

Proving Conditional Termination

Byron Cook¹, Sumit Gulwani¹, Tal Lev-Ami²,
Andrey Rybalchenko³, and Mooly Sagiv²

¹ Microsoft Research

² Tel Aviv University

³ MPI-SWS

Abstract. We describe a method for synthesizing reasonable underapproximations to weakest preconditions for termination—a long-standing open problem. The paper provides experimental evidence to demonstrate the usefulness of the new procedure.

1 Introduction

Termination analysis is critical to the process of ensuring the stability and usability of software systems. Verification of liveness properties such “Will `Decode()` always return back to its call sites?” or “Is every call to `Acquire()` eventually followed by a call to `Release()`?” can be reduced to a question of program termination [12, 27]. Automatic methods for proving such properties are now well studied in the literature, *e.g.* [1, 3, 5, 8, 10, 20, 25], and are now being deployed in industrial program verification tools such as SDV [2]. But what about the cases in which code only terminates for *some* inputs? What are the preconditions under which the code is still safe to call, and how can we automatically synthesize these conditions? We refer to these questions as the *conditional termination* problem.

This paper describes a method for proving conditional termination. Our method is based on the discovery of *conditional ranking functions*—functions over program states that are bounded but not necessarily decreasing—and then finding preconditions that promote the conditional ranking functions into standard ranking functions. We describe two procedures based on this idea: `PRESYNTH`, which finds preconditions to termination, and `PRESYNTHPHASE`, which extends `PRESYNTH` with the ability to identify the phases necessary to prove the termination of phase-transition programs [7].

The challenge in this area is to find the *right precondition*: the empty precondition is correct but useless, whereas the *weakest precondition* for even very simple programs can often be expressed only in complex domains not supported by today’s tools. In this paper we seek a method that finds useful preconditions. Such preconditions need to be weak enough to allow interesting applications of the code in question, but also expressible in the subset of logic supported by decision procedures, model checking tools, etc. Furthermore, they should be computed quickly (the weakest precondition expressible in the target logic may

be too expensive to compute). Since the notion of the right precondition is difficult to define, in this paper we allow the the reader to judge the quality of preconditions computed by our procedure for a number of examples. Several of these examples are drawn from industrial applications.

Limitations In this paper, we limit ourselves to the termination property and to sequential arithmetic programs. Note that, at the cost of complicating the exposition, we could use known techniques (*e.g.*, [4] and [12]) to extend our approach to ω -regular liveness properties and programs with heap. Our technique could also provide assistance when analyzing concurrent programs via [14], although we suspect that synthesizing *environment* abstractions that guarantee thread-termination is a more important problem for concurrent programs than conditional termination.

Related work Until now few papers have directly addressed the problem of automatically underapproximating weakest preconditions. One exception is [18], which yields constraint systems that are non-linear. The constraint-based technique in [9] could also be modified to find preconditions, but again at the cost of non-linear constraints.

The synthesis of underapproximations for weakest preconditions has many applications (*e.g.*, disproving safety properties, termination proving, interprocedural analysis). Thus, we can find various works that briefly mention weak methods of underapproximating preconditions for given applications. The most prominent example known to the authors is [26]. In contrast to these previous works, our technique is generally applicable to all arithmetic programs, and it produces much more accurate underapproximations than can be found in the existing literature.

In contrast to methods for underapproximating weakest preconditions, techniques for weakest *liberal* preconditions are abundant (*e.g.*, [11, 21]). Note that weakest preconditions are so rarely considered in the literature that weakest *liberal* preconditions are often simply called weakest preconditions, *e.g.*, [21].

Many of the components used in our procedures are known. For example we use known techniques for computing weakest liberal preconditions [17], quantifier elimination [15], and tools for proving program termination. EM-trees [7] are related to our method of detecting phase-transitions.

2 Example

In this section we informally illustrate our method by applying it to several examples. The procedures proposed by this paper are displayed in Figures 1 and 2. They will be more formally described in §3.

We have split our method into two procedures for presentational convenience. The first procedure illustrates our method's key ideas, but fails for the class of phase-transition programs. The second procedure extends the first with support for phase-transition programs. Note that phase-change programs and precondi-

tions are interrelated (allowing us to solve the phase-change problem easily with our tool), as a phase-change program can be thought of as two copies of the same loop composed, but with different preconditions.

2.1 Finding preconditions for programs without phase-change

We consider the following code fragment:

```

1      // @requires true;
2      while(x>0){
3          x=x+y;
4      }
```

We assume that the program variables x and y range over integers. The initially given `requires`-clause is not sufficient strong to guarantee termination. For example, if $x=1$ and $y=0$ at the loop entry then the code will not terminate. The weakest precondition for termination of this program is $x \leq 0 \vee y < 0$.

If we apply an existing termination prover, *e.g.*, TERMINATOR [13] or ARMC [24], on this code fragment then it will compute a counterexample to termination. The counterexample consists of 1) a stem η , which allows for manipulating the values before the loop is reached, and 2) a repeatable cycle ρ , which is a relation on program states that represents an arbitrary number of iterations of the loop.

To simplify the presentation, we represent the stem η as an initial condition θ on the variables of the loop part. (Section 4 describes this step in more details.) In our example, the initial condition θ is `true` and the transition relation of the loop is defined by:

$$\rho(\{x, y\}, \{x', y'\}) \equiv x > 0 \wedge x' = x + y \wedge y' = y .$$

In order to prove this counterexample spurious, we need to find a ranking function f such that $\rho(X, X') \rightarrow \mathcal{R}_f(X, X')$, where \mathcal{R}_f is the *ranking relation* defined by f :

$$\mathcal{R}_f(X, X') \equiv f(X) \geq 0 \wedge f(X') \leq f(X) - 1 .$$

As the termination prover has returned the above relation ρ as a counterexample, no linear ranking function f exists.

Finding candidate bounds Due to the absence of linear ranking function for ρ , we find a *conditional ranking function*, *i.e.*, a function b such that one of the conjuncts defining $\mathcal{R}_b(X, X')$ holds for ρ . We compute a conditional ranking function for ρ by finding an expression on the variables $\{x, y\}$ that is bound from below. This can be done by considering only the domain (and not the range) of ρ , *i.e.*, find functions that are bounded when there is a successor. In other words, consider $\exists x', y'. x > 0 \wedge x' = x + y \wedge y' = y$.

We apply existential quantifier elimination techniques, *i.e.*, we have

$$\text{QELIM}(\exists x', y'. x > 0 \wedge x' = x + y \wedge y' = y) \equiv x > 0 .$$

We normalize the condition $x > 0$ as $x - 1 \geq 0$, and use the function $b = x - 1$. Because $\rho(\{x, y\}, \{x', y'\}) \Rightarrow b(\{x, y\}) \geq 0$, which is the first conjunction required by \mathcal{R}_b , we can use b as our candidate ranking function.⁴

Enforcing ranking with an invariant The function $b = x - 1$ that we found only satisfies part of the requirements for proving termination with \mathcal{R}_b (*i.e.*, $b(X) \geq 0$ but $b(X') \not\leq b(X) - 1$). We need an invariant $s(\{x, y\})$ such that

$$s(\{x, y\}) \wedge \rho(\{x, y\}, \{x', y'\}) \Rightarrow \mathcal{R}_b(\{x, y\}, \{x', y'\}) .$$

Since b is bounded, we find $s(\{x, y\})$ as follows:

$$s(\{x, y\}) \equiv \text{QELIM}(\forall x', y'. \rho(x, y, x', y') \Rightarrow b(\{x', y'\}) \leq b(\{x, y\}) - 1) .$$

We obtain $s(\{x, y\}) = x \leq 0 \vee y < 0$.

Synthesizing a precondition guaranteeing the invariant Recall that the original problem statement is to find a precondition that guarantees termination of the presented code fragment. As the invariant s guarantees termination, we now need to find a precondition that guarantees the validity of s on every iteration of ρ . The required assertion is the weakest *liberal* precondition of s wrt. the loop statement. We use known techniques for computing underapproximations of weakest liberal preconditions to find the precondition that ensures the after any number of iterations of the loop s must hold in the next iteration. Using a tool for abstract interpretation of arithmetic programs [19], we obtain $r(\{x, y\}) = x \leq 0 \vee y < 0$. In summary, our procedure has discovered the precondition proposed above.

Note that we can alternate executions of our procedure together with successive applications of a termination prover to find a precondition that is strong enough to prove the termination of the entire program. The interaction between the tools is based on counterexamples for termination, which are discovered by the termination prover and are eliminated by the precondition synthesis procedure.

2.2 Finding preconditions for phase-change programs

Consider the following code fragment:

```

1      // @requires true;
2      while(x>0){
3          x=x+y;
4          y=y+z;
5      }
```

⁴ In this simple example the result was exactly the loop condition. However, when translating the cycle returned from the termination prover to a formula, some of the conditions are not on the initial variables.

Again, the given `requires`-clause is not sufficient to ensure termination. For example, if $x=1$, $y=0$, and $z=0$ at the loop entry then the code will not terminate. However, this time, the weakest precondition is given by a non-linear assertion, which is difficult to construct automatically.

Note that the precondition $z < 0$ guarantees that the loop terminates, but it does so after passing through possibly two phases. The first phase is characterized by the assertion $y \geq 0$. In fact, during this phase the value of x may not decrease towards zero. Nevertheless, *eventually* the value of y will decrease below zero, *i.e.*, a phase transition takes place. At this point x will start decreasing towards (and eventually reaching) zero.

In this example, the termination prover returns a stem which is the identity relation and cycle relation

$$\rho(\{x, y, z\}, \{x', y', z'\}) \equiv x > 0 \wedge x' = x + y \wedge y' = y + z \wedge z' = z .$$

The first step of the precondition inference repeats the procedure presented in the previous subsection. On this example, similarly to the previous one, the procedure computes $b = x - 1$ and $s = x \leq 0 \vee y < 0$. However, when computing a precondition that ensures that s is an invariant, we obtain the following linear underapproximation⁵

$$\begin{aligned} & x \leq 0 \vee x + y \leq 0 \vee (y < 0 \wedge x + 2 * y + z \leq 0) \vee \\ & (y < 0 \wedge x + 3 * y + 3 * z \leq 0) \vee (y < 0 \wedge z \leq 0) . \end{aligned}$$

The first four disjuncts correspond to the cases when the loop terminates after 0, 1, 2, and 3 iterations, respectively, and in which $y < 0$ holds until possibly the last iteration. The last disjunct is more interesting, as it states that if $y < 0$ and $z \leq 0$ then the loop terminates. The expected precondition $z < 0$ is not included in the disjunction, since it does not guarantee that $y < 0$ from the start.

Note that each of these conditions guarantee termination when they are satisfied at any iteration of the loop, not necessarily at the first one. These conditions identify phase transition points. Once they are met, x will start decreasing until the loop terminates. The solution is to constrain the transition relation with the negated condition. Termination of the constrained loop ensures that these conditions are eventually met, and thus the original loop will terminate. Thus, we call the procedure recursively with the constrained transition relation, and disjoin the returned preconditions with the existing ones.

For example, constraining the loop with $\neg(y < 0 \wedge z \leq 0)$ and calling the procedure recursively yields the additional preconditions $x + 2 * y + z \leq 0 \vee x + 3 * y + 3 * z \leq 0 \vee z < 0$. The first two disjuncts are weakenings of the previous preconditions for the cases in which the loop terminates after 2 or 3 iteration, removing the condition $y < 0$. The last disjunct $z < 0$ is the more interesting, since it ensures eventual termination.

⁵ Any linear precondition must be an underapproximation, since the weakest liberal precondition is non-linear.

After simplification the precondition computed by the procedure is

$$\begin{aligned} & \mathbf{x} \leq 0 \vee \mathbf{x} + \mathbf{y} \leq 0 \vee \mathbf{x} + 2 * \mathbf{y} + \mathbf{z} \leq 0 \vee \\ & \mathbf{x} + 3 * \mathbf{y} + 3 * \mathbf{z} \leq 0 \vee (\mathbf{y} < 0 \wedge \mathbf{z} \leq 0) \vee \mathbf{z} < 0 . \end{aligned}$$

3 Computing preconditions for termination

This section formally describes the two methods for computing preconditions for termination, PRESYNTH and PRESYNTHPHASE. We first begin by formally defining basic concepts.

3.1 Preliminaries

We assume a program $P = (X, \theta, \rho)$ over a finite set of variables X . For simplicity, we do not single out the program counter variable that ranges over control location of the program, and assume that it is included in X . The assertion θ represents the initial condition of the program. The transition relation ρ is represented by an assertion over the program variables X and their primed versions X' . The primed variables refer to the values after the transition is executed. For practicality, we assume that the initial condition and the transition relation are represented using an expression language for which (approximative) quantifier elimination procedures exist.⁶

We refer to valuations of program variables as program states. A program computation is a finite or infinite sequence of program states s_1, s_2, \dots such that s_1 is an initial state and each pair of consecutive states s_i and s_{i+1} follows the transition relation. Formally, $s_1 \models \theta$ and for each s_i in the computation we have $(s_i, s_{i+1}) \models \rho$. The program terminates from a state s if there is no infinite computation that starts at s . An assertion r is a precondition for termination if the program P terminates from each state s that satisfy r and the initial condition, *i.e.*, $s \models \theta \wedge r$.

Our procedure uses (under-approximation of) weakest liberal precondition operator WLP, which we define as usual. Given a transition relation $\rho(X, X')$ and an assertion over program variables $\varphi(X)$, this transformer yields the following conjunction when applied on the transitive closure of ρ :

$$\text{WLP}(\rho^*(X, X'), \varphi(X)) \equiv \bigwedge_{n \geq 0} \forall X'. \rho^n(X, X') \Rightarrow \varphi(X') .$$

3.2 The Procedure PreSynth

Figure 1 shows the basic precondition synthesis procedure PRESYNTH. The input of PRESYNTH consists of an initial condition θ and a transition relation ρ . Note that in the most likely usage scenario, θ and ρ will represent a counterexample

⁶ If necessary we can over-approximate the initial condition and the transition relation using assertions from such a language.

	<pre> function PRESYNTH input $\theta(X)$: initial condition $\rho(X, X')$: transition relation begin 1 $\rho(X, X') := \rho(X, X') \wedge \text{QELIM}(\exists X_0. \theta(X_0) \wedge \rho^*(X_0, X))$ 2 $B := \{b(X) \mid \forall X. \text{QELIM}(\exists X'. \rho(X, X')) \Rightarrow b(X) \geq 0\}$ 3 foreach $b(X) \in B$ do 4 $s(X) := \text{QELIM}(\forall X'. \rho(X, X') \Rightarrow b(X') \leq b(X) - 1)$ 5 $r(X) := r(X) \vee \text{WLP}(\rho^*(X, X'), s(X))$ 6 done 7 return “precondition for termination $r(X)$ ” end. </pre>
--	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Fig. 1. The procedure PRESYNTH synthesizes a precondition for the termination of a transition relation ρ from initial states θ . PRESYNTH can be used with abstractions of lines 1 and 5—over- and under-approximation, respectively—without compromising the soundness of the procedure.

to termination reported by a termination prover. In this case, PRESYNTH is applied on a compact code fragment and not the full program, which allows one to apply precise, automated reasoning-based techniques.

We discuss the major steps of the procedure in more detail. Line 1 in Figure 1 strengthens the transition relation ρ with (possibly an over-approximation of) the states that are reachable from θ . This step provides necessary precision for the subsequent computations by taking into account reachability invariants. Here, we rely on an efficient abstract interpretation tool based on *e.g.*, the Octagon domain [22]. Thus, we can implement line 1 using a standard abstract reachability procedure [16] (hence, here we assume $\theta \Rightarrow \theta^\#$, $\rho \Rightarrow \rho^\#$, and lfp is the least fixpoint operator):

$$\rho(X, X') := \rho(X, X') \wedge \text{lfp}(\rho^\#, \theta^\#)$$

Line 2 of PRESYNTH computes a set of expressions $B = \{b_1(X), \dots, b_n(X)\}$ that are bounded by the (strengthened) transition relation ρ , *i.e.*, we have $\exists X'. \rho(X, X') \Rightarrow b(X) \geq 0$ for each $b(X) \in B$. Note that this set is always finite. To compute B we apply an existential quantifier elimination procedure to eliminate the primed variables. Then, we consider the convex hull approximation of the result. Each bound expression $b(X) \in B$ is treated as a potential ranking function, and is used to guide the search for a restriction $s(X)$ on the domain of the transition relation ρ that makes ρ well-founded. PRESYNTH only considers the well-foundedness arguments constructed using $b(X)$. Line 4 uses a quantifier elimination procedure to construct the restriction $s(X)$ by imposing the *decrease* condition

$$s(X) \wedge \rho(X, X') \Rightarrow b(X') \leq b(X) - 1 .$$

Since $b(X)$ is bounded by ρ , the assertion $s(X)$ guarantees that $b(X)$ is a ranking function for the given transition relation.

The restriction imposed by $s(X)$ is effective if $s(X)$ holds for all states that are reachable from θ by applying ρ . Line 5 in Figure 1 computes the necessary precondition that guarantees the invariance of $s(X)$. All preconditions that are found using the potential ranking functions are accumulated in an assertion $r(X)$, and reported to the programmer in line 7.

Line 5 of PRESYNTH uses known abstract interpretation-based techniques for under-approximating sets and relations for the computation of $\text{WLP}(\rho^*(X, X'), s(X))$. We pass the following program to the INTERPROC analyzer [19]:

```

assume  $\theta(X')$ ;
while (*) do;
     $X=X'; X'=*$ ;
    assume  $\rho(X, X')$ ;
od
assume  $\neg s(X)$ ;

```

Using INTERPROC we first compute an abstract fixpoint using backwards analysis starting from top element in the abstract domain. Then, the abstract element at the loop entry location represents an over-approximation of the states that fail the assertion $s(X)$. The complement of this element, which we obtain by negating the corresponding assertion, provides an under-approximation of the initial set of states of the program that guarantees the invariance of $s(X)$.

Theorem 1 (PreSynth correctness). *Let $r(X)$ be an assertion computed by the procedure PRESYNTH. r is a precondition for termination of a program with the initial condition θ and the transition relation ρ .*

Proof. Let $r(X)$ be computed by PRESYNTH. For a proof by contradiction, we assume that there is an infinite computation s_1, s_2, \dots from an initial state that satisfies $r(X)$. Let $b(X)$ be a bound expression that contributed a disjunct in $r(X)$ that holds for the state s_1 , and $s(X)$ be a corresponding restriction. From the definition of WLP, we have that $s(X)$ holds for each state s_i , where $i \geq 0$. Thus, the value of $b(X)$ decreases after each program step, while being bounded from below. We reached a contradiction to the assumption that the computation is infinite. \square

3.3 The Procedure PreSynthPhase

See Figure 2. The procedure PRESYNTHPHASE extends the applicability of the basic procedure from Figure 1 to phase-transition programs (as described in Section 2). It removes the condition that the computed strengthening needs to apply immediately. Instead, we only require that the strengthening *eventually* applies. PRESYNTHPHASE implements this eventuality requirement by finding a precondition for termination of an augmented transition relation that avoids visiting states satisfying the strengthening constraint. The inferred precondition can be enforced eventually by applying PRESYNTHPHASE recursively.

We discuss the major steps of the procedure in more detail. Line 1 computes an initial precondition $r(X)$ by applying the procedure PRESYNTHPHASE. If the

```

function PRESYNTHPHASE
input
   $\theta(X)$  : initial condition
   $\rho(X, X')$  : transition relation
begin
1    $r(X) := \text{PRESYNTH}(\theta, \rho)$ 
2   if  $\exists X. r(X)$  then
3      $\rho(X, X') := \rho(X, X') \wedge \neg r(X)$ 
4      $r(X) := r(X) \vee \text{PRESYNTHPHASE}(\theta, \rho)$ 
5   endif
6   return “precondition for phase termination  $r(X)$ ”
end.

```

Fig. 2. The procedure PRESYNTHPHASE synthesizes a precondition for the phase termination of a transition relation ρ from initial states θ . It applies the procedure PRESYNTH from Figure 1 when recursively identifying computation phases. PRESYNTHPHASE can be stopped at any time, e.g., by reaching a user-provided upper bound, and it will yield a sound precondition for phase-transition termination.

precondition $r(X)$ is non-empty, as checked in line 2, then PRESYNTHPHASE weakens it by a precondition that ensures the eventuality of $r(X)$. Here, we follow a standard technique for the verification of temporal liveness properties [27], and apply a translation a termination problem. Effectively, line 3 constructs a new transition relation from ρ that avoids $r(X)$. Thus, we can apply PRESYNTHPHASE recursively on the new transition relation, see line 4.

Theorem 2 (PreSynthPhase correctness). *Let $r(X)$ be an assertion computed by the procedure PRESYNTHPHASE. r is a precondition for termination of a program with the initial condition θ and the transition relation ρ .*

Proof. Let $r_1(X), \dots, r_n(X)$ be a sequence of preconditions that are computed by applying PRESYNTH during the execution of PRESYNTHPHASE. We observe that every computation that starts in an initial state that satisfies $r_i(X)$ for $i > 1$ either terminates or eventually reaches $r_{i-1}(X)$. Hence, eventually every computation either stops or reaches $r_1(X)$. From Theorem 1 follows that the program terminates on each state reachable from an initial state satisfying $r_1(X)$. \square

4 Implementation and experiments

We have built a prototype implementation of our method based on the following collection of tools: We apply ARMC as the termination prover [24]; for quantifier elimination we use an implementation [6] of Cooper’s procedure [15]; and for abstract interpretation we use the INTERPROC analyzer [19].

See Figures 3 and 4, which contain example programs (both hand written and drawn from industrial examples) together with the results of our tool. We

Program fragment	Precondition & notes
<pre> i = 0; if (l_var >= 0) { while (l_var < 1073741824) { i++; l_var = l_var << 1; } } </pre>	$l_var > 0 \vee l_var < 0 \vee l_var \geq 1073741824$ <hr/> <p>Example from an audio compression module: We model shift by multiplication and checked for with an extra check and subtract.</p> <p>Time: 22 seconds.</p>
<pre> while (cbSrcLength >= cbHeader) { DWORD dwHeader; cbBlockLength = (UINT)min(cbSrcLength, nBlockAlignment); cbSrcLength -= cbBlockLength; cbBlockLength -= cbHeader; dwHeader = *(DWORD HUGE_T *)pbSrc; pbSrc += sizeof(DWORD); nPredSample = (int)(short)DWORD(dwHeader); nStepIndex = (int)(BYTE)HIWORD(dwHeader); if (!imaadpcmValidStepIndex(nStepIndex)) return 0; *pbDst++ = (BYTE)((nPredSample >> 8) + 128); while (cbBlockLength-->) { bSample = *pbSrc++; nEncSample = (bSample & (BYTE)0x0F); nSz = step[nStepIndex]; nPredSample = imaadpcmSampleDecode(nEncSample, nPredSample, nSz); nStepIndex = imaadpcmNextStepIndex(nEncSample, nStepIndex); *pbDst++ = (BYTE)((nPredSample >> 8) + 128); nEncSample = (bSample >> 4); nSz = step[nStepIndex]; nPredSample = imaadpcmSampleDecode(nEncSample, nPredSample, nSz); nStepIndex = imaadpcmNextStepIndex(nEncSample, nStepIndex); *pbDst++ = (BYTE)((nPredSample >> 8) + 128); } } </pre>	$cbSrcLength < cbHeader \vee (nBlockAlignment > 0 \wedge cbHeader > 0)$ <hr/> <p>Example from another audio application.</p> <p>Time: 106 seconds.</p>
<pre> while (Length-->) { READ_PURT_UCHAR(&IOBase->DataReg); } </pre>	$length > 0$ <hr/> <p>Time: 0.7 seconds</p>

Fig. 3. Programs drawn from industrial examples, runtimes, and synthesized preconditions. The runtimes include the iteration of the procedure in Figure 2 together with ARMC as termination prover.

leave it to the reader to judge the usefulness of the synthesized preconditions. The running times in the figures include the times of iterating the procedure of Figure 2 and ARMC as a termination prover. In the remainder of this section we highlight relevant implementation details.

Note to CAV reviewers: *We are preparing a release of our tool and plan to make it available in time for the CAV conference. You can find the current version of the tool at <http://www.cs.tau.ac.il/~tla/condterm/>*

Simplification Both the quantifier elimination and the abstract interpretation tools do not give minimal formulas. In many cases the formulas have redundant conjuncts or several disjuncts that imply each other. Simplification of these formulas was important for reducing the size of the result preconditions.

Program fragment	Precondition & notes
<pre> // @requires true; while(x>0){ x=x+y; y=y+z; } </pre>	$x \leq 0 \vee x + y \leq 0 \vee x + 2y + z \leq 0 \vee x + 3y + 3z \leq 0 \vee z < 0 \vee (z \leq 0 \wedge y < 0)$ <hr/> <p>This is the example from §2. The first four disjuncts cover the case where the loop terminates after 0, 1, 2, or 3 iterations respectively. The last two disjuncts ensure that the loop eventually terminates. The condition $z \leq 0 \wedge y < 0$ is the case in which the loop starts with $y < 0$ and z does not interfere.</p> <p>Time: 24 seconds.</p>
<pre> // @requires true; while(x<=N){ if (*) { x=2*x+y; y=y+1; } else { x++; } } </pre>	$x > N \vee x + y \geq 0$ <hr/> <p>An example from [7]. We find preconditions that ensure that the loop is executed 0 times ($x > N$) and the precondition that ensures termination as it appears in the paper ($x + y \geq 0$).</p> <p>Time: 4 seconds.</p>
<pre> // @requires true; while(x>=0){ x= -2*x + 10; } </pre>	$x > 5 \vee x < 0$ <hr/> <p>Example (from [23]) of a terminating linear program with no linear ranking function. Note that the program always terminates after at least 5 iterations. Our synthesized precondition provides conditions that guarantee termination after 0 or 1 iterations.</p> <p>Time: 2 seconds.</p>
<pre> // @requires n>200 and y<9; x = 0; while (1) { if (x<n) { x=x+y; if (x>=200) break; } } </pre>	$y > 0$ <hr/> <p>Example (from [18]) after translation that allows us to disprove a safety property.</p> <p>Time: 6 seconds.</p>
<pre> while (x!=y) { if (x>y) { x = x - y; } else y = y - x; } </pre>	$x = y \vee (x > 0 \wedge y > 0)$ <hr/> <p>The Euclidean algorithm for GCD. Each case in the if requires a different ranking function. Demonstrates the interaction with disjunctive well foundedness.</p> <p>Time: 17 seconds.</p>

Fig. 4. Handwritten programs, runtimes, and synthesized preconditions. The runtimes include the iteration of the procedure in Figure 2 together with ARMC as termination prover.

Loop termination after a bounded number of iterations The recursive call in PRESYNTHPHASE is problematic in case the precondition is the result of a loop termination in a bounded number of iteration, say 3. The problem is that the next call will find precondition that come from the loop terminating after 6 iteration and so on. To solve this problem, before the recursive call, we check whether the precondition ensures termination in a fixed number of steps, and if it does then we avoid the recursive call.

Stems vs. initial conditions ARMC returns a cycle and a stem as a counterexample. The stem is not only an initial condition, but rather a sequence of primitive statements that manipulate variables. We use existential quantifier elimination to convert the stem into the initial condition for the procedure. Because there are no control-flow constructs in the stem, a straightforward application of WLP translates the precondition discovered by the procedure to a precondition applicable at the beginning of the stem.

Translating paths to formulas The cycles returned by ARMC often contain multiple assignments and conditions. The translation first converts the path to SSA form and then to a path formula (see [13] for details). The extra temporaries generated are handled in the quantifier elimination as extra variables to eliminate.

Iterating the termination prover The constrained transition relation generated in PRESYNTHPHASE can be complicated. We found that calling the termination prover on this relation and letting the PRESYNTHPHASE only deal with the counterexamples returns improves the preconditions we compute. Thus, there is a mutual recursion between calling the termination prover and PRESYNTHPHASE.

5 Conclusion

This paper has described an automatic method for finding sound underapproximations to weakest preconditions to termination. Using illustrative examples we have shown that the method can be used to find *useful* underapproximations (*i.e.* something larger than `false`, but smaller than the often complex *weakest* precondition). Beyond the straightforward use of proving termination of programs that do not quite guarantee termination, we believe that our method can also be used in several areas of program verification, including the disproving of safety properties, interprocedural analysis, interprocedural termination proving, and distributed termination proving.

References

1. I. Balaban, A. Pnueli, and L. D. Zuck. Modular ranking abstraction. *Int. J. Found. Comput. Sci.*, 2007.
2. T. Ball, E. Bounimova, B. Cook, V. Levin, J. Lichtenberg, C. McGarvey, B. Ondrusek, S. K. Rajamani, and A. Ustuner. Thorough static analysis of device drivers. In *EuroSys*, 2006.

3. T. Ball, O. Kupferman, and M. Sagiv. Leaping loops in the presence of abstraction. In *CAV*, 2007.
4. A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, and T. Vojnar. Programs with lists are counter automata. In *CAV*, 2006.
5. A. Bouajjani, A. Legay, and P. Wolper. Handling liveness properties in mega-regular model checking. *Electr. Notes Theor. Comput. Sci.*, 2005.
6. A. Bradley. Implementation of Cooper's algorithm. <http://theory.stanford.edu/~arbrad/>.
7. A. Bradley, Z. Manna, and H. Sipma. The polyranking principle. In *ICALP*, 2005.
8. A. Bradley, Z. Manna, and H. Sipma. Termination of polynomial programs. In *VMCAI*, 2005.
9. A. R. Bradley, Z. Manna, and H. B. Sipma. Linear ranking with reachability. In *CAV*, 2005.
10. M. Bruynooghe, M. Codish, J. P. Gallagher, S. Genaim, and W. Vanhoof. Termination analysis of logic programs through combination of type-based norms. *TOPLAS*, 29(2), 2007.
11. C. Calcagno, D. Distefano, P. O'Hearn, and H. Yang. Footprint analysis: A shape analysis that discovers preconditions. In *SAS*, 2007.
12. B. Cook, A. Gotsman, A. Podelski, A. Rybalchenko, and M. Vardi. Proving that programs eventually do something good. In *POPL*, 2007.
13. B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *PLDI*, 2006.
14. B. Cook, A. Podelski, and A. Rybalchenko. Proving thread termination. In *PLDI*, 2007.
15. D. C. Cooper. Theorem proving in arithmetic without multiplication. *Machine Intelligence*, 1972.
16. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, 1977.
17. P. Cousot and R. Cousot. Abstract interpretation and application to logic programs. *J. Log. Program.*, 13(2&3), 1992.
18. S. Gulwani, S. Srivastava, and R. Venkatesan. Program analysis as constraint solving. In *To appear in PLDI*, 2008.
19. G. Lalire, M. Argoud, and B. Jeannet. Interproc analyzer. <http://bjeannet.gforge.inria.fr/interproc/>, 2008.
20. C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. In *POPL*, 2001.
21. R. Leino. Efficient weakest preconditions. *Information Processing Letters*, 93(6), 2005.
22. A. Miné. The Octagon abstract domain. *Higher-Order and Symbolic Computation*, 2006.
23. A. Podelski and A. Rybalchenko. A complete method for the synthesis of linear ranking functions. In *VMCAI*, 2004.
24. A. Rybalchenko. ARMC. <http://www.mpi-sws.org/~rybal/armc/>, 2008.
25. V. Schuppan and A. Biere. Liveness checking as safety checking for infinite state spaces. *Electr. Notes Theor. Comput. Sci.*, 2006.
26. N. Suzuki and K. Ishihata. Implementation of an array bound checker. In *POPL*, 1977.
27. M. Y. Vardi. Verification of concurrent programs: The automata-theoretic framework. In *LICS*, 1987.