

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Kafle, B;Gange, G;Schachte, P;Søndergaard, H;Stuckey, PJ

Title:

Lightweight Nontermination Inference with CHCs

Date:

2021

Citation:

Kafle, B., Gange, G., Schachte, P., Søndergaard, H. & Stuckey, P. J. (2021). Lightweight Nontermination Inference with CHCs. Calinescu, R (Ed.) Pasareanu, CS (Ed.) Software Engineering and Formal Methods: 19th International Conference, SEFM 2021, Virtual Event, December 6–10, 2021, Proceedings, 13085 LNCS, pp.383-402. Springer. https://doi.org/10.1007/978-3-030-92124-8_22.

Persistent Link:

<https://hdl.handle.net/11343/324461>

Lightweight Nontermination Inference with CHCs

Bishoksan Kafle¹[0000–0001–5191–1216], Graeme Gange²[0000–0002–1354–431X],
Peter Schachte³[0000–0001–5959–3769],
Harald Søndergaard³[0000–0002–2352–1883], and Peter J.
Stuckey²[0000–0003–2186–0459]

¹ IMDEA Software Institute, Madrid

² Faculty of IT, Monash University

³ Computing and Information Systems, The University of Melbourne

Abstract. Non-termination is an unwanted program property (considered a bug) for some software systems, and a safety property for other systems. In either case, automated discovery of preconditions for non-termination is of interest. We introduce NTHORN, a fast lightweight non-termination analyser, able to deduce non-trivial sufficient conditions for non-termination. Using Constrained Horn Clauses (CHCs) as a vehicle, we show how established techniques for CHC program transformation and abstract interpretation can be exploited for the purpose of non-termination analysis. NTHORN is comparable in power to the state-of-the-art non-termination analysis tools, as measured on standard competition benchmark suites (consisting of integer manipulating programs), while typically solving problems an order of magnitude faster.

1 Introduction

Inference of preconditions for Non-Termination (NT) is of interest in program analysis, debugging and verification. For some systems, the possibility of non-termination is a bug. For other systems, premature termination is unwanted, so that non-termination becomes a safety property.

Non-termination is an archetypal undecidable problem. Assume P ranges over the set of programs expressible in some Turing complete language, and S ranges over (non-empty) sets of inputs to P . Then the problem of whether P fails to terminate on every $s \in S$ is undecidable, and not semi-decidable. This is true even when S is restricted to being a *finite* non-empty set. Moreover, a proof that P *terminates* on every element of some set S tells us nothing about P 's behaviour on (subsets of) S 's complement, and in particular it tells us nothing about non-termination. Obviously, absence of a proof of termination is no proof of the absence of termination.

Inferring sufficient conditions for NT is not always possible even for non-terminating programs. For instance, if the variable i ranges over $\mathbb{Z}$, the program `while( $i \geq 0$ )  $i = \text{nondet}()$` can be shown non-terminating by choosing always a *non-negative* value for i (*demonic non-determinism*), but no condition on the input i (apart from the trivial *false*) is sufficient to ensure non-termination.

<pre> int main(int a, int b) { if (b ≥ 100) b = b - 100; else b = 100 - b; if (b ≠ 0) //reachability query while (a ≥ 1) a = 2a + b; b = b - 1; assume(a ≥ 1); //reachability query return 0; } </pre>	<pre> c₁. init(a, b) ← true. c₂. if(a, b - 100) ← b ≥ 100, init(a, b). c₃. if(a, 100 - b) ← b ≤ 99, init(a, b). c₄. wh(a, b) ← b ≠ 0, if(a, b). c₈. en ← b ≠ 0, if(a, b). c₅. wh(a', b - 1) ← a ≥ 1, a' = 2a + b, a' ≥ 1, wh(a, b). c₆. ret ← a ≤ 0, wh(a, b). c₉. ex ← a ≤ 0, wh(a, b). c₇. ret ← b = 0, if(a, b). </pre>
--	---

Fig. 1. Original program (left) and CHC encoding of its reachable states (right)

Namely, the *loop* iteration does not depend on the initial value of *i*, only on non-deterministic assignments within the loop.

A central tool for proving non-termination is the notion of *recurrence set* [24], a set of runtime states from which flow of control cannot escape. The non-termination problem is complementary to proving termination; but while a safety violation can be witnessed by a finite trace, a failure to terminate has no such witness. Instead, a witness to non-termination is a path from an initial state to a recurrence set. So we are interested in finding some conditions on the initial state that ensure that such a path occurs.

Although finding preconditions for non-termination is a fundamental problem, it has received far less attention than other termination and non-termination problems (the work of Le et al. [32] is a notable exception). Our approach to the problem is inspired by Chen et al. [10] who reduce the problem to proving safety using a sequence of *reachability queries*. Let *bad states* be those that exit the program (or loop) under consideration, and *good states* those that get stuck in it. Then the problem is to infer preconditions that ensure all executions stay within *good states*. We achieve this in two steps: (i) compute a *necessary* precondition from the *bad states*, whose complement is a potential candidate for non-termination and (ii) refine the candidate with a *sufficient* precondition from the states that enter the program (loop). Our method is potentially applicable to large code bases, as it rests on relatively cheap program analysis and transformation. To our knowledge, the combination of necessary and sufficient precondition reasoning has not previously been applied to the task.

Before we present the approach formally, let us consider the example in Fig. 1 (left), a modified version of a program studied by Le et al. [32]. Assume the variables range over the full set $\mathbb{Z}$ of integers. Then the program fails to terminate if the input satisfies $(b > 100 \wedge a \geq 1) \vee (b < 100 \wedge a \geq 1)$ (equivalently $b \neq 100 \wedge a \geq 1$). This is because *a* and *b* are always positive when entering the loop, and the loop condition $a \geq 1$ is always maintained, as *a* increases at

a higher rate than $\mathbf{b}$ decreases. Automatic derivation of these preconditions is challenging for at least three reasons:

- (i) The desired result is a disjunction of linear constraints—so we need the ability to express disjunctive information.
- (ii) Abstract interpretation working forward or backward from the goal such as $(\mathbf{a} \leq 0)$ derives *top* as invariant for the loop. That is, without a more sophisticated approach, we lose critical information about $\mathbf{a}$ and $\mathbf{b}$.
- (iii) We use over- and under-approximations to obtain sound and precise results, since the precondition must ensure that all traces enter the loop but none exit. While the first part (*all traces enter the loop*) requires under-approximation, the second part (*none exit*) can be achieved by negating an over-approximation that exits the loop.

We address Challenge (i) via partial evaluation or control flow refinement, creating a finite number of versions of each predicate—this is essential for deriving disjunctive invariants.

Challenge (ii) is addressed via forward and backward abstract interpretation, together with constraint specialisation. Challenge (iii) is addressed by refining over-approximations with under-approximations (§4).

While our approach is inspired by the ideas behind HIP-TNT+ [32], it rests entirely on simpler methods from transformation-based program analysis of Constrained Horn Clauses: control flow refinement via *partial evaluation* [17], *constraint specialisation* [27], and *clause splitting* [19]. We make these contributions:

- We reduce the problem of precondition inference for non-termination to precondition inference for safety using a sequence of *reachability queries* inspired by the work of Chen et al. [10], which reduces proving non-termination to proving safety.
- We present an enhanced modular algorithm that combines under- and over-approximation techniques based on abstract interpretation and program transformation to derive sound and precise preconditions. It includes a novel mechanism of deriving a more general precondition through iterative refinement, which comes with refined termination criteria (Sect. 4).
- Our method uniformly handles non-linear clauses (arising from modelling function calls, recursion, and nested loops) over linear integer arithmetic.
- A proof of concept is implemented in the tool NTHORN, and we present experiments which show that our prototype implementation is competitive with state-of-the-art tools for automated proof of non-termination (Sect. 5).

2 CHCs, Recurrence Sets and Preconditions

We represent a program as a set of Constrained Horn clauses (CHCs). This is convenient for representing imperative programs and properties such as *reachability queries* in a uniform way. Analysis of the program and its properties is then done by analysing the corresponding CHCs. The translation of imperative programs to CHCs is standard [16, 23, 25, 36] so we omit the details. From here on, by ‘program’ we mean a program’s CHC encoding, unless otherwise stated.

Constrained Horn clauses (CHCs). An *atom* is a formula $p(\mathbf{x})$ with p a predicate symbol and $\mathbf{x}$ a tuple of arguments. A CHC is a first-order formula written $p_0(\mathbf{x}_0) \leftarrow \varphi, p_1(\mathbf{x}_1), \dots, p_k(\mathbf{x}_k)$, with φ a finite conjunction of quantifier-free *constraints* on variables $\mathbf{x}_i$ wrt. some constraint theory $\mathbb{T}$, and $p_i(\mathbf{x}_i)$ are atoms. A clause $p_0(\mathbf{x}_0) \leftarrow \varphi_1 \vee \varphi_2, \beta$ with a disjunctive constraint is rewritten as $p_0(\mathbf{x}_0) \leftarrow \varphi_1, \beta$ and $p_0(\mathbf{x}_0) \leftarrow \varphi_2, \beta$. A *constrained fact* is a clause of the form $p_0(\mathbf{x}_0) \leftarrow \varphi$, where φ is a constraint. A clause is *linear* if $k \leq 1$, otherwise *non-linear*. A *program* is linear if all of its clauses are linear.

Given a clause set P , we assign a unique identifier to each clause in P . Further, we assume the theory $\mathbb{T}$ is equipped with a decision procedure and a projection operator, and that it is closed under negation. The notation $\varphi|_V$ represents the constraint formula φ projected onto variable set V . $\varphi \models_{\mathbb{T}} \psi$ (or equivalently $\models_{\mathbb{T}} \varphi \Rightarrow \psi$) says that φ entails ψ over $\mathbb{T}$. We write $P \vdash_{\mathbb{T}} A$ when an atom A is derivable from program P wrt. an axiomatisation of $\mathbb{T}$. We omit the subscript $\mathbb{T}$ when it is clear from the context.

CHC encoding. Fig. 1 (right) shows the CHC representation of the example, encoding the reachable states. The clause c_1 specifies the *initial states* of the program via the predicate *init* which is always reachable. Similarly, c_2 and c_3 encode the reachability of the second *if* condition via the predicate *if*. Clauses c_4 and c_5 encode the reachability of the *while* loop via the predicate *wh*. Clause c_4 states that the loop is reachable if *if* is reachable, while c_5 states that the loop is (re-)reachable from the end of its own body (recursive case). Clauses c_6 and c_7 encode the return from the program. Clause c_6 states that the program terminates if $\mathbf{a} < 0$ upon loop exit, while clause c_7 states that the program terminates when $\mathbf{b} = 0$ and the control does not satisfy the condition of the second *if*. The coloured clauses are not part of the program, but are added to aid the analysis. We employ two *special predicates* *en* and *ex* which respectively encode the states entering the loop and exiting the loop or the program. Note that multiple clauses for these predicates are possible given multiple loop entries/exits.

Definition 1 (Initial clauses and nodes). Let P be a program with a distinguished predicate p^I which we call the *initial predicate*. The *constrained facts* of the form $p^I(\mathbf{x}) \leftarrow \theta$ are the *initial clauses* of P . We extend the term “initial predicate” and use the symbol p^I to refer also to renamed versions of the initial predicate that arise during clause transformations.

For the program in Fig. 1, *init* is the initial predicate and $\text{init}(\mathbf{a}, \mathbf{b}) \leftarrow \text{true}$ is the initial clause. We shall assume integer programs, that is, all variables take integer values. Let $\text{val} : V \rightarrow \mathbb{Z}$ map variables to their values. We overload val to also map a tuple of variables to the tuple of their values.

A set of CHCs defines a *transition system*, defined as follows (in the following we shall freely interchange these concepts):

Definition 2 (Transition system). A transition system of a linear program P is a tuple $\mathcal{T} = (S, R, I)$, where

- $S = Pd \times 2^{\mathbb{Z}^{|V|}}$ is the set of states where Pd is the set of predicates of P (including *false*) and V is a finite set of program variables.
- $R \subseteq S \times S$ is a transition relation. There is a transition from $(p, \text{val}(\mathbf{x}))$ to $(p', \text{val}(\mathbf{x}'))$ labelled by c if there is a clause $p'(\mathbf{x}') \leftarrow \varphi \wedge p(\mathbf{x}) \in P$ with identifier c and if $\text{val}(\mathbf{x}) \models \varphi$ then $\text{val}(\mathbf{x}') \models \varphi$.
- $I \subseteq S$ is a set of initial states.

Non-termination and recurrence set. A transition system $\mathcal{T} = (S, R, I)$ is non-terminating iff there is an infinite sequence $s_0, s_1, s_2 \dots$, of states, with $s_0 \in I$ and $(s_i, s_{i+1}) \in R$. Non-termination of a relation R is witnessed by the existence of an (*open*) *recurrence set* [24]: a non-empty set $\mathcal{G}$ of states such that (i) $\mathcal{G}$ contains an initial state and (ii) each $s \in \mathcal{G}$ has a successor in $\mathcal{G}$. A program is non-terminating iff its transition system contains a recurrence set [24].

Chen et al. [10] extend the notion to *closed recurrence set* which facilitates automation using established techniques like abstract interpretation or model checking. A *closed recurrence set* is an open recurrence set $\mathcal{G}$ with the additional property that, for each $s \in \mathcal{G}$, all of its successors are in $\mathcal{G}$. Our method relies on closed recurrence sets to automate the reasoning.

Preconditions. Given a transition system $\mathcal{T} = (S, R, I)$, we define functions $\text{pre} : 2^S \rightarrow 2^S$, $\text{post} : 2^S \rightarrow 2^S$ and $\widetilde{\text{pre}} : 2^S \rightarrow 2^S$ as follows.

- $\text{post}(S') = \{s' \in S \mid \exists s \in S' : (s, s') \in R\}$ returns the set of states having at least one of their predecessors in the set $S' \subseteq S$;
- $\text{pre}(S') = \{s \in S \mid \exists s' \in S' : (s, s') \in R\}$ returns the set of states having at least one of their successors in the set $S' \subseteq S$;
- $\widetilde{\text{pre}}(S') = \{s \in S \mid \forall s' \in S : (s, s') \in R \Rightarrow s' \in S'\}$ returns the set of states having all of their successors in the set $S' \subseteq S$.

With these functions, we can now state precondition inference problems.

Invariants. Given a transition system $\mathcal{T} = (S, R, I)$ and a set of initial states $S' \subseteq S$, the *invariant inference problem* consists of inferring the set of reachable states from S' as $\text{inv}(\mathcal{T}, S') = \text{lfp } \lambda X. S' \cup \text{post}(X)$.

Necessary preconditions. Given a transition system $\mathcal{T} = (S, R, I)$ and a goal set $S' \subseteq S$ of states, the *necessary precondition inference problem* consists of inferring the set of *initial states* as $\text{nec_pre}(\mathcal{T}, S') = \text{lfp } \lambda X. S' \cup \text{pre}(X)$, which guarantees that some of its executions will stay in S' .

Sufficient preconditions. Given a transition system $\mathcal{T} = (S, R, I)$ and a goal set $S' \subseteq S$ of states, the *sufficient precondition inference problem* consists of inferring the set of *initial states* as $\text{suf_pre}(\mathcal{T}, S') = \text{gfp } \lambda X. S' \cap \widetilde{\text{pre}}(X)$, which guarantees that all of its executions will stay in S' .

Note that the functions inv , nec_pre and suf_pre are not computable in general. Therefore, approximations of these functions are computed instead, which

provide “one-sided” guarantees. The state-of-the-art techniques for computing `nec_pre` use over-approximations based on abstract interpretation [12] and are given in [2, 3, 13, 30, 37], while that for computing `suf_pre` use backward under-approximation or negation of some necessary preconditions [28, 34, 35, 37]. In addition, these techniques can profitably be combined with CHC transformations such as [15, 17, 21, 27] to enhance the precision of these analyses.

Example 1. Our approach derives preconditions as follows. First $\lambda_{\text{en}} = b \neq 100$ and $\lambda_{\text{ex}} = (b \geq 101 \wedge a \leq 0) \vee (b \leq 99 \wedge a \leq 0) \equiv a \leq 0 \wedge b \neq 100$ are found as necessary preconditions for the reachability of `en` and `ex`, resp. Now $\lambda = \lambda_{\text{en}} \wedge \neg \lambda_{\text{ex}} \equiv a \geq 1 \wedge b \neq 100$ represents the initial states that might reach the loop entry but not the loop exit. We consider λ a candidate for sufficient precondition, using that to strengthen the initial clause to `init(a,b) ← a ≥ 1 ∧ b ≠ 100`. Then using backward under-approximation [34] from the goal `en`, we derive $a \geq 1 \wedge b \neq 100$ as a sufficient precondition for the reachability of `en`—which happens to be the optimal precondition for non-termination in this case. If we just used under-approximations without strengthening the initial clauses, we would obtained only $a \geq 1 \wedge b > 100$ or $a \geq 1 \wedge b < 100$, and not both. $\square$

3 CHC Transformations and Their Roles in Non-Termination Analysis

We now summarise common CHC transformations that we use, such as partial evaluation, constraint specialisation and clauses splitting. We highlight their role in non-termination analysis. They are *goal preserving* transformations (or *specialisations*): given a program P and a goal A , the transformation of P wrt. to the goal A yields another program P' such that $P \vdash_{\mathbb{T}} A \text{ iff } P' \vdash_{\mathbb{T}} A$. In our setting, the goals are `en` and `ex`. Informally, we produce a specialised version of P that preserves the derivations of `en` and `ex`, but not necessarily other goals.

1. Partial evaluation (PE). PE of a set P of CHCs wrt. goal A produces a specialised version of P preserving only those derivations that are relevant for deriving A . It produces a polyvariant specialisation, which is essential for deriving disjunctive information. The partial evaluation algorithm utilised here is an instantiation of the algorithm given in [20], which is parameterised by an “unfolding rule” unfold_P and an abstraction operation abstract_{ψ} .

The unfolding rule unfold_P takes a set S of constrained facts and “partially evaluates” each element of S , using the following unfolding rule. For each $(p(\mathbf{x}) \leftarrow \theta) \in S$, first construct the set of clauses $p(\mathbf{x}) \leftarrow \psi', \beta'$ where $p(\mathbf{x}) \leftarrow \psi, \beta$ is a clause in P , and ψ', β' is obtained by unfolding $\psi \wedge \theta, \beta$ by selecting atoms so long as they are deterministic (atoms defined by a single clause) and is not a call to a recursive predicate, and ψ' is satisfiable in $\mathbb{T}$. unfold_P returns the set of constrained facts $q(\mathbf{y}) \leftarrow \psi'|_{\mathbf{y}}$ where $q(\mathbf{y})$ is an atom in β' .

Given an initial set S_0 , the closure of the unfold_P operation can be obtained as $\text{lfp } \lambda S. S_0 \cup \text{unfold}_P(S)$. It is not computable in general; so instead we compute a set $\text{cfacts}(S_0) = \text{lfp } \lambda S. S_0 \cup \text{abstract}_{\psi}(\text{unfold}_P(S))$, where the abstraction

<pre> init_1(a, b) ← b ≥ 101. init_2(a, b) ← b ≤ 99. if_2(a, b) ← b ≥ 1, init_1(a, b + 100). if_2(a, b) ← b ≥ 1, <u>init_2(a, 100 - b)</u>. wh_3(a, b) ← b ≥ 1, a ≤ 0, if_2(a, b). ex ← a ≤ 0, wh_3(a, b). wh_3(a, b) ← b ≤ -1, a ≤ 0, if_1(a, b). </pre>	<pre> init_1(a, b) ← <u>a ≤ 0</u>, b ≥ 101. init_2(a, b) ← <u>a ≤ 0</u>, b ≤ 99. if_2(a, b) ← <u>a ≤ 0</u>, b ≥ 1, <u>init_1(a, b + 100)</u>. if_2(a, b) ← <u>a ≤ 0</u>, b ≥ 1, <u>init_2(a, 100 - b)</u>. wh_3(a, b) ← a ≤ 0, b ≥ 1, if_2(a, b). ex ← a ≤ 0, b ≥ 1, wh_3(a, b). </pre>
---	---

Fig. 2. PE of Fig. 1 wrt. **ex** (left) and its CS version (right) with inferred constraints underlined. The last clause on LHS is eliminated since its body is strengthened to *false*.

operation abstract_Ψ performs property-based abstraction [23] wrt. a finite set of properties Ψ . A set of clauses is then generated by applying unfold_P to each $\text{cfacts}(S_0)$ and renaming the predicates in the resulting clauses according to the different versions produced by abstract_Ψ . We refer to [21] for more details.

2. Constraint specialisation (CS). A CS of P wrt. goal A and set Ψ of properties [27] is a transformation in which each clause $(p(\mathbf{x}) \leftarrow \varphi, \beta) \in P$ is replaced by $p(\mathbf{x}) \leftarrow \varphi, \underline{\psi}, \beta$ (the difference from the original underlined), where $(p(\mathbf{x}) \leftarrow \psi) \in \Psi$, such that the resulting set of clauses preserves the derivation of A . As a result, all paths that are irrelevant for deriving A can be eliminated.

Example 2 (Continued from Example 1). The program in Fig. 2 (left) is obtained by PE of Fig. 1 wrt. **ex**. Observe that the recursive clause **wh** is effectively eliminated, as it cannot contribute to a derivation of **ex**. The constraint in the initial clauses $b \geq 101 \vee b \leq 99 \equiv b \neq 100$ is a necessary precondition for the reachability of **ex**. This is further strengthened to $(a \leq 0 \wedge b \geq 101) \vee (a \leq 0 \wedge b \leq 99) \equiv a \leq 0 \wedge b \neq 100$ with constraint specialisation wrt. **ex** to Fig. 2 (left), which propagates $a \leq 0$ from the goal **ex** and $b \geq 101 \vee b \leq 99$ from the constrained facts to other clauses, resulting in the program on the right. Similarly, we obtain $b \neq 100$ as a necessary precondition for the reachability of **en**. $\square$

3. Clause splitting. Given a clause $(p(\mathbf{x}) \leftarrow \varphi, \beta) \in P$ and a set Ψ of properties, clause splitting replaces the clause by $p(\mathbf{x}) \leftarrow \varphi, \underline{\psi}, \beta$ and $p(\mathbf{x}) \leftarrow \varphi, \neg \underline{\psi}, \beta$, producing P' (new constraints are underlined), where $(p(\mathbf{x}) \leftarrow \psi) \in \Psi$. This embodies case splits, allowing case-based reasoning. Fioravanti et al. [19] use a related technique for splitting clauses to achieve deterministic programs. Unlike the previous transformations, it is goal independent, that is, for all atoms A of P , $P \vdash_{\mathbb{T}} A$ iff $P' \vdash_{\mathbb{T}} A$.

Common to these transformations is the set Ψ of properties, which determine the quality of the resulting clauses. Soundness of CS also depends on the choice of Ψ . Though the above program transformation techniques are generic for CHCs and are taken from the literature, application or program specific choices of Ψ that we describe next make them surprisingly effective in practice. In addition

$\left\{ \begin{array}{l} p(a, b) \leftarrow a = b. p(a, b) \leftarrow a > 0. \\ p(a, b) \leftarrow a > b. \\ p(a, b) \leftarrow b > a. \end{array} \right\}$ $\begin{array}{l} \text{ex} \leftarrow b > a, p_1(a, b). \\ p_1(a, b) \leftarrow \text{false}. \end{array}$	$\left\{ \begin{array}{l} p(a, b) \leftarrow a = b. \\ p(a, b) \leftarrow a > b. \end{array} \right\}$ $\begin{array}{l} p(a, b) \leftarrow a = b. \\ p(a + b, b) \leftarrow a > 0, p(a, b). \\ \text{ex} \leftarrow a < b, p(a, b). \end{array}$
---	---

Fig. 4. PE of Fig. 3 wrt. ex ; respective Ψ s are shown in upper part

to this, our contribution is to put these transformation together and apply them for inferring preconditions for NT, which has not been considered before.

We now discuss the specific choices we make for each of the transformations and illustrate the differences with other choices using the synthetic but representative example shown in Fig. 3.

$$\begin{array}{l} c_1. p(a, b) \leftarrow a = b. \\ c_2. p(a + b, b) \leftarrow a > 0, p(a, b). \\ c_3. \text{ex} \leftarrow a < b, p(a, b). \end{array}$$

Fig. 3. Synthetic example

For PE. The set Ψ contains the following constrained facts, generated from each clause $p(\mathbf{x}) \leftarrow \varphi, p_1(\mathbf{x}_1), \dots, p_n(\mathbf{x}_n) \in P$.

- $p(\mathbf{x}) \leftarrow \varphi|_{\mathbf{x}}$ and for each $z \in \mathbf{x}$, $p(\mathbf{x}) \leftarrow \varphi|_{\{z\}}$
- for $1 \leq i \leq n$, $p_i(\mathbf{x}_i) \leftarrow \varphi|_{\mathbf{x}_i}$ and for each $z \in \mathbf{x}_i$, $p_i(\mathbf{x}_i) \leftarrow \varphi|_{\{z\}}$.

The effect of property-based abstraction using this choice for Ψ is to create a finite number (at most $2^{|\Psi|}$) of versions of a predicate for different call contexts and answer constraints. This choice of Ψ , *obtained syntactically from the program*, has been found to provide a good balance of speed and precision.

Example 3. Fig. 4 shows PE programs for the program P in Fig. 3 wrt. ex with two different choices of the set of properties Ψ . Ψ in Fig. 4 (left) are computed as described above, while on the right are computed as follows: $\Psi = \bigcup_{(p(\mathbf{x}) \leftarrow \varphi, B) \in P} \{p(\mathbf{x}) \leftarrow \varphi|_{\mathbf{x}}\}$. The purpose here is to show that the choice of Ψ is important in getting a right specialisation. The program on the left is an empty program since there is a vacuous base case ($p_1(a, b) \leftarrow \text{false}$), while the program on the right is identical to the original (no specialisation was performed). $\square$

For CS. The properties Ψ have to be invariants for the program to produce sound transformation. They can be obtained e.g., via forward (from the constrained facts) or backward (from the goal) abstract interpretation or their combination [3]. In our case, they are obtained from forward abstract interpretation of the query-answer transformed program [27]. Ψ thus obtained *analysing the program* produces sound transformation, which is also found to be precise.

Example 4. Forward analysis of the program in Fig. 3 yields $p(a, b) \leftarrow a \geq b$ as invariant for $p(a, b)$. This is because (i) if c_2 is not taken then we have $a = b$ from

c_1 and obviously $a \geq b$ holds, (ii) if c_2 is taken then $a > 0$ is maintained since a is incremented by b in each iteration and we initially had $a = b$. Since b is not modified in c_2 , $a \geq b$ holds. We now use $\Psi = \{p(a, b) \leftarrow a \geq b\}$ to specialise the program in Fig. 3 wrt. ex , obtaining the clauses c_1 , $p(b + c, b) \leftarrow a > 0, \underline{a \geq b}, p(c, b)$ and $\text{ex} \leftarrow a < b, \underline{a \geq b}, p(a, b)$. Note that the last clause is trivially satisfied. Instead of applying forward or backward analysis in isolation, applying forward-backward analysis will immediately detect that $p(a, b) \leftarrow \text{false}$, and the subsequent specialisation using the result yields an empty program. $\square$

For clause splitting. We describe some heuristics specific to *(non-)termination analysis*, requiring separation of terminating and non-terminating computations. The targets are recursive clauses (loops) $p(\mathbf{x}') \leftarrow \varphi, p(\mathbf{x})$. (i) Given a loop, a *potential ranking function* for the loop is an expression $e(\mathbf{x})$ over variables $\mathbf{x}$ which is *non-negative* (bounded from below) but not necessarily *decreasing* from $p(\mathbf{x})$ to $p(\mathbf{x}')$. In this case, we choose the property $\{p(\mathbf{x}) \leftarrow e(\mathbf{x}) > e(\mathbf{x}')\}$ (see Example 5). (ii) The property $\{p(\mathbf{x}) \leftarrow x \geq 0 \mid x \in \mathbf{x}, \models_{\mathbb{T}} \varphi \wedge x \geq 0\}$ is useful when we have non-deterministic branches or assignments; but care needs to be taken to control the blow-up of clauses.

Example 5. Taking $p(a, b) \leftarrow a = b$ to be the initial clause of the program in Fig. 3, the program does not terminate; e.g., for input $a = 1, b = 1$. Below we apply clause splitting which reveals which clause causes non-termination. The expression a is a potential ranking function for the loop, as it is non-negative and not necessarily decreasing. We derive $\Psi = \{p(a, b) \leftarrow a + b < a\}$. Now splitting c_2 with Ψ yields c_{2a} : $p(a + b, b) \leftarrow a > 0, \underline{b < 0}, p(a, b)$ and c_{2b} : $p(a + b, b) \leftarrow a > 0, \underline{b \geq 0}, p(a, b)$ (the new constraints are underlined). Such a splitting guarantees that every infinite run of the program must use c_{2b} as suffix. This information can be exploited by (non-)termination analysers. $\square$

4 An Algorithm for Conditional Non-Termination

We now present an algorithm for inferring sufficient preconditions for non-termination. The main method is Alg. 1. As a program can only get stuck in loops or recursive code, and the translation to CHCs replaces loops with recursion, the analysis focuses on the recursive strongly connected components (SCCs) in the CHC dependency graph. As the first step, we compute the SCCs of the input set P of CHCs. Each component is a set of (non-constraint) predicates, which is either non-recursive or a set of (possibly mutually) recursive predicates. The algorithm for computing SCCs returns the components in topologically sorted order $S_1, \dots, S_n$, such that for each S_i , no predicate in S_i depends on any predicate in S_j where $j > i$. Then it annotates the program with appropriate *reachability queries* and computes a sufficient precondition for each annotated program using Alg. 2 in a modular way. These preconditions are combined disjunctively to yield the overall result. The function `annotate_program( $P, C$ )` inserts two sets of clauses to P given an SCC C as follows.

$$\begin{aligned} & \{\mathbf{en} \leftarrow \varphi|_{\text{vars}(\beta)}, \beta \mid (H \leftarrow \varphi, \beta) \in P, \text{pred}(H) \in C, \forall a (a \in \text{pred}(\beta) \Rightarrow a \notin C)\} \\ & \cup \{\mathbf{ex} \leftarrow \varphi|_{\text{vars}(\beta)}, \beta \mid (H \leftarrow \varphi, \beta) \in P, \text{pred}(H) \notin C, \forall a (a \in \text{pred}(\beta) \Rightarrow a \in C)\} \end{aligned}$$

where $\text{pred}(H)$ is the predicate symbols of H and $\text{pred}(\beta)$ is the set of predicate symbols in β . The special clauses for $\mathbf{en}$ encode the reachability of C in P while the clauses for $\mathbf{ex}$ encode the exit condition of C in P . We explore all SCCs with the aim of obtaining a more general precondition but for proving non-termination it suffices to find a *non-trivial* precondition for an SCC.

Example 6. For the program in Fig. 1, there is a single SCC, namely the one containing the predicate $\{\mathbf{wh}\}$. The clauses c_8 and c_9 respectively encode the reachability of the entry and exit of this SCC. $\square$

Alg. 2 takes as input and annotated program P (with clauses for $\mathbf{en}$ and $\mathbf{ex}$); it returns a sufficient precondition for non-termination (a linear constraint over $\mathbb{T}$ in DNF). Recall that we derive preconditions in two steps: (i) compute a necessary precondition from the *bad states*, encoded by the predicate $\mathbf{ex}$, whose complement is a potential candidate for non-termination and (ii) refine the candidate with a sufficient precondition from the states that enter the loop, encoded by the predicate $\mathbf{en}$. Let us first focus on step (i), the generation of necessary preconditions, using backward over-approximating analysis. Since these conditions need to be negated to derive candidates, their precision is important. It is well known that program specialisations, possibly applied iteratively [28] can enhance (refine) precision of such analysis. A disadvantage of this is the blind refinement of states possibly exiting the loop without knowing its frontier with the states entering it. This misses opportunities to avoid redundant computation as well as to guide the refinement process at an early stage. We therefore choose to maintain two over-approximations, namely of the states entering the loop ($\lambda_{\mathbf{en}}$) and of the states exiting it ($\lambda_{\mathbf{ex}}$). We iteratively refine these (See Alg. 2). This enables the use of over-approximating analyses which are more developed than their under-approximating counterparts. Further, step (ii) is only applied to strengthen the candidate $\lambda = \lambda_{\mathbf{en}} \wedge \neg\lambda_{\mathbf{ex}}$ to a sufficient condition. This is achieved by applying under-approximating analysis to $\text{replace_init}(P, \lambda)$ instead of the original P , so as to retain refined initial condition λ derived from the analysis of step (i).

Definition 3 ($\text{replace_init}(P, \lambda)$). Let P be a program and λ a constraint over $\mathbb{T}$. The function replace_init returns clauses of P by replacing the initial clauses $\{(p^I(\mathbf{x}) \leftarrow \theta_i) \mid 1 \leq i \leq k\}$ by $\{(p^I(\mathbf{x}) \leftarrow \lambda)\}$.

Alg. 2. The variable σ_{nt} accumulates the result and is initialised to *false*. φ_{old} keeps track of the *initial states* that could reach both $\mathbf{en}$ and $\mathbf{ex}$; it is initialised to *true* (line 3). The following operations are carried out within the *while* loop. The formulae λ_{en} and λ_{ex} (line 5 and 6) represent the set of initial states that can reach $\mathbf{en}$ and $\mathbf{ex}$, resp., and are computed using the method nec_pre . The

Algorithm 1 Inferring sufficient precondition for non-termination of a program

```
1: Input: A program  $P$ 
2: Output: Sufficient precondition  $\sigma_{nt}$  for non-termination of  $P$ .
3: Initialisation:  $\sigma_{nt} \leftarrow false$ ;
4:  $S_1, \dots, S_n \leftarrow$  topologically sorted SCCs of  $P$ 
5: for  $i = 1 \dots n$  do
6:   if (recursive  $S_i$ ) then
7:      $P_{an} \leftarrow \text{annotate\_program}(P, S_i)$ ;
8:      $\sigma_{nt} \leftarrow \sigma_{nt} \vee \text{precond\_scc}(P_{an})$  ▷ Alg. 2
9: return  $\sigma_{nt}$ 
```

Algorithm 2 `precond_scc`: Inferring sufficient precondition wrt. a SCC

```
1: Input: Program  $P$  annotated with clauses for en and ex
2: Output: Sufficient precondition  $\sigma_{nt}$  for NT.
3: Initialisation:  $\sigma_{nt} \leftarrow false$ ;  $\varphi_{old} \leftarrow true$ ;
4: while true do
5:    $\lambda_{en} \leftarrow \text{nec\_pre}(P, \text{en})$ ;
6:    $\lambda_{ex} \leftarrow \text{nec\_pre}(P, \text{ex})$ ;
7:   if  $\lambda_{en} \equiv false$  then ▷ no state reaches en
8:     return  $\sigma_{nt}$ 
9:    $\varphi_{new} \leftarrow \lambda_{en} \wedge \lambda_{ex}$ ; ▷ states that may both reach en and ex
10:   $\lambda \leftarrow \lambda_{en} \wedge \neg \lambda_{ex}$ ; ▷ states that may reach en but not ex
11:   $P_1 \leftarrow \text{replace\_init}(P, \lambda)$ ; ▷ Def. 3
12:   $\sigma_{en} \leftarrow \text{suf\_pre}(P_1, \text{en})$ ; ▷ strengthen  $\lambda$  to sufficient condition
13:   $\sigma_{nt} \leftarrow \sigma_{nt} \vee \sigma_{en}$ ;
14:  if  $\varphi_{new} \equiv false$  or  $\varphi_{old} \models_{\mathbb{T}} \varphi_{new}$  then
15:    ▷  $\lambda_{en}$  may not reach ex, or states reaching both en and ex don't shrink.
16:    return  $\sigma_{nt}$ 
17:   $\varphi_{old} \leftarrow \varphi_{new}$ ;
18:   $P \leftarrow \text{constrain\_init}(P, \varphi_{new})$  ▷ refine  $P$ , Def. 4
```

algorithm returns when (i) no initial state can reach **en** ($\lambda_{en} \equiv false$) (line 8), or (ii) the initial states satisfying $\varphi_{new} = \lambda_{en} \wedge \lambda_{ex}$ that can reach both **en** and **ex** amount to *false*, or the set of initial states does not shrink further from its previous value φ_{old} (line 14). The set of states captured by φ_{new} is an over-approximation. The algorithm aims to reduce the slack as much as possible, to be able to separate terminating traces from non-terminating ones. To this end it (i) constructs a revised program from P focusing only on the shared region and (ii) shrinks either of the regions $(\lambda_{en}, \lambda_{ex})$ via iterative specialisations. We construct the revised program as follows.

Definition 4 (`constrain_init`(P, φ)). Let P be a program and φ a constraint over $\mathbb{T}$. `constrain_init` returns the clauses of P by replacing the initial clause set $\{(p^I(\mathbf{x}) \leftarrow \theta_i) \mid 1 \leq i \leq k\}$ by the set $\{(p^I(\mathbf{x}) \leftarrow \varphi \wedge \theta_i) \mid 1 \leq i \leq k\}$.

Proposition 1. `constrain_init`(P, φ) is an under-approximation of P .

Proof (Sketch). $P_1 = \text{constrain_init}(P, \varphi)$ contains exactly the same clauses as P , except for the *initial clauses*, which are possibly constrained. Hence for all atoms A , if $P_1 \vdash_{\mathbb{T}} A$ then $P \vdash_{\mathbb{T}} A$. That is, P_1 is an under-approximation of P .

Since non-termination is preserved by under-approximation [10], we need to ensure that the precondition does as well. This is in fact the case given that these program only differ in their initial clauses. Thus, any initial state that definitely reaches **en** and stays in the loop of P_1 also does the same in P . Before formally stating this property, let us first define (in terms of CHCs) what it means for a program P to have φ as a sufficient precondition for NT.

Definition 5 (Sufficient precondition for CHCs). Let P be a program annotated with appropriate *reachability queries* (for **en** and **ex**) as described above and φ a constraint over $\mathbb{T}$. Let $P_1 = \text{replace_init}(P, \varphi)$. Then we say φ is a sufficient precondition for NT of P if $\varphi \rightarrow (P_1 \vdash_{\mathbb{T}} \text{en} \wedge P_1 \not\vdash_{\mathbb{T}} \text{ex})$.

Proposition 2 (Lifting sufficient conditions). If φ is a sufficient precondition for non-termination of $\text{constrain_init}(P, \sigma)$ (for some σ) then it is also a sufficient precondition of P .

Proof (Sketch). Let $P_1 = \text{constrain_init}(P, \sigma)$. Since P_1 and P have identical clauses except for the *initial clauses*, $\text{replace_init}(P_1, \varphi)$ and $\text{replace_init}(P, \varphi)$ yield identical clauses. So φ is also a sufficient precondition of P (Def. 5). $\square$

Note the initial states satisfying the formula $(\lambda_{en} \wedge \neg \lambda_{ex})$ may reach **en** but definitely not **ex**, so they are seen as potential candidates for non-termination. The candidates are then strengthened to sufficient preconditions (σ_{en}) using the method `suf_pre` (line 12). If $\sigma_{en} \equiv \text{false}$, then either all traces of P_1 (line 11) are terminating or `suf_pre` loses precision. Observe that we use $\text{replace_init}(P, \lambda)$ (computed from P and λ using Def. 3) where $\lambda = \lambda_{en} \wedge \neg \lambda_{ex}$ instead of P to limit our attention to those initial states that can reach **en**. If no termination criterion is satisfied, the algorithm repeats (line 18) with $\text{constrain_init}(P, \varphi_{new})$ (Def. 4) since the states that satisfy φ_{new} are the ones whose termination status is unknown so far. Note that the construction of $\text{constrain_init}(P, \varphi_{new})$ requires φ_{new} to be converted to DNF, which may blow up the number of resulting *initial clauses*, but in our experiments we have not observed that.

Soundness and termination of the algorithms. We now study some properties, including soundness and termination of the algorithms 1 and 2.

All the components used in Alg. 2 terminate, but the algorithm itself may not, owing to the fact that φ_{new} can be decreased indefinitely (the algorithm keeps refining). This is typical of algorithms for *undecidable* problems such as non-termination. So we want to ensure a weaker property, that is, of *progress*. Progress is made, in the sense that each iteration explores a strictly smaller set of initial states whose termination status are not yet known. We state this formally:

Proposition 3 (Progress and Termination of Alg. 2). Alg. 2 either terminates or progresses.

Proof (Sketch). By induction on the number of iteration of the *while* loop.

Progress. Let φ_{old} and φ_{new} be formulas characterising the set of initial states yet to be proven non-terminating at each successive iteration respectively. Note that the algorithm iterates only if $\varphi_{new} \models \varphi_{old}$ and $\varphi_{old} \not\models \varphi_{new}$, that is, if φ_{new} is strictly smaller than φ_{old} , in the set view.

Termination. Note that each individual operation in the loop, including `nec_pre`, which is computed using abstract interpretation and program transformations, terminate. The only condition under which the algorithm diverges is when φ_{new} is *strictly smaller* than φ_{old} ; in this case the algorithm progresses. $\square$

Observe that each iteration of Alg. 2 computes a valid precondition for non-termination of P even when under-approximations are used (Prop. 2). The disjunctive combination of such preconditions is also a valid precondition for non-termination of P . Again, we state this as a proposition.

Proposition 4 (Composing Preconditions). Let Φ be a set of formulas such that each $\varphi \in \Phi$ is a sufficient precondition for non-termination of P . Then so is $\bigvee \Phi$.

Proof (Sketch). $\bigvee \Phi$ satisfies the condition of Def. 5. $\square$

Proposition 5 (Soundness of Alg. 2). Let P be a program. If Alg. 2 returns σ for P , then σ is a sufficient precondition for non-termination of P .

Proof (sketch). This follows from Prop. 2 and 4, with Def. 5: At each iteration the algorithm computes a formula that satisfies the condition of Def. 5 (the formula is a sufficient precondition). Prop. 2 allows us to lift any such formula computed for `constrain_init(P,  $\varphi$ )` (for some φ) to P itself, and Prop. 4 allows us to disjunctively combine such formulas to a valid precondition for P . $\square$

The program shown here does not terminate when $x \neq 0$. On input $x > 0$ it gets stuck in the first loop, on $x < 0$ in the second. Generally, a program P with n loops may get stuck in loops $l_1, \dots, l_n$, resp., on input satisfying formulas $\varphi_1, \dots, \varphi_n$. If each such φ_i is a sufficient precondition for non-termination of P , then so is $\bigvee_{i=1}^n \varphi_i$. Taken together, Prop. 4 and 5 ensure the correctness of Alg. 1:

```
foo(int x){
  while(x > 0)
    x = x + 1
  while(x < 0)
    x = x - 1
}
```

Theorem 1 (Soundness of Alg. 1). Let P be a program. If Alg. 1 returns σ for P , then σ is a sufficient precondition for non-termination of P .

Corollary 1. If $\varphi \not\models \text{false}$ is a precondition for non-termination of a program P , then P is non-terminating.

Proof (Sketch). $\varphi \not\models \text{false}$ implies there is at least an input (satisfying φ) to P on which it does not terminate. $\square$

If Alg. 1 returns *false* for P , then P 's non-termination status is *unknown*.

5 Implementation and Experiments

Implementation. We implemented Alg. 1 as a prototype tool, NTHORN, available from <https://github.com/bishoksan/NtHorn.git>. It is written in Ciao Prolog [8] and uses PPL [1] and Yices 2.2 [18] for constraint manipulation. While refinement of candidate preconditions to the sufficient ones can be done with a tool such as [34], currently the implementation uses a simpler approach, namely, the reachability of the respective *loop* entry from each candidate—using the *safety prover* RAHFT [29]. This gives a proof of non-termination as well as some conditions on the initial states and is used in the experiments. NTHORN handles integer programs only (the classical setting for (non-)termination work [10]), but our techniques apply beyond integer arithmetic.

We rely on abstract interpretation and CHC transformations for inferring sound and precise *necessary preconditions*. In particular, our implementation performs forward and backward constraint propagation using the constraints derived from polyhedral abstraction [14] obtaining a specialised version of the program [27]. To enhance the precision of the analysis further, we apply a sequence of program transformations including control-flow refinement using partial evaluation [21], *clause splitting*, strengthening of initial clauses, and we iteratively refine necessary preconditions for entering and exiting a loop.

Experimental setting. We evaluated the approach on benchmarks from the `C_Integer` category of TermComp’20 [38]. The benchmark suite consists of 335 programs with nondeterminism: 111 non-terminating, 223 terminating, and one (Collatz) for which termination is unknown. Evaluation was done on 111 *non-terminating* programs, ignoring the terminating ones, as NTHORN can only prove non-termination. These are typical loop programs (simple or nested) with branches. Some of these loops have non-deterministic conditions while others contain non-deterministic assignments. We used small-step encoding to translate them to CHCs, obtaining only linear clauses. We run several configurations of NTHORN, namely NTHORN(X) where X can be partial evaluation `pe`, constraint specialisation `cs`, clause splitting `csp`, or some combination. Then we compare against state-of-the-art (non-)termination tools that participated in this category: APROVE [22], IRANKFINDER [5], ULTIMATEAUTOMIZER [26] and VERYMAX [31]. We used TermComp’20 versions of these tools. We add HIPTNT+ [32] (<http://loris-7.ddns.comp.nus.edu.sg/~project/hiptnt/plus/>), as it shares with NTHORN’s the ability to infer preconditions. We run several configurations of NTHORN to study the impact of different CHC transformations. Since we could not run IRANKFINDER due to some front-end issues, we took the evaluation results from StarExec [38].

Experiments were performed on a MacBook Pro, 2.7 GHz Intel Core i5 processor, 16 GB memory, running OS X 10.11.6. Timeout was 300 seconds (the competition standard) per instance. The results are shown in Table 1. The three last columns show, in order, the number of programs proved non-terminating, the number given up within 300 seconds or timeout, and the average time taken by all instances including the “gave up” instances. Among the tools, NTHORN

Tool	Proved NT	Gave up	Avg. time (secs)
APROVE	100	11	12.72
IRANKFINDER	93	18	27.00
ULTIMATEAUTOMIZER	83	28	10.94
VERYMAX	102	9	14.48
HIPTNT+	94	17	2.34
NTHORN(cs)	58	53	0.48
NTHORN(cs·pe)	94	17	0.62
NTHORN(cs·pe·csp)	98	13	0.86

Table 1. Experimental results on the TermComp’20 (C_Integer) suite.

and HIPTNT+ are the only tools capable of deriving a precondition. It would be interesting to compare the generality of the preconditions inferred by our tools. But we could not do so due to the difference in our (non-standard) output formats. So with these experiments, we seek to answer the following questions:

- Q1. Will the proposed method allow us to derive non-termination preconditions (or prove non-termination) in practice?
- Q2. How does it compare to state-of-the-art tools for proving non-termination?
- Q3. What role do the CHC transformations play?

Results. The results show different profiles wrt. to solved instances and performance for the tools. APROVE (resp. VERYMAX), the category winner in 2020 (resp. 2019), solves two (resp. four) instances more than NTHORN, while the rest solve less. This shows a remarkable effectiveness of our particular combination of mostly off-the-shelf techniques. The configuration NTHORN(cs) solves only 58 instances, while NTHORN(cs·pe) solves 94. The best result is achieved with NTHORN(cs·pe·csp). We find that each component transformation has a positive impact. Not only can we solve more problems (at the cost of solving time), we also generate more general preconditions. The combination cs·pe·csp has been chosen based on experiments, but its effectiveness aligns with our intuition. Namely, csp derives new domain specific constraints that pe can take advantage of during polyvariant specialisation, and cs, which is based on abstract interpretation, greatly benefits from the resulting specialised form.

As for speed, NTHORN(cs·pe·csp) (from here on “NTHORN”) is an order of magnitude faster than the alternatives, solving each case in less than a second, while giving up on 13 cases. The median time was 1 second, while the instance such as NO_04 with 5-level of nesting and Lcm resp. took 135 and 58 seconds and a few other took slightly more than a second. We believe the speed is due to abstract interpretation (which in this context is relatively efficient), together with the lightweight program transformation. Also, unlike other tools, NTHORN focuses on proving just non-termination. Among the 13 cases, NTHORN fails to handle LogMult and DoubleNeg because they involve non-linear operations. They are proved non-terminating only by IRANKFINDER. The cases Narrowing and NarrowKonv are shown non-terminating only by NTHORN, VERYMAX and HIPTNT+, while ChenFlurMukhopadhyay-SAS2012-Ex2.11 only by NTHORN and HIPTNT+. NTHORN could not generate preconditions for 4 programs.

These programs contain non-deterministic assignments that affect loop conditions; it might be possible that sufficient preconditions do not exist for them, though they can be shown terminating, as discussed in Sect. 1.

In summary, the results answer Q1-Q3 positively. NTHORN can be used to derive preconditions for NT and is comparable in power to the leading non-termination analysis tools, when applied to integer programs. Notably, NTHORN solves problems several orders of magnitude faster than the state-of-the-art analyzers and CHC transformations play an important role in this.

A new non-termination prover, REVTERM [9] was published recently and has not been part of our tool comparison. Like NTHORN, it does not prove termination. The experimental evaluation [9] suggests that its precision (for non-termination) is on a par with that of VERYMAX, but obtained 2–3 times faster. Comparison data from the paper [9] are in agreement with what we have found, for both precision and performance (they use a timeout of 60 seconds, rather than 300, so the average running times reported are somewhat shorter.)

6 Related Work

There is a rich body of work on proving non-termination, e.g., [2, 7, 10, 11, 24, 31, 32, 39]. Most of these provers either provide a *stem* (a sequential part from entry to loop) and the *loop*, or some precondition from which there exists a non-terminating run as a witness to non-termination. But for some applications like web-servers, a sufficient precondition (under which no trace is finite) is more useful. To our knowledge, prior to NTHORN, HIPTNT+ [32] was the only tool able to infer sufficient conditions for non-termination. Le et al. [32] propose a specification logic and Hoare-style reasoning to infer sufficient preconditions for both termination and non-termination of programs and, unlike ours, can handle programs manipulating pointers. We infer preconditions for non-termination only, relying on reduction to precondition inference for safety. But our approach is considerably simpler, as we combine existing techniques, refined iteratively.

Many non-termination provers [22, 24, 26, 39] search exhaustively for candidate *lassos* (simple while loops without branches), and attempt to prove non-termination by deriving a recurrence set using constraint solving [22, 24, 39] or automata based approaches [26]. An orthogonal approach [33] considers *lassos* with linear arithmetic and represents infinite runs as geometric series.

We exploit the notion of closed recurrence set [10] as it is useful not only for automation using a safety prover, but also for proving non-termination of non-deterministic programs and programs involving aperiodic non-termination. The method of Chen et al. [10] inserts appropriate reachability queries and uses a safety prover to eliminate terminating paths iteratively until it finds a program under-approximation and a closed recurrence set in it. The method is likely to diverge as there can be infinitely many terminating paths. Hence we use abstract interpretation to derive initial conditions that lead to the terminating paths. We negate the conditions to bar those paths. Similar to our approach, the method [31] (implemented in VeryMax [6]) searches for witnesses to non-termination

in the form of quasi-invariants (sets of states that do not exit the loop once entered) whose reachability from initial states is checked using a safety prover. Where VeryMax infers such invariants using Max-SMT solving, we do it using abstract interpretation.

Chatterjee et al. [9] rely on syntactic program reversal (applicable only to *while* programs; corresponding to linear CHCs) to derive backward polynomial conjunctive invariants using off-the-shelf tools and prove NT. Our method is more generally applicable to programs with procedures and can also infer preconditions as disjunction of linear constraints. Unlike our method, theirs provides *relative completeness guarantees*, that is, it is guaranteed to find the proof of NT under certain conditions. Bakhirkin [2], as we do, uses forward and backward abstract interpretation to find potential recurrence sets whose reachability implies non-termination. The approach can be applied to heap manipulating programs but is limited to simple *while* program. Ben-Amram et al. [4] derive a recurrence set from a failed attempt to prove termination of *multi-phase* loops, while our approach is direct and only proves non-termination.

While the above methods target programs with linear arithmetic, Cook et al. [11] prove non-termination of programs with non-linear arithmetic and heap-based operations. The key is the notion of *live abstraction*, an abstraction heuristic that ensures that any abstract trace corresponding to a terminating concrete trace is also terminating. In other words, it does not introduce any non-termination and is a sound abstraction heuristic for non-termination. This allows over-approximating non-linear assignments and heap-based commands with non-deterministic linear assignments.

7 Concluding Remarks

We have presented a new approach to preconditions for non-termination. The problem is reduced to inference of preconditions for safety, via insertion of “reachability queries” in program loops. The reduction enables us to use existing tools and techniques for safety preconditions. A prototype implementation is competitive with the state-of-the-art tools for automated proof of non-termination.

NTHORN can only infer preconditions for non-termination and is limited to programs manipulating linear integer arithmetic, whose applicability is determined by the underlying tool for inferring preconditions for safety. In future, we plan to complement it with termination analysis as done in other tools and also extend to programs that manipulate structured data, à la Cook et al. [11].

Acknowledgements

We thank the three anonymous reviewers for their careful reading of an earlier version of the paper, and their constructive suggestions for how to improve it. Bishoksan Kafle has been partially funded by the Spanish Ministry of Research, Science and Innovation, grant MICINN PID2019-108528RB-C21 *ProCode* and Madrid P2018/TCS-4339 *BLOQUES-CM*.

Bibliography

- [1] R. Bagnara, P. M. Hill, and E. Zaffanella. The Parma Polyhedra Library. *Sci. Comput. Program.*, 72(1-2):3–21, 2008.
- [2] A. Bakhirkin. *Recurrent Sets for Non-Termination and Safety of Programs*. PhD thesis, University of Leicester, 2016.
- [3] A. Bakhirkin and D. Monniaux. Combining forward and backward abstract interpretation of Horn clauses. In F. Ranzato, editor, *Static Analysis*, volume 10422 of *LNCS*, pages 23–45. Springer, 2017.
- [4] A. M. Ben-Amram, J. J. Doménech, and S. Genaim. Multiphase-linear ranking functions and their relation to recurrent sets. In B. Chang, editor, *Static Analysis*, volume 11822 of *LNCS*, pages 459–480. Springer, 2019.
- [5] A. M. Ben-Amram and S. Genaim. Ranking functions for linear-constraint loops. *Journal of the ACM*, 61(4):26:1–26:55, 2014.
- [6] C. Borralleras, M. Brockschmidt, D. Larraz, A. Oliveras, E. Rodríguez-Carbonell, and A. Rubio. Proving termination through conditional termination. In A. Legay and T. Margaria, editors, *Proc. TACAS’17*, volume 10205 of *LNCS*, pages 99–117. Springer, 2017.
- [7] M. Brockschmidt, T. Ströder, C. Otto, and J. Giesl. Automated detection of non-termination and NullPointerExceptions for Java bytecode. In B. Beckert, F. Damiani, and D. Gurov, editors, *Formal Verification of Object-Oriented Software*, volume 7421 of *LNCS*, pages 123–141. Springer, 2012.
- [8] F. Bueno, D. Cabeza, M. Carro, M. Hermenegildo, P. López-García, and G. Puebla. The Ciao Prolog system: Reference manual. Technical Report CLIP 3/97.1, UPM, 1997. <http://www.clip.dia.fi.upm.es/>.
- [9] K. Chatterjee, E. K. Goharshady, P. Novotný, and D. Žikelić. Proving non-termination by program reversal. In *Proc. PLDI 2021*, pages 1033–1048. ACM, 2021.
- [10] H. Y. Chen, B. Cook, C. Fuhs, K. Nimkar, and P. W. O’Hearn. Proving nontermination via safety. In E. Ábrahám and K. Havelund, editors, *Proc. TACAS’14*, volume 8413 of *LNCS*, pages 156–171. Springer, 2014.
- [11] B. Cook, C. Fuhs, K. Nimkar, and P. W. O’Hearn. Disproving termination with overapproximation. In *Proc. FMCAD’14*, pages 67–74. IEEE, 2014.
- [12] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *Proc. POPL’77*, pages 238–252. ACM, 1977.
- [13] P. Cousot, R. Cousot, M. Fähndrich, and F. Logozzo. Automatic inference of necessary preconditions. In R. Giacobazzi, J. Berdine, and I. Mastroeni, editors, *Proc. VMCAI’13*, volume 7737 of *LNCS*, pages 128–148. Springer, 2013.
- [14] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Proc. POPL’78*, pages 84–96. ACM, 1978.

- [15] E. De Angelis, F. Fioravanti, A. Pettorossi, and M. Proietti. Program verification via iterated specialization. *Sci. Comput. Program.*, 95:149–175, 2014.
- [16] E. De Angelis, F. Fioravanti, A. Pettorossi, and M. Proietti. Semantics-based generation of verification conditions via program specialization. *Sci. Comput. Program.*, 147:78–108, 2017.
- [17] J. J. Doménech, J. P. Gallagher, and S. Genaim. Control-flow refinement by partial evaluation, and its application to termination and cost analysis. *Theory Pract. Log. Program.*, 19(5-6):990–1005, 2019.
- [18] B. Dutertre. Yices 2.2. In A. Biere and R. Bloem, editors, *Computer-Aided Verification*, volume 8559 of *LNCS*, pages 737–744. Springer, 2014.
- [19] F. Fioravanti, A. Pettorossi, and M. Proietti. Specialization with clause splitting for deriving deterministic constraint logic programs. In *Proc. IEEE Conf. Systems, Man and Cybernetics*. IEEE Press, 2002.
- [20] J. P. Gallagher. Tutorial on specialisation of logic programs. In *PEPM’93: Proc. 1993 ACM SIGPLAN Symp. Partial Evaluation and Semantics-Based Program Manipulation*, pages 88–98. ACM, 1993.
- [21] J. P. Gallagher. Polyvariant program specialisation with property-based abstraction. In A. Lisitsa and A. P. Nemytykh, editors, *Proc. Seventh Int. Workshop on Verification and Program Transformation*, volume 299 of *EPTCS*, pages 34–48, 2019.
- [22] J. Giesl et al. Proving termination of programs automatically with AProVE. In S. Demri, D. Kapur, and C. Weidenbach, editors, *Automated Reasoning*, volume 8562 of *LNCS*, pages 184–191. Springer, 2014.
- [23] S. Grebenschikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In J. Vitek, H. Lin, and F. Tip, editors, *Proc. PLDI’12*, pages 405–416. ACM, 2012.
- [24] A. Gupta, T. A. Henzinger, R. Majumdar, A. Rybalchenko, and R. Xu. Proving non-termination. In *Proc. 35th ACM Symp. Principles of Programming Languages*, pages 147–158. ACM, 2008.
- [25] A. Gurfinkel, T. Kahsai, A. Komuravelli, and J. A. Navas. The SeaHorn verification framework. In D. Kroening and C. S. Păsăreanu, editors, *Computer-Aided Verification*, volume 9206 of *LNCS*, pages 343–361. Springer, 2015.
- [26] M. Heizmann, J. Hoenicke, and A. Podelski. Termination analysis by learning terminating programs. In A. Biere and R. Bloem, editors, *Computer-Aided Verification*, volume 8559 of *LNCS*, pages 797–813. Springer, 2014.
- [27] B. Kafle and J. P. Gallagher. Constraint specialisation in Horn clause verification. *Sci. Comput. Program.*, 137:125–140, 2017.
- [28] B. Kafle, J. P. Gallagher, G. Gange, P. Schachte, H. Søndergaard, and P. J. Stuckey. An iterative approach to precondition inference using constrained Horn clauses. *Theory Pract. Log. Program.*, 18:553–570, 2018.
- [29] B. Kafle, J. P. Gallagher, and J. F. Morales. RAHFT: A tool for verifying Horn clauses using abstract interpretation and finite tree automata. In S. Chaudhuri and A. Farzan, editors, *Computer-Aided Verification*, volume 9779 of *LNCS*, pages 261–268. Springer, 2016.

- [30] B. Kafle, G. Gange, P. Schachte, H. Søndergaard, and P. J. Stuckey. Transformation-enabled precondition inference. *Theory Pract. Log. Program.*, 2021. To appear.
- [31] D. Larraz, K. Nimkar, A. Oliveras, E. Rodríguez-Carbonell, and A. Rubio. Proving non-termination using Max-SMT. In A. Biere and R. Bloem, editors, *Computer-Aided Verification*, volume 8559 of *LNCS*, pages 779–796. Springer, 2014.
- [32] T. C. Le, S. Qin, and W.-N. Chin. Termination and non-termination specification inference. In D. Grove and S. M. Blackburn, editors, *Proc. PLDI’15*, pages 489–498. ACM, 2015.
- [33] J. Leike and M. Heizmann. Geometric nontermination arguments. In D. Beyer and M. Huisman, editors, *Proc. TACAS’18*, volume 10806 of *LNCS*, pages 266–283. Springer, 2018.
- [34] A. Miné. Inferring sufficient conditions with backward polyhedral under-approximations. *Electronic Notes in Theor. Comp. Sci.*, 287:89–100, 2012.
- [35] Y. Moy. Sufficient preconditions for modular assertion checking. In F. Logozzo, D. A. Peled, and L. D. Zuck, editors, *Proc. VMCAI’08*, volume 4905 of *LNCS*, pages 188–202. Springer, 2008.
- [36] J. C. Peralta, J. P. Gallagher, and H. Sağlam. Analysis of imperative programs through analysis of constraint logic programs. In G. Levi, editor, *Static Analysis*, volume 1503 of *LNCS*, pages 246–261. Springer, 1998.
- [37] M. N. Seghir and P. Schrammel. Necessary and sufficient preconditions via eager abstraction. In J. Garrigue, editor, *Proc. APLAS’14*, volume 8858 of *LNCS*, pages 236–254. Springer, 2014.
- [38] Termination competition 2020: C Integer. https://termcomp.github.io/Y2020/job_41519.html. Accessed: 2021-06-01.
- [39] H. Velroyen and P. Rümmer. Non-termination checking for imperative programs. In B. Beckert and R. Hähnle, editors, *Tests and Proofs*, volume 4966 of *LNCS*, pages 154–170. Springer, 2008.