

Predicate Abstraction and CEGAR for Disproving Termination of Higher-Order Functional Programs

Takuya Kuwahara¹(✉), Ryosuke Sato², Hiroshi Unno³, and Naoki Kobayashi²

¹ Knowledge Discovery Research Laboratories, NEC, Minato, Japan
t-kuwahara@me.jp.nec.com

² The University of Tokyo, Bunkyo, Japan
{ryosuke,koba}@kb.is.s.u-tokyo.ac.jp

³ University of Tsukuba, Tsukuba, Japan
uhiro@cs.tsukuba.ac.jp

Abstract. We propose an automated method for disproving termination of higher-order functional programs. Our method combines higher-order model checking with predicate abstraction and CEGAR. Our predicate abstraction is novel in that it computes a mixture of under- and overapproximations. For non-determinism of a source program (such as random number generation), we apply underapproximation to generate a subset of the actual branches, and check that some of the branches in the abstract program is non-terminating. For operations on infinite data domains (such as integers), we apply overapproximation to generate a superset of the actual branches, and check that every branch is non-terminating. Thus, disproving non-termination reduces to the problem of checking a certain branching property of the abstract program, which can be solved by higher-order model checking. We have implemented a prototype non-termination prover based on our method and have confirmed the effectiveness of the proposed approach through experiments.

1 Introduction

We propose an automated method for disproving termination of higher-order functional programs (i.e., for proving that a given program does not terminate for *some* input). The method plays a role complementary to the automated method for proving termination of higher-order programs (i.e., for proving that a given program terminates for *all* inputs) [18]. Several methods have recently been proposed for proving non-termination of programs [7–9, 11, 13, 14, 19], but most of them have focused on *first-order* programs (or, while programs) that can be represented as finite control graphs. An exception is work on term rewriting systems (TRS) [9, 11]; higher-order programs can be encoded into term rewriting systems, but the definition of non-termination is different: TRS is non-terminating if there exists a term that has a non-terminating rewriting sequence, not necessarily the initial term.

Our approach is based on a combination of higher-order model checking [15, 21] with predicate abstraction and CEGAR (counterexample-guided abstraction refinement). Values of a base type (such as integers) are abstracted to (tuples of) Booleans by using predicates, and higher-order functions are abstracted accordingly. Higher-order model checking is then used to analyze the abstracted program. A combination of predicate abstraction and higher-order model checking has been previously proposed for verifying safety properties of higher-order programs (i.e., for proving that a program does not reach an error state in *all* execution paths) [16]. With respect to that work, the approach of the present paper is novel in that we combine *overapproximation* and *underapproximation*. Note that predicate abstraction [3, 12, 16] usually yields an overapproximation, i.e., an abstract program that contains a *superset* of possible execution paths of the original program. With such an abstraction, non-termination of the abstract program (the existence of a non-terminating path) does not imply that of the original program. To address this problem, we use both under- and overapproximations. For a deterministic computation step of the original program, we apply overapproximation but check that *every* branch of the overapproximation has a non-terminating path. For a non-deterministic branch of the original program (such as random number generation and an input from the environment), we apply *under*-approximation, and check that *some* branch of the underapproximation has a non-terminating path.

Figure 1 illustrates how under- and overapproximations are combined. The program considered here is of the form:

$$\text{let } x = * \text{ in let } y = x + 1 \text{ in let } z = * \text{ in } \dots .$$

Here, $*$ generates a random integer. Thus, the program has the execution tree shown on the top of Fig. 1. The first and third steps are non-deterministic, while the second step (corresponding to $y = x + 1$) is deterministic. Suppose that the predicates used for abstracting the values of x , y , and z are $x > 0$, $y > 0$, and $0 \leq z < x + y$ (these predicates do not necessarily yield a good abstraction, but are sufficient for explaining the combination of under- and overapproximations). Then, the abstract program has the execution tree shown on the bottom of the figure. Due to the predicate abstraction, the infinitely many branches on the value of x have been replaced by two branches $x > 0$ and $\neg x > 0$. The node \exists means that only one of the branches needs to have an infinite path (for the original program having a non-terminating path). The deterministic path from $x = n$ to $y = n + 1$ has now been replaced by non-deterministic branches $y > 0$ and $\neg y > 0$. The node \forall indicates that *every* child of the node must have an infinite path. Below the node $x > 0$, however, we do not have a node for $\neg y > 0$, as $x > 0$ and $y = x + 1$ imply $y > 0$. The infinite branches on z have been replaced by non-deterministic branches on $\neg(0 \leq z < x + y)$ or $0 \leq z < x + y$. As is the case for x , the branches are marked by \exists , meaning that one of the branches needs to have an infinite path. Note that below the node $\neg x > 0$, we only have a branch for $\neg(0 \leq z < x + y)$. This is because, when $x \leq 0$, there may be no z that satisfies $0 \leq z < x + y$; so, even if there may be an infinite

execution sequence along that path, we cannot conclude that the source program is non-terminating. Thus, this part of the tree provides an *under*-approximation of the source program.

An abstract program is actually represented as a tree-generating program that generates an execution tree like the one shown on the bottom of Fig. 1. Higher-order model checking is then used for checking, informally speaking, that *every* child of each \forall -node has a non-terminating path, and that *some* child of each \exists -node has a non-terminating path.

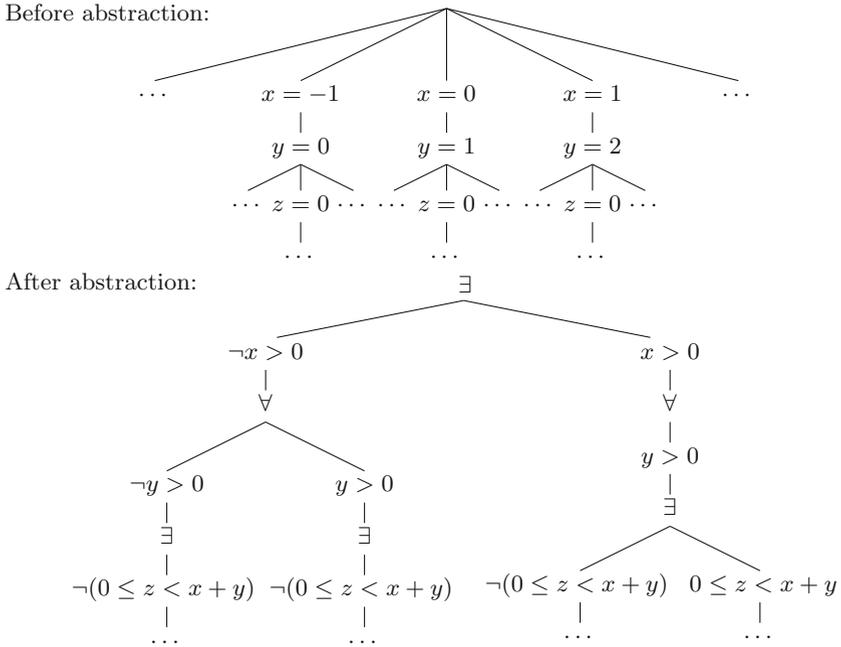


Fig. 1. Execution trees before/after abstraction

The use of overapproximation for disproving termination has also been proposed recently by Cook et al. [8]. Although their theoretical framework is general, their concrete method for automation is limited to first-order programs. They also propose a restricted form of combination of underapproximation and overapproximation, but underapproximation can be followed by overapproximation, but not vice versa.

The rest of this paper is structured as follows. Section 2 defines the language used as the target of our verification. Sections 3 and 4 describe predicate abstraction and CEGAR respectively. Section 5 reports experiments. Section 6 discusses related work and Sect. 7 concludes the paper.

2 Language

In this section, we introduce the language of source programs, used as the target of non-termination verification. It is a simply typed, call-by-value higher-order functional language. Throughout the paper, we often use the following abbreviations: \tilde{e} for a possibly empty sequence e_1, \dots, e_n , and $\{e_i\}_{i \in \{1, \dots, n\}}$ for the set $\{e_1, \dots, e_n\}$.

$$\begin{aligned}
 P \text{ (programs)} &:= \{f_i \tilde{x}_i = e_i\}_{i \in \{1 \dots n\}} \\
 e \text{ (expressions)} &:= () \mid y \tilde{v} \mid \text{if } a \text{ then } e_1 \text{ else } e_2 \\
 &\quad \mid \text{let } x = a \text{ in } e \mid \text{let } x = *_{\text{int}} \text{ in } e \\
 a \text{ (simple expressions)} &::= x \mid n \mid \text{op}(\tilde{a}) \qquad v \text{ (values)} := n \mid y \tilde{v} \\
 \frac{f \tilde{x} = e \in P \quad |\tilde{x}| = |\tilde{v}|}{f \tilde{v} \longrightarrow_P [\tilde{v}/\tilde{x}]e} \quad \frac{\llbracket a \rrbracket = n}{\text{let } x = a \text{ in } e \longrightarrow_P [n/x]e} \quad \text{let } x = *_{\text{int}} \text{ in } e \longrightarrow_P [n/x]e \\
 \text{if } n \text{ then } e_0 \text{ else } e_1 \longrightarrow_P e_0 \text{ if } n \neq 0 \quad \text{if } 0 \text{ then } e_0 \text{ else } e_1 \longrightarrow_P e_1
 \end{aligned}$$

Fig. 2. The syntax and operational semantics of the language

The syntax and operational semantics of the language is given in Fig. 2. The meta-variable f_i ranges over a set of function names, and x, y range over the set of function names and ordinary variables. The meta-variable n ranges over the set of integers, and op over a set of integer operations. We omit Booleans and regard a non-negative integer as true, and 0 as false. We require that $y \tilde{v}$ in the definition of e is a full application, i.e., that all the necessary arguments are passed to y , and $y \tilde{v}$ has a base type. In contrast, $y \tilde{v}$ in the definition of v is a partial application. Whether $y \tilde{v}$ is a full or partial application is actually determined by the simple type system mentioned below.

The expression $\text{let } x = *_{\text{int}} \text{ in } e$ randomly generates an integer, binds x to it, and evaluates e . The meanings of the other expressions should be clear. A careful reader may notice that we have only tail calls. This is for the simplicity of the presentation. Note that it does not lose generality, because we can apply the standard continuation-passing-style (CPS) transformation to guarantee that all the function calls are in this form. We assume that for every program $\{f_i \tilde{x}_i = e_i\}_{i \in \{1 \dots n\}}$, $\text{main} \in \{f_1, \dots, f_n\}$.

We assume that programs are simply-typed. The syntax of types is given by: $\tau ::= \text{int} \mid \star \mid \tau_1 \rightarrow \tau_2$. The types int and \star describe integers and the unit value $()$ respectively. The type $\tau_1 \rightarrow \tau_2$ describes functions from τ_1 to τ_2 . The typing rules for expressions and programs are given in the full version [17], which are standard except that the body of each function definition can only have type \star . This does not lose generality since the CPS transformation guarantees this condition.

The one-step reduction relation $e_1 \longrightarrow_P e_2$ is defined by the rules in Fig. 2, where $\llbracket a \rrbracket$ stands for the value of the simple expression a . A program P is *non-terminating* if there is an infinite reduction sequence $\text{main} \rightarrow_P e_1 \rightarrow_P e_2 \rightarrow_P \dots$

$\text{main} \longrightarrow_D M_1 \longrightarrow_D M_2 \longrightarrow_D \dots$. For example, the program $\{\text{main} = \text{call}(\text{main})\}$ generates an infinite (linear) tree consisting of infinitely many `call` nodes.

Intuitively, the tree generated by a program of the target language describes possible execution sequences of a source program. The property that a source program has a non-terminating execution sequence is transformed to the property of the tree that (i) every child of each br_\forall node has an infinite path, and (ii) some child of each br_\exists node has an infinite path. More formally, the set of (infinite) trees that represent the existence of a non-terminating computation is the largest set **NonTermTrees** such that for every $T \in \mathbf{NonTermTrees}$, T satisfies one of the following conditions.

1. $T = \text{call}(T')$ and $T' \in \mathbf{NonTermTrees}$
2. $T = \text{br}_\forall(T_1, \dots, T_k)$ and $T_i \in \mathbf{NonTermTrees}$ for all $i \in \{1, \dots, k\}$.
3. $T = \text{br}_\exists(T_1, \dots, T_k)$ and $T_i \in \mathbf{NonTermTrees}$ for some $i \in \{1, \dots, k\}$.

The property above can be expressed by MSO (the monadic second order logic; or equivalently, modal μ -calculus or alternating parity tree automata); thus whether the tree generated by a program of the target language belongs to **NonTermTrees** can be decided by higher-order model checking [15, 21].

3.2 Abstraction

We now formalize predicate abstraction for transforming a source program to a program (of the target language) that generates a tree that approximately represents the possible execution sequences of the source program. Following Kobayashi et al. [16], we use *abstraction types* for expressing which predicate should be used for abstracting each value. The syntax of abstraction types is:

$$\begin{aligned} \sigma(\text{abstraction types}) &::= \star \mid \mathbf{int}[Q_1, \dots, Q_k] \mid x : \sigma_1 \rightarrow \sigma_2 \\ Q(\text{predicates}) &::= \lambda x. \varphi \quad \varphi ::= n_1 x_1 + \dots + n_k x_k \leq n \mid \varphi_1 \vee \varphi_2 \mid \neg \varphi \end{aligned}$$

The type \star describes the unit value, and $\mathbf{int}[Q_1, \dots, Q_k]$ describes an integer that should be abstracted by using the predicates Q_1, \dots, Q_k . For example, given an abstraction type $\mathbf{int}[\lambda x. x \leq 1, \lambda x. 2x - 1 \leq 0]$, the integer 1 is abstracted to (**true**, **false**). In the syntax above, we list only linear inequalities as primitive constraints, but we can include other constraints (such as those on uninterpreted function symbols) as long as the underlying theory remains decidable. The type $x : \sigma_1 \rightarrow \sigma_2$ describes a function whose argument and return value should be abstracted according to σ_1 and σ_2 respectively. In σ_2 , the argument can be referred to by x if x has an integer type $\mathbf{int}[Q_1, \dots, Q_k]$. For example, $x : \mathbf{int}[\lambda x. x \leq 0] \rightarrow \mathbf{int}[\lambda y. y - x \leq 0]$ describes a function from integers to integers whose argument should be abstracted using the predicate $\lambda x. x \leq 0$ and whose return value should be abstracted using $\lambda y. y - x \leq 0$. Thus, the successor function (defined by $f x = x + 1$) will be abstracted to a Boolean function $\lambda b. \text{false}$ (because the return value $x + 1$ is always greater than x , no matter whether $x \leq 0$ or not).

The predicate abstraction for expressions and programs is formalized using the relations $\Gamma \vdash e : \sigma \rightsquigarrow M$ and $\vdash P : \Gamma \rightsquigarrow D$, where Γ , called an *abstraction type environment*, is of the form $x_1 : \sigma_1, \dots, x_n : \sigma_n$. Intuitively, $\Gamma \vdash e : \sigma \rightsquigarrow M$ means that under the assumption that each free variable x_i of e is abstracted according to σ_i , the expression e is abstracted to M according to the abstraction type σ . In the judgment $\vdash P : \Gamma \rightsquigarrow D$, Γ describes how each function defined in P should be abstracted.

The relations are defined by the rules in Fig. 4. Here, we consider, without loss of generality, only if-expressions of the form **if** x **then** e_1 **else** e_2 . Also, function arguments are restricted to the syntax: $v ::= y \tilde{v}$. (In other words, constants may not occur; note that xc can be replaced by **let** $y=c$ **in** xy .) We assume that each let-expression is annotated with an abstraction type that should be used for abstracting the value of the variable. Those abstraction types, as well as those for functions are automatically inferred by the CEGAR procedure described in Sect. 4.

$$\begin{array}{c}
 \overline{\Gamma \vdash () : \star \rightsquigarrow \text{end}} \quad \text{(PA-UNIT)} \\
 \\
 \frac{\models b_1 Q_1(a) \wedge \dots \wedge b_k Q_k(a) \Rightarrow \theta_\Gamma \psi_{(b_1, \dots, b_k)} \text{ (for each } b_1, \dots, b_k \in \{\text{true}, \text{false}\})}{\Gamma, x : \mathbf{int}[Q_1, \dots, Q_k] \vdash e : \star \rightsquigarrow M} \\
 \\
 \frac{\Gamma \vdash \mathbf{let } x : \mathbf{int}[Q_1, \dots, Q_k] = a \mathbf{ in } e : \star \rightsquigarrow}{\mathbf{br}_\forall \{ \psi_{(b_1, \dots, b_k)} \rightarrow \mathbf{let } x = (b_1, \dots, b_k) \mathbf{ in } M \mid b_1, \dots, b_k \in \{\text{true}, \text{false}\} \}} \quad \text{(PA-SEXP)} \\
 \\
 \frac{\models x \neq 0 \Rightarrow \theta_\Gamma \psi_1 \quad \models x = 0 \Rightarrow \theta_\Gamma \psi_2 \quad \Gamma \vdash e_1 : \star \rightsquigarrow M_1 \quad \Gamma \vdash e_2 : \star \rightsquigarrow M_2}{\Gamma \vdash \mathbf{if } x \mathbf{ then } e_1 \mathbf{ else } e_2 : \star \rightsquigarrow \mathbf{br}_\forall \{ \psi_1 \rightarrow M_1, \psi_2 \rightarrow M_2 \}} \quad \text{(PA-IF)} \\
 \\
 \frac{\models \theta_\Gamma \psi_{(b_1, \dots, b_k)} \Rightarrow \exists x. b_1 Q_1(x) \wedge \dots \wedge b_k Q_k(x) \text{ (for each } b_1, \dots, b_k \in \{\text{true}, \text{false}\})}{\Gamma, x : \mathbf{int}[Q_1, \dots, Q_k] \vdash e : \star \rightsquigarrow M} \\
 \\
 \frac{\Gamma \vdash \mathbf{let } x : \mathbf{int}[Q_1, \dots, Q_k] = \star_{\mathbf{int}} \mathbf{ in } e : \star \rightsquigarrow}{\mathbf{br}_\exists \{ \psi_{(b_1, \dots, b_k)} \rightarrow \mathbf{let } x = (b_1, \dots, b_k) \mathbf{ in } M \mid b_1, \dots, b_k \in \{\text{true}, \text{false}\} \}} \quad \text{(PA-RAND)} \\
 \\
 \frac{\Gamma(y) = x_1 : \sigma_1 \rightarrow \dots \rightarrow x_k : \sigma_k \rightarrow \sigma}{\Gamma \vdash v_i : [v_1/x_1, \dots, v_{i-1}/x_{i-1}] \sigma_i \rightsquigarrow V_i \text{ for each } i \in \{1, \dots, k\}} \quad \text{(PA-APP)} \\
 \\
 \frac{\Gamma \vdash y v_1 \dots v_k : [v_1/x_1, \dots, v_k/x_k] \sigma \rightsquigarrow y V_1 \dots V_k}{\{f_i : \tilde{x} : \tilde{\sigma}_i \rightarrow \star\}_{i \in \{1, \dots, k\}}, \tilde{x} : \tilde{\sigma}_j \vdash e_i : \star \rightsquigarrow M_i \text{ for each } j \in \{1, \dots, k\}} \\
 \vdash \{f_i \tilde{x}_i = e_i\}_{i \in \{1, \dots, k\}} : \{f_i : \tilde{x} : \tilde{\sigma}_i \rightarrow \star\}_{i \in \{1, \dots, k\}} \rightsquigarrow \{f_i \tilde{x}_i = \mathbf{call}(M_i)\}_{i \in \{1, \dots, k\}}} \quad \text{(PA-PROG)}
 \end{array}$$

Fig. 4. Predicate abstraction rules

The rule PA-UNIT just replaces the unit value with **end**, which represents termination. The rule PA-SEXP overapproximates the value of a simple expression a . Here, θ_Γ is the substitution that replaces each variable x of type $\mathbf{int}[Q'_1, \dots, Q'_n]$ in Γ with $(Q'_1(x), \dots, Q'_n(x))$. For example, if $\Gamma =$

$x : \mathbf{int}[\lambda x.x \leq 0, \lambda x.x \leq 2], y : \mathbf{int}[\lambda y.y \leq x]$, then $\theta_\Gamma(\#_2(x) \wedge \#_1(y))$ is $\#_2(x \leq 0, x \leq 2) \wedge \#_1(y \leq x)$, i.e., $x \leq 2 \wedge y \leq x$. The formula $b_i Q_i(a)$ stands for $Q_i(a)$ if $b_i = \mathbf{true}$, and $\neg Q_i(a)$ if $b_i = \mathbf{false}$. Basically, the rule generates branches for all the possible values (b_1, \dots, b_k) for $(Q_1(a), \dots, Q_k(a))$, and combines them with node \mathbf{br}_\vee (which indicates that this branch has been obtained by an overapproximation). To eliminate impossible values, we compute a necessary condition $\psi_{(b_1, \dots, b_k)}$ for $(Q_1(a), \dots, Q_k(a)) = (b_1, \dots, b_k)$ to hold, and guard the branch for (b_1, \dots, b_k) with $\psi_{(b_1, \dots, b_k)}$. The formula $\psi_{(b_1, \dots, b_k)}$ can be computed by using an SMT solver, as in ordinary predicate abstraction [3, 16]. (The rule generates 2^k branches, leading to code explosion. This is for the sake of simplicity; the eager splitting of branches is avoided in the actual implementation.) The rule PA-IF is similar: branches for the then- and else-clauses are generated, but they are guarded by necessary conditions for the branches to be chosen.

The rule PA-RAND for random number generation is a kind of dual to PA-SEXP. It applies an *underapproximation*, and generates branches for all the possible values (b_1, \dots, b_k) for $(Q_1(x), \dots, Q_k(x))$ under the node \mathbf{br}_\exists . Each branch is guarded by a *sufficient* condition for the existence of a value for x such that $(Q_1(x), \dots, Q_k(x)) = (b_1, \dots, b_k)$, so that for each branch, there must be a corresponding execution path of the source program. The rule PA-APP for applications is the same as the corresponding rule of [16]. Finally, the rule PA-PROG for programs just transforms the body of each function definition, but adds a special node \mathbf{call} to keep track of function calls. Note that a program is non-terminating if and only if it makes infinitely many function calls.

Example 1. Let us consider the following program **LOOP**.

```
loop h x = let b = (x > 0) in
           if b then let d = *int in let y = x + d in h y (loop app) else ()
app m k = k m      main = let r = *int in loop app r
```

LOOP is non-terminating; in fact, if $*_{\mathbf{int}}$ is always evaluated to 1, then we have:

$$\mathbf{main} \longrightarrow^* \mathit{loop\ app\ 1} \longrightarrow^* \mathit{app\ 2\ (loop\ app)} \longrightarrow^* \mathit{loop\ app\ 2} \longrightarrow^* \dots$$

Let $\Gamma_{\mathbf{LOOP}}$ be an abstraction type environment:

$$\begin{aligned} \mathit{loop} &: (\mathbf{int}[\lambda \nu. \nu > 1] \rightarrow (\mathbf{int}[\lambda \nu. \nu > 1] \rightarrow \star) \rightarrow \star) \rightarrow \mathbf{int}[\lambda \nu. \nu > 1] \rightarrow \star \\ \mathit{app} &: \mathbf{int}[\lambda \nu. \nu > 1] \rightarrow (\mathbf{int}[\lambda \nu. \nu > 1] \rightarrow \star) \rightarrow \star \end{aligned}$$

By using $\Gamma_{\mathbf{LOOP}}$ and the following abstraction types for b , d , and r :

$$b : \mathbf{int}[\lambda \nu. \nu \neq 0], d : \mathbf{int}[\lambda \nu. x + \nu > 1], r : \mathbf{int}[\lambda \nu. \nu > 1],$$

the program **LOOP** is abstracted to the following program $D_{\mathbf{LOOP}}$.

```

loop h x = call(br∀{true → let b=true in M1,
                  ¬x → let b=false in M1})
app m k = call(k m)
main = call(br∃{true → let r=true in loop app r,
               true → let r=false in loop app r})
where
M1 = br∀{b → M2, ¬b → end}
M2 = br∃{true → let d=true in M3, true → let d=false in M3}
M3 = br∀{d → let y=true in h y (loop app),
          ¬d → let y=false in h y (loop app)}.
    
```

For example, $\text{let } b : \text{int}[\lambda\nu.\nu \neq 0] = x > 0 \text{ in } e$ is transformed by PA-SEXP as follows:

$$\frac{\begin{array}{l} \models (x > 0) \neq 0 \Rightarrow \text{true} \quad \models \neg((x > 0) \neq 0) \Rightarrow \neg(x > 1) (= \theta_{\Gamma}(\neg x)) \\ \Gamma, b : \text{int}[\lambda\nu.\nu = 0] \vdash e : \star \rightsquigarrow M_1 \end{array}}{\Gamma \vdash \text{let } b : \text{int}[\lambda\nu.\nu = 0] = x > 0 \text{ in } e \rightsquigarrow \text{br}_{\forall}\{\text{true} \rightarrow \text{let } b = \text{true} \text{ in } M_1, \neg x \rightarrow \text{let } b = \text{false} \text{ in } M_1\}}$$

where

$$\Gamma = \Gamma_{\mathbf{LOOP}}, h : (\text{int}[\lambda\nu.\nu > 1] \rightarrow (\text{int}[\lambda\nu.\nu > 1] \rightarrow \star) \rightarrow \star), x : \text{int}[\lambda\nu.\nu > 1].$$

Here, recall that a non-zero integer is treated as **true** in the source language; thus, $\neg((x > 0) \neq 0)$ means $x \leq 0$. Since $\mathbf{Tree}(D_{\mathbf{LOOP}}) \in \mathbf{NonTermTrees}$, we can conclude that the program **LOOP** is non-terminating (based on Theorem 1 below). \square

The soundness of predicate abstraction is stated as follows (see the full version [17] for a proof).

Theorem 1. *Suppose $\vdash P : \Gamma \rightsquigarrow D$. If $\mathbf{Tree}(D) \in \mathbf{NonTermTrees}$, then P is non-terminating.*

4 Counterexample-Guided Abstraction Refinement (CEGAR)

This section describes our CEGAR procedure to refine abstraction based on a counterexample. Here, a *counterexample* output by a higher-order model checker is a finite subtree T of $\mathbf{Tree}(D)$, obtained by removing all but one branches of each br_{\forall} node. Figure 5 illustrates $\mathbf{Tree}(D)$ and a corresponding counterexample (showing $\mathbf{Tree}(D) \notin \mathbf{NonTermTrees}$). In the figure, “...” indicates an infinite path. For each br_{\forall} node, a model checker picks one branch containing a finite path, preserving the branches of the other nodes (br_{\exists} , call , and end).

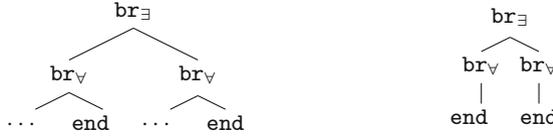


Fig. 5. $\text{Tree}(D)$ (left) and a corresponding counterexample (right)

We analyze each path of the counterexample tree to infer new abstraction types for refining abstraction. To that end, we need to distinguish between two types of paths in the counterexample tree: one that has been introduced due to overapproximation, and the other due to underapproximation. Figure 6 illustrates the two types. For each type, the lefthand side shows the computation tree of a source program, and the righthand side shows the tree generated by the abstract program. Thick lines show a path of a counterexample tree. In the example of Type I, the computation of a source program takes the then-branch and falls into a non-terminating computation, but predicate abstraction has introduced the spurious path taking the else branch, which was detected as a part of the counterexample. In the example of Type II, a source program generates a random number and non-deterministically branches to a non-terminating computation or a terminating computation. After predicate abstraction, the two branches by the random number generation have been merged; instead, the next deterministic computation step has been split into two by an overapproximation. This situation occurs, for example, for

```
let x : int[] = *int in let y : int[λy.y ≠ 0] = x in if y then loop() else ()
```

The program generated by the abstraction is

```
br∃{true → br∀{true → let y = true in ...,
               true → let y = false in ...}}
```

Thus, the branches at $*_{\text{int}}$ in the original program have been moved to the branches at br_{\forall} . The classification of the paths of a counterexample into Type I or II can be performed according to the feasibility of the path, i.e., whether there is a corresponding computation path in the source program. An infeasible path is Type I, since it has been introduced by an overapproximation, and a feasible path is Type II; it has a corresponding computation path, but the two kinds of non-determinism (expressed by br_{\exists} and br_{\forall}) have been confused by predicate abstraction. We need to refine the predicates (or, abstraction types) used for overapproximation for a Type I path, and those used for underapproximation for a Type II path. In the example program above, by refining the abstraction type for x to $\text{int}[\lambda x.x \neq 0]$, we obtain

```
br∃{true → let x = true in br∀{x → let y = true in ...},
     true → let x = false in br∀{¬x → let y = false in ...}}
```

Thus, the branches on terminating/non-terminating paths are moved to the node br_{\exists} .

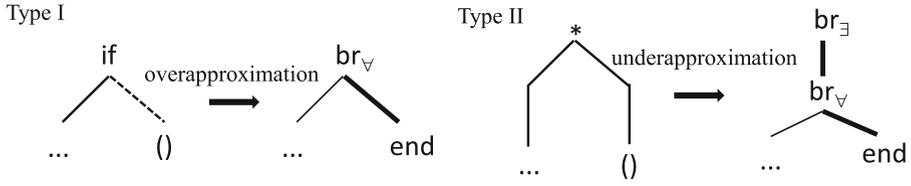


Fig. 6. Two types of paths in a counterexample

The refinement of abstraction types based on Type I (i.e., infeasible) paths can be performed in the same way as our previous work [16]. Thus, we focus below on how to deal with a Type II path.

4.1 Dealing with Type II Paths

Given a program P and a Type II path π , we first prepare fresh predicate variables R_1, \dots, R_k (called *separating predicates*), and replace each expression for random number generation $\text{let } r_i = *_{\text{int}} \text{ in } e_i$ with:¹

```
let  $r_i = *_{\text{int}}$  in assume( $R_i(r_i)$ );  $e_i$ .
```

Here, an expression $\text{assume}(\phi); e$ evaluates to e only if ϕ is **true**. Then, we instantiate R_i 's so that the following conditions hold.

- (C1) P has no longer an execution path along π .
- (C2) If the execution along π reaches $\text{let } r_i = *_{\text{int}} \text{ in assume}(R_i(r_i)); e_i$, there is at least one value for r_i such that $R_i(r_i)$ holds.

Condition C1 is for separating the path π at br_{\exists} node (recall Fig. 6; the problem of a Type II path has been that terminating/non-terminating paths are merged at br_{\exists} node). Condition C2 ensures that the paths separated from π are not empty. By C2, for example, an absurd assume statement like $\text{assume}(\text{false})$ is excluded out. We then add the instantiations of R_1, \dots, R_k to the abstraction types for r_1, \dots, r_k .

For the example

```
let  $x : \text{int}[] = *_{\text{int}}$  in let  $y : \text{int}[\lambda y. y \neq 0] = x$  in if  $y$  then loop() else ()
```

discussed above, we insert an assume statement as follows.

```
let  $x = *_{\text{int}}$  in assume( $R(x)$ ); let  $y = x$  in if  $y$  then loop() else ().
```

¹ Actually, we apply the replacement to each *instance* of $\text{let } r_i = *_{\text{int}} \text{ in } e_i$ along the execution path π , so that different assume conditions can be used for different instances of the same expression; we elide the details here.

Here, the Type II path π is the one that goes through the else-branch. Thus, a condition $R(x)$ that makes it infeasible is $x \neq 0$. As a result, $\lambda x.x \neq 0$ is added to the abstraction type for x .

We sketch below how to instantiate R_1, \dots, R_k . Using the technique of [16] condition (I) can be reduced to a set of non-recursive Horn clauses over predicate variables. Condition (II) is, on the other hand, reduced to constraints of the form

$$R_1(\tilde{x}_1) \wedge \dots \wedge R_n(\tilde{x}_n) \wedge C \Rightarrow \exists x.R(x) \wedge C'.$$

Thus, it remains to solve (non-recursive) existentially quantified Horn clauses [4]. To solve them, we first remove existential quantification by using a Skolemization-based technique similar to [4]. We prepare a linear template of Skolem function and move existential quantifiers out of universal quantifiers. For example, given

$$\forall r. (\exists \nu. \nu \leq 1 \wedge R(\nu)) \wedge \forall r. (R(r) \wedge \neg(r > 0) \Rightarrow \mathbf{false}),$$

we prepare the linear template $c_0 + c_1 r$ and transform the constraints into:

$$\exists c_0, c_1. \forall \nu, r. (\nu = c_0 + c_1 r \Rightarrow \nu \leq 1 \wedge R(\nu)) \wedge \forall (R(r) \wedge \neg(r > 0) \Rightarrow \mathbf{false}).$$

We then remove predicate variables by resolution, and get:

$$\forall \nu, r. \nu = c_0 + c_1 r \Rightarrow \nu \leq 1 \wedge \nu > 0$$

Finally, we solve constraints in the form of $\exists \tilde{x}. \forall \tilde{y}. \phi$ and obtain coefficients of linear templates that we introduced in the first step. We adopt the existing constraint solving technique based [24] on Farkas' Lemma. For the running example, we obtain $c_0 = 2, c_1 = 0$ as a solution of the constraints.

Now that we have removed existential quantification, we are left with non-recursive Horn clause constraints, which can be solved by using the existing constraint solving technique [23]. For the example above, we get

$$\forall \nu, r. (\nu = 2 \Rightarrow \nu \leq 1 \wedge R(\nu)) \wedge (R(r) \wedge \neg(r > 0) \Rightarrow \perp)$$

and obtain $R = \lambda \nu. \nu > 0$ as a solution.

5 Implementation and Experiments

We have implemented a non-termination verifier for a subset of OCaml, as an extension of MoChi [16], a software model checker for OCaml programs. We use HorSat [5] as the backend higher-order model checker, and Z3 [20] as the backend SMT solver. The web interface of our non-termination verification tool is available online [1]. We evaluated our tool by experiments on two benchmark sets: (1) test cases consisting of higher-order programs and (2) a standard benchmark set on non-termination of first-order programs [7, 19]. Both experiments were conducted on a machine with Intel Xeon E3-1225 V2 (3.20GHz, 16GB of memory)

Table 1. The result of the first benchmark set

Program	Cycle	Time (msec)	Program	Cycle	Time (msec)
loopH0	2	1,156	unfoldr_nonterm	3	13,540
indirect_e	1	111	passing_cond	2	9,202
indirectH0.e	1	112	inf_clos	2	12,264
foldr_nonterm	4	20,498	fib_CPS_nonterm	1	133
alternate	1	95	fixpoint_nonterm	2	168

with timeout of 60 seconds. The first benchmark set and an online demo page are available from our website [1].

Table 1 shows the result of the first evaluation. The columns ‘program’, ‘cycle’, and ‘time’ show the name of each test case, the number of CEGAR cycles, and the elapsed time (in milliseconds), respectively. For `foldr_nonterm`, we have used a different mode for a backend constraint solver; with the default mode, our verifier has timed out. All the programs in the first benchmark set are higher-order programs; so, they cannot be directly verified by previous tools. Our tool could successfully verify all the programs to be non-terminating (except that we had to change the mode of a backend constraint solver for `foldr_nonterm`).

We explain below two of the programs in the first benchmark set: `inf_clos` and `alternate`. The program `inf_clos` is:

```

is_zero n = (n = 0)          succ_app f n = f (n + 1)
f n cond = let b = cond n in if b then () else f n (succ_app cond)
main = f *int is_zero.

```

It has the following non-terminating reduction sequence:

```

main →* f 1 is_zero →* f 1 (succ_app is_zero) →* f 1 (succ_app2 is_zero)
→* f 1 (succ_appm is_zero) →* ...

```

Note that `succ_appm is_zero n` is equivalent to `n + m = 0`; hence `b` in the function `f` always evaluates to `false` in the sequence above. For proving non-termination, we need to reason about the value of the higher-order argument `cond`, so the previous methods for non-termination of first-order programs are not applicable.

The following program `alternate` shows the strength of our underapproximation.

```

f g h z = let x = *int in if x > 0 then g (f h g) else h (f h g)
proceed u = u ()    halt u = ()    main = f proceed halt ()

```

It has the following non-terminating reduction sequence:

```

main →* f proceed halt()
→* if 1 > 0 then proceed(f halt proceed) else ... →* f halt proceed()
→* if -1 > 0 then ... else proceed(f proceed halt) →* f proceed halt()
→* ...

```

Here, since the arguments g and h are swapped for each recursive call, the program does not terminate only if positive and negative integers are created alternately by \ast_{int} . Thus, the approach of Chen et al. [7] (which underapproximates a program by inserting assume statements and then uses a safety property checker to prove that the resulting program never terminates) would not be applicable. In our approach, by using the abstraction type $\mathbf{int}[\lambda x.x > 0]$ for x , f is abstracted to:

$$f\ g\ h\ z = \mathbf{br}_{\exists}\{\mathbf{true} \rightarrow \mathbf{let}\ x = \mathbf{true}\ \mathbf{in}\ \mathbf{br}_{\forall}\{x \rightarrow g(f\ h\ g)\}, \\ \mathbf{true} \rightarrow \mathbf{let}\ x = \mathbf{false}\ \mathbf{in}\ \mathbf{br}_{\forall}\{\neg x \rightarrow h(f\ h\ g)\}\}.$$

Thus, both branches of the if-expression are kept in the abstract program, and we can correctly conclude that the program is non-terminating.

For the second benchmark, we have borrowed a standard benchmark set consisting of 78 programs categorized as “known non-terminating examples” [7, 19]. (Actually, the original set consists of 81 programs, but 3 of them turned out to be terminating.) The original programs were written in the input language for T2 [2]; we have automatically converted them to OCaml programs. Our tool could verify 48 programs to be non-terminating in the time limit of 60 seconds. According to Larraz et al. [19], CPPINV [19], T2-TACAS [7], APROVE [6, 10], JULIA [22], and TNT [13] could verify 70, 51, 0, 8, and 19 programs respectively, with the same limit but under a different environment. Thus, our tool is not the best, but competitive with the state-of-the-art tools for proving non-termination of first-order programs, despite that our tool is not specialized for first-order programs. As for the comparison with T2-TACAS [7], our tool could verify 7 programs for which T2-TACAS failed, and ours failed for 10 programs that T2-TACAS could verify.

6 Related Work

Methods for disproving termination have recently been studied actively [7, 8, 13, 19]. Most of them, however, focused on programs having *finite* control-flow graphs with numerical data. For example, the state-of-the-art method by Larraz et al. [19] enumerates a strongly connected subgraph (SCSG), and checks whether there is a computation that is trapped in the SCSG using a SMT solver. Thus, it is not obvious how to extend those techniques to deal with recursion and higher-order functions. Note that unlike in safety property verification, we cannot soundly overapproximate the infinite control-flow graph of a higher-order program with a finite one.

Technically, the closest to our work seems to be the series of recent work by Cook et al. [7, 8]. They apply an underapproximation by inserting assume statements, and then either appeal to a safety property checker [7], or apply an overapproximation [8] to prove that the underapproximated program is non-terminating for all execution paths. A problem of their underapproximation [7] is that when an assume statement $\mathit{assume}(P)$ is inserted, all the computations such that $\neg P$ are discarded; so if P is wrongly chosen, they may overlook a

non-terminating computation present in the branch where $\neg P$ holds. As in the case for `alternate` discussed in Sect. 5, in the presence of higher-order functions, there may be no proper way for inserting assume conditions. In contrast, with our predicate abstraction, given a predicate P , we basically keep both branches for P and $\neg P$, and apply an underapproximation only if the satisfiability of P or $\neg P$ is not guaranteed (recall Fig. 1). In Cook et al.’s method [8], underapproximation cannot be applied after overapproximation, whereas under- and overapproximation can be arbitrarily nested in our method. Furthermore, although the framework of Cook et al. [8] is general, their concrete method can be applied to detect only non-termination in the form of lasso for programs having finite control-flow graphs. Harris et al. [14] also combine under- and overapproximation, but in a way different from ours: they use under- and overapproximation for disproving and proving termination respectively, not both for disproving termination.

There have also been studies on non-termination of term rewriting systems (TRS). Higher-order programs can be encoded into term rewriting systems, but the resulting analysis would be too imprecise. Furthermore, as mentioned in Sect. 1, the definition of non-termination is different.

Higher-order model checking has been recently applied to program verification [15, 16]. Predicate abstraction has been used for overapproximation for the purpose of safety property verification, but the combination of under- and overapproximation is new. Kuwahara et al. [18] have proposed a method for proving termination of higher-order programs; our new method for disproving termination plays a complementary role to that method.

The constraints generated in our CEGAR phase can be regarded as special instances of “existentially quantified Horn clauses” considered by Beyene et al. [4], where only acyclic clauses are allowed. Our constraint solving algorithm is specialized for the case of acyclic clauses. Incidentally, Beyene et al. [4] used existentially quantified clauses for verifying CTL properties of programs. Since non-termination can be expressed by the CTL formula $EG\neg terminated$, their technique can, in principle, be used for verifying non-termination. Like other methods for non-termination, however, the resulting technique seems applicable only to programs with finite control-flow graphs.

7 Conclusion

We have proposed an automated method for disproving termination of higher-order programs. The key idea was to combine under- and overapproximations by using predicate abstraction. By representing the approximation as a tree-generating higher-order program, we have reduced non-termination verification to higher-order model checking. The mixture of under- and overapproximations has also required a careful analysis of counterexamples, for determining whether and how under- or overapproximations are refined. We have implemented the proposed method and confirmed its effectiveness. Future work includes optimizations of the implementation and integration with the termination verifier [18].

Acknowledgments. We would like to thank Carsten Fuhs for providing us with their experimental data and pointers to related work, and anonymous referees for useful comments. This work was partially supported by Kakenhi 23220001 and 25730035.

References

1. MoCHI(Non-termination): Model Checker for Higher-Order Programs. <http://www-kb.is.s.u-tokyo.ac.jp/~kuwahara/nonterm/>
2. T2 temporal prover. <http://research.microsoft.com/en-us/projects/t2/>
3. Ball, T., Majumdar, R., Millstein, T., Rajamani, S.K.: Automatic predicate abstraction of C programs. In: PLDI 2001, pp. 203–213. ACM (2001)
4. Beyene, T.A., Popeea, C., Rybalchenko, A.: Solving existentially quantified horn clauses. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 869–882. Springer, Heidelberg (2013)
5. Broadbent, C., Kobayashi, N.: Saturation-based model checking of higher-order recursion schemes. In: CSL 2013. LIPIcs, vol. 23, pp. 129–148 (2013)
6. Brockschmidt, M., Ströder, T., Otto, C., Giesl, J.: Automated detection of non-termination and `NullPointerException` for Java bytecode. In: Beckert, B., Damiani, F., Gurov, D. (eds.) FoVeOOS 2011. LNCS, vol. 7421, pp. 123–141. Springer, Heidelberg (2012)
7. Chen, H.-Y., Cook, B., Fuhs, C., Nimkar, K., O’Hearn, P.: Proving nontermination via safety. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 156–171. Springer, Heidelberg (2014)
8. Cook, B., Fuhs, C., Nimkar, K., O’Hearn, P.W.: Disproving termination with over-approximation. In: FMCAD 2014, pp. 67–74. IEEE (2014)
9. Emmes, F., Enger, T., Giesl, J.: Proving non-looping non-termination automatically. In: Gramlich, B., Miller, D., Sattler, U. (eds.) IJCAR 2012. LNCS, vol. 7364, pp. 225–240. Springer, Heidelberg (2012)
10. Giesl, J., et al.: Proving termination of programs automatically with `aprove`. In: Demri, S., Kapur, D., Weidenbach, C. (eds.) IJCAR 2014. LNCS, vol. 8562, pp. 184–191. Springer, Heidelberg (2014)
11. Giesl, J., Thiemann, R., Schneider-Kamp, P.: Proving and disproving termination in the dependency pair framework. In: Deduction and Applications, No. 05431. Dagstuhl Seminar Proceedings (2006)
12. Graf, S., Saidi, H.: Construction of abstract state graphs with PVS. In: Grumberg, O. (ed.) CAV 1997. LNCS, vol. 1254, pp. 72–83. Springer, Heidelberg (1997)
13. Gupta, A., Henzinger, T.A., Majumdar, R., Rybalchenko, A., Xu, R.G.: Proving non-termination. In: POPL 2008, pp. 147–158. ACM (2008)
14. Harris, W.R., Lal, A., Nori, A.V., Rajamani, S.K.: Alternation for termination. In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 304–319. Springer, Heidelberg (2010)
15. Kobayashi, N.: Model checking higher-order programs. *J. ACM* 60(3) (2013)
16. Kobayashi, N., Sato, R., Unno, H.: Predicate abstraction and CEGAR for higher-order model checking. In: PLDI 2011, pp. 222–233. ACM (2011)
17. Kuwahara, T., Sato, R., Unno, H., Kobayashi, N.: Predicate abstraction and CEGAR for disproving termination of higher-order functional programs. Full version, available from the last author’s web page (2015)
18. Kuwahara, T., Terauchi, T., Unno, H., Kobayashi, N.: Automatic Termination Verification for Higher-Order Functional Programs. In: Shao, Z. (ed.) ESOP 2014 (ETAPS). LNCS, vol. 8410, pp. 392–411. Springer, Heidelberg (2014)

19. Larraz, D., Nimkar, K., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: Proving non-termination using Max-SMT. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 779–796. Springer, Heidelberg (2014)
20. de Moura, L., Bjørner, N.S.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008)
21. Ong, C.H.L.: On model-checking trees generated by higher-order recursion schemes. In: LICS 2006, pp. 81–90. IEEE (2006)
22. Spoto, F., Mesnard, F.: Étienne payet: a termination analyzer for Java bytecode based on path-length. *ACM Trans. Prog. Lang. Syst.* **32**(3), 8:1–8:70 (2010)
23. Unno, H., Kobayashi, N.: Dependent type inference with interpolants. In: PPDP 2009, pp. 277–288. ACM (2009)
24. Unno, H., Terauchi, T., Kobayashi, N.: Automating relatively complete verification of higher-order functional programs. In: POPL 2013, pp. 75–86. ACM (2013)