbracket c_0 \rrbracket) \sqcup \text{rseq}(\gamma_r(A), \llbracket c_1 \rrbracket) \quad (\because \text{rseq preserves } \sqcup) \\
&\sqsubseteq \gamma_r(\llbracket c_0 \rrbracket^\#(A)) \sqcup \gamma_r(\llbracket c_1 \rrbracket^\#(A)) \quad (\because \text{Ind. Hypo.}) \\
&\sqsubseteq \gamma_r(\llbracket c_0 \rrbracket^\#(A) \sqcup \llbracket c_1 \rrbracket^\#(A)) \quad (\because \text{Monotonicity of } \gamma_r) \\
&= \gamma_r(\llbracket c_0 \sqcup c_1 \rrbracket^\#(A)).
\end{aligned}$$

Now, it remains to prove the inductive step for the loop case, i.e., $c' = \text{while } b \text{ c}$.
Let

$$\begin{aligned}
G &= \lambda r. \text{rseq}(\llbracket \text{assume}(b); c \rrbracket, \Delta_{\text{St}} \sqcup r), & r_i &= \text{lfix } G, \\
H &= \lambda A. (\text{RFS}^\dagger \circ \llbracket \text{assume}(b); c \rrbracket^\#)(\{d_{\text{id}}\} \sqcup A), & A_i &= \text{fix } H
\end{aligned}$$

where fix is the fixpoint operator of the analysis. First, we prove that

$$\Delta_{\text{St}} \sqsubseteq \gamma_r(\{d_{\text{id}}\}) \quad \wedge \quad r_i \sqsubseteq \gamma_r(A_i).$$

The first conjunct follows from the soundness condition on d_{id} . For the second conjunct, we will show that the least fixpoint r_i of G is also the least fixpoint of $K = \lambda r. \text{rseq}(\Delta_{\text{St}} \sqcup r, r_0)$ for $r_0 = \llbracket \text{assume}(b); c \rrbracket$, and that $\gamma_r(A_i)$ is a pre-fixpoint of K . Since K is a monotone function on a complete lattice, this implies that the least fixpoint r_i of K is less than or equal to $\gamma_r(A_i)$. To show the coincidence between least fixpoints of G and K , we note that both G and K are continuous, so their least fixpoints can be computed by the limit of two countable sequences $\sqcup_{n \geq 0} G^n(\{\})$ and $\sqcup_{n \geq 0} K^n(\{\})$. Let L be $\lambda r. \text{rseq}(r_0, r)$. Then, for all $k \geq 0$, we have that $\text{rseq}(r_0, L^k(\Delta_{\text{St}})) = L^{k+1}(\Delta_{\text{St}})$ and also that $\text{rseq}(L^k(\Delta_{\text{St}}), r_0) = L^{k+1}(\Delta_{\text{St}})$. (The second equality can be proved by induction on k .) Using induction on n , we prove that for all $n \geq 0$,

$$G^n(\{\}) = K^n(\{\}) = \sqcup_{1 \leq k \leq n} L^k(\Delta_{\text{St}}).$$

The base case is $\{\} = \{\}$, so it holds. The inductive case holds as well, because

$$\begin{aligned}
G^{n+1}(\{\}) &= \text{rseq}(r_0, \Delta_{\text{St}} \sqcup G^n(\{\})) \\
&= \text{rseq}(r_0, \Delta_{\text{St}} \sqcup (\sqcup_{1 \leq k \leq n} L^k(\Delta_{\text{St}}))) \\
&= \text{rseq}(r_0, \Delta_{\text{St}}) \sqcup (\sqcup_{1 \leq k \leq n} \text{rseq}(r_0, L^k(\Delta_{\text{St}}))) \\
&= \text{rseq}(r_0, \Delta_{\text{St}}) \sqcup (\sqcup_{1 \leq k \leq n} L^{k+1}(\Delta_{\text{St}})) \\
&= L(\Delta_{\text{St}}) \sqcup (\sqcup_{1 \leq k \leq n} L^{k+1}(\Delta_{\text{St}})) \\
&= (\sqcup_{1 \leq k \leq n+1} L^k(\Delta_{\text{St}})),
\end{aligned}$$

and

$$\begin{aligned}
K^{n+1}(\{\}) &= \text{rseq}(\Delta_{\text{St}} \sqcup K^n(\{\}), r_0) \\
&= \text{rseq}(\Delta_{\text{St}} \sqcup (\sqcup_{1 \leq k \leq n} L^k(\Delta_{\text{St}})), r_0) \\
&= \text{rseq}(\Delta_{\text{St}}, r_0) \sqcup (\sqcup_{1 \leq k \leq n} \text{rseq}(L^k(\Delta_{\text{St}}), r_0)) \\
&= \text{rseq}(\Delta_{\text{St}}, r_0) \sqcup (\sqcup_{1 \leq k \leq n} L^{k+1}(\Delta_{\text{St}})) \\
&= \text{rseq}(r_0, \Delta_{\text{St}}) \sqcup (\sqcup_{1 \leq k \leq n} L^{k+1}(\Delta_{\text{St}})) \\
&= L(\Delta_{\text{St}}) \sqcup (\sqcup_{1 \leq k \leq n} L^{k+1}(\Delta_{\text{St}})) \\
&= (\sqcup_{1 \leq k \leq n+1} L^k(\Delta_{\text{St}})).
\end{aligned}$$

We have just shown that the least fixpoints of F and G are the limits of the same countable sequence. Thus, they must be the same. We move on to the proof that $\gamma_r(A_i)$

is a pre-fixpoint of K .

$$\begin{aligned}
K(\gamma_r(A_i)) &= \text{rseq}(\Delta_{\text{St}} \sqcup \gamma_r(A_i), (\text{assume}(b); c)) && (\because \text{Def. of } K) \\
&\sqsubseteq \text{rseq}(\gamma_r(\{d_{\text{id}}\}) \sqcup \gamma_r(A_i), (\text{assume}(b); c)) && (\because \Delta_{\text{St}} \sqsubseteq \gamma_r(\{d_{\text{id}}\})) \\
&\sqsubseteq \text{rseq}(\gamma_r(\{d_{\text{id}}\} \sqcup A_i), (\text{assume}(b); c)) && (\because \gamma_r \text{ is monotone}) \\
&\sqsubseteq \gamma_r(\llbracket \text{assume}(b); c \rrbracket^\#(\{d_{\text{id}}\} \sqcup A_i)) && (\because \text{Ind. Hypo}) \\
&\sqsubseteq \gamma_r((\text{RFS}^\dagger \circ \llbracket \text{assume}(b); c \rrbracket^\#)(\{d_{\text{id}}\} \sqcup A_i)) && (\because \text{Soundness of RFS}) \\
&= \gamma_r(H(A_i)) && (\because \text{Def. of } H) \\
&\sqsubseteq \gamma_r(A_i) && (\because A_i = \text{fix } H, \text{ and Condition on fix}).
\end{aligned}$$

Next, we notice that if A_i is not \top , r_i has to be a disjunctively well-founded relation on states. To see this, recall that $\text{fix } H$ is in the image of RFS^\dagger , and suppose that $A_i (= \text{fix } H)$ is not \top . Then, A_i should be a finite set of d 's, each of which denotes a well-founded relation via γ_r . Thus, $\gamma_r(A_i)$ is a finite union of well-founded relations. From this and $r_i \sqsubseteq \gamma_r(A_i)$, it follows that r_i is disjunctively well-founded.

Finally, we prove the loop case. If A_i is \top , so is $\llbracket \text{while } b c \rrbracket^\# A$. Thus, the proposition holds. Assume that A_i is not \top . Then, by what we have just shown, r_i is a disjunctively well-founded relation. With this, we prove the loop case:

$$\begin{aligned}
&\text{rseq}(\gamma_r(A), (\text{while } b c)) \\
&= \text{rseq}(\gamma_r(A), \text{rseq}(\Delta_{\text{St}} \cup r_i, (\text{assume}(\neg b)))) && (\because r_i \text{ is disj. well-founded}) \\
&\sqsubseteq \text{rseq}(\gamma_r(A), \text{rseq}(\gamma_r(\{d_{\text{id}}\}) \sqcup \gamma_r(A_i), (\text{assume}(\neg b)))) && (\because \text{Mono. of rseq}) \\
&\sqsubseteq \text{rseq}(\gamma_r(A), \text{rseq}(\gamma_r(\{d_{\text{id}}\} \sqcup A_i), (\text{assume}(\neg b)))) && (\because \mathcal{A} \text{ is disjunctive}) \\
&\sqsubseteq \text{rseq}(\gamma_r(A), \gamma_r(\llbracket \text{assume}(\neg b) \rrbracket^\#(\{d_{\text{id}}\} \sqcup A_i))) && (\because \text{Ind. Hypo.}) \\
&\sqsubseteq \gamma_r(\text{comp}^\dagger(A, \text{assume}(\neg b)(\{d_{\text{id}}\} \sqcup A_i))) && (\because \text{Soundness of comp}) \\
&= \gamma_r(\llbracket \text{while } b c \rrbracket^\#(A)). && (\because \text{Def. of the analysis}).
\end{aligned}$$

□

A.3 Proof of Theorem 1

We can now prove Theorem 1 easily. First, note that the generic analyzer overapproximates the concrete trace semantics. For all commands c ,

$$\begin{aligned}
\llbracket c \rrbracket &\sqsubseteq \gamma(\llbracket c \rrbracket) && (\because \text{Prop. 2}) \\
&= \gamma(\text{rseq}(\Delta_{\text{St}}, \llbracket c \rrbracket)) && (\because \text{Def. of rseq}) \\
&\sqsubseteq \gamma(\gamma_r(\llbracket c \rrbracket^\# \{d_{\text{id}}\})) && (\because \text{Prop. 3}) \\
&= \gamma(\llbracket c \rrbracket^\# \{d_{\text{id}}\}).
\end{aligned}$$

We use two kinds of γ above; γ in the last line is a map from \mathcal{A} to the sets of traces, and γ in all the other places is a map from Rels^\top to the set of traces.

Next, observe that $\gamma(A)$ can contain infinite traces only when A is \top . By combining these two observations, we can conclude that if $\text{ANALYSIS}(c)$ returns “Terminates”, $\llbracket c \rrbracket$ does not contain any infinite traces, that is, c terminates.