Accelerating Interpolants *

Hossein Hojjat[1], Radu Iosif[2],
Filip Konečný[2,4], Viktor Kuncak[1], and Philipp Rümmer[3]

[1] Swiss Federal Institute of Technology Lausanne (EPFL)
[2] Verimag, Grenoble, France
[3] Uppsala University, Sweden
[4] Brno University of Technology, Czech Republic

Abstract. We present Counterexample-Guided *Accelerated* Abstraction Refinement (CEGAAR), a new algorithm for verifying infinite-state transition systems. CEGAAR combines *interpolation-based predicate discovery* in counterexample-guided predicate abstraction with *acceleration* technique for computing the transitive closure of loops. CEGAAR applies acceleration to dynamically discovered looping patterns in the unfolding of the transition system, and combines over-approximation with underapproximation. It constructs inductive invariants that rule out an infinite family of spurious counterexamples, alleviating the problem of divergence in predicate abstraction without losing its adaptive nature. We present theoretical and experimental justification for the effectiveness of CEGAAR, showing that inductive interpolants can be computed from classical Craig interpolants and transitive closures of loops. We present an implementation of CEGAAR that verifies integer transition systems. We show that the resulting implementation robustly handles a number of difficult transition systems that cannot be handled using interpolation-based predicate abstraction or acceleration alone.

1 Introduction

This paper contributes to the fundamental problem of precise reachability analysis for infinite-state systems. Predicate abstraction using interpolation has emerged as an effective technique in this domain. The underlying idea is to verify a program by reasoning about its *abstraction* that is easier to analyse, and is defined with respect to a set of predicates [17]. The set of predicates is refined to achieve the precision needed to prove the absence or the presence of errors. A key difficulty in this approach is to automatically find predicates to make the abstraction sufficiently precise [2]. A breakthrough technique is to generate predicates based on *Craig interpolants* [13] derived from the proof of unfeasibility of a spurious trace [19].

While empirically successful on a variety of domains, abstraction refinement using interpolants suffers from the unpredictability of interpolants computed by provers, which can cause the verification process to diverge and never discover a sufficient set

* Supported by the Rich Model Toolkit initiative, `http://richmodels.org`, the Czech Science Foundation (projects P103/10/0306 and 102/09/H042), the Czech Ministry of Education (COST OC10009 and MSM 0021630528), the BUT project FIT-S-12-1 and the Microsoft Innovation Cluster for Embedded Software.

of predicates (even in case such predicates exist). The failure of such a refinement approach manifests in a sequence of predicates that rule out longer and longer counterexamples, but still fail to discover inductive invariants.

Following another direction, researchers have been making continuous progress on techniques for computing the transitive closure of useful classes of relations on integers [7, 10, 14]. These *acceleration* techniques can compute closed form representation of certain classes of loops using Presburger arithmetic.

A key contribution of this paper is an algorithmic solution to apply these specialized analyses for particular classes of loops to rule out an infinite family of counterexamples during predicate abstraction refinement. An essential ingredient of this approach are interpolants that not only rule out one path, but are also *inductive* with respect to loops along this path. We observe that we can start from any interpolant for a path that goes through a loop in the control-flow graph, and apply a postcondition (or, equivalently a weakest precondition) with respect to the transitive closure of the loop (computed using acceleration) to generalize the interpolant and make it inductive. Unlike previous theoretical proposals [12], our method treats interpolant generation and transitive closure computation as black boxes: we can start from any interpolants and strengthen it using any loop acceleration. We call the resulting technique Counterexample-Guided *Accelerated* Abstraction Refinement, or CEGAAR for short. Our experience indicates that CEGAAR works well in practice.

Motivating Example To illustrate the power of the technique that we propose, consider the example in Figure 1. The example is smaller than the examples we consider in our evaluation (Section 6), but already illustrates the difficulty of applying existing methods.

Note that the innermost loop requires a very expressive logic to describe its closed form, so that standard techniques for computing exact transitive closure of loops do not apply. In particular, the acceleration technique does not apply to the innermost loop, and the presence of the innermost loop prevents the application of acceleration to the outer loop. On the other hand, predicate abstraction with interpolation refinement also fails to solve this example. Namely, it enters a very long refinement loop, considering increasingly longer spurious paths with CFG node sequences of the form $\mathbf{0}(\mathbf{12})^i\mathbf{1e}$, for

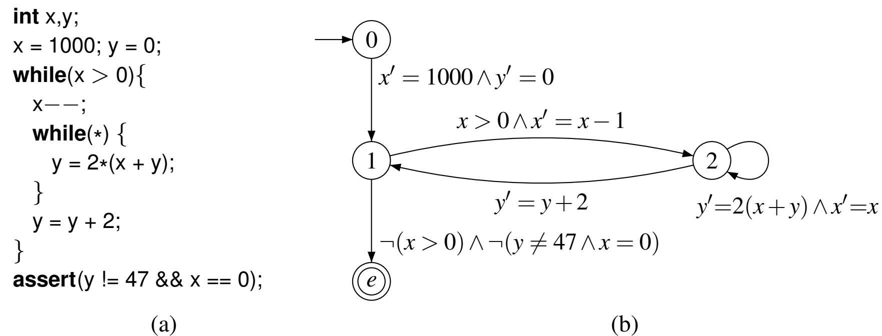

Fig. 1. Example Program and its Control Flow Graph with Large Block Encoding

$0 \leq i < 1000$. The crux of the problem is that the refinement eliminates each of these paths one by one, constructing too specific interpolants.

Our combined CEGAAR approach succeeds in proving the assertion of this program by deriving the loop invariant $y\%2 == 0 \wedge x \geq 0$. Namely, once predicate abstraction considers a path where the CFG node **1** repeats (such as **0121e**), it applies acceleration to this path. CEGAAR then uses the accelerated path to construct an inductive interpolant, which eliminates an infinite family of spurious paths. This provides predicate abstraction with a crucial predicate $y\%2 = 0$, which enables further progress, leading to the discovery of the predicate $x \geq 0$. Together, these predicates allow predicate abstraction to construct the invariant that proves program safety. Note that this particular example focuses on proving the absence of errors, but our experience suggests that CEGAAR can, in many cases, find long counterexamples faster than standard predicate abstraction.

Related Work Predicate abstraction has proved is a rich and fruitful direction in automated verification of detailed properties of infinite-state systems [17, 19]. The pioneering work in [3] is, to the best of our knowledge, the first to propose a solution to the divergence problem in predicate abstraction. More recently, sufficient conditions to enforce convergence of refinement in predicate abstraction are given in [2], but it remains difficult to enforce them in practice. A promising direction for ensuring completeness with respect to a language of invariants is parameterizing the syntactic complexity of predicates discovered by an interpolating *split prover* [21]. Because it has the flavor of invariant enumeration, the feasibility of this approach in practice remains to be further understood.

To alleviate relatively weak guarantees of refinement in predicate abstraction in practice, researchers introduced *path invariants* [5] that rule out a family of counterexamples at once using constraint-based analysis. Our CEGAAR approach is similar in the spirit, but uses acceleration [7,10,14] instead of constraint-based analysis, and therefore has complementary strengths. Acceleration naturally generates precise *disjunctive invariants*, needed in many practical examples, while constraint-based invariant generation [5] resorts to an ad-hoc unfolding of the path program to generate disjunctive invariants. Acceleration can also infer expressive predicates, in particular modulo constraints, which are relevant for purposes such as proving memory address alignment.

The idea of generalizing spurious error traces was introduced also in [18], by extending an infeasible trace, labeled with interpolants, into a finite interpolant automaton. The method of [18] exploits the fact that some interpolants obtained from the infeasibility proof happen to be inductive w.r.t. loops in the program. In our case, given a spurious trace that iterates through a program loop, we *compute* the needed inductive interpolants, combining interpolation with acceleration. The method that is probably closest to CEGAAR is proposed in [12]. In this work the authors define *inductive interpolants* and prove the existence of effectivelly computable inductive interpolants for a class of affine loops, called *poly-bounded*. The approach is, however, limited to programs with one poly-bounded affine loop, for which initial and error states are specified. We only consider loops that are more restricted than the poly-bounded ones, namely loops for which transitive closures are Presburger definable. On the other hand, our method is more general in that it does not restrict the number of loops occurring in the path pro-

gram, and benefits from regarding both interpolation and transitive closure computation as black boxes. The ability to compute closed forms of certain loops is also exploited in algebraic approaches [6]. These approaches can also naturally be generalized to perform useful over-approximation [1] and under-approximation.

2 Preliminaries

Let $\mathbf{x} = \{x_1, \ldots, x_n\}$ be a set of variables ranging over integer numbers, and $\mathbf{x}'$ be the set $\{x'_1, \ldots, x'_n\}$. A *predicate* is a first-order arithmetic formula P. By $FV(P)$ we denote the set of free variables in P, i.e. variables not bound by a quantifier. By writing $P(\mathbf{x})$ we intend that $FV(P) \subseteq \mathbf{x}$. We write $\perp$ and $\top$ for the boolean constants false and true. A *linear term* t over a set of variables in $\mathbf{x}$ is a linear combination of the form $a_0 + \sum_{i=1}^{n} a_i x_i$, where $a_0, a_1, \ldots, a_n \in \mathbb{Z}$. An *atomic proposition* is a predicate of the form $t \leq 0$, where t is a linear term. *Presburger arithmetic* is the first-order logic over propositions $t \leq 0$; Presburger arithmetic has quantifier elimination and is decidable.For simplicity we consider only formulas in Presburger arithmetic in this paper. A valuation of $\mathbf{x}$ is a function $\nu : \mathbf{x} \to \mathbb{Z}$. If ν is a valuation of $\mathbf{x}$, we denote by $\nu \models P$ the fact that the formula obtained by replacing each occurrence of x_i with $\nu(x_i)$ is valid. Similarly, an arithmetic formula $R(\mathbf{x}, \mathbf{x}')$ defining a relation $R \subseteq \mathbb{Z}^n \times \mathbb{Z}^n$ is evaluated referring to two valuations ν_1, ν_2; the satisfaction relation is denoted $\nu_1, \nu_2 \models R$. The composition of two relations $R_1, R_2 \in \mathbb{Z}^n \times \mathbb{Z}^n$ is denoted by $R_1 \circ R_2 = \{(\mathbf{u}, \mathbf{v}) \in \mathbb{Z}^n \times \mathbb{Z}^n \mid \exists \mathbf{t} \in \mathbb{Z}^n \,.\, (\mathbf{u}, \mathbf{t}) \in R_1 \text{ and } (\mathbf{t}, \mathbf{v}) \in R_2\}$. Let ε be the identity relation $\{(\mathbf{u}, \mathbf{u}) \mid \mathbf{u} \in \mathbb{Z}^n \times \mathbb{Z}^n\}$. We define $R^0 = \varepsilon$ and $R^i = R^{i-1} \circ R$, for any $i > 0$. With these notations, $R^+ = \bigcup_{i=1}^{\infty} R^i$ denotes the *transitive closure* of R, and $R^* = R^+ \cup \varepsilon$ denotes the *reflexive and transitive closure* of R. We sometimes use the same symbols to denote a relation and its defining formula. For a set of n-tuples $S \subseteq \mathbb{Z}^n$ and a relation $R \subseteq \mathbb{Z}^n \times \mathbb{Z}^n$, let $post(S, R) = \{\mathbf{v} \in \mathbb{Z}^n \mid \exists \mathbf{u} \in S \,.\, (\mathbf{u}, \mathbf{v}) \in R\}$ denote the *strongest postcondition* of S via R, and $wpre(S, R) = \{\mathbf{u} \in \mathbb{Z}^n \mid \forall \mathbf{v} \,.\, (\mathbf{u}, \mathbf{v}) \in R \to \mathbf{v} \in S\}$ denote the *weakest precondition* of S with respect to R. We use *post* and *wpre* for sets and relations, as well as for logical formulae defining them.

We represent programs as control flow graphs. A *control flow graph* (CFG) is a tuple $G = \langle \mathbf{x}, Q, \to, I, E \rangle$ where $\mathbf{x} = \{x_1, \ldots, x_n\}$ is a set of variables, Q is a set of *control states*, $\to$ is a set of edges of the form $q \xrightarrow{R} q'$, labeled with arithmetic formulae defining relations $R(\mathbf{x}, \mathbf{x}')$, and $I, E \subseteq Q$ are sets of *initial* and *error* states, respectively. A *path* in G is a sequence $\theta : q_1 \xrightarrow{R_1} q_2 \xrightarrow{R_2} q_3 \ldots q_{n-1} \xrightarrow{R_{n-1}} q_n$, where $q_1, q_2, \ldots, q_n \in Q$ and $q_i \xrightarrow{R_i} q_{i+1}$ is an edge in G, for each $i = 1, \ldots, n-1$. We assume without loss of generality that all variables in $\mathbf{x} \cup \mathbf{x}'$ appear free in each relation labeling an edge of G[5]. We denote the relation $R_1 \circ R_2 \circ \ldots \circ R_{n-1}$ by $\rho(\theta)$ and assume that the set of free variables of $\rho(\theta)$ is $\mathbf{x} \cup \mathbf{x}'$. The path θ is said to be a *cycle* if $q_1 = q_n$, and a *trace* if $q_1 \in I$. The path θ is said to be *feasible* if and only if there exist valuations $\nu_1, \ldots, \nu_n : \mathbf{x} \to \mathbb{Z}$ such that $\nu_i, \nu_{i+1} \models R_i$, for all $i = 1, \ldots, n-1$. A control state is said to be *reachable* in G if it occurs on a feasible trace.

[5] For variables that are not modified by a transition, this can be achieved by introducing an explicit update $x' = x$.

Acceleration The goal of acceleration is, given a relation R in a fragment of integer arithmetic, to compute its reflexive and transitive closure, R^*. In general, defining R^* in a decidable fragment of integer arithmetic is not possible, even when R is definable in a decidable fragment such as, e.g. Presburger arithmetic. In this work we consider two fragments of arithmetic in which transitive closures of relations are Presburger definable.

An *octagonal relation* is a relation defined by a constraint of the form $\pm x \pm y \leq c$, where x and y range over the set $\mathbf{x} \cup \mathbf{x}'$, and c is an integer constant. The transitive closure of an octagonal relation has been shown to be Presburger definable and effectively computable [10]. A *linear affine relation* is a relation of the form $\mathcal{R}(\mathbf{x}, \mathbf{x}') \equiv C\mathbf{x} \geq \mathbf{d} \wedge \mathbf{x}' = A\mathbf{x} + \mathbf{b}$, where $A \in \mathbb{Z}^{n \times n}$, $C \in \mathbb{Z}^{p \times n}$ are matrices and $\mathbf{b} \in \mathbb{Z}^n$, $\mathbf{d} \in \mathbb{Z}^p$. $\mathcal{R}$ is said to have the *finite monoid property* if and only if the set $\{A^i \mid i \geq 0\}$ is finite. It is known that the finite monoid condition is decidable [7], and moreover that the transitive closure of a finite monoid affine relation is Presburger definable and effectively computable [7, 14].

Predicate Abstraction Informally, predicate abstraction computes an overapproximation of the transition system generated by a program and verifies whether an error state is reachable in the abstract system. If no error occurs in the abstract system, the algorithm reports that the original system is safe. Otherwise, if a path to an error state (counterexample) has been found in the abstract system, the corresponding concrete path is checked. If this latter path corresponds to a real execution of the system, then a real error has been found. Otherwise, the abstraction is refined in order to exclude the counterexample, and the procedure continues.

Given a CFG $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$, and a (possibly infinite) set of predicates $\mathcal{P}$, an *abstract reachability tree* (ART) for G is a tuple $T = \langle S, \pi, r, e \rangle$ where $S \subseteq Q \times 2^{\mathcal{P} \setminus \{\perp\}}$ is a set of nodes (notice that for no node $\langle q, \Phi \rangle$ in T we may have $\perp \in \Phi$), $\pi : Q \rightarrow 2^{\mathcal{P}}$ is a mapping associating control states with sets of predicates, $i \in I \times \{\top\}$ is the root node, $e \subseteq S \times S$ is a tree-structured edge relation:

- all nodes in S are reachable from the root r
- for all $n, m, p \in S$, $e(n, p) \wedge e(m, p) \Rightarrow n = m$
- $e(\langle q_1, \Phi_1 \rangle, \langle q_2, \Phi_2 \rangle) \;\Rightarrow\; q_1 \xrightarrow{R} q_2$ and $\Phi_2 = \{P \in \pi(q_2) \mid post(\bigwedge \Phi_1, R) \rightarrow P\}$

We say that an ART node $\langle q_1, \Phi_1 \rangle$ is *subsumed* by another node $\langle q_2, \Phi_2 \rangle$ if and only if $q_1 = q_2$ and $\bigwedge \Phi_1 \rightarrow \bigwedge \Phi_2$. It is usually considered that no node in an ART is subsumed by another node, from the same ART.

It can be easily checked that each path $\sigma : r = \langle q_1, \Phi_1 \rangle, \langle q_2, \Phi_2 \rangle, \ldots, \langle q_k, \Phi_k \rangle$, starting from the root in T, can be mapped into a trace $\theta : q_1 \xrightarrow{R_1} q_2 \ldots q_{k-1} \xrightarrow{R_{k-1}} q_k$ of G, such that $post(\top, \rho(\theta)) \rightarrow \bigwedge \Phi_k$. We say that θ is a *concretization* of σ, or that σ concretizes to θ. A path in an ART is said to be *spurious* if none of its concretizations is feasible.

3 Interpolation-Based Abstraction Refinement

By *refinement* we understand the process of enriching the predicate mapping π of an ART $T = \langle S, \pi, r, e \rangle$ with new predicates. The goal of refinement is to prevent spurious

counterexamples (paths to an error state) from appearing in the ART. To this end, an effective technique used in many predicate abstraction tools is that of *interpolation*.

Given an unsatisfiable conjunction $A \wedge B$, an interpolant I is a formula using the common variables of A and B, such that $A \rightarrow I$ is valid and $I \wedge B$ is unsatisfiable. Intuitively, I is the explanation behind the unsatisfiability of $A \wedge B$. Below we introduce a slightly more general definition of a *trace interpolant*.

Definition 1 ([21]). *Let $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$ be a CFG and*

$$\theta : q_1 \xrightarrow{R_1} q_2 \xrightarrow{R_2} q_3 \ldots q_{n-1} \xrightarrow{R_{n-1}} q_n$$

be an infeasible trace of G. An interpolant *for θ is a sequence of predicates $\langle I_1, I_2, \ldots, I_n \rangle$ with free variables in $\mathbf{x}$, such that: $I_1 = \top$, $I_n = \bot$, and for all $i = 1, \ldots, n-1$, $post(I_i, R_i) \rightarrow I_{i+1}$.*

Interpolants exist for many theories, including all theories with quantifier elimination, and thus for Presburger arithmetic. Moreover, a trace is infeasible if and only if it has an interpolant. Including any interpolant of an infeasible trace into the predicate mapping of an ART suffices to eliminate any abstraction of the trace from the ART. We can thus refine the ART and exclude an infeasible trace by including the interpolant that proves the infeasibility of the trace.

Note that the refinement technique using Definition 1 only guarantees that *one* spurious counterexample is eliminated from the ART with each refinement step. This fact hinders the efficiency of predicate abstraction tools, which must rely on the ability of theorem provers to produce interpolants that are general enough to eliminate more than one spurious counterexample at the time. The following is a stronger notion of an interpolant, which ensures generality with respect to an infinite family of counterexamples.

Definition 2 ([12], Def. 2.4). *Given a CFG G, a* trace scheme *in G is a sequence:*

$$\xi : q_0 \xrightarrow{Q_1} \overset{L_1}{\overset{\curvearrowright}{q_1}} \xrightarrow{Q_2} \ldots \xrightarrow{Q_{n-1}} \overset{L_{n-1}}{\overset{\curvearrowright}{q_{n-1}}} \xrightarrow{Q_n} \overset{L_n}{\overset{\curvearrowright}{q_n}} \xrightarrow{Q_{n+1}} q_{n+1} \tag{1}$$

where $q_0 \in I$ and:

- *$Q_i = \rho(\theta_i)$, for some non-cyclic paths θ_i of G, from q_{i-1} to q_i*
- *$L_i = \bigvee_{j=1}^{k_i} \rho(\lambda_{ij})$, for some cycles λ_{ij} of G, from q_i to q_i*

Intuitivelly, a trace scheme represents an infinite regular set of traces in G. The trace scheme is said to be *feasible* if and only if at least one trace of G of the form $\theta_1; \lambda_{1i_1} \ldots \lambda_{1i_{j_1}}; \theta_2; \ldots; \theta_n; \lambda_{ni_n} \ldots \lambda_{ni_{j_n}}; \theta_{n+1}$ is feasible.

The trace scheme is said to be *bounded* if $k_i = 1$, for all $i = 1, 2, \ldots, n$. A bounded[6] trace scheme is a regular language of traces, of the form $\sigma_1 \cdot \lambda_1^* \cdot \ldots \cdot \sigma_n \cdot \lambda_n^* \cdot \sigma_{n+1}$, where σ_i are acyclic paths, and λ_i are cycles of G.

Definition 3 ([12], Def. 2.5). *Let $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$ be a CFG and ξ be an infeasible trace scheme of the form (1). An* interpolant *for ξ is a sequence of predicates $\langle I_0, I_1, I_2, \ldots, I_n, I_{n+1} \rangle$, with free variables in $\mathbf{x}$, such that:*

[6] This term is used in analogy with the notion of bounded languages [16].

1. $I_0 = \top$ *and* $I_{n+1} = \bot$
2. $post(I_i, Q_{i+1}) \rightarrow I_{i+1}$*, for all* $i = 0, 1, \ldots, n$
3. $post(I_i, L_i) \rightarrow I_i$*, for all* $i = 1, 2, \ldots, n$

The main difference with Definition 1 is the third requirement, namely that each interpolant predicate (except for the first and the last one) must be *inductive* with respect to the corresponding loop relation. It is easy to see that each of the two sequences:

$$\langle \top,\ post(\top, Q_1 \circ L_1^*),\ \ldots,\ post(\top, Q_1 \circ L_1^* \circ Q_2 \circ \ldots Q_n \circ L_n^*) \rangle \quad (2)$$

$$\langle wpre(\bot, Q_1 \circ L_1^* \circ Q_2 \circ \ldots Q_n \circ L_n^*),\ \ldots,\ wpre(\bot, Q_n \circ L_n^*),\ \bot \rangle \quad (3)$$

are interpolants for ξ, provided that ξ is infeasible (Lemma 2.6 in [12]). Just as for finite trace interpolants, the existence of an inductive interpolant suffices to prove the infeasibility of the entire trace scheme.

Lemma 4. *Let* $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$ *be a CFG and* ξ *be an infeasible trace scheme of* G *of the form (1). If* $T = \langle S, \pi, r, e \rangle$ *is an ART for* G*, such that there exists an interpolant* $\langle I_i \in \pi(q_i) \rangle_{i=0}^{n+1}$ *for* ξ*, then no path in* T *concretizes to a trace in* ξ*.*

4 Counterexample-Guided Accelerated Abstraction Refinement

This section presents the CEGAAR algorithm for predicate abstraction with interpolant-based accelerated abstraction refinement. Since computing the interpolant of a trace scheme is typically more expensive than computing the interpolant of a finite counterexample, we apply acceleration in a demand-driven fashion. The main idea of the algorithm is to accelerate only those counterexamples in which some cycle repeats a certain number of times. For example, if the abstract state exploration has already ruled out the spurious counterexamples $\sigma \cdot \tau$, $\sigma \cdot \lambda \cdot \tau$ and $\sigma \cdot \lambda \cdot \lambda \cdot \tau$, when it sees next the spurious counterexample $\sigma \cdot \lambda \cdot \lambda \cdot \lambda \cdot \tau$, it will accelerate it into $\sigma \cdot \lambda^* \cdot \tau$, and rule out all traces which comply to this scheme. The maximum number of cycles that are allowed to occur in the acyclic part of an error trace, before computing the transitive closure, is called the *delay*, and is a parameter of the algorithm (here the delay was 2). A smaller delay results in a more aggressive acceleration strategy, whereas setting the delay to infinity is equivalent to performing predicate abstraction without acceleration.

The main procedure is CONSTRUCTART which builds an ART for a given CFG, and an abstraction of the set of initial values (Fig. 2). CONSTRUCTART is a worklist algorithm that expands the ART according to a certain exploration strategy (depth-first, breadth-first, etc.) determined by the type of the structure used as a worklist. We assume without loss of generality that the CFG has exactly one initial vertex *Init*. The CONSTRUCTART procedure starts with *Init* and expands the tree according to the definition of the ART (lines 11 and 12). New ART nodes are constructed using NEWARTNODE, which receives a CFG state and a set of predicates as arguments. The algorithm backtracks from expanding the ART when either the current node contains $\bot$ in its set of predicates, or it is subsumed by another node in the ART (line 13). In the algorithm (Fig. 2), we denote logical entailment by $\phi \vdash \psi$ in order to avoid confusion.

```
1  input CFG G = ⟨x, Q, →, {Init}, E⟩
2  output ART T = ⟨S, π, Root, e⟩
3  WorkList = [], S, π, e = ∅, Root = nil
4  def ConstructART(Init, initialAbstraction) {
5    node = newARTnode(Init, initialAbstraction)
6    if (Root = nil) Root = node
7    WorkList.add(⟨Init, node⟩)
8    while (!(WorkList.empty)) {
9      ⟨nextCFGvertex, nextARTnode⟩ = WorkList.remove()
10     for (child = children(nextCFGVertex)) {
11       Let R be such that nextCFGvertex -R-> child in G
12       Φ = {p ∈ π(child) | POST(⋀nextARTnode.abstraction, R) ⊢ p}
13       if (⊥ ∈ Φ or (∃ an ART node⟨child, Ψ⟩ . ∧Φ ⊢ Ψ))
14         continue
15       node = newARTnode(child, Φ)
16       S = S ∪ {node}
17       e = e ∪ {(nextARTnode, node)}
18       if (child ∈ E and checkRefineError(node))
19         report "ERROR"
20       WorkList.add(⟨child, node⟩)
21       WorkList.removeAll(nodes from WorkList subsumed by node) }}}
```

Fig. 2. The CEGAAR algorithm (a) - High-Level Structure

The refinement step is performed by the CHECKREFINEERROR function (Fig. 3). This function returns true if and only if a feasible error trace has been detected; otherwise, further predicates are generated to refine the abstraction. First, a minimal infeasible ART path to *node* is determined (line 4). This path is generalized into a trace scheme (line 6). The generalization function FOLD takes *Path* and the delay parameter δ as input and produces a trace scheme which contains *Path*. The FOLD function creates a trace scheme of the form (1) out of the spurious path given as argument. The spurious path is traversed and control states are recorded in a list. When we encounter a control state which is already in the list, we identified an elementary cycle λ. If the current trace scheme ends with at least δ occurrences of λ, where $\delta \in \mathbb{N}_\infty$ is the delay parameter, then λ is added as a loop to the trace scheme, provided that its transitive closure can be effectivelly computed. For efficiency reasons, we syntactically check the relation on the loop, namely we check whether the relation is syntactically compliant to the definition of octagonal relations. Notice that a relation can be definable by an octagonal constraint even if it is not a conjunction of octagonal constraints, i.e. it may contain redundant atomic propositions which are not of this form. Once the folded trace scheme is obtained, there are three possibilities:

1. If the trace scheme is not bounded (the test on line 7 passes), we compute a bounded overapproximation of it, in an attempt to prove its infeasibility (line 8). If the test on line 9 succeeds, the original trace scheme is proved to be infeasible and the ART is refined using the interpolants for the overapproximated trace scheme.

```
1  def checkRefineError(node): Boolean {
2    traceScheme = []
3    while (the ART path Root→···→node is spurious) {
4      Let Path = ⟨q1,Φ1⟩→...→⟨qn,Φn⟩ be the (unique) minimal ART path with
5      pivot = ⟨q1,Φ1⟩ and ⟨qn,Φn⟩ = node such that the CFG path q1→···→qn is infeasible
6      newScheme = Fold(Path, delay)
7      if (!isBounded(newScheme)) {
8        absScheme = Concat(Overapprox(newScheme), traceScheme)
9        if (interpolateRefine(absScheme, pivot) return false
10         else newScheme = Underapprox(newScheme, Path)}
11     traceScheme = Concat(newScheme, traceScheme)
12     if (interpolateRefine(traceScheme, pivot) return false
13     node = Path.head}
14   return true }
```

Fig. 3. The CEGAAR algorithm (b) - Accelerated Refinement

2. Else, if the overapproximation was found to be feasible, it could be the case that the abstraction of the scheme introduced a spurious error trace. In this case, we compute a bounded underapproximation of the trace scheme, which contains the initial infeasible path, and replace the current trace scheme with it (line 10). The only requirement we impose on the UNDERAPPROX function is that the returned bounded trace scheme contains *Path*, and is a subset of *newScheme*.
3. Finally, if the trace scheme is bounded (either because the test on line 7 failed, or because the folded path was replaced by a bounded underapproximation on line 10) and also infeasible (the test on line 12 passes) then the ART is refined with the interpolants computed for the scheme. If, on the other hand, the scheme is feasible, we continue searching for an infeasible trace scheme starting from the head of *Path* upwards (line 13).

Example Let $\theta : q_1 \xrightarrow{P} q_2 \xrightarrow{Q} q_2 \xrightarrow{R} q_1 \xrightarrow{P} q_2 \xrightarrow{R} q_1$ be a path. The result of applying FOLD to this path is the trace scheme ξ shown in the left half of Fig. 4. Notice that this path scheme is not bounded, due to the presence of two loops starting and ending with q_2. A possible bounded underapproximation of ξ, containing the original path θ, is shown in the right half of Fig. 4. □

The iteration stops either when a refinement is possible (lines 9, 12), in which case CHECKREFINEERROR returns false, or when the search reaches the root of the ART

$$q_1 \xrightarrow{P} \overset{Q}{\curvearrowright}{q_2} \xrightarrow{R} q_1 \qquad\qquad q_1 \xrightarrow{P} \overset{Q}{\curvearrowright}{q_2} \xrightarrow{\varepsilon} q_2 \xrightarrow{R} q_1$$

$$P \uparrow\downarrow R \qquad\qquad P \uparrow\downarrow R$$

$$q_1 \qquad\qquad q_1$$

Fig. 4. Underapproximation of unbounded trace schemes. ε stands for the identity relation.

```
1  def InterpolateRefine(traceScheme, Pivot) : Boolean {
2    metaTrace = TransitiveClosure(traceScheme)
3    interpolant = InterpolatingProverCall(metaTrace)
4    if (interpolant = ∅) return false
5    I = AccelerateInterpolant(interpolant)
6    for (ψ ∈ I) {
7      let v be the CFG vertex corresponding to ψ
8      π = π[v ← (π(v) ∪ ψ)]
9    }
10   ConstructART(Pivot,Pivot.abstraction)
11   return true }
```

Fig. 5. The Interpolation Function

and the trace scheme is feasible, in which case CHECKREFINEERROR returns true (line 14) and the main algorithm in Figure 2 reports a true counterexample. Notice that, since we update *node* to the head of *Path* (line 13), the position of *node* is moved upwards in the ART. Since this cannot happen indefinitely, the main loop (lines 3-13) of the CHECKREFINEERROR is bound to terminate.

The INTERPOLATEREFINE function is used to compute the interpolant of the trace scheme, update the predicate mapping π of the ART, and reconstruct the subtree of the ART whose root is the first node on *Path* (this is usually called the *pivot* node). The INTERPOLATEREFINE (Fig. 5) function returns true if and only if its argument represents an infeasible trace scheme. In this case, new predicates, obtained from the interpolant of the trace scheme, are added to the nodes of the ART. This function uses internally the TRANSITIVECLOSURE procedure (line 2) in order to generate the meta-trace scheme (5). The ACCELERATEINTERPOLANT function (line 5) computes the interpolant for the trace scheme, from the resulting meta-trace scheme. Notice that the refinement algorithm is recursive, as CONSTRUCTART calls CHECKREFINEERROR (line 18), which in turn calls INTERPOLATEREFINE (lines 9,12), which calls back CONSTRUCTART (line 10). Our procedure is *sound,* in the sense that whenever function CONSTRUCTART terminates with a non-error result, the input program does not contain any reachable error states. Vice versa, if a program contains a reachable error state, CONSTRUCTART is guaranteed to eventually discover a feasible path to this state, since the use of a work list ensures fairness when exploring ARTs.

5 Computing Accelerated Interpolants

This section describes a method of refining an ART by excluding an infinite family of infeasible traces at once. Our method combines interpolation with acceleration in a way which is oblivious of the particular method used to compute interpolants. For instance, it is possible to combine proof-based [23] or constraint-based [26] interpolation with acceleration, whenever computing the precise transitive closure of a loop is possible. In cases when the precise computation fails, we may resort to both over- and under-approximation of the transitive closure. In both cases, the accelerated interpolants are at

least as general (and many times more general) than the classical interpolants extracted from a finite counterexample trace.

5.1 Precise Acceleration of Bounded Trace Schemes

We consider first the case of bounded trace schemes of the form (1), where the control states $q_1, \ldots, q_n$ belong to some cycles labeled with relations $L_1, \ldots, L_n$. Under some restrictions on the syntax of the relations labeling the cycles L_i, the reflexive transitive closures L_i^* are effectively computable using acceleration algorithms [7, 9, 14]. Among the known classes of relations for which acceleration is possible we consider: *octagonal relations* and *finite monoid affine transformations*. These are all conjunctive linear relations. We consider in the following that all cycle relations L_i belong to one of these classes. Under this restriction, any infeasible bounded trace scheme has an effectivelly computable interpolant of one of the forms (2),(3).

However, there are two problems with applying definitions (2),(3) in order to obtain interpolants of trace schemes. On one hand, relational composition typically requires expensive quantifier eliminations. The standard proof-based interpolation techniques (e.g. [23]) overcome this problem by extracting the interpolants directly from the proof of infeasibility of the trace. Alternatively, constraint-based interpolation [26] reduce the interpolant computation to a Linear Programming problem, which can be solved by efficient algorithms. Both methods apply, however, only to finite traces, and not to infinite sets of traces defined as trace schemes. Another, more important, problem is related to the sizes of the interpolant predicates from (2), (3) compared to the sizes of interpolant predicates obtained by proof-theoretic methods (e.g. [22]), as the following example shows.

Example Let $R(x,y,x',y') \; : \; x' = x+1 \wedge y' = y+1$ and $\phi(x,y,\ldots)$, $\psi(x,y,\ldots)$ be some complex Presburger arithmetic formulae. The trace scheme:

$$q_0 \xrightarrow{z=0 \wedge z'=z \wedge \phi} \overset{z'=z+2 \wedge R}{\overset{\curvearrowright}{q_1}} \xrightarrow{z=5 \wedge \psi} q_2 \tag{4}$$

is infeasible, because z remains even, so it cannot become equal 5. One simple interpolant for this trace scheme has at program point q_1 the formula $z\%2 = 0$. On the other hand, the strongest interpolant has $(z = 0 \wedge z' = x \wedge \phi) \circ (z' = z+2 \wedge R)^*$ at q_1, which is typically a much larger formula, because of the complex formula ϕ. Note however that ϕ and R do not mention z, so they are irrelevant. □

To construct useful interpolants instead of the strongest or the weakest ones, we therefore proceed as follows. Let ξ be a bounded trace scheme of the form (1). For each control loop $q_i \xrightarrow{R_i} q_i$ of ξ, we define the corresponding *meta-transition* $q_i' \xrightarrow{R_i^*} q_i''$ labeled with the reflexive and transitive closure of R_i. Intuitively, firing the meta-transition has the same effect as iterating the loop an arbitrary number of times. We first replace each loop of ξ by the corresponding meta-transition. The result is the *meta-trace*:

$$\overline{\xi} : q_0 \xrightarrow{Q_1} q_1' \xrightarrow{L_1^*} q_1'' \xrightarrow{Q_2} q_2' \; \ldots \; q_{n-1}'' \xrightarrow{Q_n} q_n' \xrightarrow{L_n^*} q_n'' \xrightarrow{Q_{n+1}} q_{n+1} \tag{5}$$

Since we supposed that ξ is an infeasible trace scheme, the (equivalent) finite meta-trace $\overline{\xi}$ is infeasible as well, and it has an interpolant $I_{\overline{\xi}} = \langle \top, I_1', I_1'', I_2', I_2'', \ldots, I_n', I_n'', \bot \rangle$

in the sense of Definition 1. This interpolant is not an interpolant of the trace scheme ξ, in the sense of Definition 3. In particular, none of I'_i, I''_i is guaranteed to be inductive with respect to the loop relations L_i. To define compact inductive interpolants based on $I_{\overline{\xi}}$ and the transitive closures L_i^*, we consider the following sequences:

$$\begin{aligned} I_\xi^{post} &= \langle \top, post(I'_1, L_1^*), post(I'_2, L_2^*), \ldots, post(I'_n, L_n^*), \bot \rangle \\ I_\xi^{wpre} &= \langle \top, wpre(I''_1, L_1^*), wpre(I''_2, L_2^*), \ldots, wpre(I''_n, L_n^*), \bot \rangle \end{aligned}$$

The following lemma proves the correctness of this approach.

Lemma 5. *Let $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$ be a CFG and ξ be an infeasible trace scheme of the form (1). Then I_ξ^{post} and I_ξ^{wpre} are interpolants for ξ, and moreover $I_{\xi_i}^{wpre} \rightarrow I_{\xi_i}^{post}$, for all $i = 1, 2, \ldots, n$.*

Notice that computing I_ξ^{post} and I_ξ^{wpre} requires n relational compositions, which is, in principle, just as expensive as computing directly one of the extremal interpolants (2),(3). However, by re-using the meta-trace interpolants, one potentially avoids the worst-case combinatorial explosion in the size of the formulae, which occurs when using (2), (3) directly.

Example Let us consider again the trace scheme (4). The corresponding infeasible finite trace $\overline{\xi}$ is:

$$q_0 \xrightarrow{z=0 \wedge z'=z \wedge \phi} q'_1 \xrightarrow{\exists k \geq 0 \ . \ z'=z+2k \ \wedge \ x'=x+k \ \wedge \ y'=y+k} q''_1 \xrightarrow{z=5 \wedge \psi} q_2$$

A possible interpolant for this trace is $\langle \top, z = 0, \exists k \geq 0 \ . \ z = 2k, \bot \rangle$. An inductive interpolant for the trace scheme, derived from it, is $I_\xi^{post} = \langle \top, post(z = 0, \exists k \geq 0 . z' = z + 2k \wedge x' = x + k \wedge y' = y + k), \bot \rangle = \langle \top, \ z\%2 = 0, \ \bot \rangle$. □

5.2 Bounded Overapproximations of Trace Schemes

Consider a trace scheme (1), not necessarily bounded, where the transitive closures of the relations L_i labeling the loops are not computable by any available acceleration method [7, 9, 14]. One alternative is to find abstractions $L_i^\sharp$ of the loop relations, i.e. relations $L_i^\sharp \leftarrow L_i$, for which transitive closures are computable. If the new abstract trace remains infeasible, it is possible to compute an interpolant for it, which is an interpolant for the original trace scheme. However, replacing the relations L_i with their abstractions $L_i^\sharp$ may turn an infeasible trace scheme into a feasible one, where the traces introduced by abstraction are spurious. In this case, we give up the overapproximation, and turn to the underapproximation technique described in the next section.

The overapproximation method computes an interpolant for a trace scheme ξ of the form (1) under the assumption that the abstract trace scheme:

$$\xi^\sharp : q_0 \xrightarrow{Q_1} \overset{L_1^\sharp}{\underset{q_1}{\curvearrowright}} \xrightarrow{Q_2} \ldots \xrightarrow{Q_{n-1}} \overset{L_{n-1}^\sharp}{\underset{q_{n-1}}{\curvearrowright}} \xrightarrow{Q_n} \overset{L_n^\sharp}{\underset{q_n}{\curvearrowright}} \xrightarrow{Q_{n+1}} q_{n+1} \qquad (6)$$

is infeasible. In this case one can effectivelly compute the interpolants $I^{post}_{\xi^\sharp}$ and $I^{wpre}_{\xi^\sharp}$, since the transitive closures of the abstract relations labeling the loops are computable by acceleration. The following lemma proves that, under certain conditions, computing an interpolant for the abstraction of a trace scheme is sound.

Lemma 6. *Let G be a CFG and ξ be a trace scheme (1) such that the abstract trace scheme $\xi^\sharp$ (6) is infeasible. Then the interpolants $I^{post}_{\xi^\sharp}$ and $I^{wpre}_{\xi^\sharp}$ for $\xi^\sharp$ are also interpolants for ξ.*

5.3 Bounded Underapproximations of Trace Schemes

Let ξ be a trace scheme of the form (1), where each relation L_i labeling a loop is a disjunction $L_{i1} \vee \ldots \vee L_{ik_i}$ of relations for which the transitive closures are effectively computable and Presburger definable. A *bounded underapproximation scheme* of a trace scheme ξ is obtained by replacing each loop $q_i \xrightarrow{L_i} q_i$ in ξ by a bounded trace scheme of the form:

$$\overset{L_{i1}}{\curvearrowright}\!\!\!\!\!\!\!\! q_i^1 \xrightarrow{\varepsilon} \overset{L_{i2}}{\curvearrowright}\!\!\!\!\!\!\!\! q_i^2 \xrightarrow{\varepsilon} \ldots \overset{L_{ik_i}}{\curvearrowright}\!\!\!\!\!\!\!\! q_i^{k_i}$$

where ε denotes the identity relation. Let us denote[7] the result of this replacement by $\xi^\flat$. It is manifest that the set of traces $\xi^\flat$ is included in ξ.

Since we assumed that the reflexive and transitive closures L^*_{ij} are effectively computable and Presburger definable, the feasibility of $\xi^\flat$ is a decidable problem. If $\xi^\flat$ is found to be feasible, this points to a real error trace in the system. On the other hand, if $\xi^\flat$ is found to be infeasible, let $I_{\xi^\flat} = \langle \top, I_1^1, \ldots, I_1^{k_1}, \ldots, I_n^1, \ldots, I_n^{k_n}, \bot \rangle$ be an interpolant for $\xi^\flat$. A refinement scheme using this interpolant associates the predicates $\{I_i^1, \ldots, I_i^{k_i}\}$ with the control state q_i from the original CFG. As the following lemma shows, this guarantees that any trace that follows the pattern of $\xi^\flat$ is excluded from the ART, ensuring that a refinement of the ART using a suitable underapproximation (that includes a spurious counterexample) is guaranteed to make progress.

Lemma 7. *Let $G = \langle \mathbf{x}, Q, \rightarrow, I, E \rangle$ be a CFG, ξ be an infeasible trace scheme of G (1) and $\xi^\flat$ a bounded underapproximation of ξ. If $T = \langle S, \pi, r, e \rangle$ is an ART for G, such that $\{I_i^1, \ldots, I_i^{k_i}\} \subseteq \pi(q_i)$, then no path in T concretizes to a trace in $\xi^\flat$.*

Notice that a refinement scheme based on underapproximation guarantees the exclusion of those traces from the chosen underapproximation trace scheme, and not of all traces from the original trace scheme. Since a trace scheme is typically obtained from a finite counterexample, an underapproximation-based refinement still guarantees that the particular counterexample is excluded from further searches. In other words, using underapproximation is still better than the classical refinement method, since it can potentially exclude an entire family of counterexamples (including the one generating the underapproximation) at once.

[7] The choice of the name depends on the ordering of particular paths $L_{i1}, L_{i2}, \ldots, L_{ik_i}$, however we shall denote any such choice in the same way, in order to keep the notation simple.

6 Experimental Results

We have implemented CEGAAR by building on the predicate abstraction engine Eldarica[8] [20], the FLATA verifier[9] [20] based on acceleration, and the Princess interpolating theorem prover [11, 25]. Tables in Figure 6 compares the performance of the Flata, Eldarica, *static acceleration* and CEGAAR on a number of benchmarks (the platorm used for experiments is Intel® Core™2 Duo CPU P8700, 2.53GHz with 4GB of RAM).

Model	Time [s] F.	E.	S.	D.
(a) Examples from [21]				
anubhav (C)	0.8	3.0	4.0	3.1
copy1 (E)	2.0	7.2	5.8	5.9
cousot (C)	0.6	-	6.2	5.9
loop1 (E)	1.7	7.1	5.2	5.4
loop (E)	1.8	5.9	4.8	5.4
scan (E)	3.3	-	5.1	5.0
string_concat1 (E)	5.3	-	10.1	7.3
string_concat (E)	4.9	-	7.0	7.5
string_copy (E)	4.6	-	6.3	5.7
substring1 (E)	0.6	9.4	18.2	8.3
substring (E)	2.1	3.3	6.3	3.5
(b) Verification conditions for array programs [9]				
rotation_vc.1 (C)	0.6	2.0	9.5	2.0
rotation_vc.2 (C)	1.6	2.2	18.5	2.2
rotation_vc.3 (C)	1.2	0.3	18.3	0.3
rotation_vc.1 (E)	1.1	1.3	10.2	1.3
split_vc.1 (C)	3.9	3.7	91.1	3.6
split_vc.2 (C)	3.0	2.3	74.1	2.2
split_vc.3 (C)	3.3	0.6	75.0	0.6
split_vc.1 (E)	28.5	2.3	185.6	2.4

Model	Time [s] F.	E.	S.	D.
(c) Examples from [24]				
boustrophedon (C)	-	-	-	14.4
gopan (C)	0.4	-	-	6.4
halbwachs (C)	-	-	7.3	7.0
rate_limiter (C)	31.7	6.1	8.1	5.5
(d) Examples from L2CA [8]				
bubblesort (E)	14.9	9.9	9.5	9.3
insdel (E)	0.1	1.3	2.5	1.4
insertsort (E)	2.0	4.2	5.0	4.0
listcounter (C)	0.3	-	1.9	3.7
listcounter (E)	0.3	1.4	1.6	1.4
listreversal (C)	4.5	3.0	6.0	3.3
listreversal (E)	0.8	2.7	8.1	2.8
mergesort (E)	1.2	7.7	21.3	7.4
selectionsort (E)	1.5	8.1	13.7	7.7
(e) NECLA benchmarks				
inf1 (E)	0.2	2.0	2.0	2.0
inf4 (E)	0.9	3.7	3.7	3.7
inf6 (C)	0.1	2.0	2.0	2.0
inf8 (C)	0.3	3.6	3.4	3.9

Model	Time [s] F.	E.	S.	D.
(f) Examples from [15]				
h1 (E)	-	5.1	5.6	5.1
h1.optim (E)	0.8	2.9	5.5	2.9
h1h2 (E)	-	9.4	10.1	12.2
h1h2.optim (E)	1.1	3.3	4.4	3.4
simple (E)	-	6.4	7.0	8.4
simple.optim (E)	0.8	3.0	5.1	2.9
test0 (C)	-	23.0	23.4	29.2
test0.optim (C)	0.3	3.2	5.4	3.2
test0 (E)	-	5.4	5.9	5.7
test0.optim (E)	0.6	3.0	5.8	2.9
test1.optim (C)	0.9	4.7	5.9	7.8
test1.optim (E)	1.5	4.4	5.9	4.7
test2_1.optim (E)	1.6	5.2	5.5	5.6
test2_2.optim (E)	2.9	4.6	5.9	4.6
test2.optim (C)	6.4	27.2	30.1	30.0
wrpc.manual (C)	0.6	1.2	1.4	1.2
wrpc (E)	-	7.9	8.4	8.2
wrpc.optim (E)	-	5.1	8.5	5.2
(g) VHDL models from [27]				
counter (C)	0.1	1.6	1.6	1.6
register (C)	0.2	1.1	1.1	1.1
synlifo (C)	16.6	22.1	21.4	22.0

Fig. 6. Benchmarks for **F**lata, **E**ldarica without acceleration, Eldarica with acceleration of loops at the CFG level (**S**tatic) and CEGAAR (**D**ynamic acceleration). The letter after the model name distinguishes **C**orrect from models with a reachable **E**rror state. Items with "-" led to a timeout for the respective approach.

The benchmarks are all in the Numerical Transition Systems format[10] (NTS). We have considered seven sets of examples, extracted automatically from different sources: (a) C programs with arrays provided as examples of divergence in predicate abstraction [21], (b) verification conditions for programs with arrays, expressed in the SIL logic of [9] and translated to NTS, (c) small C programs with challenging loops, (d) NTS extracted from programs with singly-linked lists by the L2CA tool [8], (e) C programs provided as benchmarks in the NECLA static analysis suite, (f) C programs with asynchronous procedure calls translated into NTS using the approach of [15] (the examples with extension .optim are obtained via an optimized translation method [Pierre Ganty, personal communication], and (g) models extracted from VHDL models of circuits following the method of [27]. The benchmarks are available from the home page

[8] `http://lara.epfl.ch/w/eldarica`
[9] `http://www-verimag.imag.fr/FLATA.html`
[10] `http://richmodels.epfl.ch/ntscomp_ntslib`

of our tool. The results on this benchmark set suggest that we have arrived at a fully automated verifier that is robust in verifying automatically generated integer programs with a variety of looping control structure patterns. An important question we explored is the importance of dynamic application of acceleration, as well as of overapproximation and underapproximation. We therefore also implemented static acceleration [12], a lightweight acceleration technique generalizing large block encoding (LBE) [4] with transitive closures. It simplifies the control flow graph prior to predicate abstraction. In some cases, such as mergesort from the (d) benchmarks and split_ vc.1 from (b) benchmarks, the acceleration overhead does not pay off. The problem is that static acceleration tries to accelerate every loop in the CFG rather than accelerating the loops occurring on spurious paths leading to error. Acceleration of inessential loops generates large formulas as the result of combining loops and composition of paths during large block encoding. The CEGAAR algorithm is the only approach that could handle all of our benchmarks. There are cases in which the Flata tool outperforms CEGAAR such as test2.optim from (f) benchmarks. We attribute this deficiency to the nature of predicate abstraction, which tries to discover the required predicates by several steps of refinement. In the verification of benchmarks, acceleration was exact 11 times in total. In 30 case the over-approximation of the loops was successful, and in 15 cases over-approximation failed, so the tool resorted to under-approximation. This suggests that all techniques that we presented are essential to obtain an effective verifier.

7 Conclusions

We have presented CEGAAR, a new automated verification algorithm for integer programs. The algorithm combines two cutting-edge analysis techniques: interpolation-based abstraction refinement and acceleration of loops. We have implemented CEGAAR and presented experimental results, showing that CEGAAR handles robustly a number of examples that cannot be handled by predicate abstraction or acceleration alone. Because many classes of systems translate into integer programs, our advance contributes to automated verification of infinite-state systems in general.

References

1. E. Albert, P. Arenas, S. Genaim, and G. Puebla. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, 46(2), February 2011.
2. T. Ball, A. Podelski, and S. K. Rajamani. Relative completeness of abstraction refinement for software model checking. In *TACAS'02*, volume 2280 of *LNCS*, page 158, 2002.
3. S. Bensalem and Y. Lakhnech. Automatic generation of invariants. *Form. Methods Syst. Des.*, 15(1):75–92, July 1999.
4. D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software model checking via large-block encoding. In *FMCAD*, pages 25–32, 2009.
5. D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Path invariants. In *PLDI*, pages 300–309, 2007.
6. R. Blanc, T. A. Henzinger, T. Hottelier, and L. Kovács. ABC: Algebraic bound computation for loops. In *LPAR (Dakar)*, pages 103–118, 2010.

7. B. Boigelot. *Symbolic Methods for Exploring Infinite State Spaces*, volume PhD Thesis, Vol. 189. Collection des Publications de l'Université de Liège, 1999.
8. A. Bouajjani, M. Bozga, P. Habermehl, R. Iosif, P. Moro, and T. Vojnar. Programs with lists are counter automata. In *CAV*, pages 517–531, 2006.
9. M. Bozga, P. Habermehl, R. Iosif, F. Konečný, and T. Vojnar. Automatic verification of integer array programs. In *CAV*, pages 157–172, 2009.
10. M. Bozga, R. Iosif, and F. Konečný. Fast acceleration of ultimately periodic relations. In *CAV*, pages 227–242, 2010.
11. A. Brillout, D. Kroening, P. Rümmer, and T. Wahl. An interpolating sequent calculus for quantifier-free Presburger arithmetic. In *IJCAR*, LNCS. Springer, 2010.
12. N. Caniart, E. Fleury, J. Leroux, and M. Zeitoun. Accelerating interpolation-based model-checking. In *TACAS'08*, pages 428–442, 2008.
13. W. Craig. Linear reasoning. A new form of the Herbrand-Gentzen theorem. *The Journal of Symbolic Logic*, 22(3):250–268, September 1957.
14. A. Finkel and J. Leroux. How to compose presburger-accelerations: Applications to broadcast protocols. In *FST TCS '02*, pages 145–156. Springer, 2002.
15. P. Ganty and R. Majumdar. Algorithmic verification of asynchronous programs. *CoRR*, abs/1011.0551, 2010.
16. S. Ginsburg and E. Spanier. Bounded algol-like languages. *Trans. of the AMS*, 113(2):333–368, 1964.
17. S. Graf and H. Saidi. Construction of abstract state graphs with PVS. In *CAV*, pages 72–83, 1997.
18. M. Heizmann, J. Hoenicke, and A. Podelski. Refinement of trace abstraction. In *SAS*, SAS'09, pages 69–85, 2009.
19. T. A. Henzinger, R. Jhala, R. Majumdar, and K. L. McMillan. Abstractions from proofs. In *31st POPL*, 2004.
20. H. Hojjat, F. Konečný, F. Garnier, R. Iosif, V. Kuncak, and P. Rümmer. A verification toolkit for numerical transition systems (tool paper). In *FM*, 2012.
21. R. Jhala and K. L. McMillan. A practical and complete approach to predicate refinement. In *TACAS*, 2006.
22. D. Kroening, J. Leroux, and P. Rümmer. Interpolating quantifier-free Presburger arithmetic. In *Proceedings, LPAR*, volume 6397 of *LNCS*, pages 489–503. Springer, 2010.
23. K. L. McMillan. An interpolating theorem prover. *Theor. Comput. Sci.*, 345(1), 2005.
24. D. Monniaux. Personal Communication.
25. P. Rümmer. A constraint sequent calculus for first-order logic with linear integer arithmetic. In *LPAR*, volume 5330 of *LNCS*, pages 274–289. Springer, 2008.
26. A. Rybalchenko and V. Sofronie-Stokkermans. Constraint solving for interpolation. In *Proceedings, VMCAI*, volume 4349 of *LNCS*, pages 346–362. Springer, 2007.
27. A. Smrcka and T. Vojnar. Verifying parametrised hardware designs via counter automata. In *Haifa Verification Conference*, pages 51–68, 2007.