Explicit-Value Analysis Based on CEGAR and Interpolation

Dirk Beyer and Stefan Löwe

University of Passau, Germany

arXiv:1212.6542v1 [cs.SE] 28 Dec 2012

Technical Report, Number MIP-1205
Department of Computer Science and Mathematics
University of Passau, Germany
December 2012

Explicit-Value Analysis Based on CEGAR and Interpolation

Dirk Beyer and Stefan Löwe

University of Passau, Germany

Abstract—**Abstraction, counterexample-guided refinement, and interpolation are techniques that are essential to the success of predicate-based program analysis. These techniques have not yet been applied together to explicit-value program analysis. We present an approach that integrates abstraction and interpolation-based refinement into an explicit-value analysis, i.e., a program analysis that tracks explicit values for a specified set of variables (the precision). The algorithm uses an abstract reachability graph as central data structure and a path-sensitive dynamic approach for precision adjustment. We evaluate our algorithm on the benchmark set of the Competition on Software Verification 2012 (SV-COMP'12) to show that our new approach is highly competitive. In addition, we show that combining our new approach with an auxiliary predicate analysis scores significantly higher than the SV-COMP'12 winner.**

I. Introduction

Abstraction is one of the most important techniques to successfully verify industrial-scale program code, because the abstract model omits details about the concrete semantics of the program that are not necessary to prove or disprove the program's correctness. Counterexample-guided abstraction refinement (CEGAR) [14] is a technique that iteratively refines an abstract model using counterexamples. A counterexample is a witness of a property violation. In software verification, the counterexamples are error paths, i.e., paths through the program that violate the property. CEGAR starts with the most abstract model and checks if an error path can be found. If the analysis of the abstract model does not find an error path, then the analysis terminates, reporting that no violation exists. If the analysis finds an error path, the path is checked for feasibility, i.e., if the path is executable according to the concrete program semantics. If the error path is feasible, the analysis terminates, reporting the violation of the property, together with the feasible error path as witness. If the error path is infeasible, the violation is due to a too coarse abstract model and the infeasible error path is used to automatically refine the current abstraction. Then the analysis proceeds. Several successful tool implementations for software verification are based on abstraction and CEGAR (cf. [4], [6], [10], [13], [16], [25]). Craig interpolation is a technique from logics that yields for two contradicting formulas an interpolant that contains less information than the first formula, but still enough to contradict the second formula [17]. In software verification, interpolation can be used to extract information from infeasible error paths [21], where the resulting interpolants are used to refine the abstract model. Predicate abstraction is a successful abstraction technique for software model checking [18], because its symbolic state representation blends well with strongest post-conditions, and abstractions can be computed efficiently with solvers for satisfiability modulo theories (SMT) [3]. CEGAR and lazy refinement [22] together with interpolation [21] effectively refine abstract models in the predicate domain. The recent competition on software verification (SV-COMP'12 [5], Table 3) shows that these advancements had a strong impact on the success of participating tools (cf. [6], [10], [25], [26]).

Despite the success of abstraction, CEGAR, and interpolation in the field of predicate analysis, these techniques have not yet been combined and applied together to explicit-value analysis. We integrate these three techniques into an explicit-value analysis, a rather unsophisticated analysis that tracks for each program variable its current value explicitly (like constant propagation [1], but without join). First, we have to define the notion of abstraction for the explicit-value domain, and the precision of the analysis (i.e., the level of abstraction) by a set of program variables that the analysis has to track. Second, in order to automatically determine the necessary precision (i.e., a *small* set of program variables that *need* to be tracked) we use CEGAR iterations to discover finer precisions from infeasible error paths. Third, we define interpolation for the explicit-value domain and use this idea to construct an algorithm that efficiently extracts such a parsimonious precision that is sufficient to eliminate infeasible error paths.

Example. Consider the simple example program in Fig. 1. This program contains a $while$ loop in which a system call occurs. The loop exits if either the system call returns 0 or a previously specified number of iterations x was performed. Because the body of the function $system_call$ is unknown, the value of $result$ is unknown. Also, the assumption $[ticks > x]$ cannot be evaluated to $true$, because x is unknown. This program is correct, i.e., the error location in line 10 is not reachable. However, a simple explicit-value model checker that always tracks every variable would unroll the loop, always discovering new states, as the expression $ticks = ticks + 1$ repeatedly assigns new values to variable $ticks$. Thus, due to extreme resource consumptions, the analysis would not terminate within practical time and memory limits, and is bound to give up on proving the safety property, eventually.

The new approach for explicit-value analysis that we propose can efficiently prove this program safe, because it tracks only those variables that are necessary to refute the infeasible error paths. In the first CEGAR iteration, the precision of the analysis is empty, i.e., no variable is tracked. Thus, the

```
1 extern int system_call();
2 int main(int x) {
3   int flag, ticks, result;
4   flag = 0; ticks = 0;
5   while(1) {
6     ticks  = ticks + 1;
7     result = system_call();
8     if(result == 0 || ticks > x) { break; }
9   }
10  if (flag > 0) { ERROR: return 1; }
11 }
```

Fig. 1: Example program to illustrate the effectiveness of CEGAR-based explicit-value analysis

error location will be reached. Now, using our interpolation-inspired method to discover precisions from counterexample paths, the algorithm identifies that the variable $flag$ (more precisely, the constraint $flag = 0$) has to be tracked. The analysis is re-started after this refinement. Because $ticks$ is not in the precision (the variable is not tracked), the assignment $ticks = ticks + 1$ will not add new abstract states. Since no new successors are computed, the analysis stops unrolling the loop. The assume operation $[flag > 0]$ is evaluated to $false$, and thus, the error label is not reachable. The analysis terminates, proving the program correct.

In summary, the crucial effect of this approach is that only relevant variables are tracked in the analysis, while unimportant information is ignored. This greatly reduces the number of abstract states to be visited.

Contributions. We make the following contributions:

- We integrate the concepts of abstraction, CEGAR, and lazy abstraction refinement into explicit-value analysis.
- Inspired by Craig interpolation for predicate analysis, we define a novel interpolation-like approach for discovering relevant variables for the explicit-value domain. This refinement algorithm is completely self-contained, i.e., independent from external libraries such as SMT solvers.
- To further improve the effectiveness and efficiency of the analysis, we design a combination with a predicate analysis based on dynamic precision adjustment [9].
- We provide an open-source implementation of all our concepts and give evidence of the significant improvements by evaluating several approaches on benchmark verification tasks (C programs) from SV-COMP'12.

Related Work. The explicit-state model checker SPIN [23] can verify models of programs written in a language called Promela. For the verification of C programs, tools like MODEX[1] can extract Promela models from C source code. This process requires to give a specification of the abstraction level (user-defined extraction rules), i.e., the information of what should be included in the Promela model. SPIN does not provide lazy-refinement-based CEGAR. JAVA PATHFINDER [20] is an explicit-state model checker for Java programs. There has been work [24] on integrating CEGAR into JAVA PATHFINDER, using an approach different from interpolation.

Program analysis with dynamic precision adjustment [9] is an approach to adjust the precision of combined analyses on-the-fly, i.e., during the analysis run; the precision of one analysis can be increased based on a current situation in another analysis. For example, if an explicit-value analysis stores too many different values for a variable, then the dynamic precision adjustment can remove that variable from the precision of the explicit-value analysis and add a predicate about that variable to the precision of a predicate analysis. This means that the tracking of the variable is "moved" from the explicit domain to the symbolic domain. One configuration that we present later in this paper uses this approach (cf. III-F).

The tool DAGGER [19] improves the verification of C programs by applying interpolation-based refinement to octagon and polyhedra domains. To avoid imprecision due to widening in the join-based data-flow analysis, DAGGER replaces the standard widen operator by a so called *interpolated-widen* operator, which increases the precision of the data-flow analysis and thus avoids false alarms. The algorithm VINTA [2] applies interpolation-based refinement to interval-like abstract domains. If the state exploration finds an error path, then VINTA performs a feasibility check using bounded model checking (BMC), and if the error path is infeasible, it computes interpolants. The interpolants are used to refine the invariants that the abstract domain operates on. VINTA requires an SMT solver for feasibility checks and interpolation.

More tools are mentioned in our evaluation section, where we compare (in terms of precision and efficiency) our tool implementation with tools that participated in SV-COMP'12.

There is, to the best of our knowledge, no work that integrates abstraction, CEGAR, lazy refinement, and interpolation into explicit-state model checking. We make those techniques available for the explicit-value domain.

II. Preliminaries

Our approach is based on several existing concepts, and in this section we remind the reader of some basic definitions.

A. Programs, Control-Flow Automata, States

We restrict the presentation to a simple imperative programming language, where all operations are either assignments or assume operations, and all variables range over integers[2]. The following definitions are taken from previous work [11]: A program is represented by a *control-flow automaton* CFA. A CFA $A = (L, G)$ consists of a set L of program locations, which model the program counter, and a set $G \subseteq L \times Ops \times L$ of control-flow edges, which model the operations that are executed when control flows from one program location to another. The set of program variables that occur in operations from Ops is denoted by X. A *verification problem* $P = (A, l_0, l_e)$ consists of a CFA A, representing the program, an initial program location $l_0 \in L$, representing the program entry, and a target program location $l_e \in L$, which represents the error.

A *concrete data state* of a program is a variable assignment $cd : X \to \mathbb{Z}$, which assigns to each program variable an

[1]http://cm.bell-labs.com/cm/cs/what/modex/

[2]Our implementation is based on CPACHECKER, which operates on C programs; non-recursive function calls are supported.

integer value. A *concrete state* of a program is a pair (l, cd), where $l \in L$ is a program location and cd is a concrete data state. The set of all concrete states of a program is denoted by $\mathcal{C}$, a subset $r \subseteq \mathcal{C}$ is called *region*. Each edge $g \in G$ defines a labeled transition relation $\xrightarrow{g} \subseteq \mathcal{C} \times \{g\} \times \mathcal{C}$. The complete transition relation $\rightarrow$ is the union over all control-flow edges: $\rightarrow = \bigcup_{g \in G} \xrightarrow{g}$. We write $c \xrightarrow{g} c'$ if $(c, g, c') \in \rightarrow$, and $c \rightarrow c'$ if there exists a g with $c \xrightarrow{g} c'$.

An *abstract data state* represents a region of concrete data states, formally defined as abstract variable assignment. An *abstract variable assignment* is a partial function $v : X \nrightarrow \mathbb{Z} \cup \{\top, \bot\}$, which maps variables in the definition range of function v to integer values or $\top$ or $\bot$. The special value $\top$ is used to represent an unknown value, e.g., resulting from an uninitialized variable or an external function call, and the special value $\bot$ is used to represent no value, i.e., a contradicting variable assignment. We denote the *definition range* for a partial function f as $\text{def}(f) = \{x \mid \exists y : (x, y) \in f\}$, and the *restriction* of a partial function f to a new definition range Y as $f_{|Y} = f \cap (Y \times (\mathbb{Z} \cup \{\top, \bot\}))$. An abstract variable assignment v represents the region $[\![v]\!]$ of all concrete data states cd for which v is valid, formally: $[\![v]\!] = \{cd \mid \forall x \in \text{def}(v) : cd(x) = v(x) \text{ or } v(x) = \top\}$. An *abstract state* of a program is a pair (l, v), representing the following set of concrete states: $\{(l, cd) \mid cd \in [\![v]\!]\}$.

B. Configurable Program Analysis with Dynamic Precision Adjustment

We use the framework of configurable program analysis (CPA) [8], extended by the concept of dynamic precision adjustment [9]. Such a CPA supports adjusting the precision of an analysis during the exploration of the program's abstract state space. A *composite* CPA can control the precision of its component analyses during the verification process, i.e., it can make a component analysis more abstract, and thus more efficient, or it can make a component analysis more precise, and thus more expensive. A CPA $\mathbb{D} = (D, \Pi, \rightsquigarrow, \mathsf{merge}, \mathsf{stop}, \mathsf{prec})$ consists of (1) an abstract domain D, (2) a set Π of precisions, (3) a transfer relation $\rightsquigarrow$, (4) a merge operator merge, (5) a termination check stop, and (6) a precision adjustment function prec. Based on these components and operators, we can formulate a flexible and customizable reachability algorithm, which is adapted from previous work [8], [12].

C. Explicit-Value Analysis as CPA

In the following, we define a component CPA that tracks explicit integer values for program variables. In order to obtain a complete analysis, we construct a composite CPA that consists of the component CPA for explicit values and another component CPA for tracking the program locations (CPA for location analysis, as previously described [9]). For the composite CPA, the general definitions of the abstract domain, the transfer relation, and the other operators are given in previous work [9]; the composition is done automatically by the framework implementation CPACHECKER.

The *CPA for explicit-value analysis*, which tracks integer values for the variables of a program explicitly, is defined as $\mathbb{C} = (D_{\mathbb{C}}, \Pi_{\mathbb{C}}, \rightsquigarrow_{\mathbb{C}}, \mathsf{merge}_{\mathbb{C}}, \mathsf{stop}_{\mathbb{C}}, \mathsf{prec}_{\mathbb{C}})$ and consists of the following components [9]:

1. The abstract domain $D_{\mathbb{C}} = (C, \mathcal{V}, [\![\cdot]\!])$ contains the set C of concrete data states, and uses the semi-lattice $\mathcal{V} = (V, \top, \bot, \sqsubseteq, \sqcup)$, which consists of the set $V = (X \nrightarrow \mathcal{Z})$ of abstract variable assignments, where $\mathcal{Z} = \mathbb{Z} \cup \{\top_{\mathcal{Z}}, \bot_{\mathcal{Z}}\}$ induces the flat lattice over the integer values (we write $\mathbb{Z}$ to denote the set of integer values). The top element $\top \in V$, with $\top(x) = \top_{\mathcal{Z}}$ for all $x \in X$, is the abstract variable assignment that holds no specific value for any variable, and the bottom element $\bot \in V$, with $\bot(x) = \bot_{\mathcal{Z}}$ for all $x \in X$, is the abstract variable assignment which models that there is no value assignment possible, i.e., a state that cannot be reached in an execution of the program. The partial order $\sqsubseteq \subseteq V \times V$ is defined as $v \sqsubseteq v'$ if for all $x \in X$, we have $v(x) = v'(x)$ or $v(x) = \bot_{\mathcal{Z}}$ or $v'(x) = \top_{\mathcal{Z}}$. The join $\sqcup : V \times V \to V$ yields the least upper bound for two variable assignments. The concretization function $[\![\cdot]\!] : V \to 2^C$ assigns to each abstract data state v its meaning, i.e., the set of concrete data states that it represents.

2. The set of precisions $\Pi_{\mathbb{C}} = 2^X$ is the set of subsets of program variables. A precision $\pi \in \Pi_{\mathbb{C}}$ specifies a set of variables to be tracked. For example, $\pi = \emptyset$ means that no variable is tracked, and $\pi = X$ means that every program variable is tracked.

3. The transfer relation $\rightsquigarrow_{\mathbb{C}}$ has the transfer $v \overset{g}{\rightsquigarrow} (v', \pi)$ if
(1) $g = (\cdot, \mathtt{assume}(p), \cdot)$ and for all $x \in X$:

$$v'(x) = \begin{cases} \bot_{\mathcal{Z}} & \text{if } (y, \bot_{\mathcal{Z}}) \in v \text{ for some } y \in X \\ & \quad \text{or the formula } p_{/v} \text{ is unsatisfiable} \\ c & \text{if } c \text{ is the only satisfying assignment of} \\ & \quad \text{the formula } p_{/v} \text{ for variable } x \\ \top_{\mathcal{Z}} & \text{otherwise} \end{cases}$$

where $p_{/v}$ denotes the interpretation of a predicate p over variables from X for an abstract variable assignment v, that is, $p_{/v} =$

$$p \ \wedge \bigwedge_{x \in \text{def}(v), v(x) \in \mathbb{Z}} x = v(x) \ \wedge \ \neg \exists x \in \text{def}(v) : v(x) = \bot_{\mathcal{Z}}$$

or

(2) $g = (\cdot, \mathtt{w} := exp, \cdot)$ and for all $x \in X$:

$$v'(x) = \begin{cases} exp_{/v} & \text{if } x = \mathtt{w} \\ v(x) & \text{if } x \in \text{def}(v) \\ \top_{\mathcal{Z}} & \text{otherwise} \end{cases}$$

where $exp_{/v}$ denotes the interpretation of an expression exp over variables from X for an abstract value assignment v:

$$exp_{/v} = \begin{cases} \bot_{\mathcal{Z}} & \text{if } (y, \bot_{\mathcal{Z}}) \in v \text{ for some } y \in X \\ \top_{\mathcal{Z}} & \text{if } (y, \top_{\mathcal{Z}}) \in v \text{ or } y \notin \text{def}(v) \\ & \quad \text{for some } y \in X \text{ that occurs in } exp \\ c & \text{otherwise, where expression } exp \\ & \quad \text{evaluates to } c \text{ after replacing each} \\ & \quad \text{occurrence of variable } x \text{ with } x \in \text{def}(v) \\ & \quad \text{by } v(x) \text{ in } exp \end{cases}$$

4. The merge operator does not combine elements when control flow meets: $\mathsf{merge}_{\mathbb{C}}(v, v', \pi) = v'$.

5. The termination check considers abstract states individually: $\mathsf{stop}_{\mathbb{C}}(v, R, \pi) = (\exists v' \in R : v \sqsubseteq v')$.

6. The precision adjustment function computes a new abstract state with precision based on the abstract state v and the precision π by restricting the variable assignment v to those variables that appear in π, formally: $\mathsf{prec}(v, \pi, R) = (v_{|\pi}, \pi)$. (In this analysis instance, prec only adjusts the abstract state according to the current precision π, and leaves the precision itself unchanged.)

The precision of the analysis controls which program variables are tracked in an abstract state. In other approaches, this information is hard-wired in either the abstract-domain elements or the algorithm itself. The concept of CPA supports different precisions for different abstract states. A simple analysis can start with an initial precision and propagate it to new abstract states, such that the overall analysis uses a globally uniform precision. It is also possible to specify a precision individually per program location, instead of using one global precision. Our refinement approach in the next section will be based on location-specific precisions.

D. Predicate Analysis as CPA

The abstract domain of predicates [18] was successfully used in several tools for software model checking (e.g., [4], [6], [10], [13], [16], [25]). In a predicate analysis, the precision is defined as a set of predicates, and the abstract states track the strongest set of predicates that are fulfilled (cartesian predicate abstraction) or the strongest boolean combination of predicates that are fulfilled (boolean predicate abstraction). This means, the abstraction level of the abstract model is determined by predicates that are tracked in the analysis. Predicate analysis is also implemented as a CPA in the framework CPACHECKER, and a detailed description is available [11]. The precision is freely adjustable also in the predicate analysis, and we use this feature later in this article to compose a combined analysis. This analysis uses the predicate analysis to track variables that have many distinct values — a scenario in which the explicit-value analysis alone would be inefficient. The combined analysis adjusts the overall precision by removing variables with many distinct values from the precision of the explicit-value analysis and adds predicates about these variables to the precision of the predicate analysis [9] to allow the combined analysis to run efficiently.

E. Lazy Abstraction

The concept of lazy abstraction [22] consists of two ideas: First, the abstract reachability graph (ARG) —the unfolding of the control-flow graph, representing our central data structure to store abstract states— is constructed on-the-fly, i.e., only when needed and only for parts of the state space that are reachable. We implement this using the standard reachability algorithm for CPAs as described in the next subsection. Second, the abstract states in the ARG are refined only where necessary along infeasible error paths in order to eliminate those paths. This is implemented by using CPAs with dynamic precision adjustment, where the refinement procedure operates on location-specific precisions and where the precision-adjustment operator always removes unnecessary information from abstract states, as outlined above.

Algorithm 1 $\mathsf{CPA}(\mathbb{D}, R_0, W_0)$, adapted from [9]

Input: a CPA $\mathbb{D} = (D, \Pi, \rightsquigarrow, \mathsf{merge}, \mathsf{stop}, \mathsf{prec})$,
a set $R_0 \subseteq (E \times \Pi)$ of abstract states with precision,
a subset $W_0 \subseteq R_0$ of frontier abstract states with precision,
where E denotes the set of elements of the semi-lattice of D

Output: a set of reachable abstract states with precision,
a subset of frontier abstract states with precision

Variables: two sets reached and waitlist of elements of $E \times \Pi$

reached $:= R_0$; waitlist $:= W_0$;
while waitlist $\neq \emptyset$ **do**
  choose (e, π) from waitlist; remove (e, π) from waitlist;
  for each e' with $e \rightsquigarrow (e', \pi)$ **do**
    // Precision adjustment.
    $(\hat{e}, \hat{\pi}) := \mathsf{prec}(e', \pi, \mathsf{reached})$;
    if $\mathsf{isTargetState}(\hat{e})$ **then**
      return $(\mathsf{reached} \cup (\hat{e}, \hat{\pi}), \mathsf{waitlist})$;
    for each $(e'', \pi'') \in \mathsf{reached}$ **do**
      // Combine with existing abstract state.
      $e_{new} := \mathsf{merge}(\hat{e}, e'', \hat{\pi})$;
      if $e_{new} \neq e''$ **then**
        $\mathsf{waitlist} := (\mathsf{waitlist} \cup \{(e_{new}, \hat{\pi})\}) \setminus \{(e'', \pi'')\}$;
        $\mathsf{reached} := (\mathsf{reached} \cup \{(e_{new}, \hat{\pi})\}) \setminus \{(e'', \pi'')\}$;
    // Add new abstract state?
    if $\neg\ \mathsf{stop}(\hat{e}, \{e \mid (e, \cdot) \in \mathsf{reached}\}, \hat{\pi})$ **then**
      $\mathsf{waitlist} := \mathsf{waitlist} \cup \{(\hat{e}, \hat{\pi})\}$;
      $\mathsf{reached} := \mathsf{reached} \cup \{(\hat{e}, \hat{\pi})\}$
return $(\mathsf{reached}, \emptyset)$;

F. Reachability Algorithm for CPA

Algorithm 1 keeps updating two sets of abstract states with precision: the set reached to store all abstract states with precision that are found to be reachable, and a set waitlist to store all abstract states with precision that are not yet processed, i.e., the frontier. The state exploration starts with choosing and removing an abstract state with precision from the waitlist, and the algorithm considers each abstract successor according to the transfer relation. Next, for the successor, the algorithm adjusts the precision of the successor using the precision adjustment function prec. If the successor is a target state (i.e., a violation of the property is found), then the algorithm terminates, returning the current sets reached and waitlist — possibly as input for a subsequent precision refinement, as shown below (cf. Alg. 2). Otherwise, using the given operator merge, the abstract successor state is combined with each existing abstract state from reached. If the operator merge has added information to the new abstract state, such that the old abstract state is subsumed, then the old abstract state with precision is replaced by the new abstract state with precision in the sets reached and waitlist. If after the merge step the resulting new abstract state with precision is covered by the set reached, then further exploration of this abstract state is stopped. Otherwise, the abstract state with its precision is added to the set reached and to the set waitlist. Finally, once the set waitlist is empty, the set reached is returned.

G. Counterexample-Guided Abstraction Refinement

Counterexample-guided abstraction refinement (CEGAR) [14] is a technique for automatic stepwise refinement of an abstract model. CEGAR is based on three concepts: (1) a *precision*, which determines the current level of abstraction, (2) a *feasibility check*, deciding if an abstract error path is feasible, i.e., if there exists a corresponding concrete error path, and (3) a *refinement* procedure, which takes as input an infeasible error path and extracts a precision that suffices to instruct the exploration algorithm to not explore the same path again later. Algorithm 2 shows an outline of a generic and simple CEGAR algorithm. The algorithm starts checking a program using a coarse initial *precision* π_0. It uses the reachability algorithm Alg. 1 for computing the reachable abstract state space, returning the sets reached and waitlist. If the analysis has exhaustively checked all program states and did not reach the error, indicated by an empty set waitlist, then the algorithm terminates and reports that the program is safe. If the algorithm finds an error in the abstract state space, i.e., a counterexample for the given specification, then the exploration algorithm stops and returns the unfinished, incomplete sets reached and waitlist. Now the according abstract error path is extracted from the set reached using procedure extractErrorPath and analyzed for feasibility using the procedure isFeasible for *feasibility check*. If the abstract error path is feasible, meaning there exists a corresponding concrete error path, then this error path represents a violation of the specification and the algorithm terminates, reporting a bug. If the error path is infeasible, i.e., not corresponding to a concrete program path, then the precision was too coarse and needs to be refined. The algorithm extracts certain information from the error path in order to refine the precision based on that information using the procedure Refine for *refinement*, which returns a precision π that makes the analysis strong enough to refute the infeasible error path in further state-space explorations. The current precision is extended using the precision returned by the refinement procedure and the analysis is restarted with this refined precision. Instead of restarting from the initial sets for reached and waitlist, we can also prune those parts of the ARG that need to be rediscovered with new precisions, and replace the precision of the leaf nodes in the ARG with the refined precision, and then restart the exploration on the pruned sets. Our contribution in the next section is to introduce new implementations for the feasibility check as well as for the refinement procedure.

H. Interpolation

For a pair of formulas φ^- and φ^+ such that $\varphi^- \wedge \varphi^+$ is unsatisfiable, a Craig interpolant ψ is a formula that fulfills the following requirements [17]:

1) the implication $\varphi^- \Rightarrow \psi$ holds,
2) the conjunction $\psi \wedge \varphi^+$ is unsatisfiable, and
3) ψ only contains symbols that occur in both φ^- and φ^+.

Such a Craig interpolant is guaranteed to exist for many useful theories, for example, the theory of linear arithmetic with

Algorithm 2 $\mathsf{CEGAR}(\mathbb{D}, e_0, \pi_0)$

Input: a configurable program analysis with dynamic precision adjustment $\mathbb{D} = (D, \Pi, \rightsquigarrow, \mathsf{merge}, \mathsf{stop}, \mathsf{prec})$,
an initial abstract state $e_0 \in E$ with precision $\pi_0 \in \Pi$,
where E denotes the set of elements of the semi-lattice of D
Output: verification result *safe* or *unsafe*
Variables: a set reached of elements of $E \times \Pi$,
a set waitlist of elements of $E \times \Pi$,
an error path $\sigma = \langle (op_1, l_1), ..., (op_n, l_n) \rangle$

```
reached := {(e0, π0)}; waitlist := {(e0, π0)}; π := π0;
while true do
  (reached, waitlist) := CPA(D, reached, waitlist);
  if waitlist = ∅ then
    return safe
  else
    σ := extractErrorPath(reached);
    if isFeasible(σ) then  // error path is feasible: report bug
      return unsafe
    else  // error path is not feasible: refine and restart
      π := π ∪ Refine(σ);
      reached := (e0, π); waitlist := (e0, π);
```

uninterpreted functions, as implemented in some SMT solvers (e.g., MATHSAT[3], SMTINTERPOL[4]).

CEGAR based on Craig interpolation has been proven successful in the predicate domain. Therefore, we investigate if this technique is also beneficial for explicit-value model checking. Interpolants from the predicate domain, which consist of path formulas, are not useful for the explicit domain. Hence, we need to develop a procedure to compute interpolants for the explicit domain, which we introduce in the following section.

III. Refinement-Based Explicit-Value Analysis

The level of abstraction in our explicit-value analysis is determined by the precisions for abstract variable assignments over program variables. The CEGAR-based iterative refinement needs an extraction method to obtain the necessary precision from infeasible error paths. We use our novel notion of interpolation for the explicit domain to achieve this goal.

A. Explicit-Value Abstraction

We now introduce some necessary operations on abstract variable assignments, the semantics of operations and paths, and the precision for abstract variable assignments and programs, in order to be able to concisely discuss interpolation for abstract variable assignments and constraint sequences.

The operations *implication* and *conjunction* for abstract variable assignments are defined as follows: implication for v and v': $v \Rightarrow v'$ if $\text{def}(v') \subseteq \text{def}(v)$ and for each variable $x \in \text{def}(v) \cap \text{def}(v')$ we have $v(x) = v'(x)$ or $v(x) = \bot$ or $v'(x) = \top$; conjunction for v and v': for each variable $x \in \text{def}(v) \cup \text{def}(v')$ we have

$$(v \wedge v')(x) = \begin{cases} v(x) & \text{if } x \in \text{def}(v) \text{ and } x \notin \text{def}(v') \\ v'(x) & \text{if } x \notin \text{def}(v) \text{ and } x \in \text{def}(v') \\ v(x) & \text{if } v(x) = v'(x) \\ \bot & \text{if } \top \neq v(x) \neq v'(x) \neq \top \\ \top & \text{otherwise } (v(x) = \top \text{ or } v'(x) = \top) \end{cases}$$

[3] http://mathsat4.disi.unitn.it
[4] http://ultimate.informatik.uni-freiburg.de/smtinterpol

Furthermore we define *contradiction* for an abstract variable assignment v: v is contradicting if there is a variable $x \in \text{def}(v)$ such that $v(x) = \bot$ (which implies $[\![v]\!] = \emptyset$); and *renaming* for v: the abstract variable assignment $v^{x \mapsto y}$, with $y \notin \text{def}(v)$, results from v by renaming variable x to y: $v^{x \mapsto y} = (v \setminus \{(x, v(x))\}) \cup \{(y, v(x))\}$.

The *semantics of an operation* $op \in Ops$ is defined by the strongest post-operator $\mathsf{SP}_{op}(\cdot)$ for abstract variable assignments: given an abstract variable assignment v, $\mathsf{SP}_{op}(v)$ represents the set of data states that are reachable from any of the states in the region represented by v after the execution of op. Formally, given an abstract variable assignment v and an assignment operation $s := exp$, we have $\mathsf{SP}_{s:=exp}(v) = v_{|X \setminus \{s\}} \wedge v_{s:=exp}$ with $v_{s:=exp} = \{(s, exp_{/v})\}$, where $exp_{/v}$ denotes the interpretation of expression exp for the abstract variable assignment v (cf. definition of $exp_{/v}$ in Subsection II-C). That is, the value of variable s is the result of the arithmetic evaluation of expression exp, or $\top$ if not all values in the expression are known, or $\bot$ if no value is possible (an abstract data state in which a variable is assigned to $\bot$ does not represent any concrete data state). Given an abstract variable assignment v and an assume operation $[p]$, we have $\mathsf{SP}_{[p]}(v) = v'$ and for all $x \in X$ we have $v'(x) = \bot$ if $(y, \bot) \in v$ for some variable $x \in X$ or the formula $p_{/v}$ is unsatisfiable, or $v'(x) = c$ if c is the only satisfying assignment of the formula $p_{/v}$ for variable x, or $v'(x) = \top$ in all other cases; the formula $p_{/v}$ is defined as in Subsection II-C.

A *path* σ is a sequence $\langle (op_1, l_1), ..., (op_n, l_n) \rangle$ of pairs of an operation and a location. The path σ is called *program path* if for every i with $1 \leq i \leq n$ there exists a CFA edge $g = (l_{i-1}, op_i, l_i)$ and l_0 is the initial program location, i.e., σ represents a syntactic walk through the CFA. Every path $\sigma = \langle (op_1, l_1), ..., (op_n, l_n) \rangle$ defines a *constraint sequence* $\gamma_\sigma = \langle op_1, ..., op_n \rangle$. The *semantics of a program path* $\sigma = \langle (op_1, l_1), ..., (op_n, l_n) \rangle$ is defined as the successive application of the strongest post-operator to each operation of the corresponding constraint sequence γ_σ: $\mathsf{SP}_{\gamma_\sigma}(v) = \mathsf{SP}_{op_n}(...\mathsf{SP}_{op_i}(..\mathsf{SP}_{op_1}(v)..)...)$. The set of concrete program states that result from running σ is represented by the pair $(l_n, \mathsf{SP}_{\gamma_\sigma}(v_0))$, where $v_0 = \{\}$ is the initial abstract variable assignment that does not map any variable to a value. A program path σ is *feasible* if $\mathsf{SP}_{\gamma_\sigma}(v_0)$ is not contradicting, i.e., $\mathsf{SP}_{\gamma_\sigma}(v_0)(x) \neq \bot$ for all variables x in $\text{def}(\mathsf{SP}_{\gamma_\sigma}(v_0))$. A concrete state (l_n, cd_n) is *reachable* from a region r, denoted by $(l_n, cd_n) \in Reach(r)$, if there exists a feasible program path $\sigma = \langle (op_1, l_1), ..., (op_n, l_n) \rangle$ with $(l_0, v_0) \in r$ and $cd_n \in [\![\mathsf{SP}_{\gamma_\sigma}(v_0)]\!]$. A location l is reachable if there exists a concrete state c such that (l, c) is reachable. A program is SAFE if l_e is not reachable.

The *precision for an abstract variable assignment* is a set π of variables. The *explicit-value abstraction* for an abstract variable assignment is an abstract variable assignment that is defined only on variables that are in the precision π. For example, the explicit-value abstraction for the variable assignment $v = \{x \mapsto 2, y \mapsto 5\}$ and the precision $\pi = \{x\}$ is the abstract variable assignment $v^\pi = \{x \mapsto 2\}$.

The *precision for a program* is a function $\Pi : L \to 2^X$, which assigns to each program location a precision for an abstract variable assignment, i.e., a set of variables for which the analysis is instructed to track values. A *lazy explicit-value abstraction* of a program uses different precisions for different abstract states on different program paths in the abstract reachability graph (ARG). The explicit-value abstraction for a variable assignment at location l is computed using the precision $\Pi(l)$.

B. CEGAR for Explicit-Value Model Checking

We now instantiate the three components of the CEGAR technique, i.e., precision, feasibility check, and refinement, for our explicit-value analysis. The precisions that our CEGAR instance uses are the above introduced precisions for a program (which assign to each program location a set of variables), and we start the CEGAR iteration with the empty precision, i.e., $\Pi_{init}(l) = \emptyset$ for each $l \in L$, such that no variable will be tracked.

The feasibility check for a path σ is performed by executing an explicit-value analysis of the path σ using the full precision $\Pi(l) = X$ for all locations l, i.e., all variables will be tracked. This is equivalent to computing $\mathsf{SP}_{\gamma_\sigma}(v_0)$ and check if the result is contradicting, i.e., if there is a variable for which the resulting abstract variable assignment is $\bot$. This feasibility check is extremely efficient, because the path is finite and the strongest post-operations for abstract variable assignments are simple arithmetic evaluations. If the feasibility check reaches the error location l_e, then this error can be reported. If the check cannot reach the error location, because of a contradicting abstract variable assignment, then a refinement is necessary because at least one constraint depends on a variable that was not yet tracked.

We define the last component of the CEGAR technique, the refinement, after we introduced the notion of interpolation for variable assignments and constraint sequences.

C. Interpolation for Variable Assignments

For each infeasible error path in the above mentioned refinement operation, we need to determine a precision that assigns to each program location on that path the set of program variables that the explicit-value analysis needs to track in order to eliminate that infeasible error path in future explorations. Therefore, we define an interpolant for abstract variable assignments.

An *interpolant* for a pair of abstract variable assignments v^- and v^+, such that $v^- \wedge v^+$ is contradicting, is an abstract variable assignment $\mathcal{V}$ that fulfills the following requirements:

1) the implication $v^- \Rightarrow \mathcal{V}$ holds,
2) the conjunction $\mathcal{V} \wedge v^+$ is contradicting, and
3) $\mathcal{V}$ only contains variables in its definition range which are in the definition ranges of both v^- and v^+ ($\text{def}(\mathcal{V}) \subseteq \text{def}(v^-) \cap \text{def}(v^+)$).

Lemma. For a given pair (v^-, v^+) of abstract variable assignments, such that $v^- \wedge v^+$ is contradicting, an interpolant exists. Such an interpolant can be computed in time $O(m+n)$, where m and n are the sizes of v^- and v^+, respectively.

Algorithm 3 Interpolate(γ^-, γ^+)

Input: two constraint sequences γ^- and γ^+,
with $\gamma^- \wedge \gamma^+$ is contradicting
Output: a constraint sequence Γ,
which is an interpolant for γ^- and γ^+
Variables: an abstract variable assignment v

```
v := SP_{γ⁻}(∅)
for each x ∈ def(v) do
  if SP_{γ⁺}(v|def(v)\{x}) is contradicting then
    // x is not relevant and should not occur in the interpolant
    v := v|def(v)\{x}
// construct the interpolating constraint sequence
Γ := ⟨⟩
for each x ∈ def(v) do
  // construct an assume constraint for x
  Γ := Γ ∧ ⟨[x = v(x)]⟩
return Γ
```

Proof. The variable assignment $v^-_{|\mathrm{def}(v^+)}$ is an interpolant for the pair (v^-, v^+).

Note. The above-mentioned interpolant that simply results from restricting v^- to the definition range of v^+ (common definition range) is of course not a 'good' interpolant. In practice, we strive for interpolants with minimal definition range, and use slightly more expensive algorithms to compute them. Interpolation for abstract variable assignments is a first idea to approach the problem, but since we need to extract interpolants for paths, we next define interpolation for constraint sequences.

D. Interpolation for Constraint Sequences

A more expressive interpolation can be achieved by considering constraint sequences. The *conjunction* $\gamma \wedge \gamma'$ of two constraint sequences $\gamma = \langle op_1, ..., op_n \rangle$ and $\gamma' = \langle op'_1, ..., op'_m \rangle$ is defined as their concatenation, i.e., $\gamma \wedge \gamma' = \langle op_1, ..., op_n, op'_1, ..., op'_m \rangle$, the *implication* of γ and γ' (denoted by $\gamma \Rightarrow \gamma'$) as $\mathsf{SP}_\gamma(v_0) \Rightarrow \mathsf{SP}_{\gamma'}(v_0)$, and γ is *contradicting* if $[\![\mathsf{SP}_\gamma(v_0)]\!] = \emptyset$, with $v_0 = \{\}$.

An *interpolant* for a pair of constraint sequences γ^- and γ^+, such that $\gamma^- \wedge \gamma^+$ is contradicting, is a constraint sequence Γ that fulfills the following requirements:

1) the implication $\gamma^- \Rightarrow \Gamma$ holds,
2) the conjunction $\Gamma \wedge \gamma^+$ is contradicting, and
3) Γ contains in its constraints only variables that occur in the constraints of both γ^- and γ^+.

Lemma. For a given pair (γ^-, γ^+) of constraint sequences, such that $\gamma^- \wedge \gamma^+$ is contradicting, an interpolant exists. Such an interpolant is computable in time $O(m \cdot n)$, where m and n are the sizes of γ^- and γ^+, respectively.

Proof. Algorithm Interpolate (Alg. 3) returns an interpolant for two constraint sequences γ^- and γ^+. The algorithm starts with computing the strongest post-condition for γ^- and assigns the result to the abstract variable assignment v, which then may contain up to m variables. Per definition, the strongest post-condition for γ^+ of variable assignment v is contradicting. Next we try to eliminate each variable from v, by testing if removing it from v makes the strongest post-condition for γ^+ of v contradicting (each such test takes n SP steps). If it is contradicting, the variable can be removed. If not, the variable is necessary to prove the contradiction of the two constraint sequences, and thus, should occur in the interpolant. Note that this keeps only variables in v that occur in γ^+ as well. The rest of the algorithm constructs a constraint sequence from the variable assignment, in order to return an interpolating constraint sequence, which fulfills the three requirements of an interpolant. A naive implementation can compute such an interpolant in $O((m+n)^3)$.

Algorithm 4 Refine(σ)

Input: infeasible error path $\sigma = \langle (op_1, l_1), ..., (op_n, l_n) \rangle$
Output: precision Π
Variables: interpolating constraint sequence Γ

```
Γ := ⟨⟩;
Π(l) := ∅, for all program locations l;
for i := 1 to n − 1 do
  γ⁺ := ⟨op_{i+1}, ..., op_n⟩
  // inductive interpolation
  Γ := Interpolate(Γ ∧ op_i, γ⁺)
  // extract variables from variable assignment that results from Γ
  Π(l_i) := {x | (x, z) ∈ SP_Γ(∅) and ⊥ ≠ z ≠ ⊤}
return Π
```

E. Refinement Based on Explicit-Interpolation

The goal of our interpolation-based refinement for explicit-value analysis is to determine a localized precision that is strong enough to eliminate an infeasible error path in future explorations. This criterion is fulfilled by the property of interpolants. A second goal is to have a precision that is as weak as possible, by creating interpolants that have a definition range as small as possible, in order to be parsimonious in tracking variables and creating abstract states.

We apply the idea of interpolation for constraint sequences to assemble a precision-extraction algorithm: Algorithm Refine (Alg. 4) takes as input an infeasible program path, and returns a precision for a program. A further requirement is that the procedure computes *inductive* interpolants [6], i.e., each interpolant along the path contains just enough information to prove the remaining path infeasible. This is needed in order to ensure that the interpolants at the different locations achieve the goal of providing a precision that eliminates the infeasible error path from further explorations. For every program location l_i along an infeasible error path σ, starting at l_0, we split the constraint sequence of the path into a constraint prefix γ^-, which consists of the constraints from the start location l_0 to l_i, and a constraint suffix γ^+, which consists of the path from the location l_i to l_e. For computing inductive interpolants, we replace the constraint prefix by the conjunction of the last interpolant and the current constraint. The precision is extracted by computing the abstract variable assignment for the interpolating constraint sequence and assigning the relevant variables as precision for the current location l_i, i.e., the set of all variables that are necessary to be tracked in order to eliminate the error path from future exploration of the state space. This algorithm for precision extraction yields a parsimonious precision, i.e., a precision containing *just* enough information to exclude the infeasible error path, and can be directly plugged-in as refinement routine

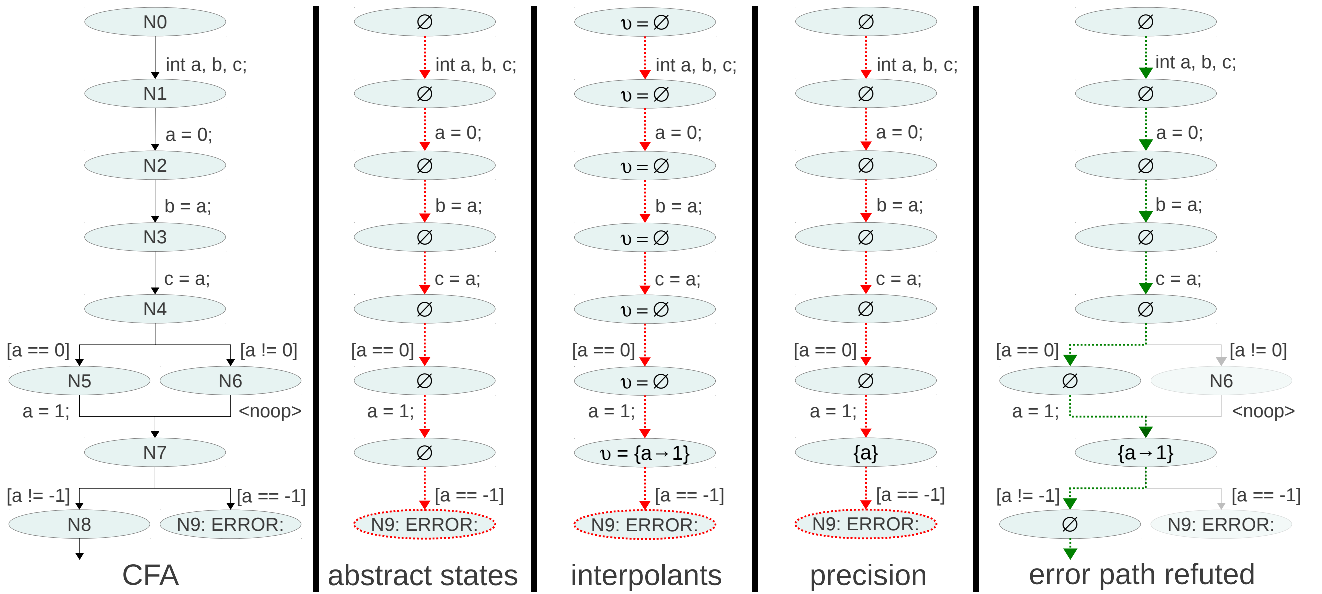

Fig. 2: Illustration of one refinement iteration; from left to right: a simple example CFA, an infeasible error path with the abstract states annotated in the nodes (precision was empty, nothing is tracked), the interpolated variable assignments annotated in the nodes, the precisions extracted from the interpolants annotated in the nodes, and finally the CFA with the abstract states annotated in the nodes according to the new precision (unreached nodes —including error— shown in gray)

of the CEGAR algorithm (cf. Alg. 2). Note that the repetitive interpolations are not an efficiency bottleneck. The path is always finite, without any loops or branching, and thus, even a full-precision check can be decided efficiently. Figure 2 illustrates the interpolation process on a simple example.

F. Optimizations

In our implementation, we added several optimizations to improve the performance of our approach.

ARG Pruning instead of Restart. Our refinement routine Refine (cf. Alg. 4) returns a set of variables (precision) that are important for deciding the reachability of the error location. One of the ideas of lazy abstraction refinement [22] is that the precision is only refined where necessary, i.e., only at the locations along the path that was considered in the refinement; the other parts of the state space are not refined. As mentioned in the discussion of the CEGAR algorithm (cf. Alg. 2), it is not necessary to restart the exploration of the state space from scratch after a refinement. Instead, we identify the descendant closest to the root of the abstract reachability graph (ARG) in which the precision was refined, and the re-exploration of the state space continues from there. In total, this significantly reduces the number of tracked variables per abstract state, which in turn leads to a more efficient analysis, because it drastically increases the chance that a new abstract state is covered by an existing abstract state.

Scoped Precision Refinement. The precision for a program assigns to each program location the set of variables that need to be tracked at that location, and the interpolation-based refinement adds new variables precisely at the locations for which they were discovered during refinement. In our experience, the number of refinements is reduced significantly if we add a variable to the precision not only at the particular location for which it was discovered, but at all locations in the local scope of the variable. This helps to avoid adding a variable twice that can occur on two different branches. By adding the variable to the precision "in advance" in the local scope, we abbreviate some refinement iterations. For example, consider Fig. 2 again. After the illustrated refinement, another refinement step would be necessary, in order to discover that variable a needs to be tracked at location N4 as well (to prevent the analysis from going through location N6). By adding variable a to the precision of all locations in the scope of variable a immediately after the first refinement, the program can be proved safe without further refinement. This effect was also observed, and used, in the software model checker BLAST [6].

Precise Counterexample Check. In order to further increase the precision of our analysis, we double-check all feasible error paths using bit-precise bounded model checking (BMC) [5], by generating a path program [7] for the error path and let the BMC confirm the bug. Since the generated path program does not contain any loop or branching, it can be verified efficiently. If both our analysis and the bit-precise BMC report *unsafe*, then we report a bug. If the BMC cannot confirm the bug, our analysis continues trying to find another error path. This additional feature is available as a command-line option in our implementation.

Auxiliary Predicate Analysis. As an additional option for further improvement of the analysis, we implemented the combination with a predicate analysis, as outlined in existing work [9]. In this combination, if the explicit-value analysis finds an error path, this path is first checked for satisfiability in the predicate domain. If the satisfiability check is positive, the result *unsafe* can be reported and the error path is returned; if negative, then the explicit-value domain is not expressive enough to analyze that program path (e.g., due to inequalities).

[5] In our implementation, we use CBMC [15] as bounded model checker.

In this case, we ask the predicate analysis to refine its abstraction along that path, which yields a refined predicate precision that eliminates the error path but considering the facts along that path in the (more precise, and more expensive) predicate domain. We need to parsimoniously use this feature because the post-operations of the predicate analysis are much more expensive than the post-operations of the explicit-value analysis. In general, after a refinement step, either the explicit-value precision is refined (preferred) or the predicate precision is refined (only if explicit does not succeed).

Using the concept of dynamic precision adjustment [9], we also switch off the tracking of variables in the explicit-value domain if the number of different values on a path exceeds a certain threshold. After this, the predicate analysis will get switched on (by the above-mentioned mechanism) and the facts on that path are further tracked using predicates. This is important if the explicit-value analysis tries to unwind loops; the symbolic, predicate-based analysis can often store a large number of values more efficiently.

Note that this refinement-based, parallel composition with precision adjustment of the explicit-value analysis and the predicate analysis is more powerful than a mere parallel product of the two analyses, because after each refinement, the explicit part of the analysis tracks exactly what it is capable of tracking, while the auxiliary predicate analysis takes care of only those facts that are beyond the capabilities of the explicit domain, resulting in a lightweight analysis on both ends. Such a combination is easy to achieve in our implementation, because we use the framework of configurable program analysis (CPA), which lets the user freely configure such combinations.

IV. Experiments

In order to demonstrate that our approach yields a significant practical improvement of verification efficiency and effectiveness, we implemented our algorithms and compared our new techniques to existing tools for software verification. In the following, we show that the application of abstraction, CEGAR, and interpolation to the explicit-value domain considerably improves the number of solved instances and the run time. Combinations of the new explicit-value analysis with a predicate-based analysis can further increase the number of solved instances. All our experiments were performed on hardware identical to that of the SV-COMP'12 [5], such that our results are comparable to all the results obtained there.

Compared Verification Approaches. For presentation, we restrict the comparison of our new approach to the SV-COMP'12 participants BLAST, SATABS, and the competition winner CPA-MEMO, all of which are based on predicate abstraction and CEGAR. Furthermore, to investigate performance differences in the same tool environment, we also compare with different configurations of CPACHECKER. The model checker BLAST is based on predicate abstraction, and uses a CEGAR loop for abstraction refinement. The predicates for the precision are learned from counterexample paths using interpolation. The central data structure of the algorithm is an ARG, which is lazily constructed and refined. BLAST won the category "DeviceDrivers64" in the SV-COMP'12, and got bronze in another category. The model checker SATABS is also based on predicate abstraction and CEGAR, but in contrast to BLAST, it constructs and checks in every iteration of the CEGAR loop a new boolean program based on the current precision of the predicate abstraction, and does not use lazy abstraction or interpolation. SATABS got silver in the categories "SystemC" and "Concurrency", and bronze in another category. The model checker CPA-MEMO is based on predicate abstraction, CEGAR, and interpolation, but extends it with the concepts of adjustable-block encoding [11] and block-abstraction memoization [26]. CPA-MEMO won the category "Overall", got silver in two more categories, and bronze in another category.

We implemented our concepts as extensions of CPACHECKER [10], a software-verification framework based on configurable program analysis (CPA). We compare with the existing explicit-value analysis (without abstraction, CEGAR, and interpolation) and with the existing predicate analysis that is based on boolean predicate abstraction, CEGAR, interpolation, and adjustable-block encoding [11]. We used the trunk version of CPACHECKER[6] in revision 6615.

Verification Tasks. For the evaluation of our approach, we use all SV-COMP'12[7] verification tasks that do not involve concurrency properties (all categories except category "Concurrency"). All obtained experimental data as well as the tool implementation are available at `http://www.sosy-lab.org/~dbeyer/cpa-explicit`.

Quality Measures. We compare the verification results of all verification approaches based on three measures for verification quality: First, we take the run time, in seconds, of the verification runs to measure the *efficiency* of an approach. Obviously, the lower the run time, the better the tool. Second, we use the number of correctly solved instances of verification tasks to measure the *effectiveness* of an approach. The more instances a tool can solve, the more powerful the analysis is. Third, and most importantly, we use the scoring schema of the SV-COMP'12 as indicator for the quality of an approach. The scoring schema implements a community-agreed weighting schema, namely, that it is more difficult to prove a program correct compared to finding a bug and that a wrong answer should be penalized with double the scores that a correct answer would have achieved. For a full discussion of the official rules and benchmarks of the SV-COMP'12, we refer to the competition report [5]. Besides the data tables, we use plots of quantile functions [5] for visualizing the number of solved instances and the verification time. The quantile function for one approach contains all pairs (x, y) such that the maximum run time of the x fastest results is y. We use a logarithmic scale for the time range from 1 s to 1000 s and a linear scale for the time range between 0 s and 1 s. In addition, we decorate the graphs with symbols at every fifth data point in order to make the graphs distinguishable on gray-scale prints.

Improvements of Explicit-Value Analysis. In the first evaluation, we compare two different configurations of the explicit-value analysis: CPA-EXPL refers to the existing implementation of a standard explicit-value analysis without abstraction and

[6]http://cpachecker.sosy-lab.org

[7]http://sv-comp.sosy-lab.org/2012

Category	CPA-EXPL			CPA-EXPL$_{itp}$		
	points	solved	time	points	solved	time
ControlFlowInt	124	81	8400	123	79	780
DeviceDrivers	53	37	63	53	37	69
DeviceDrivers64	5	5	660	33	19	200
HeapManipul	1	3	5.5	1	3	5.8
SystemC	34	26	1600	34	26	1500
Overall	217	152	11000	**244**	**164**	**2500**

TABLE I: Comparison with purely explicit, non-CEGAR approach

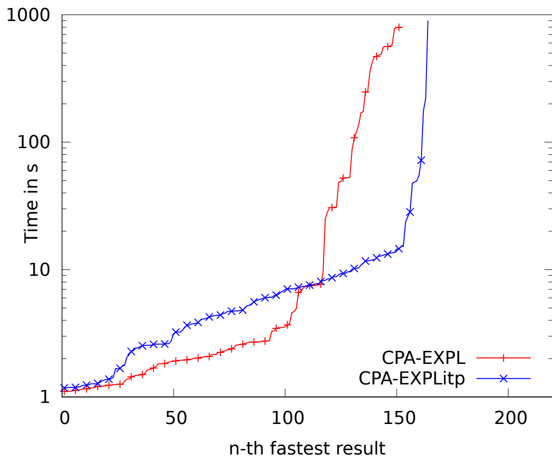

Fig. 3: Quantile plot: purely explicit analyses

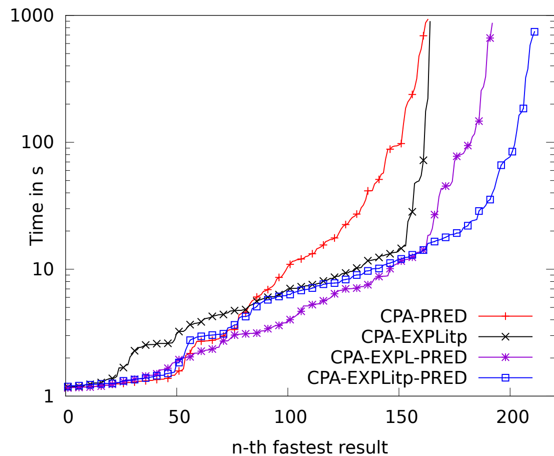

Fig. 4: Quantile plot: comparison with predicate-based configurations

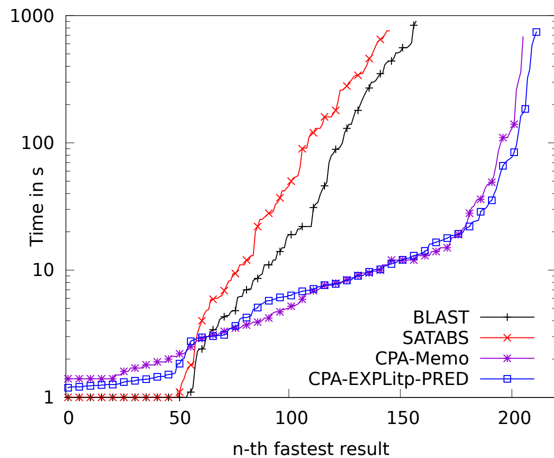

Fig. 5: Quantile plot: comparison with three existing tools

refinement, and CPA-EXPL$_{itp}$ refers to the new approach, which implements abstraction, CEGAR, and interpolation. Table I and Fig. 3 show that the new approach uses less time, solves more instances, and obtains more points in the SV-COMP'12 scoring schema.

Improvements of Combination with Predicate Analysis. In the second evaluation, we compare the refinement-based explicit analysis against a standard predicate analysis, as well as to the predicate analysis combined with CPA-EXPL and CPA-EXPL$_{itp}$, respectively: CPA-PRED refers to a standard predicate analysis that CPACHECKER offers (ABE-lf, [11]), CPA-EXPL$_{itp}$ refers again to the explicit-value analysis, which implements abstraction, CEGAR, and interpolation, CPA-EXPL-PRED refers to the combination of predicate analysis and explicit-value analysis without refinement, and CPA-EXPL$_{itp}$-PRED refers to the combination of predicate analysis and explicit-value analysis with refinement.

Table II and Fig. 4 show that the new combination approach outperforms the existing approaches CPA-PRED and CPA-EXPL$_{itp}$ in terms of solved instances and score. The comparison with column CPA-EXPL-PRED is interesting because it shows that the combination of two analyses is an improvement even without refinement in the explicit-value analysis, but switching on the refinement in both domains makes the new combination significantly more effective.

Comparison with State-of-the-Art Verifiers. In the third evaluation, we compare our new combination approach with three established tools: BLAST refers to the standard BLAST configuration that participated in the SV-COMP'12, SATABS also refers to the respective standard configuration, CPA-MEMO refers to a special predicate abstraction that is based on block-abstraction memoization, and CPA-EXPL$_{itp}$-PRED refers to our novel approach, which combines a predicate analysis (CPA-PRED) with the new explicit-value analysis that is based on abstraction, CEGAR, and interpolation (CPA-EXPL$_{itp}$). Table III and Fig. 5 show that the new approach outperforms BLAST and SATABS by consuming considerably less verification time, more solved instances, and a better score. Even compared to the SV-COMP'12 winner, CPA-MEMO, our new approach scores higher. It is interesting to observe that the difference in scores is much higher than the difference in solved instances: this means CPA-MEMO had many incorrect verification results, which in turn shows that our new combination is significantly more precise.

V. Conclusion

The surprising insight of this work is that it is possible to achieve —without using sophisticated SMT-solvers during the abstraction refinement— a performance and precision that can compete with the world's leading symbolic model checkers, which are based on SMT-based predicate abstraction. We achieved this by incorporating the ideas of abstraction,

Category	CPA-PRED			CPA-EXPL$_{itp}$			CPA-EXPL-PRED			CPA-EXPL$_{itp}$-PRED		
	score	solved	time	score	solved	time	score	solved	time	score	solved	time
ControlFlowInt	103	70	2500	123	79	780	131	85	2600	141	91	830
DeviceDrivers	71	46	80	53	37	69	71	46	82	71	46	87
DeviceDrivers64	33	24	2700	33	19	200	10	11	1100	37	24	980
HeapManipul	8	6	12	1	3	5.8	6	5	11	8	6	12
SystemC	22	17	1900	34	26	1500	62	45	1500	61	44	3700
Overall	237	163	7100	244	164	2500	280	192	5300	**318**	**211**	5600

TABLE II: Comparison with predicate-based configurations

Category	BLAST			SATABS			CPA-MEMO			CPA-EXPL$_{itp}$-PRED		
	score	solved	time	score	solved	time	score	solved	time	score	solved	time
ControlFlowInt	71	51	9900	75	47	5400	140	91	3200	141	91	830
DeviceDrivers	72	51	30	71	43	140	51	46	93	71	46	87
DeviceDrivers64	55	33	1400	32	17	3200	49	33	500	37	24	980
HeapManipul	–	–	–	–	–	–	4	9	16	8	6	12
SystemC	33	23	4000	57	40	5000	36	30	450	61	44	3700
Overall	231	158	15000	235	147	14000	280	209	4300	**318**	**211**	5600

TABLE III: Comparison with three existing tools

counterexample-guided abstraction refinement, lazy abstraction refinement, and interpolation into a standard, simple explicit-value analysis.

We further improved the performance and precision by combining our refinement-based explicit-value analysis with a predicate analysis, in order to benefit from the complementary advantages of the methods. The combination analysis dynamically adjusts the precision [9] for an optimal trade-off between the precision of the explicit analysis and the precision of the auxiliary predicate analysis. This combination out-performs state-of-the-art model checkers, witnessed by a thorough comparison on a standardized set of benchmarks.

Despite the overall success of our new approach, individual instances of benchmarks show different performance with different configurations — i.e., either with or without CEGAR. Therefore, a general heuristic for finding a suitable strategy for a single verification task would be beneficial. Also, we envision better support for pointers and data structures, because our interpolation approach can be efficiently applied even with high precision. Moreover, we so far only combined our interpolation approach with an auxiliary predicate analysis in the ABE-lf configuration, and we have not yet tried to combine this with the superior block-abstraction memoization (ABM) [26] technique. Finally, we plan to extend our interpolation approach to other abstract domains like intervals.

References

[1] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, 1986.

[2] A. Albarghouthi, A. Gurfinkel, and M. Chechik. Craig interpretation. In *Proc. SAS*, pages 300–316, 2012.

[3] T. Ball, A. Podelski, and S. K. Rajamani. Boolean and cartesian abstractions for model checking C programs. In *Proc. TACAS*, LNCS 2031, pages 268–283. Springer, 2001.

[4] T. Ball and S. K. Rajamani. The SLAM project: Debugging system software via static analysis. In *Proc. POPL*, pages 1–3. ACM, 2002.

[5] D. Beyer. Competition on Software Verification (SV-COMP). In *Proc. TACAS*, LNCS 7214, pages 504–524. Springer, 2012.

[6] D. Beyer, T. A. Henzinger, R. Jhala, and R. Majumdar. The software model checker BLAST. *Int. J. Softw. Tools Technol. Transfer*, 9(5-6):505–525, 2007.

[7] D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Path programs. In *Proc. PLDI*, pages 300–309. ACM, 2007.

[8] D. Beyer, T. A. Henzinger, and G. Théoduloz. Configurable software verification: Concretizing the convergence of model checking and program analysis. In *Proc. CAV*, LNCS 4590, pages 504–518. Springer, 2007.

[9] D. Beyer, T. A. Henzinger, and G. Théoduloz. Program analysis with dynamic precision adjustment. In *Proc. ASE*, pages 29–38. IEEE, 2008.

[10] D. Beyer and M. E. Keremoglu. CPACHECKER: A tool for configurable software verification. In *Proc. CAV*, LNCS 6806, pages 184–190. Springer, 2011.

[11] D. Beyer, M. E. Keremoglu, and P. Wendler. Predicate abstraction with adjustable-block encoding. In *Proc. FMCAD*, pages 189–197. FMCAD, 2010.

[12] D. Beyer and P. Wendler. Algorithms for software model checking: Predicate abstraction vs. IMPACT. In *Proc. FMCAD*, pages 106–113. FMCAD, 2012.

[13] S. Chaki, E. M. Clarke, A. Groce, S. Jha, and H. Veith. Modular verification of software components in C. *IEEE Trans. Softw. Eng.*, 30(6):388–402, 2004.

[14] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM*, 50(5):752–794, 2003.

[15] E. M. Clarke, D. Kröning, and F. Lerda. A tool for checking ANSI-C programs. In *Proc. TACAS*, LNCS 2988, pages 168–176. Springer, 2004.

[16] E. M. Clarke, D. Kröning, N. Sharygina, and K. Yorav. SATABS: SAT-based predicate abstraction for ANSI-C. In *Proc. TACAS*, LNCS 3440, pages 570–574. Springer, 2005.

[17] W. Craig. Linear reasoning. A new form of the Herbrand-Gentzen theorem. *J. Symb. Log.*, 22(3):250–268, 1957.

[18] S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *Proc. CAV*, LNCS 1254, pages 72–83. Springer, 1997.

[19] B. S. Gulavani, S. Chakraborty, A. V. Nori, and S. K. Rajamani. Automatically refining abstract interpretations. In *Proc. TACAS*, LNCS 4963, pages 443–458. Springer, 2008.

[20] K. Havelund and T. Pressburger. Model checking Java programs using JAVA PATHFINDER. *Int. J. Softw. Tools Technol. Transfer*, 2(4):366–381, 2000.

[21] T. A. Henzinger, R. Jhala, R. Majumdar, and K. L. McMillan. Abstractions from proofs. In *Proc. POPL*, pages 232–244. ACM, 2004.

[22] T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *Proc. POPL*, pages 58–70. ACM, 2002.

[23] G. J. Holzmann. The SPIN model checker. *IEEE Trans. Softw. Eng.*, 23(5):279–295, 1997.

[24] C. S. Pasareanu, M. B. Dwyer, and W. Visser. Finding feasible counter-examples when model checking abstracted Java programs. In *Proc. TACAS*, LNCS 2031, pages 284–298. Springer, 2001.

[25] A. Podelski and A. Rybalchenko. ARMC: The logical choice for software model checking with abstraction refinement. In *Proc. PADL*, LNCS 4354, pages 245–259. Springer, 2007.

[26] D. Wonisch. Block abstraction memoization for CPACHECKER. In *Proc. TACAS*, LNCS 7214, pages 531–533. Springer, 2012.