

Termination Analysis by Learning Terminating Programs [★]

Matthias Heizmann, Jochen Hoenicke, and Andreas Podelski

University of Freiburg, Germany

Abstract. We present a novel approach to termination analysis. In a first step, the analysis uses a program as a black-box which exhibits only a finite set of sample traces. Each sample trace is infinite but can be represented by a finite lasso. The analysis can “learn” a program from a termination proof for the lasso, a program that is terminating by construction. In a second step, the analysis checks that the set of sample traces is representative in a sense that we can make formal. An experimental evaluation indicates that the approach is a potentially useful addition to the portfolio of existing approaches to termination analysis.

1 Introduction

Termination analysis is an active research topic, and a wide range of methods and tools exist [12,14,23,27,29,36,39]. Each method provides its own twist to address the same issue: in the presence of loops with branching or nesting, the termination argument has to account for all possible interleavings between the different paths through the loop.

If the program is *lasso-shaped* (a stem followed by a single loop without branching), the control flow is trivial; there is only one path. Consequently, the termination argument can be very simple. Many procedures are specialized to lasso-shaped programs and derive a simple termination argument rather efficiently [4,5,7,11,24,31,33]. The relevance of lasso-shaped programs stems from their use as the representation of an infinite trace through the control flow graph of a program with arbitrary nesting.

We present a new method that analyzes termination of a general program $\mathcal{P}$ but has to find termination arguments only for lasso-shaped programs. In our method we see the program $\mathcal{P}$ as a blackbox from which we can obtain sample traces. We transform a sample trace π_i into a lasso-shaped program and use existing methods to compute a termination argument for this lasso-shaped program. Afterwards we construct a “larger” program $\mathcal{P}_i$ (which may have branching and nested loops) for which the same termination argument is applicable. We call this construction *learning*, because we learned the terminating program $\mathcal{P}_i$ from a sample trace π_i . Our algorithm continues iteratively until we learned a set of

[★] This work is supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center “Automatic Verification and Analysis of Complex Systems” (SFB/TR14 AVACS)


programs $\mathcal{P}_1, \dots, \mathcal{P}_n$ that forms a decomposition of the original program $\mathcal{P}$. This decomposition can be seen as a program of the form $\text{choose}(\mathcal{P}_1, \dots, \mathcal{P}_n)$, i.e., a nondeterministic choice of programs $\mathcal{P}_1, \dots, \mathcal{P}_n$, that is semantically equivalent to the original program $\mathcal{P}$.

Our technical contribution is this method, which does not only extend the existing portfolio of termination analyses but also provides a new functionality: the decomposition of a program $\mathcal{P}$ into modules $\mathcal{P}_1, \dots, \mathcal{P}_n$. This decomposition is not guided by the syntax of the program, this decomposition exploits a novel notion of modularity where a module is defined by a certain termination argument. This novel notion of modularity is the conceptual contribution of our paper.

```

program sort(int i)
 $\ell_1$ : while (i>0)
 $\ell_2$ :   int j:=1
 $\ell_3$ :   while(j<i)
        //   if (a[j]>a[i])
        //       swap(a[j],a[i])
 $\ell_4$ :     j++
 $\ell_5$ :     i--

```



Let us explain our algorithm informally using the program $\mathcal{P}^{\text{sort}}$ depicted above which is an implementation of bubblesort. Termination of $\mathcal{P}^{\text{sort}}$ can be shown, e.g., by using the quadratic ranking function $f(i, j) = i^2 - j$, or the lexicographic ranking function $f(i, j) = (i, i - j)$. Intuitively, neither of the two ranking functions is a simple termination argument.

Now, let us pick some ω -trace from $\mathcal{P}^{\text{sort}}$. We take the trace that first enters the outer while loop and then takes the inner while loop infinitely often. We denote this trace using the ω -regular expression $\text{ OUTER.INNER}^\omega$. We see that this trace is terminating. Its termination can be shown using the linear ranking function $f(i, j) = i - j$. Moreover, we see that this ranking function is not only applicable to this trace, this ranking function is applicable to all traces that eventually always take the inner loop.

$$(\text{ INNER} + \text{ OUTER})^*.\text{ INNER}^\omega \quad (1)$$

Now, let us pick another ω -trace from $\mathcal{P}^{\text{sort}}$. This time we take the trace that always takes the outer while loop. We see that this trace is terminating. Its termination can be shown using the linear ranking function $f(i, j) = i$. Moreover, we see that this ranking function is not only applicable to this trace, this ranking function is applicable to all traces that take the outer while loop infinitely often.

$$(\text{ INNER}^*.\text{ OUTER})^\omega \quad (2)$$

Finally, we consider the set of all ω -trace of the program $\mathcal{P}^{\text{sort}}$

$$(\text{OUTER} + \text{INNER})^\omega,$$

check that each trace has the form (1) or has the form (2), and conclude that $\mathcal{P}^{\text{sort}}$ is terminating.

If we are to automate the reasoning from the example above, a number of questions arise.

(A) *How does one effectively represent a set of traces that share a common reason for termination, like the sets (1) and (2) above?* The answer is given in Section 2 where we define a *module*, which is a program whose traces adhere to a certain fairness constraint.

(B) *What is a termination argument whose applicability to a whole set of traces can be checked effectively?* The answer is given in Section 3 where we present a Floyd-Hoare style annotation for termination proofs.

(C) *How can we learn a set of terminating traces (represented as a program with a fairness constraint) from a single terminating sample trace?* The answer is given in Section 4 where we construct a terminating module from a given termination proof.

(D) *How can we check that a set of modules $\mathcal{P}_1, \dots, \mathcal{P}_n$ covers the behavior of the original program $\mathcal{P}$ and can we always decompose $\mathcal{P}$ into a set of modules $\mathcal{P}_1, \dots, \mathcal{P}_n$?* One facet of the question is the theoretical completeness, which is answered in Section 5. The other facet is the practical feasibility, which is analyzed via an experimental evaluation in Section 6.

2 Fair module

Preliminaries. The key concept in our formal exposition is the notion of an ω -trace, which is an infinite sequence of program statements $\pi = s_1 s_2 \dots$. We assume that the statements are taken from a given finite set of program statements Σ . If we consider Σ as an alphabet and each statement as a letter, then an ω -trace is an infinite word over this alphabet. In order to stress the usage of statements as letters of an alphabet, we sometimes frame each statement/letter. For example, we can write the alphabet of our running example $\mathcal{P}^{\text{sort}}$ as $\Sigma_{\text{sort}} = \{ \text{ } i > 0, \text{ } j := 1, \text{ } j < i, \text{ } j++, \text{ } j \geq i, \text{ } i-- \}$ and $\pi = \text{ } j < i \text{ } j := 1 \text{ } (\text{ } j := 1 \text{ } j++ \text{ } j := 1)^\omega$ is an ω -trace.

The definition of an ω -trace as an arbitrary (infinite) sequence means that the notion is independent of the programming language semantics, which we even have not introduced yet. We will do so now. A *valuation* ν is a function that maps the program variables $\vec{v}$ to values. We use the term *valuation* instead of *state* to stress that this is independent from the program counter (and independent from control flow). We call a set of valuations a *predicate* and use the letter I to denote predicates. The letter I is used reminiscent to *invariant*, because we will use predicates to represent invariants at locations. We assume that each statement s comes with a binary relation over the set of valuations (the set of its precondition/postcondition pairs). We say that the Hoare triple $\{I\} s \{I'\}$ is valid, if the binary relation for s holds between precondition I and postcondition

I' . We use the interleaved sequences of valuations and statements $\nu_0 \xrightarrow{\mathcal{S}_1} \dots \xrightarrow{\mathcal{S}_n} \nu_n$ as a shorthand to denote that each pair of valuations (ν_i, ν_{i+1}) is contained in the transition relation of the statement $\mathcal{S}_{i+1}$.

An ω -trace may not correspond to any possible execution for one out of two reasons. First, there may be a finite prefix that does not have any possible execution, like e.g., the prefix $\mathbf{x} < 0 \ \mathbf{x} := 1 \ \mathbf{x} < 0$ of the ω -trace $(\mathbf{x} < 0 \ \mathbf{x} := 1)^\omega$. Secondly, there may be no starting valuation ν_0 for any infinite execution, although every finite prefix is executable which holds e.g., for the ω -trace $(\mathbf{x} > 0 \ \mathbf{x}++)^\omega$. In both cases we call such an ω -trace *terminating*.

The notion of an ω -trace is also independent of a program (a trace may not correspond to a path in the program's control flow graph). We introduce a program as a control flow graph whose edges are labeled with statements. Formally, a *program* is a graph $\mathcal{P} = \langle \text{Loc}, \delta, \ell_{\text{init}} \rangle$ with a finite set Loc of nodes called *locations*, a set δ of edges labeled with statements, i.e., $\delta \subseteq \text{Loc} \times \Sigma \times \text{Loc}$ and an initial node called the initial location ℓ_{init} . We call the program $\mathcal{P}$ *terminating* if each of its ω -traces is terminating.

Module: program with fairness constraint. In our method we will decompose a program into *modules* such that each module represents traces that share a common reason for termination. We now formalize our notion of a module.

Definition 1 (module). A *module* is a program together with a fairness constraint given by a distinguished location ℓ_{fin} , i.e.,

$$\mathcal{P} = \langle \text{Loc}, \delta, \ell_{\text{init}}, \ell_{\text{fin}} \rangle$$

where the set of location can be partitioned into two disjoint sets, Loc_U , and Loc_V , such that

- the initial location is contained in Loc_U ,
- the final location is contained in Loc_V , and
- no location in Loc_V has a successor in Loc_U , i.e.,


$$(\ell, \mathcal{S}, \ell') \in \delta \quad \text{implies} \quad \ell \in \text{Loc}_U \quad \text{or} \quad \ell' \in \text{Loc}_V$$

A *fair ω -trace* of a module $\mathcal{P}$ is an ω -trace that labels a fair path in the graph of $\mathcal{P}$, which is a path that visits the distinguished location ℓ_{fin} infinitely often. We call the module $\mathcal{P}$ *terminating* if each of its fair ω -traces is terminating.

A *non-fair ω -trace* of a terminating module (i.e., an ω -trace that labels a path in its control flow graph without satisfying the fairness constraint) can be non-terminating.

For the reader who is familiar with the concept of Büchi automata, a module is reminiscent of a Büchi automaton with exactly one final state. A Büchi automaton of this form recognizes an ω -regular language of the form $U.V^\omega$, where U and V are regular languages over the alphabet of statements $U, V \subseteq \Sigma^*$.

Example 1. Let us consider again our running example $\mathcal{P}^{\text{sort}}$. The sets that we gave informally by the ω -regular expressions (1) and (2) can be represented as modules. The module $\mathcal{P}_1^{\text{sort}}$ depicted on the left represents all traces that eventually only take the inner while loop. The module $\mathcal{P}_2^{\text{sort}}$ depicted on the right represents all traces that take the outer while loop infinitely often.



In this example, the decomposition of the program into modules is defined by the nestings structure of while loops. In Section 5 we present an algorithm that finds a decomposition automatically but does not rely on any information about the structure of the while loops in the program.

3 Certified Module

In this section we present a termination argument for modules that consists of two parts: a ranking function and an annotation of the module's locations.

First, we extend the usual notion of a ranking function to our definition of a module. The crux in the following definition lies in the fact that we do not require that the value of the ranking function has to decrease after a fixed number of steps. We only require that the value of the ranking function has to decrease every time the final location ℓ_{fin} is visited. As a consequence our ranking function is a termination argument that is applicable to each fair ω -trace, but does not have to take non-fair ω -traces into account.

Definition 2 (ranking function for a module). *Given a module $\mathcal{P}$, we call a function f from valuations into a well-ordered set $(\mathbb{W}, <)$ a **ranking function** for $\mathcal{P}$ if for each finite path*

$$\ell_0 \xrightarrow{\mathfrak{s}_1} \dots \xrightarrow{\mathfrak{s}_k} \ell_k \xrightarrow{\mathfrak{s}_{k+1}} \dots \xrightarrow{\mathfrak{s}_n} \ell_n$$

that starts in the initial location (i.e., $\ell_0 = \ell_{\text{init}}$) and visits the final location in the k -th step and in the n -th step (i.e., $\ell_k = \ell_n = \ell_{\text{fin}}$) and for each sequence of valuations $\nu_0, \dots, \nu_n$ such that the pair (ν_i, ν_{i+1}) is in the transition relation of the statement $\mathfrak{s}_i$, i.e.,

$$\nu_0 \xrightarrow{\mathfrak{s}_1} \dots \xrightarrow{\mathfrak{s}_k} \nu_k \xrightarrow{\mathfrak{s}_{k+1}} \dots \xrightarrow{\mathfrak{s}_n} \nu_n$$

the value of the ranking function decreases whenever ℓ_{fin} is visited, i.e.,

$$f(\nu_n) < f(\nu_k).$$

In all the following examples we take $\mathbb{Z}$ as domain of the program variables. Our well-ordered set $\mathbb{W}$ will be $(\mathbb{Z} \cup \{\infty\}, \prec)$. The ordering $\prec$ is the natural order restricted to pairs where the second operand is greater than or equal to zero (i.e., $a \prec b$ if and only if $a < b \wedge b \geq 0$).

Example 2. The function $f : \text{dom} \rightarrow \mathbb{Z} \cup \{\infty\}$ defined as $f(i, j) = i - j$ is a ranking function for the module $\mathcal{P}_1^{\text{sort}}$ depicted in Example 1.

Lemma 1. *If the module $\mathcal{P}$ has a ranking function f , then each fair trace of the module is terminating.*¹

How can we check that a function is a ranking function for a module? We next introduce a novel kind of annotation, called *rank certificate* that serves as a proof for this task. Informally, a *rank certificate* is a Floyd-Hoare annotation that ensures that the value of the ranking function has decreased whenever the final location ℓ_{fin} was visited. Therefore, we introduce an auxiliary variable **oldrnk** that represents the value of the ranking function at the previous visit of ℓ_{fin} . Initially, the auxiliary variable **oldrnk** has the value ∞ which is a value strictly greater than all other values from our well-ordered $\mathbb{W}$.

Definition 3 (certified module). *Given a module $\mathcal{P} = \langle \text{Loc}, \delta, \ell_{\text{init}}, \{\ell_{\text{fin}}\} \rangle$ and a function f from valuations into a well-ordered set $(\mathbb{W}, \prec)$, we call a mapping $\mathcal{I}$ from locations to predicates a **rank certificate** for the function f and the module $\mathcal{P}$ if the following properties hold.*

- The initial location ℓ_{init} is mapped to the predicate where the auxiliary variable **oldrnk** has the value ∞ , i. e.,

$$\mathcal{I}(\ell_{\text{init}}) \Leftrightarrow \text{oldrnk} = \infty.$$

- The accepting state is mapped to a predicate in which the value of the ranking function f over the program variables is smaller than the value of the variable **oldrnk**, i. e.,

$$\mathcal{I}(\ell_{\text{fin}}) \Rightarrow (f(\vec{v}) \prec \text{oldrnk}).$$

- The outgoing edges of non-accepting locations correspond to valid Hoare triples, i. e.,

$$\{ \mathcal{I}(\ell) \} \text{ } \mathcal{I} \{ \mathcal{I}(\ell') \} \text{ is valid for } (\ell, \mathcal{I}, \ell') \in \delta, \ell \neq \ell_{\text{fin}}$$


and outgoing edges of the final location correspond to valid Hoare triples if we insert an additional assignment statement that assigns the value of the ranking function to the auxiliary variable **oldrnk**, i. e.,

$$\{ \mathcal{I}(\ell) \} \text{ } \text{oldrnk} := f(\vec{v}) ; \mathcal{I} \{ \mathcal{I}(\ell') \} \text{ is valid for } (\ell_{\text{fin}}, \mathcal{I}, \ell') \in \delta$$

We call the triple $(\mathcal{P}, f, \mathcal{I})$ a **certified module**.

¹ An extended version of this paper that also contains the correctness proofs is available online.

Example 3. The figure on the right depicts a certified module $(\mathcal{P}_1^{\text{sort}}, f, \mathcal{I})$ where f is the ranking function $f(i, j) = i - j$ and $\mathcal{I}$ is the mapping of locations to predicates indicated by writing the predicate beneath the location.



Theorem 1 (soundness). *Each fair ω -trace of a certified module $(\mathcal{P}, f, \mathcal{I})$ is terminating.*

4 Learning a terminating program

In this section we present a method for the construction of a certified module $(\mathcal{P}, f, \mathcal{I})$. The crux of this method is that we do not construct a termination argument (a ranking function f together with a rank certificate $\mathcal{I}$) for the resulting module $\mathcal{P}$. Instead, we construct vice versa the resulting module $\mathcal{P}$ as the largest module for which a given termination argument (a ranking function f together with a rank certificate $\mathcal{I}$) is applicable. We obtain this termination proof from a single ω -trace. We call this method learning, because we learn a terminating program (given as a certified module) from a single sample trace.

The input to our method is a terminating ω -trace $\mathfrak{s}_1 \dots \mathfrak{s}_{k-1} (\mathfrak{s}_k \dots \mathfrak{s}_n)^\omega$ that is ultimately periodic. We call an ultimately periodic trace a *lasso*. We call the prefix $\mathfrak{s}_1 \dots \mathfrak{s}_{k-1}$ the *stem* of the lasso and we call the periodic part $\mathfrak{s}_k \dots \mathfrak{s}_n$ the *loop* of the lasso. For better legibility we use u (resp. v) to denote the stem (resp. loop) of the lasso. We construct a certified module $(\mathcal{P}, f, \mathcal{I})$ in the following three steps.


Step 1. Synthesize ranking function f

First, we construct a module $\mathcal{P}_{uv^\omega}$ that has only one single ω -trace, namely the lasso uv^ω . We call $\mathcal{P}_{uv^\omega}$ the *lasso module* of uv^ω and construct $\mathcal{P}_{uv^\omega} = \langle \text{Loc}, \delta, \ell_{\text{init}}, \{\ell_{\text{fin}}\} \rangle$ formally as the module that has one location for each statement (i.e., $\text{Loc} = \{\ell_0, \dots, \ell_{n-1}\}$), where ℓ_0 is the initial location, ℓ_k is the final location and the transition graph resembles the shape of a lasso, i.e., $\delta = \{(\ell_i, \mathfrak{s}_i, \ell_{i+1}) \mid i = 1, \dots, n-1\} \cup \{(\ell_{n-1}, \mathfrak{s}_n, \ell_k)\}$.

The lasso module $\mathcal{P}_{uv^\omega}$ can be seen as a program that consists of a single while loop. This allows us to use existing methods [4,5,7,11,24,31,33] to synthesize a ranking function for $\mathcal{P}_{uv^\omega}$.

Example 4. Given the ω -trace $i>0, j:=1, (j<i, j++)^\omega$, we construct the lasso module $\mathcal{P}_{uv^\omega}$ depicted

on the right and synthesize the ranking function $f(i, j) = i - j$ for this module.



Step 2. Compute rank certificate $\mathcal{I}$

Given the lasso module $\mathcal{P}_{uv^\omega}$ and the ranking function f , we now compute a rank certificate $\mathcal{I}$. Since $\mathcal{P}_{uv^\omega}$ has a “lasso shape” a mapping $\mathcal{I}$ from the locations of $\mathcal{P}_{uv^\omega}$ to predicates is a rank certificate if and only if the Hoare triples and the implication shown on the right are valid.

```

program rankDecrease()
    oldrnk := ∞
    l1 : s1
    ⋮
    lk-1 : sk-1
    lk : while (true)
        assert(f(v) < oldrnk)
        oldrnk := f(v)
        sk
    lk+1 : sk+1
    ⋮
    ln : sn

```

$$\begin{aligned}
 & \{ \text{true} \} \text{oldrnk} := \infty \{ \mathcal{I}(\ell_1) \} \\
 & \{ \mathcal{I}(\ell_i) \} s_i \{ \mathcal{I}(\ell_{i+1}) \} \quad \text{for } 1 \leq i < k \\
 & \mathcal{I}(\ell_k) \Rightarrow f(\vec{v}) < \text{oldrnk} \\
 & \{ \mathcal{I}(\ell_k) \} \text{oldrnk} := f(\vec{v}) s_k \{ \mathcal{I}(\ell_{k+1}) \} \\
 & \{ \mathcal{I}(\ell_i) \} s_{i+1} \{ \mathcal{I}(\ell_{i+1}) \} \quad \text{for } k < i < n \\
 & \{ \mathcal{I}(\ell_n) \} s_n \{ \mathcal{I}(\ell_k) \}
 \end{aligned}$$


The predicates $\mathcal{I}(\ell_i)$ for which these implications are valid, can be obtained by proving partial correctness of the program **rankDecrease**_{uvf} depicted on the left. The program **rankDecrease**_{uvf} first assigns the value ∞ (which is strictly larger than any other element in the well-ordered set $\mathbb{W}$) to the variable **oldrnk**. Afterwards the statements $s_1 \dots s_{k-1}$ are executed and the program **rankDecrease**_{uvf} enters a nonterminating **while** loop. We use an assert statement to state the correctness specification of the program **rankDecrease**_{uvf}.

The program is correct if at the beginning of the **while** loop the inequality $f(\vec{v}) < \text{oldrnk}$ holds. After this assert statement, the current value of the function f is assigned to the variable **oldrnk** and then the statements $s_k \dots s_n$ are executed.

A Floyd-Hoare annotation $\mathcal{I}(\ell_1), \dots, \mathcal{I}(\ell_n)$ that shows partial correctness of the program **rankDecrease**_{uvf} is also a rank certificate for our ranking function f and our lasso module $\mathcal{P}_{uv^\omega}$. This Floyd-Hoare annotation can be computed by static analysis [15].

Example 5. Continuing Example 4 we construct the program **rankDecrease**_{uvf} for $\mathcal{P}_{uv^\omega}$ and compute the rank certificate

depicted in the figure on the right. The rank certificate $\mathcal{I}$ is represented by the predicates denoted beneath the locations.



An alternative variant of Step 2. Some methods for the synthesis of a ranking function [7,24] also provide a *supporting invariant*. This is a predicate I such that

- I is invariant under executions of the loop $\mathcal{S}_1 \dots \mathcal{S}_{k-1}$,
- I is an overapproximation of the reachable valuations after executing the stem $\mathcal{S}_k \dots \mathcal{S}_n$,
- and each execution of the loop starting in a valuation contained I decreases the ranking function f .

If we have a supporting invariant I for the ranking function f , we do not have to construct and analyze the program $\text{rankDecrease}_{uvf}$. Alternatively, we can set the predicate $\mathcal{I}(\ell_k)$ to

$$I \wedge f(\vec{v}) < \text{oldrnk} \wedge \text{oldrnk} \geq 0$$

and obtain the remaining predicates $\mathcal{I}(\ell_0), \dots, \mathcal{I}(\ell_{k-1})$, and $\mathcal{I}(\ell_{k+1}), \dots, \mathcal{I}(\ell_n)$ as strongest postconditions by using an interpolating theorem prover.

Step 3. Construct module $\mathcal{P}$

We extend the lasso module $\mathcal{P}_{uv\omega}$ to a module $\mathcal{P}$ that also has the ranking function f and that also has the rank certificate $\mathcal{I}$. Therefore we modify $\mathcal{P}_{uv\omega}$ according to the following two rules.


Modification rule 1: merge locations If the predicates mapped to the locations ℓ_i and ℓ_j coincide (i.e., $\mathcal{I}(\ell_i) = \mathcal{I}(\ell_j)$) then we may merge both locations.

Modification rule 2: add transitions Let $\mathcal{S}$ be some program statement and let ℓ_i , and ℓ_j be locations. If $\ell_i \neq \ell_j$ and the Hoare triple $\{ \ell_i \} \mathcal{S} \{ \ell_j \}$ is valid, we may add the transition $(\ell_i, \mathcal{S}, \ell_j)$. If $\ell_i = \ell_j$ and the Hoare triple $\{ \ell_i \} \text{oldrnk} := f(\vec{v}); \mathcal{S} \{ \ell_j \}$ is valid, we may add the transition $(\ell_i, \mathcal{S}, \ell_j)$.

If we apply these modifications to a certified module we obtain again a certified module. Every strategy for applying these modifications gives rise to an algorithm that is an instance of our method.

Example 6. Continuing Example 4 we merge locations ℓ_1 and ℓ_2 . Afterwards we add for each program statement that occurs in $\mathcal{P}^{\text{sort}}$ a selfloop at ℓ_1 and

a transition between ℓ_1 and ℓ_3 . We obtain the certified module $\mathcal{P}_{\text{ext}}$ depicted on the right. The set of fair ω -traces of this module is given by the ω -regular expression $\Sigma^* \cdot (\text{j} < \text{i} \cdot \text{j}++)^\omega$. If we take the intersection of the program $\mathcal{P}^{\text{sort}}$ and the module $\mathcal{P}_{\text{ext}}$ we obtain the module $\mathcal{P}_1^{\text{sort}}$ from Example 1. In our algorithm (Section 5), we do not need to construct modules such as $\mathcal{P}_1^{\text{sort}}$ explicitly (we only use their implicit representation through $\mathcal{P}_{\text{ext}}$).



5 Overall algorithm

Until now, we have formalized (and automated) one part of our method, which is to construct a terminating module from a given sample trace. We still need to formalize (and automate) how to check that a set of modules covers all behaviours of the program. We will say that the program $\mathcal{P}$ has a decomposition into the modules $\mathcal{P}_1, \dots, \mathcal{P}_n$ if the set of ω -traces of the program $\mathcal{P}$ is the union of the set of fair ω -traces of the modules $\mathcal{P}_1, \dots, \mathcal{P}_n$.

We can automate the check that indeed all cases are covered by reducing it to the inclusion between Büchi automata. Both a program and a module are special cases of Büchi automata (where the set of states is the set of program locations and the set of final states contains all program locations respectively the final location ℓ_{fin} only). By definition, the ω -traces of the program $\mathcal{P}$ are exactly the infinite words accepted by the Büchi automaton $\mathcal{P}$ (and form the language $\mathcal{L}(\mathcal{P})$ recognized by $\mathcal{P}$), and the fair ω -traces of the module $\mathcal{P}_i$ are exactly the infinite words accepted by the Büchi automaton $\mathcal{P}_i$ (and form the language $\mathcal{L}(\mathcal{P}_i)$ recognized by $\mathcal{P}_i$), for $i = 1, \dots, n$. The inclusion

$$\mathcal{L}(\mathcal{P}) \subseteq \mathcal{L}(\mathcal{P}_1) \cup \dots \cup \mathcal{L}(\mathcal{P}_n)$$

can be checked by a model checker such as [26] or by a tool for manipulating Büchi automata such as [38].

We will use Büchi automata also in order to prove that decomposing a program into certified modules is in principle a complete method for termination analysis.

Theorem 2 (completeness). *If a program $\mathcal{P}$ is terminating then it can be decomposed into a finite set of certified modules, i.e., there are certified modules*

$$(\mathcal{P}_1, f_1, \mathcal{I}_1), \dots, (\mathcal{P}_n, f_n, \mathcal{I}_n)$$

such that the following equality holds.

$$\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{P}_1) \cup \dots \cup \mathcal{L}(\mathcal{P}_n)$$

Overall algorithm. Having reduced the check that a set of modules is a decomposition of a program, we are ready to present our algorithm for termination analysis, depicted below. The algorithm iteratively constructs certified modules $(\mathcal{P}_i, f_i, \mathcal{I}_i)$ until all ω -traces of the program are known to be terminating or we encounter an ω -trace for which we cannot find a termination argument.

At the beginning of each iteration (line 2) we check if there is an ω -trace of the program $\mathcal{P}$ that is not already a fair ω -trace of one of the modules $\mathcal{P}_1, \dots, \mathcal{P}_{n-1}$ (for which termination has already been proven). As mentioned above, we reduce this check to language inclusion of Büchi automata. Therefore we know that whenever there exists a counterexample to language inclusion there exists also a lasso-shaped counterexample. We take such a lasso-shaped ω -trace uv^ω and construct a program (called lasso module) whose only ω -trace is uv^ω (see

```

input : program  $\mathcal{P}$ 
output: certified modules  $(\mathcal{P}_1, f_1, \mathcal{I}_1), \dots, (\mathcal{P}_n, f_n, \mathcal{I}_n)$ 
1 for  $n = 0, 1, 2, \dots$  do
2   if  $\mathcal{L}(\mathcal{P}) \not\subseteq \mathcal{L}(\mathcal{P}_1) \cup \dots \cup \mathcal{L}(\mathcal{P}_{n-1})$  then
3     take  $\omega$ -trace  $u.v^\omega$  that is counterexample to inclusion;
4     construct lasso module  $\mathcal{P}_{uv^\omega}$ ;
5      $f_n := \text{synthesizeRankingFunction}(\mathcal{P}_{uv^\omega})$ ;
6     if  $f_n = \text{no ranking function found}$  then
7       return “unable to decide termination of  $\mathcal{P}$ ”
8     end
9      $\mathcal{I}_n := \text{computeRankCertificate}(f_n, \mathcal{P}_{uv^\omega})$ ;
10     $\mathcal{P}_n := \text{extendCertifiedModule}(\mathcal{P}_{uv^\omega}, f_n, \mathcal{I}_n)$ ;
11  else
12    return “ $\mathcal{P}$  is terminating,”
        “found decomposition  $(\mathcal{P}_1, f_1, \mathcal{I}_1), \dots, (\mathcal{P}_n, f_n, \mathcal{I}_n)$ ”
13  end
14 end

```

Algorithm 1: decomposition of a program $\mathcal{P}$ into certified modules

Step 1 in Section 4). Next, we analyze termination of the lasso module $\mathcal{P}_{uv^\omega}$. If we cannot find a ranking function f_n for $\mathcal{P}_{uv^\omega}$ our algorithm is unable to decide termination of $\mathcal{P}$ and returns. Otherwise we take a ranking function f_n and construct a rank certificate $\mathcal{I}_n$ for f_n and $\mathcal{P}_{uv^\omega}$ (see Step 2 in Section 4). Afterwards we use the rank certificate to construct the module $\mathcal{P}_n$. Termination of each fair ω -trace of $\mathcal{P}_n$ can be shown using the ranking function f_n and the rank certificate $\mathcal{I}_n$, i.e., $(\mathcal{P}_n, f_n, \mathcal{I}_n)$ is a certified module (see Step 3 in Section 4). If we were not able to find a counterexample to inclusion in line 2, the program $\mathcal{P}$ is already decomposed into certified modules. We have proven termination and return the certified modules $(\mathcal{P}_1, f_1, \mathcal{I}_1), \dots, (\mathcal{P}_n, f_n, \mathcal{I}_n)$.

Our approach lends itself to a variation of the above algorithm where one uses an exit condition different from the inclusion check in line 2. In that case, the algorithm returns the modules $\mathcal{P}_1, \dots, \mathcal{P}_{n-1}$ constructed so far and, in addition, a “remainder program” $\mathcal{P}_{\text{rem}}$ which is constructed via the language-theoretic difference of Büchi automata.

$$\mathcal{P}_{\text{rem}} := \mathcal{P} \setminus (\mathcal{P}_1 \cup \dots \cup \mathcal{P}_{n-1})$$

This is interesting in a variety of contexts, e.g., when we found an ω -trace that is nonterminating, or when we found an ω -trace whose termination analysis failed, or simply in case of a timeout. The remainder program can then be analyzed manually, or it can be used as a runtime monitor, etc.

6 Evaluation

It is unlikely that one approach outperforms all others on all kinds of programs, either in effectiveness (how many termination problems can be solved?) or in

efficiency (... in what time?). In this paper, we have presented the base algorithm of a new approach to termination analysis. To explore optimizations and possibilities of integration with other approaches must remain a topic of future work.

The question is whether our approach is a potentially useful addition to the portfolio of existing approaches. Therefore, the goal of the present experimental evaluation must be restricted to showing that the approach has a practical potential *in principle*, regarding effectiveness and regarding efficiency. This is not obvious since there are at least two “mission-critical” questions, namely:

- Will the algorithm just learn one terminating program $\mathcal{P}_1, \mathcal{P}_2, \dots$ after the other, going through an infinite (or just unrealistically high) number of sample traces $\pi_1, \pi_2, \dots$?
- Will the check of inclusion between Büchi automata (which is notoriously difficult and still an object of ongoing work [8,37]) be a ‘bad’ bottleneck?

We put the evaluation into the context of a previous, very thorough evaluation² in [9] that contained 260 terminating programs. Out of the 260 programs, our tool can handle 236 programs. This, we believe, indicates the potential effectiveness of our approach. In comparison regarding effectiveness, COOPERATING-T2, the “winner” of the evaluation in [9] (a highly optimized tool which integrates several approaches) can handle 14 programs that our tool cannot handle, but our tool can handle 5 programs that COOPERATING-T2 cannot handle (namely `a.10.c.t2.c`, `eric.t2.c`, `sas2.t2.c`, `spiral.t2.c` and `sumit.t2.c`). This confirms our point that no single approach provides a “silver bullet” and that it is desirable to have a large portfolio of approaches.

We implemented the algorithm presented in Section 5 in the tool ULTIMATE BUCHIAUTOMIZER that analyzes termination of C programs. The input programs and the modules are represented by Büchi automata. In order to support (possibly recursive) functions, we use Büchi automata over nested words [1] (we do not introduce the formalism in order to avoid the notational overhead) and implemented an automata library for these automata. We do not check the inclusion

$$\mathcal{L}(\mathcal{P}) \subseteq \mathcal{L}(\mathcal{P}_1) \cup \dots \cup \mathcal{L}(\mathcal{P}_n)$$

directly. Instead, we complement the modules and check the emptiness of their intersection with the program

$$\mathcal{L}(\mathcal{P}) \cap \mathcal{L}(\overline{\mathcal{P}_1}) \cup \dots \cup \mathcal{L}(\overline{\mathcal{P}_n})$$

which allow us to reuse intermediate results in further iterations. For complementing our Büchi automata we extended[41] the rank-based approach [18] to (Büchi) nested word automata. The sample ω -traces whose termination we analyze are obtained as counterexamples of an emptiness check that is implemented in our automata library. This emptiness check is purely automata theoretic, does

² <http://verify.rwth-aachen.de/brockschmidt/Cooperating-T2/>

not exploit any information about the program, but prefers short counterexamples. We use the tool LASSORANKER [24,31] to synthesize ranking functions and supporting invariants for lassos. The Floyd-Hoare annotation is obtained via interpolation (alternative variant of Step 2 in Section 4). For interprocedural ω -traces we resort to *nested interpolants* [25]. As interpolating theorem prover we use SMTInterpol [10]. While constructing the modules, we apply Modification rule 1 (merge locations) always and we apply Modification rule 2 (add transitions) lazily in the following sense. Only if the automata library queries the existence of a transition in the module, we check whether this transition can be added by applying Modification rule 2. Our tool is available as a command line version for download as well as via a web interface at the following URL.

<http://ultimate.informatik.uni-freiburg.de/BuchiAutomizer/>

The following table shows the results for a subset of the benchmarks from [9] where our tool run on a computer with an Intel Core i5-3340M CPU with 2.70GHz. Our tool and as well as LASSORANKER the SMT solver, and the automata library are written in Java. The maximum heap size of the Java virtual machine was set to 4GB (`-Xmx4G`).

For each example we list the lines of code of this example, the overall runtime that our tool needed and the time that our tool spend for analyzing lassos, constructing modules, and checking language inclusion of Büchi automata. Furthermore, we list the number of certified modules that had a trivial ranking function (e.g., $f(x) = 0$), the number of certified modules that had a non-trivial ranking function, and the number of states of the largest module that was constructed.

filename	program size	overall runtime	lasso analysis time	module constr. time	Büchi inclusion time	modules trivial rf	modules non-trivial rf	module size (maximum)
a.10.c.t2.c	183	9s	2.8s	0.7s	2.1s	2	9	5
bf20.t2.c	156	6s	0.7s	0.9s	1.9s	6	7	9
bubbleSort.t2.c	109	5s	0.7s	0.3s	1.2s	5	5	5
consts1.t2.c	40	2s	0.3s	0.1s	0.2s	2	1	5
edn.t2.c	294	119s	18.8s	7.7s	89.0s	141	15	58
eric.t2.c	53	10s	1.1s	1.7s	5.0s	4	6	14
firewire.t2.c	178	28s	3.6s	1.3s	19.0s	12	7	8
mc91.t2.c	47	12s	1.2s	0.6s	4.3s	4	10	8
p-43-terminate.t2.c	727	124s	2.1s	4.2s	110.6s	6	18	5
reverse.t2.c	1351	14s	3.1s	1.2s	2.9s	2	3	12
s3-work.t2.c	3229	28s	2.1s	4.1s	11.5s	6	12	22
sas2.t2.c	192	12s	1.3s	3.0s	5.5s	12	6	17
spiral.c	65	38s	0.9s	1.3s	32.7s	8	12	14
sumit.t2.c	83	4s	1.0s	0.2s	0.7s	4	2	4
traverse_twice.t2.c	1428	12s	1.7s	1.4s	3.2s	2	4	18
ud.t2.c	279	32s	2.1s	3.8s	22.1s	30	25	32

More results³ of our tool can be found at the SV-COMP 2014 [6] where our tool participated in the demonstration category on termination.

Discussion. A reader who is familiar with Büchi automata may wonder why it is feasible to complement Büchi automata of these sizes. The answer lies in the flexibility that our definition of a module allows. We tuned the construction of modules in a way that the “amount of nondeterminism” is kept low. However, it is still part of our future work to find a class of Büchi automata that can be easily complemented but does not hinder the module from accepting many traces.

7 Related work

Our method is related to control flow refinement [22]. There, a multi-path loop is transformed into a semantically equivalent code fragment with simpler loops. For example, following the algebraic decomposition rule

$$(a + b)^* = (b^*ab^*)^+ + b^*$$

the loop with the choice of two paths a and b is transformed into the nondeterministic choice of two loops, one where a appears and one where it does not.

We extend control flow refinement by adding fairness constraints [40] and our reasoning is based on ω -regular languages. In our running example (if we read a as the outer and b as the inner loop) we decomposed the ω -regular expression describing the nested loops as follows

$$(a + b)^\omega = (a + b)^*b^\omega + (b^*a)^\omega.$$

We do not enforce the use of a fixed set of algebraic decomposition rules. Instead, we propose an algorithm that builds a decomposition on demand from simple termination arguments. Thus, we partition a set of traces only when it is necessary and, *by construction*, we produce only modules that are guaranteed to have a simple termination argument.

There are many other termination analyses, e. g., [3,14,16,19,20,21,39]. Most related are the termination analyses based on transition invariants and termination analyses based on size-change termination.

Termination analyses based on transition invariants [9,12,13,23,27,29,34,35] combine different, independently obtained ranking functions to a termination argument. Using transition invariants it is sufficient to cover finite repetitions of the loop. In our running example, one could cover the loop by

$$(a + b)^+ = b^+ + (b^*ab^*)^+$$

using the same simple ranking functions as our method for each case. Covering only finite traces is sound, as it can be shown that

$$(a + b)^\omega = (a + b)^*b^\omega + (a + b)^*(b^*ab^*)^\omega$$

³ <http://sv-comp.sosy-lab.org/2014/results/>

using Ramsey’s Theorem. In our approach, instead of having to introduce $(a+b)^*$, we can get a more precise characterization of the code before the infinite loop; also, we can base our case-distinction on which path was taken *before* the loop was reached. Furthermore, we get smaller expressions. Compare the expression

$$(a + b)^*(b^*ab^*)^\omega$$

with our expression $(b^*a)^\omega$. Although they describe exactly the same traces, our expression is simpler and therefore leads to a simpler termination proof. Redefining the loop entry point or unfolding the loops are intrinsic techniques in our approach (as opposed to add-on heuristics). If for the program $(ab)^\omega$, it is simpler to prove the correctness of the loop $(baba)$, we use the fact that

$$(ab)^\omega = a(baba)^\omega.$$

The idea of size-change termination [4,17,28] is to track the value of (auxiliary) variables and show the absence of infinite executions by showing that one value would be decreased infinitely often in a well-ordered domain. The (auxiliary) variables can be seen as a predefined set of mutually independent termination arguments.

In contrast with the above approaches, a termination argument in our setting is a *stand alone* module (its validity is checked for the corresponding fair ω -traces, independently from all other program traces). In contrast, a component of a lexicographic ranking function, a disjunct of a transition invariant, or a size-change variable makes sense only as part of a global termination argument (whose validity has to be checked for the global program).

Finally, we use “learning” as a metaphor rather than as a technical term, in contrast with the work in [30] which uses machine learning for termination analysis.

8 Conclusion and Future Work

We have presented a algorithm for termination analysis that transforms a program into a nondeterministic choice of programs. Our transformation is not guided by the syntactic structure of the program, but by its semantics. Instead of decomposing the program into modules and analyzing termination of the modules, we construct modules that we learned from sample traces and that are terminating by construction.

The general idea of such a transformation is the same as for *trace refinement* [32]: move disjunction over abstract values to the disjunction over sets of traces. The formalization of the shared idea and the exploration of its theoretical and practical consequences for program analyses is a topic of future work.

References

1. R. Alur and P. Madhusudan. Adding nesting structure to words. *J. ACM*, 56(3), 2009.

2. K. R. Apt, F. S. de Boer, and E.-R. Olderog. *Verification of Sequential and Concurrent Programs, 3rd Edition*. Texts in Computer Science. Springer-Verlag, 2009. 502 pp, ISBN 978-1-84882-744-8.
3. D. Babic, A. J. Hu, Z. Rakamaric, and B. Cook. Proving termination by divergence. In *SEFM*, pages 93–102, 2007.
4. A. M. Ben-Amram. Size-change termination, monotonicity constraints and ranking functions. In *CAV*, pages 109–123, 2009.
5. A. M. Ben-Amram and S. Genaim. On the linear ranking problem for integer linear-constraint loops. In *POPL*, pages 51–62, 2013.
6. D. Beyer. Status report on software verification - (competition summary sv-comp 2014). In *TACAS*, pages 373–388, 2014.
7. A. R. Bradley, Z. Manna, and H. B. Sipma. Linear ranking with reachability. In *CAV*, pages 491–504, 2005.
8. S. Breuers, C. Löding, and J. Olschewski. Improved ramsey-based büchi complementation. In *FoSSaCS*, pages 150–164, 2012.
9. M. Brockschmidt, B. Cook, and C. Fuhs. Better termination proving through cooperation. In *CAV*, pages 413–429, 2013.
10. J. Christ, J. Hoenicke, and A. Nutz. Smtinterpol: An interpolating smt solver. In *SPIN*, pages 248–254, 2012.
11. B. Cook, D. Kroening, P. Rümmer, and C. M. Wintersteiger. Ranking function synthesis for bit-vector relations. In J. Esparza and R. Majumdar, editors, *TACAS*, volume 6015 of *Lecture Notes in Computer Science*, pages 236–250. Springer, 2010.
12. B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *PLDI*, pages 415–426, 2006.
13. B. Cook, A. Podelski, and A. Rybalchenko. Proving program termination. *Commun. ACM*, 54(5):88–98, 2011.
14. B. Cook, A. See, and F. Zuleger. Ramsey vs. lexicographic termination proving. In *TACAS*, pages 47–61, 2013.
15. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.
16. P. Cousot and R. Cousot. An abstract interpretation framework for termination. In *POPL*, pages 245–258, 2012.
17. S. Fogarty and M. Y. Vardi. Büchi complementation and size-change termination. *Logical Methods in Computer Science*, 8(1), 2012.
18. E. Friedgut, O. Kupferman, and M. Y. Vardi. Büchi complementation made tighter. In *ATVA*, pages 64–78, 2004.
19. P. Ganty and S. Genaim. Proving termination starting from the end. In *CAV*, pages 397–412, 2013.
20. J. Giesl, R. Thiemann, P. Schneider-Kamp, and S. Falke. Automated termination proofs with AProVE. In *RTA*, pages 210–220, 2004.
21. S. Grebenshchikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In *PLDI*, pages 405–416, 2012.
22. S. Gulwani, S. Jain, and E. Koskinen. Control-flow refinement and progress invariants for bound analysis. In *PLDI*, pages 375–385, 2009.
23. W. R. Harris, A. Lal, A. V. Nori, and S. K. Rajamani. Alternation for termination. In *SAS*, pages 304–319, 2010.
24. M. Heizmann, J. Hoenicke, J. Leike, and A. Podelski. Linear ranking for linear lasso programs. In *ATVA*, pages 365–380, 2013.
25. M. Heizmann, J. Hoenicke, and A. Podelski. Nested interpolants. In *POPL*, pages 471–482, 2010.

26. G. J. Holzmann. The model checker spin. *IEEE Trans. Software Eng.*, 23(5):279–295, 1997.
27. D. Kroening, N. Sharygina, A. Tsitovich, and C. M. Wintersteiger. Termination analysis with compositional transition invariants. In *CAV*, pages 89–103, 2010.
28. C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. In *POPL*, pages 81–92, 2001.
29. W. Lee, B.-Y. Wang, and K. Yi. Termination analysis with algorithmic learning. In *CAV*, pages 88–104, 2012.
30. W. Lee, B.-Y. Wang, and K. Yi. Termination analysis with algorithmic learning. In *Computer Aided Verification*, pages 88–104. Springer, 2012.
31. J. Leike and M. Heizmann. Ranking templates for linear loops. In *TACAS*, pages 172–186, 2014.
32. L. Mauborgne and X. Rival. Trace partitioning in abstract interpretation based static analyzers. In *ESOP*, pages 5–20. Springer, 2005.
33. A. Podelski and A. Rybalchenko. A complete method for the synthesis of linear ranking functions. In *VMCAI*, pages 239–251, 2004.
34. A. Podelski and A. Rybalchenko. Transition invariants. In *LICS*, pages 32–41, 2004.
35. A. Podelski and A. Rybalchenko. Transition predicate abstraction and fair termination. In *POPL*, pages 132–144, 2005.
36. C. Popeea and A. Rybalchenko. Compositional termination proofs for multi-threaded programs. In *TACAS*, pages 237–251, 2012.
37. M.-H. Tsai, S. Fogarty, M. Y. Vardi, and Y.-K. Tsay. State of büchi complementation. In *CIAA*, pages 261–271, 2010.
38. M.-H. Tsai, Y.-K. Tsay, and Y.-S. Hwang. Goal for games, omega-automata, and logics. In *CAV*, pages 883–889, 2013.
39. C. Urban and A. Miné. An abstract domain to infer ordinal-valued ranking functions. In *ESOP*, pages 412–431, 2014.
40. M. Y. Vardi. Verification of concurrent programs: The automata-theoretic framework. *Ann. Pure Appl. Logic*, 51(1-2):79–98, 1991.
41. X. Wu. Three operations on Büchi nested word automata for program verification. Master’s thesis, University of Freiburg, Germany, 2011.

9 Proofs

Proof (of Lemma 1). Let $\mathfrak{s}_1, \mathfrak{s}_2, \dots$ be a fair ω -trace of $\mathcal{P}$. By definition, the final location ℓ_{fin} is visited infinitely often, i.e., there is an infinite sequence $k_1 < k_2 < \dots$ such that after $\mathfrak{s}_{k_i}$ the final location is visited. Assume that the fair ω -trace is not terminating. Then, there exists an infinite sequence of valuations $\nu_0, \nu_1, \dots$ such that for each $i \in \mathbb{N}$ the pair (ν_i, ν_{i+1}) is contained in the transition relation of the statement $\mathfrak{s}_i$. By the definition of the ranking function f , its value decreases every time final location is visited, i.e.,

$$f(\nu_{k_1}) \succ f(\nu_{k_2}) \succ \dots$$

This is not possible since f maps into a well-ordered set. Hence and every fair ω -trace of $\mathcal{P}$ is terminating.

Proof (of Theorem 1). First, we show that f is a ranking function for the module $\mathcal{P}$, afterwards this theorem is a direct consequence of Lemma 1.

Consider a finite path of the module $\mathcal{P}$

$$\ell_0 \xrightarrow{\mathfrak{s}_1} \dots \xrightarrow{\mathfrak{s}_k} \ell_k \xrightarrow{\mathfrak{s}_{k+1}} \dots \xrightarrow{\mathfrak{s}_n} \ell_n$$

that starts in the initial location and visits the final location in the k -th step and in the n -th step, i. e.,

$$\ell_0 = \ell_{\text{init}} \quad \text{and} \quad \ell_k = \ell_n = \ell_{\text{fin}}.$$

Let $\nu_0, \dots, \nu_n$ be a sequence of valuations such that each pair of successive valuations (ν_i, ν_{i+1}) is in the transition relation of the statement $\mathfrak{s}_i$, i.e.,

$$\nu_0 \xrightarrow{\mathfrak{s}_1} \dots \xrightarrow{\mathfrak{s}_k} \nu_k \xrightarrow{\mathfrak{s}_{k+1}} \dots \xrightarrow{\mathfrak{s}_n} \nu_n.$$

Next, we show that the strict inequality $f(\nu_n) < f(\nu_k)$ holds. Therefore, we extend the valuation ν_i with a value for the auxiliary variable **oldrnk**. We define this value $oldrnk_i$ as follows.

$$oldrnk_i := \begin{cases} \infty & \text{if } \forall j < i. \ell_j \neq \ell_{\text{fin}}, \\ f(\nu_j) & \text{where } j \text{ is greatest index} \\ & \text{such that } j < i \text{ and } \ell_j = \ell_{\text{fin}} \end{cases} \quad \text{otherwise}$$

This extended valuation $(\nu_i \cup \{\mathbf{oldrnk} \mapsto oldrnk_i\})$ is denoted by $\bar{\nu}_i$. Now, we show by induction that for all indices i of our automaton run the extended valuation $\bar{\nu}_i$ is contained in the invariant $\mathcal{I}(\ell_i)$, i.e.,

$$\bar{\nu}_i \in \mathcal{I}(\ell_i) \quad \text{for } i = 0 \dots n.$$

Induction basis $i = 0$. The extended valuation $\bar{\nu}_0$ is an element of $\mathcal{I}(\ell_{\text{init}})$, because the initial value of **oldrnk** is ∞ and the predicate **oldrnk** = ∞ is equivalent to the invariant $\mathcal{I}(\ell_{\text{init}})$.

Induction step $i \rightsquigarrow i + 1$.

- Case 1: i is index of an accepting state:
By the induction hypothesis the extended valuation $\bar{\nu}_i$ is contained in $\mathcal{I}(\ell_i)$. According to the definition of a rank certificate the predicate $\mathcal{I}(\ell_{i+1})$ is a superset of the predicate $post(\mathcal{I}(\ell_i), \mathbf{oldrnk} := f(\bar{\nu}); \mathfrak{s}_i)$. Above we defined the value $oldrnk_{i+1} := f(\nu_i)$. Hence, the extended valuation $\bar{\nu}_{i+1}$ is contained in $\mathcal{I}(\ell_{i+1})$.
- Case 2: i is not index of an accepting state:
By the induction hypothesis the extended valuation $\bar{\nu}_i$ is contained in $\mathcal{I}(\ell_i)$. According to the definition of a rank certificate the predicate $\mathcal{I}(\ell_{i+1})$ is a superset of the predicate $post(\mathcal{I}(\ell_i), \mathfrak{s}_i)$. Since the auxiliary variable **oldrnk** does not appear in the program it is not modified by the statement $\mathfrak{s}_i$. According to the definition above, the value $oldrnk_{i+1}$ coincides with the value $oldrnk_i$. Hence, the extended valuation $\bar{\nu}_{i+1}$ is contained in the predicate $\mathcal{I}(\ell_{i+1})$.

Let $k_0 < k_1 < \dots < k_m$ be the ascending chain of indices such that $k_0 = k$, $k_m = n$ and $\ell_{k_j} = \ell_{\text{fin}}$ for all $j = 0, \dots, m$. Since $\nu_{k_j} \in \mathcal{I}(\ell_{k_j})$ the strict inequality $f(\nu_{k_j}) < \text{oldrnk}_{k_j}$ holds. As defined above, the value oldrnk_{k_j} is defined as $\nu_{k_{j-1}}$, hence the following sequence is a descending chain

$$f(\nu_{k_0}) > f(\nu_{k_2}) > \dots > f(\nu_{k_m})$$

and especially $f(\nu_k) > f(\nu_n)$ holds and. Therefore f is a ranking function for $\mathcal{P}$. Using Lemma 1, we conclude that each fair ω -trace of $\mathcal{P}$ is terminating.

Proof (of Theorem 2). The proof proceeds in two steps. First we show that a program can be decomposed into modules. In the second step we show that we can give a ranking function and rank certificate for each terminating module.

The theorem of Büchi says that we can decompose each ω -regular language L into a finite disjunction

$$L = \bigcup_{i=1}^n U_i.V_i^\omega$$

where each U_i and each V_i is a regular language.

Let $\mathcal{A}_i^U$ and $\mathcal{A}_i^V$ be finite deterministic automata that recognize the regular languages U_i and V_i , respectively. We construct the module $\mathcal{P}_i$ using the standard construction where $\mathcal{A}_i^U$ and $\mathcal{A}_i^V$ are combined to a Büchi automaton that recognizes the language $U_i.V_i^\omega$. The (single) final state of $\mathcal{P}_i$ is the initial state of $\mathcal{A}_i^V$ in this construction. Hence, we can decompose the program $\mathcal{P}$ into fair modules

$$\mathcal{L}(\mathcal{P}) = \mathcal{L}(\mathcal{P}_1) \cup \dots \cup \mathcal{L}(\mathcal{P}_n).$$

Since the modules contain the same executions as the program $\mathcal{P}$ they must also be terminating. Existence of a computable ranking functions f is a classical result. We take such a ranking function and extend the module by a specification that asserts that this ranking function is decreasing whenever the final location is visited. A Floyd-Hoare annotation which shows partial correctness of the extended module can be seen as rank certificate $\mathcal{I}$. The existence of this Floyd-Hoare annotation is also a classical result [2].