

Experimenting with Predicate Abstraction

Victor Cacciari Miraldo

HASLab/INESC TEC & Universidade do Minho, Portugal

Abstract. *Predicate abstraction* is a technique employed in software model checking to produce abstract models that can be conservatively checked for property violations in reasonable time. The precision degree of different abstractions of the same program may differ based on (i) the set of predicates used; or (ii) the algorithmic technique employed to generate the model. In this report we explain how we have extended the implementation of one such technique, that produces the most precise existential abstraction of a program, and we establish a common framework for both this direct technique and a second one, based on cartesian abstraction by weakest precondition calculations. This report is a product of the research grant BI22012_PTDC/EIA-CCO/117590/2010_UMINHO, in the scope of the AVIACC project, supervised by Professors Jorge Sousa Pinto and Maria João Frade

1 Introduction

Model Checking [8] has been successful in validating hardware system designs, to a point where its use has become not only common, but essential. It has long been hoped that this success will carry over to the realm of software, but this has proved exceptionally difficult, due to an aggravated state-space explosion problem. This is of course a typical problem of model checking, but the particular characteristics of software systems (like the presence of datatypes) make it particularly hard to handle.

Two main families of techniques have been employed in the last 10 to 15 years to make software model checking useful in practice. The first is *bounded model checking* [6], which in fact is not a specific technique for software applications. In a nutshell, it is employed as an alternative to BDD-based symbolic model checking that limits the exploration of executions of a system to a given bound on their length: paths longer than this bound are simply not explored. When applied to software, this amounts to limiting the number of loop iterations and recursive calls considered. Bounded model-checking problems can be encoded as logical *satisfiability* problems, solved with the help of a satisfiability solver. The drawback of this technique is that in general when a safety property is valid in the bounded model (e.g. a given error state is unreachable), it cannot be guaranteed to be valid in the original program.

The alternative is to consider an *abstract model* of the program, which may dramatically reduce the state-space. The goal of abstraction is to compute an abstract model \bar{M} from the concrete model M of a program, such that the size

of the state-space is reduced, but in a way that it is still sound to check for safety properties. Abstraction [11] is of course the fundamental technique that makes static program analyses feasible; the kind of abstraction used for software model checking safety properties is known as *existential abstraction* [9]. Informally, existential abstraction produces an abstract model \bar{M} from a concrete model M of a program so that any reachability problem that is solvable in M is also solvable in \bar{M} , which means that if the abstract model satisfies a given safety property, then so does the original program (i.e. concretization preserves safety).

This report is about the discussion and implementation of the first step of a software model checker, namely, it's about a specific existential abstraction technique known as *Predicate Abstraction* [14], which has the advantage over other abstraction methods that it can be computed algorithmically. Predicate abstraction keeps track of a given set E of predicates over the data, and registers how the truth value of these predicates changes with the concrete program steps. Computing a predicate abstraction of a program inherently takes exponential time on its length; but predicate abstractions are not unique, and there exist different approaches to computing them, with the usual trade-off between precision and efficiency.

In any working software model checker predicate abstraction is implemented as part of a refinement loop. In *Counter Example-Guided Abstract Refinement* (CEGAR), the refinement that is performed at each iteration consists of adding more predicates to E . The loop successively calculates predicate abstractions, starting from a rougher model, and refines them by generating new predicates, based on the false positive counter-examples returned by the model checker at each stage, until no counter-example is returned (in which case the program is safe) or a true counter-example is found (the program is unsafe, there exists a real property violation). Abstract refinement is not covered in the report: we focus on predicate abstraction using a fixed set of predicates, which is the fundamental (and costly!) building block of any software model checker.

Although the literature on software model checking accumulated over these last 15 years is vast, we found that no clear and uniform presentation of the different approaches to predicate abstraction can be found. In this report we take two such approaches, each of which has been proposed as part of a major software model checker. We remark that the precision degree of two predicate abstractions of the same program may differ based on (i) the set E of predicates provided (using more predicates results in a more precise model); or (ii) the algorithmic technique used to generate the model from E . Given a set E , the two methods considered here differ in the degree of precision of the generated abstract models. The first method is the one found in the MAGIC software model checker developed at CMU [7]; it is a straightforward technique that focuses on directly computing minimal existential abstractions by solving SAT problems. The second method (more efficient but less precise) is based on *cartesian abstraction*; it can be found at the core of the Microsoft SLAM tool [5].

The two methods are unfortunately described in the literature in terms that make them hard to compare or to implement based on the same backbone:

whereas the direct method works at the level of transition systems, the second method produces a *Boolean program* (so different model checking tools have to be used on the abstract models produced). In addition, cartesian abstraction is implemented using *weakest precondition* computations, whereas the direct method is based on satisfiability checks. The two methods are surprisingly similar, the goal of this report is to provide the aforementioned common backbone, for the double purpose of reasoning and implementation. We start by providing a brief background on the technical context of the project, the SPARK language and Boolean Programs are introduced in Sections 2 and 3. A high level presentation of the prototype produced follows in Section 4. Section 5 introduces the basic concepts of predicate abstraction; Sections 6 and 7 then present the specific details of each of the two methods we consider. Section 8 shows how to calculate the cartesian abstraction by SAT and points out how to create the direct abstraction inside the cartesian framework. Section 9 explains how we have implemented and optimized the direct satisfiability method to produce Boolean programs.

2 The SPARK Language

This research was mainly focused on software model checking, but for developing our prototype, we had to choose a source language. We chose SPARK for two reasons: (A) SPARK has a somewhat simple syntax/semantic compared to other commercial languages and; (B) SABS [18] is a joint effort with SPARK-BMC [17] to create both a software model checker and a bounded model checker for SPARK, therefore the development efforts could be divided.

For the rest of this document, no previous knowledge of SPARK is required. In fact, we only handle basic control-flow and operations, therefore we'll only briefly, and informally, present the SPARK toolset.

Ada programming language can help avoid errors that are common in other languages but, as said before, in the development of highly critical systems this is often insufficient. SPARK [3] appeared in 1988 to target precisely these systems. Although SPARK is based on a heavily restricted subset of Ada together with a set of annotations, it should be considered in its own right as a full language for the development of annotated high-assurance software. The great advantage of using a subset of a widely used language is that this makes possible to share compilers, instead of developing a new one. A consequence of this is that annotations in SPARK code must be written as Ada comments, ignored by compilers but not by the SPARK verification tools.

SPARK is not just a language, but also a set of tools that not only check if a program respects all the restrictions imposed on valid SPARK programs, but are also probably the most widely used tools for program verification. The Examiner is the tool responsible for performing syntactic and static semantic analyses for checking the validity of SPARK programs, as well as generating verification conditions. Yet, there are no model-checking tools targeting SPARK.

3 Boolean Programs

Independently of the domain we're performing model checking on, we need to choose a representation for our models. In the case of software model checking, the most common practice is to use *boolean programs* as this representation.

Boolean programs (BP) are a subset of the programs defined by the source language, namely those where the only type available is the boolean type. That is, every variable and parameter will have boolean type.

There are various reasons for choosing BPs as the model for which we'll abstract to. Boolean programs are sequential programs, therefore they share the same structure; The problem of model checking BPs is decidable, since they're equivalent to pushdown automaton, which accept context-free languages; And there already exists industrial strength tools for doing so, such as BEBOP[2], BOOM[4] or GETAFIX [16].

The syntax for boolean programs was derived from C, and can be found on [2]. The syntax for expressions permits the canonical Boolean operators with two extensions: (A) non-deterministic choice, denoted by $*$, and (B) next-state variables, denoted by p' .

A simple example of a boolean program is given in figure 1.

```
void main() begin
  decl p1, p2;

  p1, p2 := T, T;
  while * do
    assume(!p2);
    p1, p2 := (p1 | p2) ? F : *, !p2;
  od;

  assert(p1 & p2);
end
```

Fig. 1. A sample boolean program.

The semantics are fairly simple and standard, apart from the `assume` keyword. Whenever we're running along an execution path and find a `assume` ψ , the condition ψ is evaluated. Case ψ evaluates to true, we continue on the same execution path, whereas, case ψ evaluates to false, we cancel the execution of the current path and backtrack to a point where a different execution path can arise. Note that `assume` works as a filter over execution paths on the boolean programs.

4 SABS Description

The process of producing a boolean program, by abstraction, from a SPARK program is, for practical purposes, a pipeline of program transformations. In this section we'll briefly describe the pipeline implemented, with more detail on the *Static Single Assignment* step.

The actual abstraction is performed only after a series of simplification steps. Even SPARK is a complex enough language to justify the implementation effort, the intermediate steps or transformations applied to the target program are:

- (i) Parsing
- (ii) Loop and Conditional simplification
- (iii) Static Single Assignment (SSA) form
- (iv) Existential Abstraction

The target, annotated, program is parsed to a AST based on the SPARK grammar [3,20]. Our AST is modified to support special annotations, flagged by the token `--%` and used both to configure the following steps and to specify the predicates we wish to prove for the program (with an `--% assert p;`). We then transform every loop to its `while` equivalent and simplify the `elsif` statements.

We then follow to simplify the resulting, simplified program, to its SSA form. This transformation is very important, and we depend heavily on the fact that the Weakest Precondition of an SSA program $wp(p, \phi)$ can be shown to be equivalent to $\phi \rightarrow L_p$, where L_p is the logic encoding of p . More on the subject on [13].

The last step, the existential abstraction of a SSA, SPARK-based program is described in the rest of this document. For more details on the tool itself, we direct the reader to the SABS tutorial [18].

5 Predicate Abstraction

We start by defining formally the notions of model and existential abstraction. A *model* M is defined by a triple (S, S_0, T) , where S is the set of *states*, $S_0 \subseteq S$ is the set of *initial states*, and $T \subseteq S \times S$ is the *transition relation*. In what follows we will require *abstraction* functions mapping states of a model into states of another model, which we will extend to sets of states as expected. A *concretization function* $\gamma : \bar{S} \rightarrow S$ mapping abstract states into some concrete states is associated to each abstraction function.

Definition 1. A model $\bar{M} = (\bar{S}, \bar{S}_0, \bar{T})$ is an existential abstraction of another model $M = (S, S_0, T)$ w.r.t. an abstraction function $\alpha : S \rightarrow \bar{S}$ if

1. $\exists s \in S_0. \alpha(s) = \bar{s} \rightarrow \bar{s} \in \bar{S}_0$
2. $\exists (s, s') \in T. \alpha(s) = \bar{s} \wedge \alpha(s') = \bar{s}' \rightarrow (\bar{s}, \bar{s}') \in \bar{T}$

The minimal existential abstraction also satisfies the converse implications.

We will use one form of existential abstraction called *predicate abstraction*, where we abstract data by keeping track of predicates on it; every operation on the concrete model M will be translated to a Boolean operation on the abstract model \bar{M} . Predicate abstraction can be applied in the model checking of transition systems in general; when applied to software model checking, it produces a *Boolean program*, i.e. a program whose only data consists of a set of Boolean variables, with the same control-flow structure as the original program.

The idea is simple: given a predicate p_i on the variables of the original program, there will be a corresponding Boolean variable b_i in the abstract Boolean program; instructions in the original program will be abstracted into instructions on the Boolean variables, that reflect the effects of the original instructions on the truth values of the predicate. In particular, sequences of assignment instructions are mapped into parallel assignments. As a very simple example, the instruction $\mathbf{x} := -\mathbf{x}$ would be abstracted with the predicates $p_1 \doteq x \leq 0$, and $p_2 \doteq x > 0$ as the following parallel assignment: $\mathbf{b1}, \mathbf{b2} := \mathbf{b2}, \mathbf{b1}$.

Throughout the paper we will write E for the set of predicates used to construct the abstraction, and V for the set of Boolean variables in the Boolean program, with $\#V = \#E$. We will denote by $E(b)$ the predicate that is represented by the variable $b \in V$, and extend this notion to Boolean expressions in the natural way, for instance $E(b_1 \wedge b_2) = E(b_1) \wedge E(b_2)$.

To see how this matches our general discussion of abstraction, let us denote the set of Boolean values by $\mathbb{B} = \{T, F\}$. A concrete state consists of the program location $l \in \mathcal{L}$ and an assignment to its variables (for the sake of simplicity we consider that concrete programs manipulate only integer variables). Given a concrete state s , we will denote by $b_i(s)$ the logical value of the predicate $E(b_i)$ in s . The corresponding abstract state \bar{s} consists of the program location $l \in \mathcal{L}$ and a valuation of the propositional variables in V , ie, $\bar{S} = \mathcal{L} \times \mathbb{B}^n$, where $n = \#E$. The abstraction function α will map a concrete state into an abstract one, and is defined by: $\alpha(s) = (loc(s), b_1(s), \dots, b_n(s))$.

Suppose that the original program contains a command **assert** A , and one wants to model check the (safety) property that whenever the command is reached the Boolean expression A is true. If $A = p_i \in E$ for some i , then the command will be translated into **assert** b_i in the Boolean program, which can now be model-checked (if not, then a suitable expression constructed from the b_i must be used instead). The advantage of doing this is that the reachability problem for Boolean programs is decidable [2]. Dedicated model checkers for Boolean programs include BEPOP [2] and BOOM [4]. Note that predicate abstraction and Boolean model checking do not absolutely require working with Boolean programs as we do here. We could calculate the abstraction at the level of transition systems, and then model-check the abstract transition system using a general-purpose model checker. This will be further explained in Section 6.

As stated before, how to choose and refine a suitable set of predicates for a given program is outside the scope of this paper: we assume a fixed set E is provided, and consider two methods to construct an abstraction based on E . The methods differ only in the way that basic blocks of code (sequences

of assignment instructions) are treated: control-flow is basically preserved from the concrete to the Boolean program. Also, the treatment of data structures like arrays, structures, and pointers, is orthogonal to the choice of abstraction method. As such, in what follows we will essentially consider basic blocks as concrete programs, consisting of sequences of integer assignment instructions.

6 The Direct Method

The most straightforward way to compute a predicate abstraction is to interpret Definition 1 at the level of programs and apply it directly with the help of a satisfiability solver. The method is described at length in [10]; it is used in practice in the MAGIC tool [7].

As an abstract state is given by the values of the propositional variables in V induced by the logical values of the predicates in E (in a given concrete state), one needs to test which combinations of logical values of the predicates before and after execution of the block are feasible. For this we need first of all to have a logical encoding of the block, which can be obtained by converting it to *static single assignment* (SSA) form [12]. Take for instance the basic block:

$$P \equiv x := x + 10; y := y + 1$$

It is converted to the following: $P \equiv x_1 := x_0 + 10; y_1 := y_0 + 1$. The logical encoding L_P of P can now be written simply as a conjunction of equations, $L_P \equiv x_1 = x_0 + 10 \wedge y_1 = y_0 + 1$. This encoding can be applied to a much richer language, including arrays, structures, and pointers [10].

Note that each variable in the concrete program is now represented by a family (both of size 2, in the above example) of variables in its logical representation. Of these we are only interested in the initial and final versions of each variable; when considering the execution of a basic block we will in general denote by s and s' the program state expressed in terms respectively of the initial and final versions of the variables, therefore $b_i(s)$ and $b_i(s')$ will denote the initial and final values of the predicate $E(b_i)$. Consider again our example program and take for instance $V = \{b_1, b_2\}$ with $E(b_1) = x \geq 0$ and $E(b_2) = \text{even}(y)$. Then $b_1(s)$ is $x_0 \geq 0$, $b_2(s)$ is $\text{even}(y_0)$, $b_1(s')$ is $x_1 \geq 0$, and $b_2(s')$ is $\text{even}(y_1)$.

Let us now introduce some basic definitions and notation. A *literal* is a Boolean variable or its negation. Let $V = \{b_1, \dots, b_n\}$ be a set of Boolean variables. A *cube* over V is a conjunction of literals in which each of the variables of V appears exactly once. A *cover* is a disjunction of cubes. We let l, l_i, \dots range over literals, and c, c_i, c', \dots range over cubes. We will denote by C_V the set $\{c_1, \dots, c_{2^n}\}$ of all possible cubes over V . Note that each such cube uniquely corresponds to a valuation of the propositional variables in V , and thus to an abstract state. For instance the cube $b_1 \wedge \neg b_2$ corresponds to the abstract state in which b_1 is true and b_2 is false.

Following existential abstraction, to decide whether to include in the abstract model a transition from the state characterized by the cube c_i to the state

characterized by the cube c_j , it suffices to check the satisfiability of the formula

$$E(c_i)(s) \wedge L_P \wedge E(c_j)(s')$$

Say, for the program and predicates given above, we wish to check the existence of a transition from the state in which both predicates are false to the state in which both are true. We check the satisfiability of

$$\neg(x_0 \geq 0) \wedge \neg(\text{even}(y_0)) \wedge x_1 = x_0 + 10 \wedge y_1 = y_0 + 1 \wedge x_1 \geq 0 \wedge \text{even}(y_1)$$

Indeed the formula is satisfiable, for instance with $x_0 = -2$ and $y_0 = 0$, and the transition will thus be included in the abstract model. The abstract transition system can be constructed by exhaustively testing 2^{2n} formulas:

$$\begin{array}{c} E(c_1)(s) \wedge L_P \wedge E(c_1)(s') \\ \vdots \\ E(c_{2^n})(s) \wedge L_P \wedge E(c_{2^n})(s') \end{array}$$

The formulas can be checked by an SMT solver using a theory of (unbounded) integers, or, if one wishes to employ a fixed-size bitvector encoding of numbers (that stands closer to the machine representation), by a SAT solver after bit-blasting. Every satisfiable formula (corresponding to an abstract transition in the model) from the above family is recorded, allowing us to calculate an *assignment table*. In our example the solver would return the following table:

b_1	b_2	b'_1	b'_2
F	F	F	T
F	F	T	T
F	T	F	F
F	T	T	F
T	F	T	T
T	T	T	F

Definition 2 (Assignment Table). *Let P be a basic-block and V a set of boolean variables associated with a set of predicates, E . Then, the assignment table (transition relation) of such variables according to P is given by:*

$$AT_V(P) = \{(c, c') \in C_V \mid \exists s, s' . E(c)(s) \wedge L_P \wedge E(c')(s')\}$$

As described in [10], the method is used to produce an abstract transition system, which the authors then export to a general-purpose symbolic model checker to find property violations. But our interest is not in exporting the abstract model in the form of a transition relation; instead, we would like to produce a *Boolean program*. The reasons for this are twofold: first, specific model checkers for Boolean programs are of course fine-tuned for this problem, and thus handle it more efficiently. Second, other methods for generating predicate abstractions produce Boolean programs natively, and so do most existing software model checking tools; for the sake of uniformity (and to facilitate comparison) we also choose to follow the latter approach.

Informally, to produce boolean programs from the assignment table, we need to determine which pre-state cubes turn each variable true or false. More details on how to "divide" the assignment table will be given in section 9, formally, we want a projection of a binary relation.

Definition 3 (Projection). *Let V be a set of boolean variables, $R \subseteq C_V^2$ and b a cube over V . We'll define the projection of R in respect to V as:*

$$\pi_b(R) = \{c \in C_V \mid \exists c' \in C_V . c R c' \wedge b \subseteq c'\}$$

7 Cartesian Abstraction by WP Computations

SLAM [5], the tool that might be called the most successful software model checker (it has become a commercial product, currently shipped by Microsoft as part of the Windows Driver Development Kit), uses a different method for constructing predicate abstractions. It constructs less precise abstractions, and naturally does so more efficiently than the direct method.

Let again P be a basic block and L_P its logical encoding, E be the set of predicates used to construct the predicate abstraction, and V the set of Boolean variables. An alternative to using satisfiability tests is to employ weakest precondition (WP) calculations. Recall that the weakest precondition of a basic block with respect to a given assertion ψ is given by the following two rules:

$$wp(x := e, \psi) \doteq \psi[e/x] \quad wp(C_1; C_2, \psi) \doteq wp(C_1, (wp(C_2, \psi)))$$

Recall the example program and predicates of the previous section. Then

$$wp(P, E(b_1)) \equiv x \geq 0[y + 1/y][x + 10/x] \equiv x + 10 \geq 0$$

One way to construct an abstraction is to determine individually, for each Boolean variable $b \in V$, the sets of states in which the weakest preconditions $wp(P, E(b))$ and $wp(P, E(\neg b))$, respectively, are satisfied. In the first set of states execution of the block will make $E(b)$ hold in the final state, thus the assignment $b := T$ should be executed by the Boolean program. In the second set of states $b := F$ should be executed, and in states in which neither $wp(P, E(b))$ nor $wp(P, E(\neg b))$ are satisfied, the assignment $b := *$, signaling a non-deterministic assignment, should be executed. It is useful to employ the following function:

$$choose(pos, neg) = pos ? T : (neg ? F : *)$$

The general idea is that the basic block can be abstracted by a parallel assignment of the form $\dots, b, \dots := \dots, choose(wp(P, E(b)), wp(P, E(\neg b))), \dots$. But this is of course not a valid Boolean program, since $wp(P, E(b))$ cannot be expressed in terms of the Boolean variables. What we can do in the Boolean program is to determine the combinations of values of the Boolean variables that force each of the above WPs to hold. This can be formalized as follows. Given an assertion ψ , let S_V^ψ denote the following disjunction of cubes over V :

$$S_V^\psi = \bigvee \{c \in C_V \mid \models E(c) \rightarrow \psi\}$$

(V will be dropped when clear from context) Note that constructing this set requires 2^n validity tests, where $n = \#V$. Then P is abstracted by the following Boolean program, where ϕ_i denotes the assertion $wp(P, E(b_i))$:

$$b_1, \dots, b_n := \text{choose}(S^{\phi_1}, S^{\neg\phi_1}), \dots, \text{choose}(S^{\phi_n}, S^{\neg\phi_n})$$

Observe that computing the abstraction in this way requires testing the validity of $2n \times 2^n$ formulas. This is still exponential, but also exponentially better than the direct method. This method introduces more false positives than the direct method because the different predicates are considered independently of each other, thus contradictory states are present in the models. This is in fact what is known as *cartesian abstraction*.

To understand this, consider that E consists of the two predicates $E(b_1) \equiv x \geq 0$ and $E(b_2) \equiv x \leq 100$. This produces an *unsatisfiable cube*: $E(\neg b_1 \wedge \neg b_2) \equiv x < 0 \wedge x > 100$, which is a contradiction, and would be included in S_V^ψ for any condition ψ : there exists a transition from the state corresponding to the unsatisfiable cube to any other state. Moreover, transitions *into* this state could also be present, since the WPs are computed independently for $E(b_1)$ and $E(b_2)$. Compare this to what would happen with the direct method: any satisfiability formula involving a contradictory state, of the form

$$E(\neg b_1 \wedge \neg b_2)(s) \wedge L_P \wedge E(c_j)(s') \quad \text{or} \quad E(c_i)(s) \wedge L_P \wedge E(\neg b_1 \wedge \neg b_2)(s')$$

is UNSAT, and rejected from the assignment table. Thus the corresponding transitions will not be inserted in the construction of the abstract model.

8 Cartesian Abstraction by SAT

Cartesian abstraction may look different, in terms of description, from direct abstraction yet they share more similarities than what one would think. In fact, the actual difference is that cartesian abstraction consider the boolean variables one by one while the direct method uses all the possible combinations (cubes).

A first observation to make is that in the SSA setting weakest preconditions can be computed without substitution, based on the same logical encoding of a program used in the direct method [13]. With this in mind, we can arrive at a equivalente, SAT-based, definition of $S^{wp(P, E(b_i))}$.

$$\begin{aligned} S^{wp(P, E(b_i))} &= \bigvee \{c \in C_V \mid \models E(c)(s) \rightarrow wp(P, E(b_i))\} \\ &= \bigvee \{c \in C_V \mid \forall s. E(c)(s) \rightarrow wp(P, E(b_i))\} \\ &= \bigvee \{c \in C_V \mid \forall s, s'. E(c)(s) \rightarrow L_P \rightarrow E(b_i)(s')\} \\ &= \bigvee \{c \in C_V \mid \forall s, s'. \neg E(c)(s) \vee \neg L_P \vee E(b_i)(s')\} \\ &= \bigvee \{c \in C_V \mid \exists s, s'. E(c)(s) \wedge L_P \wedge \neg E(b_i)(s')\} \\ &= \bigvee \{c \in C_V \mid \exists s, s'. E(c)(s) \wedge L_P \wedge E(\neg b_i)(s')\} \end{aligned}$$

which provides a basis for the relation we are seeking to establish. Note that to produce boolean programs, it's irrelevant if we're using the WP or the SAT-based encoding of S .

Indeed, the above satisfiability problems are very close to the formulas $E(c_i)(s) \wedge L_P \wedge E(c_j)(s')$ checked for satisfiability in the direct method. We can, in fact, prove some interesting results:

Lemma 1. *Let P be a basic block, V a set of variables (associated with a set or predicates E) and b_i a literal over V . Then*

$$\pi_{\neg b_i}(AT_V(P)) \subseteq S_V^{wp(P, E(b_i))}$$

Proof. We know that E distributes over the propositional connectives, and since $a \wedge b \rightarrow a$, the proof idea is to remove every literal that is not needed in the post-state cubes used by the direct abstraction, thus arriving at the SAT encoding of $S_V^{wp(P, E(b_i))}$. We then have that:

$$\begin{aligned} \exists k \in C_V. \exists s, s'. E(c)(s) \wedge L_P \wedge E(k)(s') \wedge \neg b_i \in k \\ \Rightarrow \exists s, s'. E(c)(s) \wedge L_P \wedge E(\neg b_i)(s') \end{aligned}$$

Note that the first line is precisely the condition for a cube c to be an element of $\pi_{\neg b_i}(AT_V(P))$. Therefore the inclusion is proved. \square

This proof is valid in a theoretical point-of-view, yet due to simplification and optimization mechanisms we are not computing the actual covers $\pi_{\neg b_i}(AT_V(P))$ or $S_V^{wp(P, E(b_i))}$, but the set of all essential prime implicants of such covers, thus returning covers that could, wrongly, be seen as counter-examples of the above lemma.

From lemma 1 we can see that the difference between the two abstraction methods is, mainly, the number of variables considered simultaneously. Hence, we could extend the cartesian abstraction to compute $S_V^{wp(P, f)}$, for arbitrary formulas f . In case those formulas are a *complete disjunction* of literals over V , we can see that the projection of the assignment table for such formula and the calculation of the cartesian abstraction of the same formula coincide:

Theorem 1. *Let P be a basic block, $V = \{b_1, \dots, b_n\}$ a set of variables (associated with a set or predicates E) and k a complete disjunction of literals over V , then:*

$$\pi_{\neg k}(AT_V(P)) = S_V^{wp(P, E(k))}$$

Proof. The proof is straight forward from the SAT encoding of $S_V^{wp(P, E(k))}$, noting that $\neg k$ is a complete cube over V :

$$\begin{aligned} S_V^{wp(P, E(k))} &= \bigvee \{c \in C_V \mid \exists s, s'. E(c)(s) \wedge L_P \wedge E(\neg k)(s')\} \\ &= \bigvee \{c \in C_V \mid \exists j \in C_V. \exists s, s'. E(c)(s) \wedge L_P \wedge E(j)(s') \wedge \neg k = j\} \\ &= \pi_{\neg k}(AT_V(P)) \end{aligned}$$

\square

From theorem 1 we can see that it's possible to implement direct abstraction by WP calculations:

$$AT_V(P) = \{(c, c') \in C_V \mid c \in S_V^{wp(P, E(-c'))}\}$$

This common framework provides a way of combining optimizations, but some care must be taken. Due to simplification and optimization mechanisms, we should check that $c \rightarrow S_V^{wp(P, E(-c'))}$ rather than $c \in S_V^{wp(P, E(-c'))}$.

9 Implementation of The Direct Abstraction

The description of the direct method in Section 6 is just the first half of the story: we have indeed identified the valid transitions in the abstract model, but our goal is to produce a Boolean program. In this section we explain how we have implemented the direct algorithm so that it outputs a Boolean program.

Our goal is to abstract a basic block as a parallel assignment of Boolean variables of the form $b_1, \dots, b_n := e_1, \dots, e_n$. The task is then to find the right-hand side expressions e_1, \dots, e_n , given an assignment table. To this end the table is first divided into n tables, one for each variable in the final state. Each resulting table is then divided into its ON-set and OFF-set (that is, the assignments that turn each output variable to T and F, respectively). In our example this yields the two tables shown on the left below.

b_1	b_2	b'_1	b_1	b_2	b'_2	b_1	b_2	b'_1	b_1	b_2	b'_2
F	F	T	F	F	T	T	F	T	F	F	T
F	T	T	F	F	T	T	T	T	T	F	T
T	F	T	T	F	T	F	F	*	F	T	F
T	T	T	F	T	F	F	F	*	F	T	F
F	F	F	F	T	F	F	F	F	T	T	F
F	T	F	T	T	F	T	T	F	T	T	F

Note that the first table contains non-determinism: the same combination of values of b_1 and b_2 may result in different values for b'_1 . The second table on the other hand contains redundancy (repeated entries than can be removed). We rewrite and simplify the tables as shown on the right. Note that the first table now has what one might call an UNDET-set rather than an OFF-set. The ON, OFF and UNDET-sets constitute a partition of the set of assignments according to the possible results of the output variable. Each of these sets is captured by a Boolean formula which is the disjunction of the cubes that characterize each assignment in the set. We call these formulas respectively ON, OFF and UNDET-covers.

Definition 4. Let P be a basic-block and V a set of boolean variables, for each b_i in V we define three formulas:

$$\begin{aligned} \text{UNDET}_i &= \pi_{b_i}(AT_V(P)) \cap \pi_{\neg b_i}(AT_V(P)) \\ \text{ON}_i &= \pi_{b_i}(AT_V(P)) - \text{UNDET}_i \\ \text{OFF}_i &= \pi_{\neg b_i}(AT_V(P)) - \text{UNDET}_i \end{aligned}$$

A parallel assignment can be directly extracted from these tables by using these covers: $b_1, b_2 := ((b_1 \wedge \neg b_2) \vee (b_1 \wedge b_2)? T : *)$, $((\neg b_1 \wedge \neg b_2) \vee (b_1 \wedge \neg b_2)? T : F)$, which can in turn be simplified to $b_1, b_2 := (b_1? T : *)$, $(\neg b_2? T : F)$. If the UNDET-set is not empty a nested conditional expression will have to be used. In fact, although this has to our knowledge never been made explicit, the parallel assignment can be written as follows using the *choose* function of Section 7:

$$b_1, \dots, b_n := \text{choose}(\text{ON}_1, \text{OFF}_1), \dots, \text{choose}(\text{ON}_n, \text{OFF}_n)$$

It is clear from this small example that it would be infeasible to export a Boolean program without first attempting to simplify the assigned expressions; let us now describe how we have implemented this simplification.

Boolean simplification. The minimization of a Boolean function is a well-known problem in the area of logic circuit design: a circuit with a large number of logic gates (equivalent to a complex Boolean function) takes up a lot of physical space in its implementation. This problem is believed to be intractable [15], but there exist effective heuristics for it, such as Karnaugh Maps and the Quine-McCluskey algorithm. Our testbed is implemented using a functional programming language; for this reason we have opted for a recursive algorithm based on the *prime consensus theorem*, described in R. Rudell's thesis [19] (Sect. 2.5.1).

First, let us introduce some definitions and notation. In what follows a *cube* is simply a conjunction of literals. Associativity, commutativity and idempotence of conjunctions and disjunction allow us to treat each cube as a set of literals and each cover as a set of sets of literals.

Given two cubes, c, c' , we say they *differ in a variable* x if $x \in c$ and $\neg x \in c'$ (or vice-versa). The *distance* between c and c' , written $\text{dist}(c, c')$ is the number of variables where they differ. When $\text{dist}(c, c') = 0$ we say that c and c' *intersect* and the *intersecting cube* is $c \cup c'$.

The *consensus* of two non-intersecting cubes, c and c' , $\text{consensus}(c, c')$, is defined as follows: if $\text{dist}(c, c') \geq 2$, their consensus is empty; if $\text{dist}(c, c') = 1$, their consensus is $(c \cup c') - \{x, \neg x\}$, assuming c, c' differ in x . The notion of consensus is lifted to sets of cubes, as the pairwise consensus of the two sets.

Given two cubes, c, c' , we say that c' is *single-cube contained* in c if $c \subseteq c'$. Given a set of cubes C , the *single-cube containment* of C is the set $\text{SCC}(C) = \{c \mid \exists c' \in C. c \neq c' \wedge c \subseteq c'\}$. Let f be a Boolean function. A cube c is an *implicant* of f whenever $c \rightarrow f$. Moreover, we say that c is a *prime implicant* of f if c is minimal, i.e., there is no other implicant of c except itself. The set of prime implicants of f is denoted by $\text{primes}(f)$.

We can now state the fundamental theorem that stands at the heart of the simplification algorithm we have implemented.

Theorem 2 (Prime consensus theorem). *Let f be a Boolean function and let x be any input variable. The set of prime implicants of f can be partitioned into three sets: $P_x = \{c \in \text{primes}(f) \mid x \in c\}$, $P_{\neg x} = \{c \in \text{primes}(f) \mid \neg x \in c\}$ and $P_* = \{c \in \text{primes}(f) \mid x \notin c \wedge \neg x \notin c\}$. Then,*

$$\forall c \in P_*. \exists c \in P_x. \exists c' \in P_{\neg x}. c = \text{consensus}(c, c')$$

This theorem states that $P_* \subseteq \text{consensus}(P_x, P_{\neg x})$, because the consensus of $P_x, P_{\neg x}$ may contain non-prime implicants. We can get rid of such non-primes by constructing the single-cube containment of that set. We have

$$P_* = \text{SCC}(\text{consensus}(P_x, P_{\neg x}))$$

Now that we know how to generate P_* from P_x and $P_{\neg x}$, let us focus on the construction of P_x and $P_{\neg x}$ given a cover F of a Boolean function, and an input variable x .

A *cofactor* of F with respect to a literal l , written F_l , is defined as follows $F_l = \{c - \{l\} \mid c \in F \wedge l \in c\}$. In fact, $P_l \subseteq \{l\} \cup \text{primes}(F_l)$. So, as before, we have to get rid of the non-primes.

The following theorem summarizes how the prime implicants of a Boolean function f can be generated recursively, and is effectively an algorithm outline.

Theorem 3 (Recursive prime generation theorem). *Let f be a Boolean function with (ON+UNDET)-cover F and let x be any input variable. Then, the prime implicants of f can be generated as follows:*

$$\text{primes}(f) = \text{SCC}(A_x \cup A_{\neg x} \cup \text{consensus}(A_x, A_{\neg x})) , \text{ where } A_l = \{l\} \cup \text{primes}(F_l)$$

Note that in this divide and conquer approach, the choice of division point (the splitting variable x) will have major impact on the algorithm's efficiency. Clever rules for termination have been proposed that can speed up the process [19].

Handling variable initialization and optimizations. The predicate abstraction constructed by this method naturally eliminates transitions *from* and *to* states corresponding to *unsatisfiable cubes*, as shown at the end of Section 7.

We have introduced two modifications in the original algorithm, which we now describe. The first has to do with the fact that this method does not deal well with variable initialization in the presence of unsatisfiable cubes. To see this, let P be the basic block $x := 10$, with the two previous predicates. Clearly the block should be abstracted to the Boolean program $b1, b2 := T, T$. For this, the expected assignment table would be:

b_1	b_2	b'_1	b'_2
T	T	T	T
T	F	T	T
F	T	T	T
F	F	T	T

Observe that the last row will *not* be in the table, since the following is not satisfiable (x_0 cannot be smaller than 0 and greater than 100 at the same time):

$$(x_1 = 10) \wedge \neg(x_0 \geq 0) \wedge \neg(x_0 \leq 100) \wedge (x_1 \geq 0) \wedge (x_1 \leq 100)$$

Our guess is that tools based on the direct method calculate predicate abstractions after running a *constant propagation* transformation. We propose a modification of the algorithm that does not require this transformation.

Algorithm 1 Abstraction of Basic Blocks

```
for  $pos_c \in \text{cubesOf}(E)$  do
   $pos_f \leftarrow L_p \wedge \text{instantiate}(\text{poststate}, pos_c)$ 
  if  $\text{solve}(pos_f) = SAT$  then
     $preds \leftarrow \{p \in E : \text{varsOf}(p) \cap \text{dependentVariables}(L_p) \neq \emptyset\}$ 
    if  $preds = \emptyset$  then  $\text{addAllCombinationsFor}(pos_c)$ 
    else
      for  $pre_c \in \text{cubesOf}(preds)$  do
         $full_f \leftarrow pos_f \wedge \text{instantiate}(\text{prestate}, pre_c)$ 
        if  $\text{solve}(full_f) = SAT$  then
           $\text{addLines}(pos_c, \text{interpolate}(\text{independentVariables}(L_p), pre_c))$ 
        end if
      end for
    end if
  end if
end for
```

The second modification is an optimization: we initially run a battery of satisfiability checks of formulas combining the program and the post-state cubes. Admittedly this takes time 2^n , but observe that for each unsatisfiable cube found we save 2^n checks, one for each pre-state cube. Moreover, this initial round of checks also eliminates 2^n checks for every post-state corresponding to a cube that, although satisfiable, can never be attained by the program (such as $b_1 \wedge \neg b_2 \equiv x > 100$ in the example). This is trivially correct, since we are only eliminating from the assignment table (by factoring) unsatisfiable rows.

Abstraction algorithm. We have introduced simple modifications on the algorithm described in [10], which are able to prevent the erroneous abstractions produced by inconsistent states as described previously. Moreover, the resulting algorithm seems to dramatically reduce the number of solver calls.

Definition 5 (Dependent variable). *Let P be a basic block, we say that a given variable $x \in \text{Vars}(P)$ is dependent if the initial value of x in the pre-state is used in P (i.e. x is read before it is written). If P is a basic block in SSA form, x is dependent if x_0 occurs in P .*

Let P be a basic block, E a set of predicates, C_V the set of all possible cubes of E and L_P the logic encoding of P . Our algorithm computes the assignment table of P , by calculating for each possible satisfiable and attainable post-state, which pre-states can lead to it (note that for the pre-state we instantiate only the predicates where dependent variables occur). The pseudo-code is presented as Algorithm 1, where some auxiliary functions are used. $\text{addAllCombinationsFor}(pos)$ appends every possible cube with pos and appends the result to the assignment table. $\text{interpolate}(vars, cube)$ completes the cube by combining it with every possible combination of $vars$, in the correct positions. addLines simply adds rows to the assignment table. An example run is presented in the appendix.

10 Implementation of The Cartesian Abstraction

Implementing the cartesian method is more straight forward than the direct method since, by definition, cartesian abstraction produces boolean programs.

The method for calculating the cartesian abstraction is described in [1], together with optimizations that not only speed up the calculation, but solves to problem of boolean simplification too. In this chapter we'll describe those optimizations, translate them to our SAT encoding of cartesian abstraction and briefly explain how we implemented it.

Optimizing S_V^ψ . Let's drift away from predicate abstraction for a moment and recall the definition of S_V^ψ :

$$S_V^\psi = \bigvee \{c \in C_V \mid \models E(c) \rightarrow \psi\}$$

A naive implementation would check, for every cube $c \in C_V$, if it's valid that $E(c) \rightarrow \psi$. As already discussed, this would require 2^n , where $n = \#V$, solver calls. We can potentially decrease this number if we consider the cubes in increasing order by size, taking into account that for a cube $c \in C_V$:

- i) If $\models c \rightarrow \psi$, any cube c' such that $c \subseteq c'$ will imply ψ too. Therefore we add c to S and stop generating for cubes larger than c . This will actually produce a disjunction of the prime implicants of S_S^ψ . We'll denote this optimization by *priming*.
- ii) If $\models c \rightarrow \neg\psi$, any cube c' such that $c \subseteq c'$ will imply $\neg\psi$ too. Therefore we don't add c to S and immediately prune every cube c' that contains c . We'll denote this optimization by *pruning*.

It's trivial to check that *priming* and *pruning* are applicable to the SAT encoding of S too. We'll provide a proof for *priming*, the other proof is analogous.

Lemma 2. *Let $c \in S_V^{wp(P, E(b_i))}$ for a basicblock P , set of variables V and $b_i \in V$, then, for all $c' \in C_V$ such that $c \subseteq c'$, it holds that $c' \in S_V^{wp(P, E(b_i))}$.*

Proof. Without loss of generality, let's take $c' = c \wedge k$ for some $k \in C_V$ such that $k \cap c = \emptyset$. From hypothesis we have that:

$$\begin{aligned} & E(c) \wedge L_P \wedge E(\neg b_i) \\ \Leftrightarrow & E(c) \rightarrow wp(P, E(b_i)) \\ \Rightarrow & E(c) \wedge E(k) \rightarrow wp(P, E(b_i)) \\ \Leftrightarrow & E(c') \rightarrow wp(P, E(b_i)) \\ \Leftrightarrow & E(c') \wedge L_P \wedge E(\neg b_i) \end{aligned}$$

Therefore, $c' \in S_V^{wp(P, E(b_i))}$. □

With those optimizations in mind, we just need to generate the cubes and check them individually.

11 Handling Control-Flow

So far we have presented the concepts of predicate abstraction, discussed and unified two techniques to calculate the abstraction of a basic block, but a piece of the puzzle is still missing for us to be able to put it all together. In this chapter we're going to discuss different methods for handling control flow, their implications and implementations.

A trivial technique is to simply add the control flow conditions to the set of predicates. This will produce a boolean program with single-variable control flow. The downside of this technique is in the number of predicates that are going to be added, remembering that the abstraction techniques have an exponential number of solver calls on the number of predicates being abstracted.

Let's consider the abstraction of a conditional expression $\psi \rightarrow P_1, P_0$, where P_1, P_0 are basic blocks for a set of predicates E . In a similar paradigm of that of cartesian abstraction, we can think of rulling out every execution path that validates $\neg\psi$ and passes through P_1 , this condition can be represent by $S^{\neg\psi}$, that is, it holds whenever an abstract state implies $\neg\phi$. Determining the states where ψ holds is done, now we have to rule out those execution paths, in fact, boolean programs have an **assume** f ; keyword, that discards from further verification any path that satisfies f . Thus, **assume** $\neg S^{\neg\psi}$; will be evaluated to **assume false**; whenever we're in a path that satisfies $S^{\neg\psi}$. Let's define T^ψ by $\neg S^{\neg\psi}$ (it's called $G(\psi)$ in [1]),

That's how SLAM abstracts control flow statements, as described in [1]. Therefore, the abstraction of $c \rightarrow P_1, P_0$ is presented in (boolean-program) algorithm 2.

Algorithm 2 Abstraction of Conditionals

```
if * then
  assume  $T^\psi$ 
   $\bar{P}_1$ 
else
  assume  $T^{\neg\psi}$ 
   $\bar{P}_0$ 
end if
```

Clarke uses a very similar strategy in [10]. Given a condition ψ , he generates two sets:

$$pos_\psi = \{\alpha(s) \in C_V \mid \psi(s)\}$$

$$neg_\psi = \{\alpha(s) \in C_V \mid \neg\psi(s)\}$$

and whenever the model checker reaches a conditional, it checks whether we're in a state that is an element of pos_ψ or neg_ψ to decide where should the program

counter point to next. Let's take a closer look, without loss of generality, at pos_ψ :

$$\begin{aligned}
pos_\psi &= \{\alpha(s) \in C_V \mid \psi(s)\} \\
&= \{c \in C_V \mid \exists s . \psi(s) \wedge \alpha(s) = c\} \\
&= \{c \in C_V \mid \exists s . \neg(\neg\alpha(s) = c \vee \neg\psi(s))\} \\
&= \{c \in C_V \mid \exists s . \neg(\alpha(s) = c \rightarrow \neg\psi(s))\} \\
&= \{c \in C_V \mid \forall s . \alpha(s) = c \rightarrow \neg\psi(s)\}
\end{aligned}$$

And that is very close to the formulas $S_V^{-\psi}$, used in cartesian abstraction's paradigm for handling control-flow, but it mentions concrete states instead of going with the predicates that represent such state. We could use boolean-program's **assume** keyword or check if the disjunction of pos_ψ (resp. neg_ψ) holds before entering the conditional's then (resp. else) branch to replicate the model checker's behavior in the boolean program setting. Abstracting a **while** ψ block is analogous: we just place a **assume** T^ψ before the abstracted while's block.

Note that even in how control-flow is handled by both pioneers of both cartesian and direct abstraction the methods are very similar, like we expected after finding out the similarities in section 8.

SABS supports both heuristics for handling control-flow, for more information on how to change such options we refer the reader to [18].

12 Conclusion

In this report we have discussed the calculation of predicate abstractions for basic blocks of code, and how, based on the prime consensus theorem, we have adapted the direct abstraction technique to produce Boolean programs, introducing modifications for correctly handling variable initialization, as well as optimizations that substantially reduce the number of necessary calls to the solver. Note that we are not claiming to have produced an algorithm that performs better than working tools based on the direct method, because these tools also incorporate many other optimizations (not fully documented).

We have also shown that the only real difference from the cartesian to the direct method is the number of post-state variables considered at once in the abstraction formula. This observation raises some questions, that are left as future work.

Apart from providing a clear explanation of the abstraction techniques and an in-depth structural comparison we have also produced a software model checker laboratory. Our implementation turned out to be very inefficient compared to the commercial tools available.

13 What's next?

There's still work to do both on the theoretical side and optimization of the prototype. It would be interesting to see how, if at all, the framework provided

opens up room for optimizations. Given that the difference between the cartesian and direct method are the size of the cubes considered, could the precision be (dynamically) adjusted by considering different cube-sizes?

On the prototype side, it's left to implement the complete CEGAR cycle, interface with different boolean program model checkers, considering OOP aspects of SPARK into the abstraction and abstracting multi-procedural code.

References

1. Thomas Ball, Rupak Majumdar, Todd Millstein, and Sriram K. Rajamani. Automatic predicate abstraction of c programs. *SIGPLAN Not.*, 36(5):203–213, May 2001.
2. Thomas Ball and Sriram K. Rajamani. Bebop: A symbolic model checker for boolean programs. pages 113–130. Springer, 2000.
3. J. Barnes. *High Integrity Software: The Spark Approach to Safety and Security*. Addison-Wesley, 2003.
4. Gérard Basler, Matthew Hague, Daniel Kroening, C.-H. Luke Ong, Thomas Wahl, and Haoxian Zhao. Boom: Taking boolean program model checking one step further. In *TACAS*, pages 145–149, 2010.
5. Gérard Berry, Hubert Comon, and Alain Finkel, editors. *Computer Aided Verification, 13th International Conference, CAV 2001, Paris, France, July 18-22, 2001, Proceedings*, volume 2102 of *Lecture Notes in Computer Science*. Springer, 2001.
6. Armin Biere, Alessandro Cimatti, Edmund M. Clarke, Ofer Strichman, and Yunshan Zhu. Bounded model checking. *Advances in Computers*, 58:117–148, 2003.
7. Sagar Chaki, Edmund M. Clarke, Alex Groce, Somesh Jha, and Helmut Veith. Modular verification of software components in c. In *ICSE*, pages 385–395, 2003.
8. Edmund M. Clarke, E. Allen Emerson, and Joseph Sifakis. Model checking: algorithmic verification and debugging. *Commun. ACM*, 52(11):74–84, 2009.
9. Edmund M. Clarke, Orna Grumberg, and David E. Long. Model checking and abstraction. *ACM Trans. Program. Lang. Syst.*, 16(5):1512–1542, 1994.
10. Edmund M. Clarke, Daniel Kroening, Natasha Sharygina, and Karen Yorav. Predicate abstraction of ANSI-C programs using sat. *Formal Methods in System Design*, 25(2-3):105–127, 2004.
11. Patrick Cousot and Radhia Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252, 1977.
12. Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. Efficiently computing static single assignment form and the control dependence graph. *ACM Trans. Program. Lang. Syst.*, 13(4):451–490, 1991.
13. Daniela Carneiro da Cruz, Maria João Frade, and Jorge Sousa Pinto. Verification conditions for single-assignment programs. In *SAC*, pages 1264–1270, 2012.
14. Susanne Graf and Hassen Saïdi. Construction of abstract state graphs with pvs. In *CAV*, pages 72–83, 1997.
15. Valentine Kabanets and Jin yi Cai. Circuit minimization problem. In *STOC*, pages 73–79, 2000.
16. Salvatore La Torre, Madhusudan Parthasarathy, and Gennaro Parlato. Analyzing recursive programs using a fixed-point calculus. *SIGPLAN Not.*, 44(6):211–222, June 2009.

17. Cláudio Lourenço. *A Bounded Model Checker for SPARK Programs*. PhD thesis, Informatics Department, Minho University, Portugal, 2013.
18. Victor Miraldo. *SABS: Spark ABStraction, a tutorial*, 2013.
19. Richard L. Rudell. *Logic Synthesis for VLSI Design*. PhD thesis, EECS Department, University of California, Berkeley, 1989.
20. Adacore SPARK Team. SPARK the SPADE ada kernel. 2011.